Betting on the House

Subjective Expectations and Market Choices

Nicolas Bottan

Cornell University

Ricardo Perez-Truglia*

University of California, Los Angeles

This Draft: April 25, 2020.

Abstract

House price expectations play a central role in macroeconomics and finance. However, there is little *direct* evidence on how these expectations affect market choices. We provide the first experimental evidence based on a large-scale, high-stakes field experiment. We provided information about the evolution of local house prices to 57,910 U.S. individuals who had recently listed their houses on the market. Collectively, the properties listed by these subjects were worth $34 billion dollars. We randomized the information provided to create non-deceptive, exogenous variation in the subjects' expectations about the future home price growth. We use rich administrative data to measure the effects of these information shocks on the recipient's market choices. We find that, consistent with economic theory, higher home price expectations caused the owners to delay selling their properties. These effects are statistically highly significant, economically large in magnitude, and robust to a number of sharp checks. Our results indicate that market choices are highly elastic to expectations: a 1 percentage point increase in home price expectations reduces the probability of selling within 6 months by 2.45 percentage points.

JEL Classification: C81, C93, D83, D84, E31, R31.
Keywords: expectations, experiment, housing market, information.

*Bottan: Cornell University, Policy Analysis and Management, 2301 MVR Hall Ithaca, NY 14853, nicolas.bottan@cornell.edu. Perez-Truglia: University of California, Los Angeles, Anderson School of Management, Office C515, 110 Westwood Plaza, Los Angeles CA 90403, ricardo.truglia@anderson.ucla.edu. We are thankful for excellent comments from colleagues and seminar discussants. We thank funding from the UCLA's Ziman Center for Real Estate's Rosalinde and Arthur Gilbert Program in Real Estate, Finance and Urban Economics and UCLA Anderson's Behavioral Lab. Zhihao Han provided excellent research assistance. This project was reviewed and approved in advance by the Institutional Review Boards at University of California Los Angeles (#18-001496 and #19-000945) and Cornell University (#1811008440). The experiments were pre-registered in the AEA RCT Registry (#0003663).

1 Introduction

Consumer expectations play a central role in modern macroeconomics and finance and are a prime concern for policymakers (Bernanke, 2007). This is specially true for the households' expectations about the future growth in home prices, also known as home price expectations. Many of the accounts of the U.S. housing crisis of the late-2000s revolve around home price expectations (Shiller, 2005; Glaeser and Nathanson, 2015; Gennaioli and Shleifer, 2018; Kaplan et al., 2019). And given the large fraction of the country's wealth tied up in real estate, home price expectations can potentially have massive welfare and policy implications.

According to economic theory, home price expectations should have first order effects on the market choices in the real estate market. In particular, sellers should be more excited about selling a property when they expect prices to go down, and more likely to hold on to the property when they expect prices to go up. Despite this its central role, there is little *direct* evidence on whether this causal relationship exists at all. In this paper, we fill this gap in the literature through a from large-scale, high-stakes, pre-registered field experiment.

The relationship between home price expectations and market choices is plagued with causal identification challenges. Consider, for example, the time-series evidence that the average home price expectations tend to co-move with actual home prices. It is difficult to tell what the direction of causality is. Are home prices increasing because of the increase in expectations? Or are survey expectations going up in response to the higher home prices? Like this, there are other omitted-variable biases that can arise, which do not even go always in the same direction.[1] This is one of those contexts in which causal inference is impossible without experimental data.

In an ideal experiment, we would take a sample of homeowners who are considering selling their properties. Before they start receiving offers, we would like to flip a coin to randomize their expectations. If the coin lands on heads, the person expects the median home prices to appreciate by 1% over the next year. If the coin lands on tails, they expect the median home prices to appreciate by 10% over the next year. The hypothesis is that the person will hold on to the property for longer if the coin lands on tails instead of heads. Moreover, the magnitude of this effect can provide a measurement of the elasticity of market choices to home price expectations.

We designed a field experiment that resembles the ideal experiment quite closely. We mailed letters to homeowners who had recently listed their houses on the market. These letters included information on the prices for similar properties to the one the recipient had listed – i.e., the median value of homes in the same ZIP Code and with the same number of

[1] For instance, because survey data is often subject to substantial measurement error, that can lead to an attenuation bias and, thus, an under-estimation of the causal effect of expectations on behavior.

bedrooms. The information contained in the letter was randomized to create non-deceptive exogenous shocks to the recipient's home price expectations. We then use the administrative records to measure when the homeowners sold their properties. Thus, we can measure if the exogenous shocks to expectations from the letters affected the recipients' market choices.

The information experiment was designed as follows. All letters included information about the current median home price. We randomized two additional aspects of the letter. First, we randomized whether, in addition to the information on the current price level, the letter includes additional information on the price evolution – hereinafter referred to as the *disclosure-randomization*. Second, conditional on providing that additional information on the price evolution, we randomly picked one of five possible information sources – hereinafter, referred to as *source-randomization*. We used five sources that, according to prior survey experiments, individuals find relevant for the formation of home price expectations (Armona et al., 2019; Fuster et al., 2018): the average price change over the past year, the average price change over the past two years, or the forecast about the price change in the next year according to one of three different statistical models.

The disclosure-randomization and source-randomization create exogenous information shocks. For example, consider the following illustration for the source-randomization. If you were a subject selling a 2-bedroom home in ZIP Code 33308, you could be randomized into one of the following five signals of price growth: an annual growth rate of 1.2% over the past year; an annual growth rate of 3.6% over the past two years; or one of the three forecasts of 2.6% (according to statistical model 1), 4.1% (model 2) and 3.5% (model 3). Relative to the first signal, selecting the second signal would amount to an information shock of 2.4 percentage points (i.e., 3.6 minus 1.2). Relative to the first signal, selecting the third signal would amount to an information shock of 1.4 percentage points. And so on for the other two signals.

To complement the field experiment, we conducted a supplementary survey experiment with respondents recruited from an online sample. In this survey, we expose subjects to the same information treatments used in the field experiment, and we measure their home price expectations using a standard survey question. We implemented the online survey on the same date as the field experiment, and collected responses from 1,404 individuals. Using this supplementary survey, we confirm that our information shocks have the expected effects on home price expectations. A 1 percentage point information shock increases home price expectations by 0.205 percentage points. The degree to which individuals incorporate the information providing in the experiment is consistent with findings from prior survey experiments (Armona et al., 2019; Fuster et al., 2018).

Next, we return to the main focus of this study: measuring the effect of the information shocks on actual market behavior. The main implementation challenge for this experiment is to identify individuals who are actively deciding whether to sell their homes or not, and be able to mail them information as well as tracking whether the property is sold or not with administrative records. We identified active listings using a publicly-available information from a major online listing website. We use the county-level property identifiers to match those listings to the County Assessor's Office records to obtain information on the homeowners including their full names and mailing addresses. Moreover, we use those same identifiers to track whether the property was sold or not on a daily basis. We conducted the field experiment in June 2019. We mailed letters to 57,910 unique homeowners who listed properties in 36 counties across seven U.S. states. The properties were valued at a total of $34 billion dollars. On average subjects are 31.9% female, 58.7 years old, have a household income of $128,059 and the average property has 3.25 bedrooms, 2.61 bathrooms, 2,295 sq. ft. and was listed at $574,756.

The results from the field experiment indicate that the information shocks have a causal effect on market choices and in the direction predicted by economic theory: positive shocks to the home price expectations of the seller reduced the speed at which the properties are sold. This effect is not only highly statistically significant, but also economically large. A percentage point increase in information shock causes a 0.330 percentage point drop in the probability that the property is sold within 12 weeks, implying a behavioral elasticity of -0.33.

We provide a number of sharp robustness checks. The effects are persistent during the whole 6 months of post-treatment period. For example, the behavioral elasticity estimated at 28 weeks post-treatment (-0.325) is close to the corresponding elasticity at 12 weeks post-treatment (-0.330). We use an event-study analysis of the experiment to shows that the timing of the effects follow the the timing of the delivery of letters as expected. We estimate the effects on the outcomes right before the letters started being delivered. Since the letters had not reached the subjects yet, they should have no effect on the sales outcome. As expected, the effects of the information shocks are precisely estimated around zero in the pre-treatment period. Next, we should expect that the effects should intensify during the 7-week period while the letters were being delivered to and read by the subjects and stabilize a few weeks later. The event-study analysis fits exactly that prediction.

As an additional falsification test, we estimate placebo regressions that are identical to the baseline specification only that they use as dependent variables the pre-treatment characteristics such as the number of days the property had been listed for prior to our experiment or the original listing price. Because those outcomes were determined prior to the letter delivery, it should be impossible for the randomization of the information shocks contained in the letter to affect them. As expected, we find effects that are close to zero, statistically insignificant and precisely estimated. And we provide a number of additional robustness checks. For example, we show that the results are similar if we focus on the disclosure-randomization only or on the source-randomization only.

The effects reported above give an attenuated estimate of the true effects of expectations on market choices. Thus, we corrections for two main sources of attenuation. First, conditional on reading the letter, the information shock introduced in the letter may not materialize in the readerâs expectations, for example, because the recipient already knew the information. We use the results from the survey experiment to estimate this rate of pass-through from information shocks to expectations: i.e., a 1 percentage point increase in the information shock increases expectations by 0.205 percentage points. The second source of attenuation bias is that some subjects may not have received the letters, may have discarded them without reading them, or may have read the letter but after they already sold the property. Our back of the envelope calculations suggest that 35.1% of subjects belong to this category.

As a result, the scaled-up elasticity between expectations and probability of selling 28 weeks post-treatment boils down to -2.45 ($=\frac{-0.328}{0.205 \cdot 0.649}$). This estimate indicates that subjects are highly elastic to their expectations: increasing home price expectations by 1 percentage point would cause a reduction in the sales probability of 2.45 percentage points. This channel also has potential to explain a significant fraction of the market choices: an increase of one standard deviation in expectations (5.39 pp, according to the survey expectations in our auxiliary sample) would reduce the sales probability by 13.21 pp ($= 2.45 \cdot 5.39$).

Finally, while the above estimates pertain to the population as a whole, they may mask meaningful heterogeneity. Some individuals may care about their updated expectations but have little flexibility to change their market decision. For instance, some individual maybe they need to sell the home immediately because they already bought another home or because they need to move for work. Also, some individuals may care about expectations but not specifically about the expectation we manipulate. For example, if after selling the property you plan on buying another home in the same market, then the expected appreciation in the market should have little effect on the decision to sell, because you'll benefit from the expected appreciation either way. To asses how large these two frictions may be, we would estimate the model on a subset of the subject pool who did not face those frictions. We provide an approximation, by looking at the (pre-registered) heterogeneity of effects between owner-occupied (78.4% of subjects) versus non-owner occupied (21.6% of subjects). We argue that, relative to the owner-occupied, the non owner-occupied are less subject to the two frictions described above.

We find that the effects of information shocks are statistically significant for both samples. And consistent with the above conjecture, the effects are almost three times as strong for the non owner-occupied as for the owner-occupied. For example, at 28 weeks after the start of the letter delivery, the coefficient on the information shock is -0.637 (p-value=0.003) for the non owner-occupied and -0.225 (p-value=0.066) for the owner-occupied. For the sample of non owner-occupied, the behavioral elasticity boils down to -4.79 pp $=(\frac{-0.637}{0.205 \cdot 0.649})$. This large elasticity shows that, in absence of frictions, the effects of expectations on behavior can be even more substantial than reported above.

This study is related and contributes to various strands of literature.

This study relates to literature on the role of subjective expectations on the housing market. In particular, there are two studies that are most closely related to this paper. First, Armona et al. (2019) use survey experiments to show how the provision of information about home prices affects the expectation formation of recipients. The authors also measure the effects of this information on a laboratory experiment. We contribute by looking at high-stakes decisions in a natural context. Second, Bailey et al. (2018) shows that individuals whose geographically distant friends experienced larger recent home price increases are more likely to transition from renting to owning. We contribute to this paper by using experimental evidence.

Our study is related to a broader and growing literature on subjective expectations. These information-provision experiments have been used in the contexts of home price expectations (Armona et al., 2019; Fuster et al., 2018) as well as other macroeconomic expectations such as inflation expectations (Armantier et al., 2016; Cavallo et al., 2017; Coibion et al., 2018). These studies typically provide a random subset of respondents with a piece of information and measure the corresponding effects on expectations. These studies also measure the effects of information on other outcomes such as survey measures of intended behavior, or small-stakes laboratory experiments (Armantier et al., 2015; Armona et al., 2019).

We contribute to this literature by measuring the effects of information on market behavior in a nearly-ideal context. The dependent variable in these studies are also based on survey data, which is subject to many criticisms such as experimenter demand effect. Instead, we combine the information experiments with rich behavioral data from administrative records to measure the effects on real, high-stakes market behavior. This is a naturally-occurring context that a large fraction of Americans will face at some point in their lives.[2] This is a high-stakes context insofar the decision being made accounts for a large fraction of the net worth of the decision maker. Indeed, buying or selling a home is arguably one of the biggest

[2]According to the U.S. Census data for the fourth quarter of 2019, the homeownership rate was 65.1%. And according to data from the National Association of Realtors, 5.34 homes were sold in 2019.

decisions that Americans have to make at some point in their lives (Brooks, 2017). Our large-scale experiment involving nearly sixty thousand subjects, allows us to provide estimates that are precisely estimated and also gives us the ability to provide sharp falsification tests – for example, providing an event-study analysis of the information provision experiment.

