

Betting on the House: Subjective Expectations and Market Choices[†]

By NICOLAS BOTTAN AND RICARDO PEREZ-TRUGLIA*

Home price expectations play a central role in macroeconomics and finance. However, there is little direct evidence on their effect on market choices. We provide causal evidence using a large-scale, high-stakes field experiment in the United States. We sent letters about home price trends to 57,910 homeowners who listed their homes. These letters contained randomized information creating nondeceptive, exogenous variation in home price expectations. Consistent with economic theory, higher expectations reduced selling probability. Behavior was highly elastic: a 1 percentage point increase in expectations caused a 2.63 percentage point reduction in the probability of selling the property within 12 weeks. (JEL C93, D83, D84, R31)

Consumer expectations play a central role in modern macroeconomics and finance, and they are of special interest to policymakers (Bernanke 2007). In the context of the housing market, there is growing interest in the expectations for the future growth of home prices, also known as home price expectations (Stroebel 2016; Armona, Fuster, and Zafar 2019; Bailey, Cao, et al. 2018; Bailey, Dávila, et al. 2018; Fuster et al. 2022). Because homes account for a large fraction of households assets, home price expectations can have major welfare and policy implications. For example, home price expectations are believed to have played a key role in the 2008 US housing crisis (Shiller 2005). Despite their central role, there is little *direct* evidence on whether home price expectations have a causal effect on market choices. In this paper, we fill this gap in the literature using a large-scale, high-stakes, preregistered natural field experiment.

More specifically, we study how home price expectations affect homeowners' decisions to sell their homes. The decision to sell a home is arguably one of the

*Bottan: Cornell Jeb E. Brooks School of Public Policy, Cornell University (email: nicolas.bottan@cornell.edu); Perez-Truglia: Anderson School of Management, University of California, Los Angeles (email: ricardo.truglia@anderson.ucla.edu). Xavier Jaravel was coeditor for this article. We are thankful to the anonymous referees for their excellent feedback; and thank comments from several colleagues and seminar discussants at UC Berkeley, Cornell University, Chicago-Booth, Brown University, Berkeley-Haas, CEBI Workshop, Jinan University, NTA Annual Meeting, Universidad de San Andres, and the NBER Summer Institute. 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). This RCT was registered as AEARCTR-0003663 (Bottan and Perez-Truglia 2019). We are thankful for 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. Sofia Shchukina and Zhihao Han provided outstanding research assistance.

[†]Go to <https://doi.org/10.1257/app.20230641> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

biggest decisions that households make, financially and otherwise (Brooks 2017). Homes are typically the most valuable asset that households own, accounting for nearly 30 percent of the assets of the average US homeowner (Eggleson et al. 2020). The decision to sell a home is both costly and nonreversible; it is costly because sellers typically pay between 8 percent and 10 percent of the home value (Zillow 2022) in various selling costs such as realtor fees, attorney fees, transfer taxes, and other expenses, and it is irreversible because, once sold, there is nothing that the seller can do to get the home back short of buying it back from the new owner. The selling process can bring a lot of uncertainty too. Although some homes are sold shortly after being listed, a significant share of homes remains unsold long after being listed. Further, although many homes are sold for the listing price, many others end up being sold for significantly less or significantly more. For all these reasons, selling a home can be quite stressful. According to survey data, most sellers (69 percent) reported feeling stressed about the time it would take to sell their homes (Carter 2019), and a significant proportion of sellers (36 percent) reported crying during the home-selling process.

Estimating the relationship between home price expectations and market choices is plagued with challenges to causal identification. Consider the following correlation as an example: in markets with more optimistic home price expectations, homeowners are less willing to sell their homes. This correlation could be spuriously driven by an omitted variable. For example, an improvement in fundamentals may simultaneously generate increases in home price expectations and homeowners' unwillingness to sell. The causality may also run in the opposite direction: homeowners' reluctance to sell may cause home price expectations to increase. Without exogenous variation in expectations, it is short of impossible to credibly identify the causal effect of home price expectations on market choices.

To better illustrate the link between subjective expectations and market choices, consider a sample of homeowners who have listed a home for sale. Those sellers receive a sequence of offers. For each offer, they must decide whether to accept it, or to wait until a better offer comes along. According to economic theory, subjective home price expectations could be a key input in the decision to accept the offer. More precisely, when homeowners become more optimistic about future home prices, they should be less excited about selling their property (i.e., they should increase their reservation price). As a result, the property should take longer to sell. Conversely, homeowners who become more pessimistic about future home values are expected to sell the property faster.

In an ideal experiment, we would take this group of sellers and flip a coin to randomize their expectations. For example, if the coin lands on heads, a homeowner would be persuaded that median home prices will appreciate by 1 percent over the next year. If the coin lands on tails, the homeowner would be persuaded that median home prices will appreciate by 10 percent over the next year. Then, we would verify which owners sold their homes and which ones did not sell their homes in the months following the randomization. Our main hypothesis is that, relative to the homeowners who are randomly assigned to the 1 percent home price expectation, the homeowners assigned to the 10 percent home price expectation would take longer to sell their properties. Moreover, the magnitude of those

effects can shed light on how elastic homeowners are to changes in their home price expectations.

We designed a field experiment that closely resembled the above-mentioned ideal experiment. Our subject pool consisted of a sample of US homeowners who had listed a home on the market. To better understand this context, it is important to note that the decision to list a home is not nearly as high-stakes as the decision to sell the home. When a homeowner lists a home, they typically do not have to pay for anything: the vast majority of the selling costs (e.g., realtor fees, attorney fees, and transfer taxes) are realized if and when they decide to sell the home. The decision to list a home is also easily reversible, as it typically only takes a call to the realtor to take the listing down. For these reasons, homeowners can list a property to test the waters but then end up not selling it because of a lack of a sufficiently attractive offer. Indeed, a significant fraction of homes that are listed are not sold in the next 12 months.¹

We mailed letters to our sample of US homeowners who had listed a property for sale. These letters included information on the current median home price for comparable homes (i.e., homes in the same zip code and with the same number of bedrooms). We randomized additional factual information in the letters to create nondeceptive, exogenous shocks to the subjects' home price expectations. By chance, some owners received optimistic shocks to their expectations, whereas other owners received pessimistic shocks. We then used publicly available administrative records to determine whether and when the homeowners sold the listed property. We measured whether the exogenous shocks to home price expectations induced by our letters affected the subjects' subsequent market choices. Most importantly, we also measured whether the shocks to expectations affected whether the home was sold in the subsequent months. Finally, we measured the exogenous shocks' effects on listing prices, which constituted our proxy for the reservation price.

Information-provision experiments typically require data on prior beliefs. For example, imagine that an individual is given information that home prices will grow by 5 percent per year. Whether that individual should adjust their expectations upward, downward, or not at all depends on the individual's prior beliefs. Individuals who believe that home prices will grow by less than 5 percent should adjust their expectations downward, whereas individuals who believe that homes will appreciate by more than 5 percent should update their expectations upward. We devised an experimental design to leverage exogenous shocks to beliefs even in the absence of data on prior beliefs. This design created two distinct sources of exogenous shocks to the subjects expectations: *source randomization* and *disclosure randomization*.

Source randomization consists of randomizing the source used for the additional information on the evolution of home prices. We used three sources that prior survey experiments (Armona, Fuster, and Zafar 2019; Fuster et al. 2022) have shown

¹For instance, among all homes listed in 2019, only 64.6 percent ($= 628/(628 + 344)$) of them were ultimately sold that same year (Federal Reserve Bank of St. Louis 2019). Some of those homes are delisted without being sold, and some others remain active, with sellers waiting for the right offer. With some slight adjustments, our field experiment could also be used to study the decision to list a home. However, we focus on the decision to sell a home precisely because it is the higher-stakes decision.

to significantly affect the formation of home price expectations: the average price change over the past year, the average price change over the past two years, and one of three forecasts about the price change in the next year based on different statistical models.

To illustrate the intuition behind the source-randomization design, consider a group of subjects selling a two-bedroom home in zip code 33308. Each homeowner could be randomly assigned to one of the following five signals of home price growth: an annual growth rate of 1.2 percent over the past 1 year, an annual growth rate of 3.6 percent over the past 2 years, an annual growth forecast of 2.6 percent (according to statistical model 1), an annual growth forecast of 4.1 percent (model 2), or an annual growth forecast of 3.5 percent (model 3). For each subject, we randomize which one of the five signals is shown to them. The randomization of the source generates an exogenous information shock. By chance, some subjects in this group are shown a more optimistic signal and others a more pessimistic signal. Relative to receiving a signal from the first source (1.2 percent), receiving a signal from the second source (3.6 percent) constitutes a positive information shock, which should result in more optimistic home price expectations. More precisely, relative to being assigned to the first signal, being assigned to the second signal amounts to an information shock of +2.4 pp ($= 3.6 - 1.2$). Likewise, receiving the third, fourth, or fifth signal should amount to an information shock of +1.4 pp, +2.9 pp, or +2.3 pp, respectively.

The disclosure-randomization variation creates additional exogenous variation in beliefs. We cross-randomized whether the additional information on the evolution of home values was included in the letter or not. The disclosure-randomization followed a similar logic to the source-randomization variation, but instead of leveraging variation in signals across sources, it exploited variation in signals across different markets. These two sources of variation, disclosure randomization and source randomization, could be analyzed separately—indeed, comparing the results across the two identification strategies provided a valuable robustness check. Additionally, the two sources of variation could be combined into a single estimator to maximize statistical power.

To identify a sample of individuals who had listed their homes for sale, we used publicly available information from a major online listing website. Using unique identifiers for the property, we merged those records with rich administrative data from the corresponding county assessors office. These public records included detailed information about the property and its owners, such as their full names and mailing addresses. We used that contact information to mail a letter to the owners of the listed properties. We then used public records to track whether and when each property was sold during the six months following the mailing intervention. Moreover, the public data from the online listing website allowed us to track changes to the listing price. In June 2019, we mailed the letters to 57,910 unique homeowners in 36 counties across 7 US states. The homes were collectively valued at \$34 billion dollars.

The results from the field experiment confirmed that the information shocks affected actual, high-stakes market choices and were in the direction predicted by economic theory. A larger information shock (i.e., making expectations more optimistic) reduced the speed at which properties were sold. This effect was highly

statistically significant and large in magnitude: a 1 pp higher information shock causes a 0.350 pp drop in the probability that the property sold within 12 weeks (p -value = 0.001), implying a behavioral elasticity of -0.35 .²

The results from the field experiment were robust to a number of checks. We used an event-study analysis to exploit variation in the timing of when subjects received and read our letters. First, we estimated the effects on the outcomes immediately before the letters were delivered, at which point they should have had no effect on the sales outcome. As expected, the effects of the information shocks were precisely estimated around zero in the pretreatment period. Moreover, we leveraged the fact that not all letters were delivered and read at the same time but instead were gradually opened over a period of seven weeks. We showed that, as expected, the effects of the letters intensified during that period and stabilized thereafter. Moreover, the event-study analysis showed that the effects of our information shocks were highly persistent. For instance, the behavioral elasticity estimated at 28 weeks posttreatment (-0.295 , p -value = 0.006) was close to the corresponding elasticity estimated at 12 weeks posttreatment (-0.350 , p -value = 0.001).

For an additional falsification test, we estimated placebo regressions that were identical to the baseline specification, except that they used pretreatment characteristics as dependent variables. For example, some of those placebo outcomes were the number of days the property had been listed prior to our experiment or the original listing price. Because those outcomes were determined prior to our mailing intervention, the randomized information shocks contained in the letters should not have affected them. As expected, we found that the placebo effects were close to zero, statistically insignificant, and precisely estimated.

We provide several additional robustness checks. In our baseline specification, we combined the disclosure-randomization and source-randomization variation to maximize statistical power. However, the results were almost identical if we used those two sources of variation separately. The results were almost identical in alternative specifications, such as including a rich set of control variables. For a less parametric look at the data, we used a binned scatterplot version of the baseline specification. The results confirmed that the findings were not driven by nonlinearities or outliers.

We provide suggestive evidence about the underlying mechanisms. Ideally, we would like to study how the information shocks affect homeowners' reservation prices. However, by construction, reservation prices are unobservable to researchers—only the seller knows how much they are *truly* willing to accept for their home. As an approximation, we can use changes in listing prices as a proxy for (unobserved) changes in reservation prices. If information shocks affect reservation prices, those effects may be reflected in changes to the listing prices. Indeed, we find suggestive evidence consistent with this mechanism.³

To provide complementary evidence on the effects of information shocks, we designed a supplemental survey experiment. This supplemental experiment was

²We chose the 12-week horizon arbitrarily to summarize the information, but the results are similar for other horizons—for more details, see Section IVA.

