

STOCK MARKET CRASH AND EXPECTATIONS OF AMERICAN HOUSEHOLDS

PÉTER HUDOMIET,^a GÁBOR KÉZDI^{b,c} AND ROBERT J. WILLIS^{a,d*}

^a *Department of Economics, University of Michigan, Ann Arbor, MI, USA*

^b *Department of Economics, Central European University, Budapest, Hungary*

^c *Institute of Economics, Hungarian Academy of Sciences, Budapest, Hungary*

^d *Population Studies Center and Survey Research Center, University of Michigan, Ann Arbor, MI, USA*

SUMMARY

This paper utilizes data on subjective probabilities to study the impact of the stock market crash of 2008 on households' expectations about the returns on the stock market index. We use data from the Health and Retirement Study that was fielded in February 2008 through February 2009. The effect of the crash is identified from the date of the interview, which is shown to be exogenous to previous stock market expectations. We estimate the effect of the crash on the population average of expected returns, the population average of the uncertainty about returns (subjective standard deviation), and the cross-sectional heterogeneity in expected returns (disagreement). We show estimates from simple reduced-form regressions on probability answers as well as from a more structural model that focuses on the parameters of interest and separates survey noise from relevant heterogeneity. We find a temporary increase in the population average of expectations and uncertainty right after the crash. The effect on cross-sectional heterogeneity is more significant and longer lasting, which implies substantial long-term increase in disagreement. The increase in disagreement is larger among the stockholders, the more informed, and those with higher cognitive capacity, and disagreement co-moves with trading volume and volatility in the market. Copyright © 2010 John Wiley & Sons, Ltd.



Supporting information may be found in the online version of this article.

1. INTRODUCTION

The stock market crash of 2008 and the subsequent financial crisis constitute a rare episode whose scope and implications fall outside the life experience of American households. Whether and how those events affect people's expectations is an important question. To the extent that expectations guide investment behavior, substantial changes in expectations due to the financial crash can lead to substantial changes in investment. Besides average beliefs of 'the representative household', the crisis may have an impact on heterogeneity of such beliefs.

This study uses data from the 2008 wave of the Health and Retirement Study (HRS) to study the impact of the crisis on people's expectations. We estimate the effect of the crash on the population average of expected returns, the population average of the uncertainty about returns (subjective standard deviation), and the cross-sectional heterogeneity in expected returns (an indicator of disagreement). We show estimates from simple reduced-form regressions on probability answers as well as from a more structural model that focuses on the parameters of interest and separates survey noise from relevant heterogeneity. The measurement strategy makes use of the fact that the respondents of HRS 2008 answered the survey during 12 months from February 2008 to February 2009, a time period that includes the time of the stock market crash in early October. We show that the date of interview is largely independent of the respondents' past expectations about the stock market, so even if the date of interview is non-random it is unlikely to bias our results.

* Correspondence to: Robert J. Willis, University of Michigan, 3254 ISR, PO Box 1248, 426 Thompson Street, Ann Arbor, MI 48106, USA. E-mail: rjwillis@umich.edu

Our analysis looks at changes in expectations during the HRS sampling period of February 2008 to February 2009. It may be useful to recall some of the important events during this period. The subprime mortgage crisis began well before 2008, but the Dow Jones peaked in October 2007 above 14,000. By early 2008, though, the Dow was down to 12,000, and the rest of the year was characterized by a general decline until the crash of October. March 2008 saw the failed bailout of Bear Sterns and its subsequent sale to JP Morgan, but the rest of the Spring and the Summer went relatively quietly. On 15 September Lehman Brother filed for bankruptcy. The financial system was thought to be in severe danger, and it took a few weeks of uncertainty and heated debates before the US Congress passed the TARP bill on 3 October. The fall of 2008 also witnessed the run-up to the Presidential election on 4 November, which focused many people's attention towards economic issues, but it also led to a natural uncertainty about future economic policy.

Figure 1 shows time series of four stock market variables over the course of the HRS sampling period. We divided the sampling period into four sub-periods on the figure: February to June, July to September, October to November, and December to February 2009. We shall use these sub-periods throughout our analysis; their definition was based on the stock market time series we discuss below.

Figure 1(a) shows the level of the Dow Jones Industrial Average and the VXD annualized volatility index.¹ After initial ups and downs, the level of the index started a substantial but gradual decline in June that stopped in August. The stock market crash hit in early October with a 3000-point drop in the Dow. The stock market experienced large swings in October and November, and the Dow reached a 6-year low of 7500 in late November. After some recovery and a brief period of stability, the Dow experienced another period of steady decline in the first months of 2009. During the entire period, volatility showed the mirror image of the time series in levels, except that its increase started in September, and it reached its maximum in October and November. Figure 1(b) shows the weekly volume of trade of the shares of the Dow Jones Industrial Average (DJIA) together with the trend of searches for the term 'Dow' on Google.² The latter variable is an indicator for the attention people give to news about the stock market. The figure shows a strong co-movement of the two time series: increased attention to stock market news coincided with increased volumes in March and July of 2008, February of 2009, and, especially, October of 2008. The Google index is normalized so that its 5-year average is one. The maximal 8.8 value in the first week of October means that almost nine times as many searches were made from the USA for the Dow Jones Industrial Average than in normal times. Looking at the two panels together, we can see that the volume of trade was the highest at times when the stock market index was decreasing, when uncertainty was increasing and when people paid a lot of attention to news about the market.

The main question of this paper is whether and how expectations changed during the stock market crash in early October 2008 and the following months. We compare post-crash expectations to those earlier in 2008. It is important to keep in mind that the baseline period was characterized by early signs of the crisis and a depressed stock market. Nevertheless, the comparison can shed light on the effect of a large and perhaps qualitatively different event compared to the more 'normal' declining market.

The crash may affect the population average of expected returns for various reasons. If people are unsure about the parameters of the returns process, they may use recent realizations to update their beliefs. In such a case, the crash would have a negative effect on everyone's expectations. If, on the other hand, people believe in mean reversion in stock market prices, the effect may be of

¹ The VXD index is derived from prices of options on the DJIA, and it measures the future (30-day) expected volatility of investors. Details of this index can be found at <http://www.cboe.com/micro/vxd/>.

² <http://www.google.com/trends>.

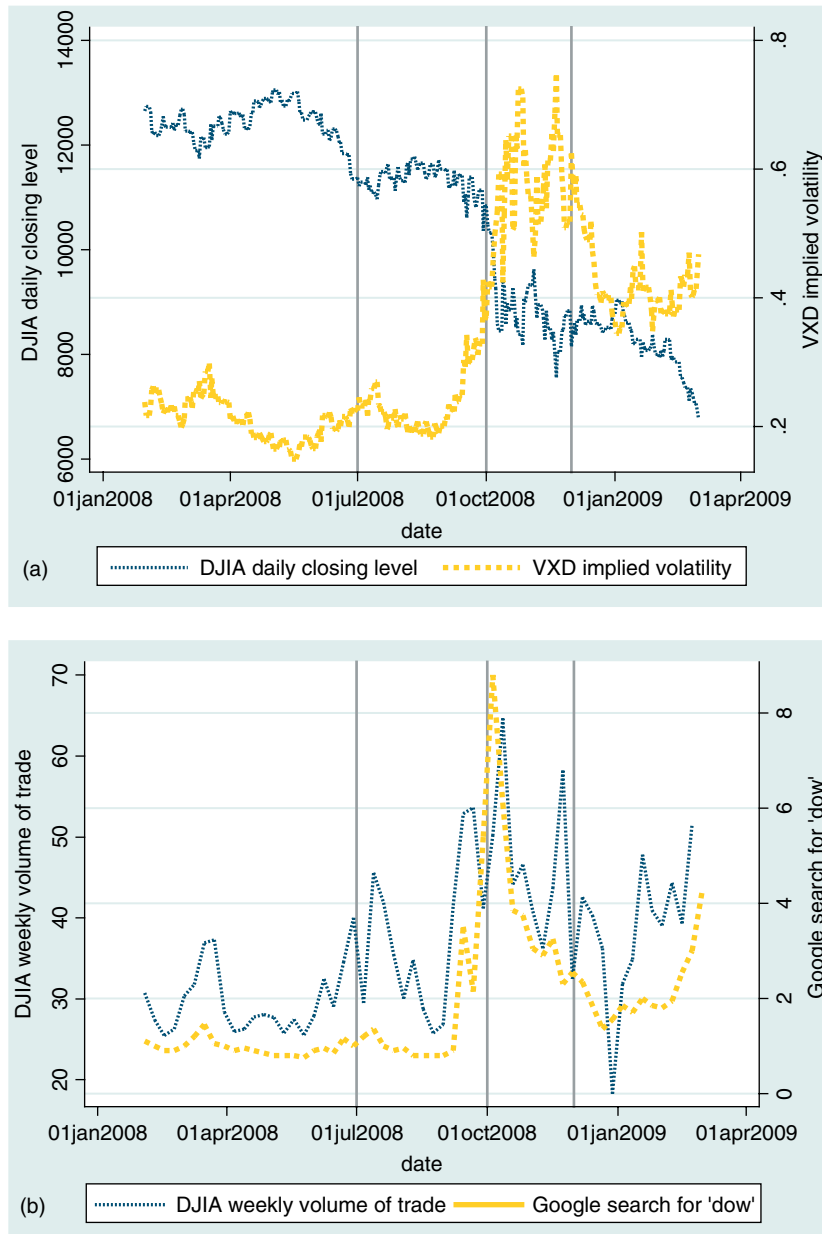


Figure 1. Level of the Dow Jones Industrial Average (daily closing), the VXD annualized volatility index, weekly volume of trade in billions of dollars and Google search for “Dow” from the USA in the sampling period of HRS 2008 (February 2008 to February 2009). This figure is available in color online at wileyonlinelibrary.com/journal/jae

