

ECONOMIC GROWTH CENTER

YALE UNIVERSITY

Box 1987, Yale Station
New Haven, Connecticut

CENTER DISCUSSION PAPER NO. 286

ATTITUDES TOWARDS RISK: IMPLICATIONS FOR ECONOMIC AND
PSYCHOLOGICAL THEORIES OF AN EXPERIMENT IN RURAL INDIA

Hans P. Binswanger

June 1978

Note: Center Discussion Papers are preliminary materials circulated to stimulate discussion and critical comment. References in publications to Discussion Papers should be cleared with the author to protect the tentative character of these papers.

ATTITUDES TOWARDS RISK: IMPLICATIONS FOR ECONOMIC
AND PSYCHOLOGICAL THEORIES OF AN EXPERIMENT IN RURAL INDIA*

An earlier paper (Binswanger, 1978a) described an experiment carried out with around 320 people in rural India to measure their pure attitudes towards risk. It also measured correlations between individual characteristics such as wealth, age, etc., and the measured risk attitudes. This paper confronts one basic set of results from these experiments with various theories of behavior under uncertainty to check their empirical relevance for the rural households studied, who belong to the poorest of the world. The theories have been developed by statisticians, economists and mathematical psychologists.

The experiment makes practically no theoretical restrictions; individuals choose among alternatives where increasing expected returns can only be purchased by increasing risk or dispersion of outcomes, and the alternatives would be ranked more or less risky almost regardless of which definition of risk is used. The reason for wanting to make a commitment to a specific theory is that only with such a commitment can the experimental results be used to make predictions of behavior in different risky situations such as the individual's farming decisions. The weaker the theoretical restrictions, the weaker the predictions which can be made. Without a theory, it is as if the human mind was perfectly divided into different compartments, with observed behavior in one compartment not allowing us to make predictions of how decisions would be made in another compartment.

The first section briefly recalls the key experimental results.

*Hans P. Binswanger is an Associate of the Agricultural Development Council presently stationed at the Economic Growth Center of Yale University. The experiment on which this paper is based was carried out while the author was stationed at the International Crops Research Institute for the Semi-Arid Tropics, Hyderabad, India, and with its generous support. I would like to thank B. C. Barah, R. D. Ghodake, S. S. Badhe, M. J. Bhende, V. Bhaskar Rao, T. Balaramaiah, N. B. Dudhane, Rekha Gaiki, K. G. Kshirgar, Madhu Nath and Usha Rani, who helped in carrying out the experiment. Harvey Lapan made particularly helpful comments on an earlier draft.

Section II considers safety-based rules of thumb and finds that only one of them is not inconsistent with the data. (It is not inconsistent because it offers no prediction of how people faced with the experiment should behave). Section III is a broad overview of utility based theories used by economists and psychologists, and the empirical evidence is used to show that only some of these models are consistent with the observed behavior. Section IV then tests and rejects the hypothesis of asset integration, i.e., it finds that one cannot write a stable utility function across wealth states but only across gains and losses, with the zero point of that function shifting as wealth changes. Section V proposes a functional form for the utility function which is consistent with the experimental evidence and a final section tries to pull the threads together.

THE KEY EXPERIMENTAL RESULT

The experiment--carried out with over 300 individuals selected at random from six villages of the semi-arid tracts of Maharashtra and Andhra Pradesh--consisted of a sequence of games with real and high payoffs of the following nature: People were offered a set of 8 choice alternatives in which higher expected return could only be "purchased" for a larger standard deviation. The alternatives A to F are described in the top panel of Table 1. Each consists of a "good luck" and a "bad luck" outcome with probability of $1/2$ which is decided on a toss of a coin. Alternative zero is a certain outcome in which the individual is simply paid Rs 50 whereas alternative F pays nothing or Rs 200 with equal probability. The alternatives D* and D are stochastically dominated by alternatives B, C and E respectively. Each alternative is given a name classifying the extent of risk aversion of the person who chooses it. These names are arbitrary and more precise measurements of risk aversion are discussed below.

The game was played--and actually played--7 or 8 times over a period of about 6 weeks with much time left for reflection. The game sequence starts with 5 games at the 0.50 Rs level, in which all amounts in Table 1 were divided by 100. The payoffs are then increased to the 5 Rs level at which all amounts of Table 1 are divided by 10. After two weeks, the game is played at the level shown and hypothetical questions are asked of each participant how he would play at the 500 Rs level, in which all amounts of Table 1 are multiplied by 10. Note that monthly wages of unskilled laborers in this area are roughly 60 to 80 Rs. The amounts were therefore large for these people. For a detailed des-

1. Effects of Payoff Size on Distribution of Risk Aversion
and on the Partial, Absolute and Relative Risk Aversion Coefficients.

	Extreme O	Severe A	Inter- mediate B	Moderate C	Slight-to- Neutral E	Neutral-to- Preferred F	-Inefficient D*	D	N.OBS
1) THE ALTERNATIVES AT THE 5 RS LEVEL									
Bad Luck (50%)	50	45	40	30	10	0	35	20	
Good Luck (50%)	50	95	120	150	190	200	125	160	
2) FREQUENCIES OF CHOICES AT DIFFERENT LEVELS									
GAME LEVEL									
0.50 Rs (No. 2)	1.7	5.9	28.5	20.2	15.1	18.5		10.1	119
5 Rs (No. 7)	0.9	8.5	25.6	36.8	12.0	8.5		7.7	117
50 Rs (No. 12)	2.5	5.1	34.8	39.8	6.8	1.7		9.3	118
500 Rs (No. 16) (no payment)	2.5	13.6	51.7	28.8	0	0.9		2.5	118
3) TRADEOFF BETWEEN E AND SE: $Z = \Delta E / \Delta SE^a$									
	.90	.735	.585	.415	.165	≤ 0			
4) PARTIAL RISK AVERSION S^a									
All levels	≥ 7.5	3.61	1.20	.51	.158	≤ 0			
5) ABSOLUTE RISK AVERSION A									
0.50 Rs	≥ 10.7	5.17	1.71	.728	.226	≤ 0			
5 Rs	≥ 1.07	.517	.171	.0728	.0226	≤ 0			
50 Rs	$\geq .107$.0517	.0171	.00728	.00226	≤ 0			
500 Rs	$\geq .0107$.00517	.00171	.000728	.000226	≤ 0			
5000 Rs	$\geq .00107$.000517	.000171	.0000728	.0000226	≤ 0			
6) RELATIVE RISK AVERSION (R) AT WEALTH = Rs 10000									
0.50 Rs	≥ 10700	5170	1710	728	226	≤ 0			
500 Rs	≥ 10.7	5.17	1.71	.728	.226	≤ 0			

a) Risk aversion measures can only be computed for indifference points between any two efficient alternatives. Therefore, one can only assign an interval to each of the alternatives O to F. To compute a unique value for each alternative, one can take the mean of the measures at the endpoints of each interval. In the case of Z, the interval length did not vary greatly and the arithmetic mean was used. For S and the other measures, the interval length increases sharply from alternative O to F and therefore the geometric mean was used (with the exception of alternative E which has a zero endpoint and where the arithmetic mean was used). Partial risk aversion was computed by solving the equation for indifference between alternative X and Y, using the constant partial risk aversion function $U = (1-S)M^{1-S}$, where M is certain income.

cription of the method and tests of its reliability, see Binswanger (1978a).

The second panel of Table 1 shows the pattern of behavior of those 118 individuals who played up to the 50 Rs level, (many more played only up to the 5 Rs level but their behavior at those low levels is fully consistent with the behavior of the smaller sample). When payoffs are small, (0.50 Rs), we find nearly 50% of individuals in the intermediate and moderate risk aversion categories (B and C). Over a third of individuals show a nearly neutral or risk preferring behavior pattern (E and F) and less than 10% are extremely or severely risk averse (O and A). When game levels rise, the proportion of individuals in the intermediate and moderate categories rises till it reaches 80% of individuals in these two classes. Near neutral and risk-preferring behavior virtually disappears, only one out of 118 individuals chose F. On the other hand, the fraction of extreme and severely risk-averse choices stays virtually constant, only at the 500 Rs level does it rise by roughly 5%. At higher game levels, the risk aversion distribution is thus single peaked with most of its weight in the two intermediate and moderate risk aversion classes.

MODELS BASED ON SECURITY MOTIVES

This class of models has recently been reviewed by Anderson (1975), whose exposition I will follow. Anderson also gives references to the authors who proposed or worked with the various rules. In all models the individual has an overridingly important security motive, either in terms of minimum income goals or in terms of critical probabilities of experiencing losses below a critical line.

Some of the rules to be considered below assume that individuals

have specific probability target (P^*) in their mind and are concerned primarily with the income level which can be achieved with that target probability. Other rules assume that individuals have specific target incomes or subsistence incomes in their mind and are concerned primarily with reducing the probability of falling below their target incomes. Finally lexicographic rules operate with constraints incorporating both target probabilities and target incomes.

The same basic rules can have different predictions depending on how the income stream is defined over which these targets are measured. When the probability distribution on overall income--including any new prospects--is considered, the models assume what we may call income integration: the individual integrates any new prospect with his old ones and considers only final income states. Alternatively we can have the same rules apply to the probability distribution of only a new prospect. The individual then operates with income source specific targets, and this case is considered first.

Income Source Specific Targets

Safety-Fixed or Maximin: This rule involves the maximization of the minimum income that can be obtained with a probability of at least a crucial P^* , i.e.

$$\text{Maximize } d, \text{ subject to } P(X < d) < P^* \quad (1)$$

X is the achieved income from the game and P^* is the critical low probability.