³We use sales prices as an alternative proxy for reservation values. However, the results for that outcome have to be taken with a grain of salt for at least two reasons: they rely on additional assumptions to deal with selection bias, and they are statistically insignificant.

included in the same randomized control trial preregistration as the field experiment and was around the same dates. The survey experiment exposed a separate sample of 1,404 subjects to the exact same information treatments used in the field experiment. However, instead of measuring the effects of the information shocks on their market choices, the survey experiment measured the effects on their subsequent home price expectations. The results from the supplementary experiment confirmed that our information shocks had significant effects on home price expectations and were in the expected direction. A 1 pp information shock increased one-year-ahead home price expectations by 0.205 pp (p -value = 0.001). Indeed, this degree of pass-through from information shocks to posterior beliefs was consistent in magnitude with the results from other survey experiments about home price expectations (Armona, Fuster, and Zafar 2019; Fuster et al. 2022) and other macroeconomic expectations (Cavallo, Cruces, and Perez-Truglia 2017; Roth and Wohlfart 2020).

Next, we discuss the magnitude of the effects of home price expectations on sellers behavior. Due to the presence of two forms of noncompliance, the behavioral elasticity of -0.350 reported above constituted an intention-to-treat effect. First, some subjects may not have received or read the letter in time. Second, conditional on reading the information provided in the letter, some recipients may have ignored the information or may have updated their expectations only partially. We provide estimates of the treatment-on-the-treated effect by correcting for both forms of noncompliance. We use statistics from the US Postal Office to correct for the first form of noncompliance and estimates from the supplemental survey experiment to correct for the second. These calculations suggest that the subjects were highly elastic to their home price expectations: increasing home price expectations by 1 pp would cause a reduction of 2.63 pp in the probability of selling the home within the next 12 weeks.

Given the large variation in home price expectations observed across households and over time, our findings imply that subjective expectations can be a major factor driving decision-making in the housing market. For example, in a given cross section of households, some households tended to be more optimistic than other households. Our treatment-on-the-treated estimates implied that a one standard deviation increase in home price expectations would cause a reduction of 14.18 pp in the probability of selling the property within the next 12 months.

Last, we measured heterogeneity in the effects of information shocks on home price expectations using characteristics of the owner, property, and local housing market. We found similar effects across the board. For example, the point estimates were similar and statistically indistinguishable between male and female homeowners and between more and less expensive homes. Two notable exceptions were stronger effects when the owner was a non-occupant and when the owner was aged 59 and above. These two differences must be taken with a grain of salt, however, as they were large in magnitude but borderline statistically insignificant. Our preferred interpretation for these sources of heterogeneity is based on selling frictions: some owners may want to react to the information on home price expectations, but face constraints on when to sell due to changes in job, school, or family composition.

This study relates and contributes to various strands of literature. Most important, it relates to literature on the role of subjective expectations in the housing

market. Some studies measure the relationship between home price expectations and market choices, such as whether to buy or rent, or the mortgage leverage choice (Bailey, Cao, et al. 2018; Bailey, Dávila, et al. 2018). For example, Bailey, Cao, et al. (2018) present evidence that individuals are more likely to transition from renting to owning after geographically distant friends experience large recent increases in home prices. Others study home price expectations through survey experiments (Fuster et al. 2022) or laboratory experiments (Armona, Fuster, and Zafar 2019).

We contribute to this literature by providing the first field experiment to measure the effects of home price expectations. Indeed, our contribution extends to broader literature on subjective macroeconomic expectations.⁴ Our context was nearly ideal in many dimensions. In terms of causal identification, our study used the gold standard in economic research: a randomized controlled trial. Our study exploited a naturally occurring context in that our subjects (like millions of homeowners in the country each year) had already decided to put their homes on the market. We employed a high-stakes context wherein the decision accounted for a large fraction of the net worth of the decision-maker. Our study also involved a large-scale experiment with 57,910 subjects, which allowed us to provide precise estimates and sharp falsification tests (e.g., event-study analysis). Furthermore, rather than using survey data to measure behavior, which is subject to many criticisms such as measurement error and experimenter demand effects, we measured the *actual* market outcomes using administrative records. Finally, our methodology provided a quantitative measure of the causal effects of expectations—through an elasticity—that provided an intuitive sense of the magnitude of the effects.

This study also contributes to the broader literature on the housing market. Deviations from the full-information rational expectations model play a prominent role in many accounts of the housing market (Case and Shiller 1989; Shiller 2005; Glaeser and Nathanson 2015; Anenberg 2016; Glaeser and Nathanson 2017; Gennaioli and Shleifer 2018; Bailey, Cao, et al. 2018; Bailey, Dávila, et al. 2018; Kaplan, Mitman, and Violante 2020). We provide experimental evidence that information frictions exist and are consequential—i.e., providing households with publicly available information has significant effects on their market choices.

The rest of the paper proceeds as follows. Section I presents the conceptual design and econometric model. Sections II and III discuss the details of the design and implementation of the field experiment. Section IV presents the main results from the field experiment. Section V presents the results from the supplemental survey experiment. Section VI discusses the magnitude of the estimates. The last section concludes.

⁴This broader literature includes other topics such as inflation (Armantier et al. 2016; Cavallo et al. 2017; Coibion et al. 2018) and GDP growth (Roth and Wohlfart 2020). These studies typically provide a random subset of respondents with a piece of information and measure the corresponding effects on their subsequent survey responses, including their posterior beliefs, attitudes, or even small-stakes laboratory choices (Armantier et al. 2015; Armona, Fuster, and Zafar 2019). Our paper adds to recent work studying the effects of macroeconomic expectations on *actual* behavior in high-stakes and naturally occurring contexts (Coibion et al. 2024).

I. Research Design and Econometric Model

Information-provision experiments typically rely on data about prior beliefs. Consider randomizing subjects to receive—or not receive—a signal that the annual growth rate of home prices in the next 12 months will be 5 percent. The effect of the signal on the subsequent home price expectations (and market choices) depends on the prior beliefs of the individual. For subjects whose prior home price expectations were below 5 percent, we expect the signal to cause them to update their home price expectations upward (and subsequently take longer to sell their homes). On the contrary, subjects whose prior home price expectations exceeded 5 percent should update their home price expectations downward (and subsequently sell their homes faster). If the prior expectations were exactly 5 percent, then recipients should not update their home price expectations (and their behavior should not change either).

The challenge in our high-stakes, large-scale context is that it is not feasible to measure the prior beliefs of tens of thousands of homeowners. Thus, we designed an information-provision experiment that does not rely on information about prior beliefs (see e.g., Bérigolo et al. 2023; Perez-Truglia and Troiano 2018; Bontan and Perez-Truglia 2022). In fact, our experimental design creates not only one but two sources of exogenous shocks to the subjects' expectations, which we call *source randomization* and *disclosure randomization* and are described in detail below. These two sources of variation can be combined in a single estimator to maximize power, but they can be analyzed separately too—indeed, this comparison provides a unique robustness check of the identification strategy.

A. Source Randomization

Let subscript i index subjects. Each subject belongs to a specific market m , denoted by the combination of property location (in the field experiment, the five-digit zip code) and property type (the number of bedrooms).

Let E_m^j be a signal about the future growth of home prices, where the subscript m notes that this signal pertains to market m (e.g., three-bedroom homes in zip code 33308) and the superscript $j = 1, \dots, J$ corresponds to the information source. For example, $j = 1$ could be the annual price change over the past one year, and $j = 2$ could be the annual price change over the past two years. For this design, it is important that the signals produced by different information sources cannot be exactly equal: e.g., even if two sources are equally optimistic on average, for some markets one source must provide a more optimistic signal, while for other markets another source must provide a more pessimistic signal. In this design, all subjects are provided with a signal, only that from a randomly chosen source. We use j_i^* to denote the source that was chosen for individual i .

Let $Y_{i,m}^{post}$ be the outcome of interest. In the field experiment, $Y_{i,m}^{post}$ is an indicator variable for whether the homeowner i who has a property in market m sold the property within a given number of months after the experiment. We use superscript *post* to make it salient that the outcome was measured in the posttreatment period

(for falsification tests, we will use pretreatment outcomes, which will be denoted by superscript *pre* instead). The regression of interest is as follows:

$$(1) \quad Y_{i,m}^{post} = \nu_0 + \nu_1 \cdot E_m^{j_i^*} + \phi_m + \varepsilon_{i,m},$$

where ϕ_m are market-specific fixed effects, to make sure that we are always comparing between pairs of individuals from the same market (i.e., comparing between individuals who could have been assigned to the exact same set of signals). One intuitive way of thinking about this experiment is that for individuals in a given market m , they could be shown one of several signals (E_m^j for $j = 1, \dots, J$), some of them more optimistic and others more pessimistic. So we flip a coin to decide which individuals in market m get the more optimistic signal and which individuals get the more pessimistic signals, and then we measure whether those who were randomly assigned to more optimistic signals behave differently than those assigned to more pessimistic signals.

The coefficient of interest, ν_1 , measures the effects of the information shock. To gain more intuition about this coefficient, and without loss of generality, we can normalize $E_m^{j_i^*}$ by taking the difference with respect to the first information source ($j = 1$) and reexpress equation (1) as follows:

$$(2) \quad Y_{i,m}^{post} = \nu_0 + \nu_1 \cdot (E_m^{j_i^*} - E_m^1) + \phi'_m + \varepsilon_{i,m}.$$

Note that since E_m^1 is the same for every individual within market m , this transformation amounts to a simple translation of the market fixed effects (from ϕ_m to ϕ'_m). Consider a subject selling a two-bedroom home in zip code 33308. That homeowner could be randomly presented with one of the following five signals: an annual growth rate of 1.2 percent over the past 1 year ($j = 1$), an annual growth rate of 3.6 percent over the past 2 years ($j = 2$), an annual growth forecast of 2.6 percent according to statistical model 1 ($j = 3$), an annual growth forecast of 4.1 percent according to model 2 ($j = 4$), or an annual growth forecast of 3.5 percent according to model 3 ($j = 5$). The variable $(E_m^{j_i^*} - E_m^1)$ equals zero if $j_i^* = 1$ and 2.4 if $j_i^* = 2$. That is, relative to receiving the first signal, receiving the second signal means an information shock of 2.4 pp ($= 3.6 - 1.2$). The information shock equals 1.4 pp, 2.9 pp, and 2.3 pp when $j_i^* = 3$, $j_i^* = 4$, and $j_i^* = 5$, respectively. Thus, the coefficient ν_1 measures the effect of a 1 pp increase in the information shock.

B. Disclosure Randomization

The disclosure randomization creates additional exogenous shocks to expectations. We randomly assigned one information source to every subject, but then cross-randomized whether the chosen signal would be actually disclosed to the subject (i.e., included in the letter sent to the subject or not included).

To explain this research design, it is useful to start with a stylized example. To keep things simple, assume that we have a single information source. According to this source, some markets are expected to grow strongly in the future, but other markets should expect weak price growth. Again, for the sake of simplicity, let's

consider the extreme case where subjects do not have access to the information source. As a result, there will be no relationship between the value of the signal and the home price expectations: subjects who could have been shown the more optimistic signal will be equally optimistic as those subjects who could have been shown the more pessimistic signal, because they never got to actually see the signal. For subjects who *are* shown the signal, however, there should be a significant correlation between the value of the signal and their home price expectations: individuals who find out that the signal is optimistic in their market should react by forming more optimistic expectations; while individuals who find out that the signal is pessimistic in their market should form less optimistic expectations. Our research design seeks to exploit that prediction—i.e., that the correlation between expectations and the signal should become stronger when the signal is disclosed versus when it is not disclosed.

Let D_i be an indicator variable that equals 1 if the chosen signal is disclosed to subject i , and 0 otherwise. The regression of interest is as follows:

$$(3) \quad Y_{i,m}^{post} = \mu_0 + \mu_1 \cdot E_m^{j_i^*} \cdot D_i + \mu_2 \cdot E_m^{j_i^*} + \mu_3 \cdot D_i + \varepsilon_{i,m}.$$

The parameter μ_2 measures the relationship between the signal ($E_m^{j_i^*}$) and the outcome ($Y_{i,m}^{post}$) for individuals who are not shown the signal (i.e., when $D_i = 0$). The key parameter is μ_1 : i.e., how much stronger that relationship is for individuals who were shown the signal ($D_i = 1$) relative to individuals who were not shown the signal ($D_i = 0$). The coefficient of interest, μ_1 , measures the effect of a 1 pp increase in the information shock.

The exogenous shocks induced by disclosure randomization operate similarly to the source randomization, except that they exploit heterogeneity in signals across markets rather than across information sources. Note that the disclosure randomization requires that there is heterogeneity in signals across subjects. If, for example, all subjects had properties in the same market (i.e., same zip code and property type), there would be no variation to identify μ_1 .