We believe that our experimental instructions can be followed by other researchers to study not only home price expectations, but countless questions from macroeconomics, urban economics, finance, real estate economics and behavioral economics. Our experimental framework does not have any barriers to entry – we rely on data sources that are publicly available and easily accessible. The experiment can be implemented in just a few weeks, for less than \$0.25 per subject, and potentially with up to 1 million subjects at a time.

The rest of the paper proceeds as follows. Section 2 describes the research design. Section 3 discusses the institutional context and data sources. Sections 4 presents the results. The last section concludes.

2 Research Design

2.1 Econometric Model

We test a simple prediction of economic theory: higher home price expectations should increase the reservation value of the seller and thus, on average, should make him take longer to sell the home. As explained in the model of Appendix A, this is a basic prediction from an asset pricing model, requiring minimal assumptions.

In the experiment, we can randomize the provision of a signal. If we provide a signal of X%, the effect on the time taken to sell the home should depend on the prior beliefs: if prior <X%, we would update upward (and sell the home slower); if prior =X%, we should not update; If prior >X% should update downward (and sell the home faster). The usual way of dealing with people's beliefs is by measuring their prior beliefs directly. However, in the context of our field experiment, it would be impossible to measure the prior beliefs of tens of thousands of specific homeowners who have listed their homes. Instead, we use a simple econometric framework based on disclosure- and source-randomization to identify the effects of expectations.

To illustrate the disclosure-randomization, consider the case of a single signal. For instance, the 1-year past change. Let Y be the outcome of interest (it could be the posterior belief about home price expectations, or whether the home was sold 3 months later). E_i^{signal} is the value of the signal that can be shown or concealed. T_i is an indicator variable that takes the value 1 if the signal was shown in the letter and 0 otherwise. The regression of interest is the following:

$$Y = \nu_0 + \nu_1 \cdot E_i^{signal} \cdot T_i + \nu_2 \cdot E_i^{signal} + \nu_3 \cdot T_i + \varepsilon_i \tag{1}$$

The coefficient ν_2 shows the relationship between the signal and the outcome when the signal is not included in the letter. The coefficient ν_1 measures how the signal affects the outcome when included in the letter. A $\nu_1 > 0$ would indicate that at least some people incorporate the signal shown to them. Note that this approach relies on heterogeneity in signals across subjects. Imagine for example that all subjects lived in the same neighborhood, so the signal takes the same value for every i. Then there would be no variation to identify ν_1 (or ν_2, for that matter).

A second source of variation comes from exploiting source-randomization of signals within subjects (e.g., Bottan and Perez-Truglia, 2017). Intuitively, we exploit the fact that a given individual can be assigned to different sources of information about home price changes and that, depending on the source, the underlying signal may end up being more or less optimistic. Let $E_i^{signal(j)}$ with $j \in 1..J$ be the set of possible signals. For example, $j = 1$ could be the average growth in the last year, while $j = 2$ could be the average growth in the past two years. Let $E_i^{signal(j^*)}$ be the signal that was randomly chosen for individual i. That signal might be included in the letter. Whether the signal is included again indicated by the indicator variable T_i. Consider the following regression:

$$Y = \nu_0 + \nu_1 \cdot E_i^{signal(j^*)} \cdot T_i + \nu_2 \cdot T_i + \sum_j \beta_j \cdot E_i^{signal(j)} + \varepsilon_i \tag{2}$$

This is the key regression we use in the implementation. For the sake of brevity, the equation above controls linearly for $E_i^{signal(j)}$ but in practice we control for it flexibly by including a set of decile dummies in addition to the linear term. Note also that the orthogonality assumption, required for causal inference, is $E[E_i^{signal(j^*)} \cdot T_i \cdot \varepsilon_i] = 0$. The random assignment of both $E_i^{signal(j^*)}$ and T_i should be jointly sufficient to warrant this assumption.

This regression is exploiting two distinct sources of variation. The first compares individuals who could have been shown signal j, but were not, versus individuals who were shown that signal. That is the same variation discussed in the previous regression. To see this, note that if allow for a single signal, then equation (2) is identical to the previous specification from equation (1).

The second source of variation exploited by equation (2) from Section 2 is that we could have shown an individual J different signals, but we randomly chose one of those. This variation comes within the group of individuals who have been shown a signal. To see this, note that we can still identify ν_1 even if T_i took the value 1 for the entire subject pool.

For example, for home A, we could disclose two signals: the past 1-year says 5%, the other says 10%. Relative to the first signal, choosing the second signal constitutes a "shock" of 5 percentage points in the information shown. This regression asks: for each percentage points of additional shock, what is the effect on the outcome of interest? A $\nu_1 > 0$ would arise if A's outcome is higher when shown the 10% signal instead of the 5% signal.

We can estimate the model using the two sources of variation separately, to see their contributions to the final estimates. To isolate the across-source variation, we can simply re-estimate the model restricting to the sample who were shown signals (i.e., $T_i = 1$), which will kill all the across-subject variation. To isolate the across-individual variation, we can estimate equation (2) from Section 2 but including $E_i^{signal(j^*)}$ as an additional control variable. Note that this will kill the across-source variation. Conditional on receiving the information, the effect of receiving source A versus source B will be absorbed by the parameter on $E_i^{signal(j^*)}$. Thus, the only remaining variation identifying ν_1 will be the fact that the signal is shown to some individuals but not to others.

2.2 Mailing Design

All subjects in the experiment were sent a letter by mail. Figure 1 shows a sample of the envelope. We took a number of measures to communicate that even though this letter was not solicited, it came from a legitimate source. The top-left corner of the envelope included the logo of University of California at Los Angeles (UCLA) as well as a note indicating that this was part of a research study. Additionally, the top-right corner of the envelope included the non-profit organization postage.

A sample letter (for a fictitious subject) is included in Appendix C. Because the content of the letter is different for different treatment groups, Figure 2 shows the content of the letter that was common to all subjects. The front and reverse of the letter first page is shown in Figures 2.a and b. In the letter, we also took measures to reassure the recipient that the communication was legitimate. The letter includes UCLA's official logo in the header, as well as contact information in the footer: a physical address that they could write, as well as the URL for the study's website with additional information about the study. A copy of this website, that was hosted on UCLA's official website, is attached in Appendix D. The website provided general information about the study (without specifying any hypotheses) as well as contacts for the researchers and the Institutional Review Board.

There are two components of the letter that were different across treatment groups. Figure 2 contains placeholders (marked as *«Information»* and *«Information Details»*) indicating the placement of the different components of the letter. A table with information about home prices goes in the first placeholder. Figure 3 shows a sample table for each of the treatment groups (discussed below). The second placeholder corresponds to the methodological notes that accompany the table, such as the data sources and statistical models used – samples of these notes are available in Appendix Figure B.1.

Homeowners are randomized into one of six treatment groups. This is the only aspect of the letter that was randomized. These treatment groups differ in the information included in the main table. All of these six groups provide information on the current median home value of similar properties in the same area. For example, if the recipient had listed a 3-bedroom home in ZIP Code 90210, the letter would indicate Zillow's estimate of the median home value for 3-bedroom homes in 90210.[3]

What may be different in different treatment groups is that the letter may contain additional information that the individual may want to incorporate in their expectations about home prices: the price change in recent years, or forecasts for the price change in the following year. We chose these two information sources because they have been proven to affect home expectations in survey experiments. For example, Fuster et al. (2018) show that, upon being shown those one of these types of information (past price changes and professional forecasts), they update their expectations in the direction of the information. Fuster et al. (2018) also show that most households are willing to pay in order to acquire these two types of information, which also suggests that they find these two types of sources useful. Additionally, these are the two types of information used in economic modelling. On the one hand, according to the the backward-looking expectations model individuals form beliefs looking at recent price changes (Case and Shiller, 1989; Shiller, 2005). On the other hand, rational expectation models indicate that households should prefer information from expert forecasts because experts have rational expectations that use all available information optimally (Carroll, 2003).

All letters include the information on the current price *level*. The difference between the six treatment groups is whether the table includes additional information on price *changes*. The following is the additional information included in each of the six treatment groups:

- **Baseline:** none.
- **Past-1:** price change over the past year.
- **Past-2:** price change over the past two years.
- **Forecast-1:** price change forecast over the next year using statistical model 1.

[3] We used the same property types used by the online real market platforms: 1-bedroom, 2-bedroom, 3-bedroom, 4-bedroom and 5+ bedroom. In a small minority of properties, the number of bedrooms is not available and thus we use a separate category ("all homes") that aggregates all the previous categories. Also, Zillow does not produce estimates of median home values for all combinations of ZIP codes and bedrooms. Thus, if the property falls into that category, we would also assign the "all homes" category.

- **Forecast-2:** price change forecast over the next year using statistical model 2.

- **Forecast-3:** price change forecast over the next year using statistical model 3.

All our forecasts are models estimated year/ZIP Code-level data on the Zillow Home Value Index for the period 1997-2019. All three treatment arms are based on simple auto-regressive models, with the difference lying in the set of explanatory variables chosen. The first model uses five lags of the dependent variable. The second model uses five lags of the dependent variable plus five lags of the state-level average of the dependent variable. The third model uses three lags of the dependent variable, three lags of the city-level average of the dependent variable and three lags of the city-level employment rate. Appendix B.3 presents more details, including a comparison of the out-of-sample predictive power between the three models and that of Zillow's official forecast.

Each panel of Figure 3 corresponds to the hypothetical table that a given individual would received if assigned to each of the six treatment groups. These are real examples based on an individual selling a 2-bedroom home in ZIP Code 33308. Panel (a) shows the baseline letter, which only includes the current median price level. The following five panels add information on the evolution of that median price: panel (b) shows an annual growth rate of 1.2% over the past year (Past-1); panel (c) shows an annual growth rate of 3.6% over the past two years (Past-2); panel (d) shows the annual growth of rate 2.6% projected by Model 1 (Forecast-1); panel (e) shows the annual growth rate of 4.1% projected by Model 2 (Forecast-2); and panel (f) shows the annual grwoth rate of 3.5% projected by Model 3 (Forecast-3).

Based on the discussion in Section 2, for identification we require heterogeneity in these signals across individuals and across sources. Figure 3 shows that there is plenty of variation across sources: relative to Past-1, selecting the Past-2 would amount to an information shock of 2.4 percentage points (i.e., 3.6 minus 1.2); relative to the Past-1, selecting Forecast-1 would amount to an information shock of 1.4 percentage points; and so on. Section 3 discusses the variation in more detail and for the whole sample.

2.3 Survey Design

While the main hypothesis relates to the effect of the information provided in the letters on the market behavior, we would like to complement that evidence with survey data. Specifically, we would like to use survey data to validate the identification strategy: whether the information shocks affected expectations in the expected direction. Moreover, we would like to quantify the strength of the effects on expectations, to aid in the interpretation of the magnitude of the effects on behavior.

A natural approach is to survey the recipients of our letters, to measure if the information contained in the letter shifted their home price expectations in the expected direction. In the letter, we included a URL to an online survey (see the blue box at the bottom of Figure 2.a). To link survey responses to each different respondent, we included a 5-letter survey code in the letter, next to the URL. Respondents had to enter the survey code in the first screen of the survey (a copy of which is attached as Appendix E) before they could answer any questions. The survey consisted of a handful of questions about home price expectations.

Based on the results from other mailing studies, we anticipated a tiny response rate to the survey. In a different context, on tax compliance, Perez-Truglia and Troiano (2018) included a link to an online survey similar to ours. Only 0.2% of the subjects who received the letter responded to the survey. This extremely low response rate, in addition to power limitations, introduces the additional concern that selection into the survey could be very extreme and also affected by treatment assignment. Despite the high likelihood that the survey data would not be useful to measure experimental effects, we still wanted to deploy this survey for an alternative use: survey responses can serve as a proxy for the dates when recipients were opening the letters, as in Perez-Truglia and Cruces (2017) and Perez-Truglia and Troiano (2018).

In anticipation of the likely event that the letter survey would not provide useful data, we collected survey data from a different source that was not subject to the concerns described above. We designed a separate survey experiment that randomized the same information included in the letter experiment, but was designed to be conducted on an auxiliary sample of respondents recruited from online platforms. The full survey instrument was included in the pre-registration, and this supplemental survey was conducted simultaneously with the main field experiment.

The full survey instrument is included in Appendix F. The structure of the online survey was as follows:

- **Step 1 (Elicit Property Details):** To provide randomized information relevant to the participant's local housing market, we first asked all respondents about their current living situation. This included whether they lived with their parents/legal guardians, whether they rent or own the property, the type of residence (single family home, apartment, etc.), number of bedrooms, and 5-digit ZIP code.

- **Step 2 (Elicit Prior Belief):** Respondents were then shown the current median price (in May 2019) for a similar home (measured by number of rooms) in their ZIP code. We then asked the subject what they thought the median price will be 1 year later (in May 2020). We also measured the level of confidence in their answers on a 4 point scale (ranging from not at all confident to very confident).