the opposite sign. Of course, people may not want to update their beliefs if they don't learn from the returns. Besides stock prices, the political and policy news may have also affected people's expectations about the future of the economy and the financial sector in general, and the stock market in particular.

Empirical papers about stock market expectations usually find that average expectations track recent changes in the level of the stock market. When the stock market is increasing, average beliefs become more optimistic and conversely. See, for example, Kezdi and Willis (2008) about American households and Hurd *et al.* (2009) about Dutch households. According to Kezdi and Willis (2008), it took a 500-point gain in the Dow Jones to generate a one percentage point gain in expected yearly returns in 2002. With such a relationship, expected returns of respondents in November 2008 should be more than five percentage points lower than expected returns of respondents 2 or 3 months earlier. On the other hand, the financial crisis of 2008 may have affected people's expectations in qualitatively different ways from the more gradual changes witnessed in 2002, especially if people had different views about the condition of the economy in 2002 and in 2008. People may expect asset prices to change in different ways after large sudden changes than gradual trends. This is the conclusion of Calvet *et al.* (2009b) who, using Swedish data, found that people tend to invest in well-performing mutual funds but also tend to dispose of winning individual stocks at the same time.

The effect of the crisis on average uncertainty is more predictably positive. Stock market risk increased dramatically, as indicated by the trend in volatility on Figure 1. Even those who do not follow the stock market could become more uncertain about the future of the economy in general and the stock market, in particular, as general uncertainty has been 'in the air' throughout the crisis.

The crisis may also affect the cross-sectional heterogeneity in households' beliefs. Heterogeneity and potential subjectivity of people's beliefs about future stock market returns has been the focus of recent developments in finance theory (see Hong and Stein, 2007, for an overview of disagreement models in finance). Harris and Raviv (1993) and Kandel and Pearson (1995) show that public announcements can increase disagreement about the fundamental value of assets if people interpret the news in different ways (see also Kondor, 2009). As Hong and Stein (2007) observe, this pattern is precisely the opposite of what one would expect based on a simple rational expectations model with heterogeneous priors, where public information should have the effect of reducing disagreement, rather than increasing it. Similar mechanisms may increase disagreement after the stock market crash as well. Dominitz and Manski (2010), for example, assume that the population is a mix of people who believe in the random walk hypothesis, who believe in the mean reversion of stock-prices, and who believe in the persistence of trends on the financial markets. When the crash hit the economy and stock prices fell sharply, people holding these various views should have interpreted its implications in different ways, and consequently the disagreement among them should have increased. Indeed, a potential explanation of the trading pattern shown in Figure 1 is that the increase of disagreement created space for trade as more optimistic traders wanted to buy and more pessimistic traders wanted to sell. Note that potential heterogeneity in the effect of the crash implies that the average effect could go either way.

Our results imply a temporary increase in the population average of expectations right after the crash. At the same time, average uncertainty increased, perhaps as the result of increased stock market volatility. Our most robust finding is that cross-sectional heterogeneity in expected returns, an indicator of the amount of disagreement, increased substantially with the stock market crash. The effects are found to be largest among stockholders, those who follow the stock market, and those with higher than average cognitive capacity. The result on average expectations thus masks a wide distribution of effects of opposing signs. We also document the co-movement of stock market expectations with *ex post* returns, implied volatility and volume of trade.

Our finding suggests that there is heterogeneity in the cognitive processes (or mental models) people use to convert public news into personal probability beliefs, in accordance with some of the disagreement literature we mentioned above. The results on changes in heterogeneity complement recent empirical investigations that show substantial heterogeneity in stock market expectations of

individual investors (Vissing-Jorgensen, 2003) as well as households (Calvet *et al.*, 2007, 2009a,b; Dominitz and Manski, 2007; Kezdi and Willis, 2008; Hurd *et al.*, 2009; Gouret and Hollard, 2010). This paper adds new results to this empirical literature by showing that the stock market crash and the financial crisis had significant effects on average expectations, average uncertainty, and, perhaps most importantly, the heterogeneity of expectations.

2. DATA

We use stock market expectations data from HRS-2008. Before turning to our analysis, it is helpful to provide some background on the evolution of the HRS stock market expectation questions.

In 2002, HRS introduced probabilistic expectations questions about returns in the stock market to the battery of subjective expectation questions that have been asked in HRS since it began in 1992. One motivation for adding these questions is that expectations about stock returns are a key component in determining retirement saving and portfolio choice. In addition, stock market expectations are of methodological interest because the history of stock returns and their daily realizations are public information, enabling researchers to investigate how news affects the updating of beliefs without the need to adjust for differences in private information.

Like other HRS probability questions, stock market expectations are asked as a percent chance based on a '0' to '100' scale where the respondent is told that:

'0' means that you think there is absolutely no chance, and '100' means that you think the event is absolutely sure to happen.

The instruction goes on to say:

For example, no one can ever be sure about tomorrow's weather, but if you think that rain is very unlikely tomorrow, you might say that there is a 10 percent chance of rain. If you think there is a very good chance that it will rain tomorrow, you might say that there is an 80 percent chance of rain.

Beginning in 2002, the HRS introduced a question about stock market expectations that has been asked in every wave of HRS since 2002. We call this the P_0 question. It reads:

We are interested in how well you think the economy will do in the next year. By next year at this time, what is the percent chance that mutual fund shares invested in blue chip stocks like those in the Dow Jones Industrial Average will be worth more than they are today?

Like other HRS subjective probability questions, many answers to the HRS stock market questions are heaped on '50' (Hurd and McGarry, 1995) and, unlike most other probability questions, a substantial number of people fail to answer the stock expectation questions at all. A number of researchers have suggested that '50' is an indicator of 'epistemic uncertainty' or imprecise probability beliefs (Fischhoff and Bruine de Bruin, 1999; Lillard and Willis, 2001). Of course, it is also possible that some people who answer '50' mean that the event in question has a 50% chance of occurring or that they think that the probability falls within some range such as 40–60% and give '50' as a rounded approximation (Manski and Molinari, 2010).

Beginning in 2006 the HRS added an 'epistemic' follow-up question to several probability questions, including the P_0 question, to help understand the meaning of '50' answers:

Do you think that it is about equally likely that these mutual fund shares will increase in worth as it is that they will decrease in worth by this time next year, or are you just unsure about the chances?

We now turn to a discussion of the 2008 data that we use in this paper. In 2008, HRS continued to ask the ‘epistemic’ follow-up to persons who answered ‘50’ to P_0 . For those who did not respond ‘50’ to P_0 or, if they answered ‘50’, indicated that the shares were equally likely to increase or decrease in value, HRS added a follow-up question:

By next year at this time, what is the chance they will have grown by x percent or more?
(For negative values of x : By next year at this time, what is the chance they will have declined by $-x$ percent or more?)

where the probability of a gain of $x\%$ from the set of $\{+10, +20, +30, +40\}$ or a loss of $-x\%$ from the set of $\{-40, -30, -20, -10\}$ is randomly assigned.³ We denote the answer to this question as P_{x+} if the random value of x is positive and as P_{x-} if it is negative. Note that P_{x+} denotes the probability that returns would be greater than x , while P_{x-} denotes the probability that they would be less than $-x$.

