Journal of Financial Economics 128 (2018) 403–421

Contents lists available at ScienceDirect

Journal of Financial Economics

journal homepage: www.elsevier.com/locate/jfec

Time varying risk aversion[☆]

Luigi Guiso[a,b], Paola Sapienza[b,c,d], Luigi Zingales[b,d,e,*]

[a] EIEF, Italy
[b] CEPR, UK
[c] Northwestern University, USA
[d] NBER, USA
[e] University of Chicago, USA

ARTICLE INFO

Article history:
Received 16 July 2015
Revised 17 October 2016
Accepted 18 October 2016
Available online 19 March 2018

JEL codes:
G11
G41
D91

ABSTRACT

Exploiting portfolio data and repeated surveys of an Italian bank's clients, we test whether investors' risk aversion increases following the 2008 crisis. We find that, after the crisis, both qualitative and quantitative measures of risk aversion increase substantially and that affected individuals divest more stock. We investigate four explanations: changes in wealth, expected income, perceived probabilities, and emotion-based changes of the utility function. Our data are inconsistent with the first two channels, while they suggest that fear is a potential mechanism underlying financial decisions, whether by increasing the curvature of the utility function or the salience of negative outcomes.

1. Introduction

As Campbell and Cochrane (1999) show, to fit historical data, asset pricing models require large fluctuations in the aggregate risk aversion. Yet, what is the direct evidence (i.e., not from stock prices) that aggregate risk aversion indeed fluctuates over time?

Aggregate risk aversion can fluctuate either because the risk aversion of the typical investor changes or because the distribution of wealth among investors with different risk aversion changes. In this paper we test the first channel and analyze whether *individual* risk aversion increases following the major financial crisis of the last 80 years—the 2008 one. We do so by exploiting portfolio choices and some survey-based measures of risk aversion elicited in a sample of clients (labeled *investors* from now on) of a large Italian bank (hereafter, *the bank*) in 2007 and repeated on the same set of people in 2009.

We find that both qualitative and quantitative measures of risk aversion exhibit large increases following the crisis. The risk premium required to accept a risky gamble with a 50% chance of winning 10,000 euros increases from 1000 euros to 2500 euros. Similarly, the fraction of investors who say they do not want to take any financial risk goes from 16% to 43%. Individuals who experience an increase in risk aversion are four times more likely to sell their stock holdings during the worst moment of the crisis than people who do not.

☆ We thank Nick Barberis, John Campbell, John Cochrane, James Dow, Stefan Nagel, Andrei Shleifer, Ivo Welch, Jeffrey Wurgler, and an anonymous referee for very helpful comments. We also benefited from comments from participants at seminars at the University of Chicago Booth, Boston College, University of Minnesota, University of Michigan, Hong Kong University, London Business School, Statistics Norway, European Central Bank, University of Maastricht, Warwick University, University of Montreal, the 2011 European Financial Association Meetings, the 2012 European Economic Association Meetings, the April 2013 NBER Behavioral Finance Meeting, UCLA Behavioral Finance Association, and Stanford University. Luigi Guiso gratefully acknowledges financial support from PEGGED, Paola Sapienza from the Zell Center for Risk and Research at Kellogg School of Management, and Luigi Zingales from the Stigler Center and the Initiative on Global Markets at the University of Chicago Booth School of Business. We thank Filippo Mezzanotti for excellent research assistantship, and Peggy Eppink for editorial help.

* Corresponding author. Present address: The University of Chicago Booth School of Business, 5807 S. Woodlawn Avenue, Chicago IL 60637, USA.

E-mail addresses: luigi@chicagobooth.edu (L. Guiso), luigi.zingales@chicagobooth.edu (L. Zingales).

https://doi.org/10.1016/j.jfineco.2018.02.007
0304-405X/© 2018 Published by Elsevier B.V.

There are many theories that can explain large changes in individual risk aversion. The best way to classify them is according to the channel that leads from the shock (large drop in stock prices) to the outcome (increase in individual risk aversion). The most prominent mechanism in the literature is changes in wealth, as predicted by the classical von Neumann-Morgenstern utility function and by the habit persistence model (Campbell and Cochrane, 1999). Prospect theory (Barberis et al., 2001) can also work through this channel.

Changes in background risk are the second most common explanation. Changes in the outside environment can affect an individual's expected income (Heaton and Lucas, 2000; Guiso and Paiella, 2008) and in so doing modify the risk aversion of the value function.

A third possibility is that a major shock has an effect on the expected distribution of returns as in Bordalo et al. (2012). In their model individuals' attention is directed to some particular realization that receives disproportionate weighting (salience). Therefore, individuals may become more risk averse because the financial meltdown has made the worst stock market realization more salient.

Finally, a major shock can affect the emotions of investors and alter their decisions about their willingness to take risks because it changes their perceived utility loss of bad outcomes (Loewenstein, 2000). In economic language, a major shock leads to a state-contingent increase in the curvature of the utility function.

Consistent with the wealth channel, we find that individuals who experience extraordinarily big losses exhibit a greater increase in the quantitative measure of risk aversion. Yet, we also find that risk aversion increases substantially even among those individuals who did not experience any loss, suggesting that not all the changes in risk aversion occur via changes in wealth.

We do not find much support for the changes-in-expected-income channel. For example, the increase in risk aversion of retirees (who in Italy enjoy a public pension) and of public employees (who at the time faced little or no risk of layoffs) is no smaller than that of the rest of the population.

We test the salience theory by looking at the individual responses on the expected distribution of returns. For those subjects willing to answer the question in both periods, we do find evidence of changes in the expected distribution of returns. Furthermore, we do find a significant increase in the fraction of people unwilling to answer that question.

Our evidence is also consistent with the Loewenstein (2000) hypothesis that, faced with a negative shock, individuals are affected by an emotion (fear) that alters their willingness to take risk in both financial and nonfinancial domains. However, with naturally occurring data it is difficult to design a direct test with any power to reject this hypothesis. For this reason, we run a laboratory experiment. While previous experiments have already shown that emotions can increase risk aversion (Kuhnen and Knutson, 2005, 2011; Knutson et al., 2008), our goal is to test whether the fear associated with a negative shock can indeed change our measures of risk aversion by a magnitude similar to what we observe in naturally occurring data.

To simulate in the lab this change in state, we rely on a fear conditioning model. As in the classical Pavlov (1927) experiment, the fear response can be triggered by conditioning factors, which have little or nothing to do with the experience itself. Kinreich et al. (2011) show that watching a horror movie stimulates the amygdala in a way consistent with the arousal of fear. Thus, to generate the fear produced by a stock market crash, we treat a sample of students with a five-minute excerpt from the movie *Hostel* (2005, directed by Eli Roth), characterized by stark and graphic images. It shows a young man inhumanly tortured in a dark basement.

We find that students treated with the horror movie exhibit a higher risk aversion (both according to the quantitative and the qualitative measure) that is very similar to the one experienced by the Italian bank's investors in 2009. The treated subjects' risk premium is $672 (27%) higher than the untreated ones'. Interestingly, the effect is entirely concentrated among students who dislike horror movies. The ones who like them seem unaffected.

Such an experiment shows that fear *causes* an increase in our measures of risk aversion, even in the absence of any change in the outside environment (which is the same for the treated and non-treated sample) and in their endowment (which is unaffected by the treatment). Obviously, the experiment cannot prove in any way that such a causal link exists among bank investors in our sample. Nevertheless, it does provide evidence that such a large increase in measured risk aversion can indeed occur even when not mediated by wealth changes and in absence of background risk. The psychology model based on fear is consistent with both the survey and the experimental data.

Our result is consistent with Cohn et al. (2015). In a lab experiment with a sample of financial professionals, they show that those "treated" with a stock market crash scenario become more risk averse and report an increase in fear, even though they do not experience any direct financial loss. This nice result is complementary to ours. Like us, they show that risk aversion can fluctuate with the stock market performance. Yet, we can show that an *actual* stock market crash, caused by the financial crisis, increases risk aversion and induces a change in portfolio allocation. Since Cohn *et al.* are limited to lab data, they are only able to show changes in the lab. However, they can successfully establish a causal link between the fear induced by the crash and a more conservative portfolio allocation, while we can only establish a correlation.

Our paper is also related to Weber et al. (2013). They survey online customers of a brokerage account in England between September 2008 and June 2009, asking them how they would allocate 100,000 pounds between a risk free asset and the UK stock market index and a few measures of risk attitudes. Similarly to us, they find that risk taking decreases between September and March, but, unlike us, their measures of risk attitudes do not change. One likely explanation for this difference is that their baseline measures are taken in September 2008 when the situation is already problematic, while our baseline measures are taken long before the inception of the crisis.

Finally, our paper is also related to the literature on market sentiments (see Baker and Wurgler, 2007 for a summary), on fear and risk aversion (e.g., Lerner and Keltner, 2000, 2001), and that on the effect of emotions and anxiety on risk attitudes, portfolio choice, and stock returns (Kamstra et al., 2003; Kramer and Weber, 2012; Kandasamy et al., 2014; and Bassi et al., 2013). Several of these papers establish that risk preferences vary over time and with emotions.

The rest of the paper continues as follows. Section 2 presents the data and our measures of risk aversion. Section 3 reports the results about the changes in risk aversion. Section 4 tests for possible explanations of these changes. Section 5 discusses how fear can be induced in a lab experiment and reports its results. Section 6 concludes.

2. Data description and measures

2.1. Sample

Our main data source is the second wave of the investors' survey run between June and September 2007 done by a large Italian bank. The survey is comprised of interviews with a sample of 1686 Italian customers. The sample was stratified according to three criteria: geographical area, city size, and financial wealth. To be included in the survey, customers must have had at least 10,000 euros worth of assets with the bank at the end of 2006. The survey is described in greater detail in Section A1 of the Internet Appendix, in which we also compare it to the Bank of Italy survey, which is representative of the Italian population.

Besides collecting detailed demographic information, data on investors' financial investments, information on beliefs, expectations, and risk perception, the survey collected data on individual risk attitudes by asking both qualitative questions on people's preferences regarding risk/return combinations in financial decisions as well as their willingness to pay for a (hypothetical) risky prospect.

For the sample of investors who participated in the 2007 survey, the bank gave us access to the administrative records of the assets that these investors have with them. Specifically, we can merge the survey data with administrative information on the stocks and on the net flows of 26 asset categories that investors have at the bank. We describe in detail the data set and its content in the Internet Appendix, Section A2. These data are available at monthly frequency for 35 months beginning in December 2006, and we use them to obtain measures of variation in wealth and portfolio investments over time. Since some households left the bank after the interview, the administrative data are available for 1541 households instead of the 1686 in the 2007 survey.

To study time variations in risk attitudes, in the spring of 2009, we asked the same company that ran the 2007 survey to run a telephone survey on the sample of 1541 investors interviewed in 2007 who were still clients of the bank. The telephone survey was fielded in June 2009 and asked a much more limited set of questions in a short 12-minute interview.[1] Specifically, investors were asked the two risk aversion questions (see below), a question about trust in their bank advisor or financial broker, and a question about stock market expectations using exactly the same wording that was used to ask these questions in the 2007 survey. Before asking the questions, the interviewer made sure that the investor was the same person who answered the 2007 survey by collecting a number of demographic characteristics and matching them with those from the 2007 survey.

