Appendix 2: Analytical Tools for Assessing the Effectiveness of Risk Management Strategies

Different modeling approaches are used by economists to capture decisionmaking in risky situations. These approaches are based on the idea that each risky strategy offers farmers a different probability distribution of income, and that determining the best strategy involves describing the different distributions and developing rules to choose among them. These approaches differ, however, in the ways in which they incorporate risk attitudes, and in the degree of flexibility allowed in specifying risk-return trade-offs.

The following sections examine selected approaches that are commonly used by economists in analyzing decisionmaking under risk. The approaches discussed here range from one of the simplest (the "safety-first" approach) to one of the more complex (the use of "expected utility"). "Certainty equivalence" also is discussed, which involves measuring risk in terms of differences in expected income. The use of these different approaches allows researchers to rank alternative strategies, and to help producers make optimal choices in different situations.

The "Safety-First" Approach

The "safety-first" approach to risk management applies if a decisionmaker first satisfies a preference for safety (such as minimizing the probability of bankruptcy) when making choices as to the firm's activities. Only when the safety first goal is met at a threshold level can other goals (such as maximizing expected returns) be addressed. Thus, attaining the highest-priority goal serves as a constraint on goals that have successively lower priorities (Robison, Barry, Kliebenstein, and Patrick). Safety-first methods are particularly applicable where survival of an individual or business is the paramount concern. However, in most business risk management situations, the use of safety-first methods is somewhat arbitrary because no single goal is clearly dominant.

The safety-first criteria can be specified in various ways in empirical applications. One of the first uses of this approach was developed in 1952, and involves choosing the set of activities with the smallest probability of yielding an expected return (Y) below a specified disaster level of return (Y-min) (Roy). To aid in understanding the various safety-first criteria, appendix table 2 shows the expected income and the probability of income less than the disaster level, which is assumed to be $50,000, for three hypothetical strategies, A, B, and C. For strategy B, for example, the expected return is $500,000, and the probability that returns under this strategy will fall below $50,000 is 4 percent. Strategy A has the highest expected return and the highest risk among the strategies illustrated, while strategy C has the lowest expected return and lowest risk.

Appendix table 2—Comparison of Roy's safety-first approach with Telser's safety-first approach

Strategy	Expected income E(Y)	Minimum disaster return (Y-min)	Probability of falling below minimum disaster return (P)
	Dollars		*Percent*
A	1,000,000	50,000	5
B	500,000	50,000	4
C	200,000	50,000	3

Source: Hypothetical example developed by ERS.

Under Roy's safety-first criteria, the optimal activity choice occurs where the probability of expected return falling below the $50,000 threshold is minimized. Strategy C, which has the lowest probability of disaster, best meets this criteria. When returns are normally distributed, the solution occurs where the disaster level (Y-min) is the greatest number of standard deviations away from the expected income. Roy's criteria can be expressed mathematically as:

$$\text{Minimize Prob } (Y < Y\text{-min})$$

A second type of approach, introduced by Telser in 1955, assumes that the decisionmaker maximizes expected returns, E(Y), subject to the constraint that the probability of a return less than or equal to a specified minimum disaster level (Y-min) does not exceed a given probability (P). Mathematically, Telser's approach is expressed as:

$$\text{Maximize } E(Y)$$
$$\text{subject to: Prob } (Y < Y\text{-min}) \leq P$$

To apply the Telser criterion to the example in appendix table 2, suppose that the critical probability is 4 percent. Then, the Telser criterion would choose strategy B, which maximizes expected income among those strategies for which P is not greater than 4 percent. Alternatively, if the critical probability were 3 percent, then strategy C would be selected, while if it were 5 percent, strategy A would be selected. This example illustrates that safety-first results can be quite sensitive to initial assumptions about what constitutes a critical loss.

The topics that have been addressed by these various types of safety-first criteria vary widely. They include: optimal hedging (Telser), dynamic cropping decisions in southeastern Washington (Van Kooten, Young, and Krautkraemer), farm extension programs (Musser, Ohannesian, and Benson), and attitudes toward risk regarding fertilizer applications among peasants in Mexico (Moscardi and de Janvry).