The full sample consists of 17,217 individuals from 11,897 households. We restricted the sample to those 14,735 persons who participated in the last three waves of HRS (2004, 6, and 8). In 2004, the sample was refreshed by a new, younger cohort. Out of the 14,735 people, 2,850 (~19%) did not answer the P_0 question (the majority answered ‘I do not know’), leaving us a sample of 11,885 people. As we indicated earlier, HRS did not ask the P_{x+} or P_{x-} questions from those who stated that they were unsure in response to the ‘epistemic’ follow-up question (2005 individuals). Answer to the P_{x+} or the P_{x-} question is missing for another 486 individuals, and education was missing for an additional 45 individuals. Putting all these restrictions together, we ended up with a sample of 9348 individuals. The average age is 68 years, and 90% of the sample is 55–89 years old. We divide the sample into four subsamples based on the date of the interview (see Figure 1). These subsamples are very unbalanced in terms of the number of observations. 6285 respondents gave interview between February and June 2008, 2286 between July and September 2008, 556 in October and November 2008, and 211 between December 2008 and February 2009.

As we see, there are many missing values in the HRS stock market expectation data. Of the 14,735 people asked, only 9348 (63%) gave adequate answers to both questions. The two main sources of missing values are the ‘I do not know’ answers to any of the questions and being ‘unsure’ after giving a 50% answer to P_0 . In the analysis we shall ignore the missing values. We think that their omission does not invalidate our results for two reasons. First, people who ‘do not know’ or are ‘unsure’ might not have meaningful expectations about the stock market and thus they are not part of the population we would like to represent. These questions are not easy to answer, and if someone has no stocks and is sure that she will never have to deal with financial assets, she does not have to form expectations about the 1-year-ahead returns asked in the survey. Second, our goal is to analyze the changes in expectations after the crash. As long as the crash itself did not result in an increase or decrease of missing answers, the sample selection problem does not influence our main results. Analysis of the time series of missing answers reveals that the stock market crash did not bring about more ‘don’t know’ or ‘unsure’ answers. There is a small temporary decrease in the fraction of ‘don’t know’ answers in October, but the decrease is both quantitatively small and statistically insignificant.

³ Randomization of x was not complete in the survey: those who gave 0% for the P_0 question were assigned to get a random x with $x < 0$ but not $x > 0$, while those who answered $P_0 = 100\%$ were assigned to get a random x with $x > 0$ but not $x < 0$.

The distribution of P_0 answers is shown in the histograms of Figure 2. In part (a) we see the above-mentioned heaping at 50. Part (b) shows that the heaping disappears if we only leave in those 50 respondents who think that shares were equally likely to increase or decrease in value (rather than being unsure). Note that HRS did not ask the follow-up P_x questions from the ‘unsure’ people, so later in the analysis we will only use people from the panel (b).

Figure 2 also highlights another interesting issue: excessive rounding. Table I shows that almost 99% of the answers are multiples of 5, and more than 80% are multiples of 10. The fact that people give approximate answers to these probability questions is not surprising, since it is very hard to compute these numbers more precisely. A careful analysis should therefore incorporate this feature of the data.




Figure 2. Histogram of P_0 answers in the total sample (a) and in the final sample (b). This figure is available in color online at wileyonlinelibrary.com/journal/jae

Table I, however, highlights an even more important problem: that of inconsistent answer pairs to the probability questions. Strongly inconsistent answers are those that contradict the laws of probability: $P_0 < P_{x+}$ or $P_0 + P_{x-} > 100$. Zero mass answers are the ones that imply zero probability of returns between the asked probabilities: $P_0 = P_{x+}$ and $P_0 + P_{x-} = 100$. Nearly 17% of the answers are strongly inconsistent, and more than 21% imply zero mass. On top of these problems, Kezdi and Willis (2008) document that many HRS respondents do not give the same answer to the same probability question (say, P_0) when it is asked twice within the survey 20 minutes apart. Analyzing stock market expectations in another dataset, Gouret and Hollard (2010) show that few people give answers that imply the same expectations if they are asked in two slightly different ways within the same survey. Perhaps surprisingly, both Kezdi and Willis (2008) and Gouret and Hollard (2010) find no relationship between personal characteristics and the propensity to give problematic answers, with the potential exception of income and expectations themselves.

We argue that such answers are due primarily to question-specific survey noise due to inattention. Survey responses are the results of individual behavior under circumstances that differ from circumstances when making an actual investment decision. Answers are given in a matter of seconds and there are practically no incentives to get the answers right. Therefore, we would be wrong to assume that the survey answers are equivalent to the probabilities that represent people's subjective return distribution which forms the basis for their investment decisions. In Section 4 we propose a method to separate survey noise from relevant heterogeneity in expectations.

3. DESCRIPTIVE ANALYSIS

In this section we analyze the answers to the probability questions in a direct way. This should be viewed as preliminary descriptive analysis that cannot estimate the magnitude of the effect of the stock market crash, for two reasons: first, the probabilities themselves are not the objects of interest; second, survey noise can lead to biased estimates (especially on the heterogeneity of beliefs). At the same time, the descriptive analysis is free of additional assumptions that we need to make in order to recover more meaningful statistics.

Before the descriptive analysis, it is instructive to discuss how probabilities P_0 and P_x are related to the parameters of interest. Standard portfolio choice models include first and second moments of the (perceived) distribution of future returns as opposed to the probabilities themselves. With the help of additional distributional assumptions, answers to two probability questions can help identify the subjective mean and variance of the returns. Recall that the object of interest is the distribution of the 1-year-ahead returns of the stock market as viewed by the respondent. If we assume that people believe that the distribution of percentage returns is normal, two points in the

Table I. Fraction of rounded and inconsistent probability answers, HRS 2008

	Fraction
<i>Rounding</i>	
P_0 is a multiple of 10	0.806
P_0 is 5, 25, 75 or 95	0.140
P_0 is a multiple of 5	0.986
<i>Inconsistency</i>	
Strongly inconsistent answers*	0.169
Zero mass answers**	0.215
<i>N</i>	9348

Note: * $P_0 < P_{x+}$ or $P_0 + P_{x-} > 100$; ** $P_0 = P_{x+}$ or $P_0 + P_{x-} = 100$.

subjective distribution identify the entire distribution and thus both the mean and the variance. Figure 3 shows a normal cumulative distribution function (c.d.f.) that is identified by the two points. The figure depicts the case where the mean of returns is 0.07 and standard deviation is 0.15—numbers close to the post-war moments of nominal yearly returns on the Dow Jones (ending with year 2007). The probability of positive returns is around 68% ($1 - 0.32$), while the probability of returns of at least 20% (0.2) is around 20% ($1 - 0.80$ s). A respondent with the postwar–pre-2008 distribution in mind would answer P_0 to be 68% and P_{20} to be 20%.

Using answers to the two probability questions, one can in principle derive the mean (μ) and the standard deviation (σ) of the beliefs of individual i . Intuitively, the mean is identified from the level of the answers, while the standard deviation is identified from the distance between the two answers (larger distance means smaller variance). Formally, we can take the inverse of the appropriate probabilities:

$$P_{0i} = \Phi\left(\frac{\mu_i}{\sigma_i}\right), \quad P_{x+,i} = \Phi\left(\frac{\mu_i - x/100}{\sigma_i}\right), \quad P_{x-,i} = \Phi\left(\frac{x/100 - \mu_i}{\sigma_i}\right)$$

where P_0 is the answer to the probability of positive returns, P_{x+} is the answer to the probability of returns at least $x\%$, and P_{x-} is the answer to the probability of losses of at least $x\%$. Note that a mean-preserving spread in uncertainty (σ_i) pushes the probabilities towards 0.5, because an increase in σ_i moves the index towards zero. This is very much in line with the casual interpretation of a ‘fifty-fifty’ answer as reflecting ignorance. Using the example of positive x returns, inverting the probabilities would give this simple nonlinear but exactly identified system of two equations and two unknowns (μ_i and σ_i):

$$\Phi^{-1}(P_{0i}) = \frac{\mu_i}{\sigma_i}, \quad \Phi^{-1}(P_{x+,i}) = \frac{\mu_i - x/100}{\sigma_i}$$

Unfortunately, survey answers to the probability questions are not suited for such a direct transformation at the individual level. The excessive rounding and the relatively high fraction

