

Stock market crash and expectations of American households^{*}

Peter Hudomiet
University of Michigan

Gabor Kezdi
Central European University

Robert J. Willis
University of Michigan

Preliminary and incomplete.

Abstract

This paper utilizes data on subjective probabilities of stock market returns from the Health and Retirement Study that was fielded in 2008 to study the impact of the crash on households' expectations about the one-year returns on the stock market index. Using a structural model that separates survey noise from relevant heterogeneity, 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. Interviews in the survey were distributed February 2008 through February 2009, in a way that is exogenous to previous stock market expectations. We use the interview date for identifying the short-run effects of the stock market crash that took place in September. We find no significant effect on the population average of expectations, but we do find some evidence for increased average uncertainty. A more robust finding is that cross-sectional heterogeneity in expectations increased substantially with the crash. Heterogeneity increased most among stockholders, those that follow the stock market and higher than average cognitive scores. We find that heterogeneity increased more than what increased volatility would have implied. The results are preliminary and incomplete.

^{*} Support from the National Institute of Aging grants PO1 AG026571 and RO3 AG29469) is gratefully acknowledged.

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. Using a structural model that incorporates survey noise, we estimate the effect of the crash on the population average of expected returns, the population average of the uncertainty about returns (the subjective standard deviation), and heterogeneity in expected returns. The measurement strategy makes use of the fact that the respondents of HRS-2008 answered the survey during twelve months in 2008 and 2009. Figure 1 shows the time series of the Dow Jones Industrial Average over the course of the sampling period, February 2008 through February 2009. Besides the level of the index (daily closing levels), Figure 1 shows daily volatility (the standard deviation of daily log returns over 30-day periods). The stock market started to decline in June, volatility started to increase in August, the crash hit in early October, and volatility peaked in early October as well. November saw a decrease in volatility and no recovery of the levels.

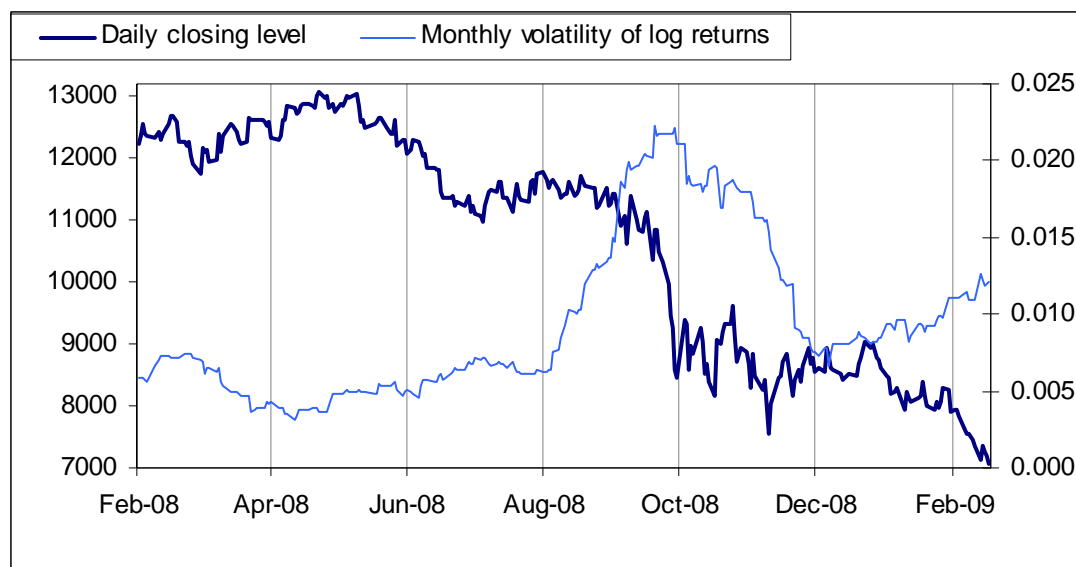


Figure 1. Level of the Dow Jones Industrial Average (daily closing) and volatility of log returns (standard deviation of daily log returns over 30-day periods) in the sampling period of HRS 2008 (February 2008 through February 2009).

We identify the effect of the crisis from the date of the interview: some of the respondents of HRS-2008 answered the survey before the stock market crash in September, while others are answering it after. The first identifying assumption is the exogeneity of interview date. We show

that the date of interview is largely independent of past expectations about the stock market. The second identifying assumption is be that observed changes in the probability answers are due to changes in relevant expectations as opposed to changes in survey response behavior. Since the questions are asked within the framework of the same survey with the same circumstances and incentives, we find this assumption reasonable.

The crisis may affect the population average of expected returns for various reasons. If people are unsure about the moments of the returns process, they may use recent realizations to update their beliefs. In such a case, the crash would have a negative effect. If, on the other hand, people believe in mean reversion in stock market prices, the effect may be of opposite sign. Of course, people may not want to update their beliefs if they don't learn from the returns series and the crash does not change their perceptions on the way the stock market works. Note that we are after the unweighted average of households' expectations. Heterogeneity can make average expectations go either way. Using within-year variation in interview dates in HRS 2002, Kezdi and Willis (2008) document that in their subsample (55 to 65 year-old individuals dealing with their households' finances), expectations track recent changes in the level of the stock market index. According to their estimates, in 2002, it took a five hundred point gain in the Dow Jones to result in one percentage point gain in expected yearly returns. With such a relationship, expected returns of respondents in November 2008 should be more than five percentage points lower than expected returns of respondents two or three months earlier. On the other hand, the financial crisis of 2008 may have affected people's expectations in qualitatively different ways than the more gradual (albeit also large) changes witnessed in 2002.

The effect of the crisis on average uncertainty is more predictably positive. Stock market risk increased dramatically, as shown by the trend in empirical 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 finance theory. The disagreement models surveyed by Hong and Stein (2007) focus on gradual information flows, limited attention, or heterogeneous prior beliefs. They do not offer an obvious effect of the crash that was sudden, loud, and large so very few people could avoid learning about it. On the other hand, public announcements are known to increase disagreement about future returns of assets if people interpret the news in different ways (see, for example, Evans and Lyons, 2008, and Kondor, 2005). Similar mechanisms may increase disagreement after the stock market crash.