The safety-first approach has both advantages and disadvantages. It does not require the specification of a farmer's risk aversion coefficient (see accompanying box on risk aversion, p. 119), and it is not limited to specific distributional assumptions, other than that utility increases with returns (subject to varying constraints depending on the specification) (appendix table 3). As a result, it is straightforward to use. On the downside, however, it is limited in its ability to address producers' varying levels of aversion to risk (although the threshold probability can serve as a proxy for risk aversion), and difficulties can also arise in choosing the critical cutoff level for disaster returns. In addition, any outcome below the cutoff is treated as equivalent to any other. In reality, observations far below the cutoff disaster level are more adverse to the farmer than those that are nearer the cutoff point.

The "E-V" Approach and Quadratic Programming

A classic problem in risk analysis involves determining an optimal allocation of resources across an array of risky alternatives. The problem was first solved by Markowitz in the context of selecting optimal stock portfolios. His solution was to find the set of allocations that maximize expected total return for different levels of variance of total return. This is called the "expected value-variance (or E-V) efficient" set or frontier. The heavy line in appendix figure 5 represents an E-V efficient frontier. On this frontier, expected return can be increased only by accepting a larger variance of return. The optimal portfolio is presumed to come from this frontier and depends on the decision-maker's preference tradeoffs between expected return and variance of return.[33]

[33]The E-V efficient set will include the portfolio that maximizes expected utility if the decisionmaker's utility function is quadratic or returns on all activities are normally distributed. See subsequent section on "The Expected Utility Approach."

Appendix table 3—Methods for ranking probability distributions, and assumptions and parameters required for each

Method	Assumptions required about utility	Parameters required to use method
Roy's safety-first	Increases with prob (Y > L)	Prob (Y < L)
Telser's safety-first	Increases linearly with Y if prob (Y < L) < P, is zero otherwise	E(Y) and prob (Y < L)
E-V efficiency	Increases with Y at a decreasing rate plus normality or quadratic utility	E(Y) and Var(Y)
1st degree stochastic dominance	Increases with Y	Complete distribution of Y
2nd degree stochastic dominance	Increases with Y at a decreasing rate	Complete distribution of Y
Expected utility and certainty equivalents	Known function of Y	Complete distribution of Y and utility function

Note: Y = income; E(Y) = expected income; L = critical level of income; P = critical probability.

Source: Compiled by ERS.

Appendix figure 5
Example E-V efficient frontier and indifference curves

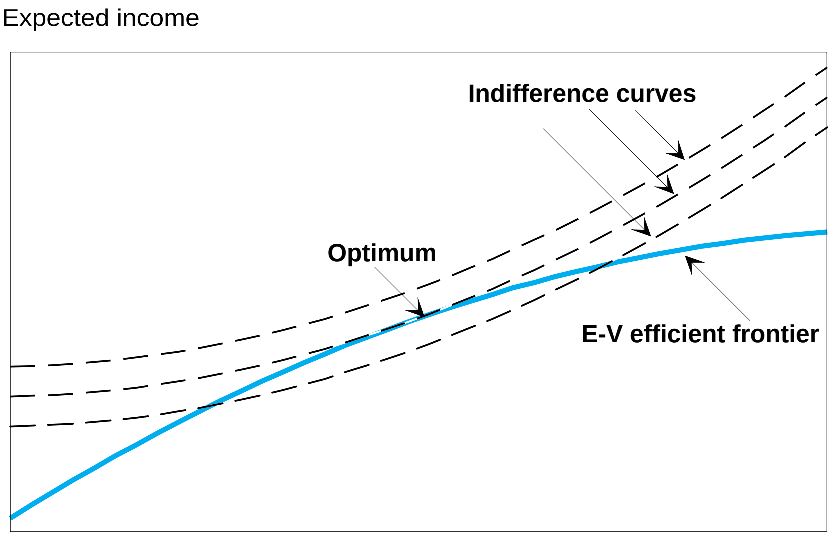

Source: Hypothetical example developed by ERS.

The E-V efficiency criterion can be used in allocating a farm's resources among alternative risky enterprises. A risk-averse farmer desires high expected return and low variance of return, which involves moving upward and/or to the left in the figure. The optimal combination of activities for the farmer occurs at the point on the E-V frontier that provides the preferred combination of expected return and variance of return. To illustrate the farmer's preferences, three indifference curves are shown as dashed lines. Each connects combinations of risk and expected return that are equally desirable to the producer. The optimal point on the E-V frontier is the point that touches the highest attainable indifference curve.

