Risk Averters that Love Risk? Marginal Risk Aversion in Comparison to a Reference Gamble

David R. Just and Travis J. Lybbert

We propose an analytical distinction between standard risk aversion based on the valuation of a single gamble and marginal risk aversion based on the change in valuation between two gambles. We measure marginal risk aversion in two dimensions—mean and variance. Data from a field experiment is used to study marginal risk aversion. Our results suggest that individuals rely on a reference gamble when assessing marginal risk. Individual responses to marginal changes in mean and variance are nearly identical in direction and magnitude—suggesting that information on both standard and marginal risk aversion is needed to accurately model behavior.

Key words: certainty effect, marginal risk aversion, probability weighting, risk aversion.

Applied economists have developed a vibrant literature that seeks to quantify the impacts of risk on production decisions. This technical literature posits an expected utility objective function and relies on advanced econometric techniques to back risk preferences out of observed input decisions. The accuracy of the assumption that individuals maximize expected utility is often threatened by pervasive and arbitrary decision heuristics, which could therefore cause severe bias when estimating risk preferences. Consider, for example, an investor who receives and evaluates new information, then makes slight adjustments in her investment positions accordingly. Such an incremental approach with marginal adjustments could imply different behavior than if she reevaluated her entire portfolio from scratch in response to new information or new investment options. Although applied economists generally assume that individuals globally reevaluate their alternatives when estimating risk preferences, individuals often seem to favor incremental adjustments based on comparisons between options. In such contexts methods that explicitly assess changes in risk preferences on the margin may yield more robust estimates of behavioral parameters.

In this article, we distinguish between standard risk aversion and marginal risk aversion. The conceptual distinction between these measures of risk aversion is simple. Standard measures of risk aversion such as Arrow–Pratt coefficients are generally based on isolated, stand-alone gambles and indicate risk aversion whenever an individual's certainty equivalent for a gamble is below its expected value. Marginal risk aversion, on the other hand, involves a comparison between gambles and is displayed when the individual responds conservatively to *changes* in the gamble. Thus, an individual is risk averse on the margin if her valuation increases (decreases) when the gamble changes to become less (more) risky.

In an expected utility maximization world the distinction between standard and marginal measures of risk aversion may be trivial, but these differences quickly become important when behavior deviates from expected utility. To illustrate, consider one of the best-known

---

David R. Just and Travis J. Lybbert are associate professor in Applied Economics and Management, Cornell University and assistant professor in Agricultural and Resource Economics, University of California-Davis and member of the Giannini Foundation for Agricultural Economics.

Data collection for this research was funded by the Agricultural Biotechnology Support Program II (USAID-Cornell University). Data collection support was provided by Dr. Balasubramanian, Vice Chancellor Ramasamy and others at Tamil Nadu Agricultural University and by K. Vijayaraghavan at Sathguru Management Consultants. Brent and Usha Zehr, Ritesh Mishra, R.M. Pillai, Datta Shirale, and others at the Maharashtra Hybrid Seed Company (MAHYCO) provided key research and field support. G. Anoop, R. Mahendran, S. Bavan, A. Gayathri, K. Kalaichelai, V. Kavitha, S. Ashok Kumar, S. Ramesh, A. S. Shanmuganandhan, and P. Vetrivel Kalai did an excellent job collecting the data. Kent Messer, Cornell University, provided invaluable expertise and assistance to the design of the experiment. Seminar audiences at the 2007 AAEA annual meeting and the October 2007 Economic Sciences Association meeting, as well as three anonymous referees, made useful suggestions for this article. The authors own all remaining errors.

Amer. J. Agr. Econ. 91(3) (August 2009): 612–626

DOI: 10.1111/j.1467-8276.2009.01273.x

expected utility deviations: the tendency to overweight small probabilities and underweight large probabilities, which is responsible in part for the certainty effect and other cognitive biases. Such probability weighting drives a wedge between standard and marginal risk aversion. To see how, suppose an individual has a simple linear utility function, which implies risk neutrality in the expected utility framework but systematically underweights large probabilities. This individual will value a gamble with a large probability of gain below its expected value, suggesting risk aversion. The same individual, however, may be indifferent between this gamble and another one with the same expected value and similar probabilities but greater variability in potential outcomes, suggesting risk neutrality. Such a disparity between behavior taken on average and behavior on the margin could help to explain why empirical estimates of risk aversion are widely variant and inconsistent (see Just and Peterson 2003). In particular, slightly different approaches to the same problem may produce very different risk estimates if one approach considers behavioral responses in aggregate, and the other implies marginal behavioral adjustments.

We derive simple measures of standard and marginal risk aversion and apply these measures to data collected from an economic experiment conducted among Indian farmers in the state of Tamil Nadu. In this experiment farmers stated their willingness-to-pay (WTP) for several different payoff distributions. Our derived measures show a substantial discrepancy between standard and marginal risk attitudes, with many farmers simultaneously appearing to be risk averse on average and risk loving on the margin. The evidence points to an anchoring-and-adjustment process in which individuals anchor on their willingness to pay for a specific gamble, then make adjustments to this value when similar gambles are presented. Rather than using a reference point (e.g., Kahneman and Tversky 1979),[1] individuals appear to use a reference gamble. Importantly, it is impossible to reconcile this behavior with a probability weighted model, since the estimates of a value function would have to account for both local convexity and global concavity.

[1] Bell (1985) proposes a related model, whereby individuals anchor on the expected outcome, with fear of disappointment (a result below that expectation), resulting in risk behavior.

Literature Review

To the best of our knowledge, no one has yet distinguished between standard and marginal risk aversion or suggested that both of the aversions may be the result of separate cognitive processes. There are, however, a couple of good reasons in the existing risk literature for why this distinction may matter to applied economists. First, expected utility theory attributes risk aversion entirely to the curvature of the utility function. Empirical applications in agriculture have generally sought to determine the level of curvature that best describes the input decisions of farmers (see Just and Just forthcoming). The outcomes of such estimation rely heavily on the assumption of expected utility maximization. Anomalies such as the certainty effect—the discounting of uncertain outcomes more than the probability would imply—could therefore cause severe bias in risk aversion estimates, because the psychological effect can be interpreted as severe risk aversion. Alternatively, the certainty effect should be absent when comparing two nontrivial gambles, potentially leading to very different estimates of risk aversion. This is fundamentally a failure of the independence axiom, which supposes that preferences should be independent of the addition of identical lottery components.

A second reason why the standard versus marginal risk aversion distinction may matter is highlighted by Rabin (2000) and Just and Peterson (2003), who show that an absurd degree of utility function curvature is required in an expected utility framework to rationalize many individuals' responses to relatively small gambles. Responses to small changes in gambles may not elicit the same radical responses. Thus, distinguishing between risk aversion and marginal risk aversion may provide insight into this apparent anomaly.

In the balance of this review section, we highlight in greater detail four strands in the risk literature that are relevant to this distinction. The first relates to Arrow's (1971) hypothesis that absolute risk aversion decreases as wealth increases—identical to shifting the mean of a distribution holding all else constant. Second, we discuss the certainty effect that implies a distinction between standard and marginal risk aversion with the point of departure being certainty. Third, the probability distortion literature highlights a potential practical benefit of considering risk aversion on the margin. Since the certainty effect can be understood as a result of a systematic misperception of probabilities, these second and third strands are closely related. Finally, we discuss a pair of papers that hypothesize a delineation between diminishing marginal utility of wealth and risk preferences.

