

Imputing Risk Tolerance From Survey Responses

Miles S. KIMBALL, Claudia R. SAHM, and Matthew D. SHAPIRO

Economic theory assigns a central role to risk preferences. This article develops a measure of relative risk tolerance using responses to hypothetical income gambles in the Health and Retirement Study. In contrast to most survey measures that produce an ordinal metric, this article shows how to construct a cardinal proxy for the risk tolerance of each survey respondent. The article also shows how to account for measurement error in estimating this proxy and how to obtain consistent regression estimates despite the measurement error. The risk tolerance proxy is shown to explain differences in asset allocation across households.

KEY WORDS: Interval regression; Measurement error; Ordered probit with known bounds; Proxy variable; Response error; Risk aversion; Risk tolerance; Surveys.

1. INTRODUCTION

Choices with uncertain outcomes, such as financial investments, career paths, and health practices, are numerous and important to welfare. Empirical studies of these behaviors often suffer from a common weakness—the inability to take into account heterogeneity in preferences. In this article, we develop a quantitative proxy for risk tolerance based on responses from a large-scale survey to account for this heterogeneity. We then use the proxy to study asset allocation.

Our measurement of risk tolerance is based on individuals' responses to questions about hypothetical risky choices. In particular, we ask them to choose between a job with a certain lifetime income and a job with a random, but higher mean lifetime income. We show how to translate these ordinal responses into a cardinal proxy for risk tolerance. To construct this proxy and use it to study behavior, we confront a number of issues. First, the survey responses about gambles over lifetime income imply a range instead of a point value for the unobserved cardinal preference parameter. Second, the survey responses are likely to be subject to measurement error. We develop a statistical model addressing both issues. Multiple responses from some individuals and refinements to the survey questions isolate the true variation in risk preferences. With the maximum likelihood estimates, we compute the proxy value—the expectation of risk tolerance conditional on survey responses—for each individual. Because it is based on a small set of survey questions, the proxy may not fully capture the systematic variation in risk preferences. This induces a nonstandard errors-in-variables problem in regression estimates that use the proxy as an explanatory variable. We provide an estimator using the proxy that is consistent despite errors in variables.

Miles S. Kimball (E-mail: mkimball@umich.edu) and Matthew D. Shapiro (E-mail: shapiro@umich.edu) are Professors, Department of Economics, and Research Professors, Survey Research Center, University of Michigan, Ann Arbor, MI 48109, and Research Associates, National Bureau of Economic Research, Cambridge, MA 02138. Claudia R. Sahn is Economist, Division of Research and Statistics, Federal Reserve Board, Washington, DC 20551 (E-mail: claudia.r.sahn@frb.gov). Kimball and Shapiro gratefully acknowledge the support of National Institute on Aging grants 2-P01-AG10179 and 5-R01-AG020638. Sahn gratefully acknowledges the support of a National Institute on Aging Pre-Doctoral Training Fellowship and an Innovation in Social Research Award from the University of Michigan Institute for Social Research. The authors thank Dan Benjamin, John Bound, Yuriy Gorodnichenko, John Laitner, Atsushi Inoue, Lutz Kilian, Tyler Shumway, Martha Starr-McCluer, Serena Ng, Gary Solon, seminar participants at Osaka University, and anonymous referees for helpful comments. The views presented are solely those of the authors and do not necessarily represent those of the Federal Reserve Board or its staff. This article subsumes an earlier working paper by Kimball and Shapiro, "Estimating a Cardinal Attribute From Ordered Categorical Responses Subject to Noise."

The plan of the article is as follows. Section 2 discusses the survey questions on lifetime income gambles and the distribution of responses in the Health and Retirement Study. (See <http://hrsonline.isr.umich.edu> for information on the survey.) Section 3 shows how to construct the cardinal proxy for risk tolerance from these survey responses, and Section 4 addresses the presence of survey response error. Researchers will be able to use such a proxy as an explanatory variable in studying a wide range of behaviors. In Section 5 we show how to estimate consistently the effect of the preference parameter on behavior. Section 6 applies these procedures to study the asset allocation decision. Our results show that our improved measure of risk preference significantly alters the estimated effects of risk tolerance and other observable characteristics on asset allocation. The final section offers conclusions.

2. SURVEYING RISK PREFERENCES

The Health and Retirement Study (HRS) is a large-scale, biennial survey, which began in 1992 with a representative sample of individuals between the ages of 51 and 61 and their spouses. In addition to detailed financial and demographic information, the study elicits risk preferences using a battery of questions developed by Barsky, Juster, Kimball, and Shapiro (1997). The Panel Study of Income Dynamics, National Longitudinal Study, Surveys of Consumers, Dutch CentERpanel, and Chilean Social Protection Survey have also fielded these gambles over lifetime income. In hypothetical scenarios, respondents choose between a certain job and a risky job. With equal chances, the risky job will double lifetime income or cut lifetime income by a specific fraction (or downside risk). Varying the downside risk on the new job in subsequent questions refines the measure of risk preferences.

Specifically, in 1992 the HRS poses the following scenario:

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?

Individuals accepting this new, risky job then consider one with a higher downside risk:

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?

Those initially declining the new job consider one with a lower downside risk:

© 2008 American Statistical Association
Journal of the American Statistical Association
September 2008, Vol. 103, No. 483, Applications and Case Studies
DOI 10.1198/01621450800000139

Table 1. Risk tolerance response categories

Response category	Downside risk of risky job		Bounds on risk tolerance	
	Accepted	Rejected	Lower	Upper
1	None	1/10	0	.13
2	1/10	1/5	.13	.27
3	1/5	1/3	.27	.50
4	1/3	1/2	.50	1.00
5	1/2	3/4	1.00	3.27
6	3/4	None	3.27	∞

NOTE: Respondents choose between a job with a certain income and a job with risky income. With equal chances, the risky job will double lifetime income or cut it by the specific fraction shown in the columns labeled downside risk. The largest risk accepted and the smallest risk rejected across gambles define a response category. In 1992, there are four categories 1–2, 3, 4, and 5–6. In 1994 and later surveys, there are six response categories. The last two columns show the bounds on relative risk tolerance consistent with these response categories in the absence of response error.

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?

These two responses order individuals in four categories: unwilling to risk a one-fifth income cut, willing to risk at most a one-third cut, willing to risk a one-third to a one-half cut, and willing to risk at least a one-half cut. In 1994, a randomly selected subsample answered the questions again. In 1994 and later implementations, there were additional questions about the willingness to accept one-tenth and three-quarter cuts. With these additional gambles, there are six distinct response categories. The second and third columns of Table 1 relate these response categories to the downside risks of the new jobs. In Section 3 we will discuss the last two columns of Table 1, which relate the response categories to the preference parameter.

In general, the gambles over lifetime income reveal a low tolerance for risk. As reported in Table 2, almost two-thirds of the respondents in 1992 are in the least risk tolerant category 1–2. The remaining one-third of respondents divide almost equally among the other three categories. The distribution of risk categories in 1994 follows a similar pattern. Over 60% of respondents fall in category 1 or 2 with most choosing the least risk tolerant category 1.

