Macroeconomic Experiences and Expectations: A Perspective on the Great Recession*

Stefan Nagel†
Stanford University, NBER, and CEPR

April 2012

Abstract

Life-time experiences shape macroeconomic expectations. Young people, endowed a small experience set, are particularly prone to rely on very recent data when forming expectations. I illustrate these patterns with micro survey data on inflation, stock return, and house price appreciation expectations. In the case of asset return expectations, the tendency to extrapolate from recent data makes asset return expectations pro-cyclical. This pro-cyclicality supported the house price boom leading up to the Great Recession and it reinforced the subsequent bust. The optimistic expectations during the boom help understand the rise in household leverage, the buildup of excess stocks of durables and housing during those years, and hence the fragility that set the stage for the ensuing household balance sheet recession. I also discuss the outlook for households' asset return expectations given their experiences with the Great Recession and the subsequent slump until the end of 2011.

*Paper prepared for Academic Consultants meeting of the Board of Governors of the Federal Reserve System, May 14, 2012.

†Graduate School of Business, Stanford University, 655 Knight Way, Stanford, CA 94305

Consumer decisions to purchase a house, to take on debt, or to invest in the stock market reflect expectations about future personal circumstances and about the macroeconomy many years into the future. As in other realms of human decision-making, past experiences may exert a profound influence on these economic expectations. Macroeconomic events may shape the expectations of generations that live through these events. Discussing the persistent increase in the demand for money following the onset of the Great Depression in 1929, Friedman and Schwartz (1963) suggest that

> "The contraction after 1929 shattered beliefs in a 'new era' [...]. The contraction instilled instead an exaggerated fear of continued economic instability, of the danger of stagnation, of the possibility of recurrent unemployment." (p.673)

In line with Friedman and Schwartz' hypothesis, experience of macroeconomic events might have an excessive influence on individuals' expectations—excessive relative to the fiction of an agent armed with rational expectations and knowledge of the parameters of the processes driving the economy.

In this paper, I present empirical evidence from microdata on how macroeconomic experiences shape expectations, and I discuss channels through which these experience effects may have contributed to the dynamics of household spending before, during, and after the Great Recession 2007-09.

I focus my discussion on subjective expectations of macroeconomic variables. This is not to say that preferences (for risk taking, for consumption smoothing) might not also be influenced by experiences of macroeconomic events. However, microdata over long time spans is required to separate the effects of life-time experiences from other factors. The availability of clean measures of preferences is more limited in this regard than the availability of expectations data.

The set of macroeconomic variables that could play an important role in household spending decisions is vast. I focus on a few for which we either have a long history of survey data

(inflation) or which may be particularly relevant in the context of the Great Recession (stock market returns and house prices).

My analysis focuses exclusively on macroeconomic expectations of households. It is reasonable to assume that in some parts of the economy, expectations of trained professionals might matter more than expectations of households. The expectations of these professionals may also be driven less by personal experiences and more by extensive analysis and modeling of historical data.[1] On the other hand, there are many examples where the household expectations studied in this paper—about inflation, stock returns, house prices—are likely to be relevant for aggregate economic outcomes. For example, even though a high share of household financial wealth is now in the hands of professional investment managers, the decision whether to allocate funds to the stock market or the bond market is still controlled, to a large extent, by households when they allocate money to bond funds, stock funds, or funds with deterministic time pattern of allocation such as the recently popular target-date funds. Active professional investors could lean against households' allocation decisions, but this activity is risky and thus limited. Similarly, the residential housing market is driven by households, and the absence of short-selling in the housing market leaves little room for professionals to influence prices.

1 How do macroeconomic experiences influence expectations?

A framework to study the effects of experiences on expectations should be flexible enough to accommodate two features. First, it should be flexible enough to allow the experiences accumulated during an individuals life-time to carry different weight than other historical data. This distinction between life-time experiences and other historical data is different from standard implementations of learning in macro models. Second, within an individuals life-time, there should be a possibility for the effect of experiences to decay over time, as

[1] It is not obvious, though, that this is always true. For example, Greenwood and Nagel (2009) find that during the internet boom of the late 1990s', young mutual fund managers decisions were influenced by their own recent positive experiences with technology stock investments.

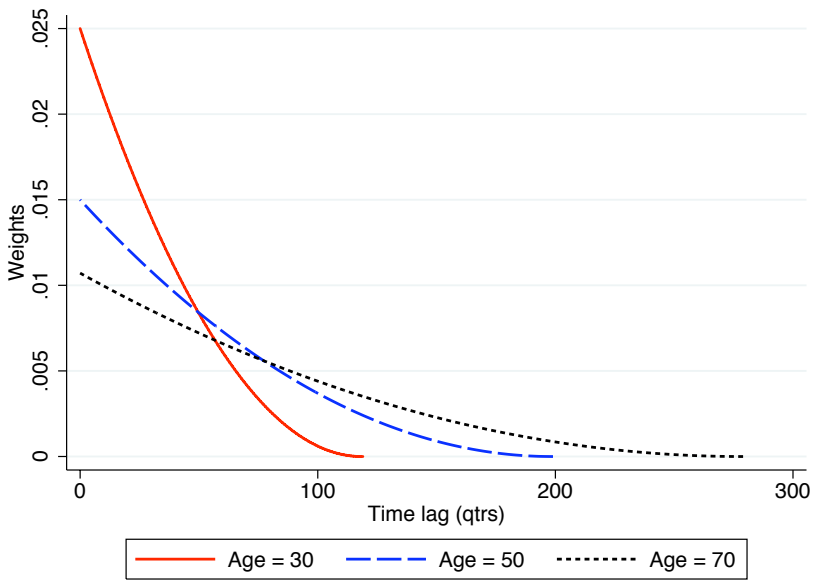

Figure 1: Weighting scheme for macroeconomic life-time experiences. Weights for quarterly observations shown, as examples, for three different ages. The weights applied to the most recent observations are on the left-hand side of the figure.

memory may be lost, or the individual becomes convinced that the economy has been subject to structural change that renders early experiences less relevant. To what extent these two features are relevant is, of course, an empirical question.

Malmendier and Nagel (2012) set up a framework that incorporates these two features. Individuals believe that a macroeconomic variable x_t follows a perceived law of motion, e.g., a first-order autoregressive process,

$$x_{t+1} = \alpha + \phi x_t + \eta_{t+1}. \tag{1}$$

To form expectations at a time point t_0 about future values of x_t for $t > t_0$, individuals need estimates of α and ϕ. They use the data realized during their life times until t_0 to estimate these parameters. Thus, younger individuals use shorter data sets and therefore come up with different expectations than older individuals.

