

PREFERENCE PARAMETERS AND BEHAVIORAL HETEROGENEITY:
AN EXPERIMENTAL APPROACH IN THE HEALTH AND RETIREMENT STUDY*

Robert B. Barsky

F. Thomas Juster

Miles S. Kimball

Matthew D. Shapiro

Survey Research Center and Department of Economics

University of Michigan

First Draft: September 1993

Revised: January 1997

*This research was supported by a program project grant from the National Institute of Aging, "Wealth, Saving, and Financial Security Among Older Households," PO1 AG10179-03. Barsky and Shapiro also acknowledge the support of the Alfred P. Sloan Foundation. Carlos Quintanilla and Lisa Sanchez provided excellent research assistance. We gratefully acknowledge the constructive and critical comments of John Campbell, Lawrence Katz, Robert Shiller, Richard Thaler, anonymous referees, and numerous seminar and conference participants.

address for correspondence:

Matthew D. Shapiro
Department of Economics
University of Michigan
Ann Arbor, MI 48109
tel. 313 764-5419
internet: shapiro@umich.edu

**Preference Parameters and Behavioral Heterogeneity:
An Experimental Approach in the Health and Retirement Study**

ABSTRACT

This paper reports measures of preference parameters relating to risk tolerance, time preference, and intertemporal substitution. These measures are based on survey responses to hypothetical situations constructed using an economic theorist's concept of the underlying parameters. The individual measures of preference parameters display heterogeneity. Estimated risk tolerance and the elasticity of intertemporal substitution are essentially uncorrelated across individuals. Measured risk tolerance is positively related to risky behaviors, including smoking, drinking, failing to have insurance, and holding stocks rather than Treasury bills. These relationships are both statistically and quantitatively significant, although measured risk tolerance explains only a small fraction of the variation of the studied behaviors.

Robert B. Barsky
F. Thomas Juster
Miles S. Kimball
Matthew D. Shapiro
Survey Research Center
and Department of Economics
University of Michigan
Ann Arbor, MI 48109
tel. 313 764-5419

I. Introduction

A recurrent theme in Amos Tversky's remarkable output is the description of individual preferences and their relation to choice behavior. In particular, Tversky work is concerned with achieving a better match between theory and empirical evidence with respect to behavior toward risk.¹ Tversky was most concerned with situations in which the evidence seems to contradict expected utility theory. While in this paper, we adhere to an expected utility benchmark, Tversky's concern with explaining individual behavior in a variety of situations also impels our work.

This paper describes the results of an experimental attempt to elicit individual preference parameters by means of direct questions closely derived from economic theory, and to study the behavioral implications of heterogeneity in the measured parameters across individuals. Participants in the Health and Retirement Study were asked to respond to hypothetical situations specifically designed to elicit information about their risk aversion, subjective rate of time preference, and willingness to substitute intertemporally. These three parameters are essential to individual choices about wealth accumulation, retirement, portfolio allocation, and insurance, as well as to policy choices that are dependent on this behavior.

Despite the analytic importance of these preference parameters, econometric studies have not fully resolved issues involving even their mean values. Indeed, even when the underlying parameter is constant across individuals, econometric estimation often relies on problematic identifying restrictions. The econometrician typically needs to posit a functional form. Instrumental variables are needed to control for potential endogeneity. The survey approach addresses these issues by constructing a survey instrument that is designed precisely to elicit the parameter of interest while asking the respondent to control for differences in economic circumstances that confound estimation. While the survey approach introduces other problems—for example, whether the respondents are giving accurate answers—it can provide a potentially important source of information about these parameters in addition to econometric evidence.

Econometric estimation of preference parameters may be particularly inadequate when heterogeneity of preferences in the population is important. In this case it may be desirable to have an

estimate of the parameters of interest for each individual in a cross-section, not just the average value of that parameter in the population. In a cross-section, one would be able to study the covariation between the estimated parameters and observed behavior with regard to saving, portfolio choice, labor supply, insurance purchases, etc. Absent enough data to estimate the econometric model for each individual (i.e. a long panel), the standard econometric approaches cannot assign values of the parameter to specific individuals.

The underlying purpose of our research is to explore the possibility of obtaining information about theoretically important parameters from direct questions involving choice in hypothetical situations, with as little departure from the theorist's concept of a parameter as possible. We obtain our measure of risk aversion by asking respondents about their willingness to gamble on lifetime income. By contrast, experiments in the existing literature typically involve stakes that have little impact on lifetime resources. A gamble whose outcome is uncorrelated with consumption should not require a risk premium.

We obtain our measures of intertemporal substitution and time preference by asking respondents to choose consumption profiles implicitly associated with different rates of return. According to the relevant economic theory, the two parameters are the solution of two equations in two unknowns. The questions typically asked about time preference in the literature do not properly distinguish between the subjective discount rate and the market rate of interest. As we emphasize in Section II, the rate at which individuals are willing to trade off present and future consumption depends on both. By asking for the preferred consumption path at more than one interest rate, we are able in principle to separate time preference from the market interest rate.

The organization of the paper is as follows. In Section II we spell out our methods for measuring risk preference, intertemporal substitution, and time preference. In Section III we report our results for the questions about risk preference. In Section IV we apply these results to the equity premium puzzle. In Section V we report our results for questions about preferred consumption paths. In Section VI we discuss some caveats about the survey and extensions of the modeling of the results. In Section VII we present our conclusions.

II. Methodology

A. Measuring Risk Aversion

The principal requirement for the question aimed at measuring risk aversion is that it must involve gambles over lifetime income. After considerable testing,² in which the precise nature of the hypothetical circumstances was refined several times to minimize misunderstandings and additional complications envisioned by respondents, we arrived at the following question.³

Suppose that you are the only income earner in the family, and you have a good job guaranteed to give you your current (family) income every year for life. You are given the opportunity to take a new and equally good job, with a 50-50 chance it will double your (family) income and a 50-50 chance that it will cut your (family) income by a third. Would you take the new job?

If the answer to the first question is “yes,” the interviewer continues:

Suppose the chances were 50-50 that it would double your (family) income, and 50-50 that it would cut it in half. Would you still take the new job?

If the answer to the first question is “no,” the interviewer continues:

Suppose the chances were 50-50 that it would double your (family) income and 50-50 that it would cut it by 20 percent. Would you then take the new job?

The questions separate the respondents into four distinct risk preference categories, depending on the answers to two questions. The categories can be ranked by risk aversion without having to assume a particular functional form for the utility function. The categorical responses (labeled I, II, III, and IV) are summarized in the first column of Table I.

The categorical responses can be thought of as resulting from the following expected utility calculation. Let U be the utility function and c be permanent consumption. An expected utility maximizer will choose the 50-50 gamble of doubling lifetime income as opposed to having it fall by factor $1 - \lambda$ if

$$(1) \quad \frac{1}{2}U(2c) + \frac{1}{2}U(\lambda c) \geq U(c),$$

that is, the expected utility of the income stream offered by the gamble exceeds the expected utility of having the current income stream with certainty.

If one is willing to assume that relative risk aversion $1/\theta = -c \cdot U''/U'$ is constant over the relevant region, the categorical responses bound the values of relative risk aversion. Table I gives the upper and lower bounds on relative risk aversion corresponding to the choices I, II, III, and IV. For most of the analysis, we work with the reciprocal of relative risk aversion, called “relative risk tolerance” in the finance literature. Risk tolerance, unlike risk aversion, aggregates linearly. Table I also gives the ranges of relative risk tolerance θ consistent with the choices. The lower bound on relative risk tolerance is the reciprocal of the upper bound on relative risk aversion and vice versa. Table I includes the mean relative risk aversion and tolerance corresponding to these ranges. This mean depends on the distribution of the preference parameter in the population. We discuss in the next section how we estimate this distribution and how we construct the estimates in the last two columns of Table I.⁴

One important criticism of this survey question is that respondents might value their current job for reasons other than the income flow associated with it and therefore might be reluctant to switch jobs even for a high expected increase in income. This “status quo bias” would tend to increase the estimate of risk tolerance because it gives a reason in addition to risk aversion for individuals to express an unwillingness to accept the gamble. In Section VI of the paper, we address the issue of status quo bias by offering a quantitative assessment of its potential impact on our results and a suggestion for eliminating status quo bias in future surveys.

