

Discussion of  
“What Have They Been Thinking?  
Homebuyer Behavior in Hot and Cold Markets:  
A Ten-Year Retrospect”  
by Robert J. Shiller and Anne K. Thompson \*

Adam M. Guren  
Boston University

May 2022

---

\*guren@bu.edu. Prepared for the Spring 2022 edition of the Brookings Papers on Economic Activity.

It is unusual to be asked to discuss a paper that is already a classic, but that is the predicament in which I find myself. This paper is a useful update on the path-breaking and influential work in the authors' prior paper, Case, Shiller, and Thompson (2012). The natural place to start discussing this 10-year retrospective is thus with a 10-year retrospective discussion of the original paper. Explaining why it has been so influential and what its impact means for the current paper will give me an opportunity to discuss its methodology and the related literature and provide a brief user's guide to their data. I will then turn to discussing the analysis of the last 10 years in housing markets and particularly the current pandemic housing market.

My overall message is that the authors should be applauded for their important contribution: The field of housing economics is unquestionably better due to their adding survey expectations to our toolkit. I also think their big idea—that high long-run expectations can be used to diagnose a housing bubble, much like a yield curve inversion is used to predict a recession—is a useful one, although I think that survey evidence on expectations should be used in conjunction with other evidence rather than on its own.

## **1 Survey Evidence on House Price Expectations**

### **1.1 The Influence of Case, Shiller, and Thompson (2012)**

When the authors of this paper began surveying home buyers on their expectations about the future path of house prices in the late 1980s, the idea that one would ask economic agents about their behavior and expectations was outlandish. Even in 2003, when they revived the survey, it was novel. Today, however, survey evidence on expectations is widespread and accepted as a crucial tool.

The authors deserve a great deal of credit for pioneering and legitimizing survey evidence on expectations in housing markets. Their survey was to my knowledge the first to

go beyond the Michigan Surveys of Consumers' question on whether it is a good or bad time to buy a house and actually ask for expectations of house price growth over various horizons as well as buyers' subjective views about the state of the market. Their work helped demonstrate the value of survey evidence for understanding housing markets and cycles and make its use commonplace, although survey evidence must be taken with a grain of salt and evaluated carefully.<sup>1</sup>

The original paper helped launch a large literature, which has been recently and comprehensively surveyed by Kuchler, Piazzesi, and Stroebel (2022). Given space constraints, I only note a few highlights relevant to the current paper here, leaving the interested reader to turn to Kuchler and others (2022) for details.

First, several key observations that Case, Shiller, and Thompson (2012) made in their initial paper, which included only 10 years of data for four metro areas, have been shown to be key features of expectations after more comprehensive analysis. In particular, Case, Shiller, and Thompson hypothesized that there was the under-reaction of short-term (1 year ahead) expectations and over-reaction of longer-run (10 years ahead) expectations, a finding that shows up again in Section V of the this update. Recently, Armona, Fuster, and Zafar (2019) used a novel informational experiment to show convincingly and causally that short-run expectations under-predict the degree of short-run momentum and long-run expectations do not fully account for mean reversion in house prices. Similarly, Case, Shiller, and Thompson postulated that sentiment about house prices spreads the media and by word of mouth. Bailey et al. (2018) validated this using Facebook data, showing that individuals' expectations are formed in part by the price appreciation of their out-of-town friends. Overall, with 10 years of hindsight, the original paper seems uncannily

---

<sup>1</sup>Influential work by Coibion and Gorodnichenko that uses survey evidence on inflation expectations to discriminate between models also played an important role in popularizing the use of survey evidence outside of housing markets. The wider acceptance of survey evidence in macroeconomics also played a role in its growing acceptance in housing economics.

prescient.

Second, since Case, Shiller, and Thompson (2012), surveys like theirs have proliferated both in the United States and abroad, which is important for two reasons. First, having more surveys and additional countries provides both more data to do the type of analysis for which the authors advocate and a richer baseline set of facts. Second, in the United States there are now high-quality surveys of house price expectations by the Michigan Surveys of Consumers and the Federal Reserve Bank of New York which come out monthly and with very little lag. These surveys make data on house price expectations a viable real-time tool for policy makers and economic forecasters. Case, Shiller, and Thompson (2012) famously pointed out that in the 2000s boom, long-run expectations ballooned to the point that they were higher than short-run expectations, which they said indicated a bubble. The presence of these sorts of data will hopefully aid in identifying bubbles as they occur.