Friedman and Savage's (1948) explanation of risk taking behavior supposed that the utility of wealth function changes its degree of concavity (or convexity) dramatically as one moves from a situation with a low level of wealth to a situation in which she/he has large amounts of wealth. In particular, they suppose that utility of wealth functions must display risk-loving behavior (convexity) for intermediate amounts of wealth and risk-averse behavior (concavity) for large or small amounts of wealth. A utility function of this shape would explain why poor individuals may buy insurance and lottery tickets simultaneously, while wealthier individuals would reject lottery tickets. Later, Arrow (1971) addressed the relationship between wealth and risk behavior, taking the mathematical theory to its limits. Arrow hypothesized that utility of wealth functions would demonstrate decreasing absolute risk aversion (DARA) and increasing relative risk aversion as wealth increased. The intuition behind these relationships is simple. As their wealth increases, individuals should be more willing to take any particular risk and less willing to risk any percentage of their wealth.

These two hypotheses, while simple, have fueled many applied studies with mixed results, although they generally support DARA. These studies employ econometric estimation to examine both production decisions facing profit risk (e.g., Bar-Shira, Just, and Zilberman 1997; Chavas and Holt 1990, 1996; Lins, Gabriel, and Sonka 1981; Pope and Just 1991; Saha, Shumway, and Talpaz 1994) and more conventional individual decisions (e.g., Bellante and Saba 1986; Cohn et al. 1975; Landskroner 1977; Morin and Fernandez-Suarez 1983; Siegel and Hoban 1982). Thus, according to the extant theoretical and empirical literature, one would expect individuals to be risk loving on the margin as the expected value of a gamble increases. DARA relates to the third derivative of utility of wealth function and occurs if and only if $-\frac{u'''(w)}{u''(w)} > -\frac{u''(w)}{u'(w)}$. The right-hand side of this inequality represents absolute risk aversion—a measure of one's willingness to take on risk, while the left-hand side represents absolute prudence (Kimball 1990)—a measure of how sensitive choice is to risk.[2] Thus, absolute prudence measures how absolute risk aversion changes as wealth level changes and may rightly be thought of as a marginal risk-aversion concept.

Next, the certainty effect may shed light on how risk aversion changes as variance of a prospect changes. The certainty effect commonly occurs when individuals must choose between some certain outcome and at least one risky choice. Individuals behave as if the probability assigned to the sample space for the risky choice sums to less than one.[3] In this case, the individual penalizes all risky choices in a way that is inconsistent with the independence axiom. In particular, the certainty effect can be found when a choice between lotteries, one being a certain outcome and the other a lottery with greater expected value, is compared to the same lotteries compounded with another lottery that receives a majority of the probability. When choosing between the original lotteries, a majority of individuals will choose the certain outcome. When presented with the compound lottery, the differences in probabilities for the best outcome become disproportionately small in the minds of individuals choosing between the lotteries, leading them to choose the riskier outcome—a violation of the independence axiom.

The certainty effect is pervasive, especially in choices involving probabilities near 1 or 0 (Camerer 1995), and has been found by many independent studies (see e.g., MacCrimmon and Larsson 1979). This pervasive effect and its associated uncertainty penalty have important implications in the present context. If, as suggested by the certainty effect, the mere presence of uncertainty causes individuals to discount their valuation of a prospect, then their valuation of a gamble will depend on whether it is evaluated relative to a certain or uncertain alternative. In the former case—which is akin to evaluating the gamble in isolation—the value is discounted by the uncertainty penalty, but in the latter the marginal value of one gamble relative to another effectively differences

[2] Constant absolute risk aversion occurs if absolute prudence and absolute risk aversion are equal, while increasing absolute risk aversion occurs if absolute prudence is less than absolute risk aversion.

[3] This provides one explanation of the certainty effect. Other explanations allow the sample space a full weight (see Camerer 1995 for a complete summary).

away this penalty.[4] If individuals formulate risk preferences by comparing gambles using a reference-based approach, this distinction may be particularly important when examining how behavior changes when underlying risk changes.

The third strand in the risk literature that is relevant for our purposes is probability distortion—one of the primary explanations of the certainty effect. Preston and Barrata (1948) were the first to note that individuals misperceive probabilities in choosing between risky outcomes. They ran auctions for various simple gambles—gambles in which there was a fixed probability of winning an amount of money. With regularity they found that individuals would bid more than the expected value for low probability wins (below 0.25) and less than the expected value for higher probability wins. The effect was robust when the experiment was run on PhD statisticians and other academics highly familiar with probability measures. This effect was exaggerated when more individuals were involved in the auction. Edwards (1953) confirmed these results using individual choices instead of auctions.

This result—that small probabilities are overweighted and large probabilities are underweighted—is one of the most robust results in risk experiments and has formed the basis for several models (Hong 1983; Kahneman and Tversky 1979). Lattimore, Baker, and Witte (1992) find that probability distortions seem to be dependent on the absolute level of payoffs involved, with higher values yielding a higher point where perceived and actual probabilities are identical. At a very rough level Tversky and Kahneman (1992) find that the weighting function differs between gains and losses. Such probability distortion suggests a potential benefit of thinking about marginal risk aversion in empirical applications. Provided that individuals' probability weighting functions are relatively stable, measures of marginal risk aversion may effectively difference away probability distortions that otherwise hamper standard measures of risk aversion. Probability distortions are unlikely to change substantially over the risk margin, thus minimizing their impact on measures of marginal risk aversion. This is similar to the way in which the certainty effect disappears when considering only uncertain prospects. In the next section we develop this potential benefit in greater detail and derive some testable implications.

A fourth literature separates the notion of diminishing marginal utility of wealth from risk preference. Dyer and Sarin (1982) introduce a framework whereby risk preference is measured with respect to diminishing marginal utility under certainty, which they call strength of preference. Their measure supposes that the utility function changes shape when facing different gambles. If when facing a risky choice the individual's von Neuman–Morgenstern utility function is more concave than when the individual is facing certainty, the individual alters her/his preferences when facing a risky choice. This individual would be considered relatively risk averse. Alternatively, an individual whose von Neuman–Morgenstern utility function becomes more convex when facing a risky prospect than when facing the same outcomes with certainty would be called relatively risk prone. The framework we suggest is similar in that it compares risk behavior between two gambles—one with risk and one without. The framework we eventually explore compares risk aversion between two nontrivial gambles, with one designated as a reference gamble, although the reference gamble payoff is not certain.

Pennings and Garcia (2001) took a more comprehensive approach by comparing the results of several risk aversion measures—both those based in expected utility and those based on more general psychological constructs. They find that, while using four different and disparate measures of risk attitude, they can explain 60% of the variation in use of futures contracts among Dutch farmers, while each individual measure can only explain between 9% and 27%. Importantly, they find that the various measures of risk aversion are significantly positively correlated, suggesting some degree of overarching validity to the notion of risk aversion. Still, individual variability in risk attitude from decision to decision is evident: many display risk aversion overall but can make individual decisions that are risk loving when evaluated in isolation. We build on this apparent inconsistency in this article without distinction between overall and marginal risk aversion.

[4] Differencing away the uncertainty penalty in this way relies on one of two assumptions: (*a*) the uncertainty penalty is not proportional to the magnitude of the gamble, or (*b*) the comparison gambles are close enough in magnitude to make the proportionality of this penalty a moot question. Our empirical application relies on the second assumption.