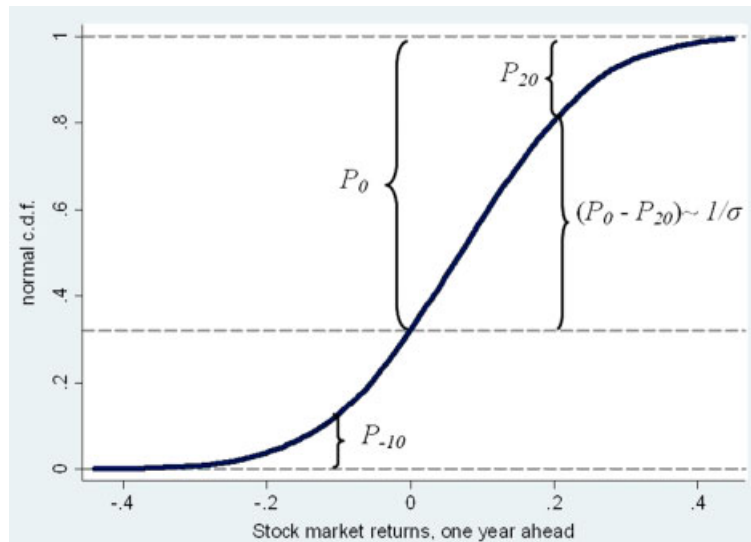


Figure 3. Standard normal c.d.f. ($\mu = 0.07$, $\sigma = 0.15$), with P_0 and P_{20} shown. This figure is available in color online at wileyonlinelibrary.com/journal/jae

of inconsistent probability answers discussed in the previous section would invalidate such an analysis. In the next section we propose a method for modeling both rounding and survey noise within a structural model. Before that we present some basic descriptive results in this section.

In order to see if the stock market crash brought about changes in expectations about stock market returns, we estimate simple ordinary least squares (OLS) regressions with crude proxies for the subjective mean (μ_i), the subjective standard deviation (σ_i) and the heterogeneity of expectations. In each regression, the right-hand side variables include three dummies for the four periods we focus on: February to June 2008 is the reference category, the first dummy is for July to September 2008, the second dummy is for October to November 2008, and the third dummy is for December 2008 to February 2009.

We estimate regressions with the probability answers themselves on the left-hand side in order to assess the effects on the population average of the level of the return distribution. If people become more pessimistic on average, we expect their answers to both the P_0 and the P_{x+} question to drop on average. If the second probability question has a negative threshold, their answer P_{x-} would go up on average. We therefore run two regressions: one with P_0 on the left-hand side, and one with P_{x+} or $1 - P_{x-}$ on the left-hand side. In order to partial out any threshold-specific factors that may bias answers to the second question, the second regression includes dummies for the different thresholds. The reference category is $x = +10$.

In order to see the effect of the crash on the cross-sectional heterogeneity of expectations (which we call disagreement), we look at regressions in which the left-hand side variables are the absolute values of the residuals from the previous regressions. If disagreement increases, the residuals from the previous regression would become more dispersed and their absolute value would therefore go up.

The effect of the crash on the population average of subjective uncertainty is approximated by a regression with the difference in the two probability answers on the left-hand side. Recall from Figure 2 that the difference between P_0 and P_{x+} is inversely related to the standard deviation of the subjective distribution. Another way to see the connection is in terms of the probability density function (p.d.f.): a larger difference would imply a larger probability mass concentrated on the support between P_0 and P_{x+} , which implies a less dispersed distribution. If the threshold of the second probability question is negative, the probability mass between P_0 and P_{x-} is given by $1 - (P_0 + P_{x-})$. In order for an increase in uncertainty to show up with a positive sign in the regressions, we used the negative of the differences for left-hand side variables: $P_{x+} - P_0$ for positive thresholds and $[(P_0 + P_{x-}) - 1]$ for negative thresholds.

Before we turn to the results of the regressions, we address the question of whether the date of the interview is exogenous to prior stock market expectations. This is our most important identifying assumption in analyzing the effect of the stock market crash. The interview date was not randomly assigned. The HRS released the names of all sample households to its national field staff of interviewers at the beginning of the field period in February 2008. Interviews were then completed in a sequence determined by each interviewer in consultation with regional field supervisors over the entire field period which ended in February 2009. Sample members who are hardest to locate, most difficult to schedule and most reluctant to be interviewed tend to receive interviews relatively late in the field period. Ultimately, over 90% of eligible sample members were interviewed.

In the 2004 and 2006 waves of the survey, HRS collected data on P_0 from respondents in our sample (but there were no second probability questions asked on stock market expectations). Using these variables we can look at whether the date at which people were interviewed in 2008 is related to their answers to the P_0 questions in previous interviews. We estimated four regressions with stock market expectation variables from 2004 and 2006 on the left-hand side and interview date in 2008 on the right-hand side. The first two regressions have P_0 on their left-hand side, while the

third and fourth regressions have the residuals from those regressions (in each pair one is for 2004 and the other is for 2006). According to the discussion above, these regressions estimate the ‘effect’ of interview date in 2008 on the average level of expectations prior to 2008 and heterogeneity of those expectations prior to 2008, respectively. The results from these ‘placebo’ regressions are shown in Table II. The only significant correlation with interview date in HRS 2008 and previous expectations is in column 2: those who answered HRS 2008 between October and November gave slightly higher P_0 answers in 2004 on average. At the same time, no such relationship was found in more recent 2006. Overall, the results suggest that the date of the interview in HRS 2008 was largely exogenous to stock market expectations prior to 2008—a result that is especially robust in terms of disagreement.

We can now turn to the effects of the interview date in 2008 on expectations in 2008. Table III shows the results. The dependent variables in columns 1 and 2 are the probability answers—our proxies for the population average of the level of the expectations. The results from the two regressions are very similar. The summer of 2008 brought no changes, and the average level was similar to the reference period December 2008 to February 2009 as well. However, October and November 2008 saw a significant, if temporary, increase in the average level of expectations.

Columns 3 and 4 report the results on the absolute value of the residual from the previous regressions, which are our proxies for disagreement. The estimates imply that disagreement stayed constant before October 2008, but it increased significantly after the crash. Contrary to the average level of expectations, the increase in disagreement lasted to the end of the sampling period. Column 5 shows the estimates on the difference between the two probability questions, which proxy the effects on the population average of uncertainty. Uncertainty seems to have increased already during the summer, and the crash brought about a substantially larger increase. Similarly to the average level of expectations, though, average uncertainty returned to its baseline level in the last period.

The results from these regressions suggest that, on average, people became more optimistic but also more uncertain after the crash, but those increases were temporary. Cross-sectional heterogeneity in expectations also increased after the crash, and that remained high a few months later as well. Unfortunately, as we highlighted earlier, these results are not suited for drawing quantitative conclusions for two reasons: they use crude proxies for the left-hand side variables of true interest, and they do not incorporate the complex survey response problems shown in the previous section.

Table II. Placebo regression results: OLS estimates with proxies for the level (columns 1 and 2) and heterogeneity (columns 3 and 4) of expectations in 2004 and 2006 as dependent variables, and the time of the interview in HRS 2008 as right-hand-side variables

Dependent variable	1	2	3	4
	P_0 in 2006	P_0 in 2004	$ u_{P_0} $ in 2006	$ u_{P_0} $ in 2004
Constant	49.376 (0.358)**	50.764 (0.351)**	20.41 (0.226)**	20.607 (0.220)**
Interview in HRS 2008 is July 08 to September 08	0.728 (0.695)	1.312 (0.680)	-0.432 (0.439)	-0.088 (0.427)
Interview in HRS 2008 is October 08 to November 08	0.878 (1.279)	2.489 (1.247)*	-0.714 (0.809)	0.196 (0.783)
Interview in HRS 2008 is December 08 to February 09	3.903 (1.992)	-1.819 (2.001)	1.394 (1.260)	1.235 (1.257)
<i>N</i>	7941	8444	7941	8444

Note: Standard errors in parentheses; * significant at 5%; ** significant at 1%. ‘ P_x ’ is defined as P_{x+} for positive thresholds and $(1 - P_{x-})$ for negative thresholds.

Table III. OLS regressions with proxies for the level (columns 1 and 2), heterogeneity (columns 3 and 4) and uncertainty (column 5) of expectations, HRS 2008

Dependent variable	1	2	3	4	5
	P_0	P_x	$ u_{P_0} $	$ u_{P_x} $	$P_x - P_0$
Constant	45.627 (0.360)**	64.813 (0.785)**	24.485 (0.184)**	21.434 (0.429)**	-24.532 (0.966)**
July 08 to September 08	-0.142 (0.698)	0.571 (0.628)	-0.39 (0.356)	0.173 (0.343)	1.702 (0.773)*
October 08 to November 08	4.922 (1.254)**	3.645 (1.128)**	1.714 (0.640)**	1.762 (0.617)**	4.343 (1.389)**
December 08 to February 09	0.041 (2.000)	1.06 (1.799)	2.562 (1.021)*	2.331 (0.983)*	-0.175 (2.215)
Dummies for the 'x' categories		Yes		Yes	Yes
N	9348	9348	9348	9348	9348

Note: Standard errors in parentheses; * significant at 5%; ** significant at 1%. ' P_x ' is defined as P_{x^+} for positive thresholds and $(1 - P_{x^-})$ for negative thresholds. ' $P_x - P_0$ ' is defined as $(P_{x^+} - P_0)$ for positive thresholds and $(P_0 + P_{x^-} - 1)$ for negative thresholds.