Maximin is a special case of safety fixed where P^* is zero. Consider Figure 1 which shows the cumulative frequency distributions of the game alternatives at three levels in semi-logarithmic scale such that equal proportional shifts of game outcomes correspond to equal horizontal distances.


It is clear from Figure 1 that the decision maker who follows this rule must, at all game levels, choose alternative 0 since that gives him the highest income with a fixed probability of $P = 0 \leq P^*$.¹ Considering Table 1 we find that at best 2.5% of the individuals can follow this decision rule since the proportion of individuals choosing alternative zero is less than that. Safety fixed without income integration thus has to be rejected as a framework to describe the observed behavior.

Safety Principle: This rule involves the minimization of the probability that income X will fall below some fixed disaster level d^* . The disaster level is usually taken as a subsistence income or alternatively as a customary income from all sources of income. This interpretation will be taken up later but let us first consider a source specific target or customary income. This may make little sense for an unusual income source such as this experiment. At best we could assume that individuals would take the unusual opportunity to obtain sufficient income from the windfall to pay back a fixed liability which they had incurred earlier, or that the experimental income would make possible unusual expense such as a pilgrimage for which they need a fixed sum which they otherwise could not get.² This interpretation will be tested below.

¹When the probability target is greater than $1/2$, the decision rule implies the choice of alternative E.

²Quite a few individuals did indeed use the income from the game for pilgrimages.

Figure 1
The Income Alternatives When Decision Rules Involve
Source Specific Income and Probability Targets



The rule can formally be written as

$$\text{Minimize } P(X \leq d^*) \quad (2)$$

Now consider figure 1. Several cases have to be distinguished.

1. Very high target income (d^*_1). In this case the theory makes no prediction for any game level, because the target income cannot be achieved in any event.

2. High target income which can be achieved by the good luck outcomes at the highest game level (d^*_2, d^*_3). With this target income we have no prediction for the 5 and 50 Rs level for the reason mentioned above. But at the 500 Rs level, alternative zero has a 100% probability of falling below d^*_3 and d^*_2 , while the probability for alternatives C and E is only 50% in the case of d^*_3 . Hence the decision maker will choose either C or E and be indifferent between the two. For d^*_2 on the other hand he will choose alternative E. We can therefore see that if a set of alternatives has a low probability of achieving an income target the safety principle pushes individuals into choosing risky alternatives. This is the case even with more complex continuous probability functions because only the risky alternatives will have a positive probability of achieving those high levels, hence they will generally have smaller probabilities of not reaching those levels and the individual minimizes those probabilities.

3. Lower target incomes which can be achieved by the good luck outcomes of all alternatives (d^*_4 to d^*_7). Under d^*_7 all 500 Rs level alternatives have a zero probability of not achieving the target income. Hence the rule gives no prediction at that level. For d^*_5 choices of 0 or C are implied and a unique choice of 0 is only implied by d^*_4 , which

is very close to the sure income. The indeterminacy present here can be overcome by lexicographic rules considered below. Note that an individual who has income target d^*_6 will have to choose E at the 50 Rs level (for reasons explained in paragraph 2) and move to the less risky alternative O or C at the 500 Rs level. Hence, a tendency towards more risk aversion at higher game levels can be implied in this decision rule.

Low Income Targets (d^*_8 to d^*_{10}). For such targets the indeterminacy of choice becomes acute for the high game levels. d^*_{10} implies no prediction anywhere, while d^*_9 and d^*_8 imply only indifference between O and C at the respective game levels and no prediction elsewhere.

The decision rule thus implies almost random choice except if the target income is either close to the sure outcome at a given game level or close to the good luck outcomes of the risky games, and even in these cases only for one game level of the sequence. It would thus imply a risk aversion distribution which is fairly close to a uniform distribution. It almost impossible to rationalize the game results of Table 1 in terms of this decision rule.

One might object that the specific game level induces in the individual such a target commensurate with the game level and with a committed expense of the individual which realistically falls into the set of outcomes. But that leads to a theory with practically no predictive power. To get a prediction for individuals one has to elicit a target income (which would change over time) for any set of alternatives which confronts them and this is an unfeasible research program.

Lexicographic Rules (LSF). Roumasset has proposed two lexicographic rules which are designed to sharpen the predictions. These rules operate with both a fixed probability target and a fixed income target and assume that the individual first wants to satisfy a safety constraint. This constraint says that he will not accept any alternative which does not give him a target income with a fixed target probability, i.e.

$$\text{Prob } (X < d^*) \leq P^* \quad (3)$$

When the constraint is satisfied the individual will maximize expected income. Note first that to implement it, we need to know both an income target and a probability target which is an ambitions information requirement.

LSF 2: The individual maximizes expected income when the safety constraint is satisfied. When it is not, he follows the safety fixed rule.

$$\begin{aligned} \text{i.e.} \quad & \max E \\ & \text{s.t. Prob } (X \leq d^*) \leq P^* \\ & \text{Otherwise} \\ & \text{maximize } d \\ & \text{s.t. Prob } (X \leq d) \leq P^* \end{aligned} \quad \begin{matrix} (3) \\ (1) \end{matrix}$$

1) High income targets ($d^*_1, d^*_2, d^*_3, d^*_4$):

No one (or only alternative zero at the 500 level) satisfies (3), and therefore the individual chooses alternative zero at all game levels.

2) Intermediate income targets (d^*_5, d^*_6, d^*_7): The high game alternatives satisfy the constraints progressively at the 500 Rs level, and the individual will choose (among those who satisfy it) the one with highest expected return. Under d^*_7 he will choose alternative E at the highest game levels, while continuing to choose zero at the lower game levels.

Hence the decision rule implies--for at least some individuals--a shift from taking no risk at low game levels to taking the highest risk at the highest game level, which is totally inconsistent with the game results.

3) Low income targets: As the income target moves down to d^*_{10} , all alternatives at the higher game levels satisfy the constraint (3) and these individuals should choose conservatively at only the lowest game levels but choose E at the higher levels.

The rule thus clearly implies that, as game levels rise, the risk aversion distributions in Table 1 should shift to the right. This is the opposite of what happened in the experiment.

LSF 1: This rule implies maximization of expected returns when the safety constraint is satisfied, but using the safety principle when it, is not:

$$\begin{aligned} &\max E \\ &\text{s.t. Prob } (X \leq d^*) \leq P^* \end{aligned} \quad (3)$$

otherwise

$$\min P(X \leq d^*) \quad (2)$$

1) High income targets (d^*_1 to d^*_3): In these cases the constraint is not satisfied and the prediction are as for the safety first rules, i.e. no prediction for d^*_1 or "risk taking" for d^*_2 and d^*_3 at the 500 Rs level without any prediction for the low game levels.

2) Intermediate income targets (d^*_4 to d^*_7): As the level goes down, more and more alternatives satisfy the constraint (3). And a small shift of the income target from d^*_4 to d^*_7 implies a shift from alternative zero to alternative E. Target d^*_7 implies a choice of either C or E at the 50 Rs level and a choice of E at the 500 Rs level. Thus we do find a possibility of observing increasing risk aversion as game levels rise. As the target income falls further, within each of the game levels the cycle from "no prediction" → choice of most risky alternative → choice of last risky alternative → choice of most risky alternative repeats itself, implying that in the measured risk aversion distributions we should observe all choices.

3) Low income targets: These imply that at all high game levels the choice of E must be made with a possibility of conservative choices only at the lowest game levels.

Evidently we do not know the distribution of target incomes and thus cannot really predict the distribution of choices in the games. But consider the three following alternatives:

Everyone has high target incomes: since that implies no predictions for the low game levels we should expect a fairly uniform distribution of risk aversion at the low and intermediate game levels. This is not the case since alternatives 0 and 1 are almost never chosen at low game levels. Furthermore, at the 5 Rs level we already observe a reduction in the proportion of nearly-risk neutral choices.

Everyone has low target incomes: this implies that at high game levels most observations should be concentrated at the risk-neutral level of the spectrum.

A fairly even distribution of target incomes over the interval d^*_{10} to d^*_1 . This implies that at least some individuals should find that all of the high payoff alternatives exceed their target incomes and choose E at high levels, which is again not what we observe.

Note how important the knowledge of the target incomes is for these theories. Slight variations in them sharply, and cyclically, alter the preference ordering of alternatives with outcomes in the range of the target incomes. The burden on accuracy of measurement of target incomes is high.

Since all the rules with income source specific targets have implications which are inconsistent with the observations we will now move to consider overall income targets.

Overall Income Targets or Income Integration

We shall see below that in many cases, to make predictions about behavior, we must know the probability distribution of income $F(I)$ with which the game prospects have to be aggregated (Income Integration). Note again, that this adds a large information requirement. Often we will have to know the--presumably subjective--probability distribution of aggregate income of the individual as well as a probability target and/ or an income target.

If $f(I)$ is the density function of an individual's income and $F(I)$ its cumulative density, and if L is the bad-luck outcome of an alternative while U is the good-luck outcome, then the cumulative distribution of income with the prospect (say C) is

$$F(I + C) = 1/2 F(I + L) + 1/2 F(I + U)$$

Graphically this can be shown in figure two: $F(I)$ is not shown, but instead $F(I + 0)$ is the distribution of income with the sure prospect of 50 Rs. (It is found by simply shifting $F(I)$ by 50 Rs). The cumulative distributions $F(I + L)$ and $F(I + U)$ are found by shifting $F(I + 0)$ by the appropriate amounts and $F(I + C)$ is the simple sum of half of the values of the two at each point. The graph also shows $F(I + E)$, i.e. the distribution of the more risky prospect E (without its corresponding distributions of bad-and good-luck outcomes).