B. Measuring Time Preference and the Elasticity of Intertemporal Substitution

Our experimental survey also sought to develop estimates of the desired slope of the path of consumption over time and the willingness of individuals to alter the slope of the consumption path in response to changes in the interest rate. These choices relate to two preference parameters: the rate of time preference and the elasticity of intertemporal substitution. To estimate these parameters, we attached an experimental set of questions to Wave I of the Health and Retirement Study (designated Module K). In contrast to the questions about gambles over lifetime income, which are part of the main

Health and Retirement Study questionnaire, the questions in the modules were asked only of a small subset of respondents.

The basic behavioral equation underlying our survey design for capturing time preference and intertemporal substitution is

$$(2) \quad \Delta \log(c) = s(r - \rho)$$

where c is consumption, r is the real interest rate, ρ is the subjective discount rate, and s is the elasticity of intertemporal substitution. Equation (2) says that if the real interest rate is greater than the rate of time preference, consumption will be growing over time. If the rate of time preference is less than the interest rate, agents start out with relatively low consumption in order to save and take advantage of the high rate of return. This effect will be larger the larger is s , which measures the strength of the willingness to intertemporally substitute in consumption. Given r and s , the larger is ρ , the less upward-sloping will be the chosen consumption path, as households discount the future more heavily. If $r < \rho$, the interest rate is not high enough to overcome the subjective discounting of the future, and agents choose consumption paths that (in expectation) fall as they age.⁵

In the module, we first posed a hypothetical set of circumstances that are meant to control for heterogeneity in economic and demographic conditions facing the household. (In particular, respondents were told to assume no inflation and that they would have no uninsured health expenses.) Then the respondents were shown charts with different profiles of consumption with constant present value at a zero interest rate and were asked to choose the preferred path. In subsequent questions, they were asked to choose among constant present value consumption paths with interest rates of 4.6 and -4.6 percent per year. From the slopes of the preferred paths and how the slopes change when the interest rate changes, one can estimate the rate of time preference and the elasticity of intertemporal substitution. The appendix contains the exact wording of the question and the charts containing the consumption profiles from which the individuals chose.

C. The Survey Instrument

The questions were included in Wave I of the Health and Retirement Study (HRS), administered to a large cross-section of households in 1992.⁶ The primary respondents were between the ages of 51 and 61 in 1992. The survey also includes their spouses. The full survey, which takes about 2 hours, was conducted face-to-face in the field. Respondents are paid for their participation. The HRS asks a wide range of questions concerning health status, retirement decisions, income, and assets. It also asks a number of behavioral questions, such as whether the individual smokes or drinks. In the case of couples, questions that pertain to individuals—including our questions on risk preference and intertemporal consumption preferences—are asked of both. Questions pertaining to the household as a whole, e.g. about wealth, are asked only of a primary respondent. The primary respondent is the member of the couple “most knowledgeable” about the family's assets, debts, and retirement planning.⁷

The survey yielded 11,707 responses to the risk preference questions, 7,278 from primary respondents and 4,429 from secondary respondents. Appendix Table I gives some summary statistics. The average age of 55.6 years reflects the sampling frame. Primary respondents, who are disproportionately male, are a little older and have a little more education.⁸ The experimental intertemporal consumption preference questions were given only to a very small subsample, yielding only 198 observations.

The risk preference questions were also included in one of the modules of Wave II of the Health and Retirement Study, which was conducted by telephone. This module was asked of roughly ten percent of the sample, most of whom responded to the questions on Wave I. We use the multiple responses to the question to estimate the measurement error associated with the responses.

III. Preference over Lifetime Income Gambles

We find substantial heterogeneity in the estimates of risk preference.

- The response exhibiting least risk tolerance is strongly modal. Hence, low risk tolerance characterizes most of the population.

- Nonetheless, there is substantial heterogeneity in risk tolerance. A significant fraction of the sample exhibited willingness to undertake substantial gambles over lifetime income.
- The measured risk tolerance has predictive power for choices over risky behaviors—the decisions to smoke and drink, to buy insurance, to immigrate, to be self-employed, and to hold stock. The behaviors studied are, nevertheless, very noisy and difficult to explain; the incremental predictive power of risk tolerance is never very high.

In this section, we present tabulations, cross-tabulations, and regressions that establish these findings.

A. The Distribution of True and Measured Risk Tolerance

The survey groups the respondents into the four risk tolerance categories detailed in Table I. The survey response is, however, likely to be subject to noise. In this subsection, we describe a procedure for estimating the distribution of the true parameter and the distribution of the noise.⁹ This procedure is possible because a subset of respondents answered the risk tolerance questions in both Wave I and Wave II of the Health and Retirement Study.¹⁰ By studying how the responses correlate across waves, we can quantify the signal and noise in the survey responses.

Consider the following model of relative risk tolerance, denoted θ_j . Let

$$(3) \quad x_j = \log(\theta_j)$$

be the logarithm of individual j 's true relative risk tolerance and let ϵ_{jk} be an independent error associated with the individual's response to the survey k ($k = \text{Wave I or Wave II}$). The true parameter is assumed to remain constant. The realized log risk tolerance

$$(4) \quad y_{jk} = x_j + \epsilon_{jk}$$

equals the true log risk tolerance plus the error. Let B_i be the range of log risk tolerance for category $i =$ I, II, III, and IV. We assume that an individual will choose response i if $y_{jk} \in B_i$. That is, there is noise in how an individual will report his or her risk tolerance on a given day, but given the noise, the individual calculates correctly which gamble to accept.¹¹ Note that this model is quite different from the standard latent variable model. In the standard model, the latent variable x_j and the cutoffs defining B_i are based on an arbitrary index. In contrast, our latent variable is a cardinal preference parameter and the cutoffs are known numbers.

To identify the statistical model requires strong assumptions. In particular, the estimation scheme presumes that the persistent component of responses represents true preferences. In Section VI we relax this assumption to allow for a persistent component to the error in the responses.

With one response per individual, we would have been able to estimate only the distribution of y_{jk} , not the distribution of the true parameter x_j . But 717 individuals responded to the risk preference questions on both waves. Appendix Table II gives the joint distribution of responses to both waves.¹² This empirical distribution allows us to estimate the distribution of the true parameter and to quantify the noise. To implement this statistical model, we assume that θ_j is distributed lognormally across individuals.¹³ The maximum likelihood estimate of the mean of log risk tolerance x_j is -1.96. The estimated standard deviation of x is 1.03 and of ϵ is 1.39.¹⁴ These parameter estimates yield a mean true risk tolerance θ of 0.24. The estimated correlation of 0.60 between the true (x) and reported (y) log risk tolerance is quite high.

Based on these assumptions, the last two columns of Table I report the expectation of relative risk aversion ($1/\theta$) and relative risk tolerance (θ) conditional on an individual responding I, II, III, or IV to the risk preference questions in Wave I.¹⁵ Table I also reports the means of the ranges for the true choices I, II, III, and IV. Were there no measurement error, the means of the ranges would be equal to the conditional expectations. Because of the measurement error, the conditional expectation of the preference parameter given the survey response reverts toward the unconditional mean.

Comparing the last two columns of Table I illustrates the importance of Jensen's inequality. The expectation of the reciprocal is substantially greater than the reciprocal of the expectation.

The following subsections report the results for the risk preference survey questions. When we report mean risk tolerance or include risk tolerance as a regressor, we use the conditional expectations shown in the last column of Table I.¹⁶ We also present cross-tabulations based on the categorical responses. These do not depend on the functional form of the utility function or our statistical model for the measurement error.

B. Heterogeneity in Individuals' Risk Tolerance

Table II gives the fraction of all respondents who fall into risk tolerance groups I, II, III, and IV. Most respondents are in category I, but a significant minority are in the higher risk tolerance categories.¹⁷ Based on the estimated underlying lognormal density, the mean risk tolerance is 0.24 and the standard deviation is 0.33.¹⁸ Table II also gives the results separately for the primary and secondary respondents. The distribution of responses across the four risk tolerance categories and mean risk tolerance are nearly identical for the two groups.

The second part of Table II shows the distribution of responses of the secondary respondent conditional on the response of the primary respondent. The diagonal elements are substantially larger than the unconditional distribution shown in the first part of the table. The simple correlation of relative risk tolerance across household members is only 0.12, but is strongly significant (t-statistic of 7.8).¹⁹

Table III examines how risk preference varies by demographic group. There are substantial differences by age in estimated risk tolerance. The youngest and the oldest cohorts are the most risk tolerant, with cohorts in the middle being less risk tolerant. The groups under 55 years old choose the least risk tolerant choice (I) relatively infrequently; the group over 70 chooses the most risk tolerant option (IV) relatively frequently. Ages 55 to 70 are relatively risk intolerant. We can reject with a high degree of confidence (p-value of 0.0001) that the mean risk tolerance of these age groups is equal.

