Decision-Making in the Presence of Risk

MARK J. MACHINA

Proposed in the 18th century by Cramer and Bernoulli and formally axiomatized in the 20th century by von Neumann and Morgenstern and others, the expected utility model has long been the dominant framework for the analysis of decision-making under risk. A growing body of experimental evidence, however, indicates that individuals systematically violate the key behavioral assumption of this model, the so-called independence axiom. This has led to the development and analysis of non-expected utility models of decision-making. Although recent work in this area has shown that the analytical results of expected utility theory are more robust than previously supposed, other important issues remain unresolved.

VIRTUALLY ALL SOCIAL, ECONOMIC, OR TECHNOLOGICAL decisions involve some degree of risk or uncertainty. In some cases, such as games of chance, the probabilities of the alternative consequences can be accurately determined. In other cases, actuarial or engineering data must be used to construct estimates of these likelihoods. To the extent that these probabilities can be quantified, however, individuals' attitudes toward risk can be subjected to theoretical analysis and empirical testing.

The 17th-century founders of modern probability theory such as Pascal and Fermat assumed that individuals would evaluate alternative monetary gambles on the basis of their expected values, so that a lottery offering the payoffs $(x_1, \ldots, x_n)$ with respective probabilities $(p_1, \ldots, p_n)$ would yield as much satisfaction as a sure payment equal to its expected value $\bar{x} = \Sigma x_i p_i$. Such an approach could be justified by appealing to the law of large numbers, which states that if a gamble is indefinitely and independently repeated, its long-run average payoff will necessarily converge to its expected value.

However, in a one-shot choice situation which cannot be replicated or averaged, individuals may well base their decisions on more than just the expected values of the alternative prospects. This point was dramatically illustrated by an example offered by Nicholas Bernoulli in 1728 and now known as the St. Petersburg Paradox: Suppose someone offers to toss a fair coin repeatedly until it lands heads, and to pay you \$1 if this occurs on the first toss, \$2 if it takes two tosses to land a head, \$4 if it takes three tosses, \$8 if it takes four tosses, and so on. What is the largest sure payment you would be willing to forgo in order to undertake a single play of this game?

This game offers a 1/2 chance of winning \$1, a 1/4 chance of winning \$2, and so on; its expected payoff is $(1/2)\$1 + (1/4)\$2 + (1/8)\$4 + \ldots = \$1/2 + \$1/2 + \$1/2 + \ldots = \$\infty$. However, even if such a well-backed offer could be made, it was felt that few individuals would forgo more than, say, \$20 for a one-shot play—a far cry from the (infinite) expected value (*1*). The resolution of this paradox, proposed independently by Gabriel Cramer and Daniel Bernoulli (Nicholas's cousin), would form the basis for the modern theory of decision-making under risk (*2*).

Arguing that a gain of \$1000 was not necessarily valued ten times as much as a gain of \$100, Cramer and Bernoulli hypothesized that individuals possess a utility of wealth function $U(x)$, and that they would value a lottery on the basis of its expected utility $\bar{u} = \Sigma U(x_i)p_i$ rather than its expected value $\bar{x} = \Sigma x_i p_i$. If utility took the logarithmic form $U(x) \equiv \ln(x)$, for example, the sure monetary gain ξ which would yield the same level of satisfaction as the St. Petersburg gamble would be given by the solution to

$$\ln(w + \xi) = (1/2)\ln(w + 1) + (1/4)\ln(w + 2) + (1/8)\ln(w + 4) + \ldots \quad (1)$$

where w denotes the individual's initial wealth. If w = \$1000, the individual would be indifferent between taking this gamble or receiving a sure gain of about \$6, if w = \$50,000 this amount is about \$9. Of course, someone with a different utility function $U^*(\cdot)$ would assign a different sure monetary equivalent.