Repeated observations from some individuals will be central to our statistical strategy for separating signal from noise in the survey responses. Among the 693 respondents who answer the gambles in both the HRS 1992 and 1994, the corre-

Table 2. Distribution of risk tolerance responses

Response category	% by HRS wave				
	1992	1994	1998	2000	2002
1					
2	64.6	18.1	19.0	18.4	18.6
3	11.6	13.5	17.0	14.4	15.3
4	10.9	14.5	10.8	8.1	9.6
5		6.3	8.0	7.5	6.1
6	12.9	4.2	7.3	5.3	5.6
Responses	11,592	717	796	884	3,591

NOTE: Tabulations use responses on the final release version of HRS 1992, 1994, 1998, 2000, and 2002 without sample weights. The sample for this article includes the 11,616 original respondents in the HRS study who answer a gamble in one of the first two waves. See Table 1 for definition of the risk tolerance response categories.

lation of the response categories across the two waves is .27, and almost half switch response categories. Altogether, the survey responses suggest substantial and persistent differences in risk preferences across individuals, but also large changes in responses within individuals across surveys.

The 1998 HRS introduced a new situational frame for the income gambles to remove the potential for status quo bias. In the original question, individuals choose between their *current* certain job and a *new* risky job. An unwillingness to switch jobs may reflect their aversion to the risky income at the new job or their desire to maintain the status quo. Status quo bias appears to be a common feature in many settings (Samuelson and Zeckhauser 1988). In the presence of status quo bias, estimates from the original question would understate individuals' true risk tolerance. Using a pilot study of undergraduates, Barsky et al. (1997) estimate average risk tolerance to be 24% lower with responses to the original question than with responses to an alternate question free of status quo bias. In 1998, 2000, and 2002, the HRS fielded a status-quo-bias-free question, in which individuals choose between two new jobs. The question wording is

Suppose that you are the only income earner in the family. Your doctor recommends that you move because of allergies, and you have to choose between two possible jobs. The first would guarantee your current total family income for life. The second is possibly better paying, but the income is also less certain. There is a 50–50 chance the second job would double your total lifetime income and a 50–50 chance that it would cut it by a third. Which job would you take—the first job or the second job?

As in the original version, follow-up questions vary the downside risk of the second job, and responses assign individuals to one of six categories. Starting in 2000, the job-related gambles are targeted to individuals less than age 65. The final three columns of Table 2 show the responses to the status-quo-bias-free question. In this article, we restrict the sample to original respondents of the HRS who answered the gambles in 1992 or 1994. The respondents in 1998 to the new question do appear more risk tolerant with only 56.9% in category 1–2 compared to 64.6% in 1992 and 61.5% in 1994. This difference disappears in the last two survey waves. Nonetheless, variation in the question wording allows us to estimate the status quo bias and question-specific response errors.

This approach to measuring risk preference from hypothetical gambles in the HRS differs fundamentally from earlier survey measurement of attitudes toward risk. Other surveys commonly use categorical responses with vague quantifiers to probe risk preferences. For example, beginning in 1983, the Survey of Consumer Finances (SCF) asks respondents:

Which of the statements comes closest to the amount of financial risk that you and your (spouse/partner) are willing to take when you save or make investments?

1. take substantial financial risks expecting to earn substantial returns
2. take above average financial risks expecting to earn above average returns
3. take average financial risks expecting to earn average returns
4. not willing to take any financial risks

While intended to order respondents by their risk tolerance, the subjective wording may generate uninterpretable variation. Because individuals must define “substantial,” “above average,” and “average” financial risks and returns, we cannot quantify differences across responses. In contrast, the income gambles on the HRS supply objective boundaries between risk categories. In the next section, we use economic theory to map survey responses to a cardinal proxy for risk tolerance.

Using the cardinal proxy has several advantages. First, it provides a unidimensional, quantitative measure of risk tolerance that allows meaningful interpersonal comparisons. Second, in many settings, such as the demand for risky assets that we study in Section 6, economic theory makes predictions that link risk preference parameters quantitatively to economic decisions. Third, by having a quantitative measure we can correct for the measurement error inevitable with proxies based on survey responses.

3. CONSTRUCTING A CARDINAL PROXY

Expected utility theory provides a cardinal metric for risk preference—the coefficient of relative risk tolerance. Denote an individual’s concave utility function over original lifetime income as $U(W)$. Faced with 50–50 gambles of doubling lifetime income or cutting it by various fractions π , an individual should accept the risky job when its expected utility exceeds the utility from the certain job—that is, if

$$.5U(2W) + .5U((1 - \pi)W) \geq U(W). \tag{1}$$

The greater the curvature of U , the smaller the downside risk π an individual will accept. Associating gamble responses more tightly with underlying risk tolerance requires a parametric utility function.

We assume that constant relative risk aversion (CRRA) well approximates individuals’ utility over lifetime income

$$U(W) = \frac{W^{1-1/\theta}}{1-1/\theta}, \tag{2}$$

where the coefficient of relative risk tolerance θ may differ across individuals. This form implies that relative risk tolerance, $\theta = -U'/WU''$ (Pratt 1964), is constant across all values of lifetime income for a given individual. Analysis of the gamble responses with household income and wealth supports this utility specification (Sahm 2007). We focus on relative risk tolerance θ rather than relative risk aversion $1/\theta$ because relative risk tolerance is linearly related to demand for risky financial assets (Breen 1979). While the survey does not directly measure risk tolerance, the responses to the income gambles with this utility function establish boundaries on the underlying preference parameter.

To illustrate how to bound risk tolerance, consider individuals in response category 3. By accepting the risky job when the downside risk is one-fifth, but declining when the downside risk is one-third, these individuals reveal risk tolerance between .27 and .50. Each bound for this category equates the expected utility of a new risky job and the current certain job:

$$.5 \frac{2^{1-1/\underline{\theta}_3}}{1-1/\underline{\theta}_3} + .5 \frac{(4/5)^{1-1/\underline{\theta}_3}}{1-1/\underline{\theta}_3} = \frac{1^{1-1/\underline{\theta}_3}}{1-1/\underline{\theta}_3} \rightarrow \underline{\theta}_3 = .27, \tag{3}$$

$$.5 \frac{2^{1-1/\bar{\theta}_3}}{1-1/\bar{\theta}_3} + .5 \frac{(2/3)^{1-1/\bar{\theta}_3}}{1-1/\bar{\theta}_3} = \frac{1^{1-1/\bar{\theta}_3}}{1-1/\bar{\theta}_3} \rightarrow \bar{\theta}_3 = .50. \tag{4}$$

Substituting the largest downside risk accepted and the smallest risk rejected from Table 1, we similarly determine the lower and upper bounds for the other categories. The last two columns

of Table 1 report the bounds for each response category. The categories exhaust the possible range of risk tolerance.

In the next section, we consider a more general model that accounts for measurement error and other features of the question. To illustrate how we map the discrete responses into a continuous distribution, assume that true risk tolerance follows a log-normal distribution,

$$\log \theta \equiv x \sim N(\mu, \sigma_x^2). \tag{5}$$

The log-normal functional form has several advantages. First, it imposes the restriction that relative risk tolerance is nonnegative. Second, it is parsimonious and computationally simple. Third, we are able to use the moment-generating function of the normal to calculate analytically the unconditional and conditional expectations of $\theta = \exp(x)$. Finally, the log-normal appears to fit the data well. It can capture the fact that the modal value of relative risk tolerance is close to 0 but that a substantial fraction of individuals have higher risk tolerance.