C. Pooled Specification

As stated in the preregistration, our baseline specification pools the two sources of random variation to maximize statistical power. This pooled specification is given by the following equation:

$$(4) \quad Y_{i,m}^{post} = \pi_0 + \pi_1 \cdot E_m^{j_i^*} \cdot D_i + \pi_2 \cdot D_i + \phi_m + \varepsilon_{i,m}.$$

The coefficient of interest is π_1 ; that measures the effect of a 1 pp increase in the information shock. To show that equation (4) pools the two sources of variation, we can show that equation (1), for source randomization, and equation (3), for disclosure randomization, are just special cases of equation (4). First, consider what happens if we disclosed information to every subject—i.e., $D_i = 1$ for every i . In that case, by construction, the only exogenous variation left would be the source randomization. To represent this case, we can replace $D_i = 1$ in equation (4). As

a result, equation (4) turns into equation (1). Second, consider what happens if we used a single information source—i.e., $J = 1$. By construction, the only exogenous variation left would be the disclosure randomization. To represent this case, we can replace $E_m^{j_i}$ by E_m^1 in equation (4). As a result, equation (4) turns into equation (3), with the only difference being that we do not need to control for E_m^1 (as in equation (3)), because that control would be absorbed by the market fixed effects (note that E_m^1 cannot take different values within a given market m).

II. Design of the Field Experiment

A. Design of the Mailings

The field experiment consisted of sending letters to a sample of homeowners who had listed a property for sale. Online Appendix Figure A.1 shows a sample of the envelope. We took a number of measures to communicate that the letter, though unsolicited, came from a legitimate source. The top-left corner of the envelope included the logo for the University of California at Los Angeles (UCLA) and a note about the research study. The top-right corner of the envelope included nonprofit organization postage. Figure 1 shows a sample letter for a fictitious subject: panel A corresponds to the front page of the letter, while panel B corresponds to the back page. All letters included an introduction, the official UCLA logo in the header, contact information in the footer, a physical correspondence address, and a URL of the study's website with additional information.⁵

One letter was mailed to every individual in the subject pool. All letters were identical, except for some of the information that was personalized (e.g., the recipient's name) and some of the information that was randomized. Figure 1 shows the placeholders (marked as «*Information*» and «*Information Details*») for the two pieces of information that differed across treatment groups. The «*Information*» portion included a table with information about home prices. Figure 2 shows a sample table for each of the six treatment groups, which are discussed in detail below. The «*Information Details*» portion included methodological notes for the table, such as data sources and statistical models used.⁶


All letters contained information on the current median home value of similar properties. For example, subjects who listed a three-bedroom home in zip code 90210 received a letter indicating Zillow's estimated median home value for three-bedroom homes in zip code 90210. We used the same property types as those used by online real estate market platforms: one-bedroom, two-bedroom, three-bedroom, four-bedroom, and five-plus-bedroom.⁷ In addition to information about the

⁵This website contained general information such as contact information for the researchers and institutional review board, but did not specify the hypotheses being tested. A copy of this website was hosted on the UCLA website. For a screenshot of the website, see online Appendix E.

⁶Online Appendix Figure A.6 shows the six corresponding methodological notes. For a sample letter with the final product, see online Appendix D.

⁷For a small minority of properties, the number of bedrooms was not available in the tax rolls, and thus we assigned them to a broader category called all homes. Zillow does not produce estimates of median home values for some combinations of zip codes and number of bedrooms. For those subjects, we used the category "all homes" as well.

Panel A. First page



Panel B. Second page



FIGURE 1. SAMPLE LETTER

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 Figure 2 and online Appendix Figure A.6, respectively. Online Appendix D shows a sample of the final product.

current median home values, the table could include information on the evolution of median home values, which individuals could use in forming their home price expectations. Homeowners were randomized into one of six treatment groups. These treatment groups differ in whether the table includes additional information on the price evolution and the source of said information:

- **Baseline:** no additional information on price evolution
- **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
- **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

We choose these information sources because they were shown to have significant effects on home price expectations in previous studies. For example, Fuster et al. (2022) show that upon being shown one of these types of information, subjects update their expectations in the expected direction. Fuster et al. (2022) also




FIGURE 2. SAMPLE INFORMATION TABLES

Notes: Each panel corresponds to the hypothetical table that a given individual would have received 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 1. See online Appendix Figure A.6 for the methodological notes accompanying each table.

show that most households are willing to pay to have access to these information sources, which suggests that they find them useful. Beyond survey experiments, there are theoretical arguments for why subjects may be interested in learning about these information sources. According to the backward-looking expectations model, individuals form beliefs by looking at recent price changes (Case and Shiller 1989; Shiller 2005). And according to models of rational expectation, households may form expectations based on professional forecasts (Carroll 2003).

The price changes over the past one year (Past-1) and past two years (Past-2) corresponded to the raw Zillow data. The three statistical forecasts (Forecast-1, Forecast-2, and Forecast-3) were based on the same Zillow data. All three forecasts were estimated using year/zip code–level data on the Zillow Home Value Index for 1997–2019. All three models are autoregressive, but they differ in the set of explanatory variables chosen. These differences in specification yielded

slightly different forecasts. Forecast Model 1 used five lags of the dependent variable. Forecast Model 2 used five lags of the dependent variable plus five lags of the state-level average of the dependent variable. Forecast Model 3 used 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.⁸ All five information sources were informative to a reasonably similar degree: all had similar predictive power and were comparable to the predictive power of Zillow's official forecasts.⁹ It is important to note that this design is nondeceptive: all letters, regardless of the source, were based on real data that homeowners could access from publicly available sources.

Each panel of Figure 2 corresponds to the hypothetical table that a subject would receive depending on the assigned treatment group. It shows real examples based on an individual selling a two-bedroom home in zip code 33308. Panel A shows the baseline letter, which includes the current median price level only. The following five panels add information on the price evolution: panel B shows an average annual growth rate of 1.2 percent over the past year (Past-1 treatment), panel C shows an annual growth rate of 3.6 percent over the past 2 years (Past-2 treatment), panel D shows the annual growth of rate 2.6 percent projected by Forecast Model 1 (Forecast-1 treatment), panel E shows the annual growth rate of 4.1 percent projected by Forecast Model 2 (Forecast-2 treatment), and panel F shows the annual growth rate of 3.5 percent projected by Forecast Model 3 (Forecast-3 treatment). As discussed in Section I, our identification strategy requires heterogeneity in signals across individuals and across information sources, which we show in Section III below.

As shown inside the blue box at the bottom of Figure 1, panel A, the letter includes a URL to an online survey and a unique survey code to verify that the response came from a legitimate subject.¹⁰ The main goal for including the survey link was to provide a proxy for the dates when recipients opened the letters, as in Perez-Truglia and Cruces (2017) and Nathan, Perez-Truglia, and Zentner (2020).

III. Data Sources and Implementation of the Field Experiment

A. Data Sources

To implement the mailing experiment, we combined two publicly available sources of data: data on active real estate listings and data for the property tax rolls from the county assessor. We scraped data on real estate listings from a major listing website. These data included rich information about the listed properties, such as address, listing price, property characteristics (number of bedrooms, bathrooms, size, days on market), and the assessor's unique parcel number (APN) for the property. We used the APN to match each listing to its corresponding record in the county assessor's tax rolls. Tax rolls contain rich information on the properties and owners. Most

⁸For more details about the three forecasting models, see online Appendix A.3.

⁹Details presented in online Appendix A.5.

¹⁰To verify that the respondents were legitimate subjects and to link survey responses at the individual level, subjects had to enter a unique survey code in the first screen of the survey before they could answer any questions. Each code consisted of a unique combination of six characters and was displayed prominently in the letter next to the URL of the survey.

important for our experiment, the tax rolls include the names of the owners and their mailing addresses (for more details on the data sources, see online Appendix A.1).

US counties typically make their property tax rolls publicly available. However, how accessible those tax rolls are can vary widely, even within a state. Some counties post the data online. For example, raw data from all counties in Florida can be easily downloaded at any time using a file transfer protocol address. Other counties, such as Alameda County in California, provide this information only in person and on a case-by-case basis. Many others, such as Los Angeles County, require filling out a short form and paying a fee to obtain a compact disc with the raw data. For the field experiment, we selected a set of 36 counties to obtain a large enough subject pool and for which all the required information from the tax rolls (e.g., owner's name and mailing address) was easily accessible. These counties are distributed across seven states and include 30 counties in Florida, Los Angeles County in California, Maricopa County in Arizona, Clark County in Nevada, Cuyahoga County in Ohio, King County in Washington, and Harris County in Texas. In practice, many other US counties likely would be feasible to include in this type of experiment.

B. Mailing Campaign

On May 28, 2019, we obtained the information on the active real estate listings and the latest available version of the secured tax rolls. Of the 173,708 active listings scraped, around 164,298 included the APN. For these listings, we merged nearly all listings (164,176 out of 164,298) with the county assessor's data. As 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 excluded properties or individuals who were not ideal for the experiment. For example, we excluded nonresidential properties and residential properties owned by businesses, because it was unclear whether our letter would be delivered to the person choosing whether to sell or not (e.g., the mailing address may correspond to the firm's lawyer). Similarly, we excluded individuals who recently moved according to the latest mail-forwarding data from the US Postal Services, and individuals who owned multiple properties in the same county (for more implementation details, see online Appendix A.2).

After applying these filters, our pool of potential subjects consisted of 61,176 individuals. From those, we selected a random sample of 60,000 individuals to receive a letter. After processing the data through the US Postal Service, we excluded a minority of subjects (2 percent) whose mailing addresses were flagged as undeliverable or vacant (1,198). An additional 892 subjects were excluded, mainly because their tax rolls were outdated and thus we mailed the letters to the previous owners.¹¹ The final subject pool comprised 57,910 individuals to receive letters. These individuals were randomly assigned to the following treatments: 20 percent to Baseline, 15 percent

¹¹ The tax rolls are updated with a lag. As a result, when we received the most up-to-date rolls, we identified 845 letters that were sent to a previous owner rather than the most current owner (i.e., the one who listed the home on the market). Another 37 subjects were dropped because they could not be matched to the administrative data. Last, ten subjects were dropped because they were fake subjects included on purpose for quality control of the mailing campaign.

each to Past-1 and Past-2, and 16.6 percent each to Forecast-1, Forecast-2, and Forecast-3. All letters were mailed on June 10, 2019.

C. Descriptive Statistics and Randomization Balance

We present descriptive statistics for the subject pool in column 1 of Table 1. The average property was listed for \$575,000 and had 3.3 bedrooms, 2.6 bathrooms, 2,300 square feet (sq. ft.) of living space, and a lot size of 13,000 sq. ft. Table 1 provides a balance test as well. Columns 2 through 7 of Table 1 break down the average property characteristics by treatment groups. The last column reports p -values for the null hypothesis that the average characteristics were equal across all six treatment groups. The results are consistent with successful random assignment: the observable characteristics are similar across all treatment groups in magnitude and not statistically distinguishable from each other.

In online Appendix A.4 we present additional details about the sample. For example, even though the tax rolls do not include owner characteristics such as gender, we were able to obtain complementary data from a private vendor. The average owner in our sample was 58.7 years old, 31.9 percent were female, 62.9 percent were White, 3 percent were African-American, 13 percent were Hispanic, 40 percent were college graduates, and their average annual household income was \$128,000. Additionally, online Appendix A.4 compares our subject pool to a representative sample of homeowners from the American Community Survey. We find that our experimental sample is fairly representative of homeowners in the same counties where the experiment was conducted. Moreover, the subject pool is representative of homeowners in the country as a whole in many dimensions such as the number of bedrooms and bathrooms and the amount of living space. However, there is one notable difference: homes in our subject pool (average listing price of \$575,000) are more expensive than in the country as a whole (average price of \$222,000 according to Zillow's Consumer Housing Trends Report). However, this difference arises mechanically because the subject pool includes counties that are more urban and more expensive than the US average.

Besides the characteristics of the properties and the owners, the subject pool is fairly representative in terms of the characteristics of the housing markets. The time to sell for properties in our subject pool (on average, 158 days) is somewhat higher than—but still in the same order of magnitude as—the corresponding average for homes in the same metropolitan areas (105 days).¹² This difference is also mechanical, due to the way in which we selected the subject pool.¹³ The owner

¹²To calculate this benchmark, we start with Zillow's housing data for June 2019 for the same metropolitan areas where the experiment was conducted. On average, it takes 60 days for a property to go from the initial listing to sale pending status. And we then add 45 days to account for the average closing time (according to Ellie Mae's Origination Insight Report).

¹³When we selected the subject pool, we did not focus on the properties that had been listed on that same day, but we also chose to include properties that had been listed weeks or even months before, to increase the sample size. Additionally, we do not make the distinction about whether the property is listed as for sale by owner, as it accounts for only 8 percent of sales nationally according to the National Association for Realtors, and of these, only 18 percent were listed through the Multiple Listing Service. This is consistent with less than 1 percent of listings in our sample being listed by the owner.