We provide descriptive evidence, and we estimate a structural model in this paper. The results imply that the average level of expectations remained largely unaffected by the crash of the stock market. At the same time, average uncertainty seems to have increased with increased volatility. Even more robust is the result that cross-sectional heterogeneity in expected returns, an indicator

of the amount of disagreement, increased substantially with the stock market crash, likely caused by increased volatility.

Recent empirical investigations show substantial heterogeneity in stock market expectations of professional investors (Vissing-Jorgensen, 2003) as well as households (Dominitz and Manski, 2005; Kezdi and Willis, 2008). Kezdi and Willis (2008) estimate heterogeneity in both subjective expected value and uncertainty, and they show that heterogeneity in expectations is strongly related to heterogeneity in stockholding. They also show evidence for causality running (at least in part) from expectations to investment behavior. This paper contributes to this empirical literature both by its substantive results and by providing an application of the structural model of Kezdi and Willis.

The results are preliminary and incomplete.

2. Data

In this paper, we analyze stock market expectations using data from HRS-2008, building on an econometric model developed by Kezdi and Willis (2008) that uses stock market expectation data from HRS-2002. Before turning to our analysis, it is helpful to provide some background on the evolution of 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 are 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?

In 2002, HRS also asked the p_{10} question:

By next year at this time, what is the chance they will have grown by 10 percent or more?

In 2004 and 2006 only the p_0 question was asked.

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 percent chance of occurring or that they think that the probability falls within some range such as 40-60 percent and give “50” as rounded approximation (Molinari and Manski, 2008).

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 responded to p_0 and, 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 per cent or a loss of $-x$ percent is randomly assigned from the set of $\{-40, -30, -20, -10, +10, +20, +30, +40\}$.¹ 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.

We use the 2008 wave of the Health and Retirement Study. The sample consists of 17217 individuals from 11897 households. We restricted the sample to those 14735 persons who participated in the last three waves of HRS, when the sample was refreshed by a new, younger cohort. The average age is 69 years, and 90 percent of the sample is aged 55 to 89 years old. 53 percent of the interviews were made in the spring of 2008, 33 percent in the summer and 13 percent after September.

As we highlighted earlier, an important identifying assumption in analyzing the effect of the crisis is the exogeneity of the date of the interview. The interview date, however, was not randomly assigned. The HRS staff tried to conduct the interviews as early as possible, thus, those were interviewed at the second half of 2008 who were not available before. In order to see whether selection into interview date may be correlated with stock market expectations, we check how expectations in previous waves vary with the interview date in 2008. Figure 2 shows the answers to the 2004 and 2006 p_0 questions and their confidence intervals by the interview date in 2008.

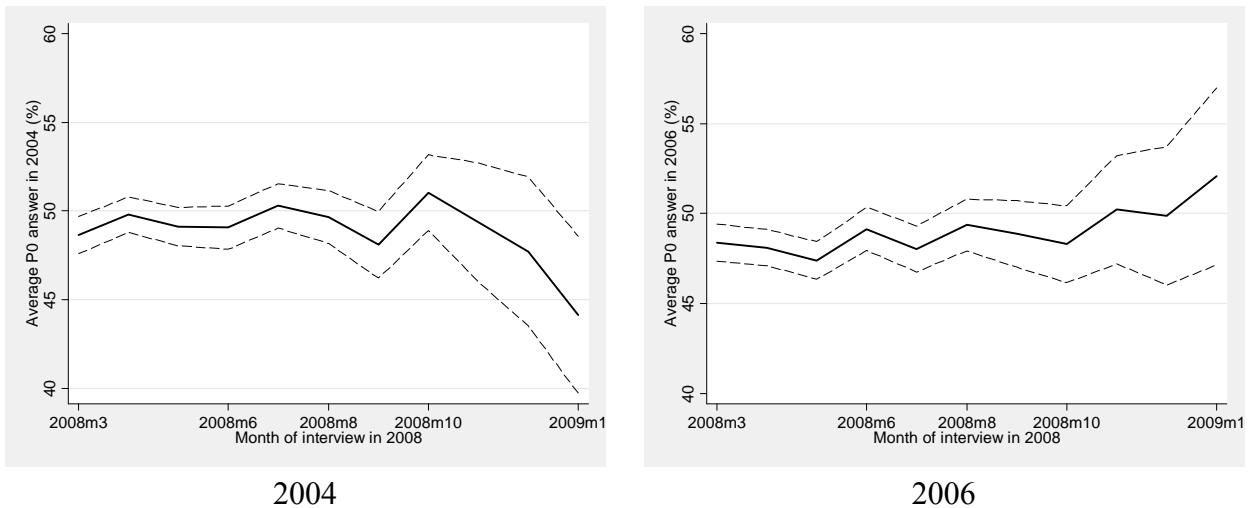


Figure 2. Answer to p_0 in 2004 and 2006 by month of interview in 2008. Average answers with 95% confidence intervals

¹ Randomization of x was not complete in the survey: those who gave 0 percent for the p_0 question were assigned to get a random x with $x < 0$ but not $x > 0$, while those who answered $p_0 = 100$ percent were assigned to get a random x with $x > 0$ but not $x < 0$. For some of the analysis (when randomness of x was crucial, see later), we re-established randomization by taking a random half of the $p_0 = 0$ and the $p_0 = 100$ subsamples and assigned $p_{x^-} = 0$ instead of the originally answered p_x .

We see a considerable widening of the confidence intervals after October 2008. This is due to a large decline in the number of respondents, from around 600 to 145 by January 2009. Past survey answers to the p_0 questions show no clear pattern with respect to interview date in 2008. While the 2004 answers exhibit a mild increase for the October 2008 respondents and a steady decline for later respondents, the 2006 show nothing for the October respondents, and, if anything, an increase for later respondents.