2.2. Risk aversion questions

If we want to test whether changes in risk aversion can explain movements in asset prices, we need a way to infer risk aversion that is independent of asset prices. For this reason, we resort to survey-based measures.[2] We have two such measures. The first, patterned after a question in the US Survey of Consumer Finances, is a qualitative indicator of risk tolerance. Each participant is asked: "Which of the following statements comes closest to the amount of financial risk that you are willing to take when you make your financial investment: (1) a very high return, with a very high risk of losing money; (2) high return and high risk; (3) moderate return and moderate risk; (4) low return and no risk."

In a world in which people face the same risk-return tradeoffs and make portfolio decisions according to Merton's formula, their risk/return choice reflects their degree of relative risk aversion. In such a world, the answers to the above questions can fully characterize people's risk preferences. However, if people differ in beliefs about stock market returns and/or volatility, these differences will contaminate their answers to the above question. This bias would affect not only cross-sectional comparisons, but also intertemporal ones, possibly revealing a change in risk preferences when none is present. While we elicit expectations about stock market returns and volatility and control for them, the controls are not without errors.

As Fig. 1A shows, in 2007, only 16% of the sample chooses the "low return and no risk" answer to the qualitative risk aversion question. So most investors are willing to accept some risk if compensated by a higher return, but very few (1.8%) are ready to choose very high risk and very high return. From the answers to this question we construct a categorical variable ranging from 1 to 4 with larger values corresponding to greater risk aversion.

The second measure of risk aversion is designed to be independent of subjective beliefs. Each investor was presented with several choices between a risky prospect, which paid 10,000 euros or zero with equal probability

[1] Since the second survey was conducted during the same season as the first, the differences in risk aversion cannot be due to season variations in the length of day (see Kamstra et al., 2003).

[2] A potential alternative, followed by Friend and Blume (1975), is to infer an individual's relative risk aversion from his share of investments in risky assets. This method is not appropriate to study time series changes in risk aversion because the necessary maintained assumption is that portfolio shares are instantaneously adjusted. If not, any adjustment costs will be reflected in the estimated changes in risk aversion (Bonaparte and Cooper, 2009).

A. Qualitative measure of risk aversion

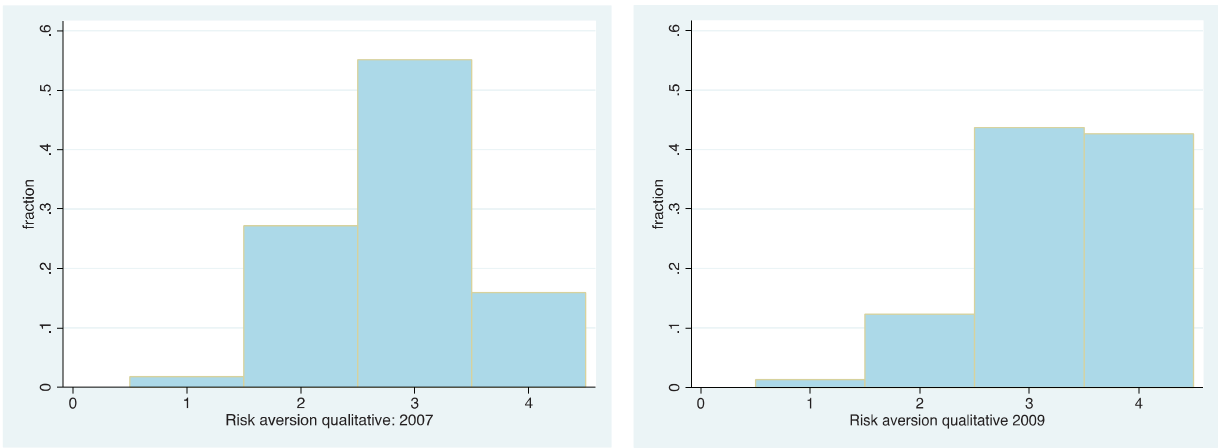

B. Quantitative measure of risk aversion (risk premium)

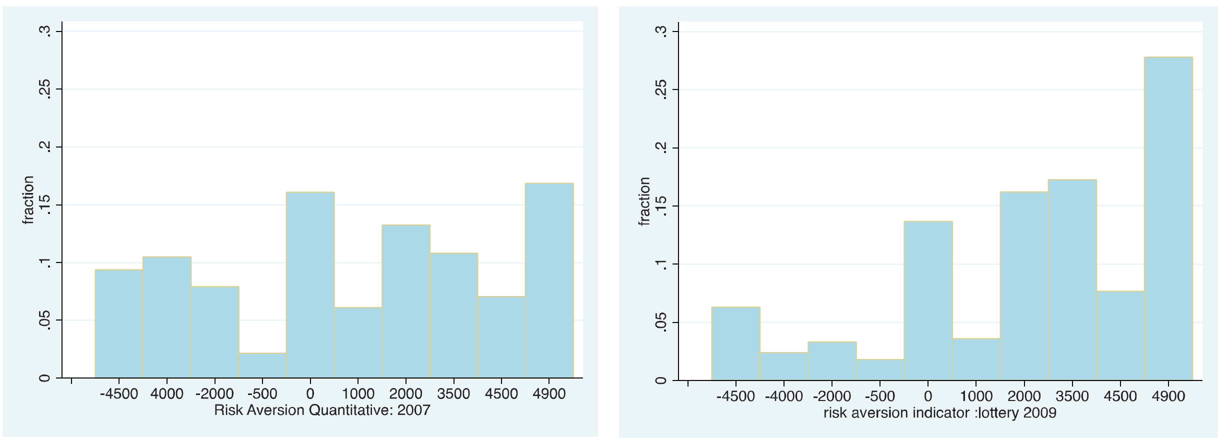

C. Quantitative measure of risk aversion (risk premium) over time

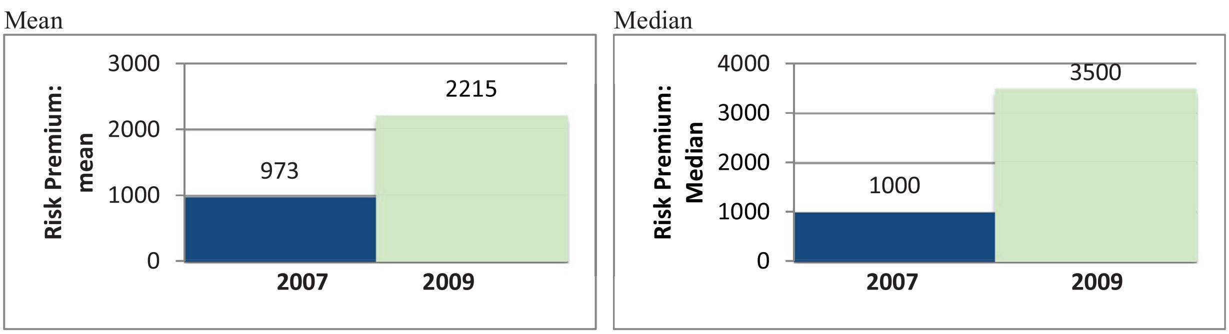

Fig. 1. Frequency distribution of the level of risk aversion indicators in 2007 and 2009. Panel A reports the frequency distribution of the qualitative measure of risk aversion in 2007 and 2009. The qualitative indicator elicits the investment objective of the respondent, offering them the choice among “Very high returns, even at the risk of a high probability of losing part of my principal”; “A good return, but with an ok degree of safety of my principal”; “An ok return, with good degree of safety of my principal”; “Low returns, but no chance of losing my principal.” Responses are coded with integers from 1 and 4, with a higher score indicating a higher aversion to risk. Panel B shows the frequency distribution of the risk aversion indicator based on the answers to the lottery that delivers 10,000 euros or zero with equal probability in 2007 and 2009. The risk premium for this gamble is computed as the difference between the expected value of the lottery (5000 euros) and each respondent’s certainty equivalent. Panel C reports the average and median risk premium for this gamble in the two years.

Table 1
Comparing the sample of non-participants and participants to the second interview.
This table shows summary statistics for a set of variables for the two samples of respondents to the 2007 bank survey: those who did not participate in the 2009 survey and those who did. The third column reports the *p*-value for the null that the means of the two samples are the same. The variables are defined in the Internet Appendix, Section A4.

	Nonparticipants (N. 1020)	Participants (N. 666)	*p*-value of test of equality
Age	55.02	54.5	0.39
Male	0.7	0.7	0.77
Married	0.69	0.67	0.40
North	0.53	0.49	0.12
Center	0.24	0.25	0.61
Education	12.44	13.18	0.00
Trust advisors 2007	3.74	3.83	0.05
Qualitative measure of risk aversion	2.88	2.85	0.31
Quantitative measure of risk aversion: risk premium	865	792	0.66
Financial assets 2007	150,980	158,950	0.22
Financial assets 2009	139,720	142,290	0.73
Stock ownership Jan 2007	0.44	0.44	0.93
Stock ownership Jun 2009	0.41	0.42	0.80
Share in stocks Jan 2007	0.1	0.11	0.54
Share in stocks Jun 2009	0.08	0.08	0.51
Risky assets ownership Jan 2007	0.79	0.81	0.41
Risky assets ownership Jun 2009	0.74	0.73	0.63
Risky assets share Jan 2007	0.56	0.58	0.29
Risky assets share Jun 2009	0.50	0.50	0.90

and a sequence of certain sums of money. These sums were progressively increasing between 100 euros and 9000 euros. More risk-averse people will give up the risky prospect for lower certain sums. Thus, the first certain sum at which an investor switches from the risky to the certain prospect identifies (an upper bound for) his/her certainty equivalent.

Specifically, investors were asked: “Imagine being in a room. To get out you have two doors. Behind one of the two doors there is a 10,000 euro prize, behind the other nothing. Alternatively, you can get out from the service door and win a known amount. If you were offered 100 euros, would you choose the service door?”

If the investor accepted 100 euros, the interviewer moved on to the next question. Otherwise the interviewer asked whether the investor would accept 500 euros to exit the service door and if not, 1500 and if not, progressively 3000, 4000, 5000, 5500, 7000, 9000, and more than 9000 euros.

The question was framed so as to resemble a popular TV game (*Affari Tuoi*, the Italian version of the TV game *Deal or No Deal*), analyzed by Bombardini and Trebbi (2012). Incidentally, it is similar to the Holt and Laury (2002) strategy which has proved particularly successful in overcoming the under/over-report bias implied when asking willingness to pay/accept.

We code answers to this question as the certainty equivalent value required by the investor to give up the risky prospect. We then compute a risk premium as the difference between the expected value of the gamble and an investor’s certainty equivalent. Fig. 1B presents these risk premia. Interestingly, roughly one-third of the investors appear risk loving in 2007. The extreme risk-averse people (with a risk premium equal to 4900) are 17% in 2007.

