

Comments on “Shifting Confidence in Home Ownership: The Great Recession” by Anat Bracha and Julian C. Jamison

Stefan Nagel

Stanford University, NBER, and CEPR

1. Introduction

For the typical household, the purchase of a house is probably the most significant financial decision over the life cycle. In addition to purely financial considerations, homeownership seems to carry added significance for individuals as well as policymakers, as is evident in the notion of a “homeownership society” or the frequent association of homeownership with the “American Dream.” This paper by Anat Bracha and Julian Jamison asks the important question whether the recent housing crash has shattered individuals’ confidence in homeownership. Are we going to see a transition from the “homeownership society” to “Generation Rent”?

Bracha and Jamison have assembled a unique dataset that can shed some light on these questions. They added several questions to the Michigan Survey of Consumers in 2011. The data allows them to use between-ZIP code area variation in the severity of the housing crash to separate the effect of the housing crash from other effects of the Great Recession. Moreover, they use an individual’s age (based on the idea that younger people’s attitudes will be more sensitive to recent experience) and whether an individual (or someone they know well) was directly affected financially by the housing crash in order to further isolate the housing crash effect. Bracha and Jamison find that a more severe local housing crash leads to a loss of confidence in homeownership only among the young and those who had direct personal experience with the consequences of the housing crash. For older people, in contrast, the housing crash seems to have increased confidence in the desirability of homeownership. These results are a nice contribution to an emerging literature that shows how the path of economic events experienced by

individuals throughout their lives influences their attitudes, beliefs, and economic decisions.

My comments focus on four issues. First, I discuss the meaning of the notion of “confidence” in homeownership. The paper would benefit from greater conceptual clarity about the potential drivers behind the survey responses to the confidence question. The housing crash may have affected preferences in various ways, and it may have shifted individuals’ beliefs about the future path of events in the macroeconomy and the housing market.

Second, I argue that the authors could further refine the measurement of individuals’ experience of the housing crash. At present, they use a peak-to-trough house price variable to summarize individuals’ housing-crash experiences. However, it is not clear that this is the best variable to examine. In particular, I argue that how one should measure individuals’ experience depends on the specific hypothesis about the drivers of “confidence” that one is interested in testing.

Third, I comment on the regression specification. The current specification could potentially understate the effect of the housing crash on homeownership confidence.

Finally, I discuss some limitations of using cross-sectional variation in local house prices to estimate experience effects.

2. What is “Confidence in Homeownership”?

Bracha and Jamison elicit respondents’ confidence about homeownership by asking them “whether, financially speaking, it is better to buy a home or to rent a home?” “Confidence” is a somewhat fuzzy concept, though. One can think of several drivers that may influence individuals’ response to this question. It would be difficult to disentangle the effects of these potential drivers empirically, but the paper would nevertheless benefit from more discussion of what is underlying the notion of “confidence,” not least because this is relevant for how one should measure past house price experiences.

A first possibility is that the experience of a housing crash influences preferences, in particular risk preferences. For example, an individual may not be able to fully imagine the utility consequences of a severe housing downturn. Direct personal experience with a housing crash, or knowledge about its effects obtained from relatives and friends, may

make the potential utility loss in such bad states of the world more salient, leading to an increase in risk aversion. An increase in risk aversion would imply that not only decisions about homeownership should be affected, but also other financial decisions that depend on particular individual's degree of risk aversion. The authors utilize this latter feature to generate a control variable for a risk preference effect: they elicit individuals' risk aversion from a hypothetical investment gamble. If the housing crash affected risk aversion, the effect should also show up in this risk aversion measure. Bracha and Jamison find that it does not. One could perhaps quibble about potential problems with the risk aversion measure (it is noisy; it is based on a gamble that is a relatively small stake for wealthier people, and so on), but it seems clear that the housing crash does not have a strong effect on risk aversion.

A second set of factors that could drive confidence about homeownership has to do with beliefs, most likely beliefs about the path of future house prices. There are several ways in which individuals' subjective distribution of future house price changes could be affected by the experienced path of past house prices. These stories can be broadly separated into two groups: effects on the subjectively expected house price appreciation, and effects on the perceived riskiness of an investment in housing. Among the risk stories, one could further distinguish risk perceptions based on the second moment and perceived crash risk, meaning perceptions about the tail of the distribution.

As I will argue, thinking through these different possible channels is important, because it matters for how one should measure the experience of past house price changes.

3. Measuring the Relevant Experience

The authors measure experience as the fall in local house prices from peak to trough. Thus, this is a measure of the size of the crash. Thinking through the possible beliefs channels discussed above, it becomes clear that the size of the crash may not always represent the relevant experience.

