EUI Working Papers
ECO 2007/47

Risk Aversion, Wealth, and Background Risk

Luigi Guiso and Monica Paiella

EUROPEAN UNIVERSITY INSTITUTE

DEPARTMENT OF ECONOMICS

Risk Aversion, Wealth, and Background Risk

LUIGI GUISO

and

MONICA PAIELLA

Risk Aversion, Wealth, and Background Risk[*]

Luigi Guiso [†]   Monica Paiella [‡]

(This draft 15.03.2007)

[†]European University Institute, University of Rome, Tor Vergata & Ente Luigi Einaudi; email: guiso@tin.it

[‡]Bank of Italy & Ente Luigi Einaudi

Abstract

We use household survey data to construct a direct measure of absolute risk aversion based on the maximum price a consumer is willing to pay for a risky security. We relate this measure to consumers' endowments and attributes and to measures of background risk and liquidity constraints. We find that risk aversion is a decreasing function of the endowment—thus rejecting CARA preferences. We estimate the elasticity of risk aversion to consumption at about 0.7, below the unitary value predicted by CRRA utility. We also find that households' attributes are of little help in predicting their degree of risk aversion, which is characterized by massive unexplained heterogeneity. We show that the consumer's environment affects risk aversion. Individuals who are more likely to face income uncertainty or to become liquidity constrained exhibit a higher degree of absolute risk aversion, consistent with recent theories of attitudes toward risk in the presence of uninsurable risks.

1 Introduction

The relationship between wealth and a consumer's attitude towards risk—as indicated, for instance, by the degree of absolute risk aversion or of absolute risk tolerance—is central to many fields of economics. As argued by Kenneth Arrow more than 35 years ago, "the behavior of these measures as wealth varies is of the greatest importance for prediction of economic reactions in the presence of uncertainty" (1970, p. 35).

Most inference on the nature of this relation is based on common sense, introspection, casual observation of behavioral differences between the rich and the poor, and a priori reasoning. Such inference concerns the sign of the relation, but there is no (even indirect) evidence concerning its curvature. The consensus view is that absolute risk aversion should decline with wealth.[1] Furthermore, if one agrees that preferences are characterized by constant relative risk aversion (a property of one of the most commonly used utility functions, the isoelastic), then absolute risk aversion is decreasing and convex in wealth, while risk tolerance is increasing and linear. The curvature of absolute risk tolerance has been shown to be relevant in a number of contexts. For example, Gollier and Zeckhauser (2002) show that it determines whether the portfolio share invested in risky

---

[1]It is on these grounds that quadratic and exponential utility, though often analytically convenient, are regarded as misleading representations of preferences; the first implies increasing absolute risk aversion and the second posits constant absolute risk aversion.

assets increases or decreases over the consumer life cycle, an issue that is receiving increasing attention. If risk tolerance is concave, then wealth inequality can help elucidate the risk premium puzzle (Gollier 2001). Furthermore, the curvature of risk tolerance and the nature of risk aversion may explain why the marginal propensity to consume out of current resources declines as the level of resources increases (Carroll and Kimball 1996; Gollier 2001) even as the elasticity of risk tolerance to the endowment influences the size of the precautionary saving motive (Kimball and Weil 2004).

The aim of this paper is to provide empirical evidence on the nature of the relationship between risk aversion and wealth. Using data from the Bank of Italy Survey of Household Income and Wealth (SHIW) on household willingness to pay for a hypothetical risky security, we recover a measure of the Arrow–Pratt index of absolute risk aversion of the consumer lifetime utility function. We then relate it to indicators of consumers endowment and also to a set of demographic characteristics as a control for individual preference heterogeneity.

Our findings show that absolute risk tolerance is an increasing function of consumers' resources, thus rejecting constant <u>absolute</u> risk aversion (CARA) preferences. We estimate that the elasticity of absolute risk tolerance to consumers' resources is between 0.6 and 0.75. This value is smaller than that implied by constant <u>relative</u> risk aversion (CRRA) preferences and suggests that

4

risk tolerance is a concave function of wealth. This result is robust to accounting for endogeneity of the endowment and for unobserved cognitive ability. When we instrument the measure of the endowment with exogenous variation in the education of the household head's father and with windfall gains, the estimated elasticity is only marginally different than the OLS estimate. We also show that our results are unaffected by different interpretations of the hypothetical security question or by the presence of classical measurement error in the household's willingness to pay for the security.

The usual definition of risk aversion and tolerance developed by Arrow (1970) and Pratt (1964) is based on the assumption that initial wealth is non-random. This definition is also constructed in a static setting or in settings where full access to the credit market is assumed. Recently it has been shown that attitudes toward risk can be affected by the prospect of being liquidity constrained and by the presence of additional uninsurable, nondiversifiable risks. For instance, Gollier (2000) shows that consumers who may be subject to future liquidity constraints are less willing to bear present risks (i.e., their risk aversion increases). Moreover, depending on the structure of preferences, the presence of risks that cannot be avoided or insured against (background risks) may make individuals less tolerant toward other, avoidable risks (Pratt and Zeckhauser 1987; Kimball 1993; Eekhoudt, Gollier, and Schlesinger 1996). One important

consequence is that individuals facing high exogenous labor income risk—which is normally uninsurable—will be more risk averse and will thus avoid exposure to portfolio risk by holding less or no risky assets.

The evidence presented in this paper sheds light also on the empirical relevance of these concepts. The availability of information on measures of background risk and on proxies of borrowing constraints allows us to relate our index of risk aversion to indicators of income-related risk and of liquidity constraints.

We find that risk aversion is positively affected by background risk and also by the possibility of being credit constrained. The effects of background risk and exposure to liquidity constraints on household willingness to bear risk are also economically relevant: increasing our measure of background risk by a single standard deviation lowers absolute risk tolerance by about 19%; being liquidity constrained lowers it by 4.4%. Overall, however, our estimates suggest that these variables can explain only a small amount of the sample variability in attitudes toward risk. Even after controlling for individual exogenous characteristics (e.g., age, gender, region of birth) there remains a large amount of unexplained variation that reflects, in part, genuine differences in tastes.

The rest of the paper is organized as follows. Section 2 describes our measure of risk aversion; we describe when wealth is nonrandom and when there is background risk. Section 3 presents descriptive evidence on absolute risk aver-

sion in our cross-section of households. In Section 4 we discuss the empirical specification used to relate absolute risk aversion to the consumer's endowment, attributes, and environment. Section 5 presents the results of the estimates. In Section 6 we check the robustness of the main findings with respect to the endogeneity of consumption and wealth, to nonresponses, and to the possible presence of outliers. Section 7 presents evidence regarding the effects of background risk on the propensity to bear risk. Section 8 discusses the consistency with observed behavior of our findings on the shape of the wealth–risk aversion relation. Section 9 summarizes and concludes.

2   Measuring risk aversion

2.1   No background risk

To measure absolute risk aversion and tolerance, we exploit the 1995 wave of the Survey of Household Income and Wealth, which is run biannually by the Bank of Italy. The 1995 SHIW collects data on income, consumption, real and financial wealth, and several demographic variables for a representative sample of 8,135 Italian households. Balance-sheet items are end-of-period values. Income and flow variables refer to 1995.[2]

---

[2]See the Appendix for a detailed description of the survey contents, its sample design, interviewing procedure, and response rates.

The 1995 survey has a section designed to elicit attitudes toward risk. Each participant is offered a hypothetical security and is asked to report the maximum price that he would be willing to pay for it. Specifically:

We would like to ask you a hypothetical question that we would like you to answer as if the situation was a real one. You are offered the opportunity of acquiring a security permitting you, with the same probability, either to gain 10 million lire or to lose all the capital invested. What is the most that you are prepared to pay for this security?

Ten million lire corresponds to just over 5,000 euros (or roughly $6,500). The ratio of the expected gain from the investment to average household total consumption is 16%; thus, the investment represents a relatively large risk. We consider this as an advantage, because expected utility maximizers behave as risk-neutral individuals with respect to small risks even if they are averse to larger risks (Arrow 1970). Thus, presenting consumers with a relatively large investment is a better strategy for eliciting risk attitudes when one relies (as we do) on expected utility maximization to characterize risk aversion.[3] The interviews

---

[3]In this vein, Rabin (2000) argues that if an expected utility maximizer refuses a small risk at all levels of wealth than he must exhibit unrealistic levels of risk aversion when faced with large-scale risks. This again suggests that offering large investments is a better way to characterize the risk aversion of expected utility maximizers.

are conducted personally at the consumer's home by professional interviewers. In order to help the respondent understand the question, the interviewers show an illustrative card and are ready to provide explanations. The respondent can answer in one of three ways: (i) declare the maximum amount he is willing to pay for the security, which we denote Z_i; (ii) don't know; (iii) unwilling to answer.

Notice that the hypothetical security's design is such that with probability $1/2$ the respondent gets 5,000 euros and with probability $1/2$ he loses Z_i. So the expected value of the security is $(5000 - Z_i)/2$. Clearly, $Z_i < 5000$, $Z_i = 5000$, and $Z_i > 5000$ euros imply (respectively) risk aversion, risk neutrality, and risk loving. This characterizes attitudes toward risk qualitatively. But we can do more: within the expected utility framework, a measure of the Arrow–Pratt index of absolute risk aversion can be obtained for each consumer. Let w_i denote household i's endowment, which for the moment we assume to be nonrandom. Let $u_i(\cdot)$ be the (lifetime) utility function and let $\widetilde{P}_i$ be the security random return for individual i, taking values 5000 and $-Z_i$ with equal probability. The maximum purchase price is thus given by

$$u_i(w_i) = \frac{1}{2}u_i(w_i + 5000) + \frac{1}{2}u_i(w_i - Z_i) = Eu_i(w_i + \widetilde{P}_i), \qquad (1)$$

where E is the expectations operator.