Two centuries later, this approach was formally axiomatized by Frank Ramsey in his treatise on the philosophy of belief, by John von Neumann and Oskar Morgenstern in their development of the theory of games, and by Leonard Savage in his work on the foundations of statistical inference (*3*). The simplicity and intuitive appeal of its axioms, the elegance of its representation of risk attitudes in terms of properties of the utility function, and the tremendous number of theoretical results it has produced have led the expected utility model to become the dominant, and indeed, almost exclusive model of decision-making under risk in economics, operations research, philosophy, and statistical decision theory (*4*).

However, these theoretical advances have been accompanied by an accumulating body of empirical evidence suggesting that individuals systematically violate the predictions of the expected utility model. The largest and most systematic class of these violations concerns the key behavioral assumption of the model, the so-called independence axiom. This has led to a growing tension in the field of decision theory, with defenders of the expected utility approach stressing the "rationality" of this axiom and the theoretical power of the model, and others emphasizing the importance of the empirical evidence and developing alternatives to the expected utility model.

My purpose in this article is to give an overview of these developments in the field of decision-making under risk. The following sections provide a description of the expected utility model both as a theoretical tool and as a behavioral hypothesis, a summary of the evidence on systematic violations of the independence axiom, and a report on the newer non-expected utility models of decision-making currently being developed. This latter work has shown that the basic results of expected utility analysis are in fact quite robust to violations of the independence axiom. However, separate evidence suggests that some of the other standard assumptions in the theory of choice under uncertainty may also be suspect.

The author is associate professor of economics, University of California, San Diego, La Jolla, CA 92093.

The Expected Utility Model

The expected utility model follows standard economic theory by specifying a set of objects of choice and assuming that the individual's preferences can be represented by a real-valued function over this choice set. Since it is a model of decision-making under risk, the objects of choice are not the ultimate outcomes that might obtain (for example, alternative wealth levels) but rather probability distributions over these outcomes. Given a set $\{x_1, \ldots, x_n\}$ of potential outcomes, the choice set thus consists of all probability distributions $P = (p_1, \ldots, p_n)$ over $\{x_1, \ldots, x_n\}$, where p_i denotes the probability of obtaining x_i and $\Sigma p_i = 1$.

The model then assumes that the individual's preferences can be represented by a real-valued maximand or preference function $V(\cdot)$ over probability distributions, in the sense that the distribution $P^* = (p_1^*, \ldots, p_n^*)$ is preferred to $P = (p_1, \ldots, p_n)$ if and only if $V(P^*) > V(P)$, and is indifferent to P if and only if $V(P^*) = V(P)$. The essence of the expected utility approach is that $V(\cdot)$ takes the linear form $V(P) \equiv \Sigma U(x_i)p_i$ for some set of coefficients $\{U(x_i)\}$, so that expected utility preferences can be described as being linear in the probabilities. When the outcome set is a continuum such as the interval $[0,M]$, the probability distribution of a random variable $\tilde{x}$ over $[0,M]$ can be represented by its density function $f(\cdot)$, or more generally, by its cumulative distribution function $F(\cdot)$ [where $F(x) \equiv \text{prob}(\tilde{x} \leq x)$], and preferences over such distributions are assumed to be representable by linear preference functionals of the form $V(f) \equiv \int U(x)f(x)dx$ or $V(F) \equiv \int U(x)dF(x)$, which can again be interpreted as the expectation of $U(\cdot)$ (*5*). Since it is clear that the transformed utility function $aU(\cdot) + b$ $(a > 0)$ will generate the same ranking over distributions as $U(\cdot)$, utility functions are often normalized so that $U(0) = 0$ and $U(M) = 1$.

Figure 1 illustrates how this model can be used to represent various attitudes toward risk. The monotonicity of the utility functions $U(\cdot)$ and $U^*(\cdot)$ in the figures reflect the property of stochastic dominance preference, where one probability distribution is said to stochastically dominate another if it can be obtained from the latter by a sequence of rightward shifts of probability mass. Since such shifts raise the probability of obtaining at least x for all values of x, stochastic dominance preference is the probabilistic analogue of the view that "more is better."