Figure 3 shows the average and the standard deviation of the answers to the p_0 question by month of interview, and their 95% confidence intervals.²

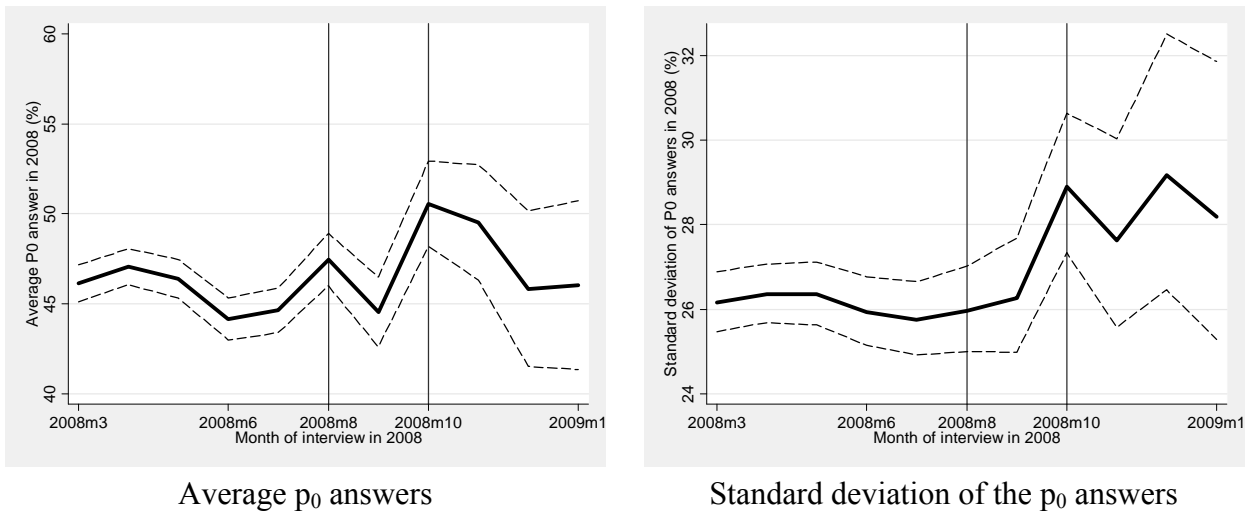


Figure 3. Answer to p_0 in 2008 by month of interview in 2008. Average and standard deviation of the answers, each with 95% confidence intervals.

Perhaps surprisingly, the effects of the stock market crash on the average of the probability answers do not seem to be robust. The average of the p_0 answers show a statistically significant decline for early summer (when the Dow started its decline), bounced back in August (when the decline of the Dow stopped), and declined again in September. Surprisingly, however, it increased considerably in October but that increase seems to have been temporary, although later changes are hard to assess because of the wide confidence intervals.

To the contrary, the heterogeneity of the p_0 answers shows a clear pattern: it increased significantly right around the stock market crash, and it remained high afterwards.

Finally, Figure 4 shows the fraction of missing answers among all respondents, and the fraction of those fifty answers that reflect ignorance (“unsure”) among all valid (non-missing) answers.

² For the confidence interval of standard deviation, the chi-squared approximation was used:

$$\sqrt{\frac{ns^2}{\chi^2(n, 0.025)}} \leq s \leq \sqrt{\frac{ns^2}{\chi^2(n, 0.975)}}$$

Since the vast majority of missing answers are “don’t know” answers, one can interpret those as reflecting large uncertainty as well (see Kezdi and Willis, 2008, for additional justification for this interpretation). Perhaps surprisingly, the descriptive results suggest that the stock market crash did not bring about more uncertainty on average, at least not in the form of more “don’t know” or “unsure” answers.

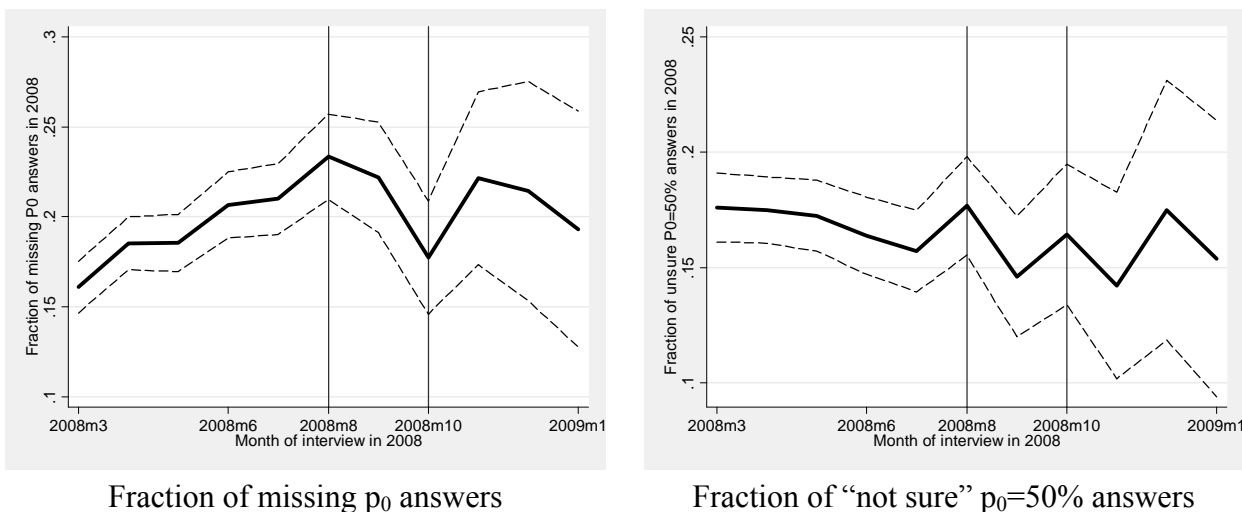


Figure 4. Fraction of missing answers to the p_0 question and the fraction of “not sure” answers in 2008, by month of interview in 2008. Each with its 95% confidence intervals.

3. Structural estimation

Answers to specific probability questions can show interesting patterns (or the lack of them), but the probabilities themselves are rarely the objects of interest. In particular, standard portfolio choice models include first and second moments of the (perceived) distribution of future returns as opposed to the probabilities themselves. Variation in the level of p_0 answers may be informative about variation in subjective means, but subjective variances are even harder to recover. With the help of additional distributional assumptions, however, 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 one-year-ahead returns of the stock market as viewed by the respondent. If we assume that people believe that distribution of percentage returns is normal, two points in the subjective distribution identify the entire distribution and thus both the mean and the variance. Figure 8 shows a normal c.d.f. that is identified by the two points. The figure depicts the case where mean of log returns is 0.07 and standard deviation is 0.15, numbers close to the post-war moments of nominal yearly log returns on the Dow Jones (ending with year 2007). The probability of positive returns is around 68 per cent ($1-0.32$), while the probability of returns of at least 20 per cent (0.2) is around 20 per cent ($1-0.80$). A respondent with the postwar-pre-2008 distribution in mind would answer p_0 to be 68 per cent and p_{20} to be 20 per cent.