The basic feature of the new cumulative probability density functions $F(I + 0)$ to $F(I + E)$ is that they will cross. Because the expected return of each succeeding alternative from 0 to E is higher, the cumulative density functions will cross before or when they reach $P = 1/2$. But this is about all what we can say about the new functions without knowledge of the old ones. Figure 3 shows the lower portions of cumulative distributions for three alternatives at two of the game levels.

Integrating the prospects with total income changes the predictions for the Safety-Fixed rules: Let D be the income achievable with a fixed probability P^* and let the individual behave according to the safety fixed rule

$$\text{Max } D$$

$$\text{s.t. Prob } (I \leq D) \leq P^* \quad (1a)$$

Figure 2
Cumulative Distributions of Game Alternatives with an Overall Income Goal





Figure 3
Cumulative Distribution of Game Alternatives at Different Levels
With Overall Income Goal



Without "income integration" this rule implied choice of alternative zero in all cases where $P^* \leq 1/2$. But with income integration (figure 3) the choice depends on target probabilities. When they are low, such as P^*_1 , the riskless alternatives will be chosen. As they rise, the choice shifts to more and more risky alternatives. This rule, therefore, is unable to give us any prediction about the behavior of individuals when confronted with the experimental at any game level. Hence it cannot be falsified by the experiment. It is difficult to imagine an experiment which could falsify it, since knowledge of personal probability targets and personal probability distributions of overall income is required. Given the difficulties faced in eliciting certainty equivalents by interviews (see Binswanger 1978a), the prospects for falsification or support of this model are not good.

Income integration has less impact on the qualitative predictions under the safety principle, i.e.

$$\text{Min } P(I \leq D^*) \quad (2a)$$

Note that, if the notion of target income is based on physical subsistence requirements, variations in the target income relative to the distributions shown in Figure 1 have a clear interpretation: Poor people should have high subsistence incomes and rich people low ones relative to their probability distributions of income. Or said otherwise, poor people should have high probabilities of not achieving their subsistence income while rich people should have low ones.

- 1) Very high subsistence income D_1^* , which cannot be reached in any case (very poor people): The model implies no predictions since all probabilities of not reaching D^* are equal to 1.
- 2) High subsistence D_2^* to D_5^* (poor people). The model implies that poor people should choose the risky alternatives in some of the low level games. In the highest-level game, some of the poor people (those with subsistence incomes between D_2^* and the right endpoint of the distribution of E_{500}) should choose the most risky alternatives. Evidently this is inconsistent with the experimental evidence. As the subsistence income falls from D_2^* to D_5^* the choice will shift rapidly to the least risky alternative zero. The model continues to imply that some of the poorest people behave in a risky manner at high payoffs, and that they should do so at higher frequencies at the low game levels. Yet we find very few poor people choosing E and F at any game level.
- 3) Low income targets (rich people). The model quickly loses predictive power, first for the high game level (D_7^*) and then the lower game levels (D_9^*).

Lexicographic rules: Both lexicographic rules imply the maximization of expected returns when the safety constraint is satisfied, i.e.

$$\max E$$

$$\text{s.t. Prob. } (I \leq D^*) \leq P^* \quad (3a)$$

These rules imply that we should observe the richer individuals (low D^*) at the risk-neutral end of the distribution of risk attitudes. Furthermore, since the probability that all alternatives satisfy the constraint

(3a) rises as the game level rises, we should expect an increasing frequency of risk neutrality as the game level rises, or at least not a decreasing frequency. This is totally opposite to what we observe in the experiment and the lexicographic rules can be rejected.

Note that this ground for rejection is independent of how the target income is defined, i.e. does not depend on a physiological need interpretation of the target income. Consider the following most favorable case for the lexicographic models: Individuals, on the basis of their assets and incomes form a target income and probability such as D_7^* and P_1^* for which, under "usual" conditions, they have a fairly low probability of falling short of. This gets around the objections that a safety based theory should not predict risk taking on the part of the poorest groups, which arises either by a high target probability or a physiological subsistence income. Nevertheless the lexicographic rules with customary income still predict that the proportion of risk-neutral individuals cannot fall as the game level rises. This is contradicted by the evidence. A customary subsistence income model also poses the additional difficulty that one has to know three elements to make individual predictions: personal subsistence income, personal probability target and subjective probability distribution of income. This is a tall order. Add to this the evidence that people have serious difficulties in evaluating low probabilities, (i.e. the tails of distributions) and you have created a model more complex than all utility based models.

The only security based model which is not inconsistent with the experimental evidence is the safety-fixed model with income integration. The

only reason for its survival is that it offers no prediction whatsoever unless personal probability targets and subjective probability distributions of overall incomes are known. The advocates of the model have yet to propose how to measure these before the model can become operational.

Finally note that all models which operate with subsistence income targets imply predictions which have been contradicted in this study. The sample includes some of the poorest households of the world. If subsistence income models do not operate here, it is hard to imagine where else they could operate.

One way to defend these rules against the experimental result just reported is to say that the experimental situation--being too simple --falls into a wholly different compartment of behavior than production or labor supply decisions, and that the evidence is therefore inconclusive. That, of course, makes it much harder to subject the models to empirical tests.¹ A modified version of this objection would not reject the experimental evidence but reason that humans divide decisions into usual decisions and unusual decisions with the game clearly being unusual. Rules of thumb would then apply only to the usual decisions. If we were to categorize all windfall gains as unusual, this would deprive the security based theories of much of their usefulness. New technologies offered at subsidized costs by the government, or employment opportunities on short run rural projects are windfall gains and one would like to predict technology adoption or labor supply to such projects using the theoretical frameworks. Usual behavior

¹For econometric evidence against subsistence income as an important determinant of wage rates in rural India, see Bardhan, (1978).

can often be observed and need not be predicted.

Finally one might object that the experiment is not a good test because it does not subject the individual to losses. In an earlier paper I showed that, when people were given the money for the game one day in advance and had to bring it back to play and put at risk, their decisions did not differ statistically from the ones when the payouts came only after the game was played. They treated opportunity losses much like real losses. Even without this, to maintain that losses must occur to test the theories is one form of compartmentalization of the mind, i.e. it says that any windfall income is not counted when it comes to compute subsistence income. This brings us back to assume multiple independent income targets for different sources of income, which reduces the predictive power of the rules and increases the difficulty in using them for predictions. In the decades since theorizing with subsistence income targets started, little progress has been made at measuring one single target and to make it empirically operational. Thinking of measuring multiple targets is a nightmare to an empirical investigator.

MODELS BASED ON UTILITY COMPARISONS

Utility-based models of behavior under uncertainty have been developed by statisticians, economists and mathematical psychologists. An very careful review of these models is Luce and Suppes (1965). Relatively few major new theoretical proposals have been made since that review. In this section the basic forms of utility functions are reviewed in their deterministic version. For a review of probabilistic choice theories the reader is referred to Luce and Suppes.¹

The basic tenet of all utility models is an attempt to associate with each action or prospect a_j a unique utility value U_j such that a decision maker will choose or prefer a_1 over a_2 ($a_1 \succeq a_2$) if, and only if, the utility value of a_1 exceeds the utility value of a_2 or is equal to it; i.e.,

$$a_1 \succeq a_2 \Leftrightarrow U(a_1) \geq U(a_2) \quad (4)$$

where \succeq indicates a relationship of preference or indifference. The outcome of each action depends on which event E_i will occur out of an Exclusive Exhaustive set of Events (EEE). An EEE is a set of events from which one event E_i must happen but more than one cannot happen. In all formulations the decision maker is assumed to associate objective probabilities or subjective probabilities (or decision weights) with each event of the EEE.

Furthermore, the action a_j associates an outcome x_{ij} with each

¹In most probabilistic models the assertion is made that the probability of choosing a_1 will be larger than that of choosing a_2 if $U(a_1) > U(a_2)$. In some versions the probability of choosing a_1 is related to the difference between $U(a_1)$ and $U(a_2)$. In random models the utility function is assumed to be random, while the choice is deterministic.

event E_i . The outcome X_{ij} can be any object providing satisfaction or utility but for this discussion we will consider it as the money income (or wealth) accruing to the individual if he chooses action a_j and the event E_i occurs. Money income is evaluated according to a utility function $U(X)$ which associates a real number to the income X .

For simplicity, the discussion which follows will be restricted to discontinuous probability functions. All the theories discussed below have the following structure of utility function in common:

$$U_j = \sum_i k_i U(X_{ij}) \quad (5)$$

where k_i is a "probability" measure in an objective or subjective sense. Note that in this formulation utility and probability combine multiplicatively for each individual event and that these products are summed over the set of events. Amos Tversky (1967) has tested this basic formulation experimentally for a large class of more specific models which can be derived from 5 and found his experimental results to be consistent with additivity.

Basic disagreements among theorists arise about the form of the utility function associated with outcomes, about whether to use objective or subjective probabilities, about how subjective probabilities are formed and whether subjective probabilities over a set of EEE should add up to one or not. (All agree that objective probabilities add up to one.)