This approach has on many occasions been applied to farming decisions, particularly to decisions about enterprise choice and diversification. E-V efficient combinations of crop and livestock enterprises can be identified and the combination that offers the preferred mix of expected return and variability of returns can be chosen. Determining E-V efficient combinations requires estimates of the variances in returns and the correlations of returns for those enterprises under consideration, as well as estimates of expected returns for those enterprises.

The attractiveness of E-V analysis is that it leads to relatively convenient solutions using quadratic programming. The exact formulation of the problem can vary. One approach is to maximize a quadratic function of activity levels subject to linear constraints as follows:

$$\text{Maximize} \sum_{i=1}^{n} x_i u_i - \lambda \sum_{i=1}^{n} \sum_{j=1}^{n} x_i x_j \sigma_{ij}$$

subject to:

$$\sum_{i=1}^{n} a_{ji} x_i \leq b_j \quad j = 1,2,\ldots m$$

$$x_i \geq 0, \quad i = 1,2,\ldots n$$

where:

x_i = the level of the ith activity;

u_i = the expected return per unit of the ith activity;

λ= the risk-return tradeoff;

σ_{ij} = the covariance of return on activities i and j;

a_{ji} = coefficients in m linear constraints on the activity levels.

b_j = levels of the linear constraints.

To trace out the E-V frontier, the quadratic programming problem must be solved parametrically as the risk aversion coefficient, λ, varies from 0 to ∞. This is a rather complicated problem, but computer algorithms are available. If the farmer is risk neutral (λ=0), the problem collapses to a an income maximization problem, which can be solved with ordinary linear programming. As risk becomes increasingly important, the risk aversion coefficient increases and the E-V portfolio becomes increasingly diversified (Anderson, Dillon, and Hardaker). Other factors, such as limitations imposed by resource constraints, may also lead to a diversified portfolio.

The quadratic programming approach has been applied to farm enterprise selection by many researchers. The first application, for example, involved the evaluation of four production activities and several resource constraints on a representative farm in eastern North Carolina (Freund). Quadratic risk programming has since been applied to many other evaluations of optimal farm enterprise choice, including studies by Barry and Willman, and Musser and Stamoulis.[34]

The use of the E-V approach has both advantages and disadvantages. As in the case of "safety-first" analysis, E-V analyses may or may not include an explicit measure of the producer's risk aversion (as illustrated in the example above). E-V analysis is limited, however, in that it assumes that the producer has an outcome distribution that is normal[35] or, alternatively, a utility function (which expresses risk preferences) that is quadratic. In addition, the farmer is assumed to always prefer more (rather than less) of the variable in question (such as income), and is assumed universally not risk preferring with respect to that variable (Hardaker, Huirne, and Anderson). As with various approaches to risk analysis, estimation of the variance-covariance matrix can present methodological pitfalls (Mapp and Helmers).

The Stochastic Dominance Approach

Unlike E-V analysis, which is based on the mean and variance of a distribution, stochastic dominance involves comparing points on two or more entire distributions. That is, when stochastic dominance is used, alternatives are compared in terms of the full distributions of outcomes. Because comparisons must be made at each specified point along each distribution in a pairwise fashion, the conceptual complexity and computational task associated with this approach are greater than when E-V analysis is used (Hardaker, Huirne, and Anderson).

The first concept of stochastic efficiency was formalized in the early 1960's, and is known as "first-degree" stochastic dominance. This approach rests on the notion that decisionmakers prefer more of a given variable (such as income) to less. Using an example, suppose there are three plans, A, B, and C, each having a probability distribution of income outcomes, "x." The cumulative density functions (CDFs) associated with the plans are $F_A(x)$, $F_B(x)$, and $F_C(x)$, respectively, as shown in appendix figure 6. The CDFs reflect the "accumulated" area under the probability density function (PDF) at each level of income for each plan. At the extreme right side of the chart, the entire PDF is summed, and the probability of realizing an income for any of the plans that is equal to or less than the amount designated on the axis is 1.00.