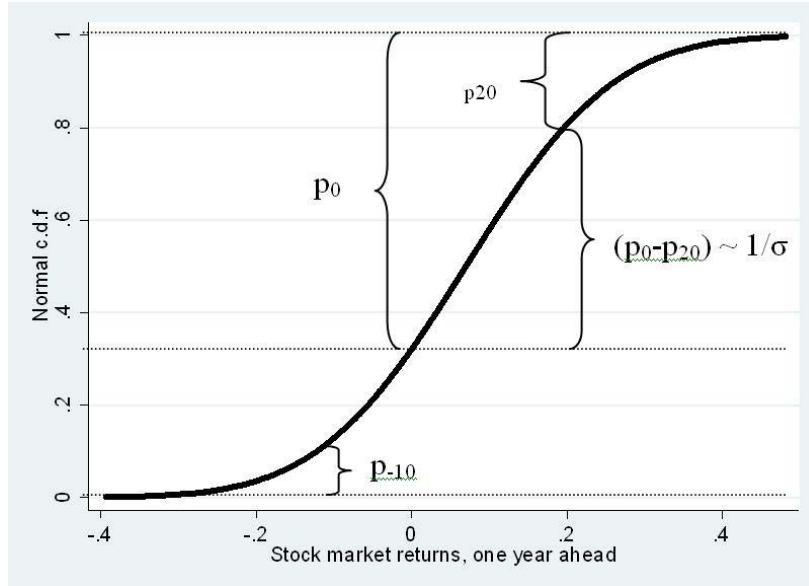


Figure 5. Standard normal c.d.f. ($\mu=0.07$, $\sigma=0.15$), with p_0 and p_{20} shown.

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 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 per cent, and p_{x-} is the answer to the probability of losses of at least x per cent. Note that 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 in 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. Survey responses are the results of individual behavior under circumstances that differ from circumstances when making an actual investment decision. There is considerably less time allowed, and there are practically no incentives to get the answers

right. Therefore, we would be wrong to assume that answers are equivalent to the probabilities that represent people's subjective return distribution which forms the basis for their investment decisions.

We can think of survey answers to probability questions as variables that contain considerable amount of non-standard noise. The histogram of the answers reveals pervasive rounding. Kezdi and Willis (2008) document additional evidence for survey noise. Many respondents give the same answer to the two probability questions, and quite a few give answers that contradict basic laws of probability ($p_{x^+} > p_0$ or, in 2008, $p_0 + p_{x^-} > 1$). Kezdi and Willis (2008) argue that such answers are due primarily to question-specific survey noise due to inattention. The sample from HRS 2002 has the fortunate element that some respondents answered the same pair of p_0 and p_x (there $p_x = p_{10}$ for all) questions twice in the course of the survey, some twenty minutes and many questions apart. The distribution of the answers to the same question (say, p_0 and $p_{0'}$) is the same, but their correlation is low, below 0.5. Kezdi and Willis (2008) show that neither the level, the absolute value nor the square of the "test-retest error" (the difference in the answers to the same question) is related to any observable variable, including stockholding.

Following Kezdi and Willis (2008) and in line with the observed characteristics of survey noise, we model the survey noise in two steps. First, we assume that the instead of μ_i , the answers are based on $\mu_i + v_{ji}$, where v_{ji} is a mean-zero additive noise variable specific to question p_{ji} ($j=0, x^+, x^-$). Variance of v is assumed to be constant across people and questions, and the noise terms can be correlated across questions $\text{Corr}(v_{ij}, v_{ik}) = \rho$. Second, we consider interval responses instead of the reported probabilities themselves. If the reported probability (p_{ij}) is in a pre-specified interval or 'bin' $[b_1, b_2]$ then the "true" probability (including the noise component v_{ij}) is assumed to be in the same bin but not necessarily the reported probability itself. One consequence of that 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

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

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

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

where, like before, p_0 is the probability of positive returns, p_{x^+} is the probability of returns at least x per cent, and p_{x^-} is the probability of losses of at least x per cent.

Within this model, there is no one-to-one mapping from the observed probability answers (p_{ij}) to the moments of interest (μ_i, σ_i). As a result, μ_i and σ_i are not identified at the individual level. At the same time, the population moments (both unconditional and conditional moments) are identified with the help of a structural model. Kezdi and Willis (2008) build a latent variable model and estimate its parameters by Maximum Simulated Likelihood.

In this paper, we use a slightly modified version of that model. 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

$$(4) \quad \mu_i = \alpha_\mu' w_i + \beta_\mu' x_i + \gamma_\mu z_{\mu i} + u_i$$

$$(5) \quad \log(\sigma_i) = \alpha_\sigma' w_i + \beta_\sigma' x_i + \gamma_\sigma z_{\sigma i}$$

In the equations, w is the vector of variables related to the stock market crash (quarter of interview or stock market variables related to the date of interview); x is the vector of covariates such as race, gender, age, education, and cognitive capacity; the z vectors are equation-specific variables. z_μ contains proxies for general optimism, while z_σ contains proxies for general uncertainty. Kezdi and Willis (2008) model portfolio choice jointly with expectations, but we do not look at portfolio choice in this paper. Importantly, if the date of interview is exogenous, the parameters of interest α_μ and α_σ are consistently estimated both with and without the other right-hand side variables.

An important extension considered in this paper addresses the possibility of increased cross-sectional heterogeneity in expectations, which may be labeled as disagreement. Disagreement may be captured in two ways. First, in some specification re-estimate the model in various subsamples or allow for interactions. Second, we let unobserved heterogeneity in μ vary with the date of the interview and/or events on the stock market. 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 variance of u to be related to the stock market crash variables denoted by w (quarter of interview or interview-date-specific stock-market variables):

$$(6) \quad \log \text{Std}(u_i) = \eta' w_i$$

Note that we assume that the moments of the noise components (v_0 and v_x) do not change throughout the sampling period. This is an important identifying assumption, and it implies that increased heterogeneity in answers can be entirely attributed to increased relevant heterogeneity in expectations (as opposed to increased survey noise). The assumption is motivated by the assumed independence of the date of the interview and the fact that all respondents answered exactly the same questionnaire, regardless of the date of the interview.