Finally, Case, Shiller, and Thompson (2012) helped motivate a literature that uses non-standard expectations to explain the 2000s cycle. At this point, essentially every legitimate explanation of the 2000s housing cycle ascribes a significant role to overoptimistic or out-of-line expectations, although there remains disagreement on the relative role of out-of-line expectations relative to other explanations like a credit supply expansion and on the source of the out-of-line expectations. Furthermore, many papers use their data to help discipline their explanations and models of what happened in the boom and bust (see, e.g., Burnside, Eichenbaum, and Rebelo, 2016, Kaplan, Mitman, and Violante, 2020, Chodorow-Reich, Guren, and McQuade, 2022) and to discriminate between various models of non-rational beliefs, a point to which I return below.

## 1.2 What Is Unique About Their Survey?

The success of the original paper in inspiring several similar surveys somewhat limits the novelty of the findings about the last 10 years in this update. While the authors point out that their survey is unique and preferable to others because it has the longest panel and because it covers recent active market participants rather than the public at large, the surveys from the Federal Reserve Bank of New York and Michigan Survey of Consumers come out more frequently, are weighted to be representative of a full population, and use modern best practices in survey design.

Figure 1 compares these three data sources. The top panel shows 1-year expectations and the bottom panel shows long-run expectations. In both panels, the thick solid lines show data from the Case-Shiller-Thompson (CST) data in this paper (I add the late and great Chip Case’s name to the data to acknowledge his contribution), the medium dashed lines show data from the Federal Reserve Bank of New York Survey of Consumer Expectations (NY Fed), and the thin dash-dot lines show data from the Michigan Surveys of Consumers (Michigan).

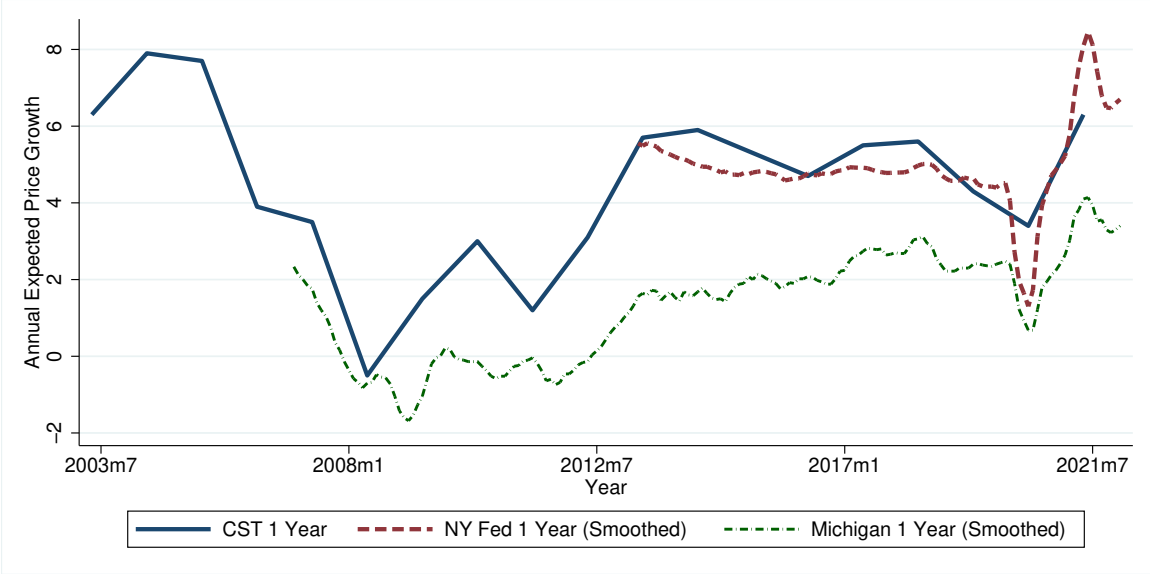
These three surveys have different survey methodologies, phrase the questions differently, ask about various time horizons, and survey different groups. These differences in survey design and phrasing can matter immensely. Specifically, CST mail surveys to a random selection of recent home buyers in 4 distinct markets in the spring of each year asking them for 1-year and 10-year ahead annual average appreciation beginning in 2003.<sup>2</sup> The NY Fed uses an internet survey of a nationally-representative group of household heads and asks them about national house price appreciation in the next year and from 24 to