The size of the crash seems most relevant if one takes the view that confidence about homeownership is driven by perceived crash risk. It would seem natural that individuals' perception of the likelihood and size of future housing price crashes is

affected most strongly by the experience of past house price crashes (rather than other aspects of the experienced distribution of past house price changes). Experiencing a severe house price crash may have shattered some individuals' beliefs that an investment in housing is a "safe" and "crash-proof" investment.

However, the size of the crash may not represent the relevant experience if confidence about homeownership is driven by individuals' beliefs about future mean price appreciation rather than perceived crash risk. One simple model of subjective belief formation would be one in which individuals' beliefs about the mean of future house price changes are a function of the experienced mean of past house price changes. For example, an individual's experience might be represented well by a weighted average of the house price changes experienced over the course of his or her life, as in Malmendier and Nagel (2011a) in a different context. Alternatively, given that there seems to be quite a high degree of momentum and mean reversion in house prices, one could perhaps get a better approximation of individuals' belief formation if one views them as estimating a simple ARMA model on the house price data in their experience set, as in Malmendier and Nagel (2011b), in the spirit of the literature on boundedly rational learning—but with the twist that an individual's attention is limited to experienced data as opposed to all available historical data.

To shed light on the role of subjective beliefs about expected house price appreciation in influencing confidence about homeownership, it would be interesting to experiment with such alternative summary measures of individuals' house price experiences. Clearly, with the cross-sectional dataset at hand, it would be difficult for Bracha and Jamison to follow the Malmendier-Nagel approach and estimate individuals' weighting of past data. But they could perhaps simply use the Malmendier-Nagel estimate of the weights (their estimated weighting function is roughly triangular, with weights declining linearly from the current period back to birth) to construct a summary measure of the house price appreciation experienced in the past.

Constructing such an alternative measure is important, in my view, because the authors' crash size measure is likely to be uncorrelated with the mean house price change experienced by individuals over the whole boom/bust cycle. The top panel in figure 1 presents some evidence in this regard from the Standard & Poor's/Case-Shiller 20-city

house price indices. The mean house price appreciation in index points from January 2000 to July 2011 (all indices start with a value of 100 in January 2000) is plotted on the horizontal axis, the size of the crash from July 2007 to January 2011 (all indices were at or close to their peak in July 2007) is shown on the vertical axis. Their correlation is -0.04 . The reason for this near-zero correlation is shown in the bottom panel. Here the price appreciation during the boom period from January 2000 to July 2007 is plotted on the horizontal axis. As the scatter plot shows, those cities that had the biggest crash also had the biggest price appreciation during the boom. In fact, the appreciation was of slightly bigger magnitude than the subsequent crash. Thus, it is possible that individuals living in areas with the biggest house price crashes actually still have a moderately positive view about future house price appreciation, because mentally they average the losses experienced during the crash period with the gains made during the preceding boom.

These patterns in the data may also help to explain what seems, at the surface, to be a rather puzzling result in the authors' empirical analysis: While younger people's confidence is adversely affected by the size of the crash, older people living in areas that witnessed the most severe housing crash seem to be *more* confident about homeownership. Under the view that confidence about homeownership is driven by expected house price appreciation rather than perceived crash risk, the potential explanation for this pattern would be that individuals form their beliefs about future house price appreciation from their experience of past (weighted) average house price appreciation, and that younger individuals put more weight on recent data than older individuals, consistent with the evidence in Malmendier and Nagel (2011a, b). The long-term house price appreciation experienced by older individuals might actually be highest in those areas that experienced the biggest drop during the recent crash. Alternatively, it would also be interesting to explore empirically whether individuals' could be viewed as estimating the degree of mean reversion from past data, and whether differences between young and old can be explained by differences in their views about mean reversion.