Equations (3) to (6) describe the parameters of interest as effects on (or correlations with) the latent variables (μ_i, σ_i). Those latent variables are then mapped to the probability answers as specified by the interval response model with additive question-specific noise components (v_0

and v_x) as described above. The model is completed by distributional assumptions on unobservables u and v . Using data from HRS 2002, Kezdi and Willis assume joint normality, independence of u and v , and $\text{Corr}(v_0, v_{10}) = \rho$. This correlation is a measure of how the survey noises for the two subsequent probability questions are related. In 2008, the second probability question was not the same for everybody. Instead, people were asked about the chance that returns will be higher or smaller than x or $-x$, where x was randomized. One can argue that the correlation can be different for different types of follow-up questions. As a consequence we estimated two correlation coefficients, one for a “return higher than x ” and one for a “returns smaller than $-x$ ” type question.

Unfortunately the variance of the survey noise is not identified; what’s identified is the variance of the sum of the noise (v) and the unobserved heterogeneity in μ (u). On the other hand, we can use outside information to calibrate the noise variance. Kezdi and Willis (2002) estimate this moment in two different ways. The first estimate uses the 2002 subsample with responses to the same probability question asked twice (see above). The second estimate uses the fact that in their structural model investment decision was also modeled. Because investment decision is independent of survey noise by assumption, the moments of the noise could be identified in that framework. Kezdi and Willis find that these two approaches give similar results. In this paper we rely on their estimates for the noise variance, with some modification. While they assumed that the noise variance is constant, we assume that it is a constant fraction of the variance of the subjective return. The reason for our deviation is that our sample is substantially more heterogeneous than that of Kezdi and Willis in terms of age. If R_n denotes the noisy expected returns and R denotes the noiseless one, then in our framework the quotient of their variances are kept constant: $\text{Var}(R_n)/\text{Var}(R) = (\text{Var}(v) + \sigma^2)/\sigma^2 = c + 1$ where c is calibrated to be 0.85, the result obtained by Kezdi and Willis for their sample. With these elements the model is complete and can be estimated using Maximum Likelihood.

4. Results on the date of the interview

In this section we look at how structural measures of stock market expectations changed through the sample period. Using the structural model outlined in the previous section, we estimate changes in average μ (the subjective expected value of returns), average σ (the subjective standard deviation of returns), and the standard deviation of u (unexplained heterogeneity in the subjective expected value). We estimate those moments for four periods that roughly define four quarters: February to May 2008 (the reference period, characterized by high level and low volatility); June to August 2008 (gradual decline, still relatively low volatility); September to November 2008 (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).

We first show results for the entire sample from two models, one with the quarter of the interview as the only right-hand side variable, and one adding the full set of covariates. Recall

that our assumption of exogenous interview date implies that the coefficients on the quarter dummies should be the same. The results are shown in table 1.

Table 1. Quarter of interview and average subjective expected value of yearly stock returns (μ), average subjective standard deviation (σ) and unobserved cross-sectional heterogeneity in expectations (Std(u)). Results from structural regressions. HRS 2008.

	Without other covariates			With other covariates		
	μ	$\log(\sigma)$	$\log[\text{Std}(u)]$	μ	$\log(\sigma)$	$\log[\text{Std}(u)]$
June 08 to August 08	-0.017 [0.010]	0.033 [0.037]	-0.013 [0.063]	-0.020 [0.009]*	0.056 [0.032]	0.000 [0.078]
September 08 to November 08	0.015 [0.016]	0.108 [0.059]	0.346 [0.082]**	0.004 [0.016]	0.145 [0.054]**	0.542 [0.088]**
December 08 to February 09	0.017 [0.050]	0.015 [0.162]	0.377 [0.219]	-0.023 [0.046]	-0.033 [0.145]	0.501 [0.218]*
female				-0.062 [0.009]**	0.198 [0.030]**	
single				-0.001 [0.010]	-0.014 [0.032]	
black				0.006 [0.017]	0.341 [0.065]**	
Hispanic				0.015 [0.021]	0.211 [0.077]**	
age				-0.002 [0.000]**	-0.005 [0.001]**	
educy				0.003 [0.002]	-0.016 [0.006]**	
numeracy08				0.012 [0.005]*	-0.065 [0.016]**	
memory08				0.015 [0.005]**	0.027 [0.016]	
followD3				-0.060 [0.009]**	0.091 [0.030]**	
prece0406				-0.003 [0.000]**		
depr0406				-0.018 [0.005]**		
fifty0406					1.691 [0.171]**	
Constant	-0.058 [0.006]**	-0.706 [0.022]**	-1.423 [0.039]**	0.249 [0.042]**	-0.576 [0.138]**	-1.693 [0.052]**
Log Likelihood		-212,000				
N		14,674				
Mean[μ]		-0.06			-0.09	
Mean[σ]		0.51			0.57	
Mean[Std(u)]		0.25			0.20	
Mean[Std(v)]		0.43			0.48	
Rho[v0,vx-]		0.00			0.00	
Rho[v0,vx+]		0.48			0.51	

Standard errors in brackets. * significant at 5%; ** significant at 1%
Reference categories: Interview date March to May, male, non-Black and non-Hispanic, married.

Coefficient estimates on the quarter of interview dummies are very similar across the two specifications, but they are stronger in magnitude when covariates are controlled for. This supports our main identifying assumption: the date of interview is approximately exogenous; if anything, it is correlated with unobservables in ways that bias the estimated effects of the stock market crash towards zero.

The estimates on the demographic covariates are broadly in line with the results of Kezdi and Willis (2008). Women tend to have lower and more uncertain stock market expectations. African Americans are more uncertain but do not differ in the level of their expectations, a finding different from Kezdi and Willis. Hispanics are more uncertain, education decreases uncertainty but its effect on the level is not significant. Numeracy and memory, on the other hand, increase μ , and people with higher numeracy scores have significantly smaller uncertainty. Finally, general measures of pessimism (the subjective probability of economic recession in 2004 and 2006, and the clinical depression variables in 2004 and 2006) are negatively related to stock market expectations, and the proxy for general uncertainty is strongly positively related to uncertainty about stock returns. The correlation between the survey noise components at a “return higher than x” type follow-up question is positive and high. It means that if somebody assigned too high value for the probability of positive returns than he will, on average, assign higher values for the probabilities that the return will be higher than x percent. The correlation is zero for the “returns lower than x” type follow-up questions.