For one plan to dominate another in the first-degree sense, the CDF for the first plan must nowhere be higher than the CDF for the second plan, and it must lie below the CDF for the second plan at some point. Mathematically, first degree stochastic dominance (FSD) can be expressed, for two representative plans A and B, as:

$F_A(x) \leq F_B(x)$, for all levels of x

[34]Other risk programming methods are available. MOTAD programming, for example, minimizes the mean absolute deviation in net income, which simplifies the problem to one of linear programming (Hazell). Target MOTAD, developed in 1983, minimizes deviations from a target level of income (Tauer). Discrete stochastic programming can also be used to determine efficient enterprise choices for farmers. For a discussion of the advantages and disadvantages of these approaches, as well as examples, see Hardaker, Huirne, and Anderson; Musser, Mapp, and Barry; Mapp, Hardin, Walker, and Persaud; and Walker and Helmers.

[35]The normality assumption may be reasonable, particularly if the number of risky prospects is not too small and the risky prospects are diverse (Anderson, Dillon, and Hardaker). In addition, several studies have concluded that the E-V approach is quite robust to violations of the normality assumption (Levy and Markowitz; Kroll, Levy, and Markowitz).

In appendix figure 6, F_C (x) dominates F_A (x) in the first-degree sense, meaning that the probability of exceeding any given level of income is greater under plan C than plan A, and that plan A cannot be a member of the first-degree stochastic dominant (FSD) set. Because the CDF for plan B crosses both plans A and C, plan B does not dominate, and is not dominated by, either A or C in the first degree sense (Hardaker, Huirne, and Anderson).

Second-degree stochastic dominance (SSD) is applicable if the decisionmaker is risk averse and prefers higher incomes to lower incomes. In contrast to first-degree stochastic dominance, SSD involves the comparison of areas under the CDFs for various plans, and, in general, has more discriminatory power than does FSD (King and Robison). For a representative plan, A, to be SSD over another plan, B, requires that:

$$\int_0^\infty F_A(x) \le \int_0^\infty F_B(x), \text{ for all values of x.}$$

In appendix figure 6, for example, note that the accumulated area under F_B(x) is less than under F_A(x) at all levels of income. Thus, B exhibits SSD over A. C also exhibits SSD over A. However, C does not exhibit SSD over B because there is a range in the lower tail where the accumulated area under C exceeds that under B. This example illustrates the stronger discriminatory power of SSD compared to FSD while showing that SSD does not discriminate among all distributions.

As with other approaches to risk analysis, the use of stochastic dominance methods has both pros and cons. Although it provides a rigorous assessment, the number of efficient sets may remain unduly large. SSD is more discriminating than FSD, but nearly one-half of randomly generated farm plans in one study, for example, were found to be within the SSD set (King and Robison). The assumption of risk aversion required by SSD may not always hold, and the pairwise comparisons that are necessary in determining the efficient set can be computationally burdensome.

Appendix figure 6
Cumulative income distributions under three alternative strategies

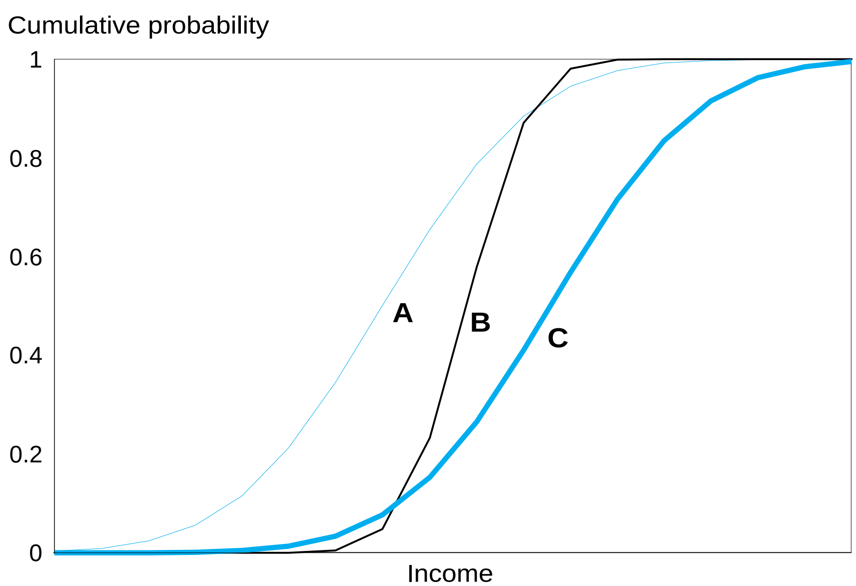

Source: Hypothetical example developed by ERS.