The Effect of Probability Distortions on Measures of Risk Aversion

In this section, we develop a framework to evaluate the impacts of probability weighting on estimates of risk aversion resulting from production decisions. We use this framework to derive our research hypotheses regarding marginal and standard risk aversion measures. Consider the case where a farmer must allocate land between two crops (following Marra and Carlson 1990). The farmer's decision could be written as

$$(1)\qquad \max_{L_1 \le \bar{L}} EU(\Pi_1 L_1 + \Pi_2(\bar{L} - L_1) + \bar{w})$$

where L_1 is land devoted to crop 1, $\bar{L}$ is total land available, Π_i is the per acre net return from crop i, $\bar{w}$ is exogenous wealth, and $EU(\cdot)$ is the expected utility function. The standard procedure with such a set-up has been to use first-order conditions and structural assumptions to estimate risk aversion (see Just and Just forthcoming for a review and critique of this literature).

The first-order condition is given by

$$(2)\qquad E[U'(\Pi_1 L_1 + \Pi_2(\bar{L} - L_1) + \bar{w}) \times (\Pi_1 - \Pi_2)] = 0.$$

To estimate this relationship, we would need to know the parameters of the per acre profit distribution for each crop, the parameters of the utility function at initial wealth, the amount of land utilized for each crop, and the net return per acre. The mean per acre return is given by $L_1(\Pi_1 - E(\Pi_1)) + (\bar{L} - L_1)(\Pi_2 - E(\Pi_2))$. The first-order condition can then be approximated using a Taylor series expansion around this mean per acre return as

$$(3)$$
$$E[U'(E(\Pi_1)L_1 + E(\Pi_2)(\bar{L} - L_1) + \bar{w}) \times (\Pi_1 - \Pi_2)] + E\begin{bmatrix} U''(E(\Pi_1)L_1 + E(\Pi_2) \\ \times (\bar{L} - L_1) + \bar{w}) \\ \times [L_1(\Pi_1 - E(\Pi_1)) + (\bar{L} - L_1) \\ \times (\Pi_2 - E(\Pi_2))](\Pi_1 - \Pi_2) \end{bmatrix} = 0.$$

Suppose that land devoted to crop 2 (i.e., productive activity) yields a certain return (e.g., leasing out the land, placing it in conservation reserve, etc.). Then, the first-order condition in (3) can be rewritten as

$$(4)\qquad L_1 = \frac{E(\Pi_1) - \Pi_2}{\rho \sigma_1^2}$$

where ρ is the coefficient of absolute risk aversion. So far, we have implicitly assumed that the farmer knows the probabilities in her/his profit distributions and takes them as given. If instead she/he used a probability weighting function of the sort proposed in Kahneman and Tversky (1979), the perceived mean would be diminished, and the perceived variance would be lower as long as the probability density of any given outcome remained below the fixed point of the weighting function (i.e., where the perceived probability is equal to the actual probability). We will assume that probability density is below this fixed point, thus resulting in a perceptual lowering of probabilities. This is consistent with the certainty effect, where crop 2 is given added weight because it does not involve uncertainty. Thus, the land allocation contains information about both risk aversion and the degree to which risky choices are diminished by probability weighting. Let the perceived profit and variance be given by $E(\Pi_1) - \gamma_1 > 0$ and $\sigma_1^2 - \eta_1 > 0$, respectively, where $\gamma_1 > 0$, $\eta_1 > 0$ are parameters representing the impact of probability weighting, so that

$$(5)\qquad L_1 = \frac{E(\Pi_1) - \Pi_2 - \gamma_1}{\tilde{\rho}\left(\sigma_1^2 - \eta_1\right)}.$$

Estimating ρ in equation (4) when equation (5) is the true relationship biases the estimate of absolute risk aversion as follows:

$$(6)\qquad \rho = \tilde{\rho}\,\frac{(E(\Pi_1) - \Pi_2)\left(\sigma_1^2 - \eta_1\right)}{(E(\Pi_1) - \Pi_2 - \gamma_1)\sigma_1^2}.$$

Thus, the estimate of ρ, which ignores probability weighting, may be higher or lower than the true coefficient $\tilde{\rho}$ depending on the sign of the first term in the denominator. Specifically, the sign of estimated risk aversion will be wrong if $\gamma_1 > E(\Pi_1) - \Pi_2$. This is troubling since we typically know too little about probability distortions to account for them empirically and must rely instead on potentially misleading estimates of ρ.

So far, these decision functions reflect the risk aversion embodied by a single decision over a stand-alone risk. Small changes in the net return distribution will alter land allocation in a way that primarily responds to risk aversion but not probability weighting. Thus, estimates based on decision changes may net out the effect of probability weighting and may therefore be more stable than estimates based on a stand-alone decision. To demonstrate, consider an isolated change first in the expected value of the net return distribution, then in its variance. With either isolated change, we could estimate risk aversion by assessing how an individual modifies her/his production decisions in response to this distributional change. If the expected value changes and absolute risk aversion when facing both the original and the modified gamble is the same,[5] then equations (4) and (5) imply

$$(7) \qquad \Delta L_1 = \frac{\Delta E(\Pi_1)}{\rho\sigma_1^2}$$

$$(8) \qquad \Delta L_1 = \frac{\Delta E(\Pi_1)}{\tilde{\rho}\left(\sigma_1^2 - \eta_1\right)}.$$

If we cannot account empirically for probability weighting and must estimate equations (7) when (8) is the true model, we would again bias our estimate of risk aversion but in a different way than before

$$(9) \qquad \rho = \tilde{\rho}\frac{\left(\sigma_1^2 - \eta_1\right)}{\sigma_1^2}.$$

Now, suppose that means were fixed, and the variance was allowed to change by θ (and resulting in an additional perceived change in variance θ_η due to probability weighting). As before

$$(10) \qquad \Delta L_1 = \frac{E(\Pi_1)}{\rho}\left(\frac{1}{\sigma_1^2} - \frac{1}{\left(\sigma_1^2 + \theta\right)}\right)$$

$$(11) \qquad \Delta L_1 = \frac{E(\Pi_1)}{\tilde{\rho}} \times \left(\frac{1}{\sigma_1^2 - \eta_1} - \frac{1}{\left(\sigma_1^2 - \eta_1 + \theta - \theta_\eta\right)}\right)$$

where $|\theta_\eta| < |\theta|$. Equations (10) and (11) imply

$$(12) \qquad \rho = \tilde{\rho}\left[\frac{\theta}{\theta - \theta_\eta}\right]\left[\frac{\sigma_1^2 - \eta_1 + \theta - \theta_\eta}{\sigma_1^2 + \theta}\right] \times \left[\frac{\sigma_1^2 - \eta_1}{\sigma_1^2}\right].$$

Equations (9) and (12) share an important property; namely, in both cases the estimate of risk aversion ρ, while biased, must have the correct sign. To see this, note that our assumptions regarding the extent of probability weighting implies that each of the braced terms in equation (12) must have the same sign numerator and denominator, ensuring that $\tilde{\rho}$ has the same sign as ρ. This suggests an important potential benefit of assessing risk preferences using comparisons between gambles: when probability distortion is present but cannot be extracted from empirical estimates of risk aversion, measures of marginal risk aversion may difference away this distortion and may therefore be more stable and more representative of behavior than standard measures based on stand-alone gambles.