4. STRUCTURAL ESTIMATION

In the previous section we derived the relation between the probability answers and the first two central moments of the subjective return distribution under the assumption of normally distributed returns. Because of rounding and response error, as discussed earlier, these relations cannot be mechanically applied to the data. We incorporate rounding and survey noise in our model in two steps.

Assume that when making an investment decision individual i thinks of 1-year-ahead returns as R_i^* with mean μ_i and standard deviation σ_i . Throughout the analysis we assume that R_i^* is normally distributed. (Results are robust to alternative distributional assumptions of Student- t and log-normal as presented online in supporting information, Appendix B.) The survey answers of individual i are, however, based on a noisy version of R_i^* that we denote as R_{ji} (where j denotes the question so that $j = 0, x^+$ or x^-). The noise is assumed to be additive: the mean of R_{ji} is $\mu_i + v_{ji}$, where v_{ji} is a mean-zero noise variable specific to question and individual. The idea behind this assumption is that in a survey situation individuals have little time and no incentives to retrieve their subjective distribution of stock market returns. As a result, the subjective distribution they have in mind when answering the questions is likely to be different from the subjective distribution they would consider in an investment situation. We allow the noise terms to be different for the two probability questions (P_0 and P_x) but correlated across questions: $\text{corr}(v_{0i}, v_{xi}) = \rho$. The estimation model will allow for estimating both the variance of the survey noise and the correlation. When estimating the noise variance, we assume that it is proportional to subjective uncertainty σ_i . The intuition behind this assumption is those who have more diffuse expectations are likely to have a harder time retrieving those expectations. A consequence of this assumption is that $\text{var}[R_i^*]/\text{var}[R_{ji}]$ is constant. That is, this assumption ensures that the signal-to-noise ratio is constant in terms of perceived stock market returns.

A second feature of our model is that we consider interval responses instead of the reported probabilities themselves. If the reported probability (P_{ji}) is in a pre-specified interval or 'bin' $[b_1, b_2]$ then the 'true' probability (including the noise component v_{ji}) is assumed to be in the same bin but not necessarily the reported probability itself. Because a large fraction of the answers are multiples of 10 (see Section 2), we have defined 10 percentage point wide bins: $[0, 5)$; $[5, 15)$; ... $[95, 100)$. One consequence of this assumption is that a round answer can represent any expectation that would lead to probabilities around the particular round number.

The two assumptions are combined to

$$P_{0i} \in [b_1, b_2) \Leftrightarrow b_1 \leq \Phi \left(\frac{\mu_i + v_{0i}}{\sigma_i} \right) < b_2 \quad (1)$$

$$P_{x+,i} \in [b_1, b_2) \Leftrightarrow b_1 \leq \Phi \left(\frac{\mu_i + v_{x+,i} - x/100}{\sigma_i} \right) < b_2 \quad (2)$$

$$P_{x-,i} \in [b_1, b_2) \Leftrightarrow b_1 \leq \Phi \left(\frac{x/100 - \mu_i - v_{x-,i}}{\sigma_i} \right) < b_2 \quad (3)$$

where, as before, P_0 is the probability of positive returns, P_{x+} is the probability of returns of at least $x\%$, and P_{x-} is the probability of losses of at least $x\%$.

Using interval responses is quite common in the literature dealing with subjective probabilities, but the explicit modeling of survey noise and the maximum likelihood approach is not. For example, Manski and Molinari (2010) argue that, because of rounding, the parameters of interest are only partially identified, and they propose an alternative estimator based on the theory of partial identification and set estimation. Their conservative strategy resulted in very wide estimated parameter sets, especially on the HRS data, probably because of excessive rounding. To avoid this problem we have chosen instead to specify the model fully with distributional assumptions on all the unobserved random variables (see later).

We specify heterogeneity in the subjective mean and variance of returns by equations in two latent left-hand side variables μ_i , σ_i , of the form

$$\mu_i = \alpha'_\mu w_i + \beta'_\mu x_i + \gamma'_\mu z_{\mu i} + u_i \quad (4)$$

$$\log(\sigma_i) = \alpha'_\sigma w_i + \beta'_\sigma x_i + \gamma'_\sigma z_{\sigma i} \quad (5)$$

where w is the vector of date of interview dummies, x is the vector of covariates such as race, gender, age, education, and cognitive capacity, and the z vectors are equation-specific variables. We say more about them later when we discuss identification.

An important issue addressed in this paper is the possibility of increased cross-sectional heterogeneity in expectations, which may be labeled as disagreement. In order to capture disagreement, we let unobserved heterogeneity in μ vary with the date of the interview. Variance in u (unobserved heterogeneity in μ) measures the heterogeneity of expected returns among individuals who share the same x and z_μ variables. Formally, we let the standard deviation of u be related to the date of interview dummies (w) and the other covariates (x):

$$\log(\text{SD}(u_i)) = \alpha'_u w_i + \beta'_u x_i \quad (6)$$

The last equation is for the standard deviation of the noise, v , which is assumed to be proportional to σ_i :

$$\text{SD}(v_{ji}) = \lambda \sigma_i \quad (7)$$

Equations (4) and (5) describe the parameters of interest as effects on (or correlations with) the expected value of latent variables (μ_i , σ_i), and equation (6) captures the effects on (or correlations with) the standard deviation of the latent variable μ_i . These latent variables are mapped to the probability answers as specified by the interval response model in equations (1)–(3), which include additive question-specific noise components (v_0 and v_x) as well. The model is completed by distributional assumptions on unobservables u and v . We assume that u , v_0 and v_x

are jointly normally distributed and that unobserved heterogeneity, u , is independent of survey noise. However, we allow for v_0 and v_x to be correlated, and we estimate their correlation. One can argue that the correlation can be different for positive versus negative thresholds in the second question, and thus we estimate two correlation coefficients: one for v_0 and v_{x+} and one for v_0 and v_{x-} .

With these elements the model is complete and can be estimated using maximum likelihood. Before we turn to the results, it is worthwhile spending some time on identification issues. For simplicity, assume for a moment that there are no covariates on the right-hand side of (4)–(6). In this unconditional model we would have six parameters to estimate: μ , σ , $SD(u)$, λ , $\text{corr}(v_0, v_{x+})$ and $\text{corr}(v_0, v_{x-})$. In order to estimate them we need at least six moments. Interesting moments are $E(P_{ji})$, $V(P_{ji})$, $E(P_{0i} - P_{xi})$, $V(P_{0i} - P_{xi})$ and the fraction of inconsistent answers. Intuitively $E(P_{ji})$ and $E(P_{0i} - P_{xi})$ help identify $E[\mu_i]$ and $E[\sigma_i]$, while $V(P_{ji})$, $V(P_{0i} - P_{xi})$ and the fraction of inconsistent answers help identify $SD(u)$, λ and the correlations.

The estimation models include covariates and some exclusion restrictions as well. We use two instruments for μ (z_μ in equation (4)) and one for σ (z_σ in equation (5)). The first instrument for μ is the average probability that respondents assigned to the possibility of an economic recession in the near future in the previous two waves—2004 and 2006—of the survey. The second instrument is an average score on nine questions about depressive symptoms of the interviewees in 2004 and 2006, such as feeling lonely or feeling sad, again from the previous two waves of the survey. The instrument for σ is the fraction of 50 probability answers in 2004 and 2006. The idea behind using this variable is that people who are generally uncertain tend to give a lot of fifty-fifty answers to probability questions.⁴

5. RESULTS OF THE STRUCTURAL MODEL

The main question addressed by our analysis is how structural parameters of stock market expectations changed through the sample period. Using the model outlined in the previous section, we estimate changes in the population average of μ_i (the subjective expected value of returns), the population average of σ_i (the subjective standard deviation of returns), and the population standard deviation of u_i (unexplained heterogeneity in the subjective expected value). We capture the change of expectations in time by dummy variables for the four periods: February to June 2008 (the reference period, characterized by relatively high level of stock market indices and low volatility); July to September 2008 (gradual decline, relatively low but increasing volatility); October to November 2008 (the aftermath of the stock market crash and subsequently low levels and high volatility); and December 2008 to February 2009 (low levels with some further decline, and lower volatility). In order to help interpret the coefficients, all right-hand-side variables except the interview date dummies are normalized to have zero mean. As a consequence, the regression constant shows the expected value of the left-hand-side variable in the reference period (February to June 2008) for an average respondent in the sample. Note that the mean of the right-hand-side variables in the reference period is very close to the overall sample mean. As a result, the regression constant is very close to the actual average response in the reference period. The results are shown in Table IV.