- **Step 3 (Information-Provision Experiment):** Respondents are told that some survey participants will be randomly chosen to receive information about home prices. Respondents were assigned to one of 6 treatment groups (described above). In the following screen, respondents found out the information selected for them. Notice that the information is presented to respondents in a table format, as in the letter experiment, with the methodological notes appearing after clicking on a button below the table.

- **Step 4 (Elicit Posterior Belief):** In the following screen, the respondent is told that all respondents are given the opportunity to reassess their guess about the future home prices, regardless of their initial responses or whether they received information. We re-elicited the same question on the median price 1 year later. After that question, we measured some additional post-treatment beliefs: their confidence on their expectations, longer-term expectations (5 years instead of 1 year later) and, as a placebo outcome, respondents were also asked about their stock market expectations (respondents were told the closing price of the Dow Jones on May 31st and asked their belief about this price 1 year later).

After this main module, we included some secondary questions. First, we measured some secondary outcomes that we expected could potentially be useful to disentangle mechanisms. Next, we included a battery of demographic questions to be able to compare this auxiliary sample to the main sample from the main experiment and also in case we needed to conduct heterogeneity analysis. Last, we designed a three-question follow-up to be conducted a month later. For more details, see Appendix B.5.

3 Data Sources and Implementation Details

3.1 Data Sources

To implement the mailing experiment in our context we needed two types of data: a source of active real estate listings, and County Assessor's Office data for the secured property tax rolls. We obtained data on real estate listings from a major listing website. These data include information such as address, listing price, property characteristics (number of bedrooms, bathrooms, size, days on market), and the assessor's parcel number (APN). We use the property APN (a unique identifying number assigned to each parcel by a County Assessor's Office) to have the ability of matching the listings to County Assessor's secured tax rolls.

Even though tax rolls are publicly available, they are not always easily accessible. Secured tax rolls contain information on property owners, property characteristics and property tax information. Important for our experiment is information on owner names and their mailing address.

There are a number of necessary conditions for a county to be eligible for this field experiment. First, we must have access to the secure tax rolls from the County Assessor's office. Second, the Assessor's data must include information about the names and full mailing addresses of the owners of the property.[4] And third, the website used to identify properties on the market should have coverage in this county. We identified 36 counties that met all three conditions, distributed across 7 states: 30 counties in Florida, Los Angeles County (CA), Maricopa County (AZ), Clark County (NV), Cuyahoga County (OH), King County (WA), and Harris County (TX). We stopped at 36 counties because that gave us enough subjects for our experiment – there may be many other counties for which these conditions are met. For more details on the data sources, see Appendix B.1.

3.2 Implementation of Mailing Experiment

Combining the geographic restrictions and eligibility criteria we define the universe of eligible listings for the letter experiment. We obtained the latest versions of the Secured Tax Rolls from Assessor's Offices prior to scraping listings for the experiment on May 28, 2019. Of the 173,708 active listing scraped, around 164,298 had an associated APN. For these listings, we had a near perfect merge to the county Assessor's data (164,176). Since the number of individuals in this initial sample was substantially higher than the number of subjects needed for our experiment, we adopted a conservative approach and limited the subject pool to those individuals for whom the highest-quality information was available. For example, we excluded individuals who changed addresses recently, properties owned by businesses, non-residential properties and properties owned by individuals with multiple properties in the same county.

After applying these criteria, our potential subject pool consisted of 61,176 individuals. We selected a random sample of 60,000 of those individuals to be sent a letter. We then excluded a minority (3.4%) of the letters because the mailing addresses were flagged as undeliverable or vacant by the United States Postal Service (1,193 letters) or because updated records indicate that we sent the letter to a former owner (845 letters).[5] This leaves us with a final sample of 57,910 individuals in the subject pool. Of these individuals, 20% were

[4] In most cases (78.41%) the mailing address matches the address for the listed property.

[5] Sales data are recorded with a lag. As a result there were 845 letters that, by the time we mailed the letters, they no longer belonged to that person.

randomly assigned to the Baseline treatment, 15% to Past-1, 15% to Past-2, and 16.6% to each of Forecast-1, Forecast-2 and Forecast-3. The letters were mailed on June 10, 2019. For more implementation details, see Appendix B.2.

3.3 Descriptive Statistics and Randomization Balance

Table 1 shows some basic descriptive statistics about the properties in the subject pool. Column (1) presents the average characteristics for the whole subject pool of 57,910 recipients. The average property had been on the market for 86 days, was listed for $575,000, had 3-bedrooms, 2.6 bathrooms, 2,300 sq. ft. of living space and a lot size of 12,000 sq. ft. Additionally, Table 1 provides a check for whether there is balance in observable characteristics across treatment groups. Subjects were cross-randomized to one of six treatment groups. In columns (2) through (7) of Table 1, we break down the average characteristics by each of the six treatment groups. The last column reports p-values for the null hypothesis that each average characteristic is the same across the six treatment groups. The results show that, as is consistent with successful random assignment, the observable characteristics are balanced across treatments.

Appendix B.4 presents further details, including the characteristics of the property owners: on average, recipients were 58 years old, 31.9% female, 68.8% white, and had a household income of $128,000 per year. We also show that our sample is quite representative in these observable characteristics of the universe of homeowners in the United States. The main exception is home values, which are twice as large for the subject pool than for the country as a whole, primarily due to the subject pool comprising states with high property values such as California.

3.4 Variation in Signals

As explained in Section 2, the identification strategy exploits variation in signals across individuals and across sources. This section shows the degree of this variation.

The results are presented in Figure 4. Figure 4.a presents the heterogeneity within information sources. It shows a histogram of the signal that subjects would have received had they been assigned to a given information source: Past-1. The results shows that there is plenty of heterogeneity across subjects. Subjects in the 10th percentile were living in areas where median home values declined by -0.7% in the previous 12 months. On the other extreme, subjects in the 90th percentile were living in areas where property values increased by 8.6%. The degree of heterogeneity is also comparable for the other four sources: e.g., while the standard deviation in signals across individuals is 3.8 percentage points for Past-1, the corresponding figure ranges from 1.2 to 4.0 for the other sources (for more details, see Appendix B.6).

In turn, Figure 4.b presents the heterogeneity in signals across information sources. It is a scatterplot showing the relationship between the signals that the subjects would have received if they had been assigned to the Past-1 treatment (i.e., the annual growth rate over the past year) versus the Past-2 treatment (i.e., the annual growth rate over the past two years). As expected due to home price momentum, these two potential signals are highly correlated: on average, an extra 1% increase in the annual price change over the past 1 year is associated with an extra 0.659% increase in the annual price change over the past 2 years. This relationship is partly mechanical (the y-axis is an average that includes the x-axis), but also partly due to the well-known momentum in home prices. In any case, the most important fact is that the relationship between these two potential signals is far from perfect, based on the $R^2 = 0.659$ that is far below 1.

Another way of illustrating this is by noting that while there are plenty of observations lying close to the 45-degree line (implying that the signal would be the same for the two information sources), there are plenty of observations that are significantly off the 45-degree line. Indeed, Figure 4.b highlights one specific example, for 2-bedroom homes in ZIP Code 90210. The recipient would have been shown a price change of -2% if randomly assigned to the Past-1 treatment group, but a price change of 6% if randomly assigned to the Past-2 treatment group. The correlations between pairs of sources is quite similar for all other combinations: e.g., while the correlation between Past-1 and Past-2 is 0.81, the correlations between the other pairs range from 0.3 to 0.89 (see Appendix B.6 for more details).

3.5 Letter Delivery

The letters were shipped on June 10 2019. It may take weeks from that date for the letter to be read by the recipients. According to the U.S. Monitor Non-Profit Standard Mail Delivery Study, 95% of the mail it studied was delivered successfully, and it took the non-profit an average of 10 days to arrive in homes, with some letters arriving as far as a month later (U.S. Monitor, 2014). Even after letter delivery, it may take days or even weeks for the typical subject to open and read the letter – some subjects may not check their mail daily, or they may be travelling.

We use the distribution of dates when the surveys were completed as our proxy for the dates when the letters were actually read – herein after referred to as “read-receipt.”[6] The

[6] Our proxy probably has a bit of upward-bias because some people may have read the letter and waited a few days until responding to the survey. Another potential source for bias, which may be upwards or downwards, is that survey respondents may open the letters more or less slowly than survey non-respondents.

results are presented in Figure 5.a. The first survey response was received in June 15, thus marking the start of letter delivery. Indeed, this date coincided with the best guess provided by the mailing company.[7] Figure 5.a suggests that the letters were read progressively over the course of the following two months. The median time to read the letter was about 3 weeks (from start date of letter delivery).

3.6 Outcome Variable: Home Sales

To measure the market choices, we track these properties on the listing website mentioned previously. We scraped the administrative data for the listing on a weekly basis starting in June 1 2019 until December 21 2019. These scrapped records indicate whether the property was sold and, if so, the date when the property was sold. We use the earliest date for which there is confirmation of the property sale, either according to the records from the Multiple Listing Service or from the County Assesor's Office.[8] This is the source for our main outcome variable, as listed in the AEA RCT pre-registry: the probability that a property is sold by a given date.

The evolution of the sales outcome is shown in Figure 5.b. Note that the fraction of homes sold increases smoothly over time. Twelve weeks after the start of the letter delivery, 38.8% of homes had been sold already. Twenty weeks after the start of the letter delivery, 50.6% of the homes had been sold. And by the end of our panel data, 28 weeks after the start of the letter delivery, 57.5% of the properties had been sold.

3.7 Implementation of Online Survey

We recruited subjects from Amazonâs Mechanical Turk (AMT) online marketplace. We followed several best practices for recruiting individuals for online surveys and experiments using AMT to ensure high quality responses (see, for instance, Crump et al. (2013)). We offered residents of the United States to participate in a 5-minute survey in exchange for \$0.75. We collected responses from June 21 to June 24.[9] The final AMT sample includes 1,404 respondents, who were assigned to treatment groups using the same probabilities as in the main experiment: 20% to the Baseline treatment, 15% to each of Past-1 and Past-2, and 16.6% to each of Forecast-1, Forecast-2 and Forecast-3. Appendix B.5 presents more details about the AMT sample, and is summarized below. Relative to the field experiment sample and the universe of U.S. homeowners, the online sample is younger, less wealthy and lives

[7] This estimate was based on the location of the shipping facility (Lombard, Illinois), the location of the letter recipients and historical data.

[8] Our records usually include both dates, which are normally days apart.

[9] The survey included a short follow-up module. Details about it are presented in Appendix B.5.

in smaller and cheaper homes. And, as is consistent with successful random assignment, the observable characteristics are balanced across treatment groups.

4 Main Results

4.1 Effects on Survey Expectations

The main results are presented in Table 2. All the regressions shown in this table correspond to equation (2) from Section 2, with the only difference between columns being a different dependent variable. *Information Shock* refers to the key independent variable: $E_i^{signal(j^*)} \cdot T_i$.

While our main interest is to study the effects on market choices, we start by analyzing the effects of the information shocks on home price expectations. This provides a natural check to the identification strategy, and it is useful to interpret the magnitude of the effects on behavior. In columns (1) through (4), the dependent variables come from the AMT survey. In column (1), the dependent variable is the posterior belief about the local home price one year ahead. In other words, the coefficient on *Information Shock* measures how a 1 pp increase in the information shock affects the subsequent one-year-ahead expectation of the respondents. We find that coefficient on *Information Shock* is positive (0.205) and statistically significant (p-value=0.001). A 1 pp increase in the information shock causes the respondent to expect home prices to grow 0.205 pp more over the next 12 months. That is, there is a 20.5% "pass-through" from information shocks to expectations. The fact that the pass-through is significantly above zero means that individuals do find the signals provided to them informative to form their expectations. One of the main reasons why the pass-through is not full is that, even when shown the information, subjects still put some weight on their prior beliefs. For example, the subject may be aware of the information already, or may not trust it. For an extended discussion of learning, see Appendix B.7.

Our baseline specification makes some implicit functional form assumptions: the relationship between outcomes and beliefs is linear and symmetrical. This is the simplest possible specification and thus provides a good starting point. This is also the standard specification in the literature on information-provision experiments (Armantier et al., 2016; Cavallo et al., 2017; Bottan and Perez-Truglia, 2017; Fuster et al., 2018; Cullen and Perez-Truglia, 2018). As usual, we can used binned scatterplots to assess whether the baseline specification fits the data well. Figure 6.a shows a binned scatterplot version of the results from column (1) of Table 2. The results indicate that the linear specification is a reasonable approximation. This figure also suggests tha the results are not driven by outliers.

Columns (2) through (4) of Table 2 present the effects on other survey outcomes. Column

(2) is identical to column (1) only that the dependent variable is the expectation 5-years-ahead instead of the expectation 1-year-ahead. The coefficient on *Information Shock* is positive (0.167) and statistically significant (p-value=0.017). This result implies that when a subject receives an information shock, it propagates not only to the more immediate expectations (1-year-ahead) but also to the longer-term expectations (5-years-ahead). While the coefficients from column (1) and (2) are statistically indistinguishable from each other (p-value=0.560), the point estimates do suggest that the information shocks affect more immediate expectations more strongly.