Table 2 presents the six subclasses of utility models consistent with equation 5. The first three have been developed and refined by mathematicians, statisticians and economists, while the last three stem from the work of experimental and mathematical psychologists. These

Table 2: Assumptions of Models based on Utility Comparisons

	Utility Function ^a	Emphasis ^b	Adding up of Subjective Probabilities ^c	Revision of Subjective Probabilities ^d	Asset Integration ^e	Author ^f
1. Expected Income EI	$U^1 = \sum_i P_i X_i$	normative			Usually assumed	
2. Expected Utility UI	$U^2 = \sum_i P_i U(X_i)$	normative			Usually assumed	Bernoulli (1738) Von Neumann-Morgenstern (1947)
3. Subjective Expected Utility SEU Statistics-Economics version	$U^3 = \sum_i h_i U(X_i)$	normative	Yes	According to Bayes' Theorem	Usually assumed	Ramsey (1926), de Finetti (1937) Savage Mosteller and Nogee (1953)*
4. Subjective Expected Utility SEU Psychology version	$U^4 = \sum_i h(P_i) U(X_i)$	predictive	Yes	Learning Theories	No	Coombs & Beardslee* (1955), Edwards* (1953, 1954 a,b)
5. Subjective Expected Income SEI (C-E approach)	$U^5 = \sum_i h(P_i) X_i$	predictive	Yes		No	Preston & Baratta* (1948), Griffith* (1949), Sprowls* (1953), <u>Handa</u> (1977)
6. Nonadditive Subjective Expected Utility NASEU Prospect Theory version	$U^6 = \sum_i h(P_i) U(X_i)$ $U^6 = U(X_1) + \sum_{i=2}^n h(P_i) [U(X_i) - U(X_1)]$	predictive	No	Learning Theories	No	Edwards (1955) Kahnemann and Tversky (1977)

Notes:

- a/ U = Utility index, X = Money outcome, P = Objective probability, h = Subjective probability.
b/ Indicates whether the emphasis of the authors proposing the theory was predictive or normative-prescriptive.
c/ Indicates whether subjective probabilities are assumed to sum to one or not.
d/ Indicates whether the theoretical framework proposes a theory of how the subjective probabilities are revised when the decision maker receives new information or forecasts.
e/ Indicates whether the utility function is usually assumed to be stable in final wealth states.
f/ The authors who provided axiomatic foundations are underlined. Stars refer to early experimental work using the models explicitly or implicitly.

two traditions of modeling have their origins in different questions, although these questions later were mixed up. The statistical-economic tradition started by asking how a person should behave in a situation of uncertainty if his decisions were to be consistent with his preferences and with basic tenets of logic and consistency. All writers seem to agree that it is in a person's interest to behave according to the statistical-economic model which fits his preferences and state of information, i.e., they are generally regarded as the superior normative models. The psychological theories can result in various inconsistencies of choice. But the psychologist's basic interest is in finding regularities in how people actually behave, i.e., the basic purpose is predictive or positive. Economists use normative models for some purposes and predictive models for others. The dominant tendency has been to use the normative models for prediction purposes as well—at least as a first approximation. There has been surprisingly little interest in economics to experimentally test whether the normative models are useful for predictive purposes.

Probability Formation or Preference

In Table 2 the symbol P_i stands for objective probabilities while h_i stands for subjective probabilities. The subjective probabilities are called by different names: personal probabilities, probability preferences (Edwards, 1954b; Preston and Baratta, 1948), Certainty Equivalent Utility Index (Handa, 1977), or simply weights (Kahnemann and Tversky, 1977).

The Expected Income (EI) model was the starting point for all later

work. Few authors argue that people behave according to this model since it implies no aversion to risk, but it is often stated that the proper set of economic institutions (contingent markets, insurance, financial instruments) allow people and firms to behave as if they were maximizing expected income. The expected utility model was proposed by Bernoulli (1738) in terms of objective probabilities and its axiomatic foundations in terms of objective probabilities were derived by von Neumann and Morgenstern (1947). Savage (1954) provided an axiomatic treatment of the Subjective Expected Utility (SEU) theory which is based on personal probabilities. Despite cautionary remarks by Savage¹, the SEU model has been widely used for deriving predictions of how people behave when confronted with certain situations (see, for example, Diamond and Stiglitz 1974, or Rothschild and Stiglitz 1971, 1970).

The SEU theory also has developed a framework for analyzing how personal probabilities are revised when the decision maker receives additional evidence. They are revised according to Bayes theorem which explains the posterior probabilities as a combination of prior probabilities and the likelihood of an observed event.²

The psychological models have their roots in experimental work on gambling. In a series of articles Preston and Baratta (1948), Griffith (1949), Sprowls (1953) and Edwards (1953, 1954 a,b) showed that in gambling

¹The title of Savage's book is Foundations of Statistics which implies that--at least initially--he had a normative theory in mind. Many of his comments (see also Savage, 1972) stress the benefit of teaching a decision maker to use the model and learn to introspect about his personal probabilities, and Savage fully expected that actual behavior would often diverge from the postulated behavior.

²For a good description of this method see Anderson et al., 1977.

situations people act as if they had clear preferences for certain probabilities rather than others (the probability preference literature). In most of these experiments the respondents were told the objective probabilities of winning and losing and attempts were made to relate the objective probabilities to subjective ones. Most experimenters found a relationship which looks roughly as follows:




Figure 4: General shape of probability preference functions.

Figure 4 indicates an overestimation of low probabilities and an underestimation of high probabilities with a crossover point varying from 0.05 to 0.25 on the objective probability scale.¹ Edwards' experiments, however, showed marked preference for probabilities of $1/2$ when the expected income from the alternatives considered was positive. He found marked preferences for low probabilities of losing large amounts when expected income from the games considered were negative. The probability preference literature thus shows that a unique and stable functional relationship between the subjective and objective probabilities could not be found, but Edwards hoped that a few fairly stable functional relation-

¹Crossover points in the neighborhood of 0.20 were found in a bidding game for play money by Preston and Baratta, and by Griffith and McGlothlin (1956), using data from betting on horses. Sprowls (1953) finds crossover points between 0.05 to 0.25. Shuford's (1959) experiments and Davidson, Suppes and Siegel also found underestimation of high probabilities but crossover points varied.

ships could be identified for different "types" of situations: one for the case when all outcomes are positive, one when all are negative, and a few more for various situations with some outcomes positive and some negative.

The early work of Preston and Baratta , Griffith and Sprowls--implicitly or explicitly--was based on the Subjective Expected Income (SEI) model which weighs probabilities of events but assumes linear utility for money outcomes. When measuring probability preference functions experimentally this assumption is usually necessary, or otherwise the same set of choices can be interpreted as arising out of an Expected Utility Model and be used to measure the curvature of the utility function. In fact Mosteller and Noguee (1951) computed utility weights as well as probability preference in alternative interpretations of their poker-dice experiment. Edwards recognized this ambiguity early and his experimental techniques roughly adjust for nonlinear utility by having individuals choose among bets with equal expected value.¹ His early theoretical work then developed what may be called the psychological version of the subjective expected utility model (Edwards, 1955). This psychological version differs from the statistical-economic version only in the stress on the stable functional relationship between subjective and objective probabilities.

¹ Another technique for measuring both expected utility and subjective probabilities simultaneously was developed by Gordon M. Becker (1962).

The Subjective Expected income model has recently been revived by Handa (1977) under the name of Certainty Equivalent approach and provided with a set of axiomatic foundations.¹ Note that--well in the economic tradition--the axioms which Handa uses do not require a stable functional relationship between subjective and objective probabilities, in fact they could also be revised according to Bayes' rule. Nevertheless, all of Handa's illustrations and predictions are based on two functions between subjective and objective probability: one function for probabilities attached to gains and one function attached to losses.

The probability functions $h(P_i)$ underlying the SEI and the psychological SEU approach confront one difficulty discussed in detail by Edwards (1954, p. 398, 1962): If a subjective probability function such as the one in Figure 4 is stable and is to be used for event sets with more than two possible events, then the subjective probabilities over the full event set cannot sum to one. This may best be illustrated with an example: In evaluating an action a_j which has three equally likely outcomes with objective probability of $1/3$ each, the subjective probabilities of Figure 4 would all either exceed one-third or fall short of it, except at the crossover point. They could therefore sum to one only at the place where objective probability equals subjective probability. This problem is more general and Edwards has shown that subjective probabilities associated with a stable probability function can sum to one if and only if subjective probability equals objective probability everywhere. To get any further with the idea of a stable functional relationship between

¹For a critique of Handa's approach, see Fishburn (1978).

objective and subjective probabilities Edwards (1962) relaxed the additivity assumption for the subjective probabilities and developed the Non-Additive Subjective Expected Utility (NASEU) model number 6. As Kahnemann and Tversky (1977) point out, this model faces a fundamental difficulty: consider the prospect $a = (x + \epsilon, p; x, 1 - p)$, i.e., a two outcome gamble where the first outcome is equal to the second outcome plus ϵ . The utility of this prospect, according to the NASEU model, will be $U^6 = h(p) U(X + \epsilon) + h(1 - p) U(X)$. As ϵ goes to zero, i.e., as we approach a certain income of amount X , the utility will tend to $U^6 = U(X) [h(p) + h(p - 1)]$. Unless the sum of the subjective probabilities is one, this will not approach $U(X)$, which is inadmissible. To overcome this problem Kahnemann and Tversky (who call their version Prospect Theory) look at the utility function from its lowest possible outcome X_1 , which is certain. The utility of the action or prospect is then computed as

$$U^6 = U(X_1) + \sum_{i=2}^n h(p_i) [U(X_i) - U(X_1)] \quad (6)$$

The decision maker first evaluates the difference in utilities between each outcome and the lowest possible outcome (except for X_1 , of course). He then weighs these differences by the subjective probabilities and adds them up. This sum is then added to the utility of the lowest possible outcome. It is easily verified that the utility of a prospect such as (a) discussed above will approach $U(X)$ as ϵ approaches zero even if the subjective probabilities do not sum to one.