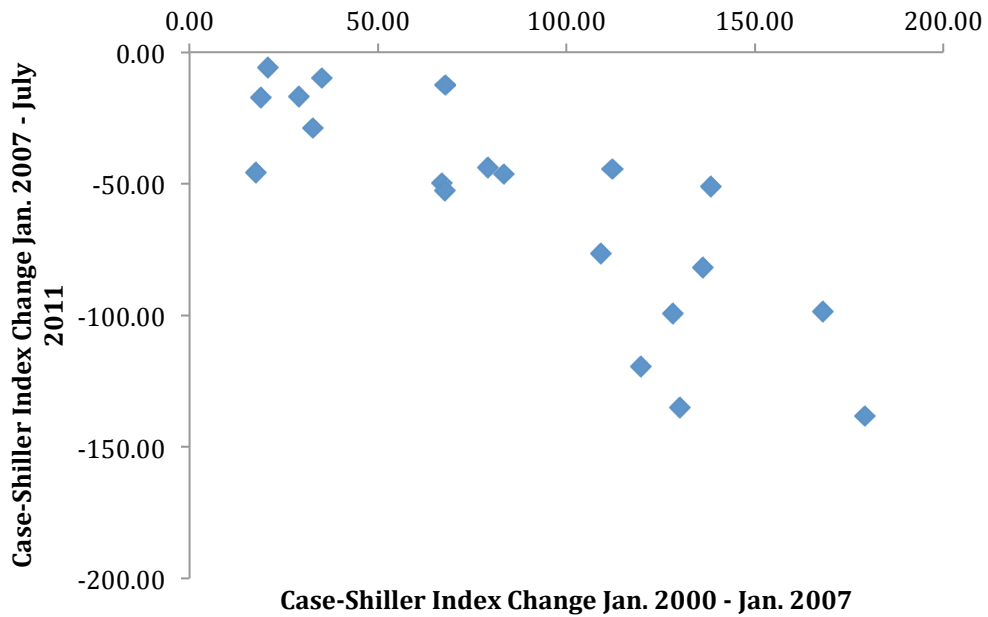
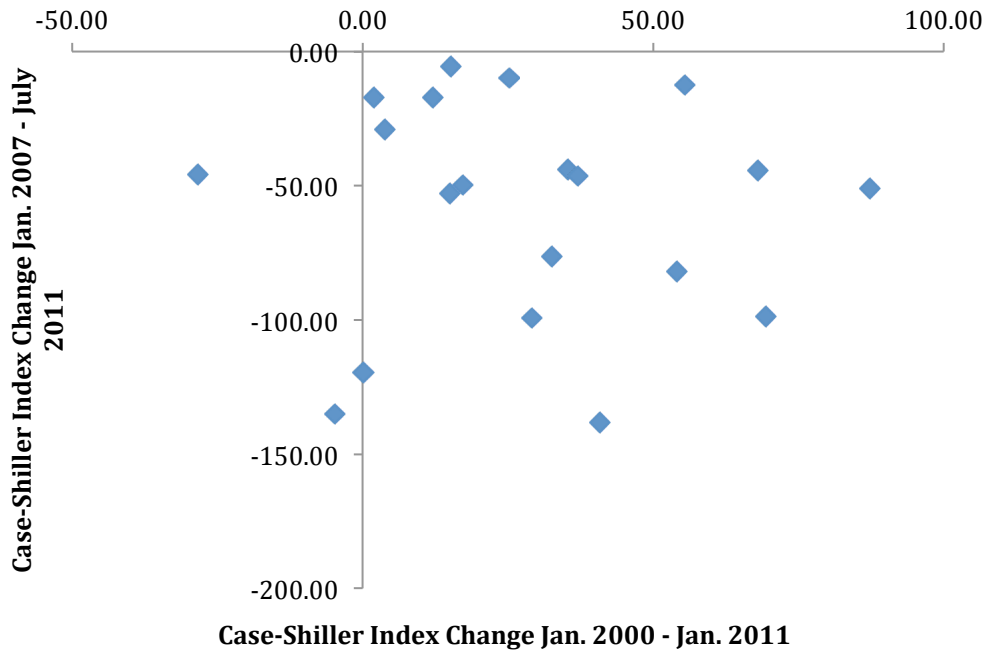


Figure 1: Correlation of Price Increase During Full Period (Top)/Boom Period (Bottom) and Size of the Crash for Standard & Poor's/Case-Shiller Home Price Indices (20 City Indices).

Source: Standard & Poor's and author's calculations.

On a different note, I also want to highlight a nice feature of the authors' dataset that allows them to investigate whether it makes a difference that an individual just lived through a crash episode or whether the individual has any direct personal experience with events occurring during such an episode. It seems a highly plausible hypothesis (and it is confirmed by the result in this paper) that experiencing an economic event has stronger effects if the event results in direct financial consequences for an individual (or persons close to the individual) rather than just acts as a news flow does not have direct personal relevance. In our work on the role of experience in influencing households' portfolio allocation decisions over the last 50 years (Malmendier and Nagel 2011a), we would have loved to make an attempt to separate personal experience (for example, the experience of losses in a person's own stock portfolio) from broader macroeconomic experiences that do not have direct financial consequences for an individual but we could not. This type of information is simply not available in a long-term dataset on household portfolio allocation. Therefore, I believe that the addition of a "personal experience" question to the Michigan Survey was a great step to making progress on this issue.

4. Regression Specification

Turning to the econometrics, I think that the authors should give some more thought to the regression specification. In particular, the choice of some of the control variables on the right-hand side (RHS) of the regression equation is not well motivated.