One way to derive a measure of the implied risk aversion would be to take a second-order Taylor expansion of the second equality in (1) around w_i and

then obtain an estimate of the Arrow–Pratt measure of absolute risk aversion in terms of the parameters of the hypothetical security of the survey. The problem with this approach is that the risk aversion at low levels of the price Z_i would be greatly underestimated, biasing toward zero the estimated relation between risk aversion and the consumer's endowment.[4]

In order to avoid this problem, we assume a specific functional form for the utility function such that the coefficient of absolute risk aversion tends to infinity as the maximum reported price tends to zero; we use this form to compute the implied absolute risk aversion for each individual in the sample. Note that we use the specific functional form only for mapping the reported willingness to pay into a measure of risk aversion. Ultimately, we are interested in the relation between risk aversion and the consumer's endowment, which is determined not by the specific utility function used for the mapping but rather by the relation between the reported price and the level of the individual endowment.

---

[4]To see this, observe that the second-order Taylor expansion of the right-hand side of (1) around w_i gives

$$Eu_i(w_i + \widetilde{P}_i) \approx u_i(w_i) + u_i'(w_i)E(\widetilde{P}_i) + 0.5u_i''(w_i)E(\widetilde{P}_i)^2.$$

Substituting this expression into equation (1) and simplifying, the measure of absolute risk aversion, $R_i(w_i)$, will be

$$R_i(w_i) \approx -u_i''(w_i)/u_i'(w_i) = 2(5000 - Z_i)/(5000^2 + Z_i^2).$$

As Z approaches zero, this measure of R tends to 0.2, whereas the true measure of risk aversion tends to infinity.

In other words, the assumed utility function is instrumental only to avoid the bias in the estimated risk aversion that would result from a second-order Taylor approximation of (1); it has no effect on the relation between endowment and absolute risk aversion. To make this point more precise, observe that the elasticity of absolute risk aversion to the endowment, $d \log R / d \log w$, can be decomposed as $d \log R / d \log w = (d \log R / dZ)(dZ / d \log w)$. The term $d \log R / dZ$ depends on how absolute risk aversion is computed and would be biased toward zero if a second-order approximation were used. Instead, the data address the term $dZ / d \log w$. In other words, it is the dependence of the willingness to pay on consumer resources that is decisive for the relation between endowment and absolute risk aversion. We emphasize this dependence by letting $Z_i = Z(w_i)$.

To compute the measure of absolute risk aversion, we experimented with two different functions: the exponential utility and the CRRA utility. In the first case we have solved the equation

$$-\mathbf{e}^{-R_i w_i} = -\frac{1}{2}\mathbf{e}^{-R_i(w_i+5000)} -\frac{1}{2}\mathbf{e}^{-R_i(w_i- Z(w_i))} \tag{2}$$

for the unknown parameter R_i. In the second case we have solved the equation

$$\frac{w_i^{1-\gamma_i}}{1-\gamma_i} = \frac{1}{2}\frac{(w_i+5000)^{1-\gamma_i}}{1-\gamma_i} + \frac{1}{2}\frac{(w_i-Z(w_i))^{1-\gamma_i}}{1-\gamma_i} \tag{3}$$

for the relative risk aversion parameter γ_i and then computed absolute risk aversion as $R_i = \gamma_i / w_i$. For all practical purposes it makes no difference which utility function is used to obtain the risk aversion measure. The average value of R_i

is 0.01981(median 0.000708) using the exponential utility and 0.01978 (median 0.000693) using the CRRA utility, and the correlation coefficient differs little from unity. This is consistent with the idea that the only role of the assumed functional form is to obtain an unbiased estimate of risk aversion at low levels of Z. Hence, as we will show, our results are invariant to which measure of risk aversion is used.[5]

Equation (2) (or (3)) uniquely defines the Arrow–Pratt measure of absolute risk aversion in terms of the parameters of the hypothetical security of the survey. Absolute risk tolerance is defined by $T(w_i) = 1/R(w_i)$. Obviously, $R(w_i) = 0$ for risk-neutral individuals (i.e., those reporting $Z_i = 5000$) and $R(w_i) < 0$ for risk-loving individuals (those with $Z_i > 5000$). Since R_i is specific to the individual, it may vary with consumer endowment and with all the attributes correlated with consumer preferences. The loss Z_i or the gain from the investment need not benefit or be fully borne by current consumption but instead may be spread over lifetime consumption. Therefore, our measure of risk aversion is better interpreted as the risk aversion of the consumer's lifetime utility and w_i as the consumer's lifetime wealth.[6]

---

[5]Before computing the risk aversion under the CRRAA utility, one must specify the consumer endowment. Using household consumption or cash on hand (defined as disposable income plus financial wealth) has yielded almost indistinguishable measures of absolute risk aversion.

[6]Tiseno (2002) studies the relationship between the risk aversion of lifetime utility ant that of period utility, showing how one can recover the latter given information on the curvature of

A few comments are in order regarding this measure and on how it compares with those used in other studies. First, it is not restricted to risk-averse individuals but extends to the risk neutral and to risk lovers. Second, our definition provides a point estimate, rather than a range, of the degree of risk aversion for each individual in the sample. These features mark a difference between our study and that of Barsky, et al. (1997), who obtain only a range measure of (relative) risk aversion and a point estimate under the assumption that preferences are strictly risk averse and utility is of the CRRA type. Furthermore, their sample consists of individuals aged 50 or above, which makes it difficult to study the age profile of risk aversion and also to test its relationship with background risk, which is likely to matter the most for the young. The study of this relationship is one of our aims in this paper. Note, however, that their elicitation strategy yields a measure of the risk aversion of period utility instead of lifetime utility as here. In this regard, the two studies should be viewed as complementary.[7]

the former and on the slope of the consumption function. She also shows that knowledge of the maximum subjective price function for a risk is sufficient to identify the risk aversion of a consumer lifetime utility.

[7]The Barsky et al. (1997) measure of risk aversion has other advantages. Since the risk tolerance question is asked in two waves of their survey and since a subset of the respondents is common to both waves, the authors can account for measurement error in their measure of relative risk aversion. They also collect information on intertemporal substitution and thus can study its relation to risk aversion.

2.2 Risk aversion with background risk

The measure of risk aversion in (2) and (3) is for nonrandom endowment, but it is easily generalized to the case of background risk using the results of Kihlstrom, Romer, and Williams (1981) and Kimball (1993). Pratt and Zeckhauser (1987), Kimball (1993), and Eekhoudt, Gollier, and Schlesinger (1996) establish a set of conditions on preferences that define classes of utility functions whose common feature is background risk whose presence makes the individual behave in a more risk-averse manner. These classes of utility functions are termed "proper", "standard", and "risk vulnerable" in the three respective studies.[8] The main implication is that even if risks are independent, individuals who are risk averse to begin with will react to background risk by reducing their exposure to avoidable

---

[8]Pratt and Zeckhauser (1987) define as "proper" the class of utility functions that ensure that introducing an additional independent undesirable risk when another undesirable one is already present makes the consumer less willing to accept the extra risk. Kimball (1993) defines as "standard" the class of utility functions that guarantee that an additional independent undesirable risk increases the sensitivity to other loss-aggravating ones. Starting from initial wealth w, a risk $\widetilde{x}$ is undesirable if and only if it satisfies $Eu(w - \widetilde{x}) \leq u(w)$, where $u(w)$ is an increasing and concave utility function. A risk $\widetilde{x}$ is loss-aggravating if and only if it satisfies $Eu'(w + \widetilde{x}) \geq u'(w)$. When absolute risk aversion is decreasing, every undesirable risk is loss-aggravating but not every loss-aggravating risk is undesirable. Finally, Eeckhoudt, Gollier, and Schlesinger's (1996) risk vulnerability implies that adding a zero-mean background risk makes consumers more risk averse.

risks. Hence, they should hold safer portfolios and tend to buy more insurance against the risks that are insurable (Eekhoudt and Kimball 1992).[9] Furthermore, insofar as income risk evolves with age, under standardness we see that background risk may help explain the life cycle of asset holdings. Several papers have cited background risk and risk vulnerability (or standardness) to explain portfolio puzzles.[10] In all these studies, standardness or risk vulnerability is simply assumed; it is not tested because evidence on individual risk aversion is lacking. Here we can provide such a test. For this purpose we must restrict the analysis to risk-averse individuals (i.e., those reporting $Z_i < 10$).

Let $\widetilde{y}_i$ denote a zero-mean background risk for individual i, with variance σ^2. If we use E_x $(x = y, P)$ to denote the expectation with respect to the random variable $\widetilde{x}$, then our indifference condition for purchasing the risky security and paying Z_i becomes

$$E_y u_i(w_i + \widetilde{y}_i) = E_P E_y u_i(w_i + \widetilde{y}_i + \widetilde{P}_i), \tag{4}$$

where we have implicitly assumed that the background risk and the risky security are independent, which is assured by construction. If preferences are risk

---

[9]Guiso, Jappelli, and Terlizzese (1996) find that households facing greater earnings risk buy assets that are less risky; or fewer risky assets; Guiso and Jappelli (1998) show that households buy more liability insurance in response to earnings risk.

[10]See Weil (1992), Gollier (2000), Heaton and Lucas (2000), and Gollier and Zeckhauser (2002).

vulnerable as in Gollier and Pratt (1996), we can use the equivalence

$$E_y u_i(w_i + \widetilde{y}_i) = v_i(w_i); \tag{5}$$

here $v_i(w_i)$ is a concave transformation of u_i, which implies that $v_i(w_i)$ is more risk averse than $u_i(w_i)$. In other words, if consumers h and j are both risk averse and if their preferences are risk vulnerable, then (assuming $w_j = w_h$) h is more risk-averse than j if $\widetilde{y}_h$ is riskier than $\widetilde{y}_j$—that is, if h faces more background risk.

We can thus account for background risk by expressing our measure of risk aversion in terms of the utility function $v_i(w_i)$, obtaining $R_i(w_i) \approx -v_i''(w_i)/v_i'(w_i)$ from either equation (2) or (3). Risk aversion will now vary not only with the consumer's endowment and attributes but also with any source of uncertainty characterizing her environment. If measures of the latter are available, then one can directly test for standardness of preferences.