Different Approaches Can Be Used to Estimate Risk Aversion

To determine a farmer's best risk management strategy, information is needed about his or her risk preferences among the different income distributions generated by those alternative strategies. Individuals who accept a lower average return to reduce the variability of returns are said to be risk averse. Many individuals are believed to be risk averse, as evidenced by the widespread demand for automobile, property, and health insurance. Premium costs for these products generally exceed expected indemnities due to administrative costs, but buyers often find the price acceptable to mitigate potentially disastrous outcomes.

While risk aversion is acknowledged as widespread, the degree of risk aversion varies among individuals and is difficult to ascertain. Two general approaches typically have been used in empirical analyses. The first approach measures risk aversion directly by confronting the decisionmaker with a choice (either actual or hypothetical) among several alternatives, at least some of which involve risk (Newbery and Stiglitz). Such approaches have been used to determine risk aversion among farmers in northeast Brazil (Dillon and Scandizzo) and rural India (Binswanger), among others.

Measuring risk preferences directly can, however, lead to unstable results. Interview methods, in particular, are faced with the inevitable problem that individuals may not be able to reveal their attitudes toward decisions they have never taken or seriously contemplated (Binswanger). In addition, such studies have typically focused on a small-scale basis. Recently, work has been undertaken to measure farmers' risk attitudes on a large-scale basis, using rating scales of risk management questions to ascertain farmers' risk preferences (Bard and Barry).

The second approach does not involve interviews with decisionmakers or experimental determination of attitudes toward risk. Rather, this category involves: 1) focusing on testing hypotheses econometrically regarding risk preference structure; or 2) directly estimating utility functional forms or risk aversion coefficients using data on actual firm choices (Saha, Shumway, and Talpaz; Antle; Love and Buccola). Several studies have used this second category, with estimates of relative risk aversion ranging widely.

More technically, measures of either "absolute" or "relative" risk aversion can be used to quantify an individual's attitude toward risk. Both measure the curvature of the utility function, and represent the degree to which the satisfaction obtained from an additional unit of income declines as income increases. The units of measurement must be considered in interpreting estimates of absolute risk aversion, whereas relative risk aversion is unit free.

Varying estimates of relative and absolute risk aversion result from different approaches and data sets. Researchers generally agree that a reasonable relative risk aversion coefficient, for example, is in the neighborhood of 2.0, or "rather risk averse." Relatively risk-averse farmers would be likely to maintain substantial financial reserves as protection against income shortfalls, while those who are less risk averse would be inclined to borrow to near their limit in order to increase their expected incomes (Hardaker, Huirne, and Anderson).

Further, once the efficient set of plans is determined, identification of the optimal choice within this set depends on knowing more about a decisionmaker's preference than merely that an unquantified aversion to risk exists (Anderson, Dillon, and Hardaker).

The Expected Utility Approach

If a decisionmaker's risk preferences can be described mathematically and the probability distributions associated with each risky alternative are known, his or her choice among the risky alternatives can be optimized directly. Expected utility provides a convenient way to represent risk preferences. The basic idea is that decisionmakers maximize expected utility, where utility is an indicator of satisfaction measured in arbitrary units. Utility increases less than proportionately with income for decisionmakers who are risk averse. In other words, utility is an increasing, but downward bending, function of income for such persons (Anderson, Dillon, and Hardaker; Robison and Barry; Laffont; Takayama).

Many different specifications can be used to capture the curvature of the utility function, and each represents the degree to which the satisfaction obtained from an additional unit of income changes as income increases. One such utility function specification can be expressed as:

$$U = 100 - (1{,}000{,}000/Y)$$

As can be seen from this equation, an increase in Y, say from $20,000 to $30,000 (50 percent), results in an increase in utility from 50 to 67 (32 percent). This utility function exhibits constant relative risk aversion, which means that the degree of risk aversion decreases with income. The coefficient of relative risk aversion in this example is 2, which is considered by many economists to be about average.