In addition Kahnemann and Tversky give some axiomatic foundations to their theory which allow the decision-maker to be less than fully consistent in his decisions. They also describe a substantial amount of

evidence (based on sets of hypothetical choices) which is inconsistent with the standard SEU theory but consistent with Prospect Theory, i.e., the modified NASEU model.¹

The assumption of one (or several) functional relationships between objective and subjective probabilities--which is stressed so much in the psychological literature--is quite alien to the statistical-economic theories of subjective probability. Few economically relevant situations exist in which an individual knows the objective probabilities. Exceptions are simple games or bets. (In more complex card games it is difficult to know or remember the objective probabilities even for a mathematically trained person.) Most economic decisions are not of the simple game type. Furthermore, if a decision maker were to try to make use of additional information to revise his subjective probabilities, using Bayes or any other rule, no stable relationship could exist between objective and subjective probabilities: each additional piece of information would result in a new function over a prespecified set of objective probabilities. To be fair, it has to be stressed that psychologists were often interested in specifically predicting

¹The sets of hypothetical questions were asked from about 100 individuals and resembled the following question: Choose between A and B, where

	2500 with probability 0.33	
A:	2400 with probability 0.66	B: 2400 with certainty
	0 with probability 0.1	

The amounts were in Israeli pounds. These types of questions are somewhat simpler than those asked in the interview method which I used to elicit utility functions. In particular they do not attempt to elicit a numerical scale value of utility. Eliciting utility values has been shown to be quite unreliable and potentially misleading in systematic ways (see Appendix of Binswanger 1978a).

gambling behavior and not, as statisticians and economists, in a universal theory of behavior under uncertainty.

But even if the stable link between objective and subjective probabilities is de-emphasized, it still is important to find out whether subjective probabilities do indeed add to one in actual decision situations. The psychological literature contains evidence on other behavioral regularities which are inconsistent with the subjective expected utility models to warrant further experimental efforts.¹ In view of the strong impact of payoff size on choice found in the experiment, one can only hope that future experiments will be based on real decisions rather than hypothetical ones and that experimentation will occur at larger payoffs than the extremely small ones of most of the psychological experiments.

When I started my own experiments of risk aversion I was unaware of the psychological models, some of which have only recently received renewed attention. It turns out, however, that the simple method for evaluating the utility function using 50% probabilities for two positive outcomes is the correct approach. The subjective probabilities of the two outcomes, even if not equal to one-half, would be equal to each other and therefore probability preferences cannot influence the measures of the curvature of the utility function. In fact, considering something like the SEU model, Ramsey (1926) had already proposed to use 50-50 probabilities to first evaluate utility function and then use the utility

¹In particular, Tversky (1969) has shown that one can construct alternatives for choice which will result in the observation of predictable intransitivities of choice. Grether and Plott (1977) have carefully replicated and expanded a sequence of experiments by psychologists which uncovered strong preference reversal phenomena.

function so derived to estimate personal probabilities for other choices.

Choosing 50-50 probabilities has the obvious disadvantage that component of aversion to--or preference for--risky prospects which arises out of probability preferences cannot be measured, much less can we test anything about such preferences with the experimental results.

But from the experimental results we can at least reject some of the theoretical frameworks of Table 2. The utility function is nonlinear and risk-averse in money income for all but one out of 118 individuals. This rules out the Expected Income and the Subjective Expected Income (Certainty Equivalent) Approaches as predictive models for the rural households considered.

The Subjective Expected Utility approach is in fact rejected by a direct contradiction of one of its basic axioms which Handa (1977) postulated.¹ This axiom (called "enhanced prospects" by Handa) says that the ranking of bets should be unaffected by multiplicative transformation of all of their outcomes by the same constant. In a sense it does not rule out risk aversion, but it assumes what amounts to constant partial risk aversion. Handa is uneasy about this assumption but defends it by saying that it may hold for games in the neighborhood of normal business transactions. But normal business transactions of the households considered clearly include all payoff sizes from the 0.50 to the 500 Rs game. And the ordering of prospect changes for most individuals within that range.

¹For a discussion of theoretical problems associated with Handa's approach, see Fishburn (1978).

Asset Integration and Measures of Risk Aversion

In making use of the concept of a utility function economists have usually chosen to write the utility function as a function of wealth W , i.e.,

$$U = U(W) \quad (7)$$

We must note that W must be the certainty equivalent of current wealth and that it may be difficult to measure it.¹ Economists have been postulated what Kahnemann and Tversky (1977) call Asset Integration: The action or prospect (X, P) is acceptable at asset position W if and only if $U(W + X, P) > U(W)$, where X and P are vectors of outcomes and their corresponding probabilities. The decision maker is assumed to make his decisions in terms of final wealth states and not in terms of gains and losses. A very good theoretical reason to do so is that such a theory guarantees that opportunity gains and losses are treated in the same way by the decision makers as "real" losses. It rules out all compartmentalization of decision making.

Most economists who have tried to empirically measure utility functions have, however, chosen to use functional representations of a utility function which would be stable over time in terms of income or gains and losses. (See, for example, Halter and Dean, 1971, or Anderson et al., 1977.) This is partly due to the fact that it is extremely difficult to estimate certain wealth, a point to which we will return later. In addition, Markovitz (1952) has proposed a utility

¹ Clifford Hildredth has shown that treating all risky decisions as if they were taken from a position of certain wealth, which has no random fluctuations, hides many complexities. In particular it tends to neglect the effect on decisions of covariance of the initial wealth prospects and the new venture. Other things equal, a positive covariance should tend to reduce the willingness to engage in a new venture.

function with alternating concave and convex segments which is time-stable in terms of gains and losses. He was prompted to propose this form of utility function when he tried to find a utility function which would have some of the properties of that proposed by Friedmann and Savage (1948), but which avoided some of the apparent inconsistencies with observed gambling and insurance behavior encountered by the Friedmann and Savage utility function which is time-table in terms of wealth. Markovitz proposed that gains and losses should be evaluated relative to a "usual" point of wealth, but that the utility function would adjust to new wealth positions as a person became used to it (see Figure 5 below). The psychological literature has always worked with utility functions in terms of gains and losses (or income), usually using current wealth and not customary wealth with respect to which to evaluate gains and losses.

All these approaches have therefore used a utility function of the form

$$U = V(M) \quad (8)$$

where M is certain income or the certainty equivalent of a prospect.

The difference between the two approaches is unimportant, as long as one does not make an assumption of stability over time. After all, suppose we measure a utility function $V(M)$ in terms of gains and losses (up to a linear transformation), at any given time the following relation holds under the assumption of asset integration.

$$U(W_0 + M) = V(M) \quad (9)$$

where W_0 is current (certain) wealth. We can proceed to find the func-

tional form for U which is clearly determined (up to a linear transformation) by equation (8). However, the problem is that Markovitz and the psychological tradition derive specific implications for behavior from "wiggles" of the utility function around the point of zero income, and that they assume that these wiggles remain around the value of zero income regardless of the "usual" or actual wealth position of the individual. If one writes a utility function in terms of wealth and measures "wiggles" around the present wealth position, but a smooth curvature at higher wealth levels, the "wiggles" will not be "transported" to new wealth positions. More generally, the curvature properties of the utility function in terms of wealth will remain the same at that point when one leaves that wealth position, and at the new wealth position one encounters curvature properties which were there even before one moved there.

The great advantage of the wealth formulation is that it allows one to use knowledge about the shape of a utility function measured (or mostly just assumed) before the wealth change to evaluate the behavior of an individual after a wealth change. All theoretical-analytical prediction of the effects of wealth on portfolio choice or savings behavior have been derived in this particular way. A utility function which is stable over time in terms of gains and losses cannot be used to derive such conclusions unless one also specifies how the utility function will change as wealth changes. One thus needs to measure an additional relationship. This does not imply that one cannot predict the behavior of an individual with respect to large gains and losses such as those of the order of his wealth. But that

function cannot be used once the person receives a massive income or experiences a massive loss. The new utility function may be flatter around the new wealth position or show more curvature, but it might essentially have the same "wiggles" around the point of zero income. And this point may correspond to a point on the old utility function where the latter had no "wiggles" and an entirely different curvature.

In looking at axiomatic treatments of the subjective expected utility model (see, for example, Arrow, 1971), it is clear that the set of axioms used do not imply that the utility function is one which is stable over time in terms of wealth since the theory is timeless. All consistency and transitivity axioms are specified in terms of the properties of a preference ordering over prospects or actions of the form $(\underline{X}, \underline{P})$, where \underline{X} is a vector of outcome and \underline{P} a vector of probabilities. One can always add W_0 or any other constant to all X for all prospects and obtain the same preference ordering among them with consistency and transitivity properties. Thus the axioms have no "preference" for a utility function which is stable over time in either wealth or income. Furthermore, the axioms are all about consistency of decisions over a set of prospects available now. To obtain a stable utility function in wealth we must make an additional assumption of invariance of the utility function (not the utility levels on the function) to changes in wealth or time. This assumption has usually crept in by the back door of convenience rather than being made explicit.

One way to test whether utility functions should be specified as stable in terms of wealth or in terms of gains and losses is to inspect measured utility functions for "wiggles" around zero gain and loss or for relatively larger or lower risk aversion around that point than at other points. The best way to do this would be to observe the behavior of the same individual with respect to relatively small gambles before and after a large wealth change. That is usually not possible and in the empirical section we will rely more on the evidence across individuals in different wealth classes. In an intuitive sense, if we should observe that risk aversion varies in systematic ways with payoff size and much more rapidly than with respect to equivalent wealth changes across individuals, this would tend to support the concept of a utility function in terms of income rather than wealth.