It is interesting the the shape of the relation between R (or risk tolerance) and w can have implications for the sign of the effect of background risk on absolute risk aversion. Hennessy and Lapan (2006) show that a positive and concave relation between risk tolerance and wealth is sufficient for preferences to be standard as in Kimball (1993). Similarly, Eekhoudt, Gollier, and Schlesinger (1996) show that a sufficient (but not necessary) condition for absolute risk aversion to increase with background risk is that it be a decreasing and convex

16

function of the endowment—an assumption that is satisfied, for instance, by CRRA utility. Gollier and Pratt (1996) argue that the convexity of absolute risk aversion should be regarded as a natural assumption,[11] "since it means that the wealthier an agent is, the smaller is the reduction in risk premium of a small risk for a given increase in wealth." Though plausible, this assertion is not backed by any empirical evidence. Our results lend support to this conjecture by implying that absolute risk aversion is a convex function of the endowment.

3   Descriptive evidence

The question on the risky security was submitted to the entire sample of 8,135 household heads, but only 3,458 answered as being willing to purchase the security. Out of the 4,677 others, 1,586 reported a "do not know" and 3,091 overtly refused to answer or to pay a positive price (25 offered more than 10,000 euros; we omitt these responses because such a price leads to a sure loss). This is likely due to the complexity of the question, which might have led some participants to skip it altogether because of the relatively long time required to understand its meaning and to provide an answer. Nonresponses also reflect that the question on the risky security was asked abruptly by the interviewers, without preparing the respondents with a set of "warm up" questions. However, this strategy has its

---

[11] Observe that if consumers are risk averse at all levels of wealth and if absolute risk aversion is a strictly decreasing function of wealth, then absolute risk aversion must be convex in wealth.

advantages. First, depending on how the introductory questions are framed and when they are asked, they may end up affecting the answers and thus distorting the measure of the true preference parameter; this is avoided by asking the question abruptly. Second, it avoids bringing in noisy respondents (e.g., those with a poor understanding of the question), as would probably happen with "warm up" questions. Thus, although a high nonresponse rate signals that the question is complex and there may be cognitive problems, it does not mean that those who chose to respond gave erroneous answers. To the contrary, if those who answered did so because they had a good understanding of the question (or the time to grasp and answer it), then the elicitation strategy with no "warm up" questions may have effectively screened out the noisy respondents.

Figure 1 shows the histogram of the willingness to pay for the total sample and for the subsamples of low- and high-educated individuals (defined as people with up to and more than middle schooling, respectively). The reported price ranges from a few euros to 10,000 euros or more (few observations). However, the bulk of the responses (70% of the sample) are between 300 and 5,000 euros; there is a median willingness to pay of 500 euros and a mean of 1,161, signaling a distribution with a long right tail. Not surprisingly, individual responses show spikes at "focal" prices such as 100 euros and multiples of 1,000 euros but there are also many observations at nonfocal prices. A nonnegligible number of

respondents (576) declare they are ready to pay as much as 5,000 euros for the hypothetical security, but few are willing to pay more than that.

When the figure is drawn by level of education, the distribution is shifted to the right for the highly educated. This group has the same median as the low-educated group but a higher *mean* (1,367 euros compared to 1,014); otherwise, the two distributions look similar. This is consistent with high-educated, wealthier individuals having a higher willingness to pay. In fact, the sample correlation between reported price and level of household consumption is 0.15 (standard error 0.017; Table 1, Panel B), and a simple regression of the reported price on the level of consumption shows a positive and highly significant coefficient (t-statistic $= 6.7$). As shown in Section 2, a positive correlation between willingness to pay and individual endowment is necessary for decreasing risk aversion, though this is not informative as to the shape of the risk aversion function. Panel A of Table 1 shows the moments of the distribution of the willingness to pay when the sample is grouped by gender and by age, and panel B shows sample correlation of willingness to pay with various measures of the endowment.

From the initial data set we drop those households whose total net wealth is less than 50 euros and those whose head's age is less than 21 or more than 89 years (5% of the observations). Table 2 reports descriptive statistics for the whole sample of 7,704 households, for the sample of 3,297 respondents to the

19

risky security question, and also for several subsamples of the latter. Out of 3,297 individuals willing to purchase the security, the great majority (96%) are risk averse in that they report a maximum price lower than the gain offered by the security; 140 individuals are either risk neutral (123, or 3.7% of the sample) or risk loving (17, a tiny minority). Those who responded to the question are on average six years younger than the total sample and have higher shares of male-headed households (80.1% compared to 74.9%), of married people (79.3% and 73.3%, respectively), of self-employed (18.2% and 14.6%) and of public sector employees (28.2% and 23.9%). They are also somewhat wealthier and slightly better educated (1.3 more years of schooling). These differences suggest that there are some systematic effects explaining the willingness to respond. Probit regressions, reported columns (1) and (2) of Table A.1, confirm this factual evidence and suggest that the probability of answering the question is higher among younger and more educated households. Public-sector employees are also relatively more likely to respond. Moreover, the response probability is increasing in household income but decreasing in net worth.

In order to shed more light on what is driving the answers to the hypothetical security question, we report summary statistics distinguishing between those who refused to respond and those who were unwilling to pay a positive price. If answering the risky security question requires some nontrivial compu-

tational ability in order to price risk, individuals with weaker cognitive abilities can be expected to face higher computational costs, which in turn may lead them to refusing participation or (perhaps equivalently) posting a zero price.[12] It turns out that those reporting a zero price are closer, in terms of observable characteristics, to those reporting a positive price than are the nonrespondents. Nonrespondents are on average less educated than those reporting a zero price (6.6 years of education versus 7.8 years), and both are less educated than the participants (9.4 years on average). Since education is positively correlated with cognitive ability (Cascio and Lewis 2005), this pattern is consistent with heterogeneity in cognitive ability driving heterogeneity in the willingness to answer

---

[12]One could argue that nonparticipation or participation at zero price may reflect nonstandard preferences, such as loss aversion and narrow framing. It is well known that loss aversion (especially when coupled with narrow framing) may explain why an individual turn's down even a lottery with small but positive expected return if it involves gains and losses both (i.e., winning 50 euros with probability 1/2 and losing 35 euros with probability 1/2). However, in our question the individuals can choose the size of the potential loss. Hence, even a loss-averse individual will be willing to pay *some* positive price to purchase the security. For instance, assume that utility is linear over gains and losses and that the loss aversion parameter is within the realistic interval of 2.5 and 5 (Tversky and Kahneman 1992). Then the willingness to pay for the risky security, when assessed in isolation from other risks (and thus assuming complete narrow framing), would be between 1,000 and 2,000 euros. More generally, even behavioral models have a hard time explaining why individuals reject such lotteries (Barberis, Huang, and Thaler 2006).

questions involving risky prospects. Therefore, in our estimates we need to control for the possibility that nonresponses may induce selection bias.

A perhaps more serious concern is that heterogeneity in cognitive ability may bias the willingness to pay of those who report a positive price for the security. There is some evidence that risk attitudes correlate with cognitive ability. For instance, Benjamin, Brown, and Shapiro (2006) find—in an experiment involving 90 Chilean students asked to choose between one small safe prospect and several small risky prospects with increasing expected return—individuals with lower cognitive ability are more likely to turn down the risky prospect. This contradicts the risk-neutral property of expected utility preferences with respect to small, favorable bets. In a related study involving several different samples of students, Shane (2005) finds that those who perform better in a "cognitive reflection test" tend to be systematically more willing to take risks. It is unclear whether risk preferences drive cognitive ability or vice versa, but the positive correlation between risk tolerance and (unobserved in our sample) cognitive ability may be responsible for a spurious correlation between risk tolerance and consumer endowment, since the latter is typically correlated with ability. This may lead to the conclusion that absolute risk aversion is decreasing when, in fact, it is not. We shall deal with this problem in Section 6 by relying on instrumental variables.

The subsamples of risk-loving and risk-neutral consumers on the one hand and risk-averse consumers on the other exhibit several interesting differences. The risk averse are younger and less educated; they are less likely to be male, to be married, or to live in the North of Italy. Strong differences also emerge comparing the type of occupation: only 17.6% of the risk averse are self-employed, compared to 30% among the risk prone or risk neutral. But in the public sector, 27.9% of employees are risk loving or risk neutral while 28.3% are risk averse. These differences are likely to reflect self-selection, with more risk-averse individuals choosing safer jobs. Finally, notice that risk-averse consumers are significantly less wealthy than the non–risk-averse consumers (respectively 31,500 and 56,800 euros cash on hand).

Table 2 reports also the characteristics of consumers who are modestly risk averse (at or above the sample median of the reported price Z_i) and of those who are highly risk averse (below median). Highly risk-averse consumers are on average two years older, somewhat less well educated, less likely to be married, and much more likely to live in the South. They are also less wealthy than the modestly risk averse in terms of income, consumption, and cash on hand (respective medians of 28,200 and 35,800 euros cash on hand). Finally, 15.9% (resp.20.2%) of the highly (resp. modestly) risk averse are self-employed, while 29.3% (26.8%) of the highly (modestly) risk averse are public-sector employees.

Thus, an individuals's degree of risk aversion could well explain the risk level of his occupation.

One way to assess whether our measure of risk aversion reflects mainly noise or instead reveals individual risk attitudes is to check whether it has predictive power over consumers' choices that theory suggests should be affected by individuals' risk aversion. In Table 3, the second column shows the sample correlations between measured risk aversion and nine indicators of risky choices: a dummy for self-employment and one for being an entrepreneur; indicators for being a public employee, ownership and portfolio share of risky assets, and investment in education (measured by the number of years of schooling); a dummy for whether a person has moved from his region of birth and one for whether he tends to change jobs; and an indicator for chronic disease. These are the variables on which Guiso and Paiella (2006) focus in assessing the predictive power of a risk aversion measure. Based on theory, one expects a negative correlation between risk aversion and self-employment, entrepreneurship, risky asset ownership and share, investment in education (a risky endeavor), being a mover, having a high propensity to change jobs, and incurring a chronic disease; conversely, there should be a positive correlation between risk aversion and being a public employee. The signs of the unconditional correlations are generally consistent with the priors but differ in size depending on the particular behavior.

Correlations are high (above 0.1) for being self-employed, for risky asset ownership and share, and for investment in education; they are lower for the other indicators.