TABLE 1—RANDOMIZATION BALANCE TEST

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

Notes: Average characteristics of the 57,910 properties 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 pretreatment characteristics (i.e., 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. bedrooms” is the property’s number of bedrooms. “No. bathrooms” 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.

occupancy rate is quite close to the country average: around 67 percent of the listings are owner occupied in our subject pool, which is similar to the 64 percent of the housing units that were owner occupied in the country as a whole.¹⁴

D. Variation in Signals

As explained in Section I, the identification strategy relies on variation in signals across information sources and across markets. In this section, we show that there is significant variation in both of these dimensions.

Figure 3, panel A presents the variation in signals across information sources, which is the relevant variation for the source-randomization design. This scatterplot shows the relationship between the signals that the subjects would have received had they been assigned to the Past-1 treatment (i.e., annual growth rate over the past one year) versus the Past-2 treatment (i.e., annual growth rate over the past two years). For example, for two-bedroom homes in zip code 33308, the recipient would have been shown a price change of 1.2 percent if randomly assigned to the Past-1 treatment and a price change of 3.5 percent if randomly assigned to the Past-2 treatment. The two signals are highly correlated: on average, an extra 1 percent increase in the annual price change over the past 1 year is associated with an extra 0.659 percent increase in the annual price change over the past 2 years. This relationship is partly

¹⁴The US average statistics are from the five-year estimates from the American Community Survey corresponding to the period 2015–2019.




FIGURE 3. VARIATION IN SIGNALS ACROSS INFORMATION SOURCES AND ACROSS MARKETS

Notes: Panel A 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 versus the Past-2 treatment. The size of the circles are proportional to the number of observations, and the signals are truncated at -10 percent and $+20$ percent. Panel B 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 pp and are truncated at -10 percent and $+20$ percent.

mechanical (Past-2 is the average between Past-1 and another number) and 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: the $R^2 = 0.659$ is high but substantially below 1. Moreover, we document significant variation occurs across other pairs of information sources too (see online Appendix A.5 for more details).

Figure 3, panel B shows the variation in signals across markets, which is relevant for the disclosure-randomization design. This figure shows a histogram of the signal that subjects would have received if they had been assigned to the Past-1 treatment. The figure shows plenty of variation. Subjects in the tenth percentile lived in areas where median home values declined by -0.7 percent in the previous 12 months, and subjects in the ninetieth percentile lived in areas where property values increased 8.6 percent. There is quite a bit of variation in the other four information sources too (see online Appendix A.5 for more details).

E. Letter Delivery

The letters were mailed on June 10, 2019. To make the experiment more affordable, we used nonprofit postage. According to the US Monitor nonprofit standard mail delivery study, it takes nonprofit mailings about ten days to be delivered, with some letters arriving as much as a month after mailing (US Monitor 2014).¹⁵ As such, some subjects received the letter a few days after June 10, whereas other

¹⁵This delivery time is more than twice that of first-class mail, which is handled first, followed by presort standard, and, finally, nonprofit mail.

subjects received it weeks later. Even after the envelope is delivered, it may take days or even weeks for the subjects to open the envelope and read the letter. Some subjects may have been out of town when the letter was delivered to their homes. Other subjects may have received the letter right away but put it away and did not open it until weeks later.

Following Perez-Truglia and Cruces (2017) and Nathan, Perez-Truglia, and Zentner (2020), we used the distribution of dates when the individuals entered the survey codes included in the letter, as a proxy for when the letters were actually read. Hereafter, we refer to these dates as the read receipt.¹⁶ Figure 4, panel A presents the results. The first survey response was recorded on June 15, thus marking the start of letter delivery. Indeed, this date coincided with the best guess provided by the mailing company.¹⁷ This figure suggests that the letters were opened gradually from the start of the letter delivery until eight weeks later. The median time from the start of the letter delivery until the read receipt was approximately three weeks.

F. Main Outcome: Home Sale

To measure the behavioral outcomes, we scraped the administrative data from the real estate listing website on a weekly basis, from 2 weeks before the start of letter delivery until 28 weeks after the start of letter delivery.¹⁸

Whether the property was sold by a given date is the main outcome, and one of the primary outcomes listed in the American Economic Association randomized controlled trial preregistry. The preregistry defined the primary outcome as “time to sell,” which would be truncated for properties that are not sold or delisted. However, we realized that defining the dependent variable as whether the property was sold by a given date provides a simple and elegant solution to this truncation challenge and simplifies the exposition of results. Therefore, we use it as our main dependent variable in the paper. In online Appendix B.1 we show that the results are similar when using “time to sell” as dependent variable, employing duration models that naturally account for truncation. Our second preregistered outcome was the property’s sale price, which is also subject to censoring: for properties that were not sold, we do not know what the sales price would have been had they been sold. There is no simple or elegant fix for this data limitation. We provide a detailed discussion below in Section IVC and detailed analysis of this outcome in online Appendix B.3.

Administrative records indicate whether the property was sold and on what date. Confirmation of property sale came from either the Multiple Listing Service or the county assessor records. If we had confirmation from both sources, we used the earliest date for which we had confirmation. Our records usually included both sources, and the two dates were typically close. The Multiple Listing Service date is often

¹⁶Our proxy probably had some upward bias, because some people may have read the letter and waited a few days to respond to the survey. Another potential source of bias, which may be upward or downward, is that survey respondents could open the letters more or less slowly than survey nonrespondents.

¹⁷We asked the mailing company to provide a guess for when the first letters would be delivered based on the location of the shipping facility (Lombard, Illinois) and the location of the letter recipients.

¹⁸While we have not processed or looked at the data yet, our algorithm continued to automatically scrape the data on a weekly basis. As a result, if needed, we could look at longer time horizons too.




FIGURE 4. TIMING OF READ RECEIPT AND PROPERTY SALES

Notes: The vertical red line indicates the estimated date when the first letter was delivered (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.

earlier, allowing us to detect a sale as soon as possible (which is helpful for the interpretation of the event-study analysis).¹⁹

Figure 4, panel B shows the evolution of the sales outcome, corresponding to the cumulative fraction of properties in the subject pool. Note that the fraction of homes sold increases smoothly over time. By 12 weeks after the start of letter delivery, 36.99 percent of homes had been sold. By 20 weeks after the start of letter delivery, 50.6 percent of homes had been sold. By 28 weeks, the end of our panel data, 57.5 percent of the properties had been sold.

IV Main Results

A. Effects on Market Choices

We first present the main results from the field experiment: the intention-to-treat effects of information shocks on market choices of homeowners. Our main outcome of interest is whether the property was sold a given number of weeks after the start of the mail delivery. When interpreting these estimates, it is crucial to keep in mind that the effects of the information shocks on the sales outcome are not expected to materialize instantaneously. First, while some letters were read shortly after the start of the letter delivery, most letters were not read until a few weeks after that date. Second, even after the subject has read the letter, some time needs to pass so that the owner has the opportunity to act on their updated expectations: it may take one or more weeks for the seller to receive an offer, and after the offer is accepted, it

¹⁹ Note that our definition of a sold property does not include aborted sales. Properties under contract or pending sales are not included in our definition of sold property.

may take one or more weeks for the property to appear as sold in the records of the Multiple Listing Service or the county assessor's office.

The main regression results are presented in Table 2. Each column corresponds to a different regression. All regressions in this table are based on data from the field experiment and are 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. Of the properties, 36.99 percent were sold within this time horizon. We use this time horizon for the baseline results just because it provides enough time so that all subjects could have been affected by the information shocks. According to the read receipt proxy from Section III E, virtually all subjects read our letter within eight weeks after the start of letter delivery. As a result, when looking at the sales outcome 12 weeks after the start of the letter delivery, most subjects had been exposed to the information for 4–11 weeks, thus allowing sufficient time for the information to affect the sales outcomes. In any case, we reproduce the estimates for every possible time horizon and show that the results are robust.

We first examine the effects of information shocks by separately exploiting our two distinct sources of experimental variation (source randomization and disclosure randomization). Column 1 of Table 2 corresponds to the specification that uses the source-randomization only—i.e., equation (1) from Section I, which is restricted to the subsample of subjects who received signals (i.e., $D_i = 1$). *Information Shock* in Table 2 always corresponds to the coefficient on the key independent variable (in this specification, E_m^i). The results from column 1 indicate that the information shocks significantly affected homeowners' market choices: the information shock has a large negative effect on the probability that the property is sold (-0.318) and highly statistically significant (p -value = 0.009). This negative sign is consistent with the prediction from economic theory—i.e., a positive shock to expectations should decrease the probability that the property is sold. This coefficient is also economically large: an increase in the information shock of just 1 pp causes a 0.318 pp drop in the probability that the property is sold within the following 12 weeks. Since the right-hand-side and left-hand-side variables are measured in percentage points, this coefficient can be interpreted as a behavioral elasticity of -0.318 . While this effect is already economically significant, note that it reflects an intention-to-treat effect, because the information shocks should not be expected to fully materialize in changes to expectations. We defer a more careful discussion on economic magnitude to Section VI below, where we estimate a treatment-on-the-treated effect by correcting for various sources of noncompliance.

Using the variation induced solely by disclosure randomization yields similar estimates. Results are presented in column 2 of Table 2 (corresponding to equation (3), which uses the entire subject pool). As in column 1, the results from column 2 indicate that the information shock has an effect on market choices that is negative, large (-0.419), and statistically significant (p -value = 0.014). Most notably, the coefficients are similar in the specification that uses the source randomization (-0.318 , from column 1) solely as in the specification that uses the disclosure randomization (-0.419 , from column 2) solely—in fact, we cannot reject the null hypothesis that these two estimates (-0.318 and -0.419) are equal (p -value = 0.577). The fact

TABLE 2—MAIN REGRESSION RESULTS

	Dep. var: S_{+12w}					
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Information Shock</i>	−0.318 (0.122)	−0.419 (0.170)	−0.350 (0.105)	−0.341 (0.103)	−0.349 (0.113)	−0.330 (0.103)
Experimental variation	Source	Disclosure	Pooled	Pooled	Pooled	Pooled
Additional controls				Y		
Extended dummies					Y	
Alternative specification						Y
Mean outcome	36.93	36.99	36.99	36.99	36.99	36.99
SD outcome	48.26	48.28	48.28	48.28	48.28	48.28
R^2	0.167	0.008	0.155	0.193	0.155	0.034
Observations	46,423	57,910	57,910	57,910	57,910	57,910
	(7)	(8)	(9)	(10)	(11)	
<i>Information Shock</i>	−0.353 (0.121)	−0.307 (0.139)	−0.364 (0.113)	−0.411 (0.114)	−0.331 (0.109)	
Experimental variation	Pooled	Pooled	Pooled	Pooled	Pooled	
Group left out	Past-1	Past-2	Forecast-1	Forecast-2	Forecast-3	
Additional controls						
Extended dummies						
Alternative specification						
Mean outcome	37.10	37.10	36.95	36.91	36.92	
SD outcome	48.31	48.31	48.27	48.26	48.26	
R^2	0.164	0.164	0.165	0.165	0.165	
Observations	49,238	49,241	48,092	48,275	48,281	

Notes: Heteroskedasticity-robust standard errors in parentheses. Each column corresponds to a different regression. All regressions are based on data from the field experiment and are 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 the source-randomization specification given by equation (1) from Section I, which is restricted to the sample with $D_i = 1$ and with *Information Shock* corresponding to the key independent variable: E_m^i . Column 2 corresponds to the disclosure-randomization specification given by equation (3), with *Information Shock* referring to the key independent variable: $E_m^i \cdot D_i$. Columns 3–11 correspond to the pooled specification given by equation (4), with *Information Shock* referring to the key independent variable: $E_m^i \cdot D_i$. Column 4 is identical to column 3 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 bedrooms, four dummies for number of bathrooms, the logarithm of square footage built, the logarithm of lot size, and six state dummies. Column 5 is identical to column 3 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 6 is identical to column 3 except that rather than controlling nonparametrically for market fixed effects, it controls for the potential signals in a parametric way. Columns 7–11 are identical to column 3 except that they exclude subjects for one treatment group at a time.

that the results are highly consistent across two very different experimental designs is quite reassuring about the validity of the identification strategies.

To maximize statistical power, in column 3 of Table 2 we use the specification from equation (4) that combines the shocks induced by both the source randomization and the disclosure randomization. Again, we find that the information shock has an effect on market choices that is negative, large (−0.350), and highly statistically significant (p -value = 0.001). Given that it combines both sources of variation, it should not be surprising that the pooled specification yields an estimated value (−0.350, from column 3) in between the corresponding estimates from the specification that uses the source randomization (−0.318, from column 1) and the disclosure randomization

(-0.419 , from column 2). As anticipated in the preregistration, we use this pooled specification in the remainder of the paper due to its superior statistical power.