The weighting of experienced life-time data is done according to the weighting scheme

displayed in Figure 1. The plot shows the weights (for quarterly data) that individuals of different age apply to observations in the past. For example, a 50-year old individual assigns a weight of 0.015 to the last quarterly observation, and observations more distant in the past receive lower weights, with close-to- zero weights for observations during youth. The weights shown in the figure are estimates obtained by Malmendier and Nagel (2012) from fitting the above model of expectations formation to inflation expectations microdata from the Michigan Survey of Consumers.[2] Our weighting scheme was flexible enough to permit, for example, equal weights, or weights that put much less weights on data in the distant past, but the data preferred the weights shown in Figure 1. Approximately similar weights, applied to stock and bond return histories, also do a good job in explaining household portfolio choices and financial risk-taking in the Survey of Consumer Finances (Malmendier and Nagel (2011)).

This learning-from-experience framework shares with standard learning approaches the prediction that individuals' expectations depend on the path of macroeconomic history. As new data is realized, individuals update their views about the process driving the economy (parameters α and β in the example in equation (1)). Different from standard learning models, the learning-from-experience theory makes the prediction that expectations should be heterogeneous by age: Younger individuals rely on different data sets than older individuals in forming their expectations. This heterogeneity of expectations by age represents a useful "footprint" of learning-from-experience that helps to empirically identify the learning-from-experience mechanism in expectations data. I illustrate this below with several examples.

1.1 Inflation

Panel (a) in Figure 2 plots one-year inflation expectations from the Michigan Survey of Consumers, with the sample broken up into two age groups (age < 50 and age ≥ 50). The plot shows that there was substantial generational disagreement about future inflation rates

[2]The estimation in Malmendier and Nagel (2012) is done with a recursive formulation of a learning rule in which the parameters are updated each quarter based on the newly observed inflation rate. Setting the parameter θ to 3.0 in this learning rule (which is, approximately, the point estimate in various specifications) produces an implied weighting of past data as shown in Figure 1

following the high-inflation years of the 1970s and early 1980s. Younger people expected between 2 and 3 percentage points higher inflation rates. This gap in expectations shrinks gradually until it completely disappears in the mid-1990s. In recent years, older individuals had somewhat higher inflation expectations.

Panel (b) shows that this disagreement between young and old is matched by the difference in inflation histories experienced by these age groups. To construct this plot in Panel (b), I assume that each cohort (defined by year of birth) estimates a first-order autoregressive model of inflation rates as in equation (1) with data from birth to the quarter prior to the quarter in which the forecast is made, and weighted using the weighting scheme shown in Figure 1. Through this autoregressive model, individuals use their experienced inflation history to form an opinion about the mean and the persistence of inflation rates. The resulting experience-based inflation forecast matches the actual inflation expectations in Panel (a) well. In the early 1980s, the inflation experience of young people is dominated by the persistently high inflation rates of the 1970s, while the experience set of older individuals also includes the low inflation years in the 1950s and 1960s. As a result, younger individuals in the early 1980s perceived a higher mean inflation rate, as well as a higher degree of inflation persistence. By the mid-1990s, this difference between young and old had faded away after many years of relatively low inflation rates, consistent with a similar pattern in actual expectations in Panel (a). In the last decade, the experience-based forecast also matches the fact in Panel (a) that older individuals had somewhat higher inflation expectations than younger individuals, but the magnitude of the difference in actual expectations is smaller than predicted in Panel (b).

Figure 2 illustrates how the learning-from-experience model is useful to understand expectations formation. Formal statistical inference in Malmendier and Nagel (2012) using data that extends back to the 1950s confirms these visual impressions. Furthermore, Malmendier and Nagel (2011) also show that variation in experienced real returns on long-term government bonds (which are driven largely by variation in experienced inflation histories) helps explain household portfolio allocation to long-term bonds in the Survey of Consumer

(a) One-year inflation expectations

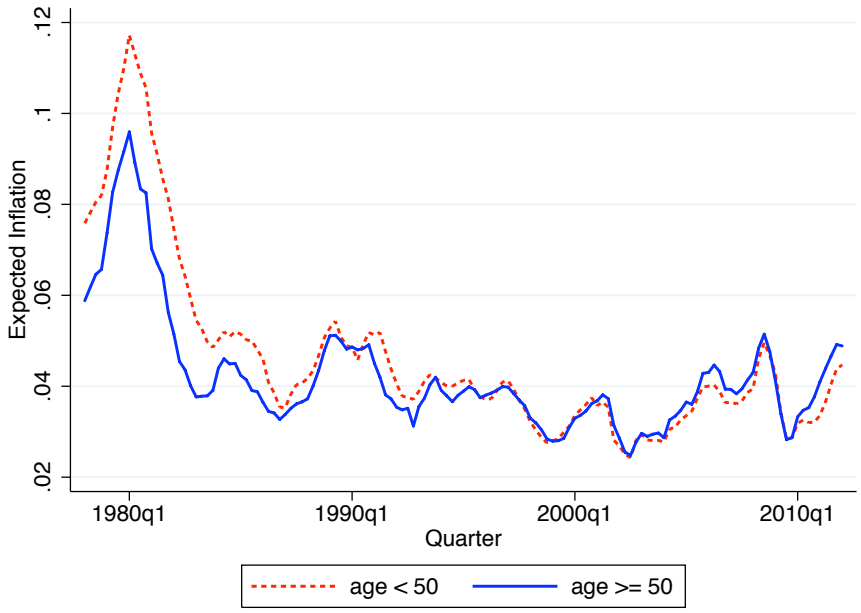

(b) Experience-based inflation forecast

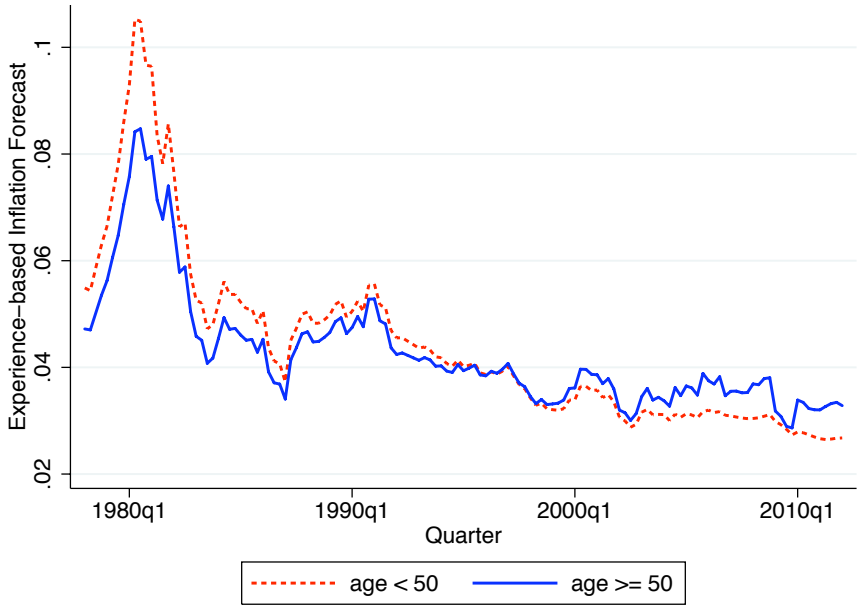

Figure 2: Inflation expectations and experienced inflation (four-quarter moving averages). Inflation expectations in Panel (a) are from the Michigan Survey of Consumers for individuals of age < 50 or age ≥ 50. The experienced-based inflation forecast in Panel (b) is the fitted value from an AR(1) model estimated from inflation data realized over each birth-year cohort's life time, weighted according to the weighting scheme illustrated in Figure 1.