I am most worried about the inclusion of the *Rent* indicator as a control variable. This is one of the controls that are highly significant, but there is a danger that it might absorb the very effect that the authors are trying to uncover. The *Rent* variable indicates whether someone is renting in 2011, meaning after the experience of the housing crash. This creates a potential problem, because whether an individual is currently renting or not can be seen as an alternative proxy for confidence in homeownership. Putting *Rent* on the RHS is therefore like putting a noisy proxy for the dependent variable on the RHS. It is therefore not surprising that it turns out to be highly significant. To see the problem with this, consider the extreme case in which a variable (*Rent*) $z = y + u$ is a noisy measurement of y (confidence) with uncorrelated error u . Suppose the true model is $y = \beta x + e$ with $\beta > 0$, where x represents the explanatory variable of interest. Clearly, z is

correlated with e , and hence endogenous. Including z into the regression yields an inconsistent estimate of β . Thus, the inclusion of *Rent* probably leads to an understatement of the effects of the housing crash on homeownership confidence.

For the same reason, the inclusion of the risk aversion variable as a control could potentially cause problems. Under the hypothesis that the housing crash affects risk preferences (one of the possible channels through which the housing crash might affect “confidence”), the risk aversion variable would be endogenous in similar ways. My guess this is less of a concern for this application because, based on the authors’ results, there does not seem to be a detectable link between the housing crash and risk aversion. Nevertheless, it might be better to exclude the risk aversion variable. If one asks, in a first step, whether the housing crash affects confidence, without taking a stand on the particular channel that may be driving the confidence effect, no such control for risk aversion is needed. In a second step (as the authors do in table A.4 in the appendix), one can then ask, by regressing risk aversion on the housing crash size measure, to what extent risk aversion might be the channel through which the housing crash influences confidence.

5. Local House Price Variation as Source of Identification

The authors have assembled a survey dataset that allows the use of local house price variation as a source of identification. This is a great step forward. Nevertheless, it is important to keep in mind the limitations of this identification strategy.

My biggest concern is that confidence and house prices are (locally) jointly determined. The level of house prices that equilibrates the local housing market presumably depends, among other things, on people’s confidence in homeownership, while confidence in turn may depend, among other things, on house prices. The latter effect is the channel that the authors have in mind, but their identification strategy does not isolate this latter effect from the first one. In other words, there could be reverse causality. As a related example, a similar issue appears in Mian and Sufi (2011) in an investigation of the effect of house price growth on household borrowing, which leads Mian and Sufi to instrument for house price growth with local housing supply elasticities. Perhaps one could use a similar approach in this setting here. To the extent that those

areas with the least elastic supply not only saw the biggest rise in house prices, but also the biggest drop, this strategy could perhaps work.

The evidence that in areas with big housing crashes only individuals who are young and personally affected exhibit a lack of confidence in homeownership makes this reverse causality implausible as a full explanation for the results, but a concern remains that reverse causality could influence the magnitude of the estimated effects: An initial drop in house prices might dent the confidence of people with particularly malleable beliefs (meaning young people) and people that are directly affected. The reduced confidence in homeownership among these subgroups then leads to a further drop in the equilibrium price of houses in the local housing market. As a consequence, the correlation between the drop in confidence among the young who are personally affected and falling house prices overstates the causal effect of house prices on their confidence. In my view, the correlation is still interesting, but one has to be cautious in the interpretation. The estimates are not suitable for counterfactual experiments.

6. Concluding Thoughts

Understanding the formation of beliefs and attitudes is important for macroeconomics and many other areas in economics. It is becoming increasingly clear that to understand how beliefs and attitudes are shaped by historical data, one has to distinguish between historical data that may be on the record somewhere, and historical data that happened during individuals' lives, allowing them to absorb the data in "real time" from the news, in the market place, and/or from their particular social environment. Beliefs and attitudes of young people, who have the smallest set of experienced data, are particularly responsive to recent events. Bracha and Jamison provide new evidence that further supports this view. They also show that a further distinction is important: it matters whether an individual just happens to live in an area or during time period in which an event occurs or whether the individual is personally affected (or someone close to this individual is) by this event.

As I have argued, the authors could do a number of things to improve the robustness and the interpretability of the results. There is still scope for squeezing more information out of their dataset. It would be great if the authors or others could, in future

research, also investigate whether the housing crash affected actual decisions in ways that are consistent with the reported attitudes in the survey.

References

Malmendier, Ulrike, and Stefan Nagel. 2011a. “Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?” *Quarterly Journal of Economics* 126(1):373–416.

Malmendier, Ulrike, and Stefan Nagel. 2011b. “Learning from Inflation Experiences.” Graduate School of Business Working Paper. Stanford, CA: Stanford University. Available at <http://faculty-gsb.stanford.edu/nagel/documents/InflExp.pdf>.

Mian, Atif, and Amir Sufi. 2011. “House Prices, Home Equity-Based Borrowing, and the U.S. Household Leverage Crisis.” *American Economic Review* 101(5):2132–2156.