The third and fourth columns of Table 3 show the economic effects of increasing risk aversion from the 10th to the 90th percentile. Values are computed from controlled regressions on the sample of risk-averse individuals by using the same specification as in Guiso and Paiella (2006).[13] They reveal that our survey measure of risk aversion has considerable predictive power on such behaviors as occupation choice: moving from the 10th to the 90th percentile of the distribution of risk aversion, lowers the probability of being self-employed by 19% of the sample mean, that of being an entrepreneur by 13%, of holding risky assets by 22% and their portfolio share by 46%, and of time invested in education by 7%. Based on this evidence we feel confident that, despite the extent of nonresponses, our risk attitude indicator captures the respondents' willingness to bear risk. This is not to say that our measure of risk aversion is free of measurement error. We defer to Section 5 a discussion of the consequences of measurement error in the reported price for our estimates.

Finally, one could object that the question asked in our survey might

---

[13]The estimates of the effects are slightly different because in the previous paper we used an alternate methodology to map the reported willingness to pay into the measure of absolute risk aversion.

have been variously interpreted by different respondents. In particular, some may have interpreted the "gain" of 5,000 euros in the question as a "get" of 5,000 euros, in which case the assumed payoff from the security was $5000 - Z$ instead of 5000. Under the "get" interpretation, the expected utility from buying the security would be $(1/2)u_i(w_i + 5000 - Z_i) + (1/2)u_i(w_i - Z_i)$ instead of $(1/2)u_i(w_i + 5000) + (1/2)u_i(w_i - Z_i)$, and the expected value of the security would be $(5000 - Z_i)/2$ instead of $1/2(5000 - Z_i)$. As a consequence, more respondents would be classified as risk lovers and so drop out of our regression sample, since we focus on the risk averse. Under the "get" interpretation, our sample of risk-averse consumers would drop to 2,533 observations from 3,157. The card that respondents were shown when asked the question, shown in Figure 2, strongly suggests that the "gain" interpretation is the correct one; furthermore, interviewers were verbally instructed on the correct interpretation of the question. But because we cannot be completely sure that respondents actually interpreted the question as expected by the questionnaire designers, we also test for results under the alternative interpretation. Obviously, one cannot completely rule out that different households interpreted the question in different ways.

4 Empirical specification

Most of the literature assumes that agents are risk averse and is interested in assessing how risk aversion varies with consumers' attributes and in particular with their endowments. Accordingly, the rest of this paper focuses on risk-averse individuals.

To estimate the relation between our index of absolute risk aversion and individual endowment, we use the following specification (the household index i is omitted for brevity):

$$R(w) = \frac{ae^{\gamma H + \eta}}{w^\beta} = \frac{\kappa}{w^\beta}, \tag{6}$$

where w denotes the (lifetime) endowment, H is a vector of consumer characteristics affecting individual preferences, η is a random shock to preferences, and a, γ, β are unknown parameters.[14] Notice that $R(\cdot)$ is always positive and is decreasing in w for all positive values of β. Furthermore, if $\beta > 0$ then $R(\cdot)$ is always convex in w. Though simple, this formulation is flexible enough to allow us to analyze the curvature of absolute risk tolerance, which is defined as

$$T(w) = \kappa^{-1} w^\beta. \tag{7}$$

Thus, if $\beta > 0$ then risk tolerance is an increasing function of w; furthermore,

---

[14]Notice that our empirical specification (6) does not allow for heterogeneity in the β- parameter. If β varies across individuals then our estimates would be affected by heteroskedasticity. However, a formal test cannot reject the null hypothesis that the error term is homoskedastic.

it will be concave, linear, or convex in w depending on whether β is less than, equal to, or greater than 1. Because β measures the speed at which $R(\cdot)$ declines with endowment, it follows that $T(\cdot)$ is a concave (resp. convex) function of w if absolute risk aversion falls as consumption increases at a speed lower (resp. greater) than 1, which is the value characterizing CRRA preferences. Since most theoretical ambiguities rest on the curvature of T, not R, our approach is not restrictive.

Although equation (6) is assumed, a utility function that gives rise to a measure of absolute risk aversion as in (6) is

$$u(w) = \int e^{-\frac{ae^{\gamma H + \eta}\widetilde{w}^{1-\beta}}{1-\beta}}d\widetilde{w} = \int e^{-\frac{\kappa \widetilde{w}^{1-\beta}}{1-\beta}}d\widetilde{w}, \qquad (8)$$

which converges to the CRRA utility $u(w) = w^{1-\kappa}/(1-\kappa)$ as β tends to 1. If β = 1, then $\kappa = ae^{\gamma H + \eta}$ measures relative risk aversion.

Taking logs on both sides of (7), our empirical specification becomes

$$\log(T) = -\log \kappa + \beta \log w = -\log a - \gamma H + \beta \log w - \eta. \qquad (9)$$

The curvature of absolute risk tolerance—as well as the relation between absolute risk aversion and endowment—is thus parameterized by the value of β. We focus our discussion on risk tolerance, rather than risk aversion, because the former aggregates cleanly in the presence of heterogeneity (as shown by Breeden 1979).

As pointed out previously, if background risk $\widetilde{y}$ is present then our measure must be interpreted as measuring the risk aversion of the indirect lifetime

utility function $v(w) = Eu(w + \widetilde{y})$. The question that arises is whether we can draw implications for the relation between the risk aversion of $u(\cdot)$ and the level of the endowment, on the one hand, from the relation between the risk aversion of $v(\cdot)$ and the endowment, on the other.[15] In the Appendix we show that taking a second-order Taylor expansion of the indirect utility function around w yields the following index of the absolute risk aversion of this approximated utility

$$R_v(w, s) = \kappa w^{-\beta} \left[\frac{1 + p_u t_u s^2/2}{1 + p_u r_u s^2/2} \right];$$

here κ is defined as before, s is the coefficient of variation of the consumer's endowment, and r_u, p_u, and t_u denote (respectively) the degrees of relative risk aversion, relative prudence, and relative tolerance of the utility function $u(\cdot)$.[16] Notice that $\kappa w^{-\beta}$ is the absolute risk aversion of $u(\cdot)$ and that $R_v(w, s) > \kappa w^{-\beta}$ if, given $s > 0$ and assuming the consumer is prudent (i.e., $p_u > 0$), relative risk tolerance is larger than relative risk aversion. Furthermore, when $t_u > r_u$, the term in square brackets is increasing in s and $R_v(\cdot)$ is also increasing in s. Taking logs of the inverse of $R_v(\cdot)$ and using the relations between r_u, p_u, and t_u spelled out in the Appendix, our empirical specification for risk tolerance when

---

[15]The indirect utility function inherits several properties of $u(\cdot)$. In particular, if $u(\cdot)$ is characterized by decreasing absolute risk aversion (DARA) then so is $v(\cdot)$. Furthermore, as shown by Kihlstrom, Romer, and Williams (1981), comparative risk aversion is preserved by the indirect utility if $u(\cdot)$ exhibits nonincreasing risk aversion.

[16]See the Appendix for definitions of relative prudence and tolerance.

there is background risk becomes

$$\log(T_v) = -\log \kappa + \beta \log w - \delta s^2, \tag{10}$$

where $\delta = \beta p_u$. This formulation allows us to test directly whether background risk affects risk attitudes. It requires two conditions to hold: consumers must be prudent ($p_u > 0$) and risk aversion must be decreasing ($\beta > 0$).

5 Results

Table 4 shows the results of our estimation of (10) using different measures of consumer resources. The analysis is conducted on the sample of risk-averse consumers. We control for sample selection related to nonresponse by including among the regressors the Mills ratio based on the probit model for the probability of responding to the survey question (see columns (3) and (4) of Table A.1), which includes among the regressors only variables that are expected to be exogenous with respect to the individual's attitude toward risk.

Because our measure of risk tolerance is best interpreted as the risk tolerance of the consumer's value function, estimating the value of β requires information on the value of consumer lifetime endowment—which is typically nonobservable. To overcome this problem we use household consumption, which is readily available from the SHIW. In a life cycle/permanent income context, consumption expenditure is a sufficient statistic for lifetime resources as perceived by the

30

consumer; hence, it is the best "guesstimate" of unobservable lifetime endowment. However, we also check our results using accumulated financial assets and human wealth, as measured by income, as proxies for the lifetime endowment.

If we assume that consumption is proportional to the endowment w (i.e., $c = \lambda w$), our empirical specification becomes

$$\log(T_v) = -\log \kappa' + \beta \log c - \delta s^2, \tag{11}$$

where $\kappa' = \kappa \lambda^\beta$. In the first three columns of Table 4 we report estimates when absolute risk aversion is computed from equation (2) using the exponential utility function. In the first column we regress $\log(T_i)$ only on (log) nondurable expenditure and do not include any consumer characteristics that can proxy for differences in tastes. The estimate of β is 0.673 and is highly statistically significant, leading to a strong rejection of preferences with CARA. The estimated value of β implies that absolute risk aversion declines with endowment but at a rate slower than that implied by constant relative risk aversion preferences. In fact, the hypothesis that $\beta = 1$ is formally rejected ($p = 0.0000$), suggesting that absolute risk tolerance is a concave function of consumer resources.

In the second column of the table we include a set of strictly exogenous individual characteristics—such as age, gender, junior education attainment (a dummy equal to 1 if the household head has completed eighth grade)—as well as dummies for the presence of siblings and for region of birth. If tastes are

impressed in our chromosomes or evolve over life in a systematic way or depend on one's education[17] or are affected by the culture of the place of birth or by the possibility of relying on the support of a brother or sister, then these variables should have predictive power. In fact, the analysis shows that only education does, with risk aversion being higher among the least educated. Furthermore, a test of the hypothesis that the coefficients for age, gender, education, and siblings are jointly equal to zero can hardly be rejected at the standard levels of significance ($p = 0.0881$). In contrast, the 19 regional dummies[18] included in the regression, capturing the region of birth, are jointly significant (see the bottom of Table 4). Furthermore, the coefficients for these dummies (not shown) reveal a pattern: compared to those born in the central and southern regions, consumers born in the North are somewhat less risk averse. One possible interpretation is that the dummies are capturing regional differences in culture, which are transmitted with the upbringing. In addition to these variables, we insert in the regression two dummies for the occupation of the father of the household head:

---

[17]Education can depend to some extent on the individual's attitude toward risk, especially when it comes to higher education enrollment. Hence, in the regressions we control only for eighth-grade attainment, which can be considered exogenous and determined by factors are independent of individual attitudes toward risk.