There are also differences in risk tolerance by sex. Males are somewhat more risk tolerant than females, with the biggest difference being in males' propensity to choose the most risk tolerant option (IV). Again, the differences are statistically significant.

There are noticeable differences in risk tolerance by the race and religion of the respondent. Whites are the least risk tolerant, blacks and Native Americans somewhat more risk tolerant, and Asians

and Hispanics the most risk tolerant. Again, the differences are easiest to see in the columns I and IV giving the extreme responses. For example, Asians are 7 percentage points less likely than whites to choose the least risk tolerant response and are 9 percentage points more likely to choose the most risk-tolerant response. Risk tolerance also varies significantly by religion. Protestants are the least risk tolerant and Jews the most. In risk tolerance, Catholics are about halfway between Protestants and Jews.

C. Is Risk Tolerance Related to Behavior?

In this subsection, we examine the extent to which measured risk tolerance predicts risky behavior. Showing that our measure of risk tolerance predicts behavior in the way one would expect partially validates the survey measure. Psychologists studying the conceptualization and measurement of personality traits have been interested in what Mischel [1971] calls the issue of “behavioral specificity.” Do individuals tend to show similar responses to all risky situations (e.g. financial, social, and health risks), or is risk-taking in one setting nearly independent of risk-taking behavior in other settings? Slovic cites a dozen studies apparently showing that “the majority of the evidence argues against the existence of risk-taking propensity as a generalized characteristic of individuals.” [See Slovic 1972a and 1972b.] More recently, however, questions measuring the characteristics of “harm avoidance,” “novelty seeking,” and “reward dependence”—no doubt closely related to risk aversion—have formed the basis of the much-used Cloninger tridimensional personality scale. [See Cloninger, Przybeck, and Svrakic 1991.] Some researchers in neurology and psychiatry [Menza, Golbe, Cody, and Forman 1993] have reported evidence of a biological basis for particular responses on the Cloninger scales, which suggests that they measure stable personality traits with some constancy across settings. The results of this section can be used to evaluate whether our survey measure captures a parameter that similarly has explanatory power across behaviors.

Moreover, the relationship of measured risk tolerance with various behaviors is something of interest in its own right. We stop far short of constructing a complete theory of all the behaviors that are potentially related to risk preference. Rather, we present cross-tabulations analogous to the ones in the previous subsection. We also estimate some simple, linear regressions in an attempt to control for some correlates of risk preference.

The risk tolerance measure does predict risky behaviors—including smoking, drinking, having insurance, choosing risky employment, and holding risky assets. These results are often strongly significant statistically and are associated with quantitatively significant coefficient estimates. We can decisively reject the null that the measured preference parameters are unrelated to behavior. The fraction of the variance of the various behaviors that our survey instrument explains is, however, quite small.²⁰

Smoking and Drinking. The first three panels of Table IV show the distribution of risk tolerance conditional on smoking and drinking—behaviors that increase health risk.²¹ The corresponding regression estimates are reported in the first rows of Table V. Individuals who have ever smoked are more risk tolerant than those who never smoked and those who smoke now are more risk tolerant than those who do not smoke now. Of particular interest are those who say they once smoked, but do not smoke now. The sample is largely composed of middle-aged to older individuals. Hence, those who quit smoking would have done so during a period of increasing public awareness of the risks associated with cigarette smoking. Those who quit smoking are somewhat more risk tolerant than those who never smoked, but less risk tolerant than current smokers.

Whether an individual drinks or not is also related to measured risk tolerance. Risk tolerance is higher for those who drink than for those who do not drink. The difference in risk tolerance between drinkers and nondrinkers is about the same as between smokers and nonsmokers. Moderate drinking is not generally believed to be a health risk. Table IV shows risk tolerance by drinks per day. Those who take less than one drink per day have a willingness to accept the moderate gambles (II and III) relatively often. As drinks per day increase, there is a monotonic increase in mean risk tolerance. For heavy drinkers, risk tolerance—measured either by willingness to choose gamble IV or by mean risk tolerance—is substantially above average.

The regressions reported in Table V show that the risk tolerance measure predicts smoking and drinking even after controlling for the demographic variables. Moreover, the risk tolerance measure has a substantial quantitative role in predicting these behaviors. For example, the most risk tolerant respondents are over 3-1/2 percentage points more likely to have ever smoked than the least risk tolerant respondents (0.0922 times the 0.42 difference between the expected risk tolerances of a respondent in categories I and IV). Risk tolerance is also a significant explanatory variable for drinking behavior.

Moving from the lowest to highest response for risk tolerance is associated with a 4 percent increase in the probability of drinking (t-statistic of 3) and a 0.1 drink increase in the number of drinks per day (t-statistic of 4-1/2).²²

Education and Employment Status. The fourth panel of Table IV shows a U-shaped relationship between years of schooling completed and the measure of risk tolerance. Individuals with exactly 12 years of schooling are the least risk tolerant. Indeed, the mean risk tolerance of 0.229 and average propensity to choose response IV of 10.2 percent are the lowest we found for any group that we categorized. Those with less than 12 or from 13 to 16 years of schooling have slightly greater than average risk tolerance. Those with some post-college education (years greater than 16) have substantially greater than average risk tolerance. In the multivariate analysis in Table V, the number of years of schooling is not associated with risk tolerance—in part because of the nonlinearity we found in the cross-tabulation.

Among the biggest risks voluntarily taken by a large segment of the population is self-employment. The self-employed generally face a riskier overall income stream than their wage-earning or salaried counterparts. [See Friedman 1957 and Carroll 1994.] Thus, one would expect risk tolerance to be positively associated with the decision to undertake self-employment. The fifth panel of Table IV shows that the self-employed are more risk tolerant than employees.²³ The multivariate analysis shows that the most risk tolerant respondents are about one percentage point more likely to select self-employment than the least risk tolerant respondents. Given that the probability of self-employment is less than 12 percent, this is a quantitatively large effect, but it is not statistically significant (t-statistic of 1.1).

Region and Immigrant Status. An epic risk is to move to a new country in search of a better life. The idea that immigrants are more daring than the average person is part of the American mythology. Migration within the United States could also entail significant risks. The western United States has in the past been an internal frontier to which one might argue the more daring have migrated. Some of the attitudes from that frontier past may have persisted to the present.

Both region of residence and immigrant status are significantly predicted by risk tolerance. Residents of the western United States are more risk tolerant than residents of other regions. Immigrants

are also substantially more risk tolerant than natives. They are especially likely to be in category IV. (See Table IV.)

Given that many recent immigrants are Hispanic and Asian and that Hispanics and Asians have high risk tolerance (see Table III), it is important to check that immigrant status is not confounded with ethnicity. The positive association of risk tolerance and immigrant status survives controlling for the demographic factors, but has a t-statistic of only 1-3/4 (Table V).

Health and life Insurance. Anyone with positive risk aversion should be fully insured against purely financial risks when insurance is actuarially fair. In the presence of a load factor, however, those who are most risk averse should be most willing to buy insurance against financial risks. A complication arises (as with smoking and drinking) because the kinds of insurance purchases on which we have information are health and life insurance, where the risks are not purely financial—the marginal utility of wealth potentially depending on health status, for instance. We appeal to financial responsibility for the support of others as the basis of our a priori expectation that the (financially) more risk averse are more likely to purchase both medical and life insurance.²⁴

Table VI examines our measure of risk tolerance according to whether or not the individual has health insurance. We do separate tabulations for employees, the self-employed, and those not working. To focus on those who have the option of having insurance or not, this tabulation excludes those in the Medicare-eligible age group.

For each of the three employment classes, more risk tolerant individuals are less likely to have health insurance. For those employed, measured risk tolerance seems to be an important factor sorting individuals into jobs with health insurance. For the not employed, risk preference is a powerful determinant of the propensity to be insured. The effect of risk tolerance on the propensity to be insured is smaller among the self-employed than among the unemployed. Between groups, the self-employed have a higher risk tolerance and have much lower average propensity to be insured than employees. Similarly, Table VII shows that individuals without life insurance are substantially more risk tolerant than those with it.

The results in the cross-tabulations for health and life insurance carry over when the demographic factors are controlled for in the regressions reported in Table V. The most risk tolerant respondents are

8.2 percentage points more likely not to have health insurance and over 6-1/2 percentage points more likely to forgo life insurance than the least risk-tolerant respondents. Both results are highly statistically significant (t-statistics in excess of 5) and are quantitatively important.