We use standard maximum likelihood methods to estimate the mean μ and variance σ_x^2 of log risk tolerance in the population. Consider first a case in which we observe one response category c for each individual. The probability of being in category j is

$$P(c = j) = P(\log \underline{\theta}_j < x < \log \bar{\theta}_j) = \Phi((\log \bar{\theta}_j - \mu)/\sigma_x) - \Phi((\log \underline{\theta}_j - \mu)/\sigma_x), \tag{6}$$

where $\Phi(\cdot)$ is the cumulative normal distribution function. Maximizing the sample log-likelihood of the individuals’ first gamble response yields a mean log risk tolerance of -1.98 and a standard deviation of 1.76 as reported in the second column of Table 3. These parameters are precisely estimated: Both have an asymptotic standard error of .03. For the maximum likelihood estimation, we use the modified method of scoring where the sample average of the outer product of the score function approximates the information matrix.

For many applications, it is valuable to assign a numerical risk tolerance proxy for each individual conditional on his or her survey responses. Using the estimated population parameters, we can impute log risk tolerance conditional on a survey response in category j as

$$E(\log \theta | c = j) = \mu + \sigma_x \frac{\phi((\log \underline{\theta}_j - \mu)/\sigma_x) - \phi((\log \bar{\theta}_j - \mu)/\sigma_x)}{\Phi((\log \bar{\theta}_j - \mu)/\sigma_x) - \Phi((\log \underline{\theta}_j - \mu)/\sigma_x)}, \tag{7}$$

where $\phi(\cdot)$ is the standard normal density function. Alternatively, from the moment-generating function we can impute risk tolerance as

$$E(\theta | c = j) = \exp\left(\mu + \frac{\sigma_x^2}{2}\right) \left(\Phi((\log \bar{\theta}_j - \mu - \sigma_x^2)/\sigma_x) - \Phi((\log \underline{\theta}_j - \mu - \sigma_x^2)/\sigma_x) \right) / \left(\Phi((\log \bar{\theta}_j - \mu)/\sigma_x) - \Phi((\log \underline{\theta}_j - \mu)/\sigma_x) \right). \tag{8}$$

Given the parameter estimates, the proxy, $h = E(\theta | \mathbf{c})$, has four values, .083, .367, .706, and 3.687, for individuals in response categories 1–2, 3, 4, and 5–6. Unlike ordinal rankings, this proxy quantifies the average difference in log risk tolerance across the risk categories.

Table 3. Distribution of log risk tolerance: maximum likelihood estimates

Parameter	Ignoring response error	Modeling response error	Including application covariates
Log risk tolerance			
Mean μ	-1.98 (.03)	-1.84 (.03)	-1.86 (.07)
Standard deviation σ_x	1.76 (.03)	.73 (.04)	.73 (.04)
Status quo bias b_o		-.11 (.04)	-.10 (.07)
Response error standard deviation			
Original question, transitory σ_{eo}		1.39 (.05)	1.40 (.05)
Original question, persistent $\sigma_{\kappa o}$.73 (.10)	.72 (.10)
SQB-free question, transitory σ_{ef}		1.43 (.03)	1.42 (.03)
SQB-free question, persistent $\sigma_{\kappa f}$.60 (.09)	.61 (.09)
Number of individuals	11,616	11,616	11,616
Number of responses	11,616	17,580	17,580
Number of parameters	2	7	19
Log-likelihood	-12,073.4	-21,208.3	-21,121.3

NOTE: The second column estimates the model in Section 3. The third column models survey response error, as described in Section 4. The model of log risk tolerance in the fourth column includes the covariates from the application in Section 6. Asymptotic standard errors are given in parentheses.

4. ADDRESSING SURVEY RESPONSE ERROR

Responses to hypothetical income gambles likely provide a noisy signal of true risk tolerance. Thus, the risk tolerance proxy from the previous section is also error prone. Statistical procedures that use the risk tolerance proxy will be subject to errors-in-variables problems. In particular, using the proxy as an explanatory variable in a regression will lead to attenuation biases and inconsistent coefficient estimates. Because a key aim of including the risk questions on large-scale surveys is to provide researchers with a means to control for heterogeneity in preferences, it is critical to address and correct for the consequences of survey response error.

That some individuals give multiple responses to the risk tolerance questions provides a lever for quantifying survey response error. By making the structural assumption that preferences are immutable, we attribute the common component in an individual’s answers to true preference and the changes to response error. Recall that $x = \log(\theta)$ is the individual’s true preference parameter. With two versions of the gamble question, we also incorporate a question-specific persistent response error. The survey response error in wave w to question type q is a normal disturbance ϵ_{qw} added to x that leads the individual to choose the gamble response category corresponding to the sum ξ_{qw} . The error ϵ_{qw} can be interpreted as either an individual’s misperception of his or her risk tolerance or an error the individual makes in calculating the bounds $(\underline{\theta}_j, \bar{\theta}_j)$ that map preferences into the gambles. Hence,

$$\xi_{qw} = x + \epsilon_{qw} = x + b_q + \kappa_q + e_{qw}, \tag{9}$$

where b_q is a common bias across individuals of question type q , κ_q is the individual’s persistent response error for question type q , and e_{qw} is the individual’s transitory response error for a particular wave w and question type q . The components are distributed as $\xi_{qw} \sim N(\mu + b_q, \sigma_q^2)$, $\kappa_q \sim N(0, \sigma_{\kappa q}^2)$, and $e_{qw} \sim N(0, \sigma_{eq}^2)$ with $\sigma_q^2 = \sigma_x^2 + \sigma_{\kappa q}^2 + \sigma_{eq}^2$. The covariance in responses across waves for different question types depends only on the variance of true risk tolerance. For the same question type, the variance of the persistent error also affects the covariance across waves. We assume that the survey response error is a purely random—or “classical”—measurement error. Specifically, the response error ϵ_{qw} is independent of an individual’s true risk tolerance and any other attributes.

We analyze the two question types, the original question o and the status-quo-bias-free question f , so $q \in \{o, f\}$. In each wave, only one question type is asked. We assume that the new version is not subject to status quo bias on average, so $b_f = 0$ and $\xi_{fw} \sim N(\mu, \sigma_f^2)$. Identification of the parameters requires that at least some individuals answer the gambles more than once and some of the multiple responders answer the same question type more than once. Of the 11,616 individuals in our sample, all answer the original question at least once, and 4,244 individuals answer a status-quo-bias-free question. There are 693 individuals who answer the original question twice. For the bias-free question, 471 individuals answer in two surveys and 278 in three surveys.

In Section 3 we discuss how an individual with true log risk tolerance x will be assigned to a category by responses to the survey questions. Survey response error can move the individual into a different category from wave to wave and affects assignment to response categories even for those who answer only in one wave. For individuals who respond in only one wave, the likelihood of category j is

$$P(c_w = j) = \Phi\left(\frac{\log \bar{\theta}_j - \mu - b_q}{\sigma_q}\right) - \Phi\left(\frac{\log \underline{\theta}_j - \mu - b_q}{\sigma_q}\right). \tag{10}$$

This likelihood depends on the variance of error-prone risk tolerance, σ_q^2 , not that of true risk tolerance, σ_x^2 . Obviously, if all individuals answered in only one wave to one question type, the problem is underidentified.