[18]The Italian territory is divided into 20 regions and 95 provinces; the latter correspond broadly to U.S. counties. We will use the provincial partitioning in Section 7, where we look at the effect of background risk and liquidity constraints on risk aversion.

the first dummy is equal to 1 if the consumer's father is/was self-employed (0 otherwise); the second dummy is equal to 1 if he is/was a public-sector employee (0 otherwise). This allows us to test whether parents' attitudes towards risk—as reflected in their choise of occupation—are transmitted to their children. The estimates show that none of these variables has a significant effect on the degree of risk aversion. However, although the signs on the dummies turn out as expected and imply relatively lower (higher) risk tolerance for those individuals whose father was a public employee (self-employed), which can be viewed as being indicative of greater (lower) aversion to risk.

The third column of the table reports results using cash on hands as an alternative proxy for the endowment. Cash on hand is defined as the sum of financial assets plus household income (excluding asset income). The basic findings are confirmed: absolute risk tolerance is an increasing and concave function of the endowment (though the coefficient is somewhat smaller than when consumption is used) and CARA and CRRA preferences are rejected.

In all cases most of the variance of observed risk tolerance is left unexplained, as shown by the low R^2 values, suggesting that most of the taste heterogeneity across consumers cannot be accounted for by the set of exogenous variables that we observe. The estimated relation between absolute risk tolerance and consumer resources is consistent with Arrow's (1965) hypothesis that

absolute risk aversion should decrease as endowment increases whereas relative risk aversion should increase. Yet the latter claim is consistent with the former only if the wealth elasticity of absolute risk aversion is less than unity, as our findings thus far indicate. Our estimated elasticity is also consistent with the evidence of Holt and Laury (2002), who find—in a set of experiments where individuals choose among risky prospects of different scale—that individuals are more risk averse when facing larger scale prospects. This finding is inconsistent with constant relative risk aversion and implies increasing relative risk aversion.[19]

The remaining two columns show estimates when absolute risk aversion is computed using equation (3), that is, the isoelastic utility specification. Given the high correlation between this measure of absolute risk aversion and the one based on equation (2), results are almost identical.

The Mills ratio, which was included in all regressions to correct for any selection bias due to systematic nonresponse, has a small insignificant coefficient. This suggests that self-selection is unlikely to be an issue and so lack of a control should not bias the estimated β.

---

[19]Holt and Laury (2002) show that behavior in their experiment is well represented by a utility function of the form $U(w) = 1 - \exp\left\{-\alpha w^{1-\beta}\right\}/\alpha$, that mixes CARA and CRRA preferences. The absolute risk aversion coefficient is given by $\beta/w + \alpha(1 - \beta)/w^{\beta}$, which is increasing in w for $0 < \beta < 1$. The authors argue that the estimate $\beta = 0.269$ is consistent with the behavior they observe.

The estimates of β presented in the table are subject to two caveats. First of all, possible misinterpretations of the survey question, as well as difficulties in figuring out the maximum price to be paid, suggest that the $\log(T_i)$ variable on the left-hand side (LHS) is likely to suffer measurement error. Since the mismeasured variable is the willingness to pay for the security, of which T_i is a nonlinear function, we have by no means escaped the case where a (classical) error-ridden left-hand-side variable affects the consistency of the estimates. In the Appendix, we derive analytically the implications of measurement error for our estimates and determine the conditions under which an error (of the classical type) in the willingness to pay would have negligible effects. We then argue that these conditions are met in our data, which makes measurement error in Z an unlikely source of serious bias for our estimates.

The second issue that deserves notice concerns the results that have been obtained assuming that consumption is proportional to endowment, so that the marginal propensity to consume out of wealth is constant. A large literature has argued that, in the presence of uncertainty and with prudent behavior, the marginal propensity to consume is large for low values of wealth and tends to the (constant) perfect foresight value as wealth increases (see Carroll and Kimball 1996). If this is true then consumption is a concave function of endowment, implying that the estimated value of β in equation (11) reflects not only the

35

elasticity of the risk aversion of lifetime utility to w but also the elasticity of consumption to w. It is easy to show that in this case our estimate of β is larger (in absolute terms) than the true value,[20] implying that, if anything, risk aversion declines with endowment at a lower speed than we estimate. Thus, without knowledge of the consumption function, we can establish only an upper bound on the value of β.

6 Robustness

6.1 Allowing for a different interpretation of the security question

The foregoing results are obtained assuming that individuals interpret the question as meaning that they "gain" 5,000 euros with probability 1/2 if they purchase the security. In Table 5 we check whether our results are affected if we assume instead that individuals have interpreted the question as meaning that they "get" 5,000 euros with probability 1/2 and so their gain is $5,000 - Z$. Since

---

[20]To see this, let $c = c(w)w$, which implies that $w = c/c(w)$. Hence $\beta \log w = \beta(\log c - \log c(w))$. We can then treat this as an "omitted variable" problem because it is as if we did not include $\beta \log c(w)$ in the regression. The estimated coefficient on $\log c$ will be given by $\widetilde{\beta} = \beta[1 - \text{Cov}(\log c, \log c(w))/\text{Var}(\log c)]$. Since the covariance is negative (the larger the level of consumption c, the smaller the propensity to consume the last unit of wealth) it follows that, $\widetilde{\beta} > \beta$.

a larger number of individuals are now classified as being risk lovers or risk neutral, our sample size decreases to 2,533 observations. However, even with this large drop in the number of observations, the parameter β in columns (1)–(5) is quite similar to the one we obtain in Table 4 and is estimated with similar precision. Conclusions about the other regressors also remain unchanged.

6.2 Endogenous consumption and wealth

The results we have reported thus far do not take into account that consumption and wealth are endogenous variables affected by consumer preferences. As a result, the estimated coefficients are potentially affected by endogeneity bias. However, the direction of that bias is not clear a priori. If more risk-averse individuals choose safer but less rewarding prospects, they may end up being poorer and so consume less than the less risk-averse. This would tend to overstate the positive relation between risk tolerance and wealth. However, if the more risk averse are also more prudent then, ceteris paribus, they will compress current consumption, save more, and end up accumulating more assets.[21] In this case our estimates of the relationship between risk tolerance and wealth

---

[21] Risk aversion and prudence usually go together. If the utility function is exponential, absolute risk aversion and prudence are measured by the same parameter; if it is CRRA, absolute prudence is equal to absolute risk aversion plus $1/c$; if preferences are described by equation (8), absolute prudence is equal to absolute risk aversion plus β/c.

will be biased toward zero, which could partly explain why these estimates show risk tolerance increasing less than proportionately with wealth.[22] Another potential source of spurious correlation between risk tolerance and the endowment is unobserved heterogeneity in cognitive ability. A positive correlation between cognitive ability and wealth may lead to the conclusion that risk tolerance is increasing with wealth when in fact it is not. However, because we find that risk tolerance increases with wealth at a speed that is lower than that implied by CRRA preferences, accounting for any cognitive effects would only strengthen our conclusion.

To address these issues we re-estimate equation (10) with instrumental variables (IV). Finding appropriate instruments for consumption and wealth is no easy task. We rely on the following sets of instruments. First, we use characteristics of the household head's father—namely, his education and year of birth—on the grounds that wealth is likely to be correlated with that of one's family as proxied by the father's education and cohort. Second, we employ mea-

---

[22]Another possible explanation of our results is the presence of measurement error in consumption or wealth; this could be a source of attenuation bias. To verify whether attenuation bias is an issue, we instrumented consumption with total wealth, liquid assets, and cash on hand and obtained an estimate of β equal to 0.75 (standard error of 0.31). Instrumenting cash on hand with total wealth and consumption yields an estimate of β equal to 0.56 (standard error of 0.23). These results suggest that attenuation bias due to erroneously measured endowments is unlikely to change our conclusions.

sures of windfall gains, such as a dummy for the house being acquired as a result of a bequest or gift, the value of insurance settlements and other transfers received, and an estimate of the capital gain on the house since it was acquired. The capital gain estimate is also interacted with three dummies for the size of the town of residence. Finally, we include in the instrument set a third-order polynomial in the age of the household head (to capture life cycle wealth effects), the interaction between age and gender, and three dummies for the size of the hometown. Overall, the instruments explain about 30% of the variance of (log) nondurable consumption and of (log) cash on hand.

Table 6 shows the results of IV estimation. We report the specification including (on the right-hand side (RHS) age, gender, education, siblings, occupation of household head's father, and region of birth. For some agents the information on some instruments is missing, so the sample size is smaller by about 30 observations with respect to ordinary least squares (OLS) analysis. Overall, the IV estimates of the parameter β are much like the OLS estimates. For instance, when consumption is used the IV estimated β is 0.606, just slightly smaller than the OLS estimate; when using cash on hand, the IV estimate of 0.667 is somewhat larger than the OLS estimate of 0.618 suggesting that neither reverse causality nor unobserved cognitive ability are likely to be driving the results. The main difference is that the IV estimates are much less precise:

the hypothesis that $\beta = 1$ cannot be rejected at standard levels of significance. However, the probability that β is larger than 0 and smaller than 1 exceeds 90%, as shown at the bottom of the table. Overall, then, even the IV estimates imply that absolute risk aversion is a decreasing function of wealth, rejecting CARA preferences, and suggest that preferences may deviate also from a CRRA representation. Figure 3 shows the risk tolerance–consumption relation when the OLS and IV estimators are used. In both cases the profile is far from being linear, as would be the case under CRRA.

The last two columns of Table 6 repeat the estimates while excluding from the sample individuals aged over 75, those with nonpositive financial assets or nonpositive income, and those with consumption or cash on hand below the 10th percentile; we thus ensure that the slope of the risk tolerance/endowment relation is not driven by some poor individuals reporting abnormally low willingness to pay. Results are robust to this choice of sample.