Finances. Cohorts that experienced low real returns on long-term government bonds (which tends to happen in times when experienced inflation rates were high) allocate a smaller portion of their financial investments to long-term government bonds. Furthermore, Piazzesi and Schneider (2012) show that the disagreement between cohorts about the path of future inflation, which implies disagreement about real interest rates, was associated with greater credit volume, as younger households found it advantageous, given their beliefs, to borrow from older households to buy housing assets. Thus, the experience-effects appear not only in elicited expectations, but also in economic choices of households.

1.2 Stock Market Returns

The public's views about future returns from stock market investments are likely to be one of the more important macroeconomic expectations. Unfortunately, data on stock market return expectations is much more patchy than the data on inflation expectations. For this reason I start first with some data on choices rather than beliefs.

Panel (a) in Figure 3 presents the detrended stock market participation rate since the 1960s from the Survey of Consumer Finances from 1960 to 2007. The plot shows the proportion of each age group (age < 50 or age ≥ 50) that owns stocks, either directly or indirectly through mutual funds. Raw stock market participation rates have grown steadily over time due to the increasing prevalence of defined-contribution retirement plans. Without the detrending, both series would show an almost linear upward trend. To focus on variation around this long-run trend, I detrended the series with a linear time trend. The resulting detrended series broadly move in line with stock prices. Stock market participation rates are high following stock market booms in the 1960s and 1990s, and they are low following periods of poor stock market performance such as the 1970s/early 1980s, and since 2000 (note that the series end before the financial crisis). The most notable fact in this figure is that the stock market participation rate of younger individuals varies much more with stock market performance than the participation rate of older individuals.

(a) Linearly detrended stock market participation rate

(b) Experienced stock market excess returns

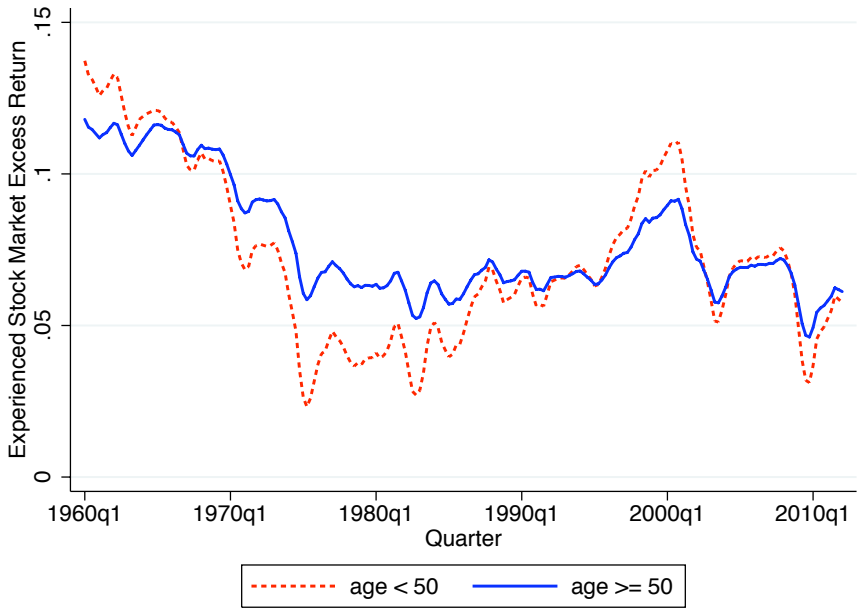

Figure 3: Stock market participation rate and experienced stock market excess returns. The stock market participation rate in Panel (a) is the proportion of households in the Survey of Consumer Finances with household head age < 50 or ≥ 50 that own stocks or stock mutual funds. Experienced stock market excess returns in Panel (b) are a weighted average of the difference between the return on the S&P500 index and treasury bill yields over a household head's life time, weighted according to the weighting scheme illustrated in Figure 1, and plotted as a four-quarter moving average.

Panel (b) shows that this pattern in the data can be explained by differences between young and old in their experienced stock market return histories. The experienced returns are calculated in the same fashion as the experienced-based inflation forecasts in Figure 2, but with the first-order autoregressive model in equation (1) replaced with a constant mean model (i.e., $\phi = 0$). The reason for this different approach to modeling the perceived law of motion is that it seems widely known that inflation often has substantial short-run persistence. Such short-run persistence does not exist in stock returns, and so I model individuals as forming a view about the mean rate of return from stock market investments as a simple weighted average of the stock market return experienced during their life-time (with weights as in Figure 1). More precisely, I calculate experienced stock market returns in excess of (one-month) Treasury bill yields as the individuals views about the desirability of stock market investments are likely to be influenced by the relative return of stocks compared with the return of safe investments. Panel (b) shows that the age-related differences in experienced excess returns broadly match the age-heterogeneity in stock market participation rates in Panel (a). Periods like the 1970s and early 1980s when young people had relatively low stock market participation rates compared with older people are also periods when the young had relatively poor experienced returns. The opposite happens following stock market booms, such as the late 1990s. Formal statistical inference of these relationships between experienced returns and stock market participation rates, as well as other measures of financial risk taking is reported in Malmendier and Nagel (2011).

A plausible explanation for these effects is that stock market booms make individuals more optimistic about the returns that can be earned from stock market investments, which increases the willingness to invest in stocks. Younger individuals form their views about stock returns based on a smaller data set of experienced returns. As a result, they are more likely to be excessively optimistic about future stock returns after a few years of rising stock prices. Their return expectations are more pro-cyclical than the expectations of older individuals.