The baseline specification from equation (4) assumes that the effects of information shocks on market choices are linear. This functional form assumption is natural for two reasons. First, in theory, this simple linear specification is consistent with a simple model of Bayesian learning (Cavallo, Cruces, and Perez-Truglia 2017). Most importantly, in practice, several information-provision experiments found this linear specification to fit the data quite neatly in a variety of contexts such as home price expectations and inflation expectations (Armantier et al. 2016; Cavallo, Cruces, and Perez-Truglia 2017; Bottan and Perez-Truglia 2022; Fuster et al. 2022; Cullen and Perez-Truglia 2022). In theory, however, this functional form specification may miss some important features. For example, this specification may miss asymmetries: e.g., individuals may find it easier to react to good news than to bad news. This specification may also miss nonlinearities: e.g., some individuals may only consider signals that are not too extreme. To take a less parametric look at the data, Figure 5, panel A presents the binned scatterplot version of the baseline results from column 3 of Table 2. The results indicate that the linear baseline specification fits the data well. Moreover, this binned scatterplot confirms that the results are not driven by nonlinearities or by outliers.


The next robustness test exploits the timing of when subjects read the letters. The letters were being received and read progressively over a specific period of time. Thus, we can verify if the timing of the estimated effects is consistent with the timing of when subjects read the letters. Figure 6, panel A presents an event-study analysis of the effects of our information intervention. Each coefficient in Figure 6, panel A corresponds to a different regression using the same baseline specification from column 3 of Table 2 but with different dependent variables. For example, the leftmost coefficient uses a binary dependent variable that takes the value 100 if the property was sold two weeks before the first letter was delivered.²⁰ The next coefficients correspond to a horizon of 1 week prior to the start of the letter delivery, then 1 week after the letter delivery, and so on and so forth until the farthest horizon for which we have data (28 weeks after the start of letter delivery). To facilitate the comparison of timing of read receipts and the effects of the information, the bottom half of Figure 6, panel A shows the evolution of the read-receipt proxy.

The evidence indicates that the timing of the experimental effects is largely consistent with the timing of letter delivery. First, the information contained in the letter should not have any effects prior to the start of the letter delivery, because the subjects had not received the letters yet. As expected, Figure 6, panel A shows estimated coefficients for the two pretreatment horizons (i.e., the two leftmost coefficients) that are close to zero, precisely estimated, and statistically insignificant.

Next, we test two predictions about the timing of the effects in the posttreatment period. First, we expect the effects of the information shocks not to materialize immediately but to build up gradually over the course of weeks as more and more letters are being read. Second, due to the nature of the sales process in the real estate

²⁰This is the earliest horizon we can use, since this is the date when the administrative data were first downloaded to create the letters.

Panel A. Effects on behavior



Panel B. Effects on expectations




FIGURE 5. EFFECTS OF INFORMATION SHOCKS ON EXPECTATIONS AND BEHAVIOR: BINNED SCATTERPLOTS


Notes: Panel A corresponds to a binned scatterplot representation of the results presented in column 3 of Table 2. This binned scatterplot focuses on the key independent variable, *Information Shock* ($E_m^i \cdot D_i$). 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. Panel B is the binned scatterplot representation of the results presented in column 1 of Table 3. This binned scatterplot focuses on the key independent variable, *Information Shock* ($E_m^i \cdot D_i$). 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. Each panel reports the slope with heteroskedasticity-robust standard errors in parentheses.

market, some of the effects of the information shocks may lag a few weeks behind the read receipt of the letters. More precisely, after reading our letter, some sellers may have been sitting on an offer already, but for other sellers it may take one or two weeks to receive their first offer and thus have the opportunity to sell. And even after a seller chooses to accept an offer, due to the standard real estate closing process, it may take one or more weeks for that property to appear as sold in the records of the Multiple Listing Service or the county assessor's office.

Figure 6, panel A shows that these two predictions about the posttreatment coefficients are borne out by the data. The effects of the information shocks build up over time at a similar rate as the read receipt of the letters and with a lag of a couple of weeks. Figure 6, panel B provides a complementary view of the same results from Figure 6, panel A. For each possible weekly horizon, Figure 6, panel B plots our proxy for the share of letters read (in the *x*-axis) against the estimated effect of the information shocks (in the *y*-axis). Again, we see that the effects of the information shocks are close to zero prior to the start of the letter delivery, and then they become negative and grow stronger as more and more subjects read the letters.

Figure 6, panel A also provides some useful information about the persistence of the effects. If the information made some sellers delay their decision by a matter of just a few weeks, we would expect the posttreatment coefficients in Figure 6, panel A to quickly revert back to zero. On the contrary, we find that the effects of the information shocks were highly persistent, remaining as strong at six months after the start of letter delivery as they were at three months after the start of the letter delivery. More precisely, Figure 6, panel A shows a coefficient for 28 weeks later that is negative (-0.295), precisely estimated and statistically highly significant

Panel A. Effects versus time



Panel B. Effects versus readership




FIGURE 6. EFFECTS OF INFORMATION SHOCKS ON BEHAVIOR: EVENT-STUDY ANALYSIS

Notes: Panel A: Each coefficient in the top half of the panel corresponds to a separate regression based on 57,910 subjects from the field experiment. Every regression corresponds to equation (4) from Section I, and the coefficient being graphed corresponds to the coefficient on the key independent variable, *Information Shock* ($E_{it}^i \cdot D_i$). 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 vertical red line indicates the estimated date when the first letter was delivered (June 15, 2019). The bottom half of the panel 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 percent confidence intervals are based on heteroskedasticity-robust standard errors. Panel B: This figure combines the two series shown in panel A, by plotting the treatment effects against the proportion of the letters from the field experiment that had been read at every point in time. The labels at each marker mark the number of weeks after the start of the letter delivery that the effect of the information shock and the share of letters read were recorded at.

(p -value = 0.006). Moreover, this coefficient at the 28-week horizon (-0.295) is similar in magnitude to—and statistically indistinguishable from—the corresponding coefficient for the 12-week horizon (-0.350).

In Section IIIC, we show that the pretreatment characteristics were balanced across the six treatment groups. However, given that our econometric model goes beyond a simple comparison of means, this balance test is useful but not sufficient. In the spirit of Chetty, Friedman, and Rockoff (2014), for a more direct falsification test we reproduce the same regression from the baseline specification (column 3 of Table 2), but using pretreatment characteristics as dependent variables. One challenge with this type of falsification analysis is that the dependent variables have different distributions and thus the coefficients are not directly comparable across different regressions. As in Chetty, Friedman, and Rockoff (2014), we standardize the coefficients by dividing them by the standard deviation of the corresponding dependent variable. Then, we multiply the standardized coefficients by 100, for readability.

The results from these falsification tests are presented in Figure 7. The two standardized coefficients on the left-hand side of the figure correspond to two of the posttreatment outcomes shown in the Figure 6, panel A: whether the property was sold at 12 or 28 weeks after the start of letter delivery. The six standardized coefficients on the right-hand side of the figure are based on the exact same regression

specification, but using placebo dependent variables based on pretreatment characteristics: the number of days the property was listed, the initial listing price, the number of bedrooms, the number of bathrooms, the square footage of the building, and the lot size. Because each of these six placebo outcomes were determined before the letters were mailed, the information shocks should have no effect on them. As expected, the coefficients on *Information Shock* for these falsification outcomes are close to zero, statistically insignificant, and precisely estimated. Moreover, these six placebo coefficients are statistically different from the corresponding effects on the post-treatment outcomes. For example, the standardized coefficients (times 100) are -0.725 ($p\text{-value} = 0.001$) for the sales probability at 12 weeks posttreatment and -0.037 ($p\text{-value} = 0.868$) for the pretreatment number of days listed. Furthermore, the difference between these two coefficients (-0.725 and -0.037) is statistically significant ($p\text{-value} = 0.021$).²¹

B. Additional Robustness Checks

Additional robustness checks are shown in columns 4–11 of Table 2. Columns 4–6 show the results under some alternative regression specifications. The specification from column 4 is identical to that of column 3, except that it includes a host of additional control variables: the number of days the property was on the market prior to the experiment, the initial listing price, a set of four indicator variables for the number of bedrooms, four indicator variables for the number of bathrooms, and the square footage of the building and the lot. Because the treatment is randomized, controlling for additional variables should not significantly affect the coefficient on *Information Shock*. As expected, the point estimate with the additional controls (-0.341 , from column 4) is almost identical to—and statistically indistinguishable from—the baseline coefficient (-0.350 , from column 3).

In the baseline specification (column 3 of Table 2), we control for one dummy variable that indicates whether the information was disclosed to the subject. Because letters disclosed information from different sources, one might worry that the source being presented may have an effect on its own, above and beyond the effect of its signal. For example, when sharing information sources like Past-1 and Past-2, the reader may be prompted to think about the past, and that may have an effect of its own. To address that concern, the specification from column 5 of Table 2 is identical to the baseline specification of column 3, except that instead of controlling for one disclosure indicator, it controls for a set of five disclosure indicators (i.e., one for each information source). The results are virtually identical in this alternative specification to those in the baseline specification. For instance, the coefficient from the alternative specification (-0.349 , from column 5) is almost identical to the coefficient from the baseline specification (-0.350 , from column 3), and their difference is statistically insignificant ($p\text{-value} = 0.999$).

²¹ This equality test between two coefficients is based on the same data but different regressions. To allow for a nonzero 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.




FIGURE 7. EFFECTS OF INFORMATION SHOCKS ON BEHAVIOR: PLACEBO OUTCOMES

Notes: All the regressions shown in this table correspond to equation (4) from Section I, based on data on the 57,910 subjects in the field experiment. Each coefficient corresponds to a separate but identical regression, with the only difference being the dependent variables. We report the coefficient on the key independent variable, *Information Shock* ($E_m^I \cdot D_i$). All coefficients have been standardized by dividing by the standard deviations of the corresponding dependent variable and then multiplying by 100 for readability. Each dependent variable is listed in the x-axis. We use blue circles to denote the posttreatment 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 pretreatment 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. Bedrooms* is the property's number of bedrooms; *No. Bathrooms* 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 percent confidence intervals are based on heteroskedasticity-robust standard errors.

The baseline specification (column 4 of Table 2) controls for market fixed effects. Those fixed effects are meant to isolate the exogenous variation of the source randomization, by controlling flexibly for the set of signals that a given subject could have been assigned to. While this baseline specification is feasible in the context of our large-scale field experiment, where there are multiple subjects per market, it has the limitation that it would not be feasible in smaller datasets. For example, this specification cannot be estimated with our auxiliary survey experiment, for which there is typically one subject per market. Columns 6 of Table 2 is identical to that of column 3, except that instead of the market fixed effects, it controls for the potential feedback in a less data-intensive way: it includes five linear terms, one for each of the five potential signals that the subject could have been assigned to (i.e., one per information source), as well as five sets of dummies breaking those variables in deciles to deal with any potential nonlinearities. The results from this less demanding specification are almost identical to the results from the baseline specification: the coefficient

from the less demanding specification (-0.330 , from column 6) is almost identical in magnitude to—and statistically indistinguishable from—the corresponding coefficient from the baseline specification (-0.350 , from column 3).

In columns 7–11 of Table 2, we explore whether the results are sensitive to dropping observations for each of the five information sources, on a one-by-one basis. Relative to the baseline specification, the resulting coefficients are a bit less precisely estimated, because dropping each treatment group entails throwing away between 15 percent and 20 percent of the observations. However, the results are always similar in magnitude, precisely estimated, and statistically significant. The six coefficients from columns 7–11 (-0.353 , -0.307 , -0.364 , -0.411 , and -0.331) are similar in magnitude to the corresponding coefficient from the baseline coefficient (-0.350 , from column 3), thus confirming that the results are not driven by any single information source.²²

Due to space constraints, some additional robustness checks and results are presented in the online Appendix. In the previous analysis, the outcome variable is the probability that property is sold at a given point in time. Our original preregistered outcome variable was the time elapsed from the start of letter delivery to the sale of the property. As we anticipated in the preregistration, a challenge with this outcome is that it is truncated: for 42 percent of properties in the sample that were not sold by the end of the sample window, we do not know when they will ultimately be sold. By looking at the probability of selling at a given point in time, we avoid having to deal with any sort of truncation bias. In any case, there is a host of duration models that can deal with this type of truncation. We report those results in online Appendix B.1, which are qualitatively and quantitatively similar to the results from the baseline specification presented above.