We will refer to the measure based on preferences for risk-return combinations as the qualitative indicator and to the one based on the lottery as the quantitative indicator. The first is a measure of relative risk aversion, while the second is a measure of absolute risk aversion. These risk-aversion measures should be thought of as measures of the risk aversion for the investor’s value function and as such are potentially affected by any variable that impacts people’s willingness to take risk, such as their wealth level or any background risk they face.

In the 2009 survey we also ask “After the stock market crash did you become more cautious and prudent in your investment decisions?” The possible answers are: “More or less like before,” “A bit more cautious,” or “Much more cautious.” Thirty-five percent of the investors declare themselves to have become “much more” and 18% “a bit more” cautious. Therefore, we create a “change in cautiousness” variable equal to zero if the response is no change, one if the response is a bit more, and two if it is much more. The summary statistics for these measures and all the other variables are contained in Table 2.

2.3. Sample attrition

Of the 1541 who were contacted, roughly one-third agreed to be re-interviewed so that we end up with a two-year panel of 666 investors. Table 1 compares the characteristics of responders and non-respondents to the 2009 survey along several dimensions. In the first part of the table, we compare the two samples according to the demographic characteristics collected in the 2007 survey such as age, gender, marital status, geographical location, and education. The differences are small and not statistically significant, with the exception of education, for which we cannot statistically reject the hypothesis that the two sam-

Table 2
Summary statistics of risk aversion measures, other variables, and controls.
Panel A reports the summary statistics for the risk aversion measures. The qualitative measure of risk aversion elicits the investment objective of the respondent, offering them the choice among "Very high returns, even at the risk of a high probability of losing part of the principal," "A good return, but with an ok degree of safety of the principal," "An ok return, with good degree of safety of the principal," "Low returns, but no chance of losing the principal." The responses are coded with integers from 1 to 4, with a higher score meaning a higher risk aversion. The quantitative measure of risk aversion is calculated by eliciting the certainty equivalent for a gamble that delivers either 10,000 euros or zero with equal probability; the risk premium is then obtained as the difference between the expected value of the gamble (5000 euros) and each respondent's certainty equivalent. Panels B, C, and D report the summary statistics for all the other variables used later in estimates and defined in the Internet Appendix, Section A4.

Panel A: Risk aversion measures in 2007 and 2009

	Quantitative measure (risk premium in euros)			Qualitative measure		
	Mean	Median	Sd	Mean	Median	Sd
Level in 2007	792	1000	3248	2.87	3	0.72
Level in 2009	2215	3500	2815	3.28	3	0.73
Change 2009/2007	1423	2500	3994	0.42	0	0.81
Fraction of people with increase in risk aversion		0.55			0.46	
Fraction of people with unchanged risk aversion		0.18			0.44	
1- Fraction of people with a decrease in risk aversion		0.73			0.90	

Panel B: Other variables: levels

	Mean	Median	Sd
Male	0.70	1	0.46
Age	54.81	57	12.3
Education (years)	12.73	13	4.25
Retired	0.33	0	0.47
Government employee	0.33	0	0.47
Size of the financial loss	−0.05	−0.01	0.07
Log net wealth: 2007	13.11	13.10	0.59
Log net wealth: 2009	13.05	13.03	0.64
Risky asset ownership 2007	0.79	1	0.40
Risky asset share 2007	0.57	0.70	0.37
Knightian uncertainty	0.29	0	0.46
Change in cautiousness	2.13	2	0.90
Trust advisors	3.78	4	0.91
Change trust advisors	−0.24	0	1.11

Panel C: Other variables: first differences

	Mean	Median	Sd
Δ Log net wealth 2009–2007	−0.06	−0.051	0.27
Δ Log net wealth 2009–2008	−0.04	−0.005	0.20
Δ Ownership of risky assets	−0.06	0	0.35
Δ Share risky assets	−0.04	0	0.24
Δ in Expected gross stock return in euros	−432	−168	3442
Δ in Range gross stock market return in euros	−249	−50	7200

Panel D: Variables in the rebalancing equation

	Mean	Median	Sd
Risk aversion ratio: $\left[\frac{\gamma}{\gamma_i'} - 1\right]\omega_i^M$	−0.05	0.0	0.21
Mean risky asset share 2007: ω_i^M	0.52	0.59	0.36
Post-shock share: $\left[\frac{p\omega_i^M}{p\omega_i^M+1-\omega_i^M}\right]$	0.65	0.72	0.28

ples differ. Still, the economic magnitude of the difference is small (less than a year of education).

In the middle part of the table, we compare the two samples according to their risk attitudes, as measured in 2007. Along this dimension, which is the most important one for our analysis, participants in the 2009 survey do not differ from non-participants. For instance, the average 2007 risk premium for the hypothetical risky prospect is 865 euros among non-respondents to the 2009 telephone survey and 792 euros among respondents (p-value = 0.66).

While the two samples do not differ in observable characteristics in 2007, they might differ in time-varying characteristics. For example, the crisis might have affected the two groups differentially, in a way that is correlated with their willingness to be re-interviewed. Fortunately, we have the 2007 and 2009 administrative data (and hence

Table 3
Correlation between the various measures of risk aversion.
The table reports the correlation between the two measures of risk aversion for the two survey waves 2007 and 2009, the correlation between their changes, and the correlations between their changes and a measure of change in cautiousness in investing defined in the Internet Appendix, Section A4.

	Correlations between measure of risk aversion				
	Qualitative and quantitative indicator: 2007	Qualitative and quantitative indicator: 2009	Change in qualitative and change in quantitative indicator: 2007–2009	Change in qualitative indicator and change in cautiousness	Change in quantitative indicator and change in cautiousness
	0.116	0.160	0.118	0.119	0.074
p-value	0.00	0.00	0.002	0.002	0.056

portfolio choices) of both the respondents and the non-respondents. Hence, the last part of Table 1 compares these choices. The stock of financial assets, before and after the crisis, does not differ between the two groups, nor does the fraction of financial wealth invested in stock or risky financial assets. Similarly, there are no differences in the percentage of people who own stock or risky financial assets. We, thus, conclude that there does not seem to be any systematic selection in the investors' decisions to be re-interviewed in June 2009.

2.4. Validating the risk aversion measure

A large and increasing literature shows that questions like the ones above predict risk taking behavior in various domains (see, for instance, Dohmen et al., 2011; Donkers et al., 2001; Barsky et al., 1997; Guiso and Paiella, 2006, 2008). The risk aversion measures elicited in this way are also robust to the specific domain of risk: using a panel of 20,000 German consumers, Dohmen et al. (2011) show that indicators of risk attitudes over different domains tend all to be correlated, with correlation coefficients of around 0.5—a feature that is consistent with the idea that risk aversion is a personal trait.

To validate our measures, we run various tests. First, in Table 3 we document that our qualitative and quantitative measures are positively correlated either when using the 2007 cross section (correlation coefficient 0.12) or the 2009 cross section (correlation 0.16) or when looking at the correlation between the changes in the two measures between 2007 and 2009 (correlation coefficient 0.12). We also find that the change in cautiousness variable has a 12% correlation (*p*-value 0.002) with the changes in the qualitative measure of risk aversion and a 7.4% correlation (*p*-value 0.056) with changes in the quantitative measure of risk aversion.

Second, we document that our measures tend to be correlated in expected ways with classical covariates of risk attitudes.[3] As panel A of Table 4 shows, risk aversion decreases with total wealth levels in both the 2007 and the 2009 cross sections. Also, as documented in the literature, men are less risk averse than women (Byrnes et al., 1999).

Third, we document that our measures have predictive power on investors' financial choices. Panel B of Table 4 shows that the qualitative indicator of risk aversion is strongly negatively correlated with ownership of risky financial assets (a dummy variable equal to one if an individual owns stocks and corporate bonds in her portfolio). The correlation with the lottery-based measure is negative but weaker. This is partly due to some investors providing noisy answers in the quantitative measure, which is more difficult to understand. When we drop inconsistent answers—those who are highly risk averse according to the first indicator (a value greater than 2), but highly risk lovers on the basis of the lottery question (a risk premium less or equal to −4000 euros)—we also find that the quantitative measure significantly predicts risky asset ownership (column 3). Furthermore, the change in risk aversion predicts the change in assets ownership: those whose risk aversion increased more between 2007 and 2009 are more likely to become non-stockholders over the same period (Table 4, panel C). In the Internet Appendix (Table A.2, panels A and B), we also document a similar pattern for the level and in the change in the share of wealth held in risky assets.

2.5. Measure of subjective probabilities

Depositors were also asked to state what (in their view) the minimum and maximum value of a 10,000 euro investment in a fully diversified stock mutual fund would be 12 months later. Next, they were asked to report the probability that the value of the stock by the end of the 12 months was above the mid-point of the reported support. Under very simple assumptions about the shape of the distribution, this parsimonious information allows computing the subjective mean and variance of stock market returns. We have computed these moments assuming the distribution is uniform, but results are the same assuming it is triangular.

Some investors claim that the *maximum* one-year-ahead value of a 10,000-euro investment in the stock market is zero. This is a sign that they might have misunderstood the question, raising some doubt on the quality of this measure. To address these concerns, we drop all the responses where the maximum is below 3000 euros (i.e., maximum

[3] These patterns of correlations have been documented in several studies, either using surveys or experiments [e.g., Croson and Gneezy, (2009) for gender; Barsky et al. (1997), Guiso and Paiella (2006, 2008), Hartog et al. (2002)].

Table 4
Validation.

Panel A, columns 1 and 2 report estimates of an ordered probit model where the dependent variable is the qualitative measure of risk aversion for the two different survey waves, 2007 and 2009. Columns 3 and 4 report interval regressions; the dependent variable is the interval of the risk premium obtained from the lottery question as described previously, divided by the expected value of the lottery.

Panel B reports marginal effects of probit models, where the dependent variable is a dummy variable equal to one if the individual holds risky assets in her portfolio. The quantitative indicator of risk aversion is the risk premium defined previously divided by 5000 (the expected value of the lottery) in 2007. The last column reports the results eliminating individuals who reported inconsistent answers to the risk aversion question (those who are highly risk averse according to the first measure—a value greater than 2—but risk lover on the basis of the quantitative question—a certainty equivalent greater or equal to 9000 euros).

Panel C reports the marginal effects for ordered probit regressions; the dependent variable is the change in a dummy variable equal to one if an individual owns risky assets between June 2009 and June 2008 (just before the financial collapse). The change in risk aversion is calculated as the difference between the reported answers in the 2009 and 2007 surveys. The change in the quantitative indicator of risk aversion is the change in risk premium divided by 5000 (the expected value of the lottery). All the other variables are defined in the Internet Appendix, Section A4. Robust standard errors are in parentheses. */**/*** indicates statistical significance at the 10%, 5%, and 1% level, respectively. In Panel C, changes in wealth have been trimmed out at the 1st and 99th percentiles.