Income and wealth. Tables VIII and IX show risk tolerance by quintiles of income and wealth. Risk tolerance decreases with income and wealth until the middle of the distributions, and then increases. Note that the pattern of risk tolerance by income and wealth is similar to that for age. Risk tolerance rises at the high end of the wealth, income, and age distributions.²⁵

Home equity is the major component of wealth for most individuals. The 20 percent of individuals who do not live in houses they own are substantially more risk tolerant than those who own their homes. The most risk tolerant individuals are over six percent less likely to own homes than the least risk tolerant individuals (see the last line of Table V). It is not obvious what correlation one would expect a priori. Although house prices are volatile and houses are often highly leveraged, owning a house insulates individuals from local changes in the cost of shelter, and thus provides some consumption insurance.

Financial Assets. Studying the demand for risky assets is an important application of our risk preference measures. Table X presents regressions of portfolio shares on the demographic variables, risk tolerance, wealth, and income for a subsample that has positive financial assets. Many households have little or no financial wealth. We limit this analysis to households that have at least \$1000 in financial wealth. This criterion excludes about one-sixth of the households. Since asset ownership depends substantially on income and wealth, we include these as controls in the regressions of portfolio variables.²⁶

The questions about assets apply to the household. In the Health and Retirement Study, they are answered by the “knowledgeable respondent”—the member of the household with the best knowledge of the household's assets. The assets are characteristics of the household (there is no information on asset ownership within the household), while risk preference is a feature of individuals. Recall that the risk tolerance measure is positively, but not strongly, correlated within couples (Table II). To study the role of potentially conflicting risk tolerances within the family, we enter the risk tolerance of members of

couples separately. Hence, the regressions include the risk tolerance of the primary respondent (R1) and the difference between the risk tolerance of the primary and secondary respondent.²⁷

The risk tolerance measure has significant predictive power for stock ownership. In households where the primary respondent gave the most risk-tolerant response, the fraction of financial assets held in equities is 4.1 percentage points higher on average than in those where the primary respondent gave the least risk-tolerant response. Since the average fraction of portfolios in stocks is only 14 percent, this effect is substantial. It is also strongly statistically significant. If the secondary respondent is less risk tolerant than the primary respondent, the stock share is lower, although this result is not statistically significant.²⁸

Similarly, relatively safe assets—Treasury bills and savings accounts—make up more of the portfolios of the less risk tolerant respondents. Bonds are too small a share of portfolios for the results to be decisive (1.4 percent of portfolios on average), although there is a marginally significant positive relationship between bond holding and risk tolerance. Ownership of other assets (trusts, collections held for investment) is powerfully related to risk tolerance.²⁹ Therefore, for assets at opposite ends of the risk spectrum—stocks at one end versus Treasury bills and savings accounts at the other end—the risk tolerance measure has substantial explanatory power for portfolio demands.

Yet, the relationship between risk tolerance and the holding of risky assets is much weaker than theory suggests it should be. In the standard capital asset pricing model, portfolio shares equal the product of risk tolerance and covariance-weighted excess returns. Hence, the elasticity for the share of each risky asset with respect to risk tolerance should be 1. When the estimated coefficients of risk tolerance in Table X are expressed as elasticities, the elasticity of the stock portfolio share with respect to risk tolerance is estimated to be 0.17; the elasticities of the bond share and of other assets are estimated to be 0.25 and 0.21. Consequently, there is inadequate sensitivity of portfolio shares to risk tolerance compared to the prediction of the standard model.

IV. Heterogeneous Risk Preferences and the Equity Premium Puzzle

In this section, we discuss how to aggregate our estimates of individuals' risk tolerance in a way that informs the demand for stockholding. The “equity premium puzzle” [Grossman and Shiller 1981 and

Mehra and Prescott 1985] is a mismatch between the low levels of risk tolerance (high levels of risk aversion) required to explain empirical facts about mean asset returns and the range of values for risk tolerance that seem reasonable to most economists. We ask whether the answer to the equity premium puzzle might simply lie in the fact that the average individual is more risk averse than an economist might have expected, as indicated by the high percentage of respondents in the least risk-tolerant category.³⁰ We demonstrate here that our findings are in fact not consistent with this story. Although most individuals are quite risk averse, there are enough risk tolerant individuals to hold the outstanding supply of equity at far less than the historically observed risk premium. Equivalently, because finance theory implies that in aggregating preferences across individuals the least risk averse receive the greatest weight, the effective risk aversion of the representative investor is actually rather low. These statements are robust to different ways of treating nonstockholders.

The literature on the equity premium puzzle finds that aggregate risk aversion must be in excess of 30—and possibly as high as 100—in order to explain the 6 percentage point per year mean excess return of equities above the Treasury bill rate over the last century or so. [See Mankiw and Zeldes 1991.] What is the appropriate way to use our estimate of the population distribution of preferences to aggregate the heterogeneous individuals and construct a single number for use in an asset demand equation? One might be tempted to take a simple (weighted or unweighted) arithmetic average of risk aversion over all individuals. But in the aggregation of a capital asset pricing model, those with greater willingness to bear risk ought to receive greater weight, since they tend to take large positions in risky assets. More specifically, individual asset demands—which can be aggregated by simple addition—involve a term that is proportional to risk tolerance. In the consumption capital asset pricing model, the level of asset demand is implicitly determined by the covariance of consumption with the return on the asset. [See Breeden 1979.] Let c_i be an individual's level of consumption, Z the return on any asset in excess of the safe rate of return, and θ_i the individual's risk tolerance. The consumption capital asset pricing model says that

$$(5) \quad E_t(Z) = \frac{1}{\theta_i} \text{Cov}_t\left(Z, \frac{\Delta c_i}{c_i}\right).$$

Consumption c_i is known at time t , so one can multiply through by $c_i\theta_i$ and sum over all households to get

$$(6) \quad E_t(Z) \sum_i c_i \theta_i = \sum_i \text{Cov}_t(Z, \Delta c_i).$$

Denote aggregate consumption as C and define aggregate risk tolerance by the consumption-weighted average $\Theta = \sum_i (c_i/C) \theta_i$. Then

$$(7) \quad E_t(Z) = \frac{1}{\Theta} \text{Cov}_t\left(Z, \frac{\Delta C}{C}\right)$$

because covariance is a linear operator. Equation (7) has the same form as equation (5), but with aggregate consumption and risk tolerance replacing the individual values.

In the first column of Table XI, we show average relative risk tolerance, computed using the numerical assignments and with the theoretically mandated consumption weights proxied by equal weights, income weights, and wealth weights, respectively. In the second column, we show the arithmetic weighted average of relative risk aversion. Jensen's inequality is strongly operative. The difference between mean risk tolerance and the reciprocal of mean risk aversion is substantial. For the entire sample (top panel of Table XI), the reciprocal of the mean risk tolerance of 0.24 equals 4.2, while mean risk aversion is 12.1. This result is not very sensitive to the weighting. Hence, the heterogeneity we find implies a dispersion in risk preferences that is large enough to make the difference between the arithmetic mean of risk aversion (12.1) and the harmonic mean of risk aversion (4.2) an important one. These levels of risk aversion are not high enough to explain the level of the equity premium.

A large group of individuals do not hold stock at all. If these individuals hold zero net positions because of fixed costs of being in the stock market and/or constraints on short sales (rather than because

their unconstrained optimum for stockholding is precisely zero), they require special treatment.

Formally, using $\tilde{\Theta}$ to represent the consumption-weighted average of risk tolerance

where nonstockholders have their risk tolerance replaced by zero,

$$(8) \quad E_t(Z) = \frac{1}{\tilde{\Theta}} \text{Cov}_t(Z, \Delta C/C)$$

if nonstockholders have consumption that is uncorrelated with the stock return Z . If nonstockholders on average have consumption that is positively correlated with stock returns, as Mankiw and Zeldes (1991) find—thereby having an implicit position in equities through the correlation of equities with macroeconomic events—then the right-hand side of (8) is an upper bound for $E_t(Z)$.

In the middle panel of Table XI, we repeat our previous calculations, assigning nonstockholders zero risk tolerance, to get an estimate of $\tilde{\Theta}$. (In the last panel of Table XI, we show risk tolerance for stockholders only. These calculations fairly closely replicate the numbers for all respondents without special treatment of nonstockholders.) Although the precise risk tolerance estimates are two to three times larger for the whole sample than they are when nonstockholders are assigned zero risk tolerance, the qualitative conclusions above continue to hold. Aggregate risk tolerance is low enough so that the equity premium remains a puzzle.