Columns (3) and (4) of Table 2 present some falsification tests. Column (3) is identical to column (1) except that the dependent variable is measured before (instead of after) the information-experiment. Since the information shock has not been administered to the subject yet, it should have no effect whatsoever on the prior expectations. As expected, the coefficient on *Information Shock* from column (3) is close to zero (-0.014), statistically insignificant (p-value=0.837) and precisely estimated. Moreover, we can reject the null hypothesis that the baseline coefficient from column (1) is equal to the falsification coefficient from column (4), with a p-value<0.001.[10]

Column (4) of Table 2 presents the other falsification test. Column (4) is identical to column (1) but the dependent variable is about the stock market expectations instead of the home price expectations. If the information shock, which is specific to home price expectations, spilled over to the stock market expectations, that would suggest that the effects may be due to spurious reasons such as numerical anchoring. On the contrary, the coefficient on *Information Shock* from column (4) is close to zero (0.017), statistically insignificant (p-value=0.899) and precisely estimated.

4.2 Effects on Market Behavior

Next, we look at the effect of the information shocks on actual market choices. The main results are presented in columns (5) through (8) of Table 2, which are estimated with the data from the field experiment. Recall that all regressions in this table correspond to equation (2) from Section 2, but with different dependent variables. In column (5) the dependent variable takes the value 100 if the property has been sold at 12 weeks after the start of letter delivery and 0 otherwise. By that horizon, around 37% of the properties had been sold. Note that the letter could not have an immediate effect – that would only happen if the seller was deciding

[10]This is an equality test between two coefficients based on the same data but different regressions. To allow for a non-zero covariance between these two coefficients, we estimate a system of Seemingly Unrelated Regressions. In the remainder of the paper, when comparing coefficients from the same data but different regressions, we always use this method.

on an offer on the same day of letter delivery. Most likely, it will take weeks since the letter is received until the first offer arrives and thus for the information to be able to affect the decision. The time horizon in column (5) is two weeks after all the letters have been read (according to our proxy). Thus, it it is arguable that the majority of the information shocks should have had enough time to affect market choices by then.

The coefficient on *Information Shock* from column (5) of Table 2 is negative (-0.330) and highly statistically significant (p-value=0.001). This negative sign is consistent with economic theory: i.e., a positive shock to expectations should decrease the probability that the property is sold. This coefficient is also economically large. A 1 pp larger information shock causes a 0.330 percentage point drop in the probability that the property is sold sold within 12 weeks, thus implying a behavioral elasticity of 0.33. This is the elasticity between the sales probability and the information shock. The elasticity between the sales probability and the home price expectation is probably a multiple of that. Intuitively, the elasticity reported above is an intention-to-treat effect, because the information shocks do not fully materialize on changes in expectations. In Section 4.4 below we provide back of the envelope calculations for the treatment-on-the-treated effect.

Again, we can visualize the data to assess whether the baseline specification fits the data well. Figure 6.b presents the binned scatterplot version of the results from column (5) of Table 2. The results indicate that the linear specification is a great approximation to the data. Moreover, this figure indicates that the results are not driven by outliers.

Column (6) of Table 2 is identical to column (5), only that the dependent variable indicates if the property had been sold at 28 weeks, instead of 12 weeks, after letter delivery. This is the longest horizon we can estimate given that we have the administrative data updated up to the end of 2019 (i.e., 28 weeks after the start of the letter delivery). These results would indicate if the effects of the information shocks were short lived (e.g., some owners waited an additional few weeks) or whether they were more persistent. The coefficient on *Information Shock* from column (6) is negative (-0.325) and statistically highly significant (p-value=0.002). Indeed, the coefficient 28 weeks later (-0.325, from column (6)) is almost identical, and statistically indistinguishable from, the corresponding coefficient 12 weeks later (-0.330, from column (5)). These results indicate that the effects of the information shocks were highly persistent, staying equally strong half a year after the start of the letter delivery.

Columns (7) and (8) of Table 2 present some falsification tests. Column (7) is identical to column (1) except that the dependent variable indicates if the property had been sold before the start of the letter delivery (i.e., one week before, rather than 12 weeks after). Since the information shock was not administered to the subject by then, it should have no effect on their market choices. As expected, the coefficient on *Information Shock* from column

(7) is close to zero (0.014), statistically insignificant (p-value=0.469) and precisely estimated. Indeed, we can reject the null hypothesis that the falsification coefficient (0.014, from column (4)) is equal to the baseline coefficient (-0.330, from column (1)) with a p-value<0.001.

Moreover, we can take this analysis even further, by means of an event-study analysis. We can assess if the timing of the effects follows as expected the timing of the delivery of letters. As discussed above, we would expect that the effects of the letters should not materialize right away, but it should start being felt a couple of weeks after the letters were sent. Moreover, we would expect the effects to intensify as more and more letters as being read, and then stabilize weeks after the last letter was read. The event-study results are presented in Figure 7. This figure reproduces column (5) of Table 2 but for every single week since the letters were produced (i.e. 2 weeks before the start of the letter delivery) until the last month of administrative data (i.e., 28 weeks after the start of the letter delivery. To ease the interpretation, we added a figure at the bottom of the event study with the evolution of the letters being opened (a reproduction of Figure 5.a). The evolution of the coefficients over time are exactly as expected. First, the effects are close to zero and statistically insignificant for the two weeks prior to the start of the letter delivery. Second, the effects start to build up at 2 weeks after the start of the letter delivery, and intensify during the period in which the letters were being gradually opened, to finally stabilize a month after the letters were finished being read.

Column (8) of Table 2 presents the other falsification test. While Table 1 shows that the pre-treatment characteristics were balanced across treatment groups, that is not the most relevant comparison given that we exploit treatment heterogeneity. The most direct balance test would consist of reproducing the same regression as in the baseline specification (column (1)) but using the pre-treatment characteristic as dependent variable. In column (8) of Table 2 we present the results with one such pre-treatment characteristic: the logarithm of the number of days that the property had been listed for before our experiment. Since that outcome was determined before the letters were even shipped, the effect should be null. As expected, the coefficient on *Information Shock* from column (8) is close to zero (0.001), statistically insignificant (p-value=0.755) and precisely estimated.

One challenge with the last falsification test is that the falsification variable has a different distribution as the main outcome of interest, thus making the quantitative comparison between the coefficients less direct. Moreover, we would like to reproduce the analysis for other pre-treatment characteristics. These two challenges are addressed in Figure 8. The two estimates to the left correspond to the two post-treatment outcomes shown in columns (5) and (6) of Table 2: whether the property was sold at 12 and 28 weeks after letter delivery. The remaining estimates are based on six different pre-treatment characteristics: the number of days the property had been listed (in logs, as in column (8) of Table 2), the initial listing price (in logs), and the number of bedrooms, bathrooms, and the square feet built and lot size. To facilitate the comparison across different outcomes, we report the standardized coefficients: i.e., the coefficients as proportion of the standard deviation of the corresponding dependent variable. As expected, we find the effects on the pre-treatment outcomes are close to zero, statistically insignificant, precisely estimated, and statistically different from the corresponding effects on the post-treatment outcomes. For example, the standardized coefficients are -0.683 (p-value=0.001) on the sales probability 12 weeks post-treatment and 0.068 (p-value=0.755) on the pre-treatment number of days listed. And the difference between the two is highly statistically significant (p-value=0.023).

4.3 Additional Results

Table 3 presents some additional robustness checks. Column (1) is the baseline specification, identical to column (5) from Table 2. The specification from column (2) is identical to that of column (1), only that it includes some additional control variables: the logarithm of the days the property was on the market prior to the experiment, the logarithm of the initial listing price, four dummies for number of beds, four dummies for number of bedrooms, the logarithm of square footage built, the logarithm of lot size, and six state dummies. Note that the R^2 goes up substantially, from 0.034 in column (1) to 0.111 in column (2), meaning that the control variables have quite a bit of explanatory power. Since the treatment is randomized, controlling for additional variables should not make a difference. Indeed, the point estimate (-0.325, from column (2)) is almost identical to the baseline coefficient (-0.330, from column (1)), with their difference being statistically insignificant (p-value=0.990).

In the baseline specification, we include a single indicator variable that takes the value 1 if the signal was shown to the subject. Since there are different types of information, we may want to control for each type of information separately. The specification from column (3) of Table 2 is identical to that of column (1), with the only exception that instead of controlling for one treatment indicator, it controls for a set of five treatment indicators (i.e., one for each of the five treatments that are not Baseline). The results indicate that this is an irrelevant detail of the specification. These additional controls do not add any explanatory power to the equation, as the R^2 is 0.034 in column (3) is identical to that of column (1). Moreover, the coefficient from this extended specification (-0.325, from column (3)) is almost identical to the baseline coefficient (-0.330, from column (1)), with difference being statistically insignificant (p-value=0.999).

Recall from Section 2 that the research design exploits two sources of heterogeneity in signals: across subjects and across sources. Column (4) and (5) of Table 2 provides estimates that rely on the two sources of variation separately, to assess whether the results are driven by one specific source of variation, or whether they are robust to both identification strategies. Column (4) provides an estimate that only exploits the variation across individuals. To do this, as discussed in Section 2, column (4) uses an specification that is identical to that of column (1) but controlling for the signal chosen for the individual, without the treatment interaction. Column (5) provides an estimate that only exploits the variation across sources, by dropping the observations corresponding to the baseline group. The results from Table 2 indicate that the results are robust across the two identification strategies. The coefficients are similar in the baseline specification (-0.330, from column (1)), in the specification that only uses the disclosure-randomization (-0.286, from column (4)), and in the specification that uses the source-randomization (-0.330, from column (5)). These three coefficients are statistically significant on their own (p-values of 0.001, 0.049 and 0.005), and statistically indistinguishable from each other.

Column (5) of Table 2 shows what happens when we drop the Baseline group. It is natural to wonder what would happen if we were to drop each of the other five treatment groups. Those results are presented in columns (6) through (10). The results indicate that the coefficient is not driven by any single one of the five treatment groups. The five different coefficients (-0.338, -0.257, -0.325, -0.382 and -0.320, from columns (6) through (10) respectively) are consistent in magnitude to the corresponding baseline coefficient (-0.330, from column (1)), and also robust in terms of statistical significance.

Some additional robustness checks are presented in the Appendix. It is possible that the two information sources had different effects. For example, subjects may have backward-looking expectations and thus be more prone to incorporating the information about the past. Subjects may be less likely to react to the information on forecasts than to the information about the past because of the lower willingness to trust information produced by researchers (i.e., our in-house forecasts) than information produced by a company they know about (i.e., the historical median values produced by Zillow). As stated in the pre-registration, we wanted to focus on the pooled estimates to maximize statistical power. However, whether the results are driven by one specific type of information may be informative for researchers planning to use this same experimental design. Appendix B.9 shows that the information about the past was more effective than the information about forecasts.

In the previous analysis, the outcome variable is the probability of sales at a given point in time. An alternative outcome variable would be the time until sold (i.e., the number of days elapsed between the start of the letter delivery and the sale of the property). As we anticipated in the pre-registration, a problem with this outcome is that it is censored in nature: for properties that have not been sold yet (a whooping 42% of the sample), we do not know if they would be sold 1 week or 1,000 weeks later. By looking at the probability of selling at a given point in time, we avoid this censoring problem. In any case, as reported in Appendix B.8, the results are qualitatively and quantitatively robust if we estimate the effects on this censored outcome using the standard duration models.

The data that we scrapped contains other information that can be used to construct secondary outcomes. The effects on these secondary outcomes are presented in the Appendix. One outcome that we listed in the pre-registration was the sales price. As anticipated in the pre-registration, however, the challenge with this outcome is that it is censored: we do not observe the sales price for the 42% of the sample of the properties that were not sold. In Appendix B.10, we that the results are quite sensitive to different assumptions related to the censoring bias. In addition to sales price, our scrapped data includes other weekly outcomes related to the listings, such as whether the property remains listed on the website we scrapped or whether there have been changes to the listing price. The results for these outcomes are presented in Appendix B.11.

4.4 Magnitude of the Effects

In this section, we discuss the interpretation of the main results. The object of interest is the causal effect of home price expectations on the speed that the property is sold. From now on we measure sales speed as the probability that the home is sold at 28 weeks after the start of letter delivery – in any case, as discussed above, the results would be similar for other time horizons. Our starting point is the effect of the information shocks on behavior. A 1 pp higher information shock decreases the sales probability by 0.325 pp, for a behavioral elasticity of -0.325. However, this elasticity is an intention-to-treat estimate, for three reasons explained below.

First, conditional on reading the letter, the information shock introduced in the letter probably does not materialize in the reader's expectations. For example, the readers may know the information shown to them, may not entirely trust the source of the information, or may have a lot of confidence in their prior belief. In Section 4.1 we used the auxiliary survey to estimate the rate of pass-through from information shocks to expectations: a 1 pp higher information shock increases home price expectations by 0.205 pp. The results from this auxiliary sample should be taken with a grain of salt, for example, because the incentives to pay attention to information about home prices may be much lower relative to the sample used for the field experiment. However, if we are willing to extrapolate the results from the auxiliary survey, then these estimates would imply that a 1 pp increase in home price expectations causes a decline of 1.59 pp $=(\frac{0.325}{0.205})$ in the sales probability – i.e., a behavioral elasticity between expectations and sales probability of 1.59.