Behavior encompassing both risk aversion and probability distortion is important in practice because the fixed effects of risk enter into the standard first-order conditions as in equation (5). Thus, the common approach that relies on first-order conditions will produce estimates that are not predictive of how behavior will change given changes in the parameters of the distribution of profits or other production characteristics. Rather, these estimates may be highly biased and potentially imply risk-averse behavior, when in fact the underlying behavior is risk loving on the margin. This forms our first research hypothesis.

> H1: In the presence of probability weighting behavior—including the certainty effect—risk behavior on the margin will not conform to standard risk measures for an individual risk.

Alternatively, estimating risk aversion using observed changes in behavior, although it will not allow us to identify probability distortions, should lead us to operationally more accurate predictions of behavior. Hence, marginal risk aversion—whether based on changes in expected value or variance—may be a more robust notion of behavior for the typical problems encountered by economists. Our second research hypothesis is derived from this theory.

---

[5] This is equivalent to assuming local constant absolute risk aversion. The gambles are assumed to differ only by expected value and thus may induce different levels of absolute risk aversion under more general assumptions. Černý (2004) notes that, unless the gambles involve large skewness, absolute risk aversion should be approximately stable regardless of the properties of the utility function.

H2: In the presence of probability weighting, measures of marginal risk aversion are more stable than standard risk aversion.

To test these hypotheses, we derive simple measures of marginal and standard risk aversion and compare them across several gambles. We define our indicator of standard risk aversion as $R \equiv 1 - \frac{WTP}{E(x)}$, where WTP is the willingness to pay for a gamble defined implicitly as $EU(x - WTP + w_0) = U(w_0)$, where x is the risky outcome, w_0 is the initial wealth, and $E(x)$ is the expected value of the risky outcome. Thus defined, this indicator is interpreted much like ρ; namely, $R > 0\,(< 0)$ indicates an individual who is risk averting (loving) in the standard sense. Since R is based on observed behavior, any probability weighting is captured by this indicator. Given probability distortion, therefore, R is akin to $\tilde{\rho}$ and may actually have the opposite sign of ρ as illustrated by equation (6) because probability weighting leads the individual to behave as if she/he was risk averse in the expected utility sense. Thus, discounting all risky choices can impact decision functions in a way that may obscure the underlying risk preferences.

In general estimates of marginal risk aversion could be obtained by estimating changes in probability distributions, then using a standard utility function as a marginal utility function in order to incorporate observed changes in behavior. We leverage the experimental structure of our data to derive an even simpler indicator of marginal risk aversion. We focus on changes in $E(x)$ holding variance constant and changes in variance holding $E(x)$ constant. In the first case we take $\Delta WTP < \Delta E(x)$ as evidence of marginal risk aversion. In the second case we take $\Delta WTP/\Delta\sigma < 0$ as evidence of marginal risk aversion. We define these two indicators of marginal risk aversion as

$$(13)\quad MR_{E(x)} = \begin{cases} 1 - \dfrac{\Delta WTP}{\Delta E(x)} & \text{if } \Delta E(x) > 0 \text{ and } \Delta\sigma = 0 \\ \dfrac{\Delta WTP}{\Delta E(x)} - 1 & \text{if } \Delta E(x) < 0 \text{ and } \Delta\sigma = 0 \end{cases}$$

$$(14)\quad MR_{\sigma} = -\frac{\Delta WTP}{\Delta\sigma} \quad \text{if } \Delta E(x) = 0 \text{ and } \Delta\sigma \neq 0.$$

These marginal risk aversion indicators are interpreted much like ρ and R: positive (negative) values indicate that an individual is risk averting (loving) on the margin.

According to MR_{σ}, an individual who is willing to pay more to increase the variance of a gamble holding expected value constant is marginally risk loving (i.e., she/he would prefer more risk on the margin). $MR_{E(x)}$ is a more subtle measure of marginal risk aversion. If the starting point is a certain payoff, $MR_{E(x)}$ will mirror standard measures of risk aversion. But if the starting point is a gamble, $MR_{E(x)}$ relates to the changes in absolute risk aversion that occur as wealth changes and thus to the notion of increasing or decreasing absolute risk aversion (IARA or DARA). For example, we could think of a shift in mean (holding all else constant) as adding wealth and offering the same gamble, where WTP is now augmented by the increase in wealth. In this case $MR_{E(x)}$ will be less than 1 if WTP to increase wealth by Δ, given the same gamble, is less than Δ. In other words, the individual has higher absolute risk aversion when wealth is increased. While the possibility of IARA or DARA has been explored thoroughly in the literature, econometric applications typically assume that IARA or DARA must prevail over the entire range of wealth, which explicitly excludes the possibility of such a switch.

Data

This article uses data from the Salem and Perambalur districts of Tamil Nadu state, India (see figure 1). These data were collected with local support from Tamil Nadu Agricultural University and funding from the Agricultural Biotechnology Support Program (USAID-Cornell University). Ten enumerators surveyed 290 households in three Perambalur villages (Annukur, Pandagapadi, and Namaiyur) and three Salem villages (Vellaiyur, Kilakku Raajapalayam, and Kavarparnai). The team collected data in two parts. In the first part, enumerators administered a detailed household questionnaire focused on farmers' management decisions, valuation of seed traits, risk exposure, and wealth. In the second part, the team conducted experiments with farmers to elicit their valuation of hypothetical yield distributions. Farmers earned money (Rupees, or Rs.) according to their performance in the experiment.

Figure 1. Map of surveyed villages in Salem and Perambalur districts of Tamil Nadu (TN), India

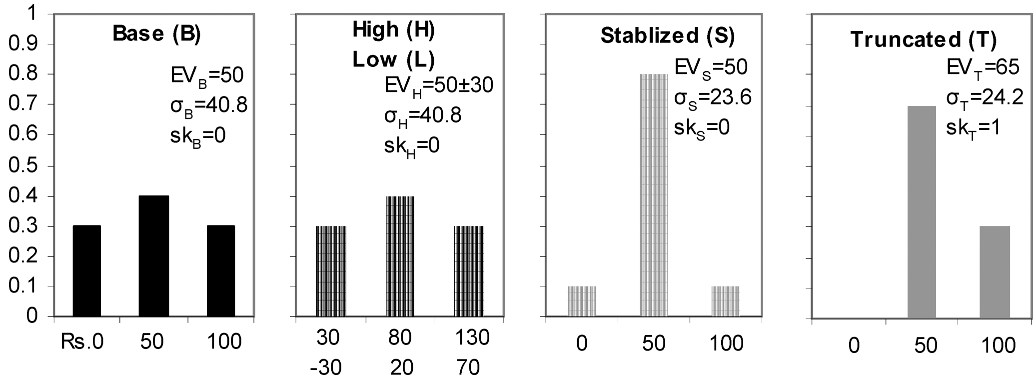

Figure 2. Marginal probability distributions for distribution types used in experiment (payoffs in Rupees (Rs.) on *x*-axis)

The experiment consisted of a series of hypothetical farming seasons. At the beginning of each season, farmers were offered a "seed" with a known Rupee-payoff distribution. This distribution was explained simply and repeatedly and shown graphically in order to facilitate farmers' understanding of the payoff distribution implied by a given "seed." The distribution of a particular "seed" was represented by ten chips in a small black bag. There were three colors of chips, each representing a "harvest" payoff: blue (high), white (average), and red (low). The distribution was modified by changing the proportion of blue, white, and red chips in the bag. Farmers' valuation of the seed was elicited using an open-ended question and the well-known Becker-DeGroot-Marschak (BDM) mechanism (1964). As shown in figure 2, we focus on four payoff distributions from the experiment: a benchmark base gamble (B), a high gamble with a higher mean payoff (H), a low gamble with a lower mean payoff (L), a stabilized gamble with lower variance (S), and a truncated gamble with positive skewness (T). Figure 2 shows the marginal probability distributions and the expected value (EV), standard deviation (σ), and skewness (sk) for each of these gambles.[6] Every farmer valued each of these payoff distributions several times, first during practice rounds, then in a final high-stakes round.