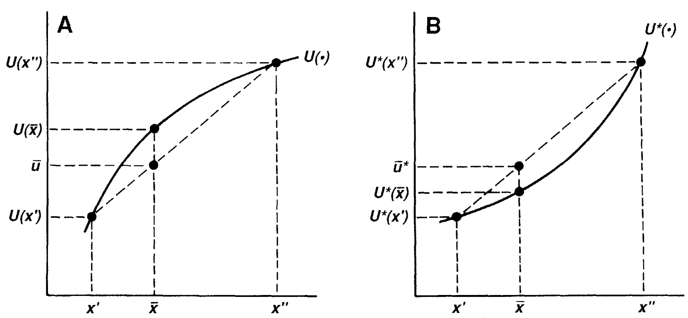

Fig. 1. (**A**) Concave utility function of a risk averse individual. (**B**) Convex utility function of a risk-preferring individual.

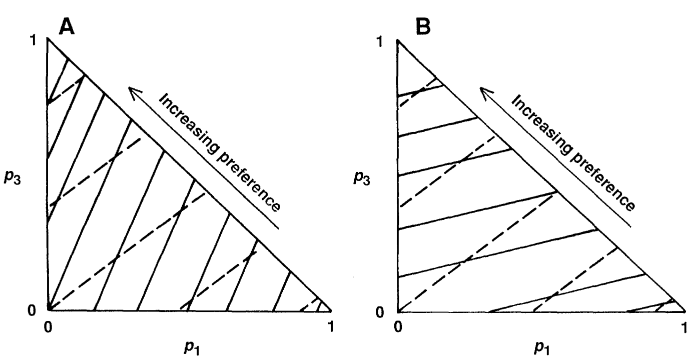

Fig. 2. (**A**) Relatively steep indifference curves of a risk averse individual. (**B**) Relatively flat indifference curves of a risk-preferring individual.

The points $\bar{x} = (2/3)x' + (1/3)x''$ in the figure denote the expected value of the gamble offering a 2/3 : 1/3 chance of the outcomes x' or x'', and $\bar{u} = (2/3)U(x') + (1/3)U(x'')$ and $\bar{u}^* = (2/3)U^*(x') + (1/3)U^*(x'')$ give the expected utilities of this gamble for $U(\cdot)$ and $U^*(\cdot)$. For the concave (that is, bowed upward) utility function $U(\cdot)$ we have $\bar{u} < U(\bar{x})$, implying that the individual would rather receive a sure payment equal to the expected value of the gamble than actually take the gamble itself. For the convex (bowed downward) utility function $U^*(\cdot)$ we have $\bar{u}^* > U^*(\bar{x})$, so that this individual would prefer to bear the risk rather than receive a sure payment of $\bar{x}$. Since Jensen's inequality (*6*) implies that these respective attitudes will extend to all risky gambles, $U(\cdot)$ is referred to as risk averse and $U^*(\cdot)$ as risk preferring (*7*). Researchers such as Arrow and Pratt have shown how the relative concavity or convexity of a utility function, as measured by the curvature index $-U''(x)/U'(x)$, can lead to theoretical predictions of how risk attitudes, and hence behavior, will vary with wealth or across individuals in a variety of different risky situations (*8*).