Panel A: Cross-sectional correlates of risk aversion

	Risk aversion qualitative		Risk aversion quantitative			
			Whole sample		Eliminate inconsistent answers	
	2007	2009	2007	2009	2007	2009
Male	−0.338***	−0.497***	0.006	0.150**	−0.031	0.011
	(0.063)	(0.109)	(0.046)	(0.075)	(0.041)	(0.059)
Age	−0.047**	−0.011	−0.001	0.026	0.010	0.031*
	(0.020)	(0.032)	(0.014)	(0.021)	(0.013)	(0.017)
Age2/100	0.049***	0.020	0.008	−0.015	−0.004	−0.022
	(0.019)	(0.031)	(0.014)	(0.021)	(0.012)	(0.017)
Education	−0.035***	−0.044***	−0.013**	−0.010	−0.021***	−0.016***
	(0.007)	(0.012)	(0.005)	(0.008)	(0.005)	(0.006)
Log net wealth: 2007	−0.139***		−0.057		−0.042	
	(0.047)		(0.036)		(0.033)	
Log net wealth: 2009		−0.147**		0.013		0.001
		(0.074)		(0.050)		(0.040)
Observations	1494	584	1494	584	1311	548

Panel B: Risk aversion and risky assets ownership

	Whole sample		Eliminate inconsistent answers
	(1)	(2)	(3)
Risk aversion qualitative: 2007	−0.122***		
	(0.032)		
Risk aversion quantitative: 2007		−0.001	−0.055***
		(0.005)	(0.004)
Male	0.129***	0.154***	0.162***
	(0.016)	(0.028)	(0.027)
Age	0.022**	0.025***	0.026***
	(0.010)	(0.008)	(0.009)
Age2/100	−0.020**	−0.023***	−0.023***
	(0.009)	(0.007)	(0.008)
Education	0.018***	0.020***	0.019***
	(0.005)	(0.002)	(0.002)
Trust advisor 2007	0.039***	0.047***	0.049***
	(0.011)	(0.007)	(0.008)
Log net wealth: 2007	0.145***	0.152***	0.137***
Observations	1494	1494	1311

Panel C: Effect of changes in risk aversion on changes in ownership of risky assets

	Whole sample		Eliminate inconsistent answers
	(1)	(2)	(3)
Δ Risk aversion: Qualitative measure	−0.172		
	(0.105)		
Δ Risk aversion: Quantitative measure		−0.154*	−0.243**
		(0.088)	(0.100)
Male	0.367**	0.379**	0.308
	(0.172)	(0.171)	(0.189)
Age	0.074	0.071	0.070
	(0.062)	(0.062)	(0.070)
Age2/100	−0.069	−0.066	−0.071
	(0.058)	(0.058)	(0.065)
Education	0.006	0.008	0.011
	(0.019)	(0.020)	(0.020)
Δ in Advisor trust	−0.065	−0.082	−0.088
	(0.072)	(0.072)	(0.082)
Δ Log net wealth 2009–2007	1.467***	1.351***	1.214***
	(0.371)	(0.366)	(0.462)
Observations	569	569	500

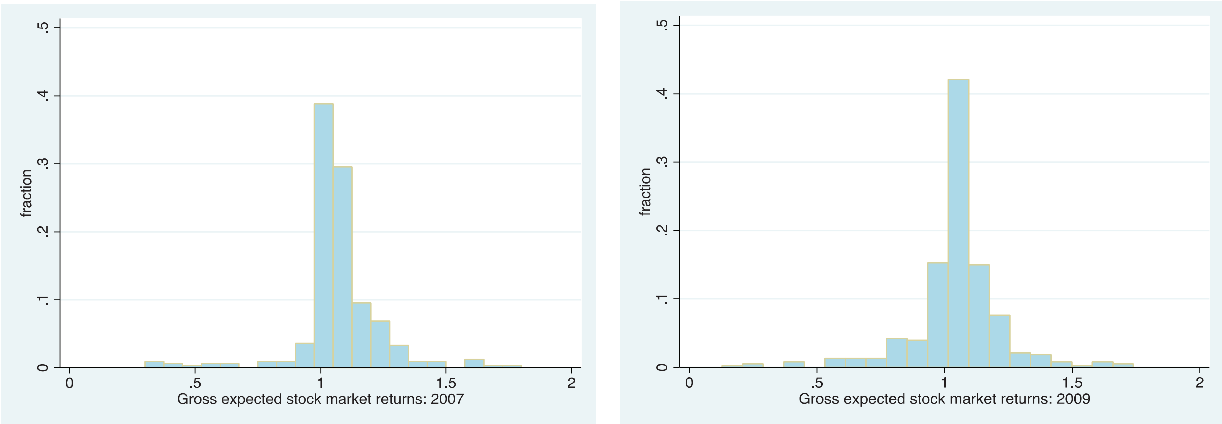

Fig. 2. Distribution of expected gross stock returns. The figure shows the cross-sectional distribution of one-year ahead subjective expected stock returns in 2007 and 2009. Expected returns are obtained from the answers to a question asking the minimum and maximum value of an investment of 10,000 euros in a fund representative of the Italian stock market one year later and the probability that the value is below the mid-point of this range. We drop the observations where the respondents claimed that the maximum one-year-ahead value of a 10,000 euros investment is 3000 euros or less. The reported distributions are for the respondents to both the 2007 and 2009 surveys.

net return of −70% or below). In so doing, the number of observations drops from 470 to 337. The resulting distribution of individual expected returns is presented in Fig. 2.

2.6. Changes in wealth and financial losses

For all the participants in the survey, we have access to the administrative data, which include the amount of deposits at the bank, the amount and composition (by broad categories) of their brokerage account at the bank, the proportion of financial wealth represented by their holdings at the bank, and the value of their house. Thanks to these data we can infer the changes in investors' total wealth and the losses incurred on their financial portfolio.

The change in total wealth is computed as the sum of the actual changes in their financial wealth held at the bank (divided by the proportion of financial wealth held at the bank to obtain an estimate of total household assets) and the imputed changes in home equity. To impute these changes we look at the variation in local indexes of real estate prices.

The losses on the financial portfolio are computed by multiplying the holdings of risky securities (stocks, stock mutual funds, corporate bonds, and corporate bond funds) before August 2008 by the proportional change in their price between September 2008 (before the Lehman collapse) and February 2009 (when the stock market started to rebound) and then scaling by the stock of financial assets before August 2008.

3. Changes in risk aversion

3.1. Changes in individual risk aversion

Fig. 1A compares the distribution of the qualitative measure of risk aversion before and after the crisis. Before the crisis the average response was 2.87, after the crisis it has jumped to 3.28 (recall, a higher number indicates higher risk aversion). This change is statistically different from zero at the 1% level. In 2007, only 16% of the respondents chose the most conservative option "low return and no risk," in 2009, 43% did. In the Internet Appendix (Table A.3) we show the transition matrix of the responses. There is a homogeneous shift toward more conservative combinations of risk and return. Of the people who chose the most aggressive option ("Very high returns, even at the risk of a high probability of losing part of the principal"), 83% change toward a more conservative one. Of those who had chosen the second more risky combination ("Moderate Risk/Medium Return"), 77% move to more conservative options, while only 2% move to the most aggressive one. Of those who chose "Small Risk/Some Return," 44% move to "Low Return and No Risk," while only 9.5% move to more aggressive options. Note that these very stark results are present in spite of a censoring in the data. The 16% of the respondents who chose the most conservative option in 2007 cannot become more risk averse.[4]

Fig. 1B compares the distribution of the risk premium before and after the crisis and Fig. 1C the mean and median risk premium before and after the crisis (the transition matrix is in table A.4 of the Internet Appendix). As Fig. 1C shows, before the crisis the average risk that premium investors were willing to pay to avoid a gamble offering 10,000 euros and zero with equal probability was 973 euros. In 2009, the risk premium for the same group of people increased to 2215 euros. The median increased from 1000 to 3500. All these changes are statistically different from zero. Interestingly, the large surge in the risk premium is driven by a much higher number of people who choose the lowest certainty equivalent (and thus the highest risk premium).

Since the risk premium is proportional to the investor risk aversion, these estimates imply that the (absolute) risk aversion of the average investor has increased by a factor of two, and that of the median investor by a factor of 3.5!

[4] The effect of this censoring is considered in the Data Appendix, Fig. A1.

Table 5
Net sales of risky assets.
The table reports the average monthly net-purchases (purchases minus sales) of risky assets (as proportion of the assets held at the beginning of the period) before the Lehman collapse (June 2008–August 2008, column 1) and after (October 2008–April 2009, column 2). The third column reports the difference in net sales in the post-and pre-Lehman collapse (*p*-values in parentheses). Row A includes all the households who responded to both surveys and own some risky assets at the beginning of the period. Rows B and C refer to the subsample of Row A, that respectively did not exhibit an increase in the qualitative measure of risk aversion and those who did. Similarly, Rows D and E report the net purchase of individuals who respectively did not and did exhibit an increase in the quantitative measure of risk aversion.

	Pre-Lehman Jun-Aug 08 (1)	Post-Lehman Oct 08-Apr 09 (2)	Difference Post – Pre Lehman (*p*-values) (3)
A. Subsample who own risky assets at the beginning of period	−0.022	−0.038	−0.06
			(0.224)
B. Subsample of A without an increase in qualitative risk aversion	−0.023	−0.021	0.002
			(0.922)
C. Subsample of A with an increase in qualitative risk aversion	−0.020	−0.052	0.032
			(0.057)
p-values of the difference between B and C	(0.596)	(0.089)	–
D. Subsample of A without an increase in quantitative risk aversion	−0.012	−0.020	−0.008
			(0.616)
E. Subsample of A with an increase in quantitative risk aversion	−0.029	−0.046	−0.017
			(0.250)
p-values of the difference between D and E	(0.594)	(0.424)	–

One benign reason why risk aversion might have increased is that from the first to the second survey our investors became older. While true, this effect is likely to be small, since only two years went by. Nevertheless, we computed the average risk aversion by age and then took the difference of risk aversion between the first and the second surveys, keeping the age constant (i.e., between the average of people who were 30 in 2009 and the people who were 30 in 2007). The results are unchanged.

Such an increase cannot be attributed solely to a worsening of expectations about the distribution of future investments since it manifests itself also in the quantitative measure, which is unrelated to the stock market. In fact, the probability distribution underlying the gamble in the quantitative measure is objective, not subjective.

3.2. Changes in portfolios

As Table 2.C shows, the share of risky assets in individual portfolios dropped between 2007 and 2009. This drop could be the mechanical effect of a decrease in the value of risky assets held in the portfolio or the consequence of an active sale of risky assets by individuals (or both). In the current section we focus on the second component: the active sale of risky assets. In a standard (Mertonian) model of portfolio choice with constant risk aversion, expected value, and volatility, individuals should buy more risky assets after a drop in their value. We showed that individuals do not exhibit a change in the expected return or volatility after the crisis. Thus, at least the individuals who expressed opinions on the return distribution should not be selling risky assets after the shock more intensively than the whole sample (indeed the data show that if anything they seem to sell less).