For those answering the income gambles in both waves, the probability of observing category j in wave w and category k in wave w' is

$$P(c_w = j, c_{w'} = k) = \vec{\Phi}(\bar{N}_{jq}, \bar{N}_{kq'}, \rho) + \vec{\Phi}(\underline{N}_{jq}, \underline{N}_{kq'}, \rho) - \vec{\Phi}(\bar{N}_{jq}, \underline{N}_{kq'}, \rho) - \vec{\Phi}(\underline{N}_{jq}, \bar{N}_{kq'}, \rho), \tag{11}$$

where $\vec{\Phi}(\cdot)$ is the bivariate normal cumulative distribution function, $\bar{N}_{jq} = (\log \bar{\theta}_j - \mu - b_q)/\sigma_q$, $\bar{N}_{kq'} = (\log \bar{\theta}_k - \mu - b_{q'})/\sigma_{q'}$, $\underline{N}_{jq} = (\log \underline{\theta}_j - \mu - b_q)/\sigma_q$, and $\underline{N}_{kq'} = (\log \underline{\theta}_k - \mu - b_{q'})/\sigma_{q'}$. When the question type is the same, that is, $q = q'$, the correlation, ρ , between the variables ξ_{qw} and $\xi_{qw'}$ is the fraction $(\sigma_x^2 + \sigma_{\kappa q}^2)/\sigma_q^2$ of the total variance of the error-prone variable due to true log relative risk tolerance plus the question-specific persistent response error. When the question types differ, that is, $q \neq q'$, the correlation is $\sigma_x^2/\sigma_q\sigma_{q'}$ where

the covariance depends only on the variation in true log relative risk tolerance. Unlike the typical multiple-indicator solution to the errors-in-variables problem, identification here does not require repeat observations from all individuals in the sample.

Maximizing the sample log-likelihood with respect to μ , σ_x , $\sigma_{\kappa o}$, $\sigma_{\kappa f}$, $\sigma_{e o}$, and $\sigma_{e f}$ yields consistent estimates of the parameters. The third column of Table 3 reports the estimates. The estimated mean of log risk tolerance -1.84 is somewhat higher in this model with multiple gamble responses and question-specific response errors. The original question type is associated with an 11% lower reported risk tolerance. While this status quo bias is relatively modest, it is statistically different from 0. A more substantial shift occurs in the estimated variation of true log risk tolerance, as the estimate of the standard deviation falls to .73 from 1.76. Most of this decline is from modeling transitory response error using multiple gamble responses of some individuals. The modeling of question-specific persistent response error also lowers the estimated heterogeneity in true preferences somewhat. Together, this implies a much lower estimate of mean risk tolerance in the population: .21 instead of .65. The variability from response error greatly exceeds that from true risk tolerance. This finding highlights the limited test-retest reliability of the gambles and the need for multiple responses from some individuals. Nonetheless, the income gambles still convey much useful information on preferences as the application in Section 6 validates.

Ignoring survey response error overstates the heterogeneity in risk preferences. As noted, this causes an upward bias in estimated average risk tolerance. This effect is not dependent on the log-normal specification. Given the nonnegativity of risk tolerance, noise will, in general, shift the mean of the distribution of $\exp(\xi)$ to the right. Figure 1 illustrates the effects of response error. The solid line is the empirical distribution of the discrete responses in 1992 from Table 2 using the bounds $(\underline{\theta}_j, \bar{\theta}_j)$ in Table 1. The solid curve is the fitted log-normal distribution of true risk tolerance $\theta = \exp(x)$. The dashed curve is the fitted log-normal distribution of the true parameter plus noise, $\exp(x + \epsilon)$. The figure shows how the distribution of the true parameter moves mass away from the extremes relative to the distribution that includes noise from response errors.

Table 4 summarizes additional features of the estimated distribution of true risk preferences based on the parameter estimates in the third column of Table 3. The second column shows the distribution of log risk tolerance. The third column shows the distribution of the level of risk tolerance. The estimated mode of .094 indicates that the bulk of respondents have very low risk tolerance. Yet, there are enough respondents with relatively high risk tolerance to pull the mean substantially above the mode. About 25% of respondents are estimated to have risk tolerance greater than or equal to .259, and about 10% have risk tolerance greater than .402. Yet, virtually no respondents have risk tolerance as high as 1 (logarithmic utility).

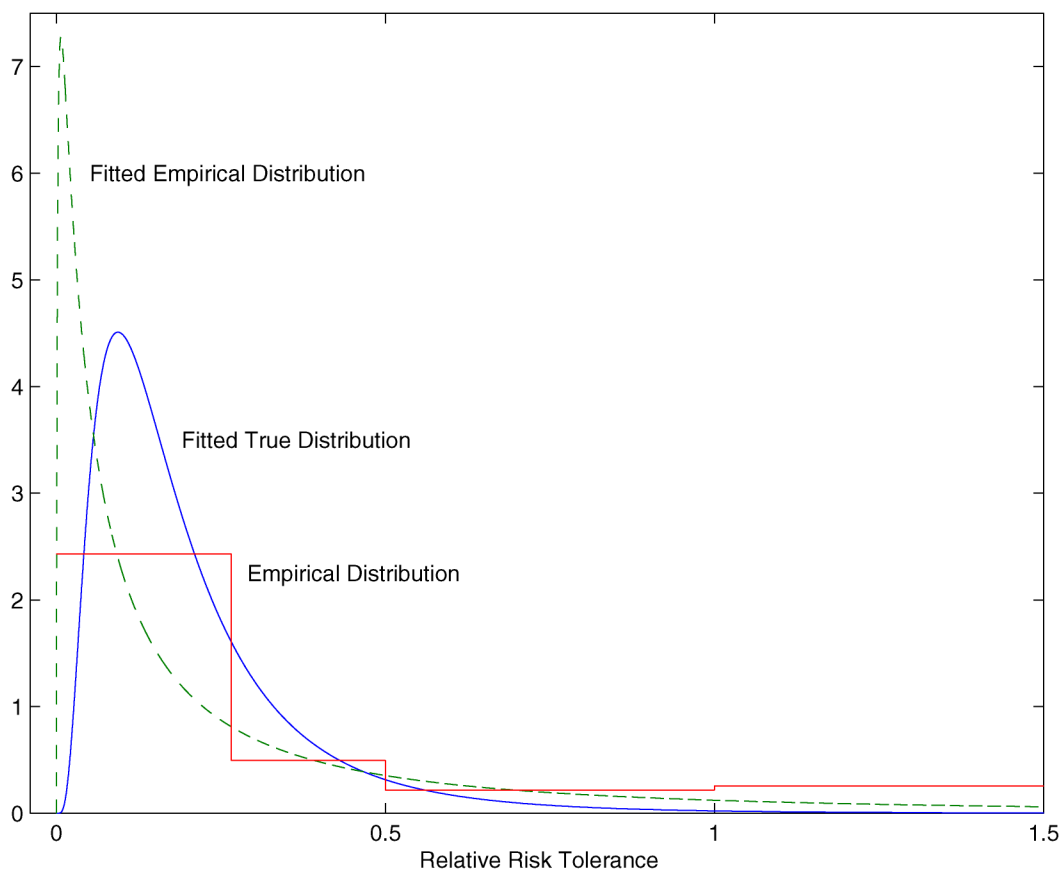


Figure 1. Distribution of relative risk tolerance. The solid line shows the empirical distribution of the survey responses. The solid curve shows the fitted distribution of the true level of risk tolerance, $\theta = \exp(x)$, using the model from Section 4. The dashed curved shows the fitted empirical distribution, $\exp(\xi) = \exp(x + \epsilon)$.