---

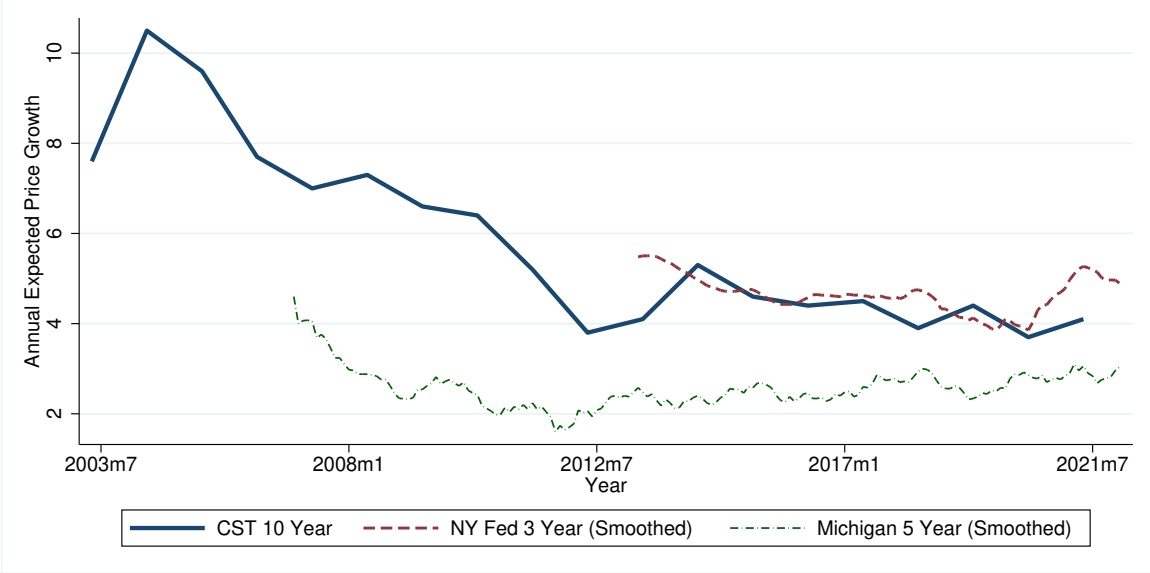
<sup>2</sup>CST ask “how much of a change do you expect there to be in the value of your home over the next 12 months” and “On average over the next ten years how much do you expect the value of your property to change **each year?**” (Their emphasis, which was added starting in the 2012 survey.) They have a response rate of 12 to 35%. Their survey asks only about house prices.

Figure 1: Comparison of Case-Shiller-Thompson, NY Fed, and Michigan Survey House Price Expectation Surveys

(a) One-Year Expectations



(b) Long-Run Expectations



Notes: The top panel shows 1-year ahead house price expectations and the bottom panel shows long-run house price expectations from Case-Shiller-Thompson (thick solid lines), the NY Fed Survey of Consumer Expectations (medium dashed lines), and Michigan Survey of Consumers (thin dash-dot lines). The Case-Shiller-Thompson data is from Table 3 of the paper under discussion and the longer-run expectation is average annual house price growth over the next 10 years. The NY Fed data is monthly data smoothed using a 5-year moving average, and the longer-run expectation is price growth between 24 and 36 months from now. The Michigan survey data are monthly data smoothed using a 5-year moving average, and the longer-run expectation is average annual price growth over the next five years. Data are accurate as of February 2022.