7   Risk aversion and background risk

In a world of incomplete markets, attitude toward risk may vary between consumers not only because of differences in taste parameters but also because consumers face different environments. In Section 2 we discussed how risk aversion can be affected by background risk. In this section we test whether attitudes

toward risk are affected by the presence of uninsurable, independent risks and by the possibility of being liquidity constrained in the future. To measure background risk, we rely on per capita gross domestic product (GDP) growth at the provincial level for the period 1952–1992, which we use to compute a measure of the variability of GDP growth in the province of residence. For each province we regress (log) GDP on a time trend and compute the residuals. We then calculate the variance of the residuals and assign this estimate to all households living in the same province. The main advantage of this variable compared with subjective measures of future income uncertainty, such as those analyzed by Guiso, Jappelli, and Pistaferri (2002), is that it is likely to be truly exogenous and thus less subject to the self-selection problems that affect subjective measures reflecting occupational choice.[23] The variance of provincial GDP growth is an estimate

---

[23] The 1995 SHIW contains a special section in which households are asked a set of questions designed to elicit the perceived probability of being employed over the twelve months following the interview and the variation in earnings if employed. Guiso, Jappelli, and Pistaferri (2002) use these data to obtain an estimate of expected earnings and their variance. They show that the subjective variance is negatively correlated with a dummy for risk aversion and interpret this as evidence of self-selection. When we use the subjective measure of earnings uncertainty as a proxy for background risk, we find that its effect on risk tolerance is positive, small, and not statistically significant. This can be read as evidence that self-selection nullifies any background risk effect and that subjective measures are inadequate for isolating the effect of background risk on an individual's willingness to bear risk.

of aggregate risk and should be largely exogenous to the individual risk attitude (unless risk-averse consumers move to provinces with low-variance GDP).

Table 7 reports the estimation results using consumption as a proxy for the endowment and instrumenting it as in Table 6; results are similar if cash on hand is used. Standard errors allow for clustering effects. The first column shows the estimates when this proxy for background risk is included in the specification. Adding this variable raises somewhat the estimated elasticity of absolute risk tolerance to the endowment, although the elasticity remains below 1. The degree of risk tolerance is decreasing in the variance of per capita GDP in the province of residence even after controlling for age, gender, education, siblings, occupation of household head's father, and region of birth; moreover, the effect is highly statistically significant. This is consistent with background risk models. Increasing our measure of background risk by a single standard deviation lowers absolute risk tolerance by about 19%. If risk-averse individuals tend to move from high-variance to low-variance provinces, then this would tend to generate a positive correlation between risk tolerance and background risk. Therefore, the estimates reported in Table 7 constitute, if anything, a lower bound (taken in absolute terms) of the true effect of background risk.

The second column of Table 7 adds to our basic specification an indicator of liquidity constraints. As argued by Gollier (2000), liquidity constraints act

to reduce the consumer horizon, thus limiting opportunities to "time diversify" any risk currently taken and accentuating risk aversion. Our measure of risk aversion is based on the notion that what matters is the threat of liquidity constraints driven by (a) the inability to borrow and (b) impediments to drawing down accumulated assets in order to increase consumption.[24] Thus, we define as *liquidity constrained* those who have been refused credit; have not asked believing their request would be turned down (2% of the sample); have a ratio of liabilities to total assets above 25% (6% of the sample); or have financial assets amounting to less than 1% of their net worth (18% of the sample). Overall, we classify as potentially liquidity constrained about 26% of the households in our sample of respondents. The estimated coefficient of this indicator has the expected negative sign and is statistically significant, which implies that being liquidity constrained (or risking to become such) makes agents more risk averse. Economically, being liquidity constrained lowers risk tolerance by 4.4%. This result is robust to the inclusion in the regression of the variance of provincial GDP, as shown in the last column of Table 7, implying that background risk is not proxying for liquidity

---

[24]If households have low liquid assets, it can be relatively difficult to smooth consumption when confronted with unexpected negative income or expenditure shocks. If liquid assets are low then, in order to increase consumption, households might choose the potentially costly options of tapping into their illiquid assets or asking for credit. Besides, if they are already heavily indebted then the cost of additional credit might well be high. Hurst and Stafford (2004) provide further arguments for using this definition of liquidity constraints for households.

constraints and that these two variables exert an independent role on individuals' willingness to bear risk.

8 Consistency with observed behavior

If our findings on the relation between risk tolerance and wealth do reflect the structure of individual preferences, then this should show up in actual behavior; that is, observed behavior should be consistent with the shape of the measured risk tolerance–wealth relation. Here we discuss some implications of our empirical characterization of this relation. First, if relative risk tolerance is decreasing in wealth, as implied by our findings, then the portfolio share of risky assets should decline as wealth increases. Second, if absolute risk tolerance is a concave function of the consumer endowment (as our results suggest), then the portfolio share of risky assets should be an increasing function of age. Third, if our estimates do indeed identify the utility function parameters, then they should be coherent with those based on the estimation of Euler equations for consumption.

8.1 Wealth–portfolio relation

The first implication is clearly contradicted by the data, since portfolio shares are found to be an increasing function of wealth. This is obviously in contrast also with constant relative risk aversion preferences. One strategy that has been

pursued in the literature is to maintain the CRRA characterization of the utility of consumption and to assume that wealth enters the utility function directly as a luxury good—for instance, through a joy-of-giving/bequest motive. As Carroll (2000) shows, the implication is that a larger proportion of lifetime wealth will be devoted to risky assets. Clearly, this mechanism can still explain the data even if the utility function exhibits increasing relative risk aversion (IRRA), provided that the joy-of-giving motive is sufficiently strong. Another explanation is that some portfolio management costs decline with the size of the investment in risky assets (which is increasing in wealth). If such costs are significant, this mechanism could overturn any IRRA-based incentive to lower the portfolio share of risky assets. A third possibility, analyzed by Peress (2004), is that households face costs of acquiring financial information. Since wealthier consumers tend to invest larger amounts in stocks, they have more incentives to invest in information acquisition. Being more informed, they tend to invest a larger share than do less-informed consumers, and this mechanism also could counteract the tendency of the portfolio share to decrease via IRRA. Guiso and Jappelli (2006) show empirically that, indeed, wealthier consumers collect more financial information and that those who invest more in collecting information have larger shares of stocks in their portfolio. Thus, our results do not conflict with the evidence.

8.2 Age/portfolio profile

To check the second implication—that with concave risk tolerance the age/portfolio profile is upward sloping—we run Tobit regressions of the portfolio share of risky assets on wealth (linear and squared), age, and a set of controls that include city size, household size, gender, region of residence, education of household head, and so on. We exclude households with zero wealth and those whose head is aged more than 60, since the elderly may have various incentives to decumulate assets—especially the reskier ones—after retirement.[25]

After these exclusions, our sample includes 4,799 households. Table 8 shows the results of the estimates separately when risky assets are divided by total wealth and by financial wealth. The first two columns use the entire sample

---

[25] As pointed out by Hurd (2000), the portfolio behavior of retired, elderly consumers may be quite different from that of the nonretired. First, the retired have a limited ability to return to the labor force and to use this possibility as a buffer against financial losses. This limitation should reduce their willingness to hold risky assets. Second, the elderly face substantial mortality risk, which increases sharply at advanced old age and leads to a decline in consumption and wealth; this, in turn, may affect portfolio composition. Third, retired consumers have large annuity income flows, and the risks associated with those flows are quite different from the risks of earnings. Finally, the elderly face a much a higher risk of healthcare consumption than the nonelderly, which discourages holdings of risky assets. Consistent with these observations, Guiso and Jappelli (2000) find that the portfolio share of risky assets declines with age after retirement.

of consumers aged less than 61. The last two columns check the results on the subsample of consumers who responded, to the security question (we control for risk aversion). All the estimates show that the share of risky assets is mildly increasing in age: the portfolio share increases by 2 percentage points for a 10-year increase in age, which is consistent with our empirical characterization of absolute risk aversion. With CRRA, the share of risky assets is independent of the investor's horizon and hence of age.

8.3   Euler equation estimates

To further check the consistency of our results with observed behavior, we followed a third route. We estimated the value of the parameter β using an Euler equation for consumption, assuming that the utility function has the form given by equation (8) and that consumption is proportional to lifetime wealth.[26] In this case, the risk aversion of the period utility is proportional to that of lifetime utility.[27]

---

[26]We thank Pierre Andrè Chiappori for suggesting this test.

[27]As shown by Tiseno (2002), the risk aversion of lifetime utility, $R_U(w)$, is related to the risk aversion of period utility, $R_u(c(w))$, by the identity

$$R_U(w) = R_u(c(w))c_w,$$

where c_w is the marginal propensity to consume out of lifetime endowment. If the latter is constant in w then, up to a rescaling factor, the risk aversion of the period utility is the same as the risk aversion of lifetime utility.

Suppose there is no background risk. Using equation (8), the Euler equation for consumption is

$$\exp\left\{-\kappa\frac{c_t^{1-\beta}}{1-\beta}\right\} = \gamma(1 + r_{t+1})\exp\left\{-\kappa\frac{c_{t+1}^{1-\beta}}{1-\beta}\right\} + \varsigma_{t+1}; \qquad (12)$$

here γ is the subjective discount factor and ς_{t+1} is an expectation error, which is orthogonal to all variables in the information sets of the agents at time t. Thus, we can use these orthogonality conditions and estimate the parameters of equation (13) using a generalized "method of moment" estimator. If we are uncovering true preference parameters, then we should obtain values of β and of the implied degree of risk aversion that are similar to those estimated in previous sections.

To estimate (12), we rely on the panel component of the SHIW and pool together the observations for the years 1989, 1991, 1993, and 1995. The panel component is such that half of the sample in a given survey is re-interviewed in the subsequent one. Thus, 4 is the maximum number of time periods that a household can be present. Because our estimator is consistent only when the number of observations along the time dimension is large, these results should be regarded as being merely suggestive. We restrict our attention to households that in 1995 answered our survey question on the risky security. Furthermore, since the Euler equation is known not to hold when credit markets are imperfect, we have excluded households that are likely to be liquidity constrained (as defined

in Section 7). Our estimation sample consists of an unbalanced panel of about 2,400 observations on about 1,300 households. As a measure of consumption we use household real expenditure on nondurable goods and services, adjusted for household size.[28] As a measure of the interest rate we use the return on bank deposits and checking accounts, which varies over time and across Italian provinces, and impute to each household in a given year and living in a given province the average rate prevailing in that year and province. To account for demographics, we let the parameter κ be a function of the age and gender of the household head. As instruments in the agents' time-t information set, we use the interest rate lagged one period, a categorical variable for the size of bank deposits, a dummy for the ownership of saving accounts, the Herfindahl index of bank concentration in the province, a categorical variable for the size of the town where the household lives, the population of the province in which the town is located, and year dummies. Results (not shown) give a point estimate of the parameter β of 1.1999 with a standard error of 0.7423, which implies a range of values that is reasonably close to the estimates that we have obtained from the relationship between our measures of absolute risk aversion and consumer endowment.