Table 4. Distribution of risk preferences

Parameter	Log risk tolerance	Risk tolerance	Risk aversion
Mean	-1.84 (.03)	.206 (.008)	8.2 (.3)
Median	-1.84 (.03)	.159 (.005)	6.3 (.2)
Mode	-1.84 (.03)	.094 (.004)	3.7 (.2)
Standard deviation	.73 (.04)	.172 (.018)	6.8 (.7)
Fractiles			
1	-1.54	.029	1.2
5	-1.32	.048	1.9
10	-1.20	.063	2.5
25	-1.01	.097	3.9
50	-.80	.159	6.3
75	-.59	.259	10.3
90	-.40	.402	16.0
95	-.28	.523	20.8
99	-.07	.858	34.1

NOTE: The values are calculated from the parameter estimates in the third column of Table 3. Asymptotic standard errors approximated with the delta method are given in parentheses.

For many applications—notably demand for risky assets—relative risk tolerance θ is the relevant preference parameter (Breedon 1979; Barsky et al. 1997). But in other applications, such as the strength of the precautionary saving motive, its reciprocal $1/\theta$, relative risk aversion, would be the parameter of interest (Carroll and Kimball, in press). When preferences are heterogeneous across individuals, the reciprocal of average relative risk tolerance is not equal to the average of its reciprocal. The last column of Table 4 gives the parameters and fractiles of the distribution of relative risk aversion. For our parameter estimates, average relative risk tolerance is .206. The estimated average of relative risk aversion is 8.2, which is far greater than $1/.206 = 4.9$. This difference between the expectation of the reciprocal and the reciprocal of the expectation is a powerful example of Jensen’s inequality. Jensen’s inequality gets its bite in this application from the substantial heterogeneity in preferences, the convexity of the $1/\theta$ function, and the concentrated mass of the probability density near 0, where the function $1/\theta$ is most curved.

Many researchers will want to impute risk tolerance for individuals. As our proxy for individual risk preference, we calculate the expected risk tolerance, conditional on an individual’s responses, using the estimated distributional parameters of our statistical model. The formula is similar to (8) in Section 3 except that it now accounts for question-specific response error and multiple responses to the gamble questions. Table 5 reports the proxy values of risk tolerance, as well as of log risk tolerance and risk aversion, for respondents to one status-quo-bias-free question. The proxy of risk tolerance for response category 1 (reject job with one-tenth downside risk) is .153. The range of relative risk tolerance corresponding to those preferences is from 0 to .13. (See Table 1.) Hence, the proxy value for this response lies slightly higher than the range. For risk category 2, the proxy of .203 lies near the center of the range from .13 to .27. With the more risk tolerant response categories, the

Table 5. Imputation of risk preference

Response category	Log risk tolerance	Risk tolerance	Risk aversion
1	-2.107	.153	10.4
2	-1.811	.203	7.6
3	-1.693	.228	6.7
4	-1.575	.257	6.0
5	-1.419	.301	5.1
6	-1.172	.387	4.0

NOTE: The proxy values are for responses to a single SQB-free question and are based on the estimates in the third column of Table 3. The values differ for persons answering in multiple surveys, the original question type, or in the combined categories 1–2 and 5–6. We provide a spreadsheet of all possible values online (http://www.amstat.org/publications/jasa/supplemental_materials).

proxy values are substantially lower than the range. For example, category 5 (accept a job with one-third downside risk but reject a job with one-half downside risk), the proxy of .301 lies far below the low end of the range from 1.0 to 3.7. The proxy values of log risk tolerance and risk aversion follow a similar pattern, as do the proxies from a response to the original question type. Hence, correcting for response error shifts the proxy toward the unconditional mean. Yet, substantial heterogeneity and meaningful quantitative differences remain even after this correction.

For those answering in multiple waves, we use all their responses to sharpen the estimate of their relative risk tolerance. These additional responses greatly widen the range of proxy values. The lowest imputed value of risk tolerance in our sample is .087 and the highest value is .732. When individuals give different responses across waves, we adjust the proxy values accordingly. Table 5 contains only a small subset of the 370 unique proxy values observed in this sample. For researchers who wish to make imputations based on our parameter values for any possible response to the HRS questions, we provide a spreadsheet of all possible values of risk tolerance and risk aversion online (http://www.amstat.org/publications/jasa/supplemental_materials).

5. STUDYING BEHAVIOR WITH THE PROXY

A major application of our proxy for risk tolerance is its use as a regressor to control for heterogeneity in preferences when studying a wide range of behaviors. The proxy $h = E(\theta|\mathbf{c})$ is the conditional expectation of true risk tolerance. Hence, the deviation of the proxy from the true variable $u = \theta - h$ is not a classical measurement error. In particular, the deviation is correlated with the true variable, not the proxy. In this section, we discuss the nonstandard errors-in-variables problem that arises from use of the proxy and present an estimator that addresses this problem.

To study the effects of risk tolerance and other regressors on behavior, consider a model

$$y = \theta\delta_\theta + \mathbf{z}\delta_{\mathbf{z}} + \nu, \tag{12}$$

where θ is true risk tolerance and \mathbf{z} is a $1 \times K$ vector of observables that also affect the behavior of interest y . To simplify later analysis, all variables are expressed as deviations from their means. We make the assumptions that, conditional on the regressors, the population error is mean 0, $E(\nu|\theta, \mathbf{z}) = 0$, and that

the expected outer product matrix of (θ, \mathbf{z}) has full rank. If we observed true risk tolerance and the other regressors, ordinary least squares (OLS) would consistently estimate the population parameters, $\delta = (\delta_\theta, \delta'_z)'$.

Now consider the use of the proxy. Substituting the proxy $h = E(\theta|\mathbf{c})$ in (12), we have

$$y = h\delta_\theta + \mathbf{z}\delta_z + \eta, \tag{13}$$

where

$$\eta = u\delta_\theta + v. \tag{14}$$

The composite error term η includes an expectation error $u = \theta - h$ and the structural error term v . Unlike a classical measurement error, the deviation u of the proxy from the true variable is uncorrelated with the proxy h and correlated with the true variable θ . This implies that in a univariate linear regression of a dependent variable y on only the proxy h , there is no attenuation bias and the OLS coefficient is consistent.

In a multivariate setting, the OLS estimator using the proxy is unlikely to provide consistent estimates of the population parameters, $\delta = (\delta_\theta, \delta'_z)'$. The proxy of risk tolerance $h = E(\theta|\mathbf{c})$ only conditions on an individual's gamble response categories, so regressors \mathbf{z} that are correlated with true risk tolerance θ would also correlate with the expectation error u . For example, men may be more risk tolerant than women. Then gender would be correlated with the expectation error. Using the proxy with a standard set of demographic regressors, the OLS coefficient estimate for men would mix the direct effects of gender with the indirect effects of risk tolerance. A more general statement of the problem is that

$$E(\mathbf{z}'h) \neq E(\mathbf{z}'\theta). \tag{15}$$

The lack of equality in (15) arises because of the correlation between the proxy's expectation error u and the regressors \mathbf{z} , which also implies that OLS is inconsistent.