To control for learning and ordering effects, all farmers started and ended with the benchmark gamble B, and each farmer's valuation of

[6] These simple typological distributions were chosen to facilitate farmers' understanding of the experiment. We used simple pictures like those in figure 2 to capture each distribution and explain the experiment to farmers.

Table 1. Descriptive Statistics for Indicators of Standard Risk Aversion (R) and Marginal Risk Aversion (MR)

		Mean	Std. Dev	Std. Error	Min	Max	% Risk Averse (>0)	N
	B	0.11	0.24	0.014	−0.70	0.71	69%	290
	S	0.11	0.27	0.016	−0.90	0.80	74%	
R	H	0.27	0.21	0.012	−0.24	0.75	89%	
	L	−0.50	0.74	0.043	−4.00	0.85	23%	
	T	0.21	0.23	0.014	−0.52	0.85	81%	
MR_{σ}	[B–S]	0.02	0.82	0.048	−2.71	4.46	49%	290
	[B–H]	0.55	0.57	0.033	−1.00	2.25	81%	
$MR_{E(x)}$	[B–L]	−0.51	0.50	0.029	−1.95	1.17	10%	

B was computed as his average valuation of B over these two rounds. Between these two B rounds, gambles H, L, S, and T were randomly ordered (see Lybbert 2004 for more details about the experiment). These data have been used elsewhere to assess poor farmers' valuation of pro-poor seeds (Lybbert 2006). That analysis of these data suggested that farmers' valuation is responsive to expected yield but not to higher moments of the yield distribution and points to the distinction between standard and marginal risk aversion that we make in this article.

Analysis

In this section, we analyze the standard and marginal risk aversion indicators defined above. We first present descriptive statistics based on these indicators and examine their consistency with theories of probability weighting and the two resulting hypotheses, H1 and H2, discussed above. We then compare the two indicators in an effort to detect deviations and patterns between standard and marginal risk aversion.

According to table 1, a majority of individuals in our data are risk averse according to our indicator R. The one exception is for gamble L, which was the only one that involved losses and an expected value below that of gamble B. This may result from a simple anchoring and adjustment mechanism. Having already assessed and valued gamble B, individuals may anchor on their WTP for this base gamble and then fail to adjust sufficiently up or down when valuing the alternatives. Such an anchoring on gamble B would create exactly the pattern we observe for gambles L and H in table 1. Without first seeing gamble B, we would expect the mean of R for H and L to be much closer to that of B. This behavior is similar to those supposed by prospect theory, albeit with a reference gamble rather than a reference point.

Since we presented the B gamble as the baseline distribution, and all other gambles were presented sequentially as variations of this baseline distribution, we compute marginal risk aversion relative to the B gamble. Table 1 also shows the percentage of marginal risk-averse individuals, based on MR, for each gamble pairing. Across these gambles, fewer individuals are marginal risk averse than standard risk averse. This provides support for H1 and thus the probability weighting model, in that marginal risk behavior does not resemble average behavior. For example, fewer individuals are marginal risk averse when moving from B to S than are standard risk averse for gamble S. The same holds for gambles H and L as well. In the case of MR_{σ}, three-quarters of our subjects are risk averse on average, while one-half are risk loving on the margin, implying that many are simultaneously standard risk averse and marginal risk loving.

While table 1 indicates that marginal risk aversion diverges substantially from standard risk aversion, it also suggests that our measure of marginal risk aversion may be no more stable than our standard risk aversion measure. To test H2 more directly, we assess the stability of these measures using the absolute value of the coefficient of variation (i.e., relative standard deviation) for each measure of risk aversion. Since the structure of our experiment allows us to compute multiple standard and marginal risk aversion indicators for each subject, we are able to generate a coefficient of variation for R and MR for each subject. Figure 3 displays the density of these subject-specific coefficients of variation. Statistical tests of the null hypothesis that R and MR are equally variable within individuals confirm what is visually evident in this figure: our marginal measure is indeed no more

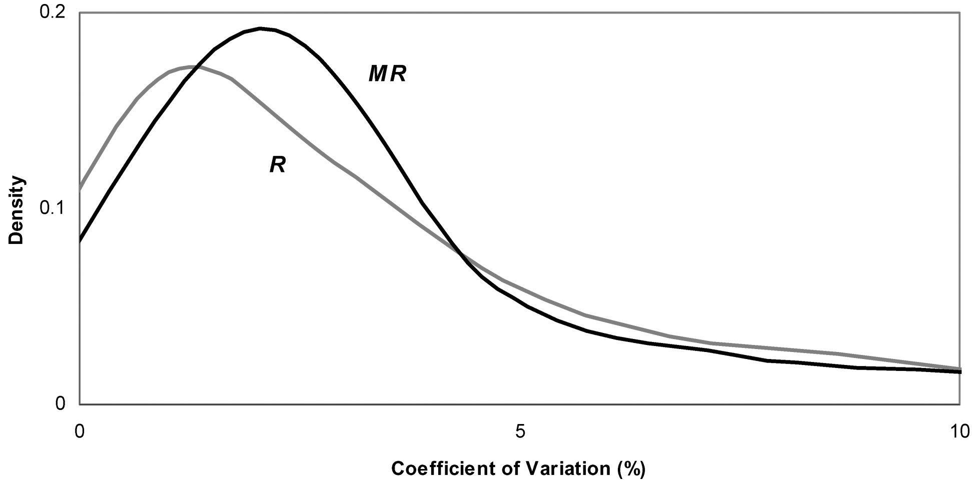

Figure 3. Distribution of subject-specific relative standard deviations (absolute value of coefficients of variation) for standard risk aversion (*R*) and marginal risk aversion (*MR*) across gambles

stable than our standard measure.[7] This refutes H2, the research hypothesis that marginal measures of risk aversion should be more stable than standard measures, and calls into question the probability weighting model of marginal versus average risk aversion. If, in fact, marginal risk aversion is nearly as volatile as overall risk aversion, we must search for an alternative or supplementary explanation for the divergence in measures of standard and marginal risk aversion. One possibility might be a tendency to anchor one's valuation of a gamble on a relevant benchmark—the base gamble B in this case—and adjust this valuation in response to changes in the gamble in reference to this benchmark. In isolation such a decision heuristic is consistent with rejecting H1 and failing to reject H2. If both probability distortion and anchoring-and-adjustment play a role in decision-making, marginal risk-aversion measures may not difference away the probability weighting bias. Instead, the bias in $R[\mathrm{B}]$ due to probability distortion may carry over to *MR*, since subsequent gambles are assessed in reference to one's valuation of B.