Overall, though we should be cautious in drawing any strong conclusions

---

[28] Household per-adult equivalent expenditure is obtained as follows: the household head is weighted 1, other adults in the household are weighted 0.8, and children are weighted 0.4.

from the evidence in this section, we take it as suggesting that our estimates of the relation between risk aversion and wealth are not contradicted by observed behavior.

9   Conclusions

In this paper we construct a direct measure of absolute risk aversion using the 1995 Bank of Italy Survey of Household Income and Wealth. The measure is based on a simple yet powerful question on the maximum price a consumer is willing to pay to buy a security, a question whose answer allows us to derive an estimate of the degree of absolute risk aversion for each individual in the sample. This estimate is then used to gather direct evidence on the nature of the relationship between individual risk predisposition (on the one hand) and (on the other hand) individual endowment, demographic characteristics, and measures of uninsurable risk exposure and liquidity.

Our findings suggest that, among risk-averse consumers, the degree of absolute risk aversion is decreasing in individual endowment—thus strongly rejecting CARA preferences—but the point estimate of risk aversion's elasticity with respect to consumption is about 0.65, below the unitary value predicted by CRRA utility. Consequently, absolute risk tolerance is a concave function of consumer endowment. How reasonable is this finding?

One way to answer this question is to run the following experiment. Suppose that a consumer with annual consumption of 10,000 euros ($8,000, roughly the 17th percentile) is willing to pay (at most) 500 euros to buy the security. Then, using equation (6), the implied value of her absolute risk aversion would be 0.001386. Suppose now that relative risk aversion is constant. In this case, if our consumer had an annual consumption of 75,000 euros (roughly 99th the percentile) then absolute risk aversion would be 0.000185 ($0.001386 \times 10000/75000$) and she would report a price of about 3,200 euros to acquire the security. This may seem an implausibly high figure: it is close to the gain offered by the security in the riskless state. Intuitively, CRRA implies that absolute risk tolerance increases "too fast" with the endowment. If instead absolute risk tolerance increases with endowment at the speed implied by our estimates, then the price that the richer consumer is willing to pay for the security would be about 1,800 euros—a figure that seems much more plausible.

As argued previously, our findings are also consistent with the empirical evidence that young households take on relatively less portfolio risk than more mature households. In fact, according to Gollier and Zeckhauser (2002), the concavity of absolute risk tolerance is a necessary and sufficient condition for such behavior to be optimal.

Individual risk aversion appears also to be characterized by a substantial

amount of unexplained heterogeneity. Consumer attributes and demographic characteristics are of little help in predicting degree of risk aversion. The only exceptions are education and the region of birth; the latter is likely to capture regional differences in risk predisposition and culture that are transmitted with upbringing within the family. Insofar as these differences reflect fundamental heterogeneity in risk preferences, our findings add to the skepticism about "representative agent" models in contexts where risky decisions are involved.

At a more general level, our findings imply that individuals sort themselves out in such a way that the highly risk averse face less risky prospects. This self-selection makes it problematic to assess the effect of risk on choice—an issue that arises, for instance, when evaluating the effect of income uncertainty on investment in risky assets or when testing for precautionary savings (Fuchs-Schündeln and Schündeln 2005). When risk preferences are observed directly, as here, one can assess the importance of self-selection for estimates of the effect of risk on behavior (see Guiso and Paiella 2006) for an application to precautionary saving).

In a world of incomplete markets, individual attitudes toward risk may vary across households because of differences not only in tastes but also in the environment. We address this issue by analyzing the impact that income uncertainty and borrowing constraints have on the degree of risk tolerance. We find

unequivocal evidence that background risk and borrowing restrictions shape consumer attitudes about accepting risk. One important implication is that imperfections in financial markets may discourage entrepreneurship and investment because they limit access to external finance and also, more directly, because they decrease the willingness of individuals to bear risk. Establishing whether this is actually an important channel is an interesting topic for future research.

A APPENDIX

A.1 The SHIW

The Bank of Italy Survey of Household Income and Wealth (SHIW) collects detailed data on demographics, household consumption, income, and balance-sheet items. The survey was first run in the mid-1960s but has been available on tape only since 1984. Over time it has undergone a number of changes in sample size and design, sampling methodology, and questionnaire form. However, sampling methodology, sample size, and the broad contents of the information collected have been unchanged since 1989. Each wave surveys a representative sample of the Italian resident population and covers about 8,000 households—although, at times, specific parts of the questionnaire are asked of only a random subsample. Sampling occurs in two stages, first at the municipality level and then at the household level. Municipalities are divided into 51 strata defined by 17 regions and 3 classes of population size (more than 40,000, 20,000 to 40,000, less than 20,000). Households, which are randomly selected from registry office records, are defined as groups of individuals related by blood, marriage, or adoption and sharing the same dwelling. The head of the household is conventionally identified with the husband, if present. If the person who would usually be considered the head of the household works abroad or was absent at the time of the interview, the head is taken to be the person responsible for managing the household's re-

sources. The net response rate (ratio of responses to households contacted net of ineligible units) was 57% in the 1995 wave. Brandolini and Cannari (1994) present a detailed discussion of sample design, attrition, and other measurement issues and also compare the SHIW variables with the corresponding aggregate quantities.

A.2 Definitions of the variables

In the empirical analysis, all demographic variables (age, education, gender, number of brothers and sisters, marital status, region of birth, occupation type, and sector) refer to the household head.

Consumption, net worth, and financial wealth. Consumption is the sum of the expenditure on food, entertainment, education, clothing, medical expenses, housing repairs and additions, and imputed rents. Net worth is the total of financial and real assets net of household debt. Financial wealth is given by the sum of cash balances, checking accounts, savings accounts, postal deposits, government paper, corporate bonds, mutual funds, and investment in fund units and stocks. Real assets include investment real estate, business wealth, primary residence, and the stock of durables.

Discouraged borrowers and rejected loan applicants. The following questions have been asked in each wave of the survey since 1989: "During the year did you or a member of the household think of applying for a loan or a mortgage

to a bank or other financial intermediary, but then changed your mind on the expectation that the application would be turned down?" Those answering Yes to this question are classified as discouraged borrowers. Those answering Yes to the following questions are classified as rejected borrowers:: "During the year did you or a member of the household apply for a loan or a mortgage to a bank or other financial intermediary and have it turned down?"

Education of the household head's father. This variable is originally coded as: no education (0); completed elementary school (5 years); completed junior high school (8 years); completed high school (13 years); completed university (18 years); graduate education (more than 20 years).

Education of the household head. This variable is originally coded exactly as the previous (father) variable.

Income. Total household after-tax income, excluding asset income.

Indicators of background risk. The variance of shocks to per capita GDP in the province of residence is obtained from time-series data on per capita GDP at the province level from 1952 to 1992. For each province we regress the logarithm of per capita GDP on a linear trend and compute the variance of the residuals from this regression. We then impute this variance to all households living in the same province.

Indicator of liquidity constraints. A dummy variable that allows for dis-

couraged borrowing and rejected loan applications and for low liquid assets or high indebtedness, which can prevent consumption smoothing in case of unexpected negative income shocks.

Risk aversion. The Arrow–Pratt measure of absolute risk aversion; the risk attitude indicators are obtained from responses to a survey's direct question concerning a hypothetical security. Each survey participant is offered a hypothetical security and is asked to report the maximum price that he would be willing to pay in order to buy it. The wording of the security question and the methodology implemented to compute risk aversion are described in the text.

Year of birth of the household head's father. This variable is used to define ten-year intervals starting from 1900. An additional interval is defined for those born in or after 1950. We then construct six indicators: the first is set equal to 1 if the household head's father was born between 1900 and 1909, the second indicator is set to 1 if he is born between 1910 and 1919, and so on.

Windfall gains measures. Six measures are used. The first is a dummy for home ownership as a result of gift or bequest. The second is the sum of the settlements received related to life (excluding annuities), health, casualty insurance. The third measure is the sum of severance payments, unemployment benefits, and redundancy allowance. The fourth is the sum of any additional financial aid from central or local governments, other public institutions, or charities. The

fifth consists of gifts/monetary contributions received from friends or family living outside the household dwelling. The last instrument is a measure of windfall gains (or losses) on housing constructed using time-series data on house prices at the province level over the years 1965–1994. For homeowners, we compute the house price change since the year the house was acquired (or since 1965 if it was acquired earlier). To non-homeowners, we assign the house price change since the year they started working (or since 1965). This can be justified on the grounds that non-homeowners start saving to buy a home as soon as they start working.

A.3   Risk aversion of the indirect utility function

Let $v(w) = Eu(w+\widetilde{y})$ denote the indirect utility function. Taking a second-order Taylor approximation of the RHS around the endowment w, we can approximate the indirect utility by

$$v(w) \simeq u(w) + u''(w)\sigma^2/2.$$

Using the first and second derivatives of this expression, we can write the degree of absolute risk aversion of $v(\cdot)$ as

$$R_v(w) = -\frac{v''(w)}{v'(w)} \simeq R_u(w)\left(\frac{1 + P_u T_u \sigma^2/2}{1 + P_u R_u \sigma^2/2}\right),$$

where $R_u(w) = -(u''/u')$, $P_u(w) = -(u'''/u'')$, and $T_u(w) = -(u^{iv}/u''')$ respectively denotes the degrees of absolute risk aversion, absolute prudence, and

absolute temperance with respect to the utility function $u(\cdot)$. From the previously displayed equation it is clear that $T_u > R_u$ is a sufficient condition for (a zero-mean) background risk to make a prudent consumer more risk averse.