Although these figures illustrate the flexibility of the expected utility model compared to the Pascal-Fermat expected value model, an alternative graphical approach can be used to highlight the behavioral restrictions implied by the hypothesis of linearity in the probabilities. Consider the set of all distributions (p_1, p_2, p_3) over the fixed outcome levels $\{x_1, x_2, x_3\}$, where $x_1 < x_2 < x_3$. Since $p_2 \equiv 1 - p_1 - p_3$, we can represent this set of distributions by the points in the unit triangle in the (p_1,p_3) plane (Fig. 2). Since upward movements in the triangle increase p_3 at the expense of p_2 (that is, shift probability mass from the outcome x_2 up to x_3) and leftward movements reduce p_1 to the benefit of p_2 (shift probability from x_1 up to x_2), these movements (and more generally, all northwest movements) result in stochastically dominating distributions and would accordingly be preferred. Since the individual's indifference curves or iso-expected utility loci in this diagram are given by the solutions to

$$U(x_1)p_1 + U(x_2)[1 - p_1 - p_3] + U(x_3)p_3 = \text{constant} \qquad (2)$$

they will consist of parallel straight lines (the solid lines in the figures), with more preferred indifference curves lying to the northwest. This implies that knowledge of an individual's indifference curves over any small region is sufficient to know their preferences over the entire set of distributions.

The dashed lines in Fig. 2 are not indifference curves but rather iso-expected value lines, that is, solutions to

$$x_1p_1 + x_2[1 - p_1 - p_3] + x_3p_3 = \text{constant} \qquad (3)$$

Since northeast movements along these lines do not change the expected value of the distribution but do increase the probabilities of the tail outcomes x_1 and x_3 at the expense of the middle outcome x_2, they represent the set of increases in risk in this diagram (*7*). When the utility function $U(\cdot)$ is concave (risk averse), its indifference curves can be shown to be steeper than the iso-expected value lines (Fig. 2A), and increases in risk will lead to lower indifference curves. When $U(\cdot)$ is convex (risk preferring), its indifference curves will be flatter than the iso-expected value lines (Fig. 2B), and increases in risk will lead to higher indifference curves. If we compare two different utility functions, the one which is more risk averse (more concave) will possess the steeper indifference curves.

The property of linearity in the probabilities can also be represented as a restriction on the individual's attitudes toward probability mixtures of distributions. Given an outcome set $\{x_1, \ldots, x_n\}$, the α:$(1-\alpha)$ probability mixture of the distributions $P^* = (p_1^*, \ldots, p_n^*)$ and $P = (p_1, \ldots, p_n)$ is defined as the distribution $\alpha P^* + (1-\alpha)P = (\alpha p_1^* + (1-\alpha)p_1, \ldots, \alpha p_n^* + (1-\alpha)p_n)$. This may be thought of as that single-stage distribution which yields the same ultimate probabilities over $\{x_1, \ldots, x_n\}$ as a two-stage lottery which offers an α:$(1-\alpha)$ chance of winning the distributions P^* or P (*9*). Since linearity of $V(\cdot)$ implies that $V(\alpha P^* + (1-\alpha)P) \equiv \alpha V(P^*) + (1-\alpha)V(P)$, expected utility preferences will exhibit the following property, known as the independence axiom (*10*): If P^* is preferred (indifferent) to P, then the mixture $\alpha P^* + (1-\alpha)P^{**}$ will be preferred (indifferent) to $\alpha P + (1-\alpha)P^{**}$ for all $\alpha > 0$ and P^{**}. This condition, which is in fact equivalent to the property of linearity in the probabilities, can be interpreted as follows: In terms of the ultimate probability of obtaining each outcome, the choice between the mixtures $\alpha P^* + (1-\alpha)P^{**}$ and $\alpha P + (1-\alpha)P^{**}$ is equivalent to being offered a coin with a $(1-\alpha)$ chance of landing tails, in which case you will receive the lottery P^{**}, and being asked before the flip whether you would rather receive the lottery P^* or P in the event of a head. Now either tails will come up, in which case your choice will not have mattered, or else heads will come up, in which case you are "in effect" back to a choice between P^* or P, and you "should" have the same preferences over them as you would before.

Even though its conclusion is prescriptive, this argument has played a large role in the widespread adoption of expected utility maximization as a descriptive model of choice under risk. However, with a few exceptions in the early 1950s, it is only recently that the expected utility hypothesis has undergone the type of empirical testing that such a widely used behavioral hypothesis might be expected to receive.