The second source of attenuation bias is related to the fact that some individuals may have not read the letters. Some letters may not have even reached the owners, for example, because the letters were lost. Some recipients may have received the letters but threw them away without reading them. Some letters may have arrived too late (e.g., after the property had been already sold) and thus had no chance of affecting the decision. Thus, we need an estimate of the reading rate (i.e., the share of recipients that actually read the letter on time). Then we could divide the intention-to-treat effects by the reading rate to scale it up to the treatment-on-the-treated, as in Perez-Truglia and Cruces (2017). We do not have a direct estimate of the reading rate, but we can come up with some back of the envelope calculations.

According to the U.S. Monitor Non-Profit Standard Mail Delivery Study, around 5% of standard non-profit mailers fail to be delivered (U.S. Monitor, 2014). Based on data from the US Postal Service Household Diary Survey (Mazzone and Rehman, 2019), we estimate that, conditional on delivery, around 26% of our letters were not read by the recipient.[11] And based on the timing of survey responses and the timing of sales in the baseline group, we estimate that roughly 7.7% of the letters were read after the property had already been sold. These three estimates combined imply a reading rate of 64.9% ($= 0.95 \cdot 0.74 \cdot 0.923$). As a result, to account for this source of attenuation bias, the scaled-up behavioral elasticity between expectations and sales probability would boil down to 2.45 ($= \frac{1.59}{0.649}$).[12] We believe the previous estimate of the reading rate is conservative, in the sense that the reading rate may be a bit smaller under alternative assumptions and thus would lead to an even larger scale-up factor (see Appendix B.12 for more details about these calculations and alternative estimate). The elasticity of 2.45 reported above suggests that subjective expectations have the potential to have sizeable effects on market choices. For example, an increase of one standard deviation in expectations (5.39 pp, according to the survey expectations in our auxiliary sample) would reduce the sales probability by 13.21 pp ($= 2.45 \cdot 5.39$).[13]

Finally, while the above estimates pertain to the population as a whole, they may mask meaningful heterogeneity. For some individuals, the elasticity between expectations and market choices may be smaller, while it may be much larger for others. More precisely, there are two factors that may curtail the importance of expectations in this specific context. First, some individuals may want to change their behavior due to their renewed expectations but may be unable to do so due to constraints. For instance, some individuals may need to sell

[11] This 26% figure is based on the 2018 HDS Recruitment Sample and corresponds to the estimate of treatment of advertising mail reported in Figure 5.3 of (Mazzone and Rehman, 2019).

[12] See Appendix B.12 for details.

[13] This 5.39 is the raw standard deviation in prior beliefs for the one-year-ahead home price expectation. The results are similar (5.51 instead of 5.39) if we use instead the standard deviation within a given property type (i.e., same ZIP code and number of beds).

the home immediately because they already bought another property to which they need to move shortly. The second factor that may hinder the importance of expectations is that some individuals may care about their expectations but not specifically the expectation that is manipulated in our experiment. Consider an individual who, after selling the property, is planning to (or already did) buy another property in the same neighborhood. In that case, the expectation about the appreciation in that neighborhood should be largely irrelevant for the decision to sell the property, because the subject will be affected by the same appreciation regardless of whether the property is sold or not or regardless of when the property is sold.

To assess how important expectations can be, we would like to estimate the model on a subset of the subject pool who did not face the two frictions described above. While we do not observe those frictions directly, we can provide a rough approximation by looking at the (pre-registered) heterogeneity of effects between owner-occupied (78.4% of subjects) versus non-owner occupied (21.6% of subjects). We argue that, relative to the owner-occupied, the non owner-occupied are less subject to these two frictions.

First, the non owner-occupied have arguably more flexibility in whether to sell the home or not and when to sell it. The decision to sell is very consequential for the daily lives of the owner-occupied, because they have to move out of the property after selling it. They may be selling the property due to time-sensitive reasons such as having to move because of work or school, because they have already bought another home, or because of marriage, divorce or the birth of a child. The non owner-occupied, on the contrary, do not need to move after selling the property. On the contrary, their decision of whether to sold and when to sold respond to purely financial incentives.

Second, the non owner-occupied should arguably care more about the expectation of appreciation in the specific neighborhood of the property. After they sell their current property, a significant fraction of the owner-occupied will buy another property in the same or similar market (results presented in Appendix B.13). For that reason, they will be affected by the same home price appreciation regardless of whether they sell the property or when they sell it. The non owner-occupied, on the contrary, have the option to take whatever they get from selling the property and invest it in a housing market that they expect to appreciate more, or to invest the funds outside of real estate.

While the owner and non-owner occupied differ markedly in the two aspects described above, they are otherwise quite similar. For example, they are quite similar in observable characteristics such as age, ethnic composition, income and education. And while there are some differences, the timing of sales outcomes does not look dramatically different: e.g., the probability of selling the property 28 weeks post-treatment was 52.89 pp for non-owner occupied versus 57.99 pp for owner-occupied (for more details, see Appendix B.13).

Figure B.6 presents the results. This figure shows the event-study analysis of the effects of the information shocks broken down by the type of owner-occupied status. We find that the effects of information shocks are qualitatively similar for the owner-occupied (depicted in blue circles) and non owner-occupied (depicted in red diamonds). However, and consistent with the prediction discussed above, the effects are almost three times as strong for the non owner-occupied as for the owner-occupied. For example, at 12 weeks after the start of the letter delivery, the coefficient on the information shock is -0.677 (p-value<0.001) for the non owner-occupied and -0.209 (p-value=0.082) for the owner-occupied. Moreover, the difference between these two coefficients is statistically significant (p-value=0.046). The results are qualitatively robust in other time horizons. For instance, at 28 weeks after the start of the letter delivery, the coefficient on the information shock is -0.637 (p-value=0.003) for the non owner-occupied and -0.225 (p-value=0.066) for the owner-occupied, and their difference is statistically significant too (p-value=0.095).

We can take the estimate for the non owner-occupied (-0.637) and apply the same corrections for pass-through rate (0.205) and reading rate (0.649) used above. These estimates imply that, for the non owner-occupied, the elasticity between expectations and market choices boils down to -4.79 pp $=(\frac{-0.637}{0.205 \cdot 0.649})$: i.e., a 1 pp increase in home price expectations causes a decline of 4.79 pp in the sales probability. This estimate shows that, in absence of frictions, home price expectations can have a large effect on market choices. Additionally, the above evidence provides support for the widespread view that non-owner occupied properties have a disproportionate influence on speculation in the housing market (Gao et al., 2020).

While our favorite interpretation for the heterogeneity by owner-occupied status is based on the two frictions discussed above, however, we cannot rule out alternative interpretations. Perhaps non owner-occupied are more responsive because they are different in other respects, like financially more sophisticated and thus more likely to understand the information and use it correctly. However, the fact that these two groups look so similar in observable characteristics, including income and education, constitute evidence against this interpretation. Or perhaps non-owner occupied were more likely to incorporate the information shocks into their expectation. For example, they may be less informed about the local home prices because they do not live in the area.

5 Conclusions

We provided the first experimental evidence on the causal effects of home price expectations on market choices. We conducted a large-scale, high-stakes field experiment involving 57,910 U.S. individuals who had recently listed homes on the market. We sent them letters and randomized the information included in those letters to generate exogenous shocks to their expectations home price expectations. We then used the administrative records to measure the effects of those shocks on the recipients' decision to sell their properties. Consistent with economic theory, higher home price expectations caused the owners to hold on to their properties for longer. Our conservative estimate for the average treatment effect suggests that market choices are highly elastic to expectations: a 1 percentage point increase in home price expectations reduces the probability of selling within 6 months by 2.45 percentage points.

Our estimates address an open question that is central for the literature on macroeconomics and finance, on the causal effect of expectations on market choices. Additionally, we believe our experimental framework can be adopted extensively by other researchers to study other questions from a range of fields such as macroeconomics, urban economics, finance, real estate economics and behavioral economics. The information provided in the mailings does not need to be related to home prices. Instead, the mailings can inform about any characteristics of the market that may be relevant for the decision of whether to sell the property, or related decisions as to whether to change the listing price.

There are some experimental methods that have been broadly adopted to study a broad range of hypotheses. Just to mention some examples, the audit study pioneered in (Bertrand and Mullainathan, 2004) has been used by countless studies to test hypotheses in labor economics; and the mailing study pioneered by (Slemrod et al., 2001) has been used extensively to study tax compliance. Our experimental framework has a number of advantages that could warrant widespread adoption too. It provides causal inference based on randomization, the effects on behavior are measured with objective data from administrative records, in a naturally-occurring context and based on high-stakes choices. Our framework has practical advantages too. Following our detailed instructions, the experiment can be implemented relatively quickly, with a few weeks of preparation and with results ready in a matter of few months. The experiment is relatively cheap, costing less than $0.25 per subject. Last, the experiment can be implemented at massive scales, with potentially up to 1 million subjects at a time.[14]

References

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

[14]For example, according to Zillow Research (https://www.zillow.com/research/data/) there were 1,705,251 listings in the the month of our experiment (May 2019).

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

Armona, L., A. Fuster, and B. Zafar (2019). Home Price Expectations and Behaviour: Evidence from a Randomized Information Experiment. *Review of Economic Studies 86*(4), 1371–1410.

Bailey, M., R. Cao, T. Kuchler, and J. Stroebel (2018). The Economic Effects of Social Networks: Evidence from the Housing Market. *Journal of Political Economy 126*(6), 2224–2276.

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

Bertrand, M. and S. Mullainathan (2004). Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination. *American Economic Review 94*(4), 991–1013.

Bottan, N. and R. Perez-Truglia (2017). Choosing Your Pond: Location Choices and Relative Income. *NBER Working Paper No. 23615*.

Brooks, D. (2017). The Home Buying Decision. *New York Times, January 6, 2017*.

Carroll, C. (2003). Macroeconomic Expectations of Households and Professional Forecasters. *Quarterly Journal of Economics 118*(1), 269–298.

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

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

Coibion, O., Y. Gorodnichenko, and S. Kumar (2018). How Do Firms Form Their Expectations? New Survey Evidence. *American Economic Review 108*(9), 2671–2713.

Crump, M. J. C., J. V. McDonnell, and T. M. Gureckis (2013). Evaluating Amazon's Mechanical Turk as a Tool for Experimental Behavioral Research. *PLOS ONE 8*(3), 1–18.

Cullen, Z. and R. Perez-Truglia (2018). How Much Does Your Boss Make? The Effects of Salary Comparisons. *NBER Working Paper No. 24841*.

Fuster, A., R. Perez-Truglia, M. Wiederholt, and B. Zafar (2018). Expectations with Endogenous Information Acquisition: An Experimental Investigation. *NBER Working Paper No. 24767*.

Gao, Z., M. Sockin, and W. Xiong (2020). Economic Consequences of Housing Speculation. *The Review of Financial Studies, forthcoming*.

Gennaioli, N. and A. Shleifer (2018). *A Crisis of Beliefs: Investor Psychology and Financial Fragility.* New Jersey: Princeton University Press.

Glaeser, E. L. and C. G. Nathanson (2015). Housing Bubbles. In G. Duranton, J. V. Henderson, and W. C. Strange (Eds.), *Handbook of Regional and Urban Economics*, Volume 5, Chapter 11, pp. 701–751. Elsevier.

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

Mazzone, J. and S. Rehman (2019). The Household Diary Study Mail Use and Attitudes in FY 2018. *Retrieved March 28, 2020, from https://www.prc.gov/dockets/document/109368*.

Perez-Truglia, R. and G. Cruces (2017). Partisan interactions: Evidence from a field experiment in the United States. *Journal of Political Economy 125*(4), 1208–1243.

Perez-Truglia, R. and U. Troiano (2018). Shaming Tax Delinquents. *Journal of Public Economics 167*, 120–137.

Shiller, R. (2005). *Irrational Exuberance.* New Jersey: Princeton University Press.

Slemrod, J., M. Blumenthal, and C. Christian (2001). Taxpayer response to an increased probability of audit: evidence from a controlled experiment in Minnesota. *Journal of Public Economics 79*(3), 455–483.

U.S. Monitor (2014). 7 Myths of Direct Mailing. *Retrieved March 28, 2020, from https://www.targetmarketingmag.com/promo/7MythsofDM.pdf*.