In Table 5 we report the average monthly net-purchases (purchases minus sales) of risky assets (as proportion of assets held at the beginning of the period) before the Lehman collapse (June to August 2008) and after it (October 2008 to April 2009). The sample in row A includes all the households who responded to both surveys and own some risky assets at the beginning of the period (so that they can sell them). In the 2008 months leading to the Lehman collapse, households were net sellers of risky assets (2.2% average net sales). It is important to notice that the level of net purchases is relatively small, because most households at any given time are inactive. After Lehman, the net sales almost doubled (3.8%), but this difference is not statistically significant (column 3 of the table).

In rows B and C we report the net purchase before and after Lehman of the individuals who did not exhibit an increase in the qualitative measure of risk aversion and of those who did. While the net purchases are almost identical before (see the first column), they are very different after Lehman (column 2). People with no increase in risk aversion sell 2.1% of their assets, while individuals with an increase sell 5.2%. This level of net purchases is statistically different from the one exhibited by the same people before Lehman (*p*-value of 0.057) and that of the investors who did not exhibit an increase in qualitative risk aversion (*p*-value of 0.089).

In rows D and E we report the net purchase before and after Lehman of the individuals who did not exhibit an increase in the *quantitative* measure of risk aversion and of those who did, respectively. While the pattern is similar to the one observed in rows B and C, the differences here are not statistically significant at the conventional levels.

3.3. A reality check on the magnitude of the changes

Our sample is representative of Italian individual investors, but not of all investors: institutional investors and professional traders are not represented. Yet, if we treat it

Table 6
Aggregate risk aversion.
This table computes measures of the aggregate absolute risk aversion for the two sample years 2007 and 2009 using as weights the total net wealth of the individuals in the sample. The absolute risk aversion measure is obtained from the elicited risk premium assuming an exponential utility function.

	Total sample *wealth weights:*		Stockholders *wealth weights:*	
	2007	2009	2007	2009
ARA 2007	1.30	1.30	1.22	1.22
ARA 2009	2.28	2.25	2.42	2.40
Change in ARA	0.98	0.95	1.20	0.98

as a representative sample, we can compute the aggregate risk aversion and check whether the change in the aggregate risk aversion is large enough to explain the large drop in stock prices.

To compute the aggregate risk aversion we start by mapping the risk premium computed from the quantitative question into a coefficient of absolute risk aversion by using a constant relative risk aversion (CRRA) utility function. Then, we compute the aggregate risk aversion by weighting these coefficients by the net total wealth of each individual. As Table 6 shows, the aggregate absolute risk aversion (ARA) in 2007 was 1.3. If we maintain the individual risk aversions estimated in 2007 and multiply them by the 2009 wealth weights, the aggregate risk aversion does not change at all. By contrast, if we use the 2009 estimated individual risk aversion, the aggregate risk aversion almost doubles. If we repeat the analysis restricting the sample to people who were stockholders in 2007, the results are the same.

Now that we have computed the variation in aggregate risk aversion, we can estimate whether this change is sufficiently large to justify the severe drop in stock prices that took place. What is relevant for asset prices is the relative risk aversion. Since the change in total wealth is small (a relatively small fraction was invested in equity), all the increase in absolute risk aversion translates into an increase in the relative risk aversion. To compute how this increase could affect stock prices we make the (strong) assumption that the only source of variation was a (temporary) increase in risk aversion. This implies that the future expected cashflow remains unchanged and that after one year even the risk aversion returns to normal. Then, next year stock price P_1 should remain unchanged and all the adjustment should take place in today's stock priceP_0. By using the Merton (1969) model, we can write

$$r^e = \frac{P_1}{P_0} - r_f = \gamma\sigma^2 \tag{1}$$

where the left-hand side is the equity premium and $\sigma^2$the variance of stock returns. If the expected variance of returns does not change and the risk aversion γ doubles to $\gamma' = 2\gamma$, as it does in our sample, then the initial stock price P_0' should be

$$\frac{P_0'}{P_0} = \frac{1}{\frac{\gamma'}{\gamma}\left[1 - \frac{P_0}{P1}r_f\right] + \frac{P_0}{P_1}r_f} \tag{2}$$

which is roughly half of what it was before. Hence, stock price roughly halves if risk aversion doubles. Thus, the sharp increase in risk aversion is quantitatively sufficient to explain the severe drop in stock prices during the crisis. Yet, it begs the question of what caused such an increase in risk aversion.

4. What causes the changes in risk aversion?

4.1. Changes in wealth

A characteristic that standard expected utility models have in common with the non-standard ones (habit formation and prospect theory) is that any change in risk aversion is mediated by changes in wealth. For this reason, we start by analyzing whether there is any relation between changes in risk aversion and changes in wealth.

Fig. 3 plots a nonparametric estimation of the relationship between changes in risk aversion and the size of financial losses incurred between September 2008 and February 2009 (if we use total wealth the results are the same). As Fig. 3A shows, there is no consistent relationship between the increase in the *qualitative* measure of risk aversion and the size of the losses in the financial portfolio during the financial crisis. For losses between zero and 20%, the increase in risk aversion is stable at 0.4, around 14% of the sample mean in 2007. For losses above 20%, the increase in risk aversion seems to first decrease and then increase.

As Fig. 3B shows, for the *quantitative* measure there seems to be a negative relation between the size of the financial loss and the relative risk premium (the risk premium divided by the expected value of the lottery), consistent with a wealth channel. Yet, even people with no losses exhibit a significant increase in the relative risk premium (by 20 percentage points), which seems to contradict the wealth channel.

In Table 7 we revisit this issue in regression format, which allows us to control for individual characteristics. The dependent variable in panel A is the change in the qualitative risk aversion between 2007 and 2009. As control variables in column 1 we use the initial level of risk aversion, gender, two dummy variables for the age groups, and education. Our explanatory variable of interest is the size of the financial loss, calculated as the loss in value of the risky investments between September 2008 and February 2009, scaled by the value of financial assets held in September 2008. As in the figure, we find no evidence of correlation between this variable and the increase in risk aversion. In column 2 we re-estimate the same specifica-

A. Qualitative measure of risk aversion

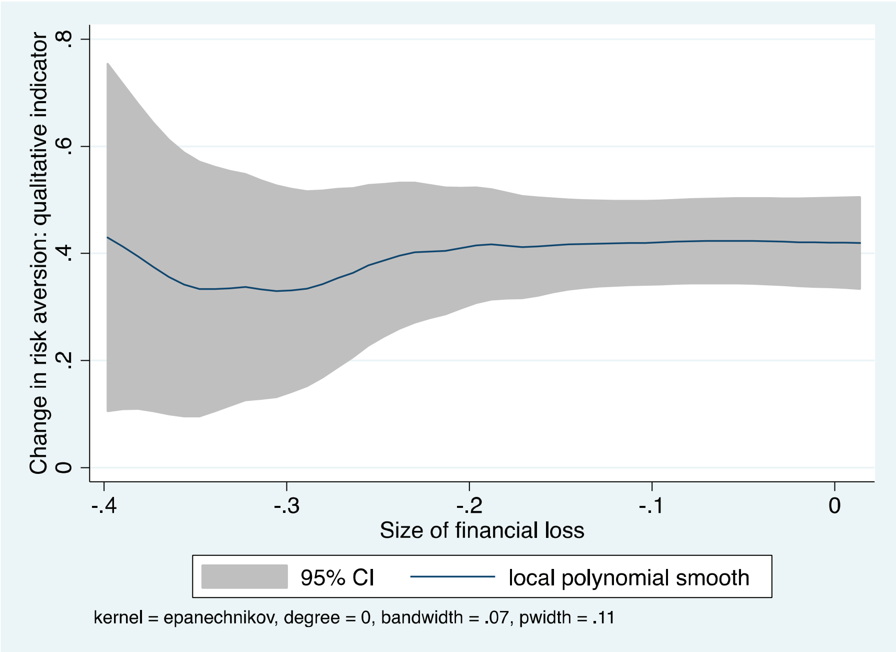

B. Quantitative measure of risk aversion (risk premium)

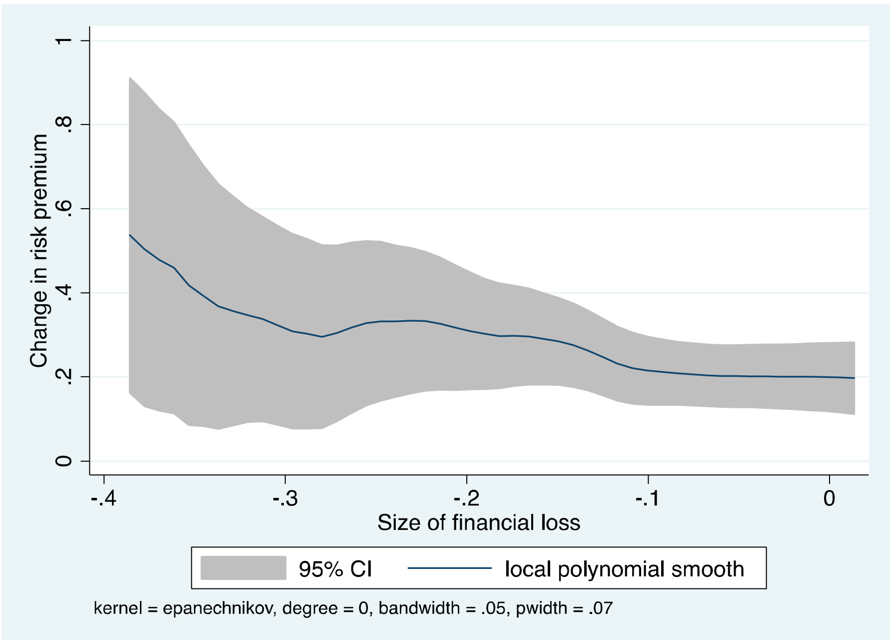

Fig. 3. Financial loss and change in risk aversion. The figure plots the relation between potential losses in the financial portfolio between September 2008 and February 2009 and the change in the qualitative indicator of risk aversion (Panel A) and in the risky premium of the quantitative lottery (Panel B). The change in risk premium is scaled by the expected value of the lottery (5000 euros). The figures show the 95% confidence interval around the estimated polynomial. The relation is estimated using a kernel-weighted local polynomial regression. The financial loss is computed as loss in value of risky investments held at the end of September 2008 between September 2008 and February 2009, scaled by the initial value of financial assets.

Table 7
Change in risk aversion, financial loss, and Knightian uncertainty.
Panel A reports ordered probit estimates for the first difference of the qualitative measure of risk aversion. Panel B reports interval regressions estimates for the changes in the risk premium scaled by the expected value of the lottery. All the other variables are defined in the Internet Appendix, Section A4. Robust standard errors are in parentheses. */**/*** indicates statistical significance at the 10%, 5%, and 1% levels, respectively. Size of financial loss is computed as loss in value of risky investments held at the end of September 2008 between September 2008 and February 2009, scaled by the initial value of financial assets. Change in wealth is computed as the log difference in the value of net wealth in the second quarter of 2009 and the value in the second quarter of 2008. Changes in net wealth have been trimmed out at the 1st and 99th percentiles.