To assess the strength of such an anchoring-and-adjustment tendency, we take a closer look at how our simple measure *R* changes across the six gambles in the experiment. Since the B gamble is the logical benchmark, we nonparametrically regress $R[i]$ for $i = \{\mathrm{H, L, S, T}\}$ on $R[\mathrm{B}]$.[8] As shown in figure 4, these four regression lines are ordered top-to-bottom by EV. Their intercepts suggest that individuals who are risk neutral when facing gamble B become risk averse when facing a more favorable gamble with higher EV (H) or lower variance (S) or both (*T*) but become risk loving when facing a gamble with lower EV (L). The slope of these regression lines for H, S, and T is clearly less than one. Parametric versions of these regressions confirm that the slope of these lines is statistically significantly less than one.[9] The regression line for L, on the other hand, has a slope closer to one. Together, the evidence in figure 4 suggests that individuals anchor on gamble B when valuing changes in risk. In the case of the Rs. 30 increase in EV entailed in gamble H, they anchor on B and increase their WTP by much less than Rs. 30, and overall standard risk aversion increases as a result. The opposite occurs when EV falls by this amount (as in gamble L), and individuals become much more risk loving according to standard measures as a result of insufficient reductions in their WTP. This is strong evidence that individuals are not evaluating each gamble in isolation, even though the gambles were presented to them separately. Rather, they seem

[7] Standard *t*-tests for differences in the mean coefficient of variation for each also fail to reject the null that the means are equal. Kolmogornov–Smirnov tests also fail to reject the null that the two distributions are equal.

[8] All of the nonparametric results that we report are computed using an Epanechnikov kernel. In all cases other kernels produce nearly identical results.

[9] The 95% confidence intervals on the slopes of these lines are as follows: [0.18, 0.38] for H, [0.32, 0.53] for T, and [0.34, 0.58] for S.

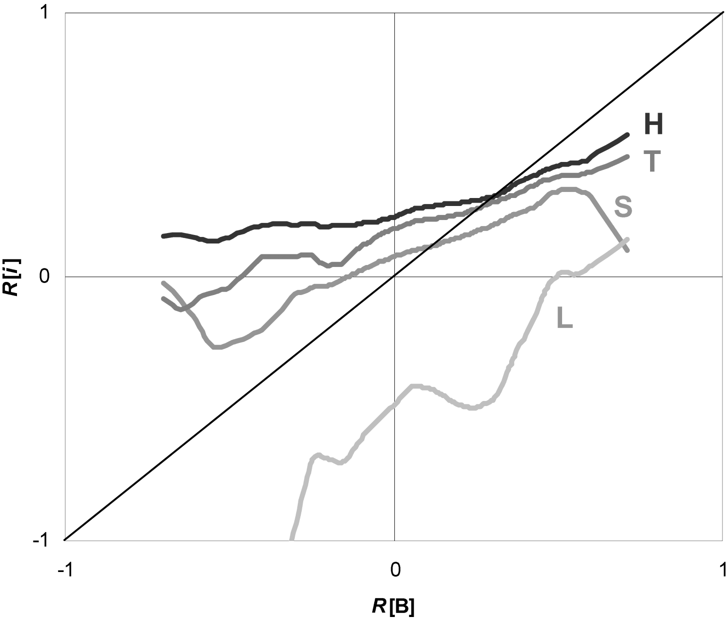

Figure 4. Standard risk aversion indicator for base gamble (R[B]) compared to the risk-aversion indicator for high, low, stable, and truncated gambles ($R[i]$, $i = \{H, L, S, T\}$)

to make incremental adjustments after comparing gambles.

Next, consider the relationship between standard and marginal risk aversion. Expected utility theory disallows simultaneously risk-averse and risk-loving behavior at a given local point on the utility function and therefore requires a simple positive relationship between the two in the case of changing variance. Thus, if we regress R on MR, expected utility theory would require the resulting line to have a positive slope and pass through the origin. Figure 5 displays nonparametric regressions of this relationship using B as the base gamble for computing MR. The top, middle, and bottom panel pairs in figure 5 assess this relationship when moving from this base gamble to S, H, and L, respectively. Thus, adjacent panels share the same MR x-axis. The only difference between adjacent panels is the gamble used to compute R: panels on the left use R[B], while those on the right use $R[i], i = \{S, H, L\}$.[10] In the positive (negative) quadrant of these graphs, standard and marginal measures of risk aversion indicate consistent risk-averse (loving) behavior. In the other quadrants these measures indicate simultaneous risk-averse and risk-loving behavior. As shown by the percentage of individuals in each quadrant, the share of individuals who are simultaneously risk averse and risk loving is heavily shaped by whether we base our stand-alone indicator R on the base gamble or comparison gamble. In the top panels 57% are simultaneously risk averse and risk loving when using R[B], but this drops to 27% when using R[S].

The patterns evident in figure 5 are worth discussing in greater detail. The top left panel regresses R[B] on MR_σ[B–S] and suggests that those who initially display greater standard risk aversion in gamble B (i.e., high R[B]) display lower marginal risk aversion—or even marginal risk seeking—when moving from B to S. The opposite pattern is evident in the top right panel, which regresses R[S] on MR_σ[B–S] and suggests that those who display greater marginal risk aversion when moving from B to S (i.e., high MR_σ[B–S]) ultimately display greater standard risk aversion in gamble S. The middle and bottom sets of panels show this pattern when regressing R[B] and $R[i]$ on $MR_{E(x)}$[B–i] for $i = \{H, L\}$. This remarkably consistent pattern provides further evidence that individuals anchor on gamble B and adjust their valuation of subsequent gambles with reference to this anchor. In particular,

[10] The left-hand panels of figure 4 might suggest to some that some sort of reversion to the mean is taking place, in that those with lower standard risk aversion display higher marginal risk aversion. However, the right-hand panels show the exact opposite relationship, contradicting the notion of reversion to the mean and making it difficult to dismiss the notion of a reference gamble.