Figure 1: Sample Envelope

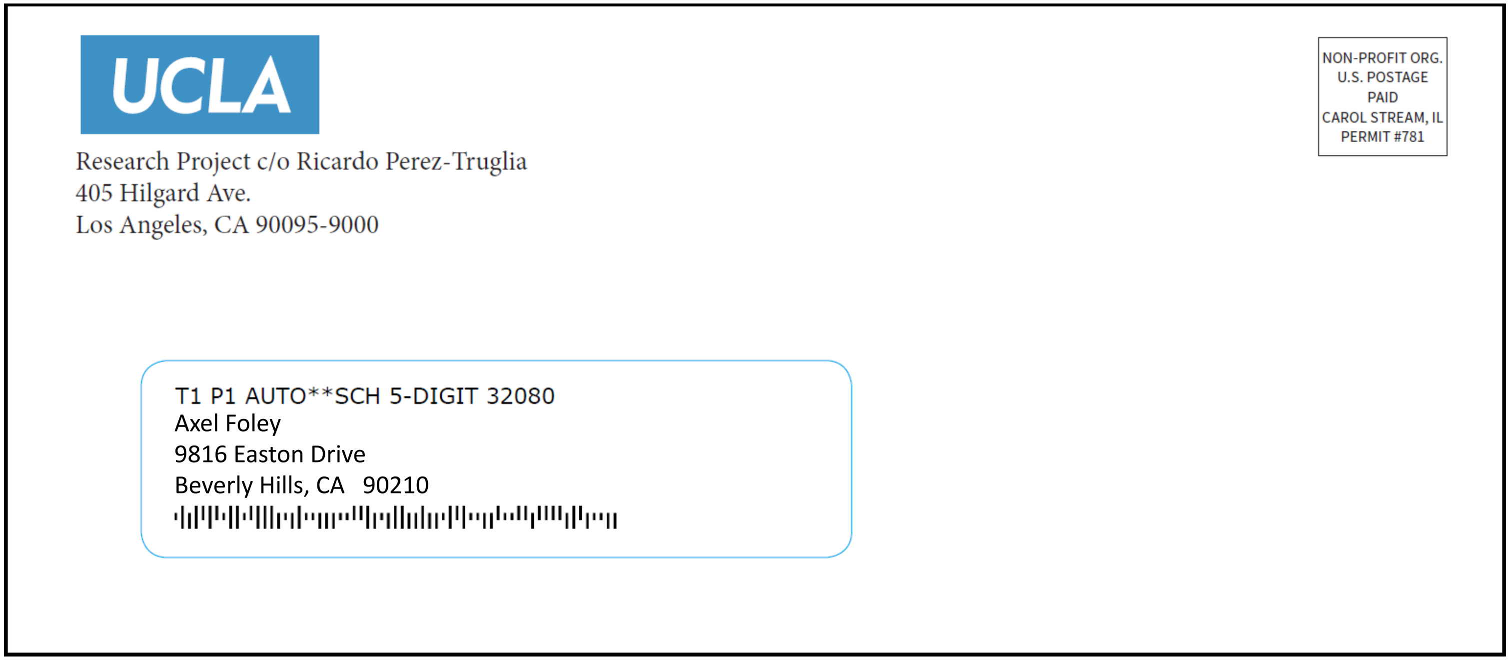

Notes: Screenshot of the outside of the envelope used in the field experiment.

Figure 2: Sample Letter

a. First Page

UCLA

Los Angeles, May 31st 2019

Dear **Axel Foley**,

We are researchers at UCLA and we are reaching out to you as part of a research study about decision making of homeowners.

According to our records, you may be considering selling a property. We know these decisions can be difficult, so we want to share some information that we hope can be helpful:

<<INFORMATION>>

If you would like to help us with our study, we kindly ask you fill out the following 2-minute survey:

Visit **www.surveyhousing.com** and enter validation code

Participation is voluntary and responses are 100% confidential. The results of this study can provide valuable insights to homeowners across the country. Your participation in the survey is greatly appreciated.

110 Westwood Plaza, Suite C515
Los Angeles, CA 90095-1481

Website: http://www.anderson.ucla.edu/housingstudy

Please recycle

b. Second Page

Your household was randomly chosen to receive this letter. *We will not send you any more letters in the future.*

If you have any questions about the study, you can find contact information on our website: www.anderson.ucla.edu/housingstudy.

Thank you for your attention!

Ricardo Perez-Truglia
Assistant Professor of Economics
University of California, Los Angeles

Nicolas Bottan
Post-Doctoral Associate
Cornell University

If you have questions about your rights as a research subject, or you have concerns or suggestions and you want to talk to someone other than the researchers, you may contact the UCLA Office of the Human Research Protection Program by phone: (310) 206-2040; by email: participants@research.ucla.edu or by mail: Box 951406, Los Angeles, CA 90095-1406.

Methodological Notes:

<<INFORMATION DETAILS>>

T1 P1 AUTO**SCH 5-DIGIT 32080
Axel Foley
9816 Easton Drive
Beverly Hills, CA 90210

Notes: Screenshot of the letter used in the field experiment. The two placeholders (marked as *«Information»* and *«Information Details»*) indicate the placement of the two components that were randomly allocated. Their samples, by treatment group, are presented in Figures 3 and B.1 respectively. Appendix C shows a sample of the final product.

Figure 3: Sample Information Tables

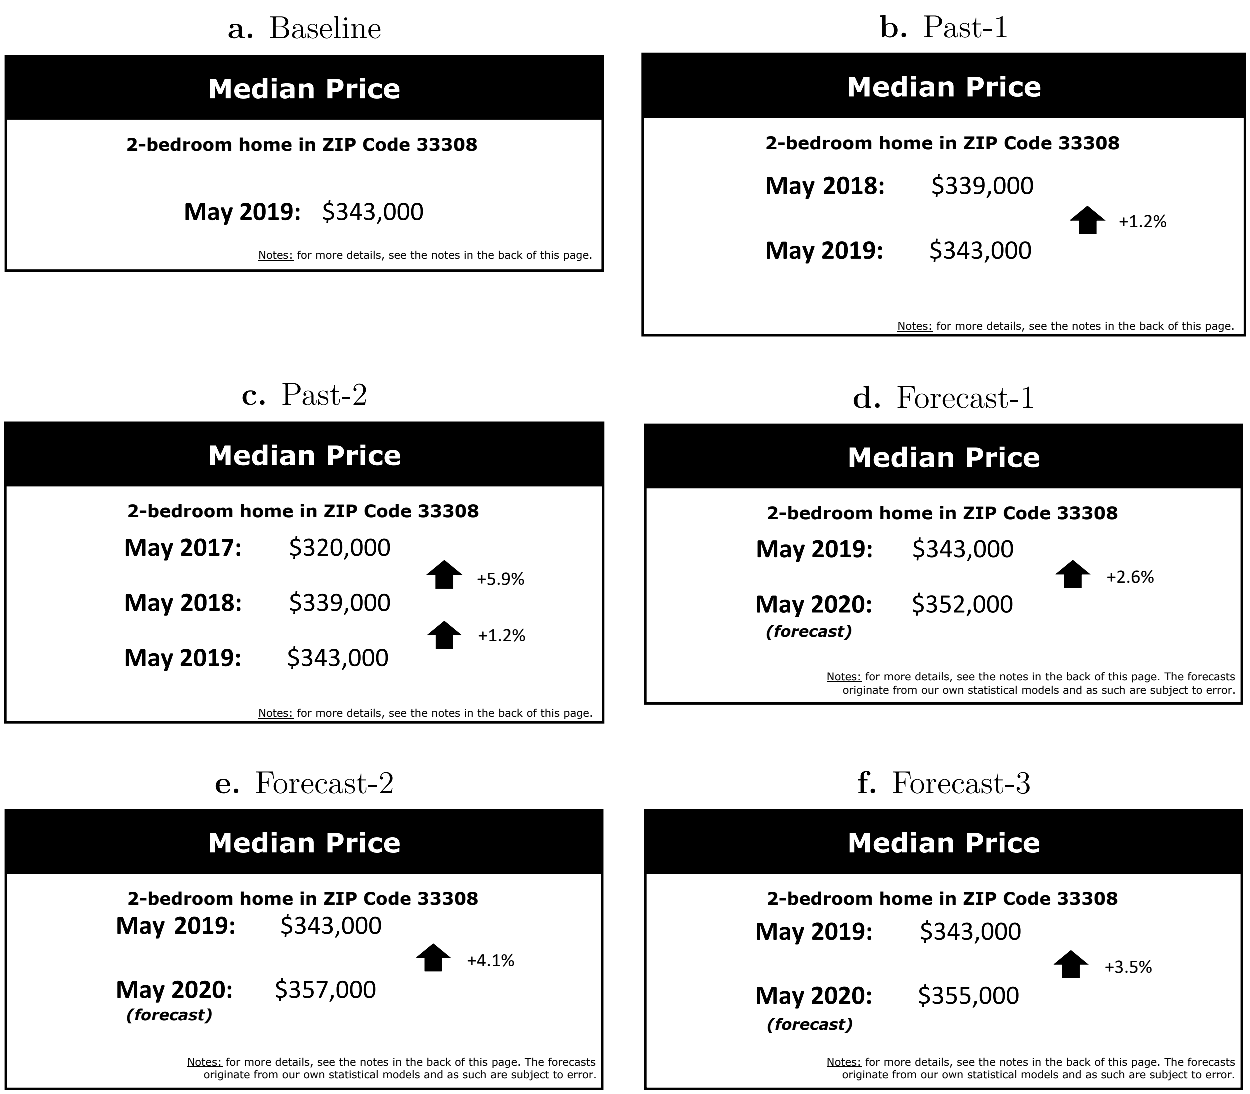

Notes: Each panel corresponds to the hypothetical table that a given individual would receive under the different treatment groups. The table is then placed in the middle of the first page of the letter, in the location of the placeholder *«Information»* from Figure 2. See Figure B.1 for the methodological notes accompanying each table.

Figure 4: Heterogeneity in Signals Within and Across Information Sources

a. Past-1

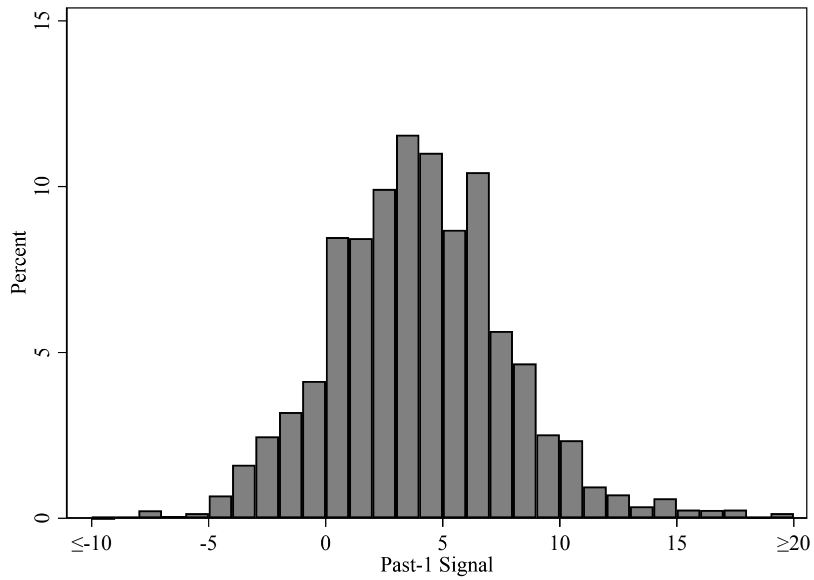

b. Past-1 vs. Past-2

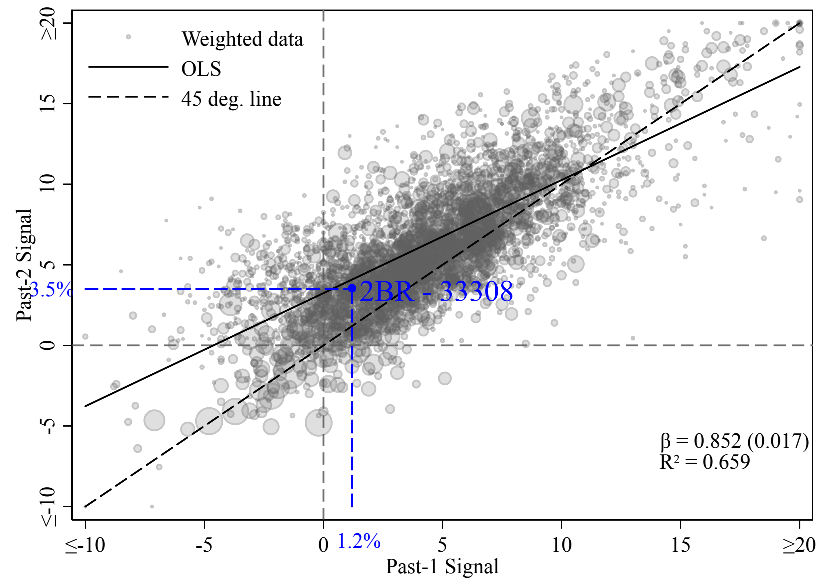

Notes: Panel (a) shows the distribution of the signals that the 57,910 subjects would have received if they had been assigned to the Past-1 treatment (i.e., the annual growth rate over the past year). The bins have a width of 1 percentage point and are truncated at -10% and +20%. Panel (b) is a scatterplot showing the relationship between the signals that the 57,910 subjects would have received if they had been assigned to the Past-1 treatment (i.e., the annual growth rate over the past year) versus the Past-2 treatment (i.e., the annual growth rate over the past two years). The size of the circles are proportional to the number of observations, and the signals are truncated at -10% and +20%.