While this explanation is plausible, it is also possible that experienced returns affect indi-

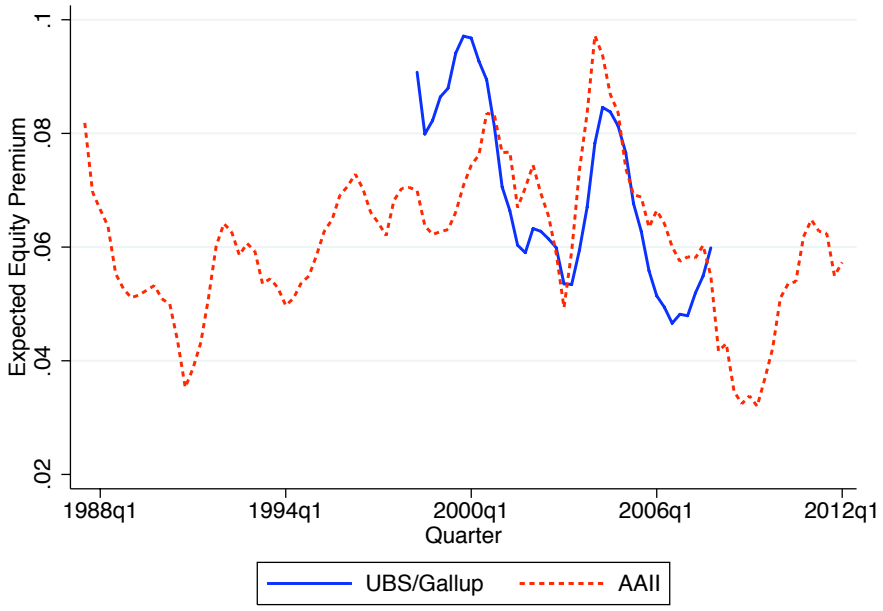

Figure 4: One-year expectation of stock market excess returns (four-quarter moving averages). The UBS/Gallup Survey data is based on responses to a question about the expected one-year return on an individuals' own portfolio, and the plotted expected excess return series is obtained by subtracting the one-year Treasury yield at the end of the interview quarter. The AAII return expectations are extracted from categorial responses about "bearish", "neutral", or "bullish" stock market sentiment by assuming an underlying distribution of return expectations and fitting the mean and standard deviation of the extracted expectations to the UBS/Gallup survey in periods of overlapping coverage.

viduals risk preferences. For example, living through a severe stock market downturn could reduce individuals' risk tolerance. Based on a choice variable like stock market participation, it is impossible to tell whether the effect works through preferences or expectations. However, the limited data on stock market return expectations that is available seems to be consistent with an expectations channel (which does not rule out, of course, that a risk preference channel exists as well). Malmendier and Nagel (2011) find that experienced stock market returns explain age-related differences in stock market return expectations in the UBS/Gallup survey from 1998-2007.

The short series of expectations data available from the UBS/Gallup survey is shown in Figure 4. The solid line shows a four-quarter moving average of the mean one-year expected

stock return of all respondents in excess of the yield on U.S. Treasury bonds with one year of remaining maturity. At a level of almost 10%, the level of expected stock market excess returns was highest around the peak of the market in early 2000. It declined towards 5% in early 2007, with some rebound later in that year. These numbers from the UBS/Gallup Survey are also consistent with the return expectations obtained from a special supplement to the Michigan survey of consumers in Amromin and Sharpe (2009).

The dashed line shows excess return expectations extracted from a survey conducted by the American Association of Individual Investors (AAII).[3] The AAII data should be interpreted with caution, because the AAII survey is not based on a representative random sample. It represents the views of AAII members that are willing to fill out a survey form on the AAII website. Furthermore, the information collected is categorical. Survey participants declare whether they are "bullish","bearish", or "neutral" about the stock market, and the available data series contain the percentage of respondents that fall in each category. I extract an implied excess return expectation from this categorical information by assuming an underlying distribution of return expectations whose parameters shift over time. In this setting, the mean expected return should be a function is the fraction of "bulls" minus the fraction of "bears" divided by the fraction of "neutral" respondents. I calibrate this series to match the mean and standard deviation of the UBS/Gallup survey expectations in periods of overlapping coverage. Details are available in Appendix A. Figure 4 shows that stock return expectations of individual investors are strongly pro-cyclical. They rise in stock market booms (from early 1990s to 2000) and they fall in stock market downturns (around 1990, 2000/01 and 2008/09).

1.3 House Prices

The residential real estate market arguably played a bigger role than the stock market in the Great Recession of 2007-09. Unfortunately, to the best of my knowledge, no long-term time

[3]The data is available at `http://www.aaii.com/sentimentsurvey`.

series of household survey data on house price appreciation expectations exist. The Michigan Survey of Consumers collects information on home buying attitudes, but the interpretation of the responses to these questions involves some difficulties. Individuals are asked whether they think it is a "good time" or a "bad time" to buy a house. Unfortunately, a variety of factors could influence individuals' views about whether it is a good or bad time to buy a house. For example, high interest rates could lead respondents to conclude that it is a bad time to buy a house, even though they might have a positive view about expected house price appreciation. Moreover, pro-cyclicality of price appreciation expectations might not be visible in the responses to these questions. Following a period of rising housing prices, individuals might conclude that it is a bad time to buy a house—even if they expect prices to appreciate further—because they view houses as no longer affordable, given their labor income and credit constraints. Such individuals who are optimistic about price appreciation but respond in the survey that it is a bad time to buy a house will not have their views about future house prices recorded in the survey, because following a "bad time" response, the survey only asks for reasons why it is a *bad* time to buy a house.

Despite these problems, the information in the survey can be used to construct an approximate measure of house price expectations along the lines of the stock return expectations measure from the AAII data above. I calculate the fraction of individuals who respond that it is a good time to buy a house *and* mention price appreciation as a reason for their positive view minus the fraction that think it is a bad time to buy a house *and* mention price depreciation as a reason, divided by the fraction that does not fall into these two categories. As with the AAII data above, I use a series of explicit percentage expectations over a shorter subsample to calibrate the extracted house price appreciation expectation series. In this case here, I use an annual series of percentage expectations from the Case/Shiller Survey of Homebuyers reported in Case and Shiller (2010) to calibrate the Michigan Survey of Consumers series by matching the mean and standard deviation in periods of overlapping coverage. Details are provided in Appendix B. The Case/Shiller data available to me does not allow me

to distinguish age subgroups, so this analysis is done without separating younger and older individuals.

The resulting series are shown in Figure 5, Panel (a). Panel (a) shows that optimism about house price gains rose to an all-time high in the housing boom leading up to 2006. The start of the housing crash in 2006 lead to a drastic revision of expectations. Expectations quickly deteriorated to an unprecedented level of pessimism.

Panel (b) shows that these patterns are broadly in line with individuals' house price appreciation experiences, calculated from an updated series of house prices in Shiller (2005). Earlier house price booms leading up to 1980 and 1990, and, most pronounced, the long boom leading up to 2006, lead to high levels of experienced house price appreciation, and the peaks in experienced house price appreciation roughly coincide with peaks in optimism in Panel (a). The crash in expectations in Panel (a) that starts in 2006 also parallels the sharp drop in experienced house price appreciation in Panel (b). Overall, the pro-cyclicality of house price appreciation expectations is apparent.