We have enough structure on the problem to derive moment conditions that will allow for a consistent estimator using the proxy. The assumption of purely random response error and the properties of conditional expectations imply that the proxy is uncorrelated with both the structural error term v and the expectation error u . This yields the following moment condition for the proxy:

$$E(h\eta) = E(hu)\delta_\theta + E(hv) = 0. \tag{16}$$

To formulate a moment condition for the other observables, we assume that the conditional expectation of each observable z_k in the vector \mathbf{z} is linear in risk tolerance, such that

$$z_k = \theta\beta_k + \zeta, \tag{17}$$

where $E(\zeta|\theta) = 0$ and $\beta_k = E(\theta^2)^{-1}E(\theta z_k)$. The linear specification serves as a good approximation and could be extended to a risk tolerance vector that included higher order terms. With purely random response error, ζ is independent of the response error ϵ , which together with θ determines the proxy h . This implies that $E(h\zeta) = 0$. By definition, the proxy h is also uncorrelated with the expectation error $u = \theta - h$. Substituting the proxy in (17), we have

$$z_k = h\beta_k + u\beta_k + \zeta, \tag{18}$$

so the regression of z_k on the proxy h consistently estimates β_k , that is, $\beta_k = E(\theta^2)^{-1}E(\theta z_k) = E(h^2)^{-1}E(h z_k)$. We define the true-to-proxy variance ratio as

$$\lambda = E(\theta^2)/E(h^2). \tag{19}$$

It follows that

$$\begin{aligned} E(\theta z_k) &= \beta_k E(\theta^2) = \beta_k \lambda E(h^2) \\ &= \lambda E(h z_k) \quad \text{for all } z_k \in \mathbf{z}, \end{aligned} \tag{20}$$

where the first equality uses the population estimate of β_k in terms of θ , the second equality uses the definition of λ , and the third equality uses the population estimate of β_k in terms of h . We restate the model in (12) with the proxy h adjusted by λ as

$$y = \lambda h\delta_\theta + \mathbf{z}\delta_z + \omega, \tag{21}$$

where

$$\omega = (\theta - \lambda h)\delta_\theta + v. \tag{22}$$

With (16) and (20), we have two sets of orthogonality conditions that identify the model:

$$E(h\eta) = E[h(y - h\delta_\theta - \mathbf{z}\delta_z)] = 0, \tag{23}$$

$$E(\mathbf{z}'\omega) = E[\mathbf{z}'(y - \lambda h\delta_\theta - \mathbf{z}\delta_z)] = 0. \tag{24}$$

The second orthogonality condition effectively multiplies the covariance of \mathbf{z} with the proxy h by the variance ratio λ to get the implied covariance of \mathbf{z} with θ .

The estimator of δ will be based on the sample estimates of the proxy h and the true-to-proxy variance ratio λ . We can implement this generalized method of moments (GMM) estimator because we have an estimate of λ from the maximum likelihood estimation. This situation contrasts with the standard errors-in-variables setting where the true-to-proxy variance ratio is unidentified. Substituting the sample analogs into the moment conditions and solving for the estimates gives

$$\begin{aligned} \begin{bmatrix} \widehat{\delta}_\theta \\ \widehat{\delta}_z \end{bmatrix} &= \begin{bmatrix} N^{-1} \sum h_i^2 & N^{-1} \sum h_i z_i \\ N^{-1} \sum \lambda z_i' h_i & N^{-1} \sum z_i' z_i \end{bmatrix}^{-1} \\ &\quad \times \begin{bmatrix} N^{-1} \sum h_i y_i \\ N^{-1} \sum z_i' y_i \end{bmatrix}. \end{aligned} \tag{25}$$

Under the conditions specified, these will be consistent estimates of δ and have a limiting normal distribution. Note the ratio λ in the lower left block of the inverted matrix. There are three cases in which this estimator is identical to the OLS estimator: (i) when there are no regressors other than risk tolerance; (ii) when none of the other regressors is correlated with true risk tolerance; and (iii) when there is no expectation error for the proxy, that is, $\theta = h$, so $\lambda = 1$. Taking into account that the proxy variance is attenuated with respect to the true preference parameter is important in multivariate models with strong correlations between the other regressors and risk tolerance.

The asymptotic distribution of the estimator in (25) is

$$\sqrt{N}(\widehat{\delta} - \delta) \xrightarrow{D} N(0, \mathbf{A}^{-1}\mathbf{B}\mathbf{A}^{-1}), \tag{26}$$

where

$$\begin{aligned} \mathbf{A} &= E \begin{bmatrix} h^2 & h\mathbf{z} \\ \lambda\mathbf{z}'h & \mathbf{z}'\mathbf{z} \end{bmatrix}, \\ \mathbf{B} &= E \left[\begin{pmatrix} h\eta \\ \mathbf{z}'\omega \end{pmatrix} \begin{pmatrix} h\eta & \mathbf{z}'\omega \end{pmatrix} \right]. \end{aligned} \tag{27}$$

Table 6. True-to-proxy variance ratio λ

Parameter	Estimate
Variance	
Risk tolerance θ	.030
Proxy $h = E(\theta \mathbf{c})$.005
True-to-proxy ratio λ	6.319

NOTE: The estimated variance of true risk tolerance and its proxy depend on the estimated parameters in the third column of Table 3. Section 4 describes the relationship between survey responses and the proxy values. The true-to-proxy variance ratio λ is an input to the GMM estimator in (25) and the R^2 in (28).

While we do not directly observe risk tolerance, we can still compute an implied R^2 for the model in (12) based on the true values of risk tolerance. The R^2 as if true θ were observed is

$$R^2 = \hat{\delta}' \begin{bmatrix} N^{-1} \sum \widehat{\lambda} \widehat{h}_i^2 & N^{-1} \sum \widehat{\lambda} \widehat{h}_i \mathbf{z}_i \\ N^{-1} \sum \widehat{\lambda} \mathbf{z}_i' \widehat{h}_i & N^{-1} \sum \mathbf{z}_i' \mathbf{z}_i \end{bmatrix} \hat{\delta} / (N^{-1} \sum y_i^2). \tag{28}$$

Using the standard R^2 from a regression with the proxy would understate the explanatory power of the model, because the variability of the proxy understates the true variability of risk tolerance. Table 6 shows that this understatement is substantial. The ratio λ of the variance of the true risk tolerance to the proxy is 6.32. When the other regressors are strongly correlated with risk tolerance, the GMM estimator in (25) and the implied R^2 in (28) will more accurately characterize the effects of risk tolerance on behavior than standard estimators. Even in a univariate regression on risk tolerance alone, (28) is needed to calculate the implied R^2 .