Figure 5: Timing of Read-Receipt and Property Sales

a. Read-Receipt

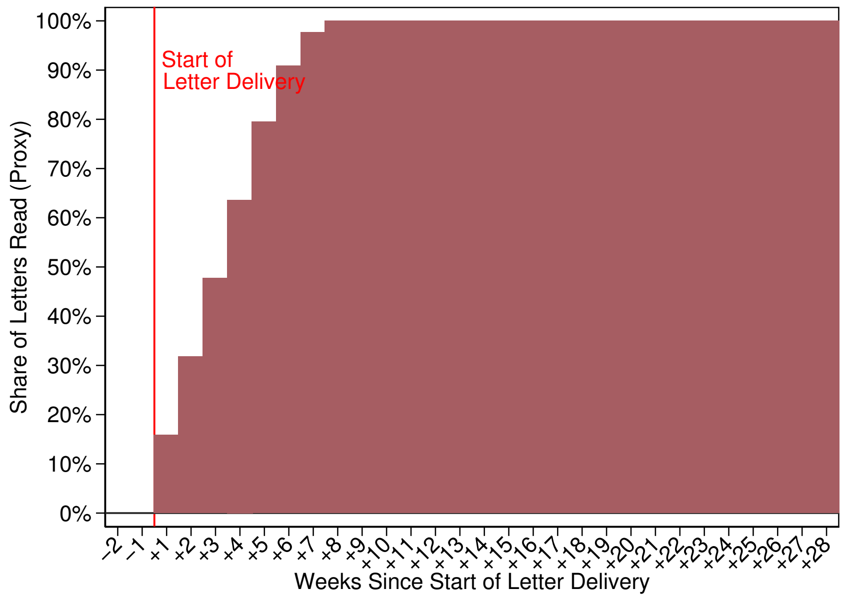

b. Property Sales

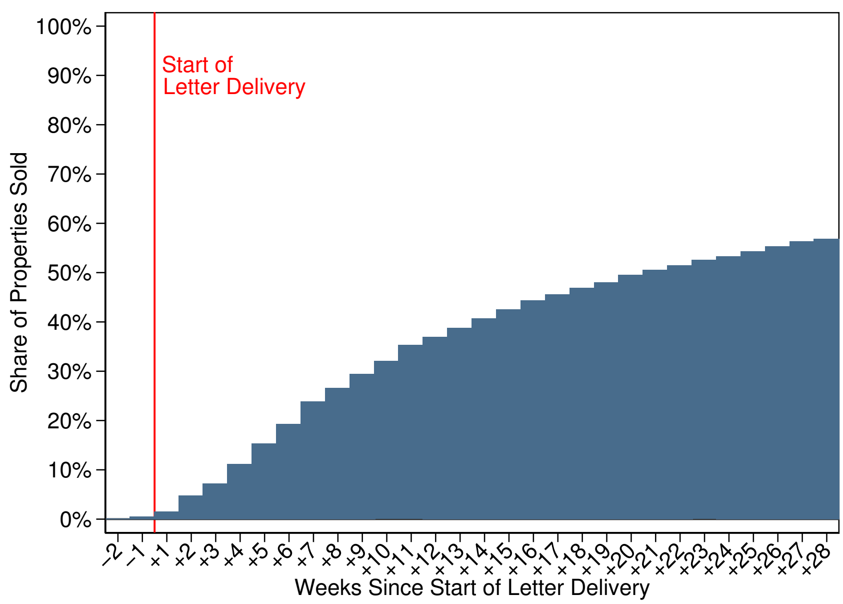

Notes: The red line indicates the estimated delivery date for the first letter (June 15 2019). Panel (a) shows the evolution of the responses to the online survey included in the letter. These dates constitute our read-receipt: i.e., our proxy for the dates when the letters were actually read. Panel (b) shows the fraction of the properties in the subject pool that were sold at each point in time, according to the administrative records.

Figure 6: Effects of Information Shocks on Expectations and Behavior: Binned Scatterplots

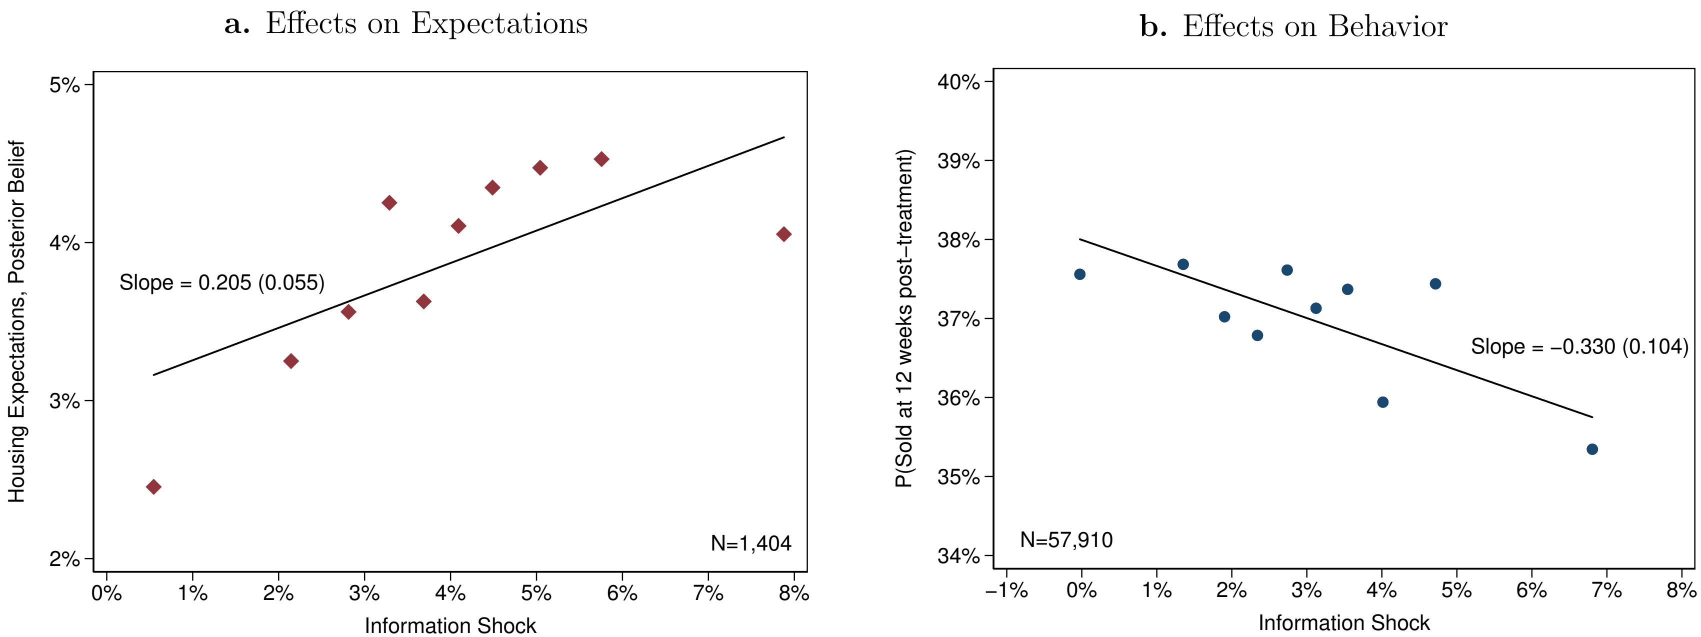

Notes: Panel (a) corresponds to a regression given by equation (2) from Section 2. This binned scatterplot focuses on the key independent variable, *Information Shock* ($E_i^{signal(j^*)} \cdot T_i$). Results are based on 1,404 subjects from the AMT supplemental survey. The dependent variable is the posterior belief (i.e., elicited after the information-provision experiment) about the expected growth rate of the median home value over the following year. Panel (b) corresponds to a regression given by equation (2). This binned scatterplot focuses on the key independent variable, *Information Shock* ($E_i^{signal(j^*)} \cdot T_i$). Results are based on 57,910 subjects from the field experiment. The dependent variable is an indicator variable taking the value 100 if the property was sold at 12 weeks after the start of the letter delivery and 0 otherwise. Each panel reports the slope with heteroskedasticity-robust standard errors in parentheses.

Figure 7: Effects of Information Shocks on Behavior: Event-Study Analysis

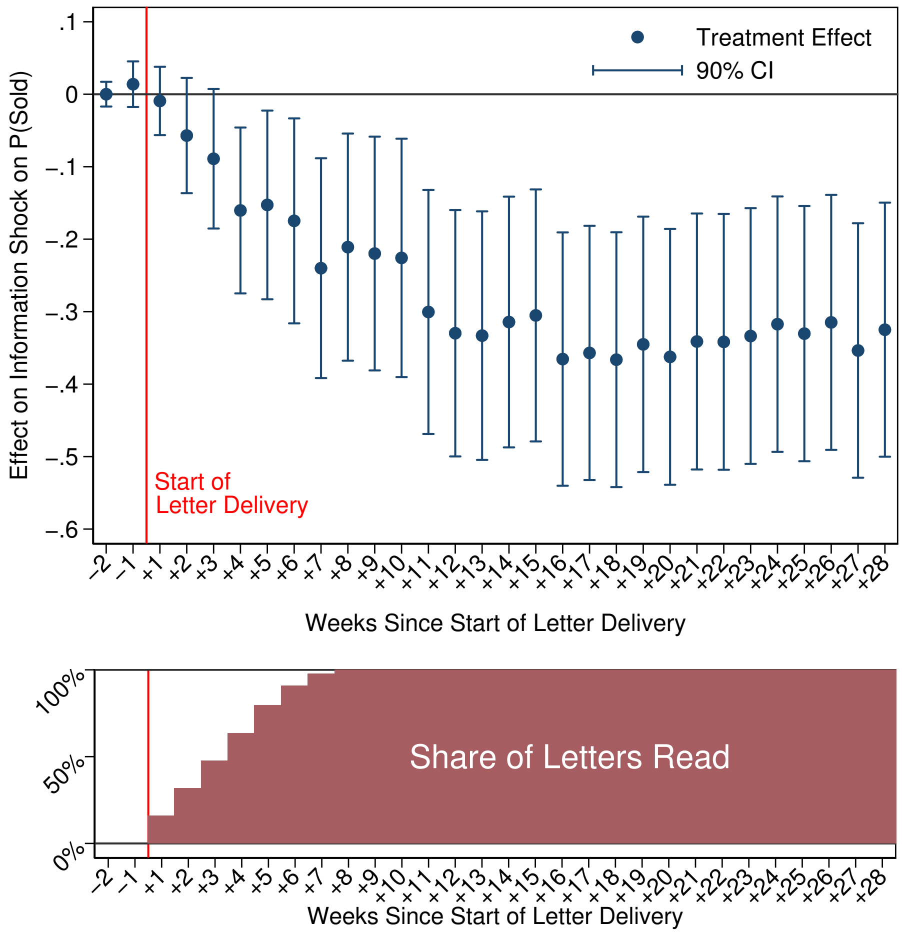

Notes: Each coefficient corresponds to a separate regression based on 57,910 subjects from the field experiment. Every regression corresponds to equation (2) from Section 2, and the coefficient being graphed corresponds to the coefficient on the key independent variable, *Information Shock* ($E_i^{signal(j^*)} \cdot T_i$). All regressions are identical except for the dependent variable. The x-axis indicates the dependent variable used, which is always an indicator variable that takes the value 100 if the property has been sold at a number of weeks after the start of the letter delivery and 0 otherwise. For example, the coefficient on *+12 weeks* is based on a dependent variable that takes the value 100 if the property was sold at 12 weeks after the start of the letter delivery. The red line indicates the estimated delivery date for the first letter (June 15 2019). The smaller figure at the bottom shows the proportion of the letters from the field experiment that had been read at every point in time according to our proxy (the responses to the online survey included in the letter). The 90% confidence intervals are based on heteroskedasticity-robust standard errors.