Measures of Risk Aversion and a Test of Asset Integration

The ideas above can be expressed more rigorously in terms of the behavior of various risk aversion measures. Assume that a utility function, as in Figure 5, has been measured:




Figure 5: Markovitz-type Utility Function on an Income and Wealth Scale

On the vertical axis utility has been measured both in terms of income $V(M)$ and wealth $U(W)$, while the horizontal axis measures certain income M and wealth W . M on the income scale corresponds to $W + M$ on the wealth scale. In what follows all derivatives and utilities will be measured at M and $W + M$ on each of the scales, i.e., we will be looking at the point B on the utility function.

Obviously it is true that

$$\begin{aligned} V(M) &= U(W_0 + M) \\ V_M &= U_W = U_M \\ V_{MM} &= U_{WW} = U_{MM} \end{aligned} \quad (10)$$

where the subscripts denote derivatives of the functions with respect to the subscripts at the point $(W_0 + M)$. Pratt has defined the following measures of risk aversion for a utility function in wealth:

$$\text{Absolute Risk Aversion } A = - \frac{V_{MM}}{V_M} = - \frac{U_{WW}}{U_W} = - \frac{U_{MM}}{U_M} \quad (11)$$

$$\text{Relative Risk Aversion } R = - W \frac{U_{WW}}{U_W} = WA \quad (12)$$

When we evaluate R at the point $(W_0 + M)$ this becomes

$$R = (W_0 + M)A$$

Finally both Menezes and Hanson (1970) as well as Zeckhauser and Keeler (1970) have defined the following measure

$$S = - MA \quad (13)$$

Menezes and Hanson used the term Partial Risk Aversion for S , which I will also follow. (Zeckhauser and Keeler used the term Size-of-Risk Aversion.)

Partial Risk Aversion is equal to Relative Risk Aversion for individuals with zero wealth. Partial risk aversion on a utility function in

gains and losses is--in a sense--equivalent to relative risk aversion on a utility function in wealth, with the certainty equivalent of the prospect replacing the certain wealth. As shown by several authors, the three measures are related to each other as follows at the point $(W_0 + M)$

$$R = W_0 A + S \quad (14)$$

In fact once A , W_0 and M are known, all three can be computed from A . Since A can be computed from both a utility function in terms of income as well as one in terms of wealth, it does not matter for measurement purposes with which specification one starts.

The three measures have the following interpretation: Consider the prospect $(\underline{X}, \underline{P})$ where \underline{X} and \underline{P} are vectors. Absolute Risk Aversion traces the behavior of an individual to the prospect $(\underline{X}, \underline{P})$, when his wealth rises and prospect remains the same. Decreasing absolute risk aversion is usually assumed and implies that an individual's willingness to accept a given fair gamble should rise as his wealth rises.

Relative risk aversion traces the behavior of an individual as both his wealth and the size of the prospect $(\underline{X}, \underline{P})$ rise. Let t be a scalar. We are considering the individual in a new position where he now owns wealth tW and is confronted with the prospect $(t\underline{X}, \underline{P})$. Increasing relative risk aversion was hypothesized by Arrow (1971) and implies that an individual's willingness to accept a given gamble decreases when both his wealth and all outcomes of a gamble are multiplied by the same constant.

Partial risk aversion traces the behavior of an individual when the scale of the prospect changes but his wealth remains the same. Increasing partial risk aversion implies a decrease in the willingness of the individual to take a gamble as the size of the prospect varies.

Diamond and Stiglitz (1974), demonstrate the behavior of the risk premium in relationship to the three measures. The risk premium Π is a function of both wealth W and the prospect $Z = (\underline{X}, \underline{P})$, and is the amount of money one would have to pay an individual to accept a gamble or the amount of insurance he would be willing to pay not to have to play a gamble, i.e., it is the amount of certain income which makes the individual indifferent between accepting or rejecting a gamble. Implicitly it is defined as follows:

$$U[W + E(Z) - \Pi(W, Z)] = EU(W + Z) \quad (15)$$

where E is the expectation operator. $\Pi(W, Z)$ is the absolute risk premium, $\Pi(tW, tZ)/t$ is the relative risk premium as a proportion of both wealth and size of prospect while $\Pi(W, tZ)/t$ is what we may call the "partial" risk premium as a proportion of the size of the prospect. Diamond and Stiglitz show that the absolute risk premium (relative; partial) rises or falls with t according to whether absolute (relative; partial) risk aversion is greater or less than zero.

Note that the behavior of the absolute risk aversion coefficient is the same with respect to income as well as initial wealth, i.e.,

$$A_{W_0} = A_M = - \frac{U_{WWW}U_W - U_{WW}^2}{U_W^2} \quad (16)$$

(By taking the derivative of equation (12) it is also clear that the behavior of the relative risk aversion coefficient is the same with respect to income and initial wealth.)

Equation (16) is a testable implication of the assumption

of asset integration and must be fulfilled for stable utility functions in terms of wealth. This will be done below.

In contrast to A and R, the partial risk aversion coefficient responds differently to changes in wealth than income.

$$\begin{aligned} S_{W_0} &= M \frac{\partial A}{\partial W} \\ S_M &= A + M \frac{\partial A}{\partial W} \end{aligned} \quad (17)$$

Since A is positive the response of the partial risk aversion coefficient will always be larger to changes in the prospect than to equal changes in wealth. Menezes and Hanson have also shown that, for an individual with nonzero wealth and who is risk averse, the partial risk aversion coefficient must be increasing with an increase in prospect size t.

The behavior of the relative risk aversion coefficient has long been controversial. Arrow has shown that for $U(W)$ to be bounded from below and above the relative risk aversion coefficient must be less than one at low wealth levels and greater than one at high wealth levels. If it were monotonic in between, it would have to rise from below one to above one.

One observation which would be difficult to reconcile with a stable utility function in terms of final wealth states would be if we found the relative risk aversion coefficient to drop very rapidly from high levels as prospect size rises by fairly small amounts relative to wealth (remember again that $R_W = R_M$). This would indicate that the risk aversion function would have to have a hump just below the current wealth level and in a fairly small neighborhood around current wealth. In other words,

the utility function would indeed have much stronger curvature around zero income (= current wealth) than away from it. If this were true for most individuals one would immediately ask how that strong curvature segment ended up in the neighborhood of current wealth and be led to a specification of a utility function in terms of current income or gains and losses.

In panels 4, 5 and 6 of Table 1 the approximate measures of partial, absolute and relative risk aversion are given for the alternatives. A utility function with constant partial risk aversion was used to approximate these measures ($U = (1-S)M^{1-S}$). Each indifference point between two alternatives, say A and B, defines an equation $EU(A) = EU(B)$ which can be used to measure S at that point. A and B can then be computed once the game levels and wealth levels are given.¹ The indifference points establish the endpoints for the interval within which A, S and R must lie for any of the alternatives. The geometric mean of the endpoints was assigned to the alternatives as the approximate measure for those individuals who chose it. Partial risk aversion varies from values ≥ 7.5 for extreme risk averters to values of less than 0 for the risk-preferrers. For any given choice alternative it is, of course, invariant to the scale of the game.

From Table 1 we see that, as the game level rises the risk aversion distributions shift to the right. This implies increasing partial risk aversion. At high game levels most individuals have partial risk aversion values in the neighborhood of one. At the 0.50 Rs level the value of absolute risk aversion also centers around one, but it falls very rapidly to values of around 0.0017 at the 500 Rs level with the possible maximum around .01 at that game level. Relative risk aversion, for an individual with the approximate modal value of wealth of Rs 10000

¹The measures were evaluated for certain incomes of 0.7, 7, 70 and 700 Rs respectively which is roughly the certainty equivalent of the alternatives at the various game levels.

starts in the neighborhood of 1000 to drop to roughly one or two at the 500 Rs level.

The extreme variations of absolute and relative risk aversion across game levels (and wealth levels for R) are caused by the fact that partial risk aversion is fairly stable across such levels. If S varies only between zero and 5, then A and R must vary much more since they are related to S as follows $A = S/M$; $R = S(1 + \frac{W}{M})$. Given the empirical finding of fair stability of S, R is a particularly difficult measure to compare across individuals since it depends both on W and the game level. And except for wealth levels close to zero, the typical behavior of our 118 individuals exhibits declining relative risk aversion: A choice of E at the 0.50 level implies a relative risk aversion of 226 whereas a choice of B at the 500 Rs level implies a value of 1.71. The individuals, in making choices at the 0.50 Rs level are putting extremely small proportions of their wealth at risk, yet still most of them are not risk-neutral. As the game rises to larger proportions in terms of their wealth, they cannot, with these alternatives, move sufficiently rapidly to less risky alternatives to lead to increasing relative risk aversion. Their behavior at low game levels is far too risk-averse to be consistent with Arrow's hypothesis.

The more rigorous test of asset integration is whether the absolute (and relative) risk aversion coefficients change at approximately the same rate with changes in income as with changes in wealth. In the regression analysis of absolute risk aversion¹ on personal characteristics (Binswanger 1978a) I found that wealth tends to have a slightly negative effect on risk aversion which, however, was not always statistically significant. It was also noted there that--at the 5 Rs level--a massive change in wealth was required to make a risk-averse person who chooses B to behave in a nearly risk-neutral fashion (Table 7 of Binswanger 1978).