V. Preferences over Consumption Paths

In this section, we report the results of our experimental survey questions designed to elicit estimates of the preference parameters governing intertemporal substitution and time preference. As with the other experimental modules, Module K was administered to a very small subset of the HRS respondents: There are 198 respondents. In contrast, there are more than 11,000 responses to the risk preference questions discussed in the previous sections. In light of the small sample, the results should be regarded as tentative. Nonetheless, we can characterize our results broadly as follows:

- Most individuals have low elasticities of intertemporal substitution. Our point estimate for the average elasticity of intertemporal substitution is 0.18. Virtually no respondents have intertemporal substitution as elastic as that implied by log utility.
- Although the mean elasticity of intertemporal substitution is only slightly less than the mean risk tolerance, there is essentially no relationship between individuals' estimated elasticity of intertemporal substitution and relative risk tolerance.
- At a zero interest rate, the modal time preference is for a flat consumption path, but an upward slope is chosen more often than a downward slope. Hence, the mean preference is for an upward sloping consumption path.

The remainder of this section presents these results in more detail and discusses some of the problems that arose in the experimental survey.

Unlike the risk preference questions, the questions about the slope of the consumption path had possible responses that are either uninformative about the elasticity of intertemporal substitution or are inconsistent with utility maximization. For example, respondents who—regardless of the interest rate—chose either the extreme positive slope or the extreme negative slope could have any elasticity of intertemporal substitution. Since they are at a corner (given the range of choices we present), we do not learn anything about their willingness to substitute intertemporally when confronted with different interest rates. Of the 198 respondents to the module, 24 (12 percent) gave such uninformative responses. Since they convey no information about intertemporal substitution, they are left out of those tabulations.

The set of possible responses to the questions on the module left open the possibility of responses that were inconsistent with utility maximization. While we leave these out of the tabulations, it is important to examine the nature and extent of these inconsistent answers. The first question offered three consumption profiles. This was meant as a warm-up to familiarize the respondents with the form of the questions. The second question offered the same three choices plus two intermediate possibilities. Sixteen (8 percent) of the respondents made inconsistent choices. We eliminated these respondents from the tabulation even if their subsequent responses were otherwise consistent. Another 42 respondents (21

percent) displayed other inconsistencies in the subsequent choices. Specifically, these involved changing the slope of the desired consumption path in the direction opposite to the change in the interest rate.

There was, in particular, some tendency for respondents to react perversely when moving to a negative real rate, implying a negative elasticity of intertemporal substitution, although not one very different from zero. We excluded these observations. Including these observations in an analysis allowing for response noise would pull down the already low estimate of the elasticity of intertemporal substitution.

Intertemporal Substitution. Excluding the responses that are either uninformative about intertemporal substitution or are inconsistent, we are left with 116 useful observations. Just as with the risk preference question, the discrete nature of the survey questions leads to responses that correspond to ranges of preference parameters. Because we present the respondents with a fairly rich set of consumption profiles, the responses cannot be categorized into a few, non-overlapping groups, as they were for the risk preference questions. The respondents are presented five slopes for each of the three interest rates (zero, positive, and negative).

For each valid response, we calculate the range of possible elasticities of intertemporal substitution and most-desired slopes of the consumption path at a zero interest rate that are consistent with the responses. These calculations are analogous to the ranges for the risk preference parameters given in Table I. Table XII summarizes the preference parameters of the respondents. The first row gives the elasticity of intertemporal substitution. The second row gives time preference as measured by the slope of the desired consumption path at a zero interest rate. For each respondent's answers to the questions in the module, we calculate the lower bound and the upper bound of both parameters. The averages across respondents of these lower and upper bounds are reported in the first two columns of Table XII. The third column reports the average of the midpoints between these upper and lower bounds.

The average of the estimated lower bounds of the intertemporal elasticity of substitution is very close to zero. The average of the estimated upper bounds is 0.36. The average midpoint is 0.18. The average lower bound is as low as it is because the responses for most households (103 of the 116 valid responses) are consistent with a zero elasticity of intertemporal substitution. Indeed, the most common response was to choose a flat consumption profile regardless of whether the interest rate was zero,

positive, or negative. The next most common response was to choose the moderately upward sloping path for each interest rate. Eighty-four (72 percent) of the valid responses fell into these two groups. Because the interest rate is varying across the paths, these responses provide a fairly tight upper bound on the elasticity of intertemporal substitution. For those always preferring a flat consumption profile, the upper bound is 0.23; the upper bound is 0.29 for those always preferring the moderately upward sloping profile.

The remaining 28 percent of the valid responses did not fall into tight groupings. Many showed a higher elasticity of intertemporal substitution—with a midpoint as high as 1.08 for one respondent. Yet, even among those for whom intertemporal substitution is bounded away from zero, preferences are rarely as elastic as log utility. Only 2.5 percent of respondents had a maximum elasticity of intertemporal substitution greater than or equal to one.

Time Preference. Table XII also gives results for the time preference parameter. The overall average slope of the desired consumption path at a zero interest rate is 0.78 percent per year.³¹ Thus, we confirm the findings by experimental and cognitive psychologists that there is some evidence for a negative time-discount rate—on average, people prefer an upward sloping consumption profile, even when the interest rate is zero.³²

Intertemporal Substitution versus Risk Tolerance. Many applications assume that a representative consumer maximizes a time- and state-separable utility function with the period-by-period function having the constant elasticity functional form. In this context, relative risk tolerance (the reciprocal of relative risk aversion) equals the elasticity of intertemporal substitution. Selden [1978] and Epstein and Zin [1989] have discussed preferences where the individual's elasticity of intertemporal substitution is not equal to risk tolerance. Weil [1990], Hall [1988], and Barsky [1989] explore the implications of such preferences for consumption and asset pricing while Kimball and Weil [1992] discuss the implications of these preferences for precautionary saving.³³ For the respondents to our experimental survey, we can test directly whether there is a relationship between risk tolerance and intertemporal substitution.

Table XIII tabulates the consumption path parameters by the four responses to the risk preference question. There is no significant relationship—either statistically or economically—between risk

tolerance and intertemporal substitution. Similarly, if we regress the estimated elasticity of intertemporal substitution (measured by the midpoint of the range of possible values for each respondent) on the mean risk tolerance, we get a coefficient of 0.01 with a standard error of 0.02. Under the usual assumption that risk tolerance equals the elasticity of intertemporal substitution, the regression coefficient should be one. Hence, it appears that there is no relationship between measured intertemporal substitution and measured risk tolerance even though their means are similar.

Given the scant number of observations from this module, we do not consider tabulations with the demographic and behavioral variables.

VI. Extensions and Qualifications

Our analysis of the survey responses is based on several presumptions. First, we make assumptions about how the individuals interpret the questions. Second, in the case of the risk tolerance survey responses, we use identifying assumptions to estimate a parametric model of relative risk tolerance. In this section, we investigate the robustness of the results under departures from some of these assumptions. In particular, we investigate how a preference for an existing job might cause respondents to give responses that understate their risk tolerance and how our assumption of no persistent response error would lead to an overstatement of risk tolerance. Instead of simply noting the direction of these offsetting biases, we attempt to quantify them. This quantification requires further parametric specification.

A. Status Quo Bias

The survey question asks whether the respondent would switch to a job that is “equally good” except for the income risk. Nonetheless, the survey respondents might imagine that there is a cost to switching jobs, might have a nonpecuniary value to their job, or might simply have an excess propensity to decline the gamble because doing so preserves the status quo. This status quo bias would cause our results to understate risk tolerance to the extent that individuals are disinclined to accept the gambles for reasons other than their preference toward risk.³⁴

Status quo bias can be analyzed using the theoretical model developed above. Suppose that individuals place a premium on their current job above the consumption flow that it allows them to sustain. That is, we can imagine individuals responding to the survey based on whether

$$(9) \quad \frac{1}{2}U(2c) + \frac{1}{2}U(\lambda c) \geq U(\phi c),$$

where ϕ parameterizes the status quo bias in terms of consumption flows. ϕ equal to one is our baseline case of no status quo bias. If ϕ is greater than one, the respondents will accept a gamble that also entails switching jobs if it delivers an expected utility strictly in excess of the current job.

In Table XIV we present results for various values of ϕ . The top line gives the results for our baseline specification of no status quo bias. The next two lines show how the estimated distribution of relative risk tolerance changes if it takes a 5 percent ($\phi = 1.05$) or 10 percent ($\phi = 1.1$) consumption premium to make the respondent indifferent between his or her current job and a different job with a certain level of income. Allowing for 5 percent status quo bias raises the estimate of mean risk tolerance from 0.241 to 0.299. It also increases the estimated dispersion substantially. Allowing for 10 percent status quo bias raises mean risk tolerance further, to 0.396. As status quo bias increases, respondents expressing a willingness to switch jobs are assigned substantially higher risk tolerance conditional on their survey response; i.e., they must be quite risk tolerant if their willingness to undertake the income gamble more than compensates for a general unwillingness to switch jobs.