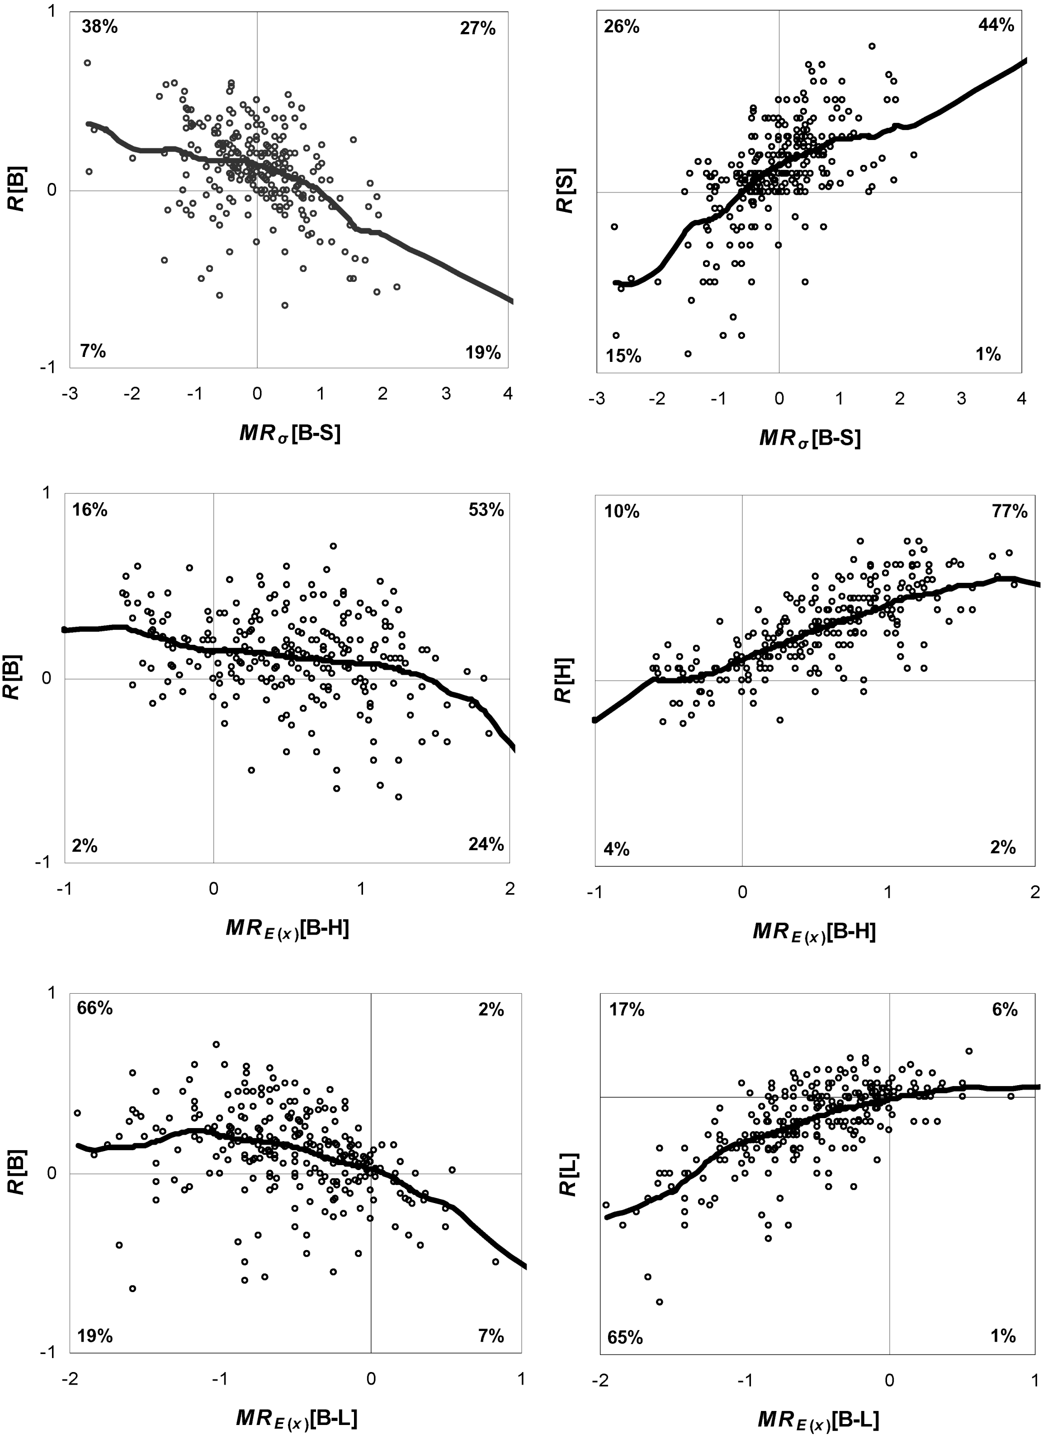

Figure 5. Kernel regressions of marginal risk aversion (*MR*) relative to gamble B on risk aversion (*R*) (percentage of observations in each quadrant noted)

risk-averse individuals with a low WTP for gamble B heavily adjust their valuation when offered variations on gamble B. In contrast those with a higher WTP for gamble B adjust their valuation very little when offered variations on gamble B. This difference in adjustment is enough to completely reverse the risk-aversion ranking of participants on average.

Last, the regression curves in the three right panels are of nearly the same location and slope, although the ranges of marginal risk aversion on the horizontal axis are different for each gamble. If this relationship holds true more generally, estimation of marginal and standard risk aversion resulting from a pair of gambles could be predictive of the measures of risk aversion over a wide range of gambles. Interestingly, the prevalence of DARA (as measured by $MR_{E(x)}$) appears to be dependent upon standard risk aversion in the base gamble. Those who display standard risk loving in the base gamble display IARA behavior. This is, of course, flipped when comparing to an alternative gamble. We see this same pattern of behavior when examining the response to shifts in variance (MR_{σ}), with greater aversion to risk on the margin resulting from standard risk-averting behavior in the base gamble. Such behavior could add to the rapidly expanding literature exploring reference points under uncertainty, thus providing one explanation for why small changes in gambles that straddle the gain/loss divide may not result in rapid changes in risk aversion. Such a change in a gamble may therefore only elicit a mild readjustment rather than the large, discrete loss aversion effect found by Kahneman and Tversky (1979).

Conclusion

In this article, we differentiate between marginal and standard risk aversion. Marginal risk aversion embodies the changes in behavior observed when the underlying gambles change. Standard risk aversion is embodied by the behavior observed for a single gamble. Empirical work has focused solely on standard risk aversion and its impacts on production and other behaviors. Well-documented behavioral anomalies suggest that standard risk aversion measures will not be predictive of behavior when underlying risky choices change. This may result from probability weighting or from reference-based behavior. If behavior is consistent with probability weighting, marginal risk aversion measures will be more robust to estimation bias and have a greater correlation with behavioral changes due to shifts in probability distributions. While the results of our experiment confirm strong differences between marginal risk aversion and standard risk aversion, the measures of marginal risk aversion appear to be just as volatile as measures of standard risk aversion, refuting probability weighting as the sole explanation for the difference between marginal and standard risk aversion. Instead, the evidence suggests a strong tendency—either in isolation or in conjunction with probability distortion—to anchor on a reference gamble and make adjustments on the margin in response to changes in risk.

The relationship between marginal and standard risk aversion depends primarily on how the gamble is adjusted from a reference gamble. In comparing two gambles, one of which involves a marginal distributional change, marginal risk aversion is negatively related to standard risk aversion measured over the reference gamble, while positively related when measured over the alternate gamble produced by this marginal change. We find remarkable regularity in this relationship, both in terms of slope and intercept, across a set of several gambles. Standard risk aversion measures display unreasonable variation across gambles, undermining researchers' ability to predict behavior in one gamble from that in another relying solely on standard risk aversion measures. The results of this experiment appear to suggest that we need *both* measures of standard risk aversion from a reference gamble (e.g., the current year yield distribution) and a measure of marginal risk aversion (e.g., marginal response to changes in year over year yield distributions) to make robust predictions of risk response.[11]

As a profession, we need to acknowledge that empirical estimates of behavioral risk parameters must be interpreted relative to the modeling construct from which they are derived. Our simple crop allocation model shows how strict interpretations of model parameters can be misleading. If, for example, a model assumes away some pervasive behavior like misperception of probabilities, these assumptions should cause us to be careful in interpreting any resulting risk aversion coefficients. Just and Just (forthcoming) examine this issue in the context of parameter identification for a wide set of production models. Further work is needed to link estimation techniques with potential biases and interpretation of results—perhaps in a way similar to how the contingent

[11] This observation applies to the use of proxies for risk as well as models employing assumed risky distributions. Variables that proxy for risk should measure marginal changes as well as absolute levels of risk.

valuation literature has dealt with response bias.