It is important to note that in practice, different individuals may react differently to changes in their home price expectations, amounting to heterogeneity in the key parameter of interest (π_1). In that case, our estimates would identify the local average treatment effect of information shocks (Imbens and Angrist 1994)—that is, a weighted average of π_1 's with a higher weight given to homeowners whose expectations are more responsive to the information shocks. This feature is not specific to this study but a very common occurrence in information-provision experiments (e.g., see the discussion in Cullen and Perez-Truglia 2022). With this feature in mind, online Appendix C.2 presents a detailed heterogeneity analysis across several characteristics of the owner, the property, and the market. We find similar effects across the board. For example, the point estimates are similar and statistically indistinguishable between male and female homeowners, between more and less expensive homes, or between hotter or colder markets. If anything, there is suggestive evidence of heterogeneity for two subgroups of the population: non-owner-occupied properties and older household owners (however, these differences must be taken

²² A related question is whether the information about the past (Past-1 and Past-2 treatments) was more or less compelling than the information about the forecasts (Forecast-1, Forecast-2, and Forecast-3 treatments). For example, if most subjects have backward-looking expectations, they may be more elastic to information about the past than to the forecasts (Case and Shiller 1989; Shiller 2005). In online Appendix C.1, we provide some suggestive evidence that the backward-looking information may have been more effective than the forecasts—however, that result must be taken with a grain of salt due to lack of sufficient statistical power.

with a grain of salt, as they are borderline statistically insignificant). Our preferred interpretation for these two sources of heterogeneity is that of selling frictions: these types of owners likely face fewer frictions when deciding whether to sell their home or when to sell it, and thus their behavior is more responsive when they revise their expectations (for more details, see online Appendix C.3).

Last, in online Appendix B.2 we show the effects of the information shocks on the listing status (e.g., active or delisted). We find no evidence of significant changes to the probability of delisting. This evidence indicates that sellers with more optimistic expectations kept their properties on the listing website, suggesting that they were still waiting for better offers to come along.

C. Effect on Listing Price

Our preferred interpretation of the effects of the information shocks on the selling decision is that homeowners revised their expectations and, as a result, changed their reservation prices. To gain further insight into this mechanism, it would be ideal to measure the effect of the information shocks on reservation prices. However, reservation prices are unobservable to researchers: only the seller knows how much he or she is *truly* willing to accept for the home. Nevertheless, we can use listing prices as a potential proxy for reservation prices. For example, if a positive information shock increases a homeowner's home price expectations, the homeowner may choose to reflect that optimism in subsequent changes to the listing price.

We scraped weekly listing prices for all properties in our sample.²³ One important fact about listing prices is that they are very sticky: in most cases, the listing prices are never changed, they very rarely increase, and sometimes they decrease, but typically by a bit (e.g., less than 5 percent drop).²⁴ The evolution of average listing price changes is presented in Figure 8, panel A. The *x*-axis denotes the number of weeks since letter delivery, while the *y*-axis shows the corresponding average percentage change in listing price. Listing prices drop over time, but rather slightly. After 28 weeks since the start of letter delivery, the listing price dropped by an average of just 2.85 percent.

We find suggestive evidence that information shocks are affecting homeowners' listing prices. In Figure 8, panel B we reproduce the analysis from the event study presented in our main analysis from Figure 6, panel A but use the percent change in listing price relative to original listing price as the dependent variable instead of a variable indicating whether the property was sold. Each coefficient in Figure 8, panel B corresponds to a different regression. The estimates suggest that a larger information shock has a positive effect on the average listing price change. The effect on listing price is largest during the first 12 weeks following the start of the letter delivery, and it starts to decline slowly afterwards. As a result, the coefficients are statistically significant during the first three months, but they become statistically insignificant during the following three months. For instance, nine weeks after letter

²³ After a property has been sold or delisted, we use the last published listing price available. By the end of our study period, around 60 percent of properties are sold and 20.4 percent are delisted.

²⁴ For more details, see online Appendix B.3.




FIGURE 8. EFFECTS ON LISTING PRICE

Notes: Panel A: This figure presents the average change in the listing price of the properties over time. Each bar averages the listing prices of the 57,910 subjects, expressed as a proportion of the original listing price, at each week relative to the start of the letter delivery. Panel B: Each coefficient corresponds to a separate regression based on 57,910 subjects from the field experiment. Every regression is based on equation (4) from Section I, and the coefficient being graphed corresponds to the coefficient on the key independent variable, *Information Shock* ($E_m^I \cdot D_t$). The dependent variable is the difference between the listing price of the corresponding week and the original listing price, in percentage points. For example, the coefficient on +12 weeks is based on a dependent variable that takes the value -5 if, 12 weeks after the start of the letter delivery, the property's listing price was 5 percent below the original listing price. The vertical red line indicates the estimated date when the first letter was delivered (June 15, 2019). The 90 percent confidence intervals are based on heteroskedasticity-robust standard errors.

delivery, a 10 percent higher information shock increases the average listing price by 0.12 percentage points (p -value = 0.096), which is 4.2 percent ($= \frac{0.12}{2.85} \times 100$) of the corresponding mean outcome.

In online Appendix B.3, we present some additional results and robustness checks. First, we examine what part of the distribution of listing price changes is driving the average effect. We show that there is a small positive effect on the probability of a price increase, no significant effect on unchanged price and for a small decrease in price (< 5 percent), and a significant decrease in the probability of a large price decrease (≥ 5 percent) in the first 12 weeks. This evidence indicates that homeowners are delaying large reductions in listing price (i.e., maintaining the higher initial listing price for longer).

An alternative proxy to homeowner's reservation prices is the sales price that was listed as primary outcome in the preregistration. As detailed in the preregistration, this outcome is also censored: by construction, the sale price is only observed if the property is sold, which was not the case for 42 percent of the properties (i.e., those that were not sold within 28 weeks of the start of the mail delivery). As a result, and despite the random assignment, estimating the effects on the censored variable is subject to potentially severe selection bias. To estimate the treatment effect on sales price, we need to impose additional assumptions. With those caveats in mind, the effects on sales prices are presented in online Appendix B.3. In a nutshell, we find that, consistent with the effect on listing prices, the information shock has a positive effect on sales prices. Indeed, the effect on the sales price seems quantitatively stronger than the effect on listing price. However, the effect on sales prices has to

be taken with a grain of salt: in addition to requiring extra assumptions to deal with censoring, the effects are imprecisely estimated and, thus, statistically insignificant at conventional levels.

V. The Supplemental Survey Experiment

We developed a supplemental survey experiment that randomized the same information included in the field experiment but was designed to be conducted on an auxiliary sample of respondents recruited from an online platform. The goal of this supplemental survey experiment was twofold. The first goal was to provide direct evidence that the information shocks had the intended effects on expectations. The second goal is to quantify the strength of the effects of information shocks on expectations, which we use to scale-up the intention-to-treat estimates.

A. Survey Design

The survey instrument was part of the same RCT preregistration used for the field experiment. Moreover, the supplemental survey was also conducted around the same dates when the main field experiment was conducted.²⁵ Online Appendix G includes the full survey instrument, which is summarized below:

Step 1: (Elicit Property Details)—To provide randomized information relevant to the respondent, we asked respondents about their current residency, such as the number of bedrooms and five-digit zip code.

Step 2: (Elicit Prior Belief)—Respondents were shown the current median home value (in May 2019) for a similar home (same number of bedrooms and zip code) and were asked to report their expectations for that median value one year later (in May 2020).

Step 3: (Information-Provision Experiment)—All respondents were told that some survey participants will be randomly chosen to receive information about home prices. On the following screen, respondents find out the information selected for them. Respondents were assigned to one of the same six treatment groups from the field experiment described in Section IIA.

Step 4: (Elicit Posterior Belief)—On the following screen, subjects were told that all participants can reassess their guess about future home prices, regardless of what they guessed initially or the information that they received from us. We

²⁵The online survey that was included in the letter, which we use to construct the read receipt proxy, included some survey questions about home price expectations (see online Appendix B.3 for the full survey instrument). However, based on previous studies, and as stated in the RCT preregistration, we anticipated that this additional survey data would be inadequate for the analysis. For example, Perez-Truglia and Troiano (2018) conducted a mailing intervention in the context of tax compliance that included a link to an online survey similar to ours, but they received a response rate of just 0.2 percent. In line with that study, the response rate was to the survey link included in the letter was extremely low: 0.79 percent (i.e., 455 out of the 57,910 subjects completed the survey).

reelicited the same question about their expectations of the median price one year later, as well as an additional question on their expectations five years later. As a placebo outcome, respondents were asked about their stock market expectations—i.e., they were told the closing price of the Dow Jones on May 31 and asked what they expected the price to be 12 months later.²⁶

We recruited subjects from Amazon's Mechanical Turk (AMT) online marketplace from June 21 to June 24. We restricted participants to US residents and followed best practices for recruiting participants and ensuring high-quality responses via AMT (Crump, McDonnell, and Gureckis 2013). The attention check was passed by 99.15 percent of respondents, and 92.17 percent of respondents reported that our survey was easy to understand. The final AMT sample included 1,404 respondents. Online Appendix A.6 presents more details about this sample. Consistent with successful random assignment, the observable characteristics were balanced across treatment groups. We also show that, relative to the field experiment sample and the universe of US homeowners, the online sample is younger, is less wealthy, and lives in smaller and less expensive homes.

B. Effects on Survey Expectations

Unlike the field experiment, in the auxiliary survey we *do* observe subjects' prior beliefs. As a result, we can use that information to estimate the updating of expectations as in more traditional information-provision experiments.

Learning models, such as the Bayesian model, predict that respondents who overestimated relative to the signal would revise their beliefs downward when shown the signal, while those who underestimated with respect to the signal would revise their beliefs upward. Figure 9 provides a simple test of this prediction. This figure provides a binned scatterplot between the prior gaps (i.e., the difference between the signal chosen for the subject and the subject's prior belief, in the x -axis) and the belief revisions (i.e., the difference between posterior and prior beliefs, in the y -axis). Intuitively, the x -axis shows the maximum revision we would expect if the respondent were to fully react to the information, and the y -axis shows the observed revision. The red diamonds from Figure 9 correspond to individuals in the treatment group (i.e., those who were shown the signal). As expected, there is a strong relationship between the belief revisions and prior gaps for individuals in the treatment group: an additional 1 pp in perception gap is associated with a 0.545 pp higher revision (p -value < 0.001). In contrast, the blue circles show a much weaker relationship among individuals in the control group: an additional 1 pp in the prior gap is associated with a revision of 0.100 pp (p -value < 0.001). This finding suggests a statistically significant but economically small degree of spurious revision, consistent with the results from other information-provision experiments (Armantier et al.

²⁶The survey included some additional questions that could be used for disentangling mechanisms and for heterogeneity analysis (see online Appendix A.6 for more details). The survey also included a short follow-up module (see online Appendix A.6 for more details).




FIGURE 9. AMT SUPPLEMENTAL SURVEY RESULTS: BELIEF UPDATING

Notes: The binned scatterplot provides a finer description of the belief updating for the 1,404 subjects from the AMT supplemental survey. The dependent variable is the belief update about the expected growth rate of the median home value over the following year (i.e., the gap between the posterior belief that was elicited after the information-provision experiment, and the prior belief that was elicited before the information-provision experiment). The independent variable is the gap between the signal (the information about home prices provided to the participant) and the prior belief about the expected home price growth rate. The blue circles represent the subjects that received no additional information about home price growth. The red diamonds represent the subjects that received additional information about home price growth. β is the slope of each line, estimated with a simple linear regression with heteroskedasticity-robust standard errors reported in parentheses.

2016; Cavallo, Cruces, and Perez-Truglia 2017; Bottan and Perez-Truglia 2022; Fuster et al. 2022; Cullen and Perez-Truglia 2022).

The results from Figure 9 indicate that when receiving a signal, the average subject is updating posterior belief, but not fully. There are many reasons why subjects, even if fully rational, may update partially or may even not update at all. For example, the subject may have already seen the information or may feel more confident in their prior beliefs than in the signal. Indeed, the degree to which individuals incorporated the information provided in the experiment is comparable in magnitude with the findings from other survey experiments. The most direct comparison is with respect to Fuster et al. (2022), who conducted a survey experiment in the 2017 Federal Reserve Bank of New York’s Survey of Consumer Expectations. They randomly provided subjects with signals from similar information sources as those used in this paper (e.g., price changes over recent years and professional forecasts for the next year) and then measured the effects of the information provision on the subsequent home price expectations. The results from our paper are not expected to be identical to those in Fuster et al. (2022), because of differences in implementation and survey design. For example, Fuster et al. (2022) concerns expectations about national home prices, while our supplemental survey revolves around local home prices. There are even more significant differences in how the respondents were recruited: while our supplemental survey relies on a convenience sample from

AMT, Fuster et al. (2022) is based on a nationally representative survey conducted by the Federal Reserve Bank of New York.