The estimates are in line with the reduced-form OLS results of Table III. Average optimism about stock market returns increased temporarily in October–November: on average, people seemed to expect a recovery during this period. By December, the average expectations returned to where

⁴ This approach was first suggested by Lillard and Willis (2001) and has also been used by Sahn (2007) and Pounder (2007).

they were prior to the crash. Average uncertainty about stock market returns increased by 11% during the summer, and it increased again in October–November, by almost an additional 20%. However, average uncertainty seemed to return to its initial level afterwards. Unobserved cross-sectional heterogeneity in expectations increased by 13% in late summer as well, and it increased substantially in the fall. By October and November, the cross-sectional standard deviation was more than 50% larger than it was at the beginning of the year. Heterogeneity decreased somewhat after December, but it remained larger than before the crash.

The coefficients on the other right-hand-side variables indicate that women are significantly more pessimistic and more uncertain about stock returns; minorities are substantially more uncertain; older people are less optimistic and less uncertain; more educated people are less uncertain; and smarter people are more optimistic and less uncertain. Stockholders and those who follow the stock market are significantly more optimistic, and the latter are also less uncertain. The results imply that heterogeneity in beliefs is also different in different groups; those that are characterized by higher uncertainty on average tend to be more diverse in their beliefs. The sign on the variables that serve for exclusion restrictions are intuitive: those who were more pessimistic about the economy in the past have lower expectations on average, and the same is true for those with more depressive symptoms. The fraction of fifty-fifty probability answers in the past is a strong predictor of uncertainty about stock market returns.

The last three lines of Table IV contain estimates for the technical parameters. The ratio of true uncertainty to total variance that includes uncertainty as well as survey noise ($\sigma^2/(\sigma^2 + V(v))$) is constant by assumption and is estimated to be 0.645. This implies that the noise variance is almost as large as true uncertainty. The noise terms are allowed to be correlated across questions. The correlation is positive when both probability questions ask about returns higher than a particular threshold value. It is negative when the second question is about the probability of returns smaller than the predefined values.

At first sight, it is surprising that the population average of expected returns is negative during the baseline period. Note, however, that male stockholders who follow the stock market and have above average cognitive capacity expect substantial positive returns on average; their average μ of 0.06 is close to the pre-2007 historical mean of 0.07.

The coefficients on the interview date dummies in Table IV show overall changes in the average level, average uncertainty, and heterogeneity of expectations. It is interesting to see whether those changes were different in different groups. In order to examine such possibilities, we estimated the model with full interaction using dummy variables that split the sample into two parts. The first model with interactions distinguishes between stockholders and non-stockholders. The second model looks at those who follow the stock market versus those who do not. The third model looks at those whose cognitive capacity is above the average versus those below average. The coefficient estimates of the three models are in supporting information Appendix A. The main results are summarized below with the help of three figures.

We first look at stockholders versus non-stockholders. Stockholders include all those who owned stocks directly, through mutual funds or in tax-sheltered accounts such as 401(k) accounts. Since asset holdings are defined at the household level, members of the same household were assigned the same stockholder status. From an asset pricing point of view, the effect of the crash on stockholders is more interesting than the effect on other households. Note that stockholding may be endogenous to the financial crisis. Therefore we used the pre-crash stock-holding status from the 2006 wave of HRS. Figure 4 shows the results from the interaction model. The figure shows the 25th, 50th and 75th percentile of the population distribution of subjective expected returns (μ_i) for the four sampling periods. The distributions are recovered from the empirical distribution of the covariates and the normal assumption for the unobservables.

Table IV. Date of interview in 2008 and average subjective expected value of yearly stock returns (μ), average subjective standard deviation (σ) and unobserved cross-sectional heterogeneity in expectations ($SD(u)$). Results from structural regressions, HRS 2008

	μ	$\log(\sigma)$	$\log(SD(u))$
Constant	-0.088 (0.006)**	-0.606 (0.023)**	-1.23 (0.079)**
July 08 to September 08	0.001 (0.010)	0.113 (0.037)**	0.131 (0.047)**
October 08 to November 08	0.062 (0.025)*	0.292 (0.088)**	0.569 (0.099)**
December 08 to February 09	-0.028 (0.033)	0.019 (0.119)	0.38 (0.135)**
Female	-0.062 (0.009)**	0.235 (0.034)**	0.146 (0.044)**
Single	0.004 (0.010)	0.04 (0.039)	0.121 (0.048)*
Black	-0.017 (0.025)	0.589 (0.093)**	0.56 (0.096)**
Hispanic	0.002 (0.027)	0.387 (0.107)**	0.332 (0.114)**
Age	-0.002 (0.000)**	-0.005 (0.002)**	-0.004 (0.002)
Years of education	0.002 (0.002)	-0.034 (0.007)**	-0.031 (0.009)**
Above average cognition	0.031 (0.010)**	-0.101 (0.037)**	-0.198 (0.048)**
Follow the stock market	0.049 (0.010)**	-0.129 (0.038)**	-0.073 (0.047)
Stockholder	0.072 (0.010)**	-0.058 (0.038)	-0.18 (0.050)**
$P(\text{economic recession})$ 2004–2006 average	-0.003 (0.000)**		
Depressive symptoms 2004–2006 average	-0.017 (0.005)**		
Ratio of fifty-fifty answers 2004–2006 average		1.512 (0.191)**	
Log-likelihood	-42 277		
N	9348		
Mean(μ)	-0.085		
Mean(σ)	0.616		
Mean ($SD(u)$)	0.343		
$\sigma^2/(\sigma^2 + V(v))$	0.645		
Rho ($v_0, vx-$)	-0.491		
Rho($v_0, vx+$)	0.252		

Note: Standard errors in parentheses; * significant at 5%; ** significant at 1%. Reference categories: interview date February to June 2008, male, non-black and non-Hispanic, married.

The results of Figure 4 indicate that stockholders have substantially higher and less uncertain expectations, consistently with standard portfolio choice models. Note that the differences are not captured in full by the interaction of stockholding status with interview date dummies presented in Table A1, as the two groups differ in terms of the covariates as well (e.g., demographics and education). The median of the expected return distribution among stockholders is positive throughout the sample period, while the median of the non-holder distribution is negative. Changes in average μ (and average σ , see supporting information Appendix A) are similar in the two groups; heterogeneity among stockholders reacted to the stock market crash more in relative terms. In October and November 2008, the estimated inter-quartile range in expected returns rose from

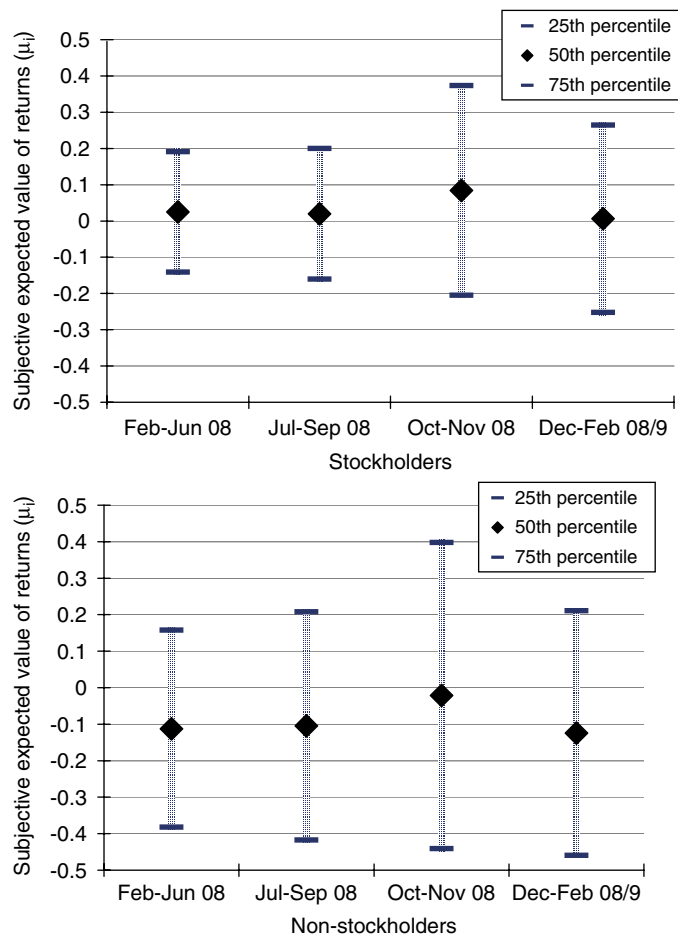


Figure 4. Cross-sectional distribution of expected returns among stockholders and non-stockholders. Quartiles of the distribution by the date of the interview, estimated from the structural model with interactions (detailed results in Table A1 in supporting information Appendix A). This figure is available in color online at wileyonlinelibrary.com/journal/jae

about 35 percentage points to almost 60 percentage points among stockholders and from around 60 percentage points to slightly more than 80 percentage points among non-stockholders.