To illustrate the use of expected utility, suppose that Farmer Smith, whose utility function is specified as above, is choosing between two strategies: (1) continuing to farm, and (2) taking a job in town. Under the first strategy, Smith has an 80-percent chance of a net income of $70,000 and a 20-percent chance of a net income of $20,000. Working in town provides a sure income of $55,000. The expected income and standard deviation in income for each strategy are shown in appendix table 4. Neither strategy dominates the other from the standpoint of E-V efficiency because the first strategy has the highest expected income, while the second has the lowest standard deviation in income (zero in this case).

Using Smith's utility function, the resulting utilities for each level of income shown in the table are:

Income	*Utility*
$20,000	50.0
$55,000	81.8
$70,000	85.7

Appendix table 4—Expected income and standard deviation of income under two strategies available to Farmer Smith

Strategy	Probabilities of expected incomes: $20,000	$55,000	$70,000	Expected income	Standard deviation of income
	Percent			*Dollars*	
1	0.2	0.0	0.8	60,000	20,000
2	0.0	1.0	0.0	55,000	0

Source: Hypothetical example developed by ERS.

Certainty Equivalence Allows Estimation of Risk in Dollars of Expected Income

Decisionmakers often would like the losses from risk or the gains from risk reduction to be measured in terms of dollars of expected income. A sure outcome that an individual finds equally desirable to a given risky prospect is called a certainty equivalent outcome (Anderson, Dillon, and Hardaker; Laffont). Knowing certainty equivalent outcomes allows one not only to rank risky alternatives, but also to estimate the cost of risk, or the premium that the individual would pay to avoid the risk. Certainty equivalence simultaneously accounts for the probabilities of the risky prospects and the preferences for the consequences (Anderson, Dillon, and Hardaker). Because decisionmakers seldom have similar attitudes toward risk, certainty equivalents vary among individuals, even for the same risky prospect (Hardaker, Huirne, and Anderson).

Certainty equivalents can be calculated when E-V analysis is used or when the utility function is known and expected utility analysis is used. In the latter case, a certainty equivalent can be calculated by first calculating expected utility and then finding the sure outcome that would provide equal utility. This involves applying the inverse of the utility function to expected utility. For the utility function above this gives:

$$CE(Y) = 1{,}000{,}000 / [100 - E(U)]$$

Applying this inverse function to the expected utilities calculated in the accompanying table gives the estimates shown in the accompanying table. These results, shown in the last column, indicate that the first strategy yields a certainty equivalent income of $46,667, compared to an expected income of $60,000. In other words, the cost of uncertainty in farming for Smith is $60,000 - $46,667 = $13,333. Thus, Smith prefers the second strategy, which gives a certainty equivalent income of $55,000, and is $8,333 higher than obtained if he had chosen the first strategy.

Certainty equivalent estimates must be used with care, primarily because they tend to convey an unwarranted sense of precision. They must be taken as rather rough approximations, and depend heavily on how accurately the underlying utility functions and probability distributions are estimated. Utility functions, in particular, differ markedly among individuals, are not directly observable, and are likely estimated with substantial errors in most cases.

Certainty equivalent income under two strategies for farmer Smith

Strategy	Expected utility	Certainty equivalent income
		Dollars
1	78.6	46,667
2	81.8	55,000

Source: Hypothetical example developed by ERS.

Appendix 2: Analytical Tools for Assessing the Effectiveness of Risk Management Strategies

Expected utility under each strategy is calculated by weighting each utility level by its probability. The results for the two strategies, shown in appendix table 5, indicate that the second strategy yields a slightly higher expected utility and is therefore preferred. Expected utility results can also be used to estimate certainty equivalents (see accompanying box, p. 121).

The use of expected utility has both advantages and disadvantages. One advantage is that this approach is quite generalizable, allowing a wide choice of utility functions and probability distributions. Unlike stochastic dominance and E-V efficiency criteria, which typically leave some alternatives unranked, expected utility generally ranks all alternatives. The major drawback of the approach is that utility functions are difficult to estimate and known only approximately, at best. Moreover, it assumes that decisionmakers exhibit a high level of rationality, which does not always seem to be the case.

Appendix table 5—Example calculation of expected utility under two strategies for farmer Smith

	Probabilities of expected utilities			Expected
Strategy	50.0	81.8	85.7	utility
		Probability		
1	0.2	0.0	0.8	78.6
2	0	1.0	0	81.8

Source: Hypothetical example developed by ERS.