Because we do not have an estimate of ϕ , the results of Table XIV are conjectural.³⁵ It would be straightforward to reword the question to eliminate the status quo bias in future surveys. Instead of asking about the current job versus a different job with risky lifetime income as we did in the first two waves of the HRS, one could ask about preferences between two new jobs, one with a certain income stream the same as the current job and the other with the risky income stream. Hence, the choice set would mandate a change in job regardless of risk preference, and the responses would be free from the bias owing to a preference for the existing job.

We are proposing to ask this status quo bias free question on future waves of the HRS. Moreover, we have conducted a pilot study of status quo bias free questions using a sample of University of Michigan undergraduates. Half the students were given the question with the original wording (choice between a hypothetical current job and a risky new job); half were given the reworded question (choice between two new jobs with current or risky income). Using the conditional expectations estimated for the HRS respondents, the mean relative risk tolerance for those responding to the original question was 0.26—remarkably close to what we found for the HRS sample. For those answering the reworded question, the mean was 0.34. Hence, status quo bias does appear to lead our results to understate the level of risk tolerance.³⁶ The results of the pilot study can be expressed in terms of our model given in equation (9). Based on the statistical model estimated with the HRS data, the difference of the responses of the two groups of students is consistent with a status quo premium of between 5 and 10 percent of lifetime consumption.

B. Persistent Measurement Error

Our statistical model of the survey responses attributes all the correlation in responses across waves to the true preference parameter. Put differently, we assume that all randomness in responses is temporary. While the model is not identified if the persistence of the measurement error is a free parameter, it is possible to calculate its implications for different assumptions about the persistence of measurement error. Consider a modification to the statistical model [equation (4)],

$$(10) \quad y_{jk} = x_j + \eta_j + \epsilon_{jk},$$

where η_j is a persistent error component. Define $\tau = \sigma_x^2 / (\sigma_x^2 + \sigma_\eta^2)$, that is, the fraction of variance of the persistent component of the response due to the true preference parameter.

In our baseline model, there is no η_j term, so $\tau = 1$. The fourth line of Table XIV presents a reestimation of the model where half of the persistent signal is error ($\tau = 0.5$). The persistent measurement error reduces the estimate of mean risk tolerance from 0.241 in the baseline model to 0.185. Moreover, it pulls the expectations conditional on the survey response toward the unconditional mean.³⁷

In the limit as τ approaches zero, the survey would contain no information about the heterogeneity of risk tolerance, although it could still provide information about mean risk tolerance. The last two lines of Table XIV present results with both status quo bias and persistent measurement error. For the parameters in the Table, the two have substantially offsetting effects on the estimated distribution of risk tolerance.

C. Discussion

The previous two subsections discuss status quo bias and persistent measurement error. We are able to quantify the effect of these biases on the estimated distribution of relative risk tolerance. The biases are offsetting: correcting for status quo bias raises the estimate of relative risk tolerance while correcting for persistent measurement error lowers it. The survey does not contain enough information to estimate the extent of either of these biases, so the quantification is based on conjectured values of the relevant parameters. Based on a pilot study, we have some information on the extent of status quo bias. This pilot study suggests that the estimate of relative risk tolerance implied by our baseline model should be adjusted upward substantially. Future research, such as asking the risk tolerance question on future waves of the HRS in the status quo bias free form, is needed to get better estimates of the relevant parameters.

As the discussion in this section and the results of Table XIV make clear, the estimates of the distribution of relative risk tolerance—based on our parametric model of the utility function and the statistical distribution of its parameter—are quite sensitive to the assumptions needed to implement it. Yet, we hasten to add that many of the results of the paper do not depend on the parametric model. The qualitative information on the distribution of risk tolerance and how it relates to demographic characteristics and behaviors reported in Tables II through IV and VI through IX does not depend on the parametric model. Moreover, our finding that the estimate of the elasticity of intertemporal substitution is independent of the risk tolerance response (Table XIII) does not depend on the parametric model.

Moreover, the modest explanatory power of the risk tolerance responses for the behaviors reported in the regressions is not affected by the choice of parametric model. The conditional expectations of relative risk tolerance according to the survey response are approximately linear across the various rows of Table XIV. The last column gives the adjustment factor by which the results of the

baseline model should be multiplied in each case. Choice of parametric model will change the regression coefficient of relative risk tolerance by the factor reported in Table XIV, but will not change the t-statistic or the goodness of fit of the regressions.

The regressions reported in Table V are simply prediction equations. A linear transformation of the risk tolerance measure has no effect on the prediction for behavior of being in one of the risk tolerance categories versus another. The regressions reported in Table X, however, have a structural interpretation that depends on the level of risk tolerance. As noted above, the coefficient of 0.097 for risk tolerance in the equation for stock (Table X) implies an elasticity of stock demand with respect to risk tolerance of 0.17, less than its theoretically mandated level of one. Based on the mean risk tolerances and adjustment factors reported in Table XIV, correcting for status quo bias of 5 percent and 10 percent reduces this elasticity to 0.16 and 0.13, respectively. Correcting for persistent measurement error makes a bigger difference. Assuming τ equal to 0.5, the estimated elasticity of the share of stocks becomes 0.37 assuming no status quo bias, 0.36 assuming a 5 percent status quo bias, and 0.30 assuming a 10 percent status quo bias. The elasticities for the other risky assets would be modified by the same factors. In essence, if one believes in substantial persistent measurement error, some fraction of the difference between the theoretically mandated elasticity and that estimated in the baseline model can be seen as a consequence of the persistent measurement error.

VII. Conclusion

This paper reports the results of experimental questions designed to elicit measures of risk tolerance, the elasticity of intertemporal substitution, and time preference. The measures concern preferences over behaviors that are central to macroeconomics and finance, namely willingness to take gambles over lifetime income and to substitute consumption over long periods. The parameters are estimated as part of the Health and Retirement Study. Estimating the preference parameters as part of a large-scale survey has several advantages. First, the estimated preference parameters can be related to the behaviors that they should predict. The economics profession is skeptical about subjective questions and answers. Being able to relate the subjectively estimated preference parameters to tangible behavior

should address some of this skepticism. Second, to the extent that the estimated parameters do predict behavior, they might be useful in many applications of the survey database.

We find that there is substantial heterogeneity in preference parameters. Although most of our respondents are in our least risk tolerant category, many are substantially more risk tolerant. Theory predicts very different behavior toward risk for agents with these varying degrees of risk tolerance. We have some success in relating these estimates to different behaviors. For example, the risk tolerance measure is related in the way one would expect to whether a respondent smokes, drinks heavily, has no health or life insurance, or holds stocks and other risky assets. Indeed, for virtually every behavior we investigate, the risk tolerance measure made qualitatively correct predictions. The regression coefficients are large in their implications for behavior. Yet, there is tremendous variability in the behaviors, so only a small fraction of their variance is explained by risk tolerance (or any covariate). This finding of a common factor in behavior, but one that leaves most of the differences between individuals unexplained, is common in the psychological literature.

University of Michigan and National Bureau of Economic Research

References

- Barsky, Robert B., "Why Don't the Prices of Stocks and Bonds Move Together?" American Economic Review, LXXIX (1989), 1132-45.
- Binswanger, Hans P., and Donald A. Sillers, "Risk Aversion and Credit Constraints in Farmers' Decision-Making: A Reinterpretation," Journal of Development Studies, XX (1983), 5-21.
- Breeden, Douglas T., "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities," Journal of Financial Economics, VII (1979), 265-96.
- Campbell, John Y., and John H. Cochrane, "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," National Bureau of Economic Research Working Paper No. 4995, January 1995.
- Carroll, Christopher D., "How Does Future Income Affect Current Consumption?" Quarterly Journal of Economics, CIX (1994), 111-47.
- Cloninger, C. Robert, Thomas R. Przybeck, and Dragan M. Svrakic, "The Tridimensional Personality Questionnaire: U.S. Normative Data," Psychological Reports, LXIX Pt 1 (1991), 1047-57.
- Epstein, Larry G., and Stanley E. Zin, "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," Econometrica, XLVII (1989), 937-69.
- Friedman, Bernard, "Risk Aversion and the Consumer Choice for Health Insurance," Review of Economics and Statistics, LVI (1973), 209-14.
- Friedman, Milton, A Theory of the Consumption Function (Princeton, NJ: Princeton University Press, 1957).
- Fuchs, Victor R., "Time Preference and Health: An Exploratory Study," in V.R. Fuchs, ed., Economic Aspects of Health (Chicago, IL: University of Chicago Press, 1982).
- Grossman, Sanford, and Robert Shiller, "Determinants of the Variability of Stock Prices," American Economic Review, LXXI (1981), 222-27.
- Hall, Robert E., "Intertemporal Substitution in Consumption," Journal of Political Economy, XCVI (1988), 339-57.
- Hansen, Lars P., and Kenneth J. Singleton, "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns," Journal of Political Economy, XCI (1983), 249-65.
- Hausman, Jerry A., "Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables," Bell Journal of Economics, X (1979), 33-54.
- Hill, Daniel H., "Respondent Rules and the Quality of Net Worth Data in the Health and Retirement Survey," Institute for Social Research, Health and Retirement Study Working Paper No. 94-002, December 1993.