The data underlying the series shown in Panel (a) are not suitable for a relative comparison between age groups. But additional evidence on age-related heterogeneity in beliefs following the housing crash is available from Bracha and Jamison (2012). Using ZIP code level variation in house prices, they find that loss of confidence in owning a house (vs. renting) in the areas hit hardest by the housing crash is concentrated among young people and those who have been directly affected by the crash in real estate prices. This could be consistent with the fact that house price appreciation experiences deteriorated the most for younger individuals, as shown in Panel (b).

(a) Expected house price appreciation

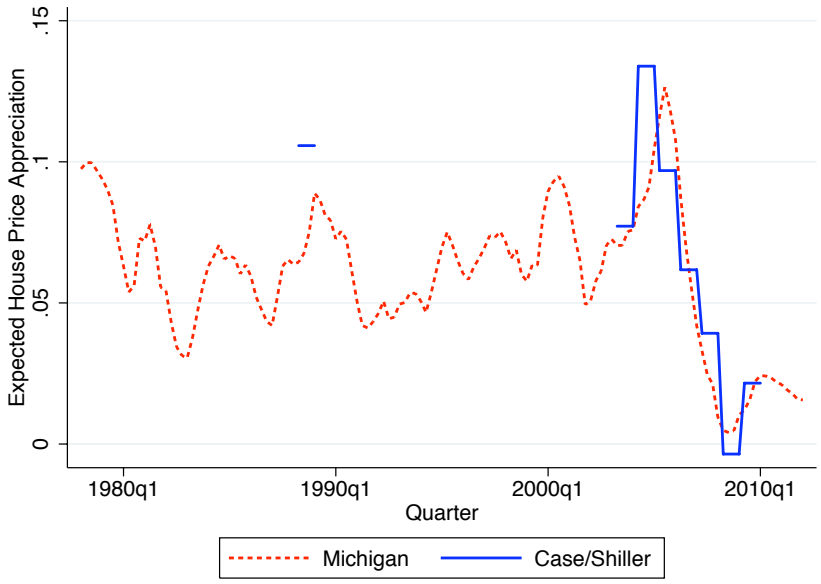

(b) Experienced house price appreciation

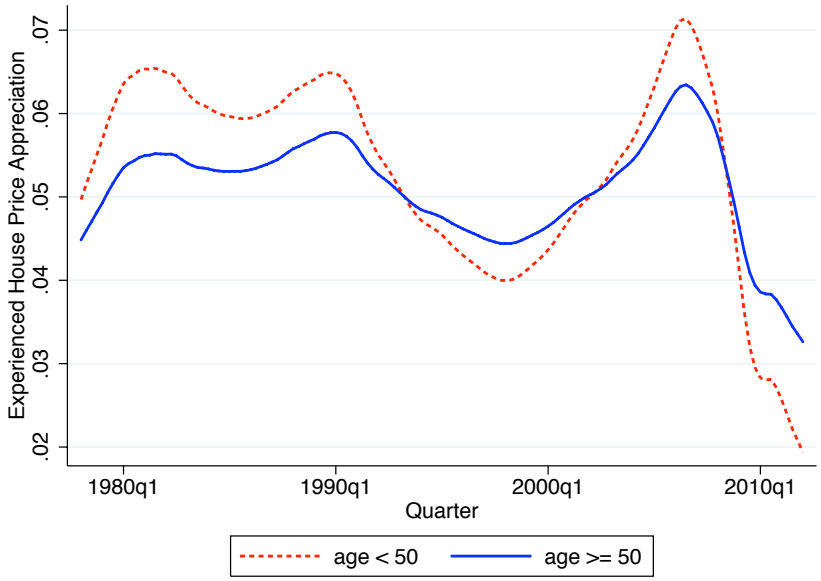

Figure 5: Expected house price appreciation and experienced house price appreciation (four-quarter moving averages). The house price appreciation expectations from the Michigan Survey of Consumers in Panel (a) are extracted from categorial responses by assuming an underlying distribution of return expectations and fitting the mean and standard deviation of the extracted expectations to the (one-year horizon) price expectations from the Case/Shiller Homebuyer Survey in periods of overlapping coverage (1988 and 2003-2009). Experienced house price appreciation in Panel (b) is calculated from the (nominal and updated) house price index in Shiller (2005), weighted according to the weighting scheme illustrated in Figure 1.

2 The role of macroeconomic experiences in the Great Recession

Among the expectations analyzed in the previous sections, the ones about asset returns are most likely to be relevant for understanding the Great Recession. Inflation experiences and expectations may have played an important role in the early 1980s, but they did not show much movement around the Great Recession. For this reason, I focus now on the pro-cyclicality of asset return expectations. I offer some thoughts on how this pro-cyclicality—induced by individuals' reliance on their own macroeconomic experiences—may help understand patterns in household spending before, during, and after the Great Recession.

Hall (2011) argues that a buildup of excess stocks of housing and consumer durables prior to the Great Recession is key to understanding the long slump that followed. This slump was characterized by a big and persistent drop in residential construction and household spending on durables and non-durables and services (Petev, Pistaferri, and Eksten (2011)). Could pro-cyclical asset return expectations have played a role in generating these spending patterns?

As a first step, it is useful to clarify the link between asset prices and pro-cyclical asset return expectations. Following the log-linearized present value model of Campbell and Shiller (1988), the log price-dividend ratio of an asset can be expressed as

$$\log \frac{P_t}{D_t} = E_t^S \sum_{i=1}^{\infty} \rho^i \left(\Delta d_{t+i} - r_{t+i}\right) = E_t \sum_{i=1}^{\infty} \rho^i \left(\Delta d_{t+i} - r_{t+i}\right)$$

where ρ is a constant close to one, Δd_{t+i} is the growth rate of future fundamentals (dividends in the case of stocks, rent or its owner-equivalent in the case of housing), r_{t+i} are future (log) returns. Suppose there is an asset-price boom, and extrapolation from experience makes a subset of individuals in the economy excessively optimistic about future Δd_{t+i}. Denote these optimists' expectations with $E_t^S[.]$. The remainder of the investor population has rational expectations, denoted $E_t[.]$. With both types of investors present in the market, the price of

the asset, and hence the $\frac{P_t}{D_t}$ ratio, will be somewhere inbetween the prices that would prevail if only optimists or only rational investors were present. Given this elevated $\frac{P_t}{D_t}$ ratio, and their non-optimistic $E_t[\Delta d_{t+i}]$, rational investors expect future returns to be low. Optimistic investors, on the other hand perceive this elevated $\frac{P_t}{D_t}$ ratio still to be low relative to their optimistic expectations $E_t^S[\Delta d_{t+i}]$, and so they expect future returns to be high.