Panel A

	(1)	(2)	(3)
Risk aversion qualitative: 2007	−1.165***	−1.172***	−1.178***
	(0.083)	(0.083)	(0.083)
Male	−0.422***	−0.423***	−0.350***
	(0.103)	(0.102)	(0.104)
Age ≤ 45	−0.695*	−0.685*	−0.604
	(0.366)	(0.372)	(0.375)
Age ≥ 65	−0.553	−0.545	−0.501
	(0.372)	(0.379)	(0.381)
Education	−0.041***	−0.042***	−0.037***
	(0.012)	(0.012)	(0.012)
Size of financial loss	0.293		
	(0.583)		
Change log net wealth 2009–Q2 2008		0.513	0.503
		(0.399)	(0.413)
Knightian uncertainty			0.441***
			(0.105)
Observations	572	572	572

Panel B

	(1)	(2)	(3)
Risk aversion quantitative: 2007	−0.860***	−0.860***	−0.860***
	(0.043)	(0.043)	(0.043)
Male	0.126**	0.126**	0.131**
	(0.057)	(0.057)	(0.058)
Age ≤ 45	−0.060	−0.069	−0.063
	(0.170)	(0.171)	(0.174)
Age ≥ 65	0.071	0.071	0.075
	(0.175)	(0.176)	(0.178)
Education	−0.003	−0.003	−0.003
	(0.005)	(0.005)	(0.005)
Size of financial loss	−0.094		
	(0.326)		
Change log net wealth 2009–Q2 2008		0.014	0.013
		(0.232)	(0.233)
Knightian uncertainty			0.029
			(0.052)
Observations	503	503	503

tion with changes in wealth in place of size of the financial loss. Also in this case, there is no effect. In Table A.5 we show the robustness of these results to use a quadratic form to model the effect of age.

In panel B we repeat the same exercise for the quantitative measure of risk aversion. As in the figure, we find no relation between the increase in risk aversion and either the size of the financial loss or the change in wealth.

To better understand the robustness of these results we focus on the 295 people who did *not* experience any financial loss between September 2008 and February 2009 (either because they gained or because they did not have any risky assets and thus did not experience any loss).[5]

[5] The people with inconsistent answers are 59. We focus on financial losses; households could have suffered losses on housing wealth. This is not the case: between the second quarter of 2008 and the first quarter of 2009, house prices increased in all local markets but one where they dropped by 3.7%. A nonparametric estimate of the relation between the change in the qualitative and quantitative measures of risk aversion and the proportional change in total wealth leads to a similar conclusion. The certainty equivalent increases by a similar amount even for households whose total wealth holdings do not change or even increase.

Before the crisis the *qualitative* measure of risk aversion of this subsample was 2.82 (statistically not different from the rest of the sample). After the crisis their qualitative measure jumped to 3.25, again not statistically different from the jump in the rest of the sample.

The same is true for the *quantitative* measure. Before the crisis, the mean risk premium for such subsample was 649 euros, not statistically different from the rest of the sample. After the crisis, the mean risk premium rises to 2260 euros, not statistically different from the rest of the sample. Therefore, investors who did not experience any loss exhibit an increase in risk aversion equal to those who did.

Our measure of financial losses is based on the wealth investors have deposited at the bank. We cannot exclude, thus, that they might have faced a loss in other investments. To check this possibility we restrict attention to individuals who declare that they only have financial wealth at the bank (184 observations). The results are very similar. The qualitative measure of the risk aversion increases from 2.89 before the crisis to 3.28 after the crisis—a change not different from that in the whole sample. The risk premium on the quantitative measure increases from 757 euros to 2382 euros, again a change similar as in the whole sample.

4.2. Changes in expected future income

What we observe is the value function risk-aversion. Thus, it can increase not only for a drop in wealth but also for an increase in the variability or a drop in the level of future income. Since the income from financial assets is generally small relative to labor income, the main suspect is the expected labor income.

With field data it is hard to test this channel, since the expected income depends upon many unobservable variables. To gain some insights on the plausibility of this hypothesis we focus on people who face very little (possibly no) labor income risk, such as government employees. Note that our second survey (June 2009) predates the Greek (October 2009) and euro crises (2010–12), and the Italian government solvency was not seriously in doubt (at that time the spread between the 5-year Italian bond and the German one was around 60 basis points).

As Table 8 shows, among people who did not experience any financial loss, government employees exhibit a surge in the qualitative measure of risk aversion higher (albeit not statistically significant) than non-government employees, while they exhibit an increase in the quantitative measure similar to that of non-government employees.

We repeat the same analyses dividing the sample of investors who did not face financial losses between retirees and non-retirees. In Italy retirees enjoy a defined benefit plan backed by government guarantee. Thus, the same considerations above apply. We find that retired people have statistically the same increase in the quantitative and qualitative measures of risk aversion as the non-retired; if anything, the magnitude is bigger for the retired, contrary to the background risk hypothesis.

If risk aversion increases because of a change in the future expected income, it should increase much more for younger people (who have most of their wealth in human capital) than for older people. Table 8, panel B compares the changes in the qualitative and quantitative measures for young (age below 45) and old (age above 65) people who did not suffer any financial loss. The change in risk aversion is not statistically different between the two groups.

Thus, we find no evidence consistent with a change in labor income or other changes in background risk being the proximate factor that leads to a surge in risk aversion. Yet, the only way to completely rule out this hypothesis is to conduct a lab experiment, where the background risk is perfectly controlled for.

4.3. Changes in probability distribution

Bordalo et al. (2012) develop a theory according to which individuals overweight the probability of salient payoffs. The collapse of Lehman and the fall in stock prices that ensued might have increased the salience of the negative payoffs, increasing their subjective probability. Given these probabilities, investors optimally behave in a more risk-averse manner, i.e., as we observe in our survey and our portfolio data.

As described in Section 2.2, our survey data allow us to calculate the subjective probability investors have about future returns. Despite the limitations of this measure, it is interesting to analyze the changes in the cross-sectional distribution of expected returns between 2007 and 2009 (see Fig. 2). The 2007 mean (median) gross return is 1.073 (1.057), and the 2009 one is 1.045 (1.042). In the Kolmogorov-Smirnov test of the equality of the 2007 and 2009 distributions of returns, the distance parameter is 0.1365, with a p-value of 0.003. Thus, we cannot reject the hypothesis that the two distributions are different. In particular, Fig. 2 shows an increase in the weight of the distribution in the negative net returns domain. This shift is consistent with the salience hypothesis. However, if this effect was the primary reason for the observed increase in risk aversion in our sample, we would expect a correlation between the change in risk aversion and the change in expected returns. When we do so, the correlation is not statistically significant (the result is also confirmed in an unreported regression).

Fig. 2 hides an important fact. That is, 27% of the households who were willing to give an answer to the distributional questions in 2007 refused to do so in 2009. This change in behavior might reflect an increase in Knightian uncertainty, which might lead investors to behave in a more risk-averse way (Caballero and Krishnamurthy, 2008).

To test this hypothesis, in the third columns of Table 7 (panels A and B) we insert an indicator variable equal to one if a household did not answer this question in 2009. When as a left-hand side variable we use our qualitative measure of risk aversion, this indicator variable has a positive coefficient, which is statistically different from zero at the 1% level. Depositors who did not answer the distributional questions in 2009 have a risk aversion twice as big as the mean. By contrast, when as a left-hand side variable we use the quantitative measure of risk aversion, we do not find any effect.

Table 8
Background risk and expectations about future income.
Panel A computes changes in the qualitative and quantitative measures of risk aversion for groups non-exposed (government-employed and retired) and exposed (non-government employees and non-retired) to background labor income risk among investors facing no financial losses. The third and sixth columns compute the difference between the two groups and the p-value of the null hypothesis that the difference is zero (in parentheses). Panel B does a similar calculation for the groups whose permanent income should respond more (the younger) or less (the older).

Panel A: Background risk

	No financial losses					
	Government employees (Nobs = 96)	Non-government employees (Nobs = 180)	Difference (p-value)	Retired (Nobs = 94)	Non-retired (Nobs = 201)	Difference (p-value)
Change in qualitative measure	0.479	0.399	0.085 (0.420)	0.468	0.423	0.045 (0.661)
Change in quantitative measure (risk premium/5000)	0.316	0.318	−0.002 (0.989)	0.423	0.275	0.149 (1.132)

Panel B: Expected future income: Sample no losses

	No financial losses		
	Age <= 45 (Nobs = 70)	Age >= 65 (Nobs = 48)	Difference (p-value)
Change in qualitative measure	0.358	0.438	−0.081 (0.585)
Change in quantitative measure (risk premium/5000)	0.189	0.388	−0.199 (0.134)

4.4. Changes in utility

The last possible channel is a change in utility. Most economists are reluctant to accept as explanations changes in the utility function because—without any specific theory for the changes—we lack testable restrictions. We need a theory of why and how utility might change after a negative shock. Loewenstein et al. (2001) recognize that emotions could affect decisions. This is tantamount to a state-contingent increase in the curvature of the utility function. We have already seen in Section 3.2 that investors increased their sales of risky assets after Lehman. In what follows we test whether the observed changes in risk aversion can explain the financial decisions to rebalance the portfolio of risky assets.

Let's assume that before the crisis individual portfolios were at the optimal Mertonian share $\omega_i^M = \frac{r^e}{\gamma_i \sigma^2}$. This assumption is realistic given that before the crisis stock prices were fairly flat for a while and thus investors did have the time to adjust. We denote with p the value of stocks after the shock relative to their value before, $p < 1$. Then, after the severe market downturn the actual share of risky assets became

$$\omega_i = \frac{p\omega_i^M}{p\omega_i^M + 1 - \omega_i^M} \tag{3}$$

We have seen that after the crisis the distribution of expected returns did not change much. If the individual risk aversion did not change either, the portfolio rebalancing of individual i would be given by

$$R_i = \omega_i^M - \frac{p\omega_i^M}{p\omega_i^M + 1 - \omega_i^M} \tag{4}$$

If the risk aversion moves to γ_i, then the portfolio rebalancing of individual i is

$$R_i = \frac{\gamma}{\gamma_i'}\omega_i^M - \frac{p\omega_i^M}{p\omega_i^M + 1 - \omega_i^M} \tag{5}$$

We can nest these two specifications as

$$R_i = \alpha\left[\frac{\gamma}{\gamma_i'} - 1\right]\omega_i^M + \beta\omega_i^M + \delta\left[\frac{p\omega_i^M}{p\omega_i^M + 1 - \omega_i^M}\right] \tag{6}$$

where if $\alpha = 0$, $\beta = 1$, $\delta = -1$ we obtain (4), i.e., the optimal rebalancing under the standard Mertonian model with no changes in risk aversion, while if$\alpha = 1$,$\beta = 0$,$\delta = -1$ we obtain (5), i.e., the optimal rebalancing when the risk aversion parameter changes.