Coefficient estimates on the quarter of interview dummies are very similar across the two specifications, but they are stronger in magnitude when covariates are controlled for. This supports our main identifying assumption: the date of interview is approximately exogenous; if anything, it is correlated with unobservables in ways that bias the estimated effects of the stock market crash towards zero.

Turning to the parameters of major interest in this paper, we can see substantial differences in the effect of the stock market crisis on the different aspects of expectations. Average optimism about stock market returns was largely unaffected by the crisis. Expectations decreased somewhat during the summer but returned to their previous levels afterwards. Average uncertainty about stock market returns increased a little during the Fall (statistically significant only in the second specification), but returned to its previous level afterwards. Heterogeneity in expectations, however, increased substantially in the Fall and remained high afterwards.

The main finding of the full-sample regressions is that heterogeneity in the level of expectations increased with the crash and remained large afterwards. In the remainder of the analysis we pool the first two quarters (March to August 2008) and the last two quarters (September 2008 to February 2009), and we focus on differences between those two periods by fixing the pre-crash period being the reference category. Table 2 shows results from the structural model estimated without covariates, with covariates, and with the covariates interacted with the post-crash

dummy. In order for the constant of these models to be comparable, all covariates (if entered) are measured as deviations from their sample mean. The table shows the parameters of major interest: the constant, the coefficient on the post-crash dummy and some additional information. The full set of estimates is in the Appendix.

Table 2. Pre-crash and post-crash estimates of average subjective expected value of yearly stock returns (μ), average subjective standard deviation (σ) and unobserved cross-sectional heterogeneity in expectations (Std(u)). Results from the structural regressions. HRS 2008.

	Without other covariates			With other covariates		
	μ	$\log(\sigma)$	$\log[\text{Std}(u)]$	μ	$\log(\sigma)$	$\log[\text{Std}(u)]$
Constant	-0.064 [0.005]**	-0.694 [0.018]**	-1.434 [0.033]**	-0.094 [0.005]**	-0.625 [0.019]**	-1.688 [0.044]**
Post-crash dummy	0.020 [0.015]	0.088 [0.055]	0.358 [0.075]**	0.009 [0.015]	0.108 [0.050]*	0.535 [0.080]**
Log likelihood		-212,124			-211,935	
Observations		14,674			14,674	

	With other covariates and their interactions with the post-crash dummy		
	μ	$\log(\sigma)$	$\log[\text{Std}(u)]$
Constant	-0.091 [0.007]**	-0.624 [0.022]**	-1.712 [0.046]**
Post-crash dummy	-0.026 [0.150]	-0.83 [0.393]*	0.543 [0.083]**
Log likelihood		-211,910	
Observations		14,674	

The first two models repeat the previous results in a more condensed way: the crash did not affect the average level of expectations, it led to a small and imprecisely estimated increase in average uncertainty, and it led to a substantial increase in unobserved heterogeneity in the level of individual expectations. The results from the model with the interactions show that it is indeed unobserved heterogeneity that increased with the crash. The parameter estimates of this model shown in table 2 are very close to those obtained without interactions. The full set of estimates shown in the Appendix table A1 reinforce this conclusion: none of the post-crash interaction variables are statistically significantly different from zero in the equation of μ . While the likelihood ratio test favors the model with the interactions (p-value is 0.0004), the only strongly significant coefficient is on the fraction of fifty answers, the proxy for general uncertainty, which is excluded from the equation of μ (that coefficient implies that uncertainty increased more among people who are less uncertain in general). The results on the level of expectations (μ) imply that we cannot identify groups whose expectations changed in a particular direction. Instead, it seems that all of the increase in disagreement is due to within-group increases.

A related but different question is whether we can identify groups in which disagreement increased more than in others. Table 3 to 5 show results from the models with covariates (but without interactions) in various subsamples, but the tables include the constant and the post-crash dummy only (right-hand side variables are again normalized to have zero mean within the relevant subsample).

We first look at stockholders versus not stockholders. The effect of the crash on the expectations of households that invest in the stock market is more interesting from the viewpoint of asset prices. Stockholding is defined at the household level. If there is more than one respondent from the household, we consider all of them as stockholders. Stockholder households include those that have stocks or mutual funds or in which at least one member has some of his or her retirement accounts invested in stocks or mutual funds. Note stockholding may be endogenous to the financial crisis. But the results using the theoretically more valid pre-crash stockholding are virtually identical to the results using the contemporaneous definition.

Table 3. By stockholding

	Stockholders			Not stockholders		
	μ	$\log(\sigma)$	$\log[\text{Std}(u)]$	μ	$\log(\sigma)$	$\log[\text{Std}(u)]$
Constant	-0.017 [0.006]**	-0.758 [0.024]**	-2.580 [0.209]**	-0.152 [0.009]**	-0.499 [0.030]**	-1.152 [0.043]**
Post-crash dummy	0.027 [0.019]	0.105 [0.064]	1.160 [0.233]**	0.002 [0.025]	0.085 [0.082]	0.246 [0.103]*
Log likelihood	- 66,274			-145,467		
Observations	5,789			8,885		

The results show that stockholders have substantially higher and less uncertain expectations, consistently with standard portfolio choice models and similarly to the results found by Kezdi and Willis (2008). While changes in μ and σ are similar (and never significant) in the two groups, heterogeneity reacted to the stock market crash in different ways. Disagreement after September 2008 increased especially among households that participate on the stock market. The increase was from $\text{Std}(u)=0.08$ to 0.24. Heterogeneity among not stockholder households was substantially larger in the pre-crash period ($\text{Std}(u)=0.32$) but increased only by a little with the crash, to 0.40. As a result, unobserved heterogeneity in the two groups became closer in the post-crash period.

Next we look at the results for the better informed versus worse informed individuals. HRS 2008 asked how closely he or she follows the stock market. 8.5 per cent answered “very closely”, and another 36 per cent 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 per cent of stockholders follow the stock market (and 30 per cent do not), while 30 per cent of non-holders follow the stock market (and 70 per cent do not).