This example makes clear that the presence of rational investors is essential for pro-cyclical asset return expectations to emerge. If only the optimistic investors were present, they would bid up asset prices, and hence $\frac{P_t}{D_t}$, so much that there would not be any reason, even under their optimistic expectations about Δd_{t+i}, to expect high returns. That a subset of the investor population is less affected by waves of optimism and pessimism than the typical household is arguably realistic. Furthermore, borrowing constraints that limit individuals' leverage are likely to limit optimistic individuals' ability to bid up asset prices to the full extent (as, e.g., in Geanakoplos (2009)). From this perspective, it is natural that asset price booms are associated with optimistic return expectations on the part of households/individuals that are typically less sophisticated than professional investors or forecasters.

Turning now to the question how pro-cyclicality in return expectations influenced spending, one possibility is that asset return expectations have a direct effect on consumers' consumption/savings decision because return expectations affect the intertemporal tradeoff between consumption now and consumption in the future. How consumers preferences for consumption in the future vs. consumption today respond to expected rates of return is governed by the intertemporal elasticity of substitution (IES). Unless the IES is very high, expectations of higher future returns will lead to higher consumption today, elevating current spending on housing, durables, nondurables, and services at the expense of saving and investment.[4]

A second channel that probably played a more important role in the run-up to the Great Recession works through the interplay of asset prices and borrowing constraints. An initial

[4]Fuster, Hebert, and Laibson (2011) present a model in which of excessive extrapolation from recent data influences macroeconomic dynamics through this intertemporal tradeoff.

rise in asset prices generates positive asset return experiences which generates optimism about further rises in asset prices. This optimism has two effects. First, it supports higher asset prices, and higher asset prices relax borrowing constraints, which opens up new opportunities for current household spending. Second, it makes households underestimate the probability that borrowing constraints might bind in the future, which increases their willingness to make use of the current relaxation of borrowing constraints and increase current spending.

Thus, the rise in housing prices in the years leading up to the peak in 2006 had a self-reinforcing character. Optimistic expectations fostered and sustained high housing prices. For existing homeowners, the resulting bubble in housing prices relaxed borrowing constraints.[5] Mian and Sufi (2010) estimate that the average homeowner extracted 25 cents for every dollar increase in home equity, and that home equity-based borrowing added $1.25 trillion in household debt. Mian and Sufi also present evidence that the funds raised were not spent on real estate or to pay down credit cards, which leaves durables spending as a likely destination of these funds.

As housing prices started falling, the whole process went into reverse. While positive house price appreciation experiences generated more optimism which further supported high house prices, the downturn lead more and more households conclude that residential real estate is not an asset that is likely to produce capital gains in the future. This reinforced (and may continue to reinforce) the fall in house prices. The fall in house prices in turn tightened borrowing constraints and forced many households to delever and reduce spending. As a result, declines in spending were concentrated among owners of real estate and securities (Petev, Pistaferri, and Eksten (2011)). Mian and Sufi (2011) show that declines in spending were particularly pronounced in counties in which house prices had risen and fallen the most and households most highly levered. The response of spending also appears to depend on the expectations about the extent to which the drop in asset prices is temporary: Chris-

[5]The macroeconomic effects of asset-price bubbles that relax borrowing constraints are analyzed in Martin and Ventura (2011), Kocherlakota (2009), and Farhi and Tirole (2012)). These papers consider rational bubbles, but similar effects are to be expected when exuberant households (and lenders) fail to realize that price appreciation will be low or negative in the future.

telis, Georgarakos, and Jappelli (2011) find that the reduction in spending is particularly pronounced among households that viewed the stock market decline as permanent.

A comparison of the housing bust with the stock market downturn in 2000-01 also supports the interpretation that experience-induced waves of optimism and pessimism affect spending mainly through the interplay of asset prices and borrowing constraints. The stock return expectations data shown earlier in Figure 4 indicates substantial optimism around the peak of the stock market. Yet the subsequent disappointment of these optimistic expectations did not lead to a deep recession. A plausible explanation is that stocks are rarely used as collateral for highly levered borrowing, and the effects of stock market movements on spending in boom and bust are therefore relatively small.

To sum up, the theory of experience-based expectations formation provides an underpinning of the household balance-sheet recession view of the Great Recession. It helps understand why the boom in asset prices, the rise in household leverage, and hence the buildup in fragility arose in the first place.

3 How persistent are the effects of the Great Recession?

The fall in real estate and stock prices during the Great Recession has substantially altered the macroeconomic experiences, particularly of the younger cohorts. It is a popular hypothesis that extreme macroeconomic events have long-lasting effects. For example, the generation of "Depression Babies" is often thought to have been shaped permanently by the experience of the Great Depression. Similarly, the zeal for inflation-fighting commonly attributed to the German Bundesbank is often seen as a consequence of German experiences with hyperinflation in 1923 and currency reform in 1948 (see the discussion in (Shiller (1997)). It is not clear, though, to what extent the German public's support for a "hawkish" Bundesbank is really driven by experiences in the first half of the 20th century and to what extent it instead reflects the perception that Germany benefited from a "hawkish" Bundesbank during the more recent 1970s and 1980s.

There is little evidence to date that macroeconomic experience effects—at least with regards to their effects on macroeconomic expectations and financial risk-taking—are really permanent. The evidence in Malmendier and Nagel (2011) and Malmendier and Nagel (2012) suggests that household investment choices and macroeconomic expectations are explained best with a weighting scheme that downweights observations earlier in life (as the weights shown in Figure 1). Therefore, the effects of experience may be persistent, but they are not permanent.[6]

I now investigate how persistent the effects of recent stock market and housing market downturns on individuals experiences are likely to be, assuming that individuals weight past experiences with the weighting scheme in Figure 1. How an individuals' experienced stock return will look like in, say, the year 2020 depends on stock market returns between 2012 and 2020 that we have not observed yet. To make a projection, I calculate an expected path of future experiences by assuming that the stocks will earn their long-run historical average excess return (from 1890 to 2011) and house prices will appreciate by their historical average price appreciation (from 1890 to 2011) during future years.

Figure 6 presents the result. Panel (a) shows the projection for experienced stock market excess returns for two age groups (age < 50 and age ≥ 50) until the year 2031. Note that the lines in this plot do not follow cohorts over time, but instead hold age constant. For example, the dotted line (age < 50) captures the below 50-year-olds in each year on the horizontal axis, i.e., the point on the line for 2012 refers to below 50-year-olds in 2012, the point on the line for 2031 shows the projected experiences for below 50-year-olds in 2031. As the plot shows, little change is expected for stock return experiences of the two age groups. The reason is that in 2012 the experienced returns of both age groups are extremely close to the long-run historical excess returns of stocks. While the market crash in 2008/09 has lead to a temporary fall below the long-run mean, especially for younger individuals, the subsequent

[6]There is the possibility, though, that extreme events may have longer lasting effects than experiences of moderate events. The analysis in Malmendier and Nagel (2011) and Malmendier and Nagel (2012) is not designed to allow for greater persistence of extreme events.