Systematic Violations of the Independence Axiom

The earliest example of systematic violation of the independence axiom (or equivalently, of linearity in the probabilities) is known as the Allais Paradox (*11*). This example involves obtaining individuals' preference rankings over each of the following pairs of gambles

$$a_1\{1.00 \text{ chance of } \$1{,}000{,}000 \text{ versus } \quad a_2\begin{cases}0.10 \text{ chance of } \$5{,}000{,}000\\ 0.89 \text{ chance of } \$1{,}000{,}000\\ 0.01 \text{ chance of } \$0\end{cases}$$

$$a_3\begin{cases}0.10 \text{ chance of } \$5{,}000{,}000\\ 0.90 \text{ chance of } \$0\end{cases} \text{versus} \quad a_4\begin{cases}0.11 \text{ chance of } \$1{,}000{,}000\\ 0.89 \text{ chance of } \$0\end{cases}$$

Setting $\{x_1,x_2,x_3\} = \{\$0;\$1{,}000{,}000;\$5{,}000{,}000\}$, these four gambles form a parallelogram in the (p_1,p_3) triangle (Fig. 3A). Under the expected utility hypothesis, a preference for a_1 in the first pair would indicate relatively steep indifference curves (as in the figure), and hence a preference for a_4 in the second pair. In the alternative case of relatively flat indifference curves, the gambles a_2 and a_3 would be preferred (*12*). However, experimenters have repeatedly found that the modal, if not majority, preference of subjects has been for a_1 in the first pair and a_3 in the second (*13–15*), suggesting that indifference curves are not parallel but rather fan out (Fig. 3B).

Although initially dismissed as an isolated example, the Allais Paradox is now known to be a special case of a general empirical phenomenon termed the common consequence effect. This effect involves pairs of probability mixtures of the form

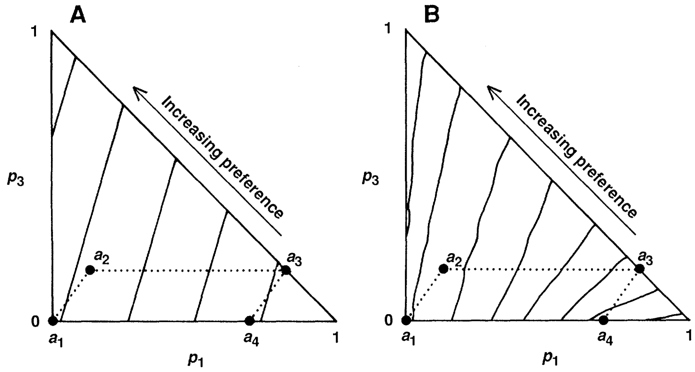

Fig. 3. (**A**) Expected utility indifference curves and the Allais Paradox. (**B**) Non-expected utility indifference curves that "fan out" and the Allais Paradox.

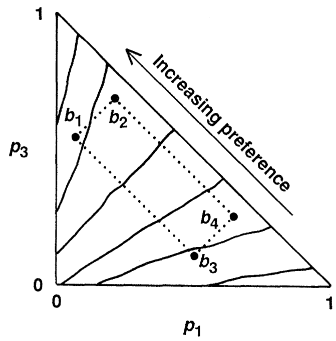

Fig. 4. Non-expected utility indifference curves that fan out and the Common Consequence Effect.