To test which expression fits the data best, we build empirical counterparts of the terms on the right-hand side of (6); the details are reported in the Internet Appendix, Section A.4. We define the shock as the drop in stock prices that occurs after August 2008, i.e., the pre-Lehman month. Since prices continue to fall until February 2009, we define various measures of the drop in risky asset prices since August 2008, computed at different months from September 2008 until February 2009. Importantly, we construct an investor-specific measure of p by taking portfolio-weighted means of the drop in different components of the risky portfolio, using as weights the risky

Table 9
Portfolio rebalance.
Panel A reports regressions where the dependent variable is the flow of risky assets bought (positive) or sold (negative) in the three months following the period specified in each column scaled by the value of total financial assets at the end of August 2008, prior to the Lehman Brothers collapse. The risk aversion ratio is the ratio between the risk aversion before and after the Lehman collapse multiplied by the average share of risky assets in 2007; the mean risky asset share is computed as the average share in risky assets over the months from January to December 2007. The post-shock share is the risky share implied by the level of the risky asset prices at the various dates specified in the columns after the collapse of Lehman. All the other variables are defined in the Internet Appendix, Section A4. Robust standard errors are in parentheses. */**/*** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. Change in wealth is computed as the log difference in the value of net wealth in the second quarter of 2009 and the value in the second quarter of 2008. Changes in wealth have been trimmed out at the 1st and 99th percentile.

	(1) September 08	(2) October 08	(3) November 08	(4) December 08	(5) January 09	(6) February 09
Risk aversion ratio: $\left[\frac{\gamma}{\gamma_i'} - 1\right]\omega_i^M$	0.048**	0.043*	0.036	0.077**	0.088**	0.083*
	(0.023)	(0.026)	(0.027)	(0.036)	(0.038)	(0.044)
Mean risky asset share 2007: ω_i^M	0.005	0.015	−0.001	0.001	0.006	−0.000
	(0.012)	(0.015)	(0.017)	(0.020)	(0.022)	(0.024)
Post-shock share: $\left[\frac{p\omega_i^M}{p\omega_i^M+1-\omega_i^M}\right]$	−0.029**	−0.028	−0.016	−0.030	−0.044*	−0.033
	(0.013)	(0.017)	(0.019)	(0.021)	(0.024)	(0.025)
Male	−0.005	−0.020	−0.026*	−0.028*	−0.034**	−0.031*
	(0.010)	(0.013)	(0.014)	(0.015)	(0.016)	(0.017)
Age	0.007**	−0.001	0.002	−0.000	−0.001	−0.003
	(0.003)	(0.006)	(0.006)	(0.006)	(0.006)	(0.007)
Age2	−0.007**	0.001	−0.002	0.000	0.001	0.003
	(0.003)	(0.005)	(0.006)	(0.006)	(0.006)	(0.007)
Education	0.001	0.001	0.002	0.002	0.003	0.002
	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)	(0.002)
Δ Log net wealth 2009–2008	0.221***	0.313***	0.309***	0.342***	0.383***	0.450***
	(0.072)	(0.087)	(0.089)	(0.088)	(0.091)	(0.099)
Observations	502	502	502	502	502	502
R-squared	0.087	0.098	0.089	0.101	0.123	0.124

portfolio compositions of each individual as of August 2008. $R_i(t)$is computed as the net flow of risky assets (positive for net purchases and negative for net sales), scaled by the value of total financial assets in August 2008.

The results are reported in Table 9. In all regressions we add some demographic controls and the change in total assets; results are invariant to these controls. The left-hand side variable represents the active reallocation in the period that goes from August 2008 to the date written at the top of each column. Thus, in column 1 the reallocation considered is the one during the period August to September 2008. In all the specifications except one, the αcoefficient is significantly bigger than zero, albeit also significantly less than one. In all the specifications, the coefficient β is not different from zero, while the coefficient $\delta$is negative and sometimes significantly different from zero, albeit always significantly different from −1. Thus, neither of the two models perfectly fits the data. Yet, considering that the noise in the data tends to downward bias the coefficient, the data seem more consistent with model (5)—in which the changes in risk aversion impact the portfolio rebalancing—than with model (4).

Loewenstein's model can explain why, in the 2008 financial crisis context, investors who did not lose any money became more risk averse even with respect to a gamble with given known odds such as our quantitative measure. The terrifying news appearing on television, the interaction with friends who lost money in the market, or the pictures of fired people leaving their failed banks might have triggered an emotional response.

5. The experiment

While suggestive, this hypothesis cannot be tested with our data because we do not have any direct measure of fear. Does the TV reporting of Lehman's fired employees trigger an emotional fear response, or does it increase the subjective probability of a very bad outcome? And if it triggers a fear response, is this response sufficiently strong to explain the increase in risk aversion that we have documented in Section 4?

To separate the emotional response from a Bayesian response and establish whether an emotional response can generate large increases in risk aversion, we rely on a laboratory experiment in which the outside environment is controlled for. As long as the treatment provides no information about the real world, the probability of an extreme event should remain constant between treated and untreated samples. To discriminate between the two hypotheses, the key feature of such an experiment is to induce fear in the lab without altering a subject's perception of her financial and economic prospects. To achieve such a goal we rely on the fear conditioning model used in psychology. Notice that our intent is not to prove whether fear causes an increase in risk aversion. This link has already been established (see, e.g., Cohn et al., 2015). Our purpose is to test whether the fear channel is powerful enough to generate an increase in risk aversion of a magnitude that resembles the one induced by the financial crisis.

As for the classical Pavlov (1927) experiment, the fear response can be triggered by conditioning factors, which

Table 10
Experimental evidence: comparison between the group of treated and untreated.
This table shows the summary statistics for treated and untreated subjects in the experiment run at Northwestern University and *t*-tests for the differences (last column). The risk aversion measures are elicited as described in Table II. The other variables are defined in the Internet Appendix, Section A4. The indicator for low risk investment is constructed from the qualitative question, setting it equal to one if the person chose “An OK return, with good degree of safety of my principal” or “Low returns, but no chance of losing my principal,” and zero otherwise. */**/*** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Variable	Obs. tot.	Mean treated	Mean non-treated	Difference
Quantitative measure of risk aversion (risk premium)	207	3198	2526	672**
Qualitative measure of risk aversion	210	2.54	2.41	0.13
Low risk investment	210	0.53	0.39	0.14*
Male	206	0.39	0.34	0.05
Age	203	19.77	19.83	−0.06
Caucasian	206	0.41	0.40	0.01
Income (thousands of dollars)	210	111.68	120.96	−9.28

have little or nothing to do with the experience itself. As Pavlov’s dog salivates when a bell rings, the fear response arises in the presence of stimuli associated to past traumatic events. This evidence suggests that a fear-based response can be triggered by fear stimuli in an unrelated domain. For example, Kinreich et al. (2011) show that watching a horror movie stimulates the amygdala in a way consistent with the arousal of fear. Yet, they do not provide evidence that this experience can alter a risk aversion measure like ours, nor that it can alter it to the extent we observe after the financial crisis. This is what we try to test.

We chose a brief horrifying scene from a movie that was sufficiently recent to be really scary for undergraduates accustomed to the scariest videogames (*Psycho* would not cut it), but sufficiently old to minimize the chance they had already seen it. We chose a five-minute excerpt from the 2005 movie, *Hostel*, directed by Eli Roth, which is characterized by stark and graphic images and that shows a young man inhumanly tortured in a dark basement. The movie won "Best Horror" at the Empire Awards in 2007.

Our experiment was run at Northwestern University in March 2011 in three different sessions. A total number of 249 students took part. The participants were recruited through an internal mailing list service that is normally employed for experiments at Northwestern.[6] A compensation of $5 was paid in cash to each subject taking part in the experiment, which in general takes around 10–15 min.

All the participants were asked to complete a questionnaire of approximately 40 questions. Its main scope is to construct some measures of risk aversion, as well as to provide other controls. To identify the effect of fear on the subjects, we relied on a simple treatment and control framework. In particular, around half of the participants were asked to watch a short video before completing the questionnaire. Since the subjects were randomly assigned to watch the video, the idea is that the difference in risk aversion between the two groups should be completely driven by this difference in the treatment.

Given the nature of the video, which potentially disturbs some of the subjects, we had to give them the option to skip the video at any moment. We dropped the observations of the subjects (27) who decided to skip the video in the first minute of the five-minute presentation, since they did not really experience much horror. This choice might underestimate the effect of the treatment, since those most sensitive to the treatment dropped out.

Another possible concern is that, if a subject has already watched the video, its perceived effect would be different from the true effect. We therefore decided to drop those 13 subjects who declared to have already watched it.

To guarantee the reliability of the results, the experiment was designed in such a way that the participants were not aware that the treatment was not identical for everyone. As measures of risk aversion, we use answers to the very same questions that were used in the bank survey, in which we translated euros into dollars at a 1:1 ratio.

As Table 10 shows, the random assignment assumption cannot be rejected: none of the main personal characteristics and demographic information has been found to be statistically different between treatment and control groups. Furthermore, around 60% of the participants were female and the average age is 20, which is not surprising given that the sample is composed of undergraduate students.

When we look at the risk aversion measure, we find that there is a large and statistically significant difference in the quantitative measure of risk aversion. Among the treated students, the risk premium they are willing to pay for avoiding the risky lottery is 672 dollars (i.e., 27%) higher. This holds true without controls and controlling for observables (Internet Appendix, Table A.6 columns 1 and 2).

In the qualitative measure we observe an increase, but this increase is not statistically significant at the conventional level (p-value = 0.11). In part, this phenomenon is due to the fact that students bunched their choices in the two central values: 96% of the responses are either two or three. Hence, the scale 1–4 is probably better reduced to a dichotomous choice: low risk aversion (1 and 2) and high risk (3 and 4). When we look at the proportion of people

[6] The students can freely enroll on the mailing list and, after they have completed an introductory demographic survey, they receive periodic communications on the experiments that are going on at the University.