¹The regressions were run on the natural log of partial risk aversion. But transforming to the log of absolute risk aversion just changes the intercept of these regressions.

This contrasts with the sharp reduction in absolute risk aversion shown in Table 1 of this paper. Consider the shifts more carefully:

	Geometric Average \bar{A}	$\ln \bar{A}$
Game No. 7, 5 Rs level	0.0662	-2.7149
Game No. 9, 50 Rs level	<u>0.0101</u>	<u>-4.4986</u>
Difference	$\Delta \bar{A} = - .0561$	$\Delta \ln \bar{A} = -1.7837$

The largest regression coefficient of $\ln A$ on wealth measured at the 5 Rs level (and at the 50 Rs level) was -0.00945 . Increasing it by twice its standard error brings it to a maximum estimate $b^* = 0.0181$. The increase in wealth required to induce the same change in absolute risk aversion than the shift from the 5 Rs to the 50 Rs game can be measured as

$$\Delta W^* = 1000 \times \Delta \ln \bar{A} / b^* = \text{Rs } 98527$$

(This is roughly 47% of the largest wealth observed in the sample). It compares with an increase in the certainty equivalent of income of between 45 and 95 Rs for the shift between the 5 Rs and the 50 Rs game levels.¹ The derivatives of absolute risk aversion with respect to wealth and income is clearly not of the same magnitude: even under the most favorable assumptions it takes an increase in wealth of roughly 1000 times the change in certain income to lead to an

¹The minimum certainty equivalents for an extreme risk averter is 5 Rs for the 5-Rs-game and 50 Rs for the 50-Rs-game. For a risk-neutral individual it is Rs 10 and 100. The smallest possible difference in M is thus 45 Rs while the largest one is 95 Rs.

equally large change in absolute risk aversion.

One objection to this test is that the coefficient of wealth is measured from the cross-sectional variation of absolute risk aversion across the households, while the difference caused by the change in game level is simply the difference in the geometric average of absolute risk aversion for the same individuals across the game scale. (Note that the difference across game scale is significant statistically at the 10% level, and also at the 1% level). Little can be done about this objection, unfortunately. A second objection is that certain wealth is relatively poorly measured and therefore its regression coefficient is not as reliably estimated than the difference of means across game scale. If W is measured with random error, its coefficient may be biased, but it is hard to imagine systematic errors in its measurement which could cause a bias of a factor of 1000. A third objection is that the utility function with constant partial risk aversion leads to poor approximation of the absolute risk aversion coefficients. Since we observe increasing partial risk aversion, a function with those characteristics should have been used. Such a function with increasing partial risk aversion (IPRAF) will be discussed below, but it has two parameters and needs equations from two indifferent points to be estimated. Which indifference points should be chosen is not clear. Absolute risk aversion for the IPRAF has been estimated for all combinations of indifference points which give a solution. The values found differ by less than 4% from those of the constant partial risk aversion function in all cases which did not include the indifference point between alternatives 0 and A. For the latter cases, absolute risk aversion is underestimated by between 10% and 20% with the CPRAF. But this is not sufficient to radically alter the regression results. In addition the experimental results tend to indicate that, at such high

risk aversion levels, partial risk aversion increases at a very slow rate.

The basic reason for rejecting asset integration (or stability over time of a utility function in forms of wealth) is that the observed behavior of individuals at low game levels is extremely cautious relative to their assets. Consider the individual choosing game B at the 5 Rs level with low and high outcome of 4 and 12 Rs. Alternative A on the other hand would give him 3 and 15 Rs. He is unwilling to risk a loss of Rs 1 with 50% probability to increase his expected income by Rs 1. If his net worth is at an average of 10000 Rs (close to the mode in the sample) then the loss with 50% is only 1/10000th of his wealth. Choosing the same alternative at the 500 Rs level implies a much higher risk relative to wealth. Stated otherwise, the curvature of the utility function is much larger at low levels of games than at high levels.

The rejection of asset integration implies that an individual's utility function is not stable relative to all wealth positions which he could achieve, but that it does adjust its "wiggles" to new wealth position when they are reached. It is a form of compartmentalization of the mind which appears not to integrate all uncertain income prospects to consider a single distribution of final wealth states. This opens the possibility that income prospects which accrue in different form may be evaluated differently depending on the form in which the income accrues. That, to some extent, limits our capacity to extrapolate the experimental findings to other situations. In particular, since the game was played only with positive payoffs, we cannot infer the shape of the utility functions for losses, which we could have if asset integration had been accepted. As it stands we have little information about the shape of a utility curve for incomes below zero.

On the other hand, lack of asset integration should not severely restrict the ability to derive comparative static results on what happens to portfolio choice and other choices when wealth changes. The experimental results give estimates of partial risk aversion of different game levels and indicate how it changes as wealth rises. The numerical bounds on partial risk aversion should in fact lead to sharper comparative static predictions.

In empirical investigations it has always been difficult to work with utility functions in the form of wealth, in part because it requires that one measure certain wealth. This may not be too difficult in portfolio analysis problems where extremely well developed markets provide portfolio valuations every day. But for problems in agriculture, imperfect land markets make estimation of certain wealth very difficult and the same applies to investment decisions in human capital.

Functional Forms for Utility Functions for Gains

This section explores what functional form for the utility function could be consistent with the experimental evidence. The power function has constant partial risk aversion (constant relative risk aversion if specified in terms of wealth) and it can be written as

$$U = (1 - S)M^{(1 - S)} \quad (18)$$

with S = partial risk aversion. Such a function would fit fairly well for those

individuals who were choosing alternatives B or C at low game levels and continued to do so throughout the sequence. The function has no upper asymptote, and that seems to be necessary to get increasing partial risk aversion. The following function may be called increasing partial risk aversion function (IPRAF) and has an upper asymptote.

$$U = 1 - e^{-aM^b} \quad \begin{matrix} 0 \leq b \leq 1 \\ 0 \leq M < \infty \end{matrix} \quad (19)$$

The limiting case for $b = 1$ is the negative exponential function $U = 1 - e^{-aM}$ which has constant absolute risk aversion $A = a$. As can be verified easily the IPRAF has $U_M > 0$, $U_{MM} < 0$ and $U_{MMM} > 0$. Partial risk aversion is equal to

$$S = abM^b + 1 - b \quad (20)$$

For $M = 0$ partial risk aversion is $(1-b)$, a value of less than one. The parameter b thus determines initial partial risk aversion, while, for given b , the parameter a determines how fast it will rise with income. The elasticity of partial risk aversion with respect to income is

$$E_{SM} = \frac{b}{1 + \frac{1-b}{ab^2M^b}} \quad (21)$$

i.e. it is zero for zero income and reaches a value of b asymptotically. This is somewhat inconsistent with the experimental evidence where we find that the proportionate increase in partial risk aversion is about the same as we move from one game level to the next for all game levels.

	Geometric Average of S	Proportional increase in S
0.50 Rs level	0.279	
5 Rs level	0.4635	+ 66%
50 Rs level	.7046	+ 52%
500 Rs level	1.0896	+ 54%

But this is based on averages, not on individual behavior, which differs quite markedly. Those whose low level choices were B or C have very slowly rising partial risk aversion while it rises much faster for those who were risk neutral initially. And there appears to be a barrier on risk aversion at the upper end.

One can use two indifference points at two game levels to define two equations which can be solved for a and b. For example, if an individual is indifferent between C and E at the 5 Rs level, this defines an equation

$$e^{-a3^b} + e^{-a15^b} = e^{-a^b} + e^{-a19^b} \quad (22)$$

and the indifference point between B and C at the 500 Rs level implies a similar equation. These equations cannot be solved analytically for a and b, but one can iteratively approximate a solution. Table 3 gives the solution for all indifference patterns between the 5 and 500 Rs levels for which solution to the equation pair exist.¹ Utility functions so estimated all imply

¹Solutions were also derived for indifference points at the 5 Rs and 50 Rs level and for indifference points at the 50 and 500 Rs level respectively. However, in all these cases partial risk aversion rises far more rapidly than implied in the experimental results. However, in the cases just mentioned it becomes clear that the IPRAF puts constraints on how fast risk aversion can increase. Since $b < 1$, the elasticity of S with respect to M cannot exceed 1, which means, for example, that one cannot be indifferent between C and E at the 5 Rs level and A and B at the 50 Rs level.

Table 3. Examples of IPRAF and the Implied Choice Patterns

	(1) ^a	(2)	(3)	(4)	(5)	(6)	(7)
COEFFICIENTS OF IPRAF							
Coefficient a	.9970	.7032	.7326	.8548	.3829	.6385	.4095
Coefficient b	.0004634	.007174	.01675	.03961	.2417	.2008	1.282

CHOICE PATTERNS IMPLIED AT DIFFERENT LEVELS
(TWO LETTERS UNDERLINED MEANS INDIFFERENCE)

LEVEL

0.50 Rs	EF	E	E	E	C	C	B
5 Rs	<u>EF</u>	<u>CE</u>	<u>CE</u>	<u>CE</u>	<u>BC</u>	<u>BC</u>	<u>AB</u>
50 Rs	E	C	C	B	B	A	A
500 Rs	<u>CE</u>	<u>BC</u>	<u>AB</u>	<u>OA</u>	<u>AB</u>	<u>OA</u>	<u>OA</u>
5000 Rs	B	A	A	O	A	O	O
50000 Rs	O	O	O	O	A	O	O

PARTIAL RISK AVERSION AT DIFFERENT LEVELS

CERTAIN INCOME

0.7	.001	.30	.28	.17	.70	.46	1.04
7	.004	.32	.32	.32	.81	.81	1.75
70	.033	.40	.54	1.4	1.1	2.3	3.6
700	.321	.81	1.76	9.3	1.76	8.8	8.3
7000	32	2.8	8.31	66	3.4	37	20

^aIn this case the payoffs of alternative F had to be changed from (0,20) to (0.01, 20) to derive a solution.

extreme or severe risk aversion (choices 0 or A) at high payoff levels, because they reach these ceiling values (where marginal utility becomes zero) fairly rapidly. We saw on page 51 that in the experiment partial risk aversion rises by about 135% between the 5 Rs and the 500 Rs level for the average of the sample. An approximately equal rise is implied in functions (2) and (5) of Table 3. However it is an open question how good the approximation is outside of the range of the game payoffs. Should we really believe that when offered a game at the 50000 Rs level, almost all individuals would choose the riskless alternative zero?