36 months from the survey date beginning in 2014.<sup>3</sup> The Michigan survey is a nationally-representative telephone survey of households that asks them about the appreciation of “homes like yours in your community” over the next year and annual averages over the next five years beginning in 2007.<sup>4</sup>

Despite the differences, one can see that both the 1-year and longer-run expectations are similar for the periods they overlap. In particular, the CST and NY Fed expectations are close to overlapping, with a more prominent drop early in the pandemic and spike late in the pandemic for the higher-frequency NY Fed data. The Michigan data, by contrast, generally gives lower average expectations but similar time paths. Nonetheless, the key patterns that Shiller and Thompson highlight in this paper, namely the fact that longer run and shorter run expectations largely overlap since 2012 and do not appear out of line in the pandemic, are both visible in the NY Fed and Michigan surveys. The fact that in the bust short-term expectations fall by more than long-term expectations is visible in both the CST and Michigan data.

What makes the CST data unique is that it is the only survey that covers the entirety of the 2000s boom and bust. This cycle is the largest and most consequential on record, so having data that shows just how out of line expectations – and in particular longer-run expectations – were in the boom is crucial to being able to use expectations data

---

<sup>3</sup>NY Fed asks “think about home prices nationwide” and say by how much “**Over then next 12 months**, by about what percent do you expect the average home price to increase/decrease?” and “**over the 12-month period between [24 months from survey date] and [36 months from survey date]**, by what percent do you expect the average home price to increase/decrease?” (Their emphasis.) The survey is weighted to be representative given response rates. The survey asks about a broad range of expectations, with the house price questions coming immediately after questions about income and credit availability in the middle of the survey.

<sup>4</sup>The Michigan Survey asks “By about what percent do you expect prices of homes like yours in your community to go (up/down), on average over the next 12 months?” and “By about what percent per year do you expect prices of homes like yours in your community to go (up/down), on average, over the next 5 years or so?” The survey is weighted to be representative given response rates. The survey asks about a broad range of expectations, with the house price questions coming after questions about inflation and specifically gas prices towards the end of the survey.

to diagnose a bubble in real time going forward. Indeed, the authors' observation that long-run expectations do not seem out of line in the pandemic boom is only revealing in comparison to their findings on the 2000s. The NY Fed survey may capture the later period, but since we do not know what this time series would look like in a significant boom and bust, it is hard to know how to interpret the COVID-era data. The same goes to a lesser extent for the Michigan data, which does not cover the 2000s boom.

That being said, given that data are released in near real time, the NY Fed and Michigan surveys are the early warning system for policy makers and economic forecasters. To maximize the impact of their research and big ideas about how survey expectations of house prices can be used to assess the direction of the housing market, I hope that Shiller and Thompson can work with the NY Fed and Michigan to compare survey designs and questions. By asking each others' questions with each others' phrasings for several years going forward – and possibly asking lab participants to answer multiple different survey questions and phrasings in multiple different scenarios – one can get to the bottom of whether these surveys behave differently due to the groups surveyed, the phrasing and sequencing of the questions, or other factors. This will help us ascertain how the “real time” surveys might look in a 2000s-like housing cycle and help policy makers assess the trajectory of the housing market with this sort of survey data. In other words, by treating these other surveys as complements not competitors, I think Shiller and Thompson can dramatically increase the influence and use of the type of survey data they pioneered.

### **1.3 A User's Guide to the CST Data**

There are two important things that users of the CST data should know.

First, the CST data suffers particularly in the boom (2003-2005) period from extremely high reported 10-year expectations. While the authors argue that 10% expected price



inflation over the next 10 years is not out of line with what actually happened over the prior 10 years, any model that attempts to match the level of expected 10-year appreciation they find will dramatically over-predict the size of the boom. Indeed, all papers that use the CST data as a calibration target that I know of find a way to artfully dodge the 10 year house price expectations in 2004 and 2005 for this very reason – either by using the 1-year expectations, by using an average of many years of the 10-year expectations, or by starting to use the CST 10-year expectations starting in 2006.