(a) Experienced stock market excess returns

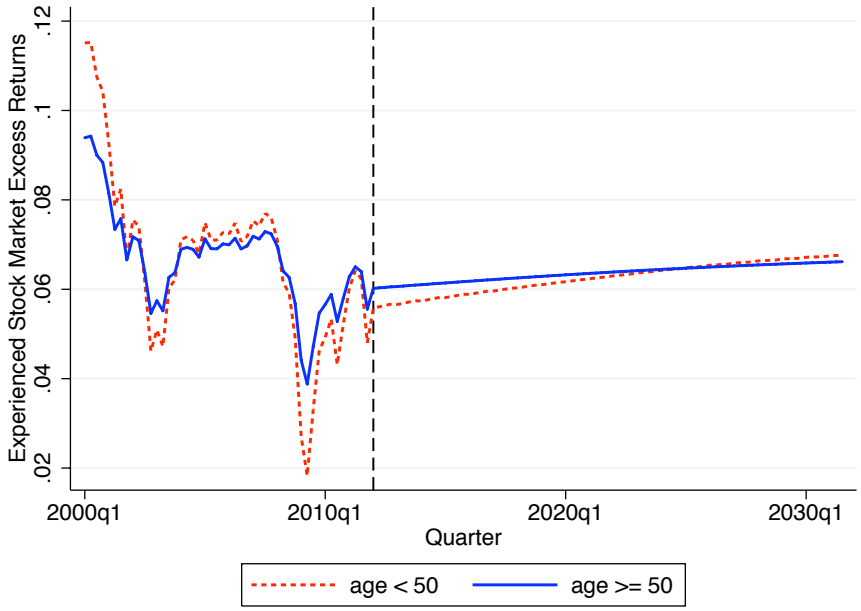

(b) Experienced house price appreciation

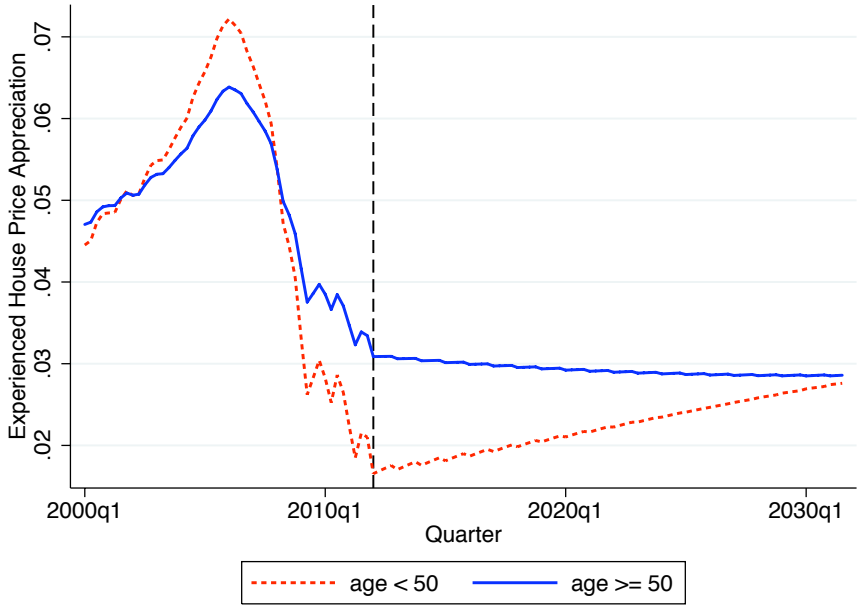

Figure 6: Expected path of future experienced stock market excess returns and house price appreciation. Experienced excess returns and house price appreciation calculated as in panels (b) of Figures 3 and 5, with sample average (from 1890 to 2011) excess returns or price appreciation substituted for future values from 2012 to 2031.

rebound in the market brought experiences back in line with the long-run historical mean.

For house price appreciation, shown in Panel (b), the situation is broadly similar, but experiences of younger individuals are currently about 1 percentage point below the long-run historical mean. If house prices do not fall further and resume growth at the historical long-run mean growth rate, experiences of young people should recover slowly, but the magnitudes are not big.

Thus, according to this analysis, there is little reason to expect much of a rebound in individuals' asset return experiences. Post-recession, the experiences are now mostly in line with long-run historical averages. The poor stock market returns over the last decade and the more recent housing crash have undone the effects of the earlier boom years, but they have not lead to unusually low asset return experiences as of 2012. As a consequence, individuals' expectations about asset returns should be close to their long-run means—neither unusually pessimistic, nor optimistic.

4 Concluding Remarks

Conditional on the realized path of house prices, economists understand fairly well why household spending rose before the Great Recession and fell afterwards: Relaxation of borrowing constraints and home-equity extraction during the boom, and deleveraging during the recession and the long slump that followed. What drove the boom in house prices and what motivated individuals to raise spending in response to relaxed borrowing constraints is less well understood. The recognition that individuals form macroeconomic expectations based on their own experiences helps shed some light on this. Rising asset prices foster optimism about future price appreciation, particularly among younger individuals, and during asset price bust this process goes into reverse.

References

Amromin, Gene, and Steven A. Sharpe, 2009, Expectations of Risk and Return among Household Investors, Working Paper, Federal Reserve Bank of Chicago and Federal Reserve Board.

Bracha, Anat, and Julian C. Jamison, 2012, Shifting Confidence in Homeownership: The Great Recession, Working Paper, Federal Reserve Bank of Boston.

Campbell, John Y., and Robert J. Shiller, 1988, The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors, *Review of Financial Studies* 1, 195–228.

Case, Karl E., and Robert J. Shiller, 2010, What Were They Thinking? Home Buyer Behavior in Hot and Cold Markets, Working Paper, Yale University.

Christelis, Dimitris, Dimitris Georgarakos, and Tullio Jappelli, 2011, Wealth Shocks, Unemployment Shocks and Consumption in the Wake of the Great Recession, Working Paper, Goethe University Frankfurt.

Farhi, Emmanuel, and Jean Tirole, 2012, Bubbly Liquidity, *Review of Economic Studies*, forthcoming.

Friedman, Milton, and Anna Jacobson Schwartz, 1963. *A Monetary History of the United States*. Princeton, NJ: Princeton University Press.

Fuster, Andreas, Benjamin Hebert, and David Laibson, 2011, Natural Expectations, Macroeconomic Dynamics, and Asset Pricing, in Daron Acemoglu and Michael Woodford, eds.: *NBER Macroeconomics Annual 2011*, Chicago, IL (University of Chicago Press).

Geanakoplos, John, 2009, The Leverage Cycle, in Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, eds.: *NBER Macroeconomics Annual*.

Greenwood, Robin, and Stefan Nagel, 2009, Inexperienced Investors and Bubbles, *Journal of Financial Economics* 93, 239–258.

Hall, Robert E., 2011, The Long Slump, *American Economic Review* 101, 431–469.