The IPRAF thus has two limitations: It is not defined for values of M less than zero and it may hit its ceiling value too rapidly if partial risk aversion increases rapidly at low levels of income and then less rapidly. For empirical applications one may eventually have to work with a utility function which has its negative segment, its low income segment and its high income segment approximated by different functional forms. As it stands we have little information about the negative segment or the very high income segment and must await more experimental work.

SUMMING UP

Experimental methods have been largely neglected in economics as a means for testing hypotheses about predictive power of various models. Psychologists have used them extensively, but usually with very small payoffs to the individuals involved and/or with very small sample sizes. Yet in this experiment it becomes clear that behavior with trivial payoffs is not at all the same as behavior at substantial payoffs, where behavior appears more predictable and regular. The direct costs of the experiment

reported here was roughly \$2500 in prize money and an equal amount in research assistance, travel and computer costs.¹ If the study had been carried out in the U.S. with payoffs roughly equal to unskilled wage rates, the direct experimental cost would probably have been \$150,000 for prizes and possibly \$50,000 for research assistance. While this is much more expensive, many research projects have budgets which far exceed this amount.

For the households studied the experiment allows us to conclude that, at substantial payoffs, almost all individuals are risk averse, but that very few are severely or extremely risk averse. In fact, risk aversion differs far less across individuals than one would have expected. Furthermore, partial risk aversion is clearly rising.

The experimental results are inconsistent with all but one of the security based models of behavior, the safety fixed model with income integration. But this model is not rejected only because it is unable to predict how rich or poor people should behave when confronted with a set of uncertain prospects as the one of the experiment.

In particular the results are inconsistent with all models which assume that behavior is strongly influenced by a goal of reaching a fixed subsistence income. If the behavioral importance of a subsistence goal cannot be shown to be important for the households studied, which belong to the poorest of the world, it is hard to imagine where subsistence based models could be important.

On the other hand, the results are not inconsistent with some of the utility maximizing models. Among the latter models, only those which employ a linear utility function can be rejected.

¹Salary costs of my assignment in India have not been counted.

Yet we do find evidence for something like bounded rationality. Individuals do not seem to base their decisions over new income opportunities on the basis of final wealth states evaluated with a time-stable utility function over such states. Such a utility function would imply a global rationality since it would evaluate all incomes, regardless of their form, in the same way and enable present decisions to be fully consistent with past decisions and future decisions. Instead, what we observe is a utility function which appears to adjust to new wealth positions and has richer people behave in much the same way as poorer ones as soon as trivial game levels are exceeded. Unfortunately, this finding implies that a game with gains only is inadequate to measure the loss branch of a utility function. New experiments will be needed to do that. Another problem which needs new experimental evidence is the issue of probability preferences or, more generally, of the formation of subjective probabilities.

- Anderson, Jock R., "Modelling Decision Making under Risk," Agricultural Development Council Conference on Risk and Uncertainty in Agricultural Development, CIMMYT, Mexico, March 1976.
- Anderson, Jock, R., John L. Dillon and Brian Hardaker, Agricultural Decision Analysis (Ames: Iowa State University Press, 1977).
- Arrow, Kenneth J., Essays in the Theory of Risk Bearing, (Amsterdam, North Holland Publishing Co., 1971).
- Bardhan, Pranab K. "Wages and Unemployment in a Poor Agrarian Economy: A Theoretical and Empirical Analysis," University of California, Berkeley, Dept. of Economics, December 1977.
- Bernoulli, Daniel, "Exposition of a New Theory on the Measurement of Risk," English Translation by Louise Sommer of the original 1738 paper, Econometrica Vol. 22, (January 1954) 23-26.
- Binswanger, Hans P. "Attitudes Towards Risk: Experimental Measurement in Rural India," Yale University, Economic Growth Center, mimeo, 1978a
- Binswanger, Hans P. "Risk-Attitudes of Rural Household in Semi-Arid India," Yale University, Economic Growth Center, Discussion Paper 275, 1978b.
- Diamond, Peter A. and Joseph E. Stiglitz, "Increases in Risk and Risk Aversion," Journal of Economic Theory, Vol. 8, (July 1974): 337-360.
- Edwards, Ward, "Probability-Preferences in Gambling," American Journal of Psychology, Vol. 66, (July 1953): 349-364.
- _____, "The Reliability of Probability-Preferences," American Journal of Psychology, Vol. 67, (March 1954a): 68-95.
- _____, "The Theory of Decision Making," Psychological Bulletin, Vol. 51 (July 1954b): 380-410.
- _____, "The Prediction of Decisions among Bets," Journal of Experimental Psychology, Vol. 50 (September 1955): 201-214.
- _____, "Subjective Probabilities Inferred from Decisions," Psychological Review, Vol. 69 (March 1962): 109-135.
- Fishburn, Peter C. "On Handa's 'New Theory of Cardinal Utility' and the Maximization of Expected Returns," Journal of Political Economy, Vol. 86 (April 1978): 321-324.
- Friedmann, Milton and Leonard J. Savage, "The Utility Analysis of Choices Involving Risk," Journal of Political Economy, Vol. 56 (August 1948): 279-304.
- Grether, David M. and Charles R. Plott, "Economic Theory of Choice and the Preference Reversal Phenomenon," American Economic Review, forthcoming.

- Griffith, R. M. "Odds Adjustment by American Horse Race Betters," American Journal of Psychology, Vol. 62 (July 1964): 387-98.
- Handa, Jagdish, "Risk, Probability and a New Theory of Cardinal Utility," Journal of Political Economy, Vol. 85, (February 1977): 97-122.
- Halter, Albert, N. and Gerald W. Dean, Decisions under Uncertainty with Research Applications. (Cincinnati: South-Western Publishing Co., 1971).
- Hildreth, Clifford, "Ventures, Bets and Initial Prospects", in Balch, Michael, Daniel McFadden and S. Wu, eds. Decision Rules and Uncertainty, (Amsterdam, North Holland Publishing Co., 1974.)
- Kahnemann, Daniel and Amos Tversky, Prospect Theory: An Analysis of Decisions Under Risk, Jerusalem, The Hebrew University, mimeo, 1977.
- Luce, Duncan R. and Patrick Suppes, "Preference, Utility and Subjective Probability," in Luce, R. Duncan, Robert R. Bush and Eugene Galanter, Handbook of Mathematical Psychology, (New York, John Wiley and Sons, Inc., 1965).
- Markovitz, Harry, "The Utility of Wealth," Journal of Political Economy, Vol. 60, (April 1952): 151-158.
- Menezes, C. F. and D. L. Hanson, "On the Theory of Risk Aversion," International Economic Review, Vol. 11, (October 1970): 181-487.
- Mosteller, F. and Nogee, P., "An Experimental Measurement of Utility," Journal of Political Economy, Vol. 59 (October 1951): 371-404.
- Von Neumann, James and Oskar Morgenstern, The Theory of Games and Economic Behavior, (Princeton, N.J., Princeton University Press, 1947)
- Pratt, John W., "Risk Aversion in the Small and in the Large," Econometrica, Vol. 32, (January-April 1964): 122-136.
- Ramsey, Frank P., "Truth and Probability," 1926, in Ramsey, F. P. The Foundations of Mathematics and other Logical Essays (New York, Harcourt, Brace and Co., 1931).
- Roumasset, James, Risk and Choice of Technique for Peasant Agriculture: The Case of the Philippine Rice Farmers, Ph.D. Dissertation, University of Wisconsin, Dept. of Economics, 1973.
- Rothschild, Michael and Joseph E. Stiglitz, "Increasing Risk I: A Definition," Journal of Economic Theory, Vol. 2 (September 1970) 225-243.
- Rothschild, Michael and Joseph E. Stiglitz, "Increasing Risk II: Its Economic Consequences," Journal of Economic Theory, Vol. 3 (March 1971): 66-84.

Savage, Leonard J., Foundations of Statistics, (New York, John Wiley and Son, 1954).

Savage, Leonard J., "Elicitation of Personal Probabilities and Expectations," Journal of the American Statistical Association, Theory and Methods Section, Vol. 66, (December 1971): 783-801.

Sprowls, R. C., "Psychological-Mathematical Probability in Relationships of Lottery Gamblers," American Journal of Psychology, Vol. 66 (January 1953): 125-30.

Tversky, Amos, "Additivity, Utility and Subjective Probability," Journal of Mathematical Psychology, Vol. 4 (June 1967): 175-201.

Tversky, Amos, "Intransitivity of Preferences," Psychological Review, Vol. 76, (Jan. 1969): 31-48.

Zeckhauser, Richard and Emmett Keeler, "Another Type of Risk Aversion," Econometrica, Vol. 38 (September 1970): 661-665.