Next we look at the results for the better informed versus less informed individuals. HRS 2008 asked how closely the respondent follows the stock market. 8.5% answered 'very closely' and another 36% answered 'somewhat closely'. The rest answered 'not at all' or did not know or refused to answer. We merged the 'very closely' and 'somewhat closely' categories and called the subsample 'informed respondents'. The rest we call 'uninformed respondents'. Being informed and stockholding are of course correlated, but the correlation is far from being perfect. 70% of stockholders claim to follow the stock market (and 30% do not), while 30% of non-holders claim to follow the stock market (and 70% do not). Note that, similarly to stockholding status, whether one follows the stock market is potentially endogenous to the stock market crash. Unfortunately, only HRS 2004 contains the information for our sample, and it is missing there for quite a few individuals. We decided to use the 2008 measures for the analysis despite its potential endogeneity. The results are very similar if one uses the 2004 measures instead.




Figure 5. Cross-sectional distribution of expected returns among those who follow the stock market (informed respondents) and those who do not (uninformed respondents). Quartiles of the distribution by the date of the interview, estimated from the structural model with interactions (detailed results in Table A2 in supporting information Appendix A). This figure is available in color online at wileyonlinelibrary.com/journal/jae

Figure 5 shows the results on the quartiles of expected returns in a way similar to the previous figure. The results are similar to the stockholder versus non-stockholder comparison, with some qualifications. The two groups differ less in terms of the initial level and heterogeneity of expectations than stockholders versus non-stockholders. The increase in the median of the distribution in October and November is more pronounced among the uninformed respondents, while the increase in the inter-quartile range is only marginally larger among the informed people. These results show that actual stockholding status is more strongly related to the effect of the stock market crash on expectations than whether one follows the stock market.

The third comparison is between those with above average cognitive capacity versus those with below average cognitive capacity. Cognitive capacity is measured by the principal component of various measures from HRS 2008. The measures include categories of self-rated memory, the score on immediate and delayed word recall and serial subtraction of seven from one hundred, and answers to three computing exercises, one of which is about compound interest rate. Cognition is correlated with whether one follows the stock market, but the correlation is not extremely strong

(66% of informed respondents are above average in terms of cognitive scores, compared to 40% of uninformed respondents). Figure 6 shows the results again in terms of the estimated quartiles of subjective expected returns. In terms of the median, the patterns are more similar to what we found for informed versus uninformed people, while the patterns in terms of the inter-quartile range are closer to the patterns by stockholding.

The results of the estimates suggest that the effect on the stock market crash on expectations was different in different groups of the population. The most pronounced differences are found between stockholders and non-stockholders. Differences between informed and uninformed people or higher cognition versus lower cognition people are similar but weaker. They all suggest, however, that disagreement increased more among those who initially agreed more.

The final question we investigate in this paper is whether the changes brought by the stock market crash are close to what one would predict by changes in different stock market indices. We seek to answer two questions. The first question is whether the patterns of stock market expectations found above are related to the evolution of the stock market. The second

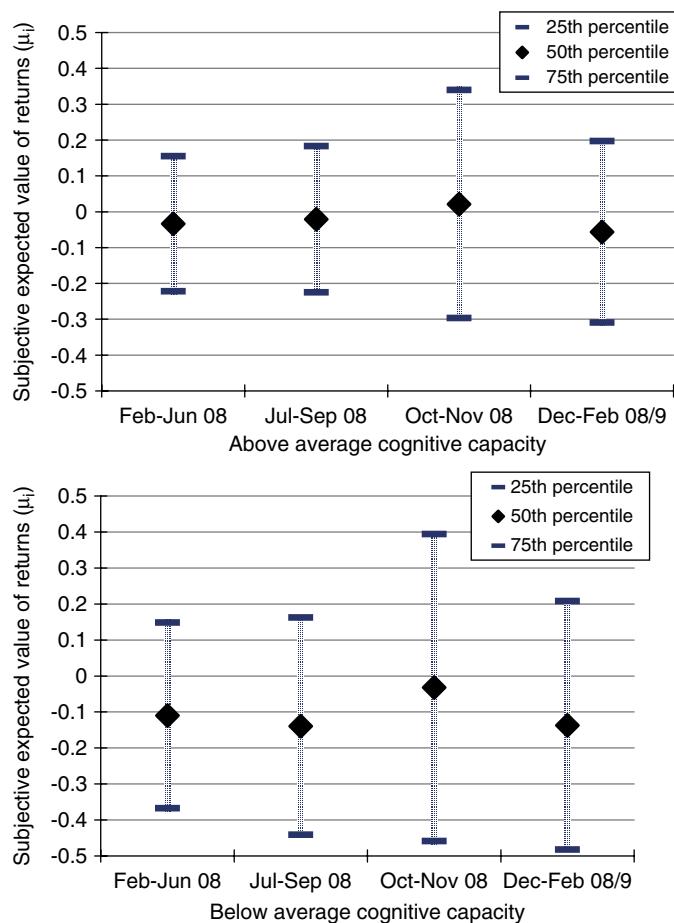


Figure 6. Cross-sectional distribution of expected returns among people with above average cognitive capacity and those with below the average. Quartiles of the distribution by the date of the interview, estimated from the structural model with interactions (detailed results in Table A3 in supporting information Appendix A). This figure is available in color online at wileyonlinelibrary.com/journal/jae

question is whether the link between the stock market indices and expectations broke after the crash.

We have created three indicators, all based on the Dow Jones index. The first is the monthly log-return, defined as the log of the average DJIA index from the 5 days before the interview minus the same lagged by 1 month. The second indicator is the average of the VXD annualized volatility measure from the 5 days before the interview. The third measure is the log of the average daily volume of trade of shares in the DJIA index, again from the 5 days before the interview. These indicators are defined from the same data as the series on Figure 1, but their exact definition is somewhat different. The first indicator enters the equation of expected returns (μ), the VDX indicator enters the equation of uncertainty (σ), and the trading volume indicator enters the equation of disagreement ($SD(u)$). The rationale for the last inclusion is that high trading periods might be the ones when traders disagree about the fundamental price of assets and thus volume patterns might be able to predict disagreement.

Table V contains the estimates from two different specifications. Specification 1 is identical to the specification of Table IV, except that the stock market indicators are entered instead of the date of interview dummies. Specification 2 differs from Specification 1 by allowing for an interaction of the stock market indicators with a dummy variable that is one if the interview date is after the crash (October 2008 to February 2009) and zero otherwise. If the relationship between the stock market indicators and expectations are the same before and after the crash, their coefficients should be stable across the specifications, and all the interaction terms should be zero.