$$b_1\text{: } \alpha\delta_x + (1-\alpha)P^{**} \text{ versus } b_2\text{: } \alpha P + (1-\alpha)P^{**}$$

$$b_3\text{: } \alpha\delta_x + (1-\alpha)P^{*} \text{ versus } b_4\text{: } \alpha P + (1-\alpha)P^{*}$$

where δ_x is the degenerate distribution yielding the outcome x with certainty, P involves outcomes both greater and less than x, and P^{**} stochastically dominates P^* (*16*). Although the independence axiom clearly implies choices of either b_1 and b_3 or else b_2 and b_4, researchers have again found a tendency for subjects to prefer b_1 in the first pair and b_4 in the second (*17–20*). When the component distributions δ_x, P, P^*, and P^{**} are each over a common outcome set $\{x_1, x_2, x_3\}$, the prospects b_1, b_2, b_3, and b_4 will again form a parallelogram in the (p_1, p_3) triangle, and a preference for b_1 and b_4 again implies indifference curves that fan out (Fig. 4). It is important to note that it is not merely the high degree of such violations (ranging from 20 to 80%) that is important, but rather the systematic nature of these departures from linearity in the probabilities: although a preference for b_2 and b_3 would also violate the expected utility hypothesis (implying indifference curves that fan in), such choices account for a very small proportion of the total violations of expected utility in these studies.

A second class of systematic violations, stemming from another early example of Allais (*11*), is known as the common ratio effect. This phenomenon involves pairs of gambles of the form:

$$c_1\begin{cases}p \text{ chance of } \$X\\ 1-p \text{ chance of } \$0\end{cases} \text{versus} \quad c_2\begin{cases}q \text{ chance of } \$Y\\ 1-q \text{ chance of } \$0\end{cases}$$

$$c_3\begin{cases}rp \text{ chance of } \$X\\ 1-rp \text{ chance of } \$0\end{cases} \text{versus} \quad c_4\begin{cases}rq \text{ chance of } \$Y\\ 1-rq \text{ chance of } \$0\end{cases}$$

where $p > q$, $0 < X < Y$, and $r \in (0, 1)$ and includes the "certainty effect" of Kahneman and Tversky (*19*) and the "Bergen Paradox" of Hagen (*21*) as special cases (*22*). Setting $\{x_1, x_2, x_3\} = \{0, X, Y\}$ and plotting these gambles in the (p_1, p_3) triangle, the line segments $\overline{c_1c_2}$ and $\overline{c_3c_4}$ are seen to be parallel, as in Fig. 5A, so that the expected

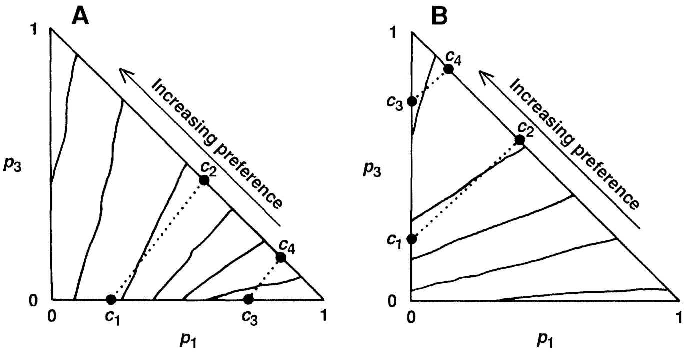

Fig. 5. (**A**) Non-expected utility indifference curves that fan out and the Common Ratio Effect. (**B**) Non-expected utility indifference curves that fan out and the Common Ratio Effect with negative payoffs.

utility model again predicts choices of c_1 and c_3 (if indifference curves are steep) or else c_2 and c_4 (if they are flat). However, investigators have found a systematic tendency to depart from these predictions in the direction of preferring c_1 and c_4 (*18, 20, 23*), which again suggests that indifference curves fan out (Fig. 5A). In a variation on this approach, Kahneman and Tversky (*19*) replaced the gains of X and Y in the above gambles with losses of these magnitudes and found a tendency to depart from expected utility in the direction of c_2 and c_3. Defining $\{x_1,x_2,x_3\}$ as $\{-Y,-X,0\}$ (to maintain the ordering $x_1 < x_2 < x_3$) and plotting these gambles (Fig. 5B), a choice of c_2 and c_3 is again seen to reflect the property of fanning out.