My concern, and the concern of David Laibson who focused his discussion of the original Case, Shiller, and Thompson (2012) paper on this issue, is that some of this is due to respondents misunderstanding the question. Prior to 2012, the survey did not underline and bold that the 10-year expectation was supposed to be “on average” and a growth rate for “each year,” and Laibson argues that some households misread this question in particular by conflating the average and total return. Given this, Laibson writes that the 10-year expectations “cannot be interpreted literally.” The authors clearly took this seriously as they explored this in the 2013 survey and found that 22 percent misunderstood the question. I share some of Laibson’s hesitance, but I am pleased to report that in this 10-year retrospective paper the authors do a better job adjusting for the sorts of survey confusion that concerned Laibson prior to 2012. Rather than reporting raw 10-percent trimmed means (dropping the highest and lowest 5 percent of responses and then calculating a mean), the authors now replace cases where the respondent gave a 10-year annual average expectation more than 10 times their 1-year expectation with the 1-year expected values and then calculate a 10-percent trimmed mean. This brings down some of the more extreme expectations – for instance, in 2004 and 2005 in Orange County, the average expectation for annual appreciation over the next 10 years is 13.3% and 10.4%, respectively, rather than 17.4% and 15.2%. While these results should still be taken with

a grain of salt and used carefully, the new figures are preferable.

Second, given the improvements in calculating the 10-year expectations in the new version, I urge researchers to use the updated expectations data from this 2022 version rather than the 2012 paper. That being said, I think there are still potential improvements. For instance, it is not clear that replacing the 10 year expectation with the 1 year expectation when the 10-year expectation is implausible is desirable. For this reason, I hope that Shiller and Thompson are able to release anonymized microdata so that researchers can implement their own trimming procedures as appropriate (not to mention analyze things like disagreement that one can only consider with microdata).

## **2 The 2012-2020 Rebound**

Part of the authors' analysis focuses on the "second" or "current" boom from 2012-2020. They compare this boom – in which short-run and long-run price expectations have remained stable and in line with mortgage interest rates – to the 2000s boom when short- and particularly long-run expectations skyrocketed. One interpretation of their findings is that expectations were more rational in the second boom.

I want to present an alternate view, which I develop with Gabriel Chodorow-Reich and Tim McQuade in a recent paper (Chodorow-Reich, Guren, and McQuade, 2022) that the 2012-2020 boom is not a second boom but instead the rebound phase of a single, 20-year boom-bust-rebound cycle. We begin with the observation that in the cross-section, areas with the largest booms (1997-2006) and busts (2006-2012) also had the largest rebounds (2012-2019). Indeed, the bust and the rebound are as highly correlated in the cross section as the bust and boom are. Furthermore, the boom is highly correlated with the overall 1997-2017 boom-bust-rebound price growth, with an R-squared of 0.62. The extremely high correlations across the three phases are indicative of a single, boom-bust-rebound

cycle rather than a boom-bust followed by a second unrelated boom. Furthermore, high correlation between the boom and longer-term price growth from 1997 to 2019 is suggestive of the boom being an overreaction to real improvements in fundamentals, an idea we explore systematically both in the data and using a model in the remainder of our paper.

We first pursue this interpretation of fundamental improvements driving long-run price growth empirically using a structural urban framework. We extract a city-level fundamental as a function of instruments for income, amenities, and supply and show that our estimated fundamental is correlated not only with long-run house price growth but also with the amplitude of the boom-bust-rebound cycle and the severity of the foreclosures crisis in the bust.

We then write down a model of a fundamentally-rooted house price cycle. In the model, a single improvement in the drift term of the “dividend” to living in a city in the late 1990s leads to a boom-bust-rebound pattern consistent with the data. Intuitively, the boom is generated by over-optimism about the fundamental improvement, the bust occurs as beliefs correct, bringing down prices and leading to price overshooting due to foreclosures. Finally, the rebound occurs as foreclosures recede and prices converge to a new, higher-growth balanced growth path. In the model, over-optimism occurs due to *diagnostic expectations*, which are non-rational expectations developed by Bordalo et al. (2019) that embed a tractable formalization of Kahneman and Tversky’s representativeness heuristic. The representativeness heuristic is that people tend to overweight the likelihood of a trait in a class when that trait has a higher likelihood in a class than in a reference population; for instance people tend to over-estimate the share of Irish with red hair because red hair is more prevalent among the Irish. In the context of asset prices, the reference population is the full history of observed dividends and the class is recently observed dividends, with inference over the dividend drift rate. As people observe higher dividends they over-weight

the probability of a very high dividend growth state, leading their long-run house price expectations to rise significantly. As people get more and more data, they realize their error and their beliefs converge to the rational belief gradually from above. Combined with the overshooting on the downside from foreclosures, this delivers a boom-bust-rebound, which we show is quantitatively consistent with the cross-section of boom-bust-rebounds across groups of cities. I see nothing in the authors' analysis of the 2012-2020 boom that is inconsistent with this story and consider it useful to think of the 2012-2020 boom this way in interpreting the authors' results.