The results are very similar to the stockholder – not stockholder comparison. Changes in the average level and average uncertainty are insignificant in both group, and heterogeneity of expectations increased substantially in the informed group while not significantly in the uninformed group. Overall, informed people have significantly higher μ and lower σ .

Table 4. By whether respondent follows the stock market

	Informed			Uninformed		
	μ	$\log(\sigma)$	$\log[\text{Std}(u)]$	μ	$\log(\sigma)$	$\log[\text{Std}(u)]$
Constant	-0.035 [0.006]**	-0.730 [0.022]**	-2.069 [0.088]**	-0.146 [0.009]**	-0.505 [0.033]**	-1.234 [0.050]**
Post-crash dummy	-0.008 [0.018]	0.082 [0.058]	0.828 [0.123]**	0.043 [0.026]	0.076 [0.094]	0.202 [0.121]
Log likelihood	- 64,388			-147,444		
Observations	6,562			8,112		

Last in this section are our results by cognitive capacity. High cognition respondents are defined to be those whose average of the numeracy and memory score is above average, while low cognition respondents are below average. Cognitive scores correlated with whether one follows the stock market, but the correlation is not very strong (66 per cent of informed respondents are above average in terms of cognitive scores, compared to 40 per cent of uninformed respondents).

Table 5. By cognitive scores

	Low cognition			High cognition		
	μ	$\log(\sigma)$	$\log[\text{Std}(u)]$	μ	$\log(\sigma)$	$\log[\text{Std}(u)]$
Constant	-0.057 [0.006]**	-0.689 [0.023]**	-2.360 [0.136]**	-0.138 [0.009]**	-0.519 [0.035]**	-1.077 [0.046]**
Post-crash dummy	0.004 [0.017]	0.138 [0.056]*	0.941 [0.170]**	0.021 [0.030]	0.098 [0.110]	0.346 [0.124]**
Log likelihood	- 93,053			-118,685		
Observations	7,609			7,065		

Despite the relatively low correlation of cognitive scores and following the stock market, these results are again close to the ones found there.

The final question we investigate is whether the changes brought by the stock market crash are close to what one would predict by changes in the level and the volatility of the stock market index. We are after whether effect of the crash is larger but qualitatively similar to the effects of smaller changes in those stock market indicators, or it is qualitatively different.

The results show that while volatility and heterogeneity are strongly positively related, disagreement after the crash increased more than predicted by increased volatility. The extra increase in heterogeneity is 23 log points, or about half of the overall increase in disagreement (54 log points, see table 2)

Table 7. Return and volatility of the stock market index in the previous month, before and after the crash.

	Stockholder		
	μ	$\log(\sigma)$	$\log[\text{Std}(u)]$
Post crash	0.004	0.038	0.231
	[0.015]	[0.062]	[0.115]*
Log returns (prev. month)	0.114		-0.293
	[0.118]		[0.878]
Volatility (prev. month)		5.817	19.793
		[3.179]	[5.600]**
Log likelihood		- 211,924	
Observations		14,674	

Table 8 shows that the previous statement is more true for stockholders: there more than half of the increase in unobserved heterogeneity is left unexplained by increased volatility..

Table 8. Return and volatility, by stockholding

	Stockholder			Not stockholder		
	μ	$\log(\sigma)$	$\log[\text{Std}(u)]$	μ	$\log(\sigma)$	$\log[\text{Std}(u)]$
Post crash	0.027	0.077	0.855	-0.005	-0.039	0.027
	[0.020]	[0.084]	[0.333]*	[0.025]	[0.098]	[0.133]
Log returns (prev. month)	0.209		0.158	0.044		-0.890
	[0.144]		[3.555]	[0.208]		[0.729]
Volatility (prev. month)		2.331	17.138		11.104	16.401
		[3.869]	[15.932]		[5.412]*	[6.987]*
Log likelihood		- 66,270			-145,462	
Observations		5789			8885	

5. Discussion

When interpreting the results, one has to keep in mind that only a small fraction of the sample was actually exposed to the financial crisis. The insignificant and somewhat sensitive effects may simply be caused by the low number of observations. The vast majority of interviews were concluded before September 2008. Although more than 1500 respondents were left for interview in or after September (about half in each subsample by sophistication), this number may not be enough to precisely estimate effects of the crash, given the amount of noise in the survey answers. Dramatic changes in expectations would show up in such small samples as well, as seems to have happened with the large increase in heterogeneity. But the insignificant effects on mean expectations and, to some extent, uncertainty, may be in part due to sample size issues. The result on increased disagreement about expected returns is all the more powerful in light of those issues.

The results on population averages could be consistent with fully informed and rational decision makers who believe in efficient markets. The lack of effect on the level of expectations might be expected for a sophisticated person who had absorbed the message of Burton Malkiel's best-selling *Random Walk Down Wall Street* (Malkiel, 1973, 2007 ninth edition). Alternatively, no effect on the level of expectations may be a plausible story if people thought the crash to be a one-time correction of the asset market bubble. This interpretation is also consistent with the estimated effects on average uncertainty: it responds to increased empirical volatility but the crash brings no extra uncertainty.

The former explanation is unable to account for substantial heterogeneity in beliefs and the large effect of the crisis on the amount of heterogeneity. There may be heterogeneity in the cognitive processes (or mental models) people use to convert public news into personal probability beliefs. That is, the "no average effect" we find might reflect the average across some people with a model of "mean reversion" who think post-crash returns will be unusually high, others with a "momentum model" who think returns will be unusually low along with "random walkers" who think the future will be like the past. Such heterogeneity may be consistent with some of the "disagreement models" of the stock market (Hong and Stein, 2007) and with empirical results reported by Dominitz and Manski (2005). This interpretation is strongly supported by our results on the large increase in heterogeneity after the crash.

6. Conclusions and directions for future research

Using the subjective probability data on households' beliefs about the one-year-ahead return on the stock market index, we built a structural model that can estimate relevant heterogeneity in subjective expectations and incorporates survey noise at the same time. Using that model, we estimated the effect of the crash on the population average of expected returns, the population average of the uncertainty about returns (subjective standard deviation), and heterogeneity in expected returns. Some of the results we show are different for more sophisticated households (those who follow the stock market) and less sophisticated ones (those who don't).

The results don't show a robust effect of the crash on the population average of expectations. They do provide some evidence for increased average uncertainty, and it is roughly in line with what one would expect with the observed increase in return volatility. Our most interesting finding is that cross-sectional heterogeneity in expectations increased with the crash. The increase is larger than what observed returns and volatility would have implied.