- Juster, F. Thomas, and Richard Suzman, "An Overview of the Health and Retirement Study," Journal of Human Resources, XXX (1995, supplement) S7-56.
- Kahneman, Daniel, and Amos Tversky, "Prospect Theory: An Analysis of Decision Under Risk," Econometrica, XLVII (1979), 263-91.
- Kahneman, Daniel, and Amos Tversky, "The Framing of Decisions and the Psychology of Choice," Science, CCXI (1981), 453-58.
- Kahneman, Daniel, and Amos Tversky, Judgment Under Uncertainty: Heuristics and Biases (New York, NY: Cambridge University Press, 1982).
- Kandel, Shmuel and Robert Stambaugh, "Asset Returns and Intertemporal Preferences," Journal of Monetary Economics, XXVII (1991) 39-71.
- Kimball, Miles S., and Matthew D. Shapiro, "Estimating a Cardinal Attribute from Ordered Categorical Responses Subject to Noise," unpublished paper, University of Michigan, June 1996.
- Kimball, Miles S., and Philippe Weil, "Precautionary Saving and Consumption Smoothing Across Time and Possibilities," National Bureau of Economic Research Working Paper No. 3976, January 1992.
- Laibson, David, "Hyperbolic Discount Functions, Undersaving, and Savings Policy," unpublished paper, Harvard University, 1996.
- Lawrance, Emily C., "Poverty and the Rate of Time Preference: Evidence from Panel Data," Journal of Political Economy, IC (1991) 54-77
- Loewenstein, George, "Anticipation and the Valuation of Delayed Consumption," Economic Journal, XCVII (1987), 666-84.
- Loewenstein, George, and Drazen Prelec, "Negative Time Preference," American Economic Review Papers and Proceedings, LXXXI (1991), 347-52.
- Loewenstein, George, and Drazen Prelec, "Anomalies in Intertemporal Choice: Evidence and an Interpretation," Quarterly Journal of Economics, CVII (1992), 573-97.
- Loewenstein, George, and Richard H. Thaler, "Anomalies: Intertemporal Choice," Journal of Economic Perspectives, III (Fall 1989), 181-93.
- Maital, Shlomo, ed. Applied Behavioral Economics (New York, NY: New York University Press, 1988).
- Maital, Shlomo, and Sharona Maital, "Time Preference, Delay, and Gratification and the Intergenerational Transmission of Economic Inequality: A Behavioral Theory of Income Distribution," in Orley P. Ashenfelter and Wallace E. Oates, eds., Essays in Labor Market Analysis in Memory of Yochanan Peter Comay (New York, NY: Wiley, 1977).

- Mankiw, N. Gregory, "The Permanent Income Hypothesis and the Real Interest Rate," Economics Letters, VII (1981), 307-311.
- Mankiw, N. Gregory, and Stephen P. Zeldes, "The Consumption of Stockholders and Non-Stockholders," Journal of Financial Economics, XXIX (1991), 97-112.
- Mehra, R., and Edward C. Prescott, "The Equity Premium: A Puzzle," Journal of Monetary Economics, XV (1985), 145-61.
- Menza, Matthew A., Lawrence I. Golbe, Ronald A. Cody, and Nancy E. Forman, "Dopamine-related Personality Traits in Parkinson's Disease," Neurology, XLIII Pt 1 (1993), 505-8.
- Mischel, Walter, Introduction to Personality (New York, NY: Holt, Rinehart, and Winston, 1971).
- Samuelson, William, and Richard Zeckhauser, "Status Quo Bias in Decision Making," Journal of Risk and Uncertainty, I (1988), 7-59.
- Selden, L., "A New Representation of Preferences over 'Certain x Uncertain' Consumption Pairs: The 'Ordinal Certainty Equivalent' Hypothesis," Econometrica, XLVI (1978), 1045-60.
- Shiller, Robert J., Maxim Boycko, and Vladimir Korobov, "Hunting for Homo Sovieticus: Situational versus Attitudinal Factors in Economic Behavior," Brookings Papers on Economic Activity, (1:1992), 127-81.
- Slovic, Paul, "Psychological Study of Human Judgement: Implications for Investment Decision Making," Journal of Finance, XXVII (1972), 779-99. (a)
- Slovic, Paul, "Information Procession, Situation Specificity, and the Generality of Risk-Taking Behavior," Journal of Personality and Social Psychology, XXII (1972), 128-34. (b)
- Szpiro, George G., "Measuring Risk Aversion: An Alternative Approach," Review of Economics and Statistics, LXVIII (1986), 156-9.
- Viscusi, W. Kip and William N. Evans, "Utility Functions that Depend on Health Status: Estimates and Economic Implications," American Economic Review, LXXX (1990), 353-74.
- Weil, Philippe, "Unexpected Utility in Macroeconomics," Quarterly Journal of Economics, CV (1990), 29-42.

Notes

1. See, for example Kahneman and Tversky [1979, 1981, and 1982].
2. We tested preliminary versions of the survey instruments on two groups. Versions of the questions were first given to undergraduate economics students. Based on the student responses, we refined the questions. They were then tested as part of the standard Survey Research Center procedure for testing survey instruments. This phase of testing is meant to uncover difficulties respondents might have in interpreting the questions.
3. The questions ask about income rather than spending or consumption. After pre-testing, we concluded that survey respondents would better understand income than consumption lotteries. Given the low levels of financial wealth of most respondents, permanent labor income and permanent income are not that different. We investigate (see below) the extent to which high-wealth and older individuals respond differently to the questions. See question L14 of the Health and Retirement Study, Wave I (page 162 of the survey instrument).
4. The numerical results are not that sensitive to the choice of a constant relative risk aversion parameterization. For example, with constant absolute risk aversion, the bounds on risk tolerance consistent with the responses would be 0.30, 0.55, and 1.04 instead of the values of 0.27, 0.5, and 1.0 given in Table I for constant relative risk aversion.
5. Equation (2) is nominally the same Euler equation routinely estimated on time series data by macroeconomists [Hansen and Singleton 1983, Mankiw 1981, and Hall 1988]. It can be derived by assuming a time-separable, constant relative risk aversion utility function, so that θ equals s . Our survey design does not depend on a particular maximization problem. In particular, θ need not equal s . Indeed, we will compare the results of the risk preference and intertemporal substitution questions to test this restriction.

6. The HRS is a representative sample within this age group, except that blacks, Hispanics, and residents of Florida are 100 percent oversampled. See Juster and Suzman [1995] for an overview of the survey. Further information and public-release data are available on the world wide web at <http://www.umich.edu/~hrswww>.

7. The first person contacted is asked to identify the most knowledgeable member of the family. There is a slight propensity for the first person contacted to over-report himself or herself as the most knowledgeable. For this point, and for a general analysis of this feature of the HRS, see Hill [1993]. Members of couples need not be married. In our sample, one-third of the secondary respondents are males (see Appendix Table I).

8. The fraction of respondents who are black and Hispanic represents the oversampling of these groups in the HRS. Further demographic breakdowns are given in the tables below where we report the results for the risk preference question.

9. See Kimball and Shapiro [1996] for a more complete discussion of this statistical issue.

10. We do not use any other information from Wave II.

11. An alternative would be that there is no noise in the preference parameter, but the individual uses noisy cutoffs. The former interpretation implies some noise in preferences. The latter places the noise in interpreting the questions. These two interpretations yield the same statistical model.

12. The univariate distributions in the two waves are nearly identical. Hence, for example, risk tolerance is not drifting as the panel ages.