Although we do not explicitly propose a new model of decision making, a model allowing a reference gamble-dependent decision rule would allow for conformity to rational axioms governing any single decision, while overall behavior could violate each of the three axioms (order, continuity, independence) when the reference gamble changes. This model would explain an important class of expected utility violations—similar to the certainty effect—where the comparison gamble is important in determining behavior.[12] This class may be of particular importance to applied economists if, for example, farmers respond to changes in risk using decision heuristics that focus only on the changes rather than on the entire risk portfolio. Our work supports the general consensus in the field of agricultural economics that risk impacts input allocation decisions. Given a static set of homogeneous risky decisions, it may be useful to employ the standard EU framework to measure risk response. In practice, however, finding homogenous and static risks may be a challenge.

Estimating risk preferences has never been an easy business, and this article suggests that reference-based decision making under risk may be pervasive and has important implications for applied economics. In particular, future empirical work should widen its focus to include measures of marginal risk aversion. The introduction of matching techniques that make use of panel data could lead to more relevant models of belief formation and adaptation over time. Behavioral models that address the changes in perceptions and in behavior stand a greater chance of capturing general—rather than idiosyncratic—behavior.

[Received February 2008; accepted January 2009.]

References

Arrow, K.J. 1971. *Essays on the Theory of Risk Bearing*. Chicago: Markham Publishing.

Bar-Shira, Z., R.E. Just, and D. Zilberman. 1997. "Estimation of Farmers' Risk Attitude: An Econometric Approach." *Agricultural Economics* 17:211–22.

Becker, D.E., M.H. DeGroot, and J. Marschak. 1964. "Measuring Utility by a Single-Response Sequential Method." *Behavioral Science* 9:226–32.

Bell, D.E. 1985. "Disappointment in Decision Making under Uncertainty." *Operations Research* 33:1–27.

Bellante, D., and R.P. Saba. 1986. "Human Capital and Life-Cycle Effects on Risk Aversion." *Journal of Financial Research* 9(1):41–51.

Camerer, C. 1995. "Individual Decision Making." In J.H. Kagel and A.E. Roth, eds. *The Handbook of Experimental Economics*. Princeton, NJ: Princeton Univerisity Press, pp. 587–617.

Cerný, A. 2004. *Mathematical Techniques in Finance: Tools for Incomplete Markets*. Princeton, NJ: Princeton University Press.

Chavas, J.P., and M.T. Holt. 1990. "Acreage Decisions under Risk: The Case of Corn and Soybeans." *American Journal of Agricultural Economics* 72:529–38.

——. 1996. "Economic Behavior under Uncertainty: A Joint Analysis of Risk Preferences and Technology." *Review of Economics and Statistics* 78:329–35.

Cohn, R.A., W.G. Lewellen, R.C. Lease, and G.G. Schlarbaum. 1975. "Individual Investor Risk Aversion and Investment Portfolio Composition." *Journal of Finance* 30(2):605–20.

Dyer, J.S., and R.K. Sarin. 1982. "Measurable Multiattribute Value Functions." *Operations Research* 27:810–22.

Edwards, W. 1953. "Probability-Preferences in Gambling." *American Journal of Psychology* 66:349–64.

Friedman, M., and L.J. Savage. 1948. "The Utility Analysis of Choices Involving Risk." *Journal of Political Economy* 56(4):279–304.

Hong, S.C. 1983. "A Generalization of the Quasilinear Mean with Applications to the Measurement of Income Inequality and Decision Theory Resolving the Allais Paradox." *Econometrica* 51(4):1065–92.

Just, D.R., and H.H. Peterson. 2003. "Diminishing Marginal Utility of Wealth and Calibration of Risk in Agriculture." *American Journal of Agricultural Economics* 85(5):1234–41.

Just, R.E., and D.R. Just. Forthcoming. "Global Identification and Tractable Specification Possibilities for Risk Preference Estimation." *Annals of Econometrics (special issue of Journal of Econometrics)*, in press.

Kahneman, D., and A. Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47(2):263–91.

[12] A large body of literature documents numerous violations of expected utility theory, including systematic violation of all three axioms. We address the certainty effect and a similar class of violations of the independence axiom where marginal and overall risk behavior diverge. Other important violations of expected utility theory (such as the common ratio effect) are not addressed directly in this article.

Kimball, M.S. 1990. "Precautionary Saving in the Small and in the Large." *Econometrica* 58(1):53–73.

Landskroner, Y. 1977. "Intertemporal Determination of the Market Price of Risk." *Journal of Finance* 32(5):1671–81.

Lattimore, P.K., J.R. Baker, and A.D. Witte. 1992. "The Influence of Probability on Risky Choice: A Parametric Examination." *Journal of Economic Behavior and Organization* 17(3):377–400.

Lins, D.A., S.C. Gabriel, and S.T. Sonka. 1981. "An Analysis of the Risk Aversion of Farm Operators: An Asset Portfolio Approach." *Western Journal of Agricultural Economics* 6:15–29.

Lybbert, T.J. 2006. "Indian Farmers' Valuation of Yield Distributions: Will Poor Farmers Value 'Pro-Poor' Seeds?" *Food Policy* 31(5):415–41.

——. 2004. "Technology and Economic Development: Three Essays on Innovation, Pricing, and Technology Adoption." Ph.D. dissertation, Cornell University.

MacCrimmon, K.R., and S. Larsson. 1979. "Utility Theory: Axioms versus Paradoxes." In M. Allais and O. Hagen, eds. *The Expected Utility Hypothesis and the Allais Paradox*. Dordrecht, Holland: D. Riedel, pp. 333–409.

Marra, M.C., and G.A. Carlson. 1990. "The Decision to Double Crop: An Application of Expected Utility Theory Using Stein's Theorem." *American Journal of Agricultural Economics* 72(2):337–45.

Morin, R.A., and A. Fernandez–Suarez. 1983. "Risk Aversion Revisited." *Journal of Finance* 38(4):1201–16.

Pennings, J.M.E., and P. Garcia. 2001. "Measuring Producers' Risk Preferences: A Global Risk-Attitude Construct." *American Journal of Agricultural Economics* 83(4):993–1009.

Pope, R.D., and R.E. Just. 1991. "On Testing the Structure of Risk Preferences in Agricultural Supply Analysis." *American Journal of Agricultural Economics* 73:743–48.

Preston, M.G., and P. Barrata. 1948. "An Experimental Study of the Auction Value of an Uncertain Outcome." *American Journal of Psychology* 61:183–93.

Rabin, M. 2000. "Risk Aversion and Expected-Utility Theory: A Calibration Theorem." *Econometrica* 68(5):1281–92.

Saha, A., C.R. Shumway, and H. Talpaz. 1994. "Joint Estimation of Risk Preference Structure and Technology Using Expo Power Utility." *American Journal of Agricultural Economics* 76:173–84.

Siegel, F.W., and J.P. Hoban. 1982. "Relative Risk Aversion Revisited." *Review of Economics and Statistics* 64(3):481–87.

Tversky, A., and D. Kahneman. 1992. "Advances in Prospect Theory: Cumulative Representation of Uncertainty." *Journal of Risk and Uncertainty* 5(4):297–323.