Despite these significant differences in methods and the sample, the results from our supplementary survey experiment are quantitatively similar to the results from Fuster et al. (2022). According to a simple Bayesian learning model, Fuster et al. (2022) estimate that their average subject forms posterior beliefs by assigning a weight of 0.380 to the signal and the remaining weight of 0.620 to the prior belief.²⁷ In comparison, when applied to our own supplementary survey, the same model indicates that the subject forms posterior beliefs by assigning a weight of 0.445 to the signal.²⁸ These two learning rates (0.380 and 0.445) are in the same ballpark. Moreover, the rate of learning we estimated in our survey experiment is also in the same order of magnitude as the corresponding estimates in studies that involve different types of macroeconomic expectations. For instance, Cavallo et al. (2017) reports that when forming inflation expectations, the average Argentine respondent assigns a weight of 0.432 to the signal about inflation.²⁹

Figure 9 presents a detailed analysis of belief updating, which relies heavily on the data on prior beliefs. For a more direct comparison to the results of the main field experiment, Table 3 presents the results of the supplemental survey experiment using the exact same regression specification used in the main field experiment, which does not require any data on prior beliefs. More precisely, all regressions in Table 3 use the exact same econometric specification as in column 6 of Table 2, but with different dependent variables.³⁰ In column 1, the dependent variable is the posterior belief about the home price expectation one year ahead (H_{1y}^{post}). The coefficient on *Information Shock* from column 1 is positive (0.205) and statistically significant (p -value = 0.001). A 1 pp increase in the information shock causes an increase in home price expectations of 0.205 pp. In other words, there is a 20.5 percent pass-through from the information shocks to the expectations. The fact that the coefficient on *Information Shock* is significantly greater than zero implies that the subjects found the information provided in the experiment relevant to form their home price expectations. The fact that the coefficient on *Information Shock* is significantly less than one is due to a variety of factors, such as the fact that some subjects felt confident about their prior beliefs, already knew the information, or were just not paying enough attention to the survey.

Columns 2 through 4 of Table 3 present the effects on other survey outcomes. The dependent variable in column 2 is identical to the dependent variable in column 1, except that it corresponds to the expectation for five years ahead instead of one year ahead. The coefficient on *Information Shock* is positive (0.167) and statistically significant (p -value = 0.017). This result implies that the information shocks affected not only the more immediate expectations (one year ahead) but also

²⁷ This estimate corresponds to the slope reported in Figure 3, panel A of Fuster et al. (2022).

²⁸ This estimate is based on the difference in slopes between the treatment and control groups from Figure 9—for more details about this methodology, see Fuster et al. (2022).

²⁹ This estimate corresponds to coefficient α -statistics reported in panel B, column 1 of Table 1 from Cavallo et al. (2017).

³⁰ As explained in the above section, we cannot use the specification from column 3 of Table 2, because the supplemental survey does not have multiple subjects with properties of the same type and in the same zip code.

TABLE 3—AMT SUPPLEMENTAL SURVEY RESULTS

	H_{1y}^{post}	H_{5y}^{post}	H_{1y}^{prior}	M_{1y}^{post}
	(1)	(2)	(3)	(4)
<i>Information Shock</i>	0.205 (0.064)	0.167 (0.070)	−0.014 (0.066)	0.017 (0.134)
Mean outcome	3.86	2.31	3.88	3.58
SD outcome	4.42	4.36	5.39	9.05
R^2	0.140	0.044	0.028	0.043
Observations	1,404	1,404	1,404	1,404

Notes: Heteroskedasticity-robust standard errors in parentheses. Each column corresponds to a different regression based on data from the AMT supplemental survey using a different dependent variable: 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 the posterior belief about the annualized expected 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. And M_{1y}^{post} is the posterior belief about the annualized expected growth rate of the stock market index over the following year.

the more distant expectations (five years ahead). The point estimate from column 2 is smaller in magnitude than the point estimate from column 1, suggesting that the information shocks affected the one-year-ahead expectations more strongly than the five-years-ahead expectations; however, that comparison must be taken with a grain of salt because the difference between the coefficients from column 1 and 2 is statistically insignificant (p -value = 0.560).

Columns 3 and 4 of Table 3 present some falsification tests. The dependent variable from column 3 is identical to that in column 1, except that while column 1 corresponds to the posterior belief, column 3 refers to the corresponding prior belief (i.e., the home price expectations elicited *before* the information-provision experiment). Since information shock had not been administered to the subject yet, it should have no effect on prior beliefs. 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 effect on posterior beliefs (0.205, from column 1) is equal to the falsification coefficient (−0.014, from column 3), with a p -value < 0.001.

Column 4 of Table 3 presents another falsification test. The dependent variables in columns 1 and 4 are both posterior beliefs (i.e., elicited after the information-provision experiment). However, whereas the dependent variable from column 1 corresponds to the home price expectations, the dependent variable from column 4 corresponds to stock market expectations. Because the information shock is specific to local home prices, we do not expect individuals to extrapolate this information to the stock market expectations.³¹ As expected, the coefficient on *Information Shock* from column 4 is close to zero (0.017), statistically insignificant (p -value = 0.899), and precisely estimated.

³¹ If, on the contrary, subjects did extrapolate to stock market expectations, that would suggest spurious motives behind the belief updating, such as numerical anchoring or experimenter demand (Cavallo, Cruces, and Perez-Truglia 2017).

As a final robustness check, Figure 5, panel B presents a binned scatterplot version of the results from column 1 of Table 3. As discussed in Section I above, the baseline specification assumes that the effects of the information shocks are linear and symmetric around zero. The results from Figure 5, panel B indicate that the linear specification fits the data quite well and also confirms that the results are not driven by nonlinearities or by outliers.

VI. Magnitude of the Effects

The results from the main field experiment presented in Section IV indicate an elasticity between the information shocks and the probability of selling the property within 12 weeks of -0.350 (column 3 of Table 2). However, the magnitude of this elasticity must be taken with a grain of salt, because it corresponds to an intention-to-treat effect. In this section, to better assess the true magnitude of the effects, we use complementary data to scale up from intention-to-treat effects to treatment-on-the-treated effects.

A. First Form of Noncompliance

The most basic form of noncompliance is that some subjects may not read the letter and, thus, whatever information was included in the letter will have no chance of changing their expectations. To correct for this form of noncompliance, following Perez-Truglia and Cruces (2017) and others, we need to estimate the *reading rate*—i.e., the share of recipients who *actually* read the letter and on time. For that, we need to figure out the share of letters that were not delivered, were delivered but not read, or were read but after it was too late (i.e., after the owner had already accepted an offer). We estimate each of these shares using various sources of data. The main estimates are summarized below, while the details are discussed in online Appendix A.7. According to the US Monitor nonprofit standard mail delivery study, around 5 percent of standard nonprofit mailers fail to be delivered (US Monitor 2014). Based on data from the US Postal Service Household Diary Survey (Mazzone and Rehman 2019), conditional on delivery, around 26 percent of the letters were ignored. And based on the timing of survey responses and home sales in the baseline group, we estimate that roughly 7.7 percent of the letters were not read until after the property was sold. Combining these three estimates leads to a reading rate of 64.9 percent ($= 0.95 \cdot 0.74 \cdot 0.923$). Following Perez-Truglia and Cruces (2017), to account for the first source of noncompliance, we can divide the intention-to-treat elasticity (-0.350 , from column 3 of Table 2) by the reading rate. This adjustment results in a scaled-up elasticity of -0.539 ($= -0.350/0.649$).

B. Second Form of Noncompliance

The second source of noncompliance is that even if they read the letter, the information shocks introduced in the letter may not change the home price expectations of the recipients. To address this second form of noncompliance, we can combine the estimates from the main field experiment and the supplemental survey experiment.

An important caveat, however, is that this exercise relies on the assumption that the reaction to the information in the supplemental survey experiment is a reasonable approximation of the reaction to the information in the main field experiment. Indeed, there are multiple reasons for why, in theory, the reaction to the information may have been somewhat stronger or weaker in one context than in the other.³² As discussed in Section V above, however, the degree to which subjects updated their expectations in our supplemental survey is remarkably similar to the corresponding degrees in related studies concerning not only home price expectations but also other types such as expectations about the inflation rate or the GDP growth. Indeed, even within a given sample, it has been documented that subjects update their expectations to a similar degree regardless of their characteristics such as gender, income, and education (Fuster et al. 2022). For these reasons, it is not unreasonable to extrapolate the findings from the supplementary survey to the main field experiment.

With these caveats in mind, we proceed with combining the two sets of results. The results from the field experiment indicate that a 1 pp higher information shock decreased the probability of selling the property within 12 weeks by 0.350 pp (column 3 of Table 2). The survey experiment indicates that a 1 pp higher information shock increased home price expectations by 0.205 pp (column 1 of Table 3). Jointly, these two results imply that a 1 pp increase in home price expectations causes a decline of 1.71 pp ($= 0.350/0.205$) in the probability of selling the property within 12 weeks—i.e., a behavioral elasticity between expectations and sales probability of -1.71 .³³ Moreover, if we adjust for the two forms of noncompliance simultaneously, we arrive to a scaled-up behavioral elasticity of -2.63 ($= -0.350/(0.205 \cdot 0.649)$), which suggests that sellers are highly elastic.

C. Other Forms of Noncompliance

There are a couple of additional sources of attenuation bias that are worth mentioning but that we do not adjust for. To the extent that it does not account for these additional sources of attenuation bias, the elasticity of -2.63 reported above constitutes a conservative estimate.

One form of attenuation bias is that the home price expectations may be irrelevant for the decisions of some sellers. More precisely, if, after selling the current property, the seller plans to buy another property in the same location, then the seller will be exposed to the same home price appreciation regardless of when the current property is sold or whether it is sold at all. Thus, for these sellers, their home price expectations should not be relevant to their selling decision. However, we present some data indicating that the scope for this form of attenuation is probably minor because only a small minority of subjects would be subject to it. For instance, using

³² On the one hand, maybe the subjects in the supplemental survey paid less attention to the information than the subjects in the field experiment, because they do not have strong financial incentives to stay informed on local home prices. On the other hand, perhaps the subjects in the supplemental survey paid more attention to the information because it is shown on the computer screen seconds before the elicitation of the posterior expectations.

³³ In practice, the decision to sell probably depends jointly on expectations at many horizons (e.g., 1 month, 12 months, 5 years). This back-of-the-envelope calculation focuses exclusively on the one-year horizon, because we do not have sufficient data or exogenous variation to disentangle the contribution of the different horizons.

the administrative data from the county assessor's records, we find that after selling the properties, only 8 percent of subjects bought another property in the same county.³⁴ We reach estimates of the same order of magnitude using an alternative data source: mail forwarding data from the US Postal Service. Moreover, our estimates are also in the same ballpark as the estimates reported by other analysts using independent sources of data (Umbrasas 2019). Under the assumption that the above 8 percent of subjects do not care about home price expectations, we would need to further scale-up our estimates by a factor of just 1.087 ($= 1/0.92$).

Another potential source of attenuation bias relates to spillovers in the information treatments. For example, one subject who received a signal from the source Past-1 could share that signal with a subject who received information from the source Forecast-1. In our baseline regression model, this type of information spillovers would lead to an attenuation bias. There are some contexts in which the room for information spillovers can be significant, such as distributing information among classmates or coworkers (Cullen and Perez-Truglia 2022). In the context of our housing experiment, however, the scope for spillovers is negligible. We distributed information to a tiny share of all the households who live in the counties where the experiment was conducted. As a result, the subjects who did receive information were geographically distant from each other. These subjects are highly unlikely to know each other, much less to have opportunities to share the information included in our letters.³⁵ Moreover, the type of information that we included in the letter was tailored to each subject (i.e., to their specific zip code and property type), which limits the scope of spillovers even further. While there are standard methods for measuring and accounting for spillovers in an experiment like ours (Cullen and Perez-Truglia 2022; Nathan, Perez-Truglia, and Zentner 2020), in our case they have a negligible effect on the estimates. In addition to the direct effects on individual behavior, beliefs could have equilibrium effects—for example, if shocks to expectations are correlated (Bailey, Cao, et al. 2018). While we did not design our experiment to explore equilibrium effects, it could be possible to adjust the research design to explore that issue.

VII. Conclusions

We provide the first experimental evidence on the causal effects of home price expectations on market choices by conducting a large-scale, high-stakes, natural field experiment involving 57,910 US individuals who listed homes on the market. To generate exogenous shocks to their home price expectations, we send each of these homeowners a letter containing randomized information about their home prices. We then use administrative records to measure the effects of those information shocks on the recipients' subsequent sales outcomes. Consistent with economic theory, we find that higher home price expectations cause homes to take longer to sell. Moreover, the magnitude of this reaction indicates that market choices are

³⁴ For details, see online Appendix B.4.

³⁵ For more details, see online Appendix B.5.

highly elastic to expectations: a 1 pp increase in home price expectations reduces by 2.63 pp the probability that the property will sell in the following 12 weeks.

Our findings have implications for macroeconomic models. The full-information rational expectations assumption is at the core of modern macroeconomics. There is, however, growing support for alternative expectation models (Case and Shiller 1989). For example, these new models play an important role in some accounts of the 2008 US housing crisis (Shiller 2005). Despite growing interest, many observers have expressed skepticism about deviations from the rational expectations framework (Guren 2018). According to this view, the stakes of housing transactions are so large that households could not possibly have biased expectations. We provide experimental evidence that providing households with publicly available information on the evolution of home prices has substantial effects on their market choices. This constitutes direct evidence that despite the high stakes, information frictions are significant.