13. The estimation uses all the observations, including those individuals who answer only one of the two waves. The likelihood is constructed from the trivariate normal distribution of x_j , y_{j1} , and y_{j2} and the bivariate normal distribution of x_j and y_{jk} for those who answered only $k = \text{Wave I or Wave II}$. Those

who answer only one wave help us estimate the distribution of y , while those who answer both allow us to identify the distributions of x and ϵ . The distributions are integrated over the truncation intervals B for y to yield the likelihood.

14. Standard errors of the estimated parameters are approximately 0.01.

15. This expectation is computed by integrating e^{-x} or e^x over the joint probability distribution of x and y . It is also possible to compute the expectation conditional on the survey response. We use these values below in the regression analysis.

16. When an individual responded to both Wave I and Wave II, we condition on both responses to assign the expected risk tolerance.

17. Shiller, Boycko, and Korobov [1992] report the results of a survey question similar to ours. They asked a small sample of respondents from different countries whether they would be willing to take a job at a 50 percent higher wage than their current job if there were a 50-50 chance of failing at the job. In the event of failure, the respondent would get his or her old job back "after some time." Since the bad outcome entails only a temporary loss, this proposition is much less risky than ours. (The aim is to elicit job-market flexibility, not risk tolerance.) They find that fifty to eighty percent of respondents would take the new job, with Russians and West Germans less willing to take the new job than those in the United States. The unwillingness of many to face even a temporary income loss for the chance of a large, permanent gain would imply a high level of risk aversion.

Binswanger and Sillers [1983] report the results of experiments where participants faced relatively large risks, ranging from 2.5 to 143 weeks of wages, but are still small relative to lifetime resources. Absent enough information to calculate the respondents' marginal propensities to consume, it is impossible to directly relate these estimates to ours. Indeed, Binswanger and Sillers emphasize the role of credit constraints in interpreting their results. Our questions about lifetime resources are designed

to circumvent the need to know the marginal propensity to consume.

18. The fractiles of the distribution of relative risk tolerance (θ) are as follows:

Fractile	0.05	0.10	0.25	0.50	0.75	0.90	0.95
θ	0.03	0.04	0.07	0.14	0.28	0.53	0.77

19. This correlation is based on the conditional expectations from our parametric model of risk tolerance. The rank correlation, which is independent of our parametric model, is 0.13. Some of this correlation arises from one spouse having heard the other's response to the same question. On the other hand, correcting for measurement error would increase the estimated correlation. We have no strong prior belief about the degree of correlation of the preference parameters of married individuals. It is not clear that risk tolerance would be a key variable upon which there is assortative mating.

20. Psychologists typically find that survey measures explain only a small fraction of individual behavior. Mischel, in connection with his discussion of the "personality coefficient," notes that the fraction of cross-sectional variation in a specific behavior that can be accounted for by responses to a survey questionnaire typically ranges from .04 to .09. [See Mischel 1971, pp. 147-8.]

21. Smoking and drinking have complicated effects on the distribution of future income. Smoking and immoderate drinking reduce mean life expectancy, and hence have a negative effect on expected income. They also increase income variance by increasing the probability that an individual will have a serious disease.

22. Researchers have studied attitudes about health-related risk and examined how these interact with economic choices. The relationship between our risk tolerance measure and smoking and drinking corroborates the findings that individuals translate health risks into pecuniary values. Viscusi and Evans [1990] estimate that workers show rather smooth, concave tradeoffs between occupational health and safety risks and consumption. Fuchs [1982] shows that the substantial heterogeneity in responses to

questions about money or commodities now versus the corresponding desiderata in the future has predictive power for the decision to smoke, and that some, but not all of this is mediated through education.

23. There is no obvious prediction about the risk tolerance of those not working—mainly retired individuals and spouses not in the labor force.

24. Researchers have used choices about insurance to elicit estimates of risk aversion. Friedman [1973] used data on choices regarding health insurance, and obtained an estimate of about 10. Szpiro [1986] returns to the idea of gauging risk aversion by studying the demand for insurance. He looks at households' willingness to pay a load factor in order to obtain insurance, using insurance company data on premiums and claims. Using these data, along with the Goldsmith data on total household wealth, Szpiro reports estimates of the coefficient of relative risk aversion between one and two. While these studies are clearly related to our results, their method is to estimate risk aversion from purchase of insurance while our survey creates an independent measure of risk aversion, which can then be related to purchase of insurance.

25. Older and high wealth individuals might interpret the survey questions differently from most respondents because labor income is a smaller fraction of their current resources. We checked for this possibility by grouping the responses by both age and wealth quintile. These groupings do not lead to the conclusion that the highly risk tolerant respondents are either old or wealthy. Moreover, we reran the regressions in Table V including the logarithms of income and wealth as regressors. Controlling for income and wealth raises some coefficients of risk tolerance and lowers others, but overall has little qualitative impact on the findings. (We report the regressions without wealth and income in Table V, owing to concern about the endogeneity of those variables.)

26. Some of the portfolio shares are zero. Tobin's Separation Theorem implies, however, that they

should all be positive. Hence, the zeros cannot be modeled as Tobit-like truncation. They may result from a fixed cost of holding a particular asset, which would imply jumps from zero to strictly positive portfolio shares. We leave such modeling to future research.

27. If there is no secondary respondent, we code the difference as zero. If the secondary respondent did not answer the risk tolerance question, we also code the difference as zero. If the primary respondent did not answer, we exclude the household from the sample. The values of the other individual-specific covariates refer to the primary respondent.

28. The regression coefficient of 0.097 corresponds to a squared correlation of 0.02 between the risk tolerance measure and the stock portfolio share—somewhat below the range commonly reported in the psychological literature for the fraction of variance explained by a battery of survey measures. We also estimated the portfolio share equations by Heckman's two-step, Tobit estimator. As would be expected, the Heckit estimates are larger than the OLS estimates. Indeed, for the equation for the stock share, the coefficient of relative risk tolerance is twice the OLS estimate.

29. The Health and Retirement Study fails to provide any information about the asset composition of retirement accounts, so we do not know their riskiness. Given the growing importance of retirement accounts and defined contribution pension plans, it is important that future surveys provide information about the composition of these accounts.

30. Kandel and Stambaugh [1991] argue for precisely this interpretation of the equity premium puzzle.

31. Because s is close to zero, we focus on the estimate of $-s \cdot \rho$ instead of dividing by s to get an estimate of ρ . See equation (2).

32. See Loewenstein [1987], Loewenstein and Prelec [1991, 1992], Loewenstein and Thaler [1989], Maital [1988], and Maital and Maital [1977]. In contrast, the econometric evidence (e.g. Hausman

[1979] and Lawrance [1991]) finds downward sloping profiles. Future research is needed to explain why the econometric and experimental evidence arrive at different conclusions. One possible explanation of the finding of high subjective discount rates in the econometric work is the difficulty of controlling for features of the economic environment facing agents, such as liquidity constraints and the need for precautionary savings.

33. Certain utility functions that display habit formation or consumption externalities also can break the link between risk tolerance and intertemporal substitution. This is an important feature of Campbell and Cochrane's [1995] analysis. Andrew Abel's discussion of that paper at the October 1994 NBER Economic Fluctuations meeting clarified this point for us. Laibson [1996] uses a novel form of the utility function which also has the effect of breaking this link.

34. See Samuelson and Zeckhauser [1988] for a study of status quo bias.

35. Samuelson and Zeckhauser [1988] quantify the extent of status quo biases. They asked student survey respondents their opinions about various public policies and their responses to hypothetical economic decisions. They pose the questions with and without status quo bias. Samuelson and Zeckhauser find that respondents were 17 percentage points more likely to choose an option that was presented as the status quo than if it was expressed neutrally. They also report that status quo bias appears to be not just a phenomenon of survey responses, but to affect economic decisions such as selection of health plans and retirement portfolios. Therefore, some of the unwillingness to make risky job changes that we find in the survey might reflect actual behavior, but not be due to risk tolerance.

36. This survey was conducted in the fall of 1996 in an intermediate macroeconomic theory class. We gratefully acknowledge the collaboration of James Andrew Kovacs in conducting this survey. It is difficult to extrapolate the results based on the students to the sample of older individuals in the HRS. In particular, status quo bias might be more severe for the HRS respondents because they are likely to have

actually been in their current jobs for a long time. Yet, the extent of status quo bias that we find in our pilot is quite close to that reported by Samuelson and Zeckhauser [1988]. In particular, 14 percent of respondents declined to take any gamble in the survey with status quo bias while only 3 percent declined any gamble in the status quo bias free survey.

37. This reduction in variance also reduces the wedge between the arithmetic and harmonic means of risk aversion arising from Jensen's inequality.