The results for μ and σ are rather clear. They indicate that their relation to the stock market indicator (monthly returns and volatility, respectively) changed dramatically after the crash. Coefficients in Specification 2 suggest that before the crash an increase in the DJIA of 1%

Table V. The effects of recent returns and volatility of the stock market index and the daily volume of trade of the shares of the DJIA, before and after the crash. HRS 2008

	Specification 1			Specification 2		
	μ	$\log(\sigma)$	$\log(SD(u))$	μ	$\log(\sigma)$	$\log(SD(u))$
Constant	-0.085 (0.006)**	-0.568 (0.041)**	-8.955 (2.213)**	-0.084 (0.005)**	-0.785 (0.089)**	-8.454 (2.234)**
Monthly log returns (log of the average of previous 5 days minus the same 1 month before)	0.048 (0.081)			0.335 (0.104)**		
VXD volatility index (average of previous 5 days)		0.013 (0.151)			0.977 (0.408)*	
Log of volume of trade (average of previous 5 days)			0.350 (0.097)**			0.327 (0.099)**
Post-crash dummy (October 08 to February 09)				-0.007 (0.032)	0.210 (0.256)	19.295 (6.101)**
Post-crash dummy interacted with monthly log returns				-0.721 (0.257)**		
Post-crash dummy interacted with VXD volatility index					-0.617 (0.601)	
Post-crash dummy interacted with log of volume of trade						-0.839 (0.271)**
Other covariates	Yes	Yes	Yes	Yes	Yes	Yes
Instruments	Yes	Yes		Yes	Yes	
Log-likelihood	-42301.2			-42268.1		
N	9347			9347		

Note: Standard errors in parentheses; * significant at 5%; ** significant at 1%.

was followed by the population average of μ higher by 0.3 percentage points (i.e., 0.003). This magnitude is broadly in line with previous findings by, for example, Kezdi and Willis (2008). The post-crash relationship is just the opposite; there, the coefficient implies that the same 1% increase would be followed by a drop of 0.4 percentage points ($= 0.335 - 0.721$). This negative relationship is most likely identified from the fact that within a 1-month window from the crash the monthly log returns indicator was large and negative, while average expectations were higher than before. Specification 1 shows no relationship between returns and average expectations, because it shows the mixed results of a positive relationship between changes in the DJIA and the average level of expectations in the pre-crash period and the temporary increase in expectations after the crash. The population average of uncertainty shows a similar pattern; it tracks the volatility index before the crash in Specification 2 closely, but its increase after the crash is smaller than what the large increase of the VDX would have implied. Again, Specification 1 mixes the two and produces an insignificant estimate.

The results in Table V are less clear on the association between disagreement and the volume of trade. Our interpretation of the results is that it was largely similar before and after the stock market crash. An increase in the volume by 1% was associated with a subsequent increase in unobserved heterogeneity by somewhat over 0.3% before the crash in both specifications. The post-crash coefficients suggest a reverse association but also a huge increase in the intercept. Taken literally, they would imply that disagreement increased by astronomical magnitudes right after the crash, and from there it tracked trading volume with a negative coefficient. Recall, though, that trading volume jumped substantially right after the crash (see Figure 1), and it decreased considerably and steadily for most of the following time in our sampling frame. At the same time, disagreement stayed substantially higher after the crash than before, and it may have even increased in the immediate aftermath of the crash when volume dropped the most (although we would not have enough power to detect that). We argue that the post-crash association between volume and disagreement was dominated by the large increase in both right after the crash (a strong positive connection) and subsequent movements are of second-order importance.

All the results presented in this section are based on the assumption of normally distributed subjective yearly returns. We checked the sensitivity of our results to this functional form assumption by considering two alternatives: the Student- t distribution with various degrees of freedom and the shifted log-normal distribution. The Student- t has fatter tails than the normal. It is motivated by the model of Weitzman (2007) who showed that, if agents have imperfect knowledge about the 'true parameters' governing the stochastic process of stock market returns, and the parameters are evolving over time, then the posterior distribution of subjective returns can be Student- t . The shifted log-normal form is motivated by finance theory. While the log-normal is practically identical to the normal for small values of the return, it is quite different for larger values. All of our important results are robust to these alternative assumptions. The detailed results are available in supporting information Appendix B.

6. CONCLUSIONS AND DIRECTIONS FOR FUTURE RESEARCH

Using survey data on households' subjective probability beliefs about the 1-year-ahead return on the Dow Jones stock market index, we estimated the effect of the stock market crash on the population average of expected returns, the population average of the uncertainty about returns (subjective standard deviation) and heterogeneity in expected returns. We presented estimates both from reduced-form OLS regressions and a structural model that can estimate relevant heterogeneity in subjective expectations and incorporates survey noise at the same time.

We used data from the Health and Retirement Study that was fielded in February 2008 to February 2009. We identified the effect of the crash from the date of the interview, which we showed to be exogenous to previous stock market expectations. The estimated effects are qualitatively similar in the reduced form regressions and from the structural model, and they are robust to the functional form assumption for the distribution of stock market returns. The results show a temporary increase in the population average of expectations and uncertainty right after the crash. The effect on cross-sectional heterogeneity is more significant and longer lasting, which implies substantial long-term increase in disagreement. Stockholders were found to have more positive, less uncertain and less heterogeneous expectations than non-stockholders, but the stock market crash led to a larger increase in disagreement among them than among non-stockholders. We found similar but smaller differences between those who follow the stock market and those who do not, as well as between those whose cognitive capacity is above the average and those whose cognition is below the average.

The large positive effect of the crash on disagreement suggests that there is heterogeneity in the cognitive processes (or mental models) people use to convert public news into personal probability beliefs, in line with the models of Harris and Raviv (1993) and Kandel and Pearson (1995). The differential effects on stockholders versus non-stockholders, and similar differences between informed and less informed or by cognitive capacity, may be due to the fact that those different groups receive different signals or process the signals in very different ways. These results provide empirical evidence for future research on heterogeneous beliefs in finance theory.

Another natural question for further research is whether the changes in expectations we document lead to changes in asset allocation. Data from HRS 2009 and 2010 will allow for a thorough analysis of the effect of the crisis on the reallocation of household portfolios and the role of expectations.

ACKNOWLEDGEMENTS

Support from the National Institute of Aging (grants PO1 AG026571 and RO3 AG29469) is gratefully acknowledged. We also thank Chuck Manski, Steve Zeldes, Julia Kiraly, participants in the ‘Workshop on Subjective Beliefs in Econometric Models’ at Laval University, in the Paris School of Economics, in the Macro Seminar at the University of Michigan and, especially, the anonymous referees for helpful comments on earlier versions of this paper.

REFERENCES

- Calvet LE, Campbell JY, Sodini P. 2007. ‘Down or out: assessing the welfare costs of household investment mistakes. *Journal of Political Economy* **115**: 707–747.
- Calvet LE, Campbell JY, Sodini P. 2009a. Fight or flight? Portfolio rebalancing by individual investors. *Quarterly Journal of Economics* February: 301–348.
- Calvet LE, Campbell JY, Sodini P. 2009b. Household heterogeneity in financial markets: measuring the financial sophistication of households. *American Economic Review* **99**(2): 393–398.
- Dominitz J, Manski CF. 2007. Expected equity returns and portfolio choice: evidence from the health and retirement study. *Journal of European Economic Association* **5**: 369–379.
- Dominitz J, Manski CF. 2010. Measuring and interpreting expectations of equity returns. *Journal of Applied Econometrics* (this issue).
- Fischhoff B, Bruine de Bruin W. 1999. Fifty-Fifty = 50%? *Journal of Behavioral Decision Making* **12**: 149–163.
- Gouret F, Hollard G. 2010. When Kahneman meets Manski: using dual systems of reasoning to interpret subjective expectations of equity returns. *Journal of Applied Econometrics* (this issue).
- Harris M, Raviv A. 1993. Differences of opinion make a horse race. *Review of Financial Studies* **6**: (473–506).

- Hong H, Stein JC. 2007. Disagreement and the stock market. *Journal of Economic Perspectives* **21**(2): 109–128.
- Hurd M, McGarry K. 1995. Evaluation of the subjective probabilities of survival in the Health and Retirement Study. *Journal of Human Resources* **30**: S268–S292.
- Hurd M, von Rooij M, Winter J. 2009. Stock market expectations of Dutch households. DNB Working Paper No. 228, Netherlands Central Bank, Amsterdam.
- Kandel E, Pearson ND. 1995. Differential interpretation of public signals and trade in speculative markets. *Journal of Political Economy* **103**: 831–872.
- Kezdi G, Willis RJ. 2008. Stock market expectations and portfolio choice of American households. Mimeo, University of Michigan.
- Kondor P. 2009. The more we know, the less we agree: Higher-order expectations, public announcements, and rational inattention. Mimeo, Central European University.
- Lillard LA, Willis RJ. 2001. Cognition and wealth: the importance of probabilistic thinking. Working Paper UM00-04, Michigan Retirement Research Center.
- Manski C, Molinari F. 2010. Rounding probabilistic expectations in surveys. *Journal of Business and Economic Statistics* **28**: 219–231.
- Pounder L. 2007. Life cycle consumption examined. PhD thesis, Department of Economics, University of Michigan.
- Sahm CR. 2007. Stability of risk preferences. Working Paper 2007-66, Federal Reserve Board.
- Vissing-Jorgensen A. 2003. Perspectives on behavioral finance: does ‘irrationality’ disappear with wealth? Evidence from expectations and actions. (Discussions by John Campbell and Owen Lamont.) *NBER Macroeconomics Annual 2003*. MIT Press: Cambridge, MA.
- Weitzman ML. 2007. Subjective expectations and asset–return puzzles. *American Economic Review* **97**(4): 1102–1130.