Our finding that subjective expectations are consequential for market choices are consistent with the results from structural models suggesting that subjective expectations play a key role in the housing market. For example, Kaplan, Mitman, and Violante (2020) estimated a structural model including shocks to aggregate productivity, shocks to credit, and shocks to beliefs about future housing demand. They found that home price expectations are the dominant force behind the observed swings in home prices. Landvoigt (2017) estimated a structural model of the housing market and found that subjective disagreement about future price growth plays a key role.

Moreover, our key parameter of interest, the causal effect of subjective expectations on market choices, plays a key role in models of housing bubbles. In these models, bubbles are generated by exogenous shocks to buyers' expectations of home values. These models are based on the assumption that the average household can be systematically overoptimistic about future growth in home prices (Glaeser, Gyourko, and Saiz 2008) and, also, that some households may be more optimistic than other households (Nathanson and Zwick 2018). These subjective expectations affect the housing market equilibrium through their effects on the households' market choices. For example, an optimistic shock to a homeowner would make them less likely to sell her home at a given price. Thus, an optimistic shock to homeowners' expectations shifts the supply curve to the left. In sum, the significance of subjective expectations in the housing market hinges on their influence on market choices.

Our evidence suggests that subjective expectations are indeed consequential. To illustrate this point, consider the well-known fact that households disagree substantially about the growth of future home prices (e.g., Fuster et al. 2022). For instance, in our own survey, two households in the same market disagree about the future growth rate of home prices on average by 5.56 pp.³⁶ If households do not act on their expectations, this disagreement would have no effect on the housing market

³⁶More precisely, we are referring to the average absolute difference of prior expected annual growth rates between a random pair of households in the same zip code and with the same number of bedrooms, calculated with responses to our supplemental survey.

equilibrium. On the contrary, the large elasticity that we document (between market choices and home price expectations) implies that disagreement between households is quite consequential. The average disagreement in expectations of 5.56 pp will cause a difference in the probability of selling the property in 12 weeks by 14.6 pp ($= 2.63 \cdot 5.56$).³⁷

Lastly, we attempt to make a methodological contribution to the use of field experiments in economics. The popularity of field experiments has grown substantially over the last few decades (Card, DellaVigna, and Malmendier 2011). These experiments explore a variety of economic choices, including hiring, tax avoidance, and even voting. Although selling a home may be one of the most important decisions a person makes in their financial life, there has not been a way to analyze this choice through a field experiment. Our experimental framework can fill that gap. Although this study focuses on the role of home price expectations, our experimental framework could be adapted to study all kinds of factors that can affect the decision to sell a home, such as the value of local amenities, property taxes, or the quality of government services, to name a few.³⁸ Indeed, our experimental design has several conceptual and practical advantages that could warrant its widespread adoption. The behavioral effects are measured using objective data from administrative records. The effects are measured in a naturally occurring context, in large-scale settings and based on high-stakes choices. All the data sources used for the experiment are publicly available, so there is no need to acquire expensive datasets or negotiate Data User Agreements. The experiment is relatively inexpensive, costing less than \$0.25 per subject, and could be scaled up to one million subjects.³⁹

REFERENCES

- Anenberg, Elliot.** 2016. "Information Frictions and Housing Market Dynamics." *International Economic Review* 57 (4): 1449–79.
- Armantier, Olivier, Wändi Bruine de Bruin, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar.** 2015. "Inflation Expectations and Behavior: Do Survey Respondents Act on Their Beliefs?" *International Economic Review* 56 (10): 505–36.
- Armantier, Olivier, Scott Nelson, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar.** 2016. "The Price Is Right: Updating Inflation Expectations in a Randomized Price Information Experiment." *Review of Economics and Statistics* 98 (3): 503–23.
- Armona, Luis, Andreas Fuster, and Basit Zafar.** 2019. "Home Price Expectations and Behaviour: Evidence from a Randomized Information Experiment." *Review of Economic Studies* 86 (4): 1371–1410.
- Bailey, Michael, Ruiqing Cao, Theresa Kuchler, and Johannes Stroebel.** 2018. "The Economic Effects of Social Networks: Evidence from the Housing Market." *Journal of Political Economy* 126 (6): 2224–76.
- Bailey, Michael, Eduardo Dávila, Theresa Kuchler, and Johannes Stroebel.** 2018. "House Price Beliefs and Mortgage Leverage Choice." *Review of Economic Studies* 86 (6): 2403–52.

³⁷ Likewise, the large effect of expectations on market behavior implies that any changes in consumer sentiment over time can become powerful drivers in the housing market.

³⁸ This setting may be attractive to researchers from diverse fields, from urban economics to real estate, finance, and behavioral economics. For instance, researchers could teach homeowners about the misaligned incentives between home sellers and their realtors, and then measure whether the information affects how quickly the homes are being sold.

³⁹ For example, Zillow Research (<https://www.zillow.com/research/data/>) shows 1,705,251 unique active listings in May 2019 (and county assessor records likely include data for most of these listings). Additionally, according to data from the National Association of Realtors, 5.34 million homes were sold in 2019.

- Bérgolo, Marcelo L., Rodrigo Ceni, Guillermo Cruces, Matías Giacobasso, and Ricardo Perez-Truglia. 2023. "Tax Audits as Scarecrows: Evidence from a Large-Scale Field Experiment." *American Economic Journal: Economic Policy* 15 (1): 110–53.
- Bernanke, Ben S. 2007. "Inflation Expectations and Inflation Forecasting." Speech, Monetary Economics Workshop of the NBER Summer Institute, Cambridge, MA, July 10, 2007.
- Bottan, Nicolas, and Ricardo Perez-Truglia. 2019. *Housing Expectations and Market Behavior*. AEA RCT Registry. <https://doi.org/10.1257/rct.3663-3.0>.
- Bottan, Nicolas L., and Ricardo Perez-Truglia. 2022. "Choosing Your Pond: Location Choices and Relative Income." *Review of Economics and Statistics* 104 (5): 1010–27.
- Bottan, Nicolas, and Ricardo Perez-Truglia. 2025. *Data and Code for: "Betting on the House: Subjective Expectations and Market Choices."* Nashville, TN: American Economic Association; distributed by Inter-university Consortium for Political and Social Research, Ann Arbor, MI. <https://doi.org/10.3886/E196421V1>.
- Brooks, David. 2017. "The Home Buying Decision." *New York Times*, January 6. <https://www.nytimes.com/2017/01/06/opinion/the-home-buying-decision.html>.
- Card, David, Stefano Della Vigna, and Ulrike Malmendier. 2011. "The Role of Theory in Field Experiments." *Journal of Economic Perspectives* 25 (3): 39–62.
- Carroll, Christopher D. 2003. "Macroeconomic Expectations of Households and Professional Forecasters." *Quarterly Journal of Economics* 118 (1): 269–98.
- Carter, Shawn M. 2019. "You're Not Alone If You Cried When You Sold Your Home—Nearly 40% of Americans Felt the Same Stress." CNBC.com, July 9. <https://www.cnbc.com/2019/07/09/almost-40percent-of-americans-cried-while-selling-their-home.html>.
- Case, Karl E., and Robert J. Shiller. 1989. "The Efficiency of the Market for Single-Family Homes." *American Economic Review* 79 (1): 125–37.
- Cavallo, Alberto, Guillermo Cruces, and Ricardo Perez-Truglia. 2017. "Inflation Expectations, Learning, and Supermarket Prices: Evidence from Survey Experiments." *American Economic Journal: Macroeconomics* 9 (3): 1–35.
- Chetty, Raj, John N. Friedman, and Jonah E. Rockoff. 2014. "Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates." *American Economic Review* 104 (9): 2593–2632.
- Coibion, Olivier, Yuriy Gorodnichenko, and Saten Kumar. 2018. "How Do Firms Form Their Expectations? New Survey Evidence." *American Economic Review* 108 (9): 2671–2713.
- Coibion, Olivier, Dimitris Georgarakos, Yuriy Gorodnichenko, Geoff Kenny, and Michael Weber. 2024. "The Effect of Macroeconomic Uncertainty on Household Spending." *American Economic Review* 114 (3): 645–77.
- Crump, Matthew J. C., John V. McDonnell, and Todd M. Gureckis. 2013. "Evaluating Amazon's Mechanical Turk as a Tool for Experimental Behavioral Research." *PLOS ONE* 8 (3): e57410.
- Cullen, Zoë, and Ricardo Perez-Truglia. 2022. "How Much Does Your Boss Make? The Effects of Salary Comparisons." *Journal of Political Economy* 130 (3): 766–822.
- Eggleston, Jonathan, Donald Hays, Robert Munk, and Briana Sullivan. 2020. *The Wealth of Households, 2017*. Washington, DC: US Department of Commerce.
- Federal Reserve Bank of St. Louis. 2019. *Table 1: New Houses Sold and For Sale*. <https://fred.stlouisfed.org/release/tables?rid=97&eid=22695> (accessed August 25, 2022).
- Fuster, Andreas, Ricardo Perez-Truglia, Mirko Wiederholt, and Basit Zafar. 2022. "Expectations with Endogenous Information Acquisition: An Experimental Investigation." *Review of Economics and Statistics* 104 (5): 1059–78.
- Gennaioli, Nicola, and Andrei Shleifer. 2018. *A Crisis of Beliefs: Investor Psychology and Financial Fragility*. Princeton, NJ: Princeton University Press.
- Glaeser, Edward L., and Charles G. Nathanson. 2015. "Housing Bubbles." In *Handbook of Regional and Urban Economics*, Vol. 5, edited by Gilles Duranton, J. Vernon Henderson, and William C. Strange, 701–51. Amsterdam: Elsevier.
- Glaeser, Edward L., and Charles G. Nathanson. 2017. "An Extrapolative Model of House Price Dynamics." *Journal of Financial Economics* 126 (1): 147–70.
- Glaeser, Edward L., Joseph Gyourko, and Albert Saiz. 2008. "Housing Supply and Housing Bubbles." *Journal of Urban Economics* 64 (2): 198–217.
- Guren, Adam M. 2018. "House Price Momentum and Strategic Complementarity." *Journal of Political Economy* 126 (3): 1172–1218.
- Imbens, Guido W., and Joshua D. Angrist. 1994. "Identification and Estimation of Local Average Treatment Effects." *Econometrica* 62 (2): 467–75.

- Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante.** 2020. "The Housing Boom and Bust: Model Meets Evidence." *Journal of Political Economy* 128 (9): 3285–3345.
- Landvoigt, Tim.** 2017. "Housing Demand during the Boom: The Role of Expectations and Credit Constraints." *Review of Financial Studies* 30 (6): 1865–1902.
- Mazzone, John, and Samie Rehman.** 2019. *The Household Diary Study: Mail Use and Attitudes in FY 2018*. Washington, DC: United States Postal Service. <https://prc.arkcase.com/portal/docket-search/advanced/filing-details/63201> (accessed March 28, 2020).
- Nathan, Brad, Ricardo Perez-Truglia, and Alejandro Zentner.** 2020. "My Taxes Are Too Darn High: Why Do Households Protest Their Taxes?" NBER Working Paper 27816.
- Nathanson, Charles G., and Eric Zwick.** 2018. "Arrested Development: Theory and Evidence of Supply-Side Speculation in the Housing Market." *Journal of Finance* 73 (6): 2587–2633.
- Perez-Truglia, Ricardo, and Guillermo Cruces.** 2017. "Partisan Interactions: Evidence from a Field Experiment in the United States." *Journal of Political Economy* 125 (4): 1208–43.
- Perez-Truglia, Ricardo, and Ugo Troiano.** 2018. "Shaming Tax Delinquents." *Journal of Public Economics* 167: 120–37.
- Roth, Christopher, and Johannes Wohlfart.** 2020. "How Do Expectations about the Macroeconomy Affect Personal Expectations and Behavior?" *Review of Economics and Statistics* 102 (4): 731–48.
- Ruggles, Stephen, Sarah Flood, Matthew Sobek, Daniel Backman, Annie Chen, Grace Cooper, Stephanie Richards, et al.** 2024. *IPUMS USA: Version 15.0 [dataset]*. Minneapolis, MN: IPUMS. <https://doi.org/10.18128/D010.V15.0> (accessed June 19, 2024).
- Shiller, Robert J.** 2005. *Irrational Exuberance*. Princeton, NJ: Princeton University Press.
- Stroebel, Johannes.** 2016. "What Can Housing Markets Teach Us about Economics?" *NBER Reporter* 4: 21–25.
- Umbrasas, Kristijonas.** 2019. "How Much Are Closing Costs for Sellers?" Zillow.com, November 26. <https://www.zillow.com/learn/closing-costs-for-sellers/>.
- US Monitor.** 2014. *7 Myths of Direct Mailing*. New York: US Monitor. <https://www.targetmarketingmag.com/promo/7MythsofDM.pdf> (accessed March 28, 2020).
- Zillow.** 2019. *Consumer Housing Trends Report 2019*. Zillow, Inc. https://wp-tid.zillowstatic.com/28/NewConstruction_CHTR_2019-1702f4.pdf.