Our paper is also a good example of how the literature has used the CST expectation survey data to discriminate between various models of non-rational beliefs and discipline macro models of the 2000s housing cycle. Indeed, one of the main reasons we use diagnostic expectations is because it is consistent with the CST observation that long-run expectations do not overshoot in the bust and instead converge smoothly from above; most other candidate models of expectations do not give this prediction.<sup>5</sup> The findings in this paper thus not only help reject rationality but also help narrow down the set of plausible non-rational models.

### **3 The Pandemic Housing Market**

Much of the analysis in the paper concerns the recent surge in house prices seen since the onset of the COVID-19 Pandemic. As mentioned above, the authors suggest that high long-run expectations can be used to diagnose a housing bubble in real time much like a yield curve inversion is used to predict a recession. They argue that 10-year expectations have not grown rapidly and so “we would not call the experience a bubble, at least not in

---

<sup>5</sup>Diagnostic expectations also creates an independence between the amplitude of the cycle and the length of the boom and bust which we observe in the cross-city data.

the classic sense.” That being said, they do hedge themselves a bit by saying the market “resembles a bubble in the sense that it is driven by a kind of excitement or fear of missing out.”

At the risk of making an inaccurate prediction, I agree with the authors: The pandemic price surge does not look like the bubble we experienced in the 2000s, so it is unlikely we will experience a correction like the one we experienced in the 2000s. In coming to this conclusion, the authors’ observation that long-run expectations have not increased significantly is an important data point, but it is not the only one. It is also important to note the lack of rapid credit expansion or speculation by short-term traders or house “flippers.” Various measures such as loan-to-value ratios, payment-to-income ratios, and credit scores of new mortgages suggest that we have not seen the type of credit expansion that occurred across the spectrum of borrower quality in the 2000s. For instance, one summary measure, the Credit Availability Index published by the Mortgage Bankers Association, saw a huge expansion and contraction in the 2000s but has barely budged in the pandemic. This is important because authors like Greenwood et al. (2022) have argued that rapid expansions in credit together with asset prices are predictive of bubbles and financial crises. Similarly, we do not seem to see a surge in the number of non-owner-occupant purchasers that we saw in the 2000s, which is a real-time proxy for the presence of speculators.<sup>6</sup> This sort of speculative activity is another hallmark of asset bubbles. Overall, I agree with the authors that policy makers and economic forecasters should use measures of long-term house price expectations in diagnosing a bubble, but I think they should be used in conjunction with other indicators rather than in a vacuum.

The other reason I am skeptical that the pandemic housing market is an expectation-fueled bubble is that there are good reasons to think that the pandemic has increased

---

<sup>6</sup>“Speculators” do not include institutional investors who are buying properties to hold and rent out. The market share of this type of investor has grown recently.

housing demand and constrained housing supply. On the demand side, the pandemic increased demand for housing space. Households began to work from home. City dwellers learned the value of additional space, particularly outdoors. Younger families decided to move to the suburbs sooner. At the same time, older households decided to age in place rather than downsizing or moving to senior living or a nursing home. All of these trends dramatically increased the demand for housing space. On the supply side, several factors have conspired to limit supply. For existing homes, the lack of downsizing by older households limited the supply of existing homes and led to record-low inventories of existing for-sale homes. Construction was also limited by material availability and supply chain disruptions as well as strong labor markets that drove up wages for construction workers and particularly for skilled construction workers who are in short supply. Immigration restrictions have also helped to drive up construction wages. The pandemic-induced expansion in demand and inelastic supply together led to a surge in house prices.