Let s denote the coefficient of variation of the consumer endowment (i.e., $s = \sigma/w$) and let r_u, p_u, and t_u denote the degrees of relative risk aversion, relative prudence, and relative temperance, respectively (obtained by multiplying the absolute degrees by w). We can then rewrite the absolute risk aversion of $v(\cdot)$ as

$$R_v(w) \simeq R_u(w) \left(\frac{1 + p_u t_u s^2/2}{1 + p_u r_u s^2/2} \right).$$

If the utility function is given by (8) in the text, then $R_u = kw^{-\beta}$, $r_u = \kappa w^{1-\beta}$, $p_u = \beta + r_u$, and $t_u = \beta + p_u$. Substituting into the preceding expression for $R_v(w)$ and taking logs yields

$$\log R_v \simeq \log \kappa - \beta \log w + \beta p_u s^2,$$

which shows that the parameter β of the utility function $u(\cdot)$ can be recovered even if there is background risk.

A.4 Implications of measurement error for the estimation of β

Let

$$\log(T) = -\log \alpha + \beta \log c + \varepsilon$$

be the true relation, where $T = f(Z)$. It can be shown that $f(Z)$ is a convex function of Z. Instead of observing T we observe $\widehat{T} = f(Z + u)$, where u is a measurement error such that $E(u) = 0$ and $E(c, u) = E(\alpha, u) = 0$. Hence, we can estimate

$$\log(\widehat{T}) = -\log a + b \log c + e.$$

Let $\log(T) = \log f(Z) = h(Z)$ and $\log(\widehat{T}) = \log(f(Z + u)) = h(Z + u)$. In general it is not clear whether the $h(\cdot)$ function is concave or convex. Taking a Taylor expansion of $h(Z + u)$ around $u = 0$ and disregarding terms of order higher than the second, we obtain:

$$\log(\widehat{T}) \approx h(Z) + h'(Z)u + \frac{1}{2}h''(Z)u^2 = \log(T) + h'(Z)u + \frac{1}{2}h''(Z)u^2.$$

Using this expression, we can write our estimate of the elasticity of risk tolerance to the endowment as

$$\begin{aligned} \widehat{b} &= (\log c' \log c)^{-1}(\log c' \log(\widehat{T}) \\ &\approx \beta + (\log c' \log c)^{-1} \log c' h'(Z)u + (\log c' \log c)^{-1} \log c' h''(Z)u^2/2, \end{aligned}$$

and

$$plim(\widehat{\beta} - \beta) \approx plim(\log c' \log c)^{-1} plim(\log c' h''(Z))\sigma_u^2/2.$$

If preferences are DARA then c and Z are positively correlated. Thus, the sign of the bias depends on the sign of $h''(Z)$. If $h(\cdot)$ is linear then there is no bias; if it is concave the bias is downward, and if it is convex the bias is upward. Since

$h''(Z) = d^2 \log f(Z)/dz^2$, we can write $h''(Z)$ as

$$h''(Z) = f(Z)^{-1} \left[f''(Z) - \frac{f'(Z)^2}{f(Z)} \right].$$

$h''(Z)$ is zero if and only if $f''(Z) = f'(Z)^2/f(Z)$; this appears to be the case in our data, where $h(\cdot)$ turns out to be approximately linear.

References

[1] Arrow, Kenneth (1965), Aspects of the Theory of Risk Bearing, Yrjo Jahnsson Lectures, The Academic Book Store, Helsinki.

[2] Arrow, Kenneth (1970), Essays in the Theory of Risk Bearing, Amsterdam: North Holland.

[3] Barberis, Nicholas, Ming Huang, and Richard H. Thaler (2006) "Individual, Monetary Gambles, and Stock Market Participation: A Case for Narrow Framing" American Economic Review, 96,4, 1069-1090.

[4] Barsky, Robert, B., Thomas F. Juster, Miles S. Kimball and Matthew D. Shapiro (1997), "Preference Parameters and Behavioral Heterogeneity: an Experimental Approach in the health and Retirement Study", Quarterly Journal of Economics, CXII, 2, 537-580.

[5] Benjamin, Daniel J., Sebastian A. Brown and Jesse M. Shapiro, (2006) "Who is "Behavioral"? Cognitive Ability and Anomalous Preferences", University of Chicago, mimeo.

[6] Brandolini, Andrea, and Luigi Cannari (1994), "Methodological Appendix", in A. Ando, L. Guiso and I. Visco (eds.) "Saving and the accumulation of wealth", Cambridge: Cambridge University Press.

[7] Breeden, Douglas T. (1979), "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities", Journal of Financial Economics, vol. 7, 265-96.

[8] Carroll, Christopher (2000), "Portfolios of the Rich", in L. Guiso, M. Haliassos and T. Jappelli (eds.) "Household Portfolios", MIT Press, Boston.

[9] Carroll, Christopher and Miles Kimball (1996), "On the Concavity of the Consumption Function", Econometrica 4, 64, 981-992.

[10] Cascio, Elizabeth and Ethan Lewis (2005), "Schooling and the AFQT: Evidence from School Entry Laws", NBER Working Paper 11113.

[11] Eekhoudt Louis, and Miles Kimball (1992), "Background Risk, Prudence and the Demand for Insurance", in G. Dionne (ed.) "Contributions to Insurance Economics", London: Kluwer Academic Press.

[12] Eekhoudt Louis, Christian Gollier, and Harris Schlesinger (1996), "Changes in Background Risk and Risk Taking Behavior", Econometrica 3, 64, 683-689.

[13] Fuchs-Schündeln, Nicola and Matthias Schündeln (2005), "Precautionary Savings and Self-Selection. Evidence from the German Reunification 'Experiment,"' Quarterly Journal of Economics, 120 (3), 1085-1120.

[14] Gollier Christian (2000), "What Does the Classical Theory Have to Say About Portfolio Choice ?", in L. Guiso, M. Haliassos and T. Jappelli (eds.) "Household Portfolios", MIT Press, Boston.

[15] Gollier Christian, (2001), "Wealth Inequality and Asset Pricing ", Review of Economic Studies, 68, 1, 181-203.

[16] Gollier Christian, and John W. Pratt (1996), "Weak Proper Risk Aversion and the Tempering Effect of Background Risk", Econometrica 5, 64, 1109-1123.

[17] Gollier Christian, and Richard Zeckhauser (2002), "Horizon Length and Portfolio Risk", Journal of Risk and Uncertainty, 24, 3, 95-212.

[18] Guiso, Luigi and Tullio Jappelli (1998),"Background Uncertainty and the Demand for Insurance Against Insurable Risks", Geneva Papers on Risk and Insurance Theory, 23, 1, 7-27.

[19] Guiso, Luigi and Tullio Jappelli (2000), "Household Portfolios in Italy ", in L. Guiso, M. Haliassos and T. Jappelli (eds.) "Household Portfolios", MIT Press, Boston.

[20] Guiso, Luigi and Tullio Jappelli (2006), "Information Acquisition and Portfolio Performance", CEPR WP 5901, October.

[21] Guiso, Luigi, Tullio Jappelli and Luigi Pistaferri, (2002), "An Empirical Analysis of Earnings and Employment Risk", Journal of Business and Economic Statistics, 20, 1-13.

[22] Guiso, Luigi, Tullio Jappelli and Daniele Terlizzese (1996), "Income Risk, Borrowing Constraints and Portfolio Choice", American Economic Review, 86, 158-172.

[23] Guiso, Luigi and Monica Paiella (2006), "The Role of Risk Aversion in Predicting Individual Behavior", in P. A. Chiappori and C. Gollier (eds.), "Insurance: Theoretical Analysis and Policy Implications", MIT Press.

[24] Heaton, John and Deborah Lucas (2000), "Portfolio Choice in the Presence of Background Risk", Economic Journal, 110, 1-26.

[25] Hennessy, David A. and Harvey E. Lapan (2006), "On the Nature of Certainty Equivalent Functionals", Journal of Mathematical Ecnomics, 43, 1, 1-10.

[26] Holt, Charles and Susan K. Laury, (2002), "Risk Aversion and Incentive Effects", American Economic Review, 92, 5, 1644-1655.

[27] Hurd, Michael D. (2000), "Portfolio Holdings of the Elderly", in L. Guiso, M. Haliassos and T. Jappelli (eds.) "Household Portfolios", MIT Press, Boston.

[28] Hurst, Erik and Frank Stafford (2004), "Home is Where Equity Is: Liquidity Constraints, Refinancing and Consumption", Journal of Money, Credit and Banking, 36, 985-1014

[29] Kihlstrom, Richard E., David Romer and Steve Williams (1981), "Risk Aversion with Random Initial Wealth", Econometrica 4,49, 911-920.

[30] Kimball, Miles (1993), "Standard Risk Aversion", Econometrica 3, 61, 589-611.

[31] Kimball, Miles S., and Philippe Weil (2004): "Precautionary Saving and Consumption Smoothing Across Time and Possibilities," Manuscript, University of Michigan.

[32] Peress, Joel (2004), "Wealth, Information Acquisition and Portfolio Choice," The Review of Financial Studies 17, 879-914.

[33] Pratt, John (1964), "Risk Aversion in the Small and in the Large", Econometrica, 32, 122-136.

[34] Pratt John W., and Richard Zeckhauser (1987), "Proper Risk Aversion", Econometrica 1, 55, 143-154

[35] Rabin, Matthew (2000), "Risk Aversion and Expected Utility Theory: A Calibration Theorem", Econometrica, 68, 1281-1292.

[36] Shane, Frederick (2005), "Cognitive Reflection and Decision Making", Journal of Economic Perspectives, 19, (4), 25-42.

[37] Tiseno, Andrea (2002), "Using Wealth, Consumption and Shadow Prices to Identify Intertemporal Preferences", University of Chicago, Department of Economics, mimeo.

[38] Tversky and Kahneman (1992), "Advances in Prospect Theory: Cumulative Representation of Uncertainty", Journal of Risk and Uncertainty, V, pp. 297-323.

[39] Weil, Philippe (1992), "Equilibrium Asset Prices with Undiversifiable Labor Income Risk", Journal of Economic Dynamics and Control, 16, 769-790.