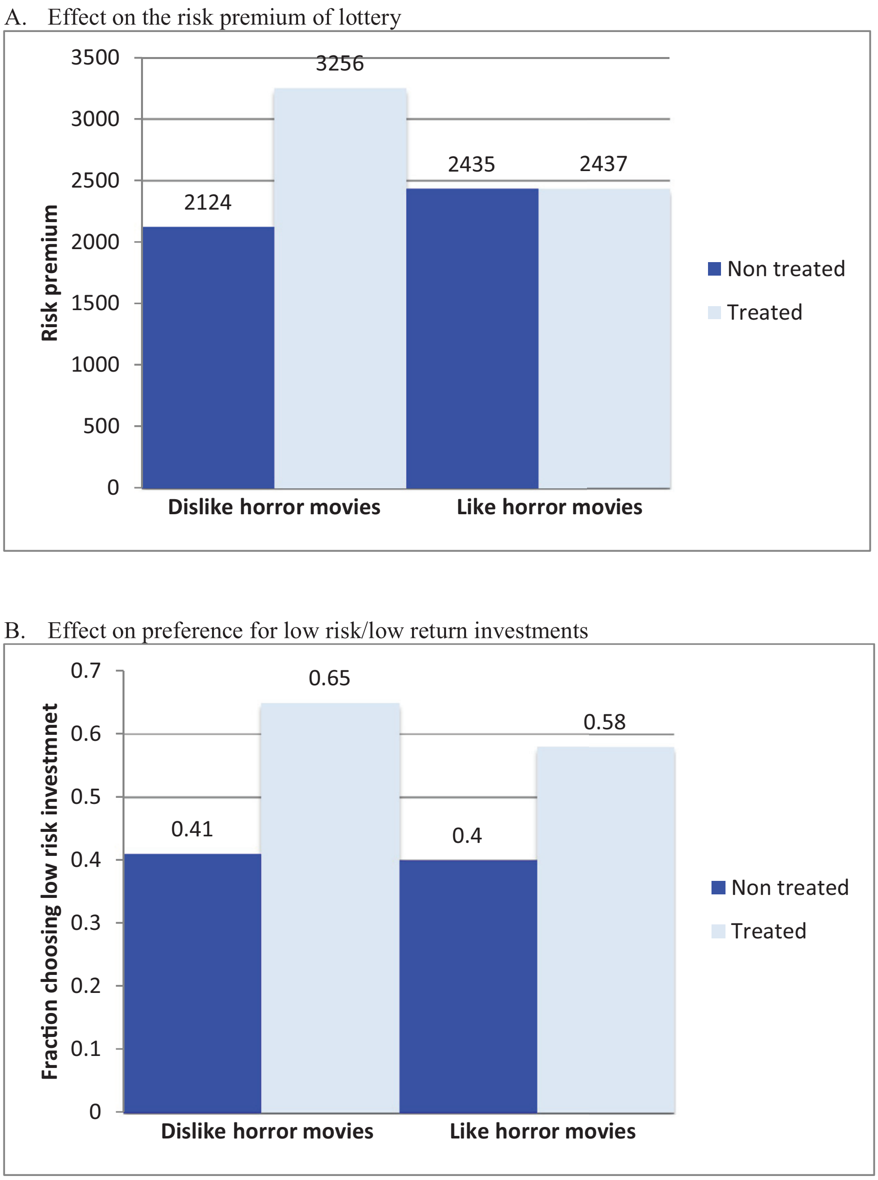

Fig. 4. Effect of fear on risk aversion. The figure shows the average risk aversion of subjects "Treated" with the horror movies or "Not treated," for groups of subjects that differ in how much they like horror movies. "Like horror movies" is calculated from a survey measure asking subjects whether they like horror movies in a scale ranging from zero to 100 increasing in liking. Subjects who replied with a value of 20 (the median) or more are classified as liking horror movies. "Dislike horror movies" is the group that reports less than 20 in liking of horror movies. Panel A shows the effect on the risk premium of the lottery; Panel B presents the effect on the risky investment choice.

choosing the low risk option, this proportion increases by 13.5% (30% of the sample mean) among the treated group (Columns 5 and 6). This difference is large and statistically significant at the 5% level.

In the second half of the sample, we asked people how much they liked horror movies on a scale from zero to 100. Roughly a third of the sample declared they do not like them at all (i.e., like = 0) and 50% report a value of liking below 20. In Fig. 4 we split the sample on this basis. In the first group, there are students who do not like horror movies (liking indicator below median). Their risk premium rises from 2124 to 3256 dollars as a result of the treatment (Panel A). This difference is statistically significant at the 1% level.

The second group is formed by those subjects who moderately like horror movies (liking indicator above 20). Here the treatment has no effect (the risk premium goes from 2435 to 2437) and this difference is not statistically significant.

We get a similar result when we look at the qualitative measure of risk aversion, where we bunched the responses into two groups. Among people who dislike horror movies, the treatment effect increases the probability of buying risky assets by almost 25%. Among those who moderately like horror movies, the increase is significantly smaller by 7%.

These results seem to be inconsistent with the background risk hypothesis and suggest that fear is a potential (and understudied) mechanism that influences financial decisions, whether it does so by increasing the curvature of the utility function or the salience of negative outcomes (Bordalo et al., 2012).

6. Conclusions

In our view, the paper has two main contributions. The first is methodological. Most papers use either naturally occurring data or lab/field data, but not both. We think that well designed lab tests can be a useful complement to the analysis of naturally occurring data, when we are faced with very important questions that are impossible to answer with naturally occurring data. In particular, it is impossible to disentangle whether the surge in risk aversion we observe in the data is due to fear or background risk. If we were unable to reproduce that surge in the lab, we could have ruled out the fear explanation. The fact we were able to does not prove that the cause of the surge in the data is fear, but it makes it more plausible. Thus, targeted lab experiments can help in identifying effects difficult to sort out in naturally occurring data.

The second contribution is to provide some evidence consistent with a fear-based explanation of the increase of risk aversion during the financial crisis.

A question we are unable to answer in this paper is how persistent such fear-induced change in risk aversion is. The evidence of Malmendier and Nagel (2011), who find a cohort effect of “Depression era babies” in the risk aversion measure of the Survey of Consumer Finances, suggests it might be long-lasting. With our sample we are unable to answer whether fear provokes long-term consequences because of the subsequent events in the eurozone, which made the 2008 shock not an isolated crisis.

Finally, our results raise an interesting question. If the behavior we document is typical and during severe downturns investors are caught by fear and sell their risky assets at the worst times, the effective return on equity investment is much lower. Can this feature explain—at least in part—the famous equity premium? Only future research would be able to tell.

References

Baker, M., Wurgler, J., 2007. Investor sentiment in the stock market. J. Econ. Perspect. 21 (2), 129–151.

Barberis, N., Huang, M., Santos, T., 2001. Prospect theory and asset prices. Q. J. Econ. 116 (1), 1–53.

Barsky, R., Juster, T., Kimball, M., Shapiro, M., 1997. Preference parameters and behavioral heterogeneity: an experimental approach in the health and retirement Study. Q. J. Econ. 112 (2), 537–579.

Bassi, A., Colacito, R., Fulghieri, P., 2013. O sole mio: an experimental analysis of weather and risk attitudes in financial decisions. Rev. Financial Stud. 26 (7), 1824–1852.

Bombardini, M., Trebbi, F., 2012. Risk aversion and expected utility theory: an experiment with large and small stakes. J. Eur. Econ. Assoc. 10 (6), 1348–1399.

Bonaparte, Y., Cooper, R., 2009. Costly Portfolio Adjustment. NBER Working Paper No. 15227.

Bordalo, P., Gennaioli, N., Shleifer, A., 2012. Salience theory of choice under risk. Q. J. Econ. 127 (3), 1243–1285.

Byrnes, J., Miller, D., Schafer, W., 1999. Gender differences in risk taking: a meta-analysis. Psychol. Bull. 125 (3), 367–383.

Caballero, R., Krishnamurthy, A., 2008. Collective risk management in a flight to quality episode. J. Finance 63 (5), 2195–2230.

Campbell, J., Cochrane, J., 1999. By force of habit: a consumption-based explanation of aggregate stock market behavior. J. Polit. Econ. 107 (2), 205–251.

Cohn, A., Engelmann, J., Fehr, E., Maréchal, M., 2015. Evidence for countercyclical risk aversion: an experiment with financial professionals. Am. Econ. Rev. 105 (2), 860–885.

Croson, R., Gneezy, U., 2009. Gender differences in preferences. J. Econ. Lit. 47 (2), 448–474.

Dohmen, T., Falk, A., Huffman, D., Sunde, U., Schupp, J., Wagner, G., 2011. Individual risk attitudes: measurement, determinants, and behavioral consequences. J. Eur. Econ. Assoc. 9 (3), 522–550.

Donkers, B., Melenberg, B., Van Soest, A., 2001. Estimating risk attitudes using lotteries: a large sample approach. J. Risk Uncertain. 22 (2), 165–195.

Friend, I., Blume, M., 1975. The demand for risky assets. Am. Econ. Rev. 65 (5), 900–922.

Guiso, L., Paiella, M., 2006. The role of risk aversion in predicting individual behavior.. In: Chiappori, P., Gollier, C. (Eds.), Competitive Failures in Insurance Markets: Theory and Policy Implications. MIT Press, Boston, MA, pp. 213–250.

Guiso, L., Paiella, M., 2008. Risk aversion, wealth and background risk. J. Eur. Econ. Assoc. 6 (6), 1109–1150.

Hartog, J., Ferrer-i-Carbonell, A., Jonker, N., 2002. Linking measured risk aversion to individual characteristics. Kyklos 55 (1), 3–26.

Heaton, J., Lucas, D., 2000. Portfolio choice in the presence of background risk. Econ. J. 110 (460), 1–26.

Holt, C., Laury, S., 2002. Risk aversion and incentive effects. Am. Econ. Rev. 92 (5), 1644–1655.

Kamstra, M., Kramer, L., Levi, M., 2003. Winter blues: a SAD stock market cycle. Am. Econ. Rev. 93 (1), 324–343.

Kandasamy, N., Hardy, B., Page, L., Schaffner, M., Graggaber, J., Powlson, A., Fletcher, P., Gurnell, M., Coates, J., 2014. Cortisol shifts financial risk preferences. Proce. Nat. Acad. Sci. 111 (9), 3608–3613.

Kinreich, S., Intrator, N., Hendler, T., 2011. Functional cliques in the amygdala and related brain networks driven by fear assessment acquired during movie viewing. Brain Connect. 1 (6), 484–495.

Knutson, B., Wimmer, E., Kuhnen, C., Winkielman, P., 2008. Nucleus accumbens activation mediates the influence of reward cues on financial risk taking. NeuroReport 19 (5), 509–513.

Kramer, L., Weber, M., 2012. This is your portfolio on winter: seasonal affective disorder and risk aversion in financial decision making. Soc. Psychol. Personal. Sci. 3 (2), 193–199.

Kuhnen, C., Knutson, B., 2005. The neural basis of financial risk taking. Neuron 47 (5), 763–770.

Kuhnen, C., Knutson, B., 2011. The impact of affect on beliefs, preferences and financial decisions. J. Financial Quant. Anal. 46 (3), 605–626.

Lerner, J., Keltner, D., 2000. Beyond valence: toward a model of emotion-specific influences on judgment and choice. Cognit. Emot. 14 (4), 473–494.

Lerner, J., Keltner, D., 2001. Fear, anger, and risk. J. Personal. Soc. Psychol. 81 (1), 146–159.

Loewenstein, G., 2000. Emotions in economic theory and economic behavior. Am. Econ. Rev. 90 (2), 426–432.

Loewenstein, G., Weber, E., Hsee, C., Welch, N., 2001. Risk as feelings. Psychol. Bull. 127 (2), 267–286.

Malmendier, U., Nagel, S., 2011. Depression babies: do macroeconomic experiences affect risk-taking? Q. J. Econ. 126 (1), 373–416.

Merton, R., 1969. Lifetime portfolio selection under uncertainty: the continuous-time case. Rev. Econ. Stat. 51 (3), 247–257.

Pavlov, I., 1927. Conditioned Reflex: An Investigation of the Psychological Activity of the Cerebral Cortex. Oxford University Press, New York.

Weber, M., Weber, E., Nosić, A., 2013. Who takes risks when and why: determinants of changes in investor risk taking. Rev. Finance 17 (3), 847–883.