Kocherlakota, Narayana, 2009, Bursting Bubbles: Consequences and Cures, Working Paper, Federal Reserve Bank of Minneapolis.

Malmendier, Ulrike, and Stefan Nagel, 2011, Depression Babies: Do Macroeconomic Experiences Affect Risk-Taking?, *Quarterly Journal of Economics* 126, 373–416.

Malmendier, Ulrike, and Stefan Nagel, 2012, Learning from Inflation Experiences, Working Paper, Stanford and UC Berkeley.

Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers, 2003, Disagreement About Inflation Expectations, in Mark Gertler and Kenneth Rogoff, eds.: *NBER Macroeconomics Annual 2003*.

Martin, Alberto, and Jaume Ventura, 2011, Theoretical Notes on Bubbles and the Current Crisis, Working Paper, European Central Bank.

Mian, Atif, and Amir Sufi, 2010, House Prices, Home Equity-Based Borrowing, and the U.S. Household Leverage Crisis, Working Paper, University of Chicago.

Mian, Atif, and Amir Sufi, 2011, What Explains High Unemployment? The Aggregate Demand Channel, Working Paper, University of Chicago.

Petev, Ivaylo, Luigi Pistaferri, and Saporta Eksten, 2011, Consumption and the Great Recession: An Analysis of Trends, Perceptions, and Distributional Effects, Working Paper, Stanford University.

Piazzesi, Monika, and Martin Schneider, 2012, Inflation and the Price of Real Assets, Working Paper, Stanford University.

Shiller, Robert J., 1997, Why Do People Dislike Inflation?, in Christina D. Romer and David H. Romer, eds.: *Reducing Inflation: Motivation and Strategy*, pp. 13–65, Chicago (University of Chicago Press).

Shiller, Robert J., 2005. *Irrational Exuberance*. Princeton, NJ: Princeton University Press.

Appendix

A Extraction of stock return expectations from categorical data

In the AAII survey, individual investors classify themselves into "bullish", "neutral," and "bearish" categories. Assume that in a survey period, subjective excess return expectations are uniformly distributed across survey respondents with mean μ, lower bound $\mu - \sigma$, and upper bound $\mu + \sigma$.[7] The parameters μ and σ can differ from one survey period to the next. Suppose that respondents sort themselves into the three categories based on fixed expected excess return thresholds: They report to be bearish if their subjectively expected return is lower than $\alpha - c$, with $c > 0$, neutral if it is between $\alpha - c$ and $\alpha + c$ and bullish if it exceeds $\alpha + c$. Denoting the fraction of respondents that are bullish with q_{bull}, those that are neutral with q_{neutral}, and those that are bearish with q_{bear}, we get

$$\frac{q_{\text{bull}} - q_{\text{bear}}}{q_{\text{neutral}}} = \frac{\mu - \alpha}{c} \tag{A.1}$$

Thus, the proportions of bullish and bearish investors allows us to uncover μ up to an additive constant α and a scale factor c. If the thresholds of return expectations at which investors consider themselves bullish or bearish are relatively constant over time, then α and a scale factor c should have little time-variation. For the period 1998-2007 a direct estimate of μ is available from the UBS/Gallup survey. I scale the ratio $(q_{\text{bull}} - q_{\text{bear}})/q_{\text{neutral}}$ so that it matches the mean and standard deviation of the excess return expectations from the UBS/Gallup survey during this period. This is the series I report as the stock market return expectations extracted from the AAII data in the main text.

B Extraction of house price appreciation expectations from categorial data

The Michigan Survey of Consumers asks respondents: "Generally speaking, do you think now is a good time or a bad time to buy a house?" The responses to this general question are difficult to interpret, because a variety of reasons besides future price appreciation (e.g., interest rates, credit availability, affordability) could lead respondents to believe that times are currently good or bad for the purchase of a house. The survey also further elicits the reasons—through an open-ended question—why a respondent thinks it is a good or bad time to buy a house. These responses are classified into a long list of categories, two of which are "Prices are going up" and "Capital appreciation: buying is a good investment". This information helps to determine whether individuals expect future price appreciation. However, these two categories will only occur after responding to the first question that it is a good time to buy a house. In contrast, the price appreciation beliefs of a respondent who

[7] Assuming a normal distribution instead leads to the formulation of Mankiw, Reis, and Wolfers (2003), which produces a virtually identical time series of expectations.

responds to the first question that it is a bad time to buy a house because of, say, high interest rates, but who thinks that prices will go up, will not be recorded, because only reasons why it is a bad time are asked for. In other words, price appreciation beliefs are only observed conditional on the respondent believing that it is a good time to buy a house.

To extract the implied price appreciation expectations, I use a similar a approach as for stock market return expectations in Appendix A, equation (A.1). The underlying assumption is that respondents sort themselves into the three categories based on fixed house price appreciation thresholds: They respond that it is a good time to buy a house *and* provide expected future price appreciation as a reason if their subjective expectation of house price appreciation exceeds $\alpha + c$ (in this case price appreciation reasons dominate other reasons why it is a good time to buy a house), they report that it is a good or bad time to buy a house *and do not* provide a future price appreciation reason for their response if their house price appreciation expectations fall between $\alpha - c$ and $\alpha + c$ (in this case price appreciation/depreciation views are dominated by other reasons), and they respond that it is a bad time to buy a house *and* provide future price depreciation as a reason if their house price appreciation expectation is lower than $\alpha - c$ (in this case price depreciation reasons dominate other reasons why it is a bad time to buy a house).

To calibrate the c and α, I use percentage expectations from Table 6 in Case and Shiller (2010). Similar to the method in Appendix A, I choose c and α to match the mean and standard deviation of home price appreciation expectations from Case and Shiller (2010) in the years when coverage overlaps (second quarters of 1988 and 2003 to 2009).

The problem in this setting is that c could be time-varying. For example, if individuals consider interest rates as attractive for a house purchase, they are, everything else equal more likely to respond that it is a good time to buy a house, but any price appreciation motive would have to be rather strong for it to dominate the interest-rate motive and be mentioned as a main reason for the belief that it is a good time to buy a house. By the same token, a price depreciation motive would have to be strong to dominate the positive interest-rate view and push the individual into the bad-time-to-buy-a-house category. Otherwise the price depreciation motive will remain unrecorded in the survey. Thus, if other factors besides the price appreciation motive are very positive or very negative as a whole, this has effects similar to a wide range for the neutral category, i.e. a large c in the stock return expectations case above in equation (A.1). For this reason, assuming c to be constant may not be a good assumption in this setting here. On the other hand, using

$$q_{\text{bull}} - q_{\text{bear}} = \frac{\mu - \alpha}{2\sigma}, \tag{A.2}$$

i.e., without scaling by q_{neutral}, produces (up to a constant scale factor) virtually identical time series. This indicates that time variation in c is probably not all that important (unless in time-variation in c exactly matches time-variation in σ).