The insignificant results may be due to small of sample size. Future research will incorporate new data from the post-crash internet survey of HRS in 2009. That survey will contain stock market expectations of a large subset of the respondents on 2008, in the form of three questions: the probability of positive returns (p_0), the probability of a stock market gain of 20 percent or more (p_{20+}) and the other about the probability of a loss of 20 percent or more (p_{20-}). The ability to measure within-person variation in probability answers from before and after the crash will allow us to investigate heterogeneity in the way in which people convert public news about stock

prices, government policy and economic conditions into updated personal probability beliefs about one-year-ahead stock returns.

REFERENCES

- Barro, Robert J. (2006), "Rare Disasters and Asset Markets in the Twentieth Century." *Quarterly Journal of Economics*, 121, 823-866.
- Dominitz, Jeff and Charles F. Manski (2005) "Measuring and Interpreting Expectations of Equity Returns" NBER Working Paper No. 11313, May 2005
- Evans, Martin D.D. & Lyons, Richard K., 2008. "How is macro news transmitted to exchange rates?," *Journal of Financial Economics*, 88(1), pages 26-50
- Fischhoff, B. and W. Bruine de Bruin (1999), "Fifty-Fifty = 50%?" *Journal of Behavioral Decision Making*, 12, 149-163
- Hong, H. and J. C. Stein (2007). "Disagreement and the Stock Market." *Journal of Economic Perspectives*, 21(2), 109-128.
- Hurd, Michael D. (1999), "Anchoring and Acquiescence Bias in Measuring Assets in Household Surveys." *Journal of Risk and Uncertainty*, 19:1-3, 111-136.
- Kezdi, Gabor, and Robert J. Willis (2008) "Stock Market Expectations and Portfolio Choice of American Households." Mimeo, University of Michigan.
- Kondor, Peter (2005). "The More We Know, The Less We Agree: Public Announcements and Higher-Order Expectations." Mimeo.
- Malkiel, Burton (1973, 2007 ninth edition) *A Random Walk Down Wall Street*, W. W. Norton & Co., New York
- Manski, C. (2004). "Measuring Expectations." *Econometrica*, 72, 1329-1376.
- Manski, C. and F. Molinari (2009), "Rounding Probabilistic Expectations in Surveys," *Journal of Business and Economic Statistics*, forthcoming.
- Rietz, Thomas A. (1988). "The Equity Risk Premium: A Solution." *Journal of Monetary Economics*, 21, 117-32.
- Tversky, A. & Kahneman, D. (1974), "Judgment under Uncertainty: Heuristics and Biases." *Science*, 185, 1124-1130.
- Weitzman, Martin L. (2007). "Subjective Expectations and Asset-Return Puzzles." *American Economic Review*, 97(4): 1102-1130.

Appendix

Table A1. Detailed estimates from the models summarized in Table 2.

	With covariates			With covariates and interactions		
	μ	$\log(\sigma)$	$\log[\text{Std}(u)]$	μ	$\log(\sigma)$	$\log[\text{Std}(u)]$
Post-crash (Sep 08 to Feb 09)	0.009 [0.015]	0.108 [0.050]*	0.535 [0.080]**	-0.026 [0.150]	-0.83 [0.393]*	0.543 [0.083]**
female	-0.062 [0.009]**	0.197 [0.030]**		-0.064 [0.009]**	0.179 [0.032]**	
single	-0.001 [0.010]	-0.017 [0.032]		-0.004 [0.010]	0.009 [0.035]	
black	0.007 [0.017]	0.333 [0.065]**		0.01 [0.019]	0.415 [0.076]**	
Hispanic	0.015 [0.021]	0.208 [0.077]**		0.019 [0.023]	0.255 [0.088]**	
age	-0.002 [0.000]**	-0.005 [0.001]**		-0.002 [0.000]**	-0.006 [0.002]**	
educy	0.003 [0.002]	-0.016 [0.006]**		0.003 [0.002]	-0.019 [0.007]**	
numeracy08	0.012 [0.005]*	-0.065 [0.016]**		0.014 [0.005]**	-0.076 [0.017]**	
memory08	0.016 [0.005]**	0.026 [0.016]		0.017 [0.005]**	0.02 [0.017]	
followD3	-0.06 [0.009]**	0.089 [0.030]**		-0.062 [0.010]**	0.061 [0.033]	
prece0406	-0.003 [0.000]**			-0.003 [0.000]**		
depr0406	-0.018 [0.005]**			-0.016 [0.005]**		
fifty0406		1.690 [0.171]**			1.865 [0.187]**	
Post-crash \times female				0.032 [0.032]	0.153 [0.085]	
Post-crash \times single				0.014 [0.035]	-0.155 [0.094]	
Post-crash \times black				-0.082 [0.054]	-0.323 [0.149]*	
Post-crash \times Hispanic				-0.036 [0.066]	-0.175 [0.185]	
Post-crash \times age				0.000 [0.002]	0.010 [0.004]*	
Post-crash \times educy				0.003 [0.006]	0.024 [0.017]	
Post-crash \times numeracy08				-0.021 [0.017]	0.080 [0.044]	
Post-crash \times memory08				-0.014 [0.018]	0.057 [0.046]	
Post-crash \times followD3				0.038 [0.033]	0.214 [0.089]*	

Post-crash × prece0406				0.000		
				[0.001]		
Post-crash × depr0406				-0.021		
				[0.017]		
Post-crash × fifty0406					-1.405	
					[0.457]**	
Constant	-0.094	-0.625	-1.688	0.286	-0.796	-1.161
	[0.005]**	[0.019]**	[0.044]**	[0.072]**	[0.172]**	[0.052]**
Log Likelihood		-211,935			-211,910	
N		14,674			14,674	
Mean[μ]		-0.01			-0.15	
Mean[σ]		0.50			0.63	
Mean[Std(u)]		0.09			0.33	
Mean[Std(v)]		0.42			0.54	
Rho[v0,vx-]		0.23			-0.27	
Rho[v0,vx+]		0.52			0.51	

Standard errors in brackets. * significant at 5%; ** significant at 1%

Reference categories: Interview date March to May, male, non-Black and non-Hispanic, married.