6. APPLICATION TO ASSET ALLOCATION

In this section, we apply the methods discussed previously to study asset allocation. Faced with uncertain asset returns, risk preferences should be central in allocating financial wealth between high-risk and low-risk assets. Individuals with greater risk tolerance should be willing to hold a larger fraction of their wealth in risky assets, such as stocks. Under complete markets, only risk tolerance and the distribution of risky asset returns affect allocations (Merton 1969; Samuelson 1969). Many individuals also anticipate labor income, which they cannot fully capitalize due to their ability to sort across contracts by their risk type (adverse selection) and to alter their post-contract behavior (moral hazard). With market incompleteness, models of asset allocation also identify a role for the determinants of future labor income, such as age and the distribution of income shocks (Heaton and Lucas 1997). Empirical studies often document substantial differences in asset allocation by gender, education, and race. Nonetheless, much of the heterogeneity in asset allocation remains unexplained.

In contrast with other empirical studies of asset allocation, our risk tolerance proxy allows us to control quantitatively for the effects of risk preference cross-sectionally. In this section, using data from the HRS, we present estimates of how the share of financial wealth held in stocks increases with risk tolerance. We also consider other regressors such as gender, education, age, race, household income, and wealth. While households typically own assets jointly, many of these attributes are person specific. We treat the respondent who is most knowledgeable

about household finances as the primary decision maker and control for his or her attributes. We limit our analysis to households with positive financial wealth and income. Because the HRS is a sample of older households that have often accumulated some wealth, this selection eliminates fewer observations than it would in an age-representative sample. Nonetheless, it does exclude approximately 20% of households. The average share of financial wealth in stocks (excluding individual retirement accounts) is .16, and a significant portion of households do not own stocks. The standard deviation of the share in stocks is .29, so there is considerable dispersion in stock allocations.

To demonstrate the usefulness of our risk tolerance proxy h and the true-to-proxy variance ratio λ , we contrast our GMM estimates with the OLS estimates that use the risk tolerance variable without taking into account response error. While focusing on the effects of risk tolerance, we also discuss the effects of gender and education. We use the estimated effects from these regressors to demonstrate the misleading inferences from failing to take into account risk tolerance heterogeneity and also failing to correct for the consequences of survey response error in the risk tolerance proxy.

As a baseline, we estimate the stock allocation model without any control for risk tolerance. This corresponds to the approach in most empirical studies. As reported in Table 7, the gender, education, and race of the financial respondent as well as the household's log income and log financial assets account for 17.0% of the variation in stock allocations. In this specification, households with men responsible for the finances have 2.4 percentage points more in stocks on average. Post-college education raises the share by 3.4 percentage points. Both are statistically significant and represent 15% and 22% of the average stock allocation.

If any of these characteristics correlate with risk tolerance, then their estimated coefficients also include the indirect effects of risk tolerance. One way to try to sort out the direct effects of risk tolerance on stock holding and to study the confounding effect of gender, education, and other regressors is to estimate a model of asset allocation controlling for the categorical survey responses to the income gambles. Based on their first response in the HRS, we assign individuals to four risk tolerance categories. This regression with categorical controls explains 17.2% of the variation in stock allocations. Households in the most risk tolerant category hold 2.5 percentage points more of their wealth in stocks than those in the least risk tolerant category. But the relationship is nonlinear as households in the second lowest risk tolerant category hold 2.6 percentage points less in stocks than those in the least risk tolerant category. Adding the categorical controls diminishes the effect of a male financial respondent to 2.3 percentage points and post-college education to 3.2. These results are consistent with the Barsky et al. (1997) finding that men and the most educated are more risk tolerant. Even partially controlling for risk preferences begins to lower the estimated effect of these attributes on asset allocation.

The last four columns of estimates in Table 7 use different versions of the cardinal proxy for risk tolerance and different estimators. In the fourth column, we use the proxy from Section 3, which ignores survey response error. All else equal, the most risk tolerant households based on one observation (risk tolerance of 3.687) average 2.9 percentage points more in stocks than the least risk tolerant households (risk tolerance of .049).

Table 7. Effect of risk preferences on the share of financial wealth in stocks

Control for log risk tolerance	None	Categorical survey response	Risk tolerance proxy			
			Ignoring response error	Modeling response error	Modeling response error	Including application covariates
Estimator	OLS	OLS	OLS	OLS	GMM	OLS
Category 3		-.026 (.010)				
Category 4		.022 (.012)				
Category 5-6		.025 (.011)				
Proxy			.008 (.003)	.146 (.054)	.162 (.060)	.152 (.056)
Male	.024 (.007)	.023 (.007)	.023 (.007)	.023 (.007)	.014 (.008)	.018 (.008)
Education						
> 16 years	.034 (.012)	.032 (.012)	.032 (.012)	.031 (.012)	.012 (.013)	.019 (.014)
13-16 years	.036 (.009)	.035 (.009)	.035 (.009)	.035 (.009)	.024 (.009)	.029 (.009)
< 12 years	-.023 (.009)	-.024 (.009)	-.023 (.009)	-.023 (.009)	-.026 (.009)	-.024 (.009)
Black	-.029 (.009)	-.029 (.009)	-.029 (.009)	-.028 (.009)	-.024 (.009)	-.027 (.009)
Hispanic	-.035 (.012)	-.036 (.012)	-.035 (.012)	-.035 (.012)	-.034 (.013)	-.038 (.012)
Age/10	-.002 (.008)	-.001 (.008)	-.001 (.008)	-.001 (.008)	.006 (.008)	.006 (.008)
Log income	.002 (.005)	.003 (.005)	.003 (.005)	.002 (.005)	.004 (.005)	.003 (.005)
Log wealth	.046 (.002)	.047 (.002)	.046 (.002)	.046 (.002)	.047 (.002)	.046 (.002)
R^2	.170	.172	.171	.171	.178	.177

NOTE: Regressions include 5,818 households with positive financial wealth and total income in 1992. Individual attributes are from the household's financial respondent. Share of wealth in stocks has a mean of .158 and a standard deviation of .286. Asymptotic standard errors are in parentheses. In the second to last column, the GMM estimates are based on the formula in (25), and the R^2 in the last two columns is based on the formula in (28). For the application subsample, the true-to-proxy variance ratio λ is 6.40. In the last column, the proxy is constructed from a model of log risk tolerance that conditions on the application covariates as well as the gamble responses.

The fifth column of Table 7 uses the proxy from Section 4, which accounts for the measurement error in the gamble responses but does not address the potential correlation between the proxy's expectation error and the other regressors discussed in Section 5. These results show that ignoring survey response error greatly understates the marginal effect of risk tolerance on stock allocations. When we use the proxy values from Section 4, the coefficient estimate for the proxy increases over tenfold. This increase shows how attenuation bias affects the estimates in the previous two columns that do not account for response error. Of course, this correction mainly scales up the coefficient estimate and does not affect the R^2 . The larger estimated effect of risk tolerance means that when risk tolerance is measured more precisely with multiple responses that the predicted differences in behavior can be substantial. The most risk tolerant households based on multiple observations (risk tolerance of .732) average 9.4 percentage points more wealth in stocks than the least risk tolerant households (risk tolerance of .087). This difference represents 60% of the average stock share. Thus, correcting for measurement error has a substan-

tial impact on the estimated responsiveness of behavior to risk tolerance.