Figure 8: Effects of Information Shocks on Behavior: Placebo Outcomes

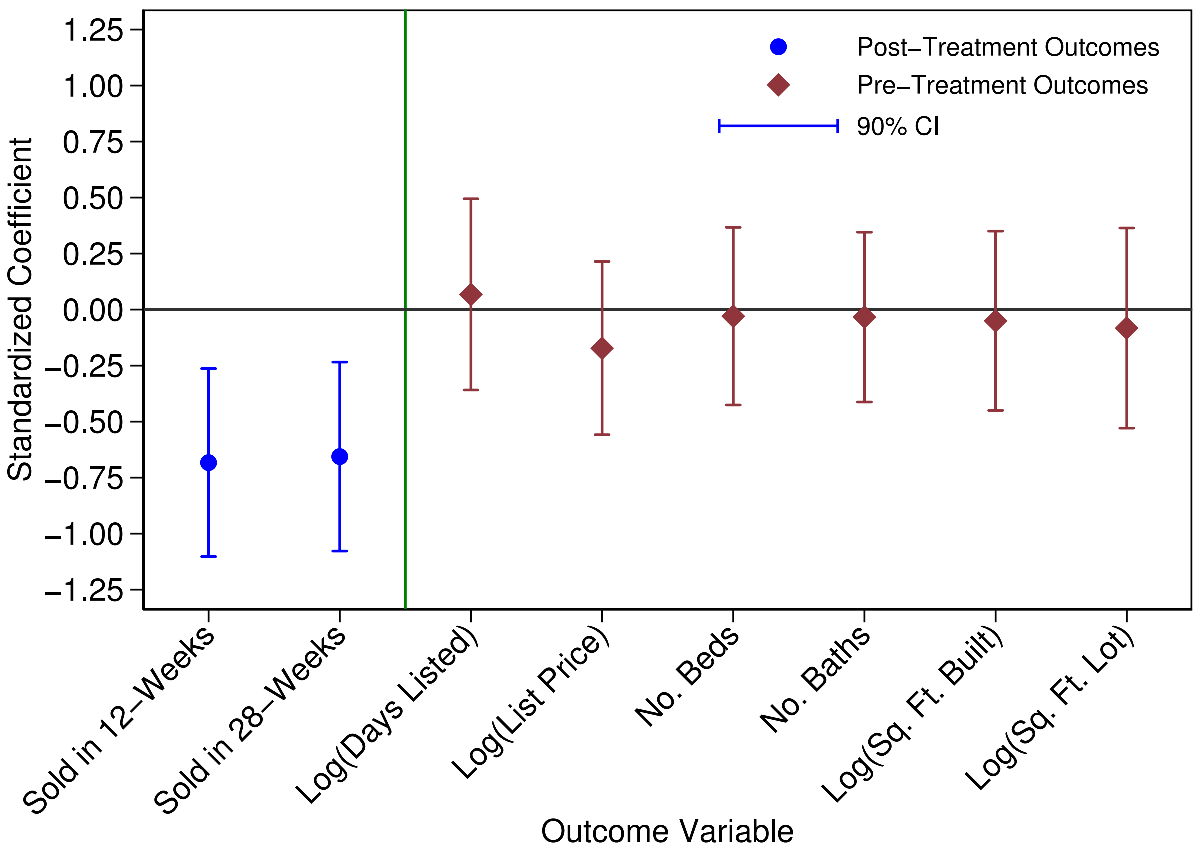

Notes: All the regressions shown in this table correspond to equation (2) from Section 2, based on data on the 57,910 subjects in the field experiment. Each coefficient corresponds to a separate but identical regressions, with the only difference being the dependent variables. We report the coefficient on the key independent variable, *Information Shock* ($E_i^{signal(j^*)} \cdot T_i$). All the coefficients have been normalized in standard deviations of the respective outcome variable. Each dependent variable is listed in the x-axis. We use blue circles to denote the post-treatment outcomes (i.e., that were determined after the start of letter delivery): *Sold in 12-Weeks* is an indicator variable that takes the value 100 if the property was sold at 12 weeks after the start of the letter delivery; and *Sold in 28-Weeks* is an indicator variable that takes the value 100 if the property was sold at 28 weeks after the start of the letter delivery. We use red circles to denote the pre-treatment outcomes (i.e., that were determined before the start of letter delivery): *Log(Days Listed)* is the logarithm of the number of days that the property had been listed for before our experiment; *Log(Listing Price)* is the logarithm of the original listing price of the property; *No. Beds* is the property's number of bedrooms; *No. Baths* is the property's number of bathrooms; *Log(Sq. Ft. Built)* is the logarithm of the property's built area in square feet; *Log(Sq. Ft. Lot)* is logarithm of the property's lot size in square feet. The 90% confidence intervals are based on heteroskedasticity-robust standard errors.

Figure 9: Heterogeneity of Behavioral Effects by Owner-Occupied Status: Event-Study Analysis

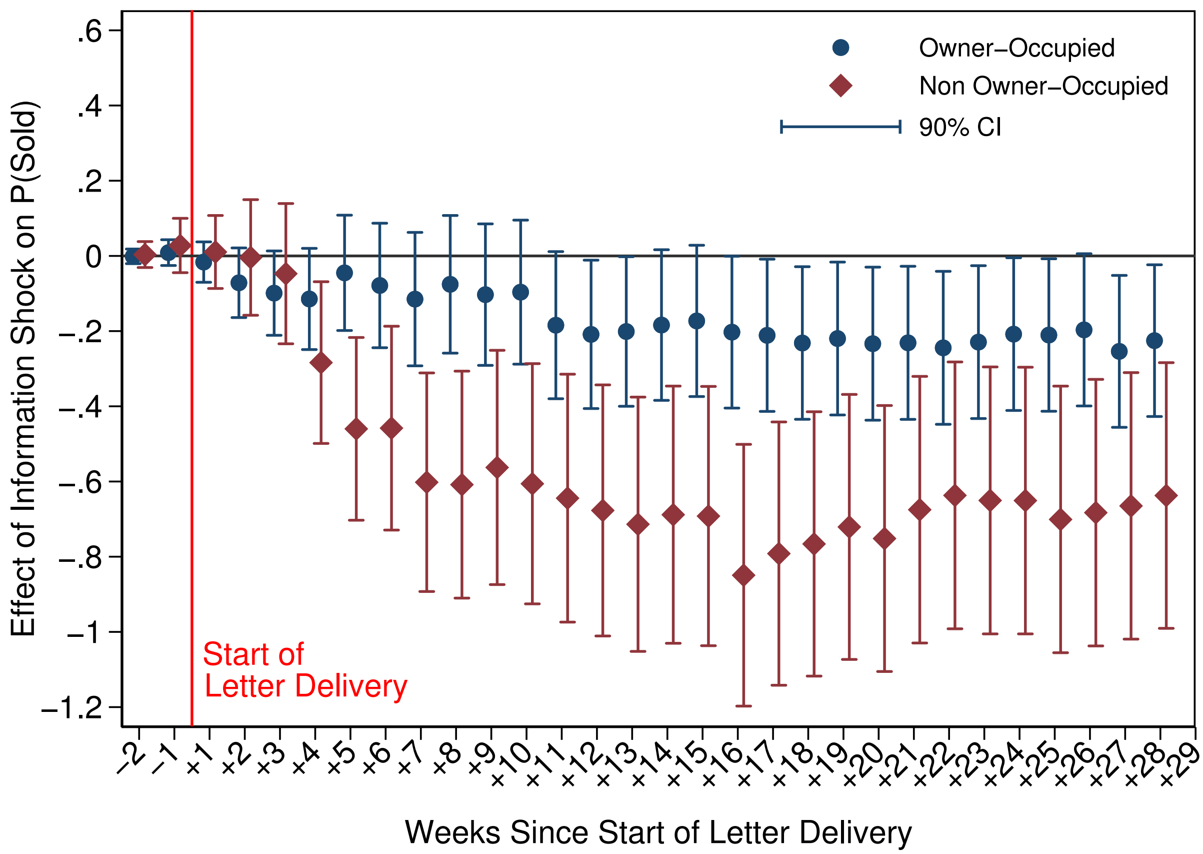

Notes: Each coefficient corresponds to a separate regression. Every regression corresponds to equation (2) from Section 2, and the coefficient being graphed corresponds to the coefficient on the key independent variable, *Information Shock* ($E_i^{signal(j^*)} \cdot T_i$). All regressions are identical except for two features: the dependent variable and the sample. Each of the blue circles are based on a regression with the 45,405 subjects from a field experiment who were living on the property while the property was listed for sale. Each of the red diamonds are based on a regression with the 12,505 subjects from the field experiment who were not living on the property while the property was listed for sale. The x-axis indicates the dependent variable used, which is always an indicator variable that takes the value 100 if the property has been sold at a number of weeks after the start of the letter delivery and 0 otherwise. For example, the coefficient on *+12 weeks* is based on a dependent variable that takes the value 100 if the property was sold at 12 weeks after the start of the letter delivery. The red line indicates the estimated delivery date for the first letter (June 15 2019). The 90% confidence intervals are based on heteroskedasticity-robust standard errors.

Table 1: Descriptive Statistics and Randomization Balance Test

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
		By Treatment Group						
	All	Baseline	Past-1	Past-2	Forecast-1	Forecast-2	Forecast-3	P-value
Days Listed	86.654	85.997	87.212	87.029	85.936	86.208	87.778	0.829
	(0.477)	(1.017)	(1.298)	(1.253)	(1.093)	(1.203)	(1.191)	
List Price ($1,000s)	574.756	575.774	586.787	559.192	574.306	586.835	565.092	0.303
	(3.914)	(8.517)	(11.445)	(9.516)	(8.147)	(10.937)	(9.081)	
No. Beds	3.256	3.249	3.259	3.243	3.254	3.269	3.264	0.582
	(0.005)	(0.010)	(0.012)	(0.012)	(0.011)	(0.011)	(0.011)	
No. baths	2.608	2.607	2.619	2.599	2.608	2.617	2.600	0.678
	(0.004)	(0.010)	(0.011)	(0.011)	(0.010)	(0.011)	(0.010)	
Sq. Ft. Built (1,000s)	2.295	2.292	2.304	2.292	2.287	2.308	2.288	0.820
	(0.005)	(0.012)	(0.014)	(0.013)	(0.012)	(0.013)	(0.013)	
Sq. Ft. Lot (1,000s)	12.958	12.992	12.666	12.934	13.223	13.163	12.730	0.389
	(0.089)	(0.199)	(0.222)	(0.230)	(0.219)	(0.221)	(0.213)	
	57,910	11,487	8,672	8,669	9,818	9,635	9,629	

Notes: Average characteristics on the 57,910 subjects in the field experiment, with standard errors reported in parentheses. Column (1) corresponds to the entire sample. Columns (2) through (7) correspond to each of the six treatment groups. Column (8) reports the p-value of the test of equal means across all six treatment groups. All the variables correspond to pre-treatment characteristics (i.e., that were determined before the start of letter delivery). *Days Listed* is the number of days that the property had been listed for before our experiment. *List Price* is the original listing price of the property. *No. Beds* is the property's number of bedrooms. *No. Baths* is the property's number of bathrooms. *Sq. Ft. Built* is the property's built area in square feet. *Sq. Ft. Lot* is the property's lot size in square feet.

Table 2: Main Regression Results

	Survey Data				Behavioral Data			
	(1) H_{1y}^{post}	(2) H_{5y}^{post}	(3) H_{1y}^{prior}	(4) M_{1y}^{post}	(5) S_{+12w}	(6) S_{+28w}	(7) S_{-1w}	(8) D_{pre}
Information Shock	0.205***	0.167**	-0.014	0.017	-0.330***	-0.325***	0.014	0.001
	(0.064)	(0.070)	(0.066)	(0.134)	(0.103)	(0.107)	(0.019)	(0.003)
Mean Outcome	3.86	2.31	3.88	3.58	36.99	56.90	0.58	3.81
Std. Dev. Outcome	4.42	4.36	5.39	9.05	48.28	49.52	7.61	1.28
Observations	1,404	1,404	1,404	1,404	57,910	57,910	57,910	57,910

Notes: Significant at *10%, **5%, ***1%. Heteroskedasticity-robust standard errors in parentheses. Each column corresponds to a different regression. All regressions correspond to equation (2) from Section 2, with *Information Shock* referring to the key independent variable: $E_i^{signal(j^*)} \cdot T_i$. The only difference between columns is that they use a different dependent variable. Columns (1) through (4) are based on data from the AMT supplemental survey. H_{1y}^{post} is the posterior belief (i.e., elicited after the information-provision experiment) about the expected growth rate of the median home value over the following year. H_{5y}^{post} is is the posterior belief about the annualized expected growth growth rate of the median home value over the following five years. H_{1y}^{prior} the prior belief (i.e., elicited before the information-provision experiment) about the expected growth rate of the median home value over the following year. M_{1y}^{post} is the posterior belief about the annualized expected growth growth rate of the stock market index over the following year. Columns (5) through (8) are based on data from the field experiment. S_{+12w} is an indicator variable that takes the value 100 if the property was sold at 12 weeks after the start of the letter delivery. S_{+28w} is an indicator variable that takes the value 100 if the property was sold at 28 weeks after the start of the letter delivery. S_{-1w} is an indicator variable that takes the value 100 if the property was sold at 1 week prior to the start of the letter delivery. And D_{pre} is the logarithm of the number of days that the property had been listed for before our experiment.

Table 3: Additional Robustness Checks

	Dep. Var: S_{+12w}									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Information Shock	-0.330***	-0.325***	-0.325***	-0.286**	-0.330***	-0.338***	-0.257*	-0.325***	-0.382***	-0.320***
	(0.103)	(0.100)	(0.111)	(0.146)	(0.119)	(0.118)	(0.137)	(0.110)	(0.111)	(0.106)
Additional Controls		Y								
Extended Dummies			Y							
Control for Feedback				Y						
Group Left Out					Baseline	Past-1	Past-2	Forecast-1	Forecast-2	Forecast-3
R^2	0.034	0.111	0.034	0.034	0.032	0.035	0.034	0.033	0.034	0.034
Observations	57,910	57,910	57,910	57,910	46,423	49,238	49,241	48,092	48,275	48,281

Notes: Significant at *10%, **5%, ***1%. Heteroskedasticity-robust standard errors in parentheses. Each column corresponds to a different regression. All regressions are based on data from the field experiment and using the same dependent variable: an indicator variable (S_{+12w}) that takes the value 100 if the property was sold at 12 weeks after the start of the letter delivery and 0 otherwise. Column (1) corresponds to equation (2) from Section 2, with *Information Shock* referring to the key independent variable: $E_i^{signal(j^*)} \cdot T_i$. Column (2) is identical to column (1) except that it includes some additional control variables: the logarithm of the days the property was on the market prior to the experiment, the logarithm of the initial listing price, four dummies for number of beds, four dummies for number of bedrooms, the logarithm of square footage built, the logarithm of lot size, and six state dummies. Column (3) is identical to column (1) except that instead of controlling for one treatment indicator, it controls for a set of five treatment indicators (i.e., one for each of the five treatments that are not Baseline). Column (4) is identical to column (1) except that it includes an additional control variable: the value of the signal chosen for the individual without the treatment interaction (i.e., $E_i^{signal(j^*)}$). Column (5) through (10) are identical to column (1) except that they exclude subjects for one treatment group at a time.