Of course, the fact that we will likely not experience a correction like the one we experienced last time does not mean there will not be a different type of correction. We tend to fight the last war but there are other factors to think about in the pandemic housing market that may not be as easily diagnosed using survey expectations. In particular, in assessing the risks of a different correction, housing economists should be paying attention to future supply and preferences for housing space.<sup>7</sup>

First, an eventual supply response could lead to a correction. The housing market currently appears to be supply constrained even in cities typically considered to be long-run elastic. Eventually supply should respond and construction should accelerate. But when and by how much? And perhaps more importantly, are market participants accounting

---

<sup>7</sup>In the very short run, inflation and interest rates are also a risk for the housing market. At the time of writing, mortgage rates have surged from near 3% in the fall to 5.3%. Rising mortgage rates will put a damper on housing demand and could cool the market off and potentially lead to a correction.

for a medium-run supply response when forming their expectations today? Or are market participants neglecting the future supply response (e.g., Greenwood and Hanson (2015))? My co-discussant Joseph Gyourko is an expert on housing supply and discusses these questions at length in his discussion so I will be brief; I share his concerns about how much we can learn from survey expectations when those expectations may not incorporate future supply responses fully.

Second, it is unclear whether the pandemic-induced change in preferences will reverse or be permanent. As I discussed above, the pandemic led to a significant increase in demand for housing space due to everything from working from home to people spending more time at home and valuing having more space to the desire to have outside space to older households wanting to age in place. One could imagine a world in which this reverses and prices decline significantly. One could just as easily imagine a world in which many of these changes in tastes are long-lived and housing demand stays strong.

In thinking about the pandemic, I am often drawn to the work of Malmendier and Nagel, who show that living through a traumatic economic event such as the Great Depression (Malmendier and Nagel, 2011) or Great Inflation (Malmendier and Nagel, 2016) has long-lasting impacts on economic agents' preferences, risk tolerances, and expectations. I suspect the COVID-19 Pandemic will have similar long-run effects on preferences, but it is hard to know exactly how and whether preferences relevant to housing markets like working from home and wanting more space will continue long-term. Economists are generally bad at forecasting changes in preferences. I thus think we need to be humble in our forecasts of housing prices and housing demand more generally.

I want to conclude where I began: The authors should be applauded for introducing and popularizing survey expectations about house prices as an important tool in analyzing housing markets. The more data on this subject that are available the better. I hope

that going forward the authors are able to open up their tremendous treasure trove of data – including microdata – to researchers so that we can better understand house price expectations and their role in shaping the dynamics of housing markets.



## References

- Armona, Luis, Andreas Fuster, and Basit Zafar (2019). “Home Price Expectations and Behaviour: Evidence From a Randomized Information Experiment”. *The Review of Economic Studies* 86 (4): 1371–1410.
- Bailey, Michael, Rachel Cao, Theresa Kuchler, Johannes Stroebel, and Arlene Wong (2018). “Social Connectedness: Measurement, Determinants, and Effects”. *Journal of Economic Perspectives* 32 (3): 259–80.
- Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer (2019). “Diagnostic Expectations and Stock Returns”. *The Journal of Finance* 74 (6): 2839–2874.
- Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo (2016). “Understanding Booms and Busts in Housing Markets”. *Journal of Political Economy* 124 (4): 1088–1147.
- Case, Karl E., Robert J. Shiller, and Anne K. Thompson (2012). “What Have They Been Thinking? Homebuyer Behavior in Hot and Cold Markets”. *Brookings Papers on Economic Activity*: 265–315.
- Chodorow-Reich, Gabriel, Adam M. Guren, and Timothy J. McQuade (2022). “The 2000s Housing Cycle With 2020 Hindsight A Neo-Kindlebergerian View”. National Bureau of Economic Research Working Paper No. 29140.
- Greenwood, Robin and Samuel G Hanson (2015). “Waves in Ship Prices and Investment”. *The Quarterly Journal of Economics* 130 (1): 55–109.
- Greenwood, Robin, Samuel G. Hanson, Andrei Shleifer, and Jakob Ahm Sorensen (2022). “Predictable Financial Crises”. *The Journal of Finance* 77 (2): 863–921.
- 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.
- Kuchler, Theresa, Monika Piazzesi, and Johannes Stroebel (2022). “Housing Market Expectations”. National Bureau of Economic Research Working Paper No. 29909.
- Malmendier, Ulrike and Stefan Nagel (2011). “Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?” *The quarterly journal of economics* 126 (1): 373–416.
- (2016). “Learning From Inflation Experiences”. *The Quarterly Journal of Economics* 131 (1): 53–87.