The sixth column of Table 7 uses the same proxy for risk tolerance as in the fifth column but replaces the OLS estimator with the GMM estimator derived in Section 5. The GMM estimates show the importance of accounting for the correlation between the expectation error of the proxy and the other regressors. When using (28) for the implied R^2 , the explained variation in stock allocations rises to 17.8% from 17.1%. The point estimate for the effect of risk tolerance rises 11% to .162. The average difference in stock allocations of the most and least risk tolerant households with multiple responses increases over 1 percentage point to 10.5. The GMM estimator has a more pronounced effect on the coefficient estimates for other regressors. As stressed in Section 5, the main issue is that in the OLS estimate the other regressors will spuriously account for variation in the dependent variable to the extent that they are correlated with risk tolerance. Having a male financial respondent now raises stock allocations by only 1.4 percentage points and the effect of a post-college education falls to 1.2 percentage points.

These coefficient estimates are 42% and 65% lower than in the regression with no measure of risk tolerance and are no longer statistically different from 0 at the 5% level.

As a check on the accuracy of the GMM estimator, the last column of Table 7 looks at an alternative estimator. Instead of basing the proxy just on the gamble response categories, we also condition on the regressors in this application. Specifically, in the first-step maximum likelihood model, we model the mean of log risk tolerance μ as a linear function of the observables \mathbf{z} . The estimated unconditional mean and variance of log risk tolerance from this alternative first-step maximum likelihood model are reported in the last column of Table 3. The estimated distribution does not differ substantially from the model that conditions on only the gamble responses. This is a direct approach to eliminate the correlation between the proxy's expectation error and the observables. Condition (15) now holds with equality, and the OLS estimator with this new proxy consistently estimates the model with true risk tolerance. The last two columns of Table 7 are very similar. This finding implies that the GMM approach, which does require on re-estimating the proxy conditional on all the covariates in the regressions, works well.

The results in Table 7 demonstrate the importance of carefully controlling for the heterogeneity in preferences. Beyond using the proxy to control understand the effect of risk tolerance, we show how the effect of other regressors can be overstated if no correction is made for the fact that the proxy is imperfectly measured and the other regressors are correlated with preferences. For researchers who want to include an individual measure of risk tolerance in their studies of other behaviors, our maximum likelihood estimates provide a valid proxy. To the extent that this proxy's expectation error is correlated with other explanatory variables of interest, the OLS estimates can be misleading. This problem can be addressed with the GMM estimator that we derived in Section 5 or can be avoided by conditioning on the other variables in the first-step maximum likelihood. While the second alternative might be the best approach, the Health and Retirement Study is currently the only dataset with a sufficient panel to correct for the survey response error in the gamble responses. When the first-step maximum likelihood is not possible (e.g., because of having only one response per individual), the proxy values we provide that condition only on the gamble responses should be used with the GMM estimator to obtain consistent estimates.

7. CONCLUSION

We demonstrate the importance of carefully controlling for risk preferences when examining asset allocation. In particular, our procedures address many issues in using survey-based measures of risk tolerance—translation of categorical responses to a cardinal metric, survey response error, and expectation error for the proxy. Our methods for constructing the proxy and estimating the effects of risk tolerance on behavior have a wide range of potential applications. A growing number of surveys, including the Panel Study of Income Dynamics in the United States, the CentERpanel in the Netherlands, and the Social Protection Survey in Chile, have fielded lifetime income gambles like those in the HRS. Our statistical procedures for constructing the risk tolerance proxy can be applied with minimal adjustment to these other surveys.

In studies of stock market participation (Hong, Kubik, and Stein 2004) and intergenerational wealth correlations (Charles and Hurst 2003), researchers have used indicator variables from income gamble responses. This approach does not fully capture the effect of heterogeneous risk preferences. According to our empirical analysis, even if the direct effects of risk tolerance are not central to the study, such indicator variables are unlikely to adequately control for risk tolerance. In other words, these partial controls are not sufficient either in theory or in practice for consistent estimates of the direct effects of other variables of interest. With survey questions and statistical techniques motivated by economic theory, we expand the options for studying the effects of risk preferences on behavior. Using the quantitative proxy for risk tolerance, we find a strong effect of risk tolerance on stock holding. Moreover, after accounting for how errors in measured risk tolerance are correlated with other variables, the estimated effects of gender and education on asset allocation are substantially reduced.

APPENDIX: BOOTSTRAP

Both the OLS and the GMM estimates in Table 7 use the risk tolerance proxy h , which is a generated regressor from the first-step maximum likelihood procedure. The variance ratio λ is another generated regressor in the GMM estimator. While the coefficient estimates from these second-step estimators are asymptotically consistent, the estimated standard errors do not reflect the sampling variation in the proxy and the variance ratio. We use a bootstrap to show this sampling variation does not qualitatively alter our inferences in Section 6.

Using a Monte Carlo experiment, we draw 199 random samples from the data and repeat the two steps of estimation in Sections 4 and 6. Sampling with replacement, we maintain the distribution of respondents to the original and status-quo-bias-free questions. We use a symmetric t test to construct the 95% bootstrap confidence interval for the proxy coefficient estimate in the asset allocation model. The OLS estimate in the fifth column of Table 7 of .146 has a confidence interval of .042 to .249. The GMM estimate in the sixth column of Table 7 of .162 has a confidence interval of .042 to .283. The OLS estimate in the seventh column of Table 7 of .152 has a confidence interval of .044 to .260. In all three cases, the estimated effect of risk tolerance on asset allocation remains statistically significant at the 5% level. As expected, sampling variation in the generated regressors has little effect on the inference of the other controls. The moderate impact of the generated regressors reflects the precision of the first-step maximum likelihood procedure.

[Received November 2005. Revised July 2007.]

REFERENCES

- Barsky, R. B., Juster, F. T., Kimball, M. S., and Shapiro, M. D. (1997), "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study," *The Quarterly Journal of Economics*, 112, 537–579.
- Breeden, D. T. (1979), "An Intertemporal Asset Pricing Model With Stochastic Consumption and Investment Opportunities," *Journal of Financial Economics*, 7, 265–296.
- Carroll, C. D., and Kimball, M. S. (in press), "Precautionary Saving and Precautionary Wealth," in *The New Palgrave: A Dictionary of Economics*, eds. S. Durlauf and L. E. Blume, New York: Macmillan.
- Charles, K. K., and Hurst, E. (2003), "Intergenerational Wealth Correlations," *Journal of Political Economy*, 111, 1155–1182.
- Heaton, J., and Lucas, D. (1997), "Market Frictions, Savings Behavior, and Portfolio Choice," *Macroeconomic Dynamics*, 1, 76–101.
- Hong, H., Kubik, J. D., and Stein, J. C. (2004), "Social Interaction and Stock Market Participation," *Journal of Finance*, 59, 137–163.
- Merton, R. C. (1969), "Lifetime Portfolio Selection Under Uncertainty: The Continuous-Time Case," *Review of Economics and Statistics*, 51, 247–257.

- Pratt, J. W. (1964), "Risk Aversion in the Small and in the Large," *Econometrica*, 32, 122–136.
- Sahm, C. R. (2007), "Stability of Risk Preference," working paper, Finance and Economic Discussion Series 2007-66, Washington: Board of Governors of the Federal Reserve System.
- Samuelson, P. A. (1969), "Lifetime Portfolio Selection by Dynamic Stochastic Programming," *Review of Economics and Statistics*, 51, 239–246.
- Samuelson, W., and Zeckhauser, R. (1988), "Status Quo Bias in Decision Making," *Journal of Risk and Uncertainty*, 1, 7–59.