A third type of systematic departure from the expected utility model involves the elicitation or assessment of subjects' utility functions, and hence has important implications for the use of such procedures in applied decision analysis. A standard assessment procedure, termed the fractile method, begins by adopting the normalization $U(0) = 0$ and $U(M) = 1$ for some positive M and then picking a fixed mixture probability, say 1/2. The next step involves determining the individual's sure monetary equivalent ξ_1 of a 1/2:1/2 chance of M or 0, which implies that $U(\xi_1) = (1/2)U(M) + (1/2)U(0) = 1/2$. Finding the sure monetary equivalents of the 1/2:1/2 chances of ξ_1 or 0 and of M or ξ_1 yields the values ξ_2 and ξ_3 which solve $U(\xi_2) = 1/4$ and $U(\xi_3) = 3/4$. By repeated application of this procedure, the utility function can in the limit be completely assessed. However, there is no reason why the mixture probability must be 1/2. Picking any other $\alpha \in (0,1)$ and defining ξ_1^*, ξ_2^*, and ξ_3^* as the sure monetary equivalents of the $\alpha:(1 - \alpha)$ chances of M or 0, ξ_1^* or 0, and M or ξ_1^* yield the equations $U(\xi_1^*) = \alpha$, $U(\xi_2^*) = \alpha^2$, $U(\xi_3^*) = \alpha + (1 - \alpha)\alpha$, and such a procedure can also be used to recover $U(\cdot)$.

Although this assessment procedure ought to recover the same (normalized) utility function for any mixture probability α, researchers have found a systematic tendency for higher values of α to lead to the "recovery" of higher valued utility functions, as in Fig. 6A (*24, 25*). By illustrating the probability distributions used to obtain the responses ξ_1, ξ_2, and ξ_3 for $\alpha = 1/2$, ξ_1^* for $\alpha = 1/4$, and ξ_1^{**} for $\alpha = 3/4$, Fig. 6B shows that, as with the common consequence and common ratio effects, this utility evaluation effect is precisely what would be expected from an individual whose indifference curves departed from expected utility by fanning out (*26*).

Some proponents of the expected utility model have criticized the above findings on grounds such as "the experimental subjects were not experienced at making decisions under uncertainty," "once subjects were made aware of their violations of the independence axiom they would correct their choices," and "the experiments did not involve real money payments and hence do not reflect real-world behavior." To the extent that these claims have been examined by the appropriate experimental modifications, however, they have not been supported (*14, 15, 17*). In an interesting extension of these studies, Battalio, Kagel, and MacDonald (*27*) have shown that laboratory rats choosing among gambles involving substantial variations in their actual daily food intake also exhibit fanning-out behavior.

Non-Expected Utility Models

The extent and systematic nature of these observed departures from expected utility have led several researchers to generalize this model by positing nonlinear functional forms for the preference function $V(\cdot)$. Examples of these include the functional forms $V(P) \equiv \Sigma v(x_i)\pi(p_i)$, $V(P) \equiv \Sigma v(x_i)\,\pi(p_i)/\Sigma\pi(p_j)$, $V(P) \equiv \Sigma v(x_i)p_i + [\Sigma\tau(x_j)p_j]^2$, $V(P) \equiv \Sigma v(x_i)p_i/\Sigma\tau(x_j)p_j$, and $V(P) \equiv \Sigma v(x_i)\,[g(p_1 + \ldots + p_i) - g(p_1 + \ldots + p_{i-1})]$, as well as their extensions to density functions $f(\cdot)$ or cumulative distribution functions $F(\cdot)$ (*19, 20, 24, 28–32*). Many (though not all) of these forms are flexible enough to exhibit the properties of stochastic dominance preference and risk aversion–risk preference in a non-expected utility framework and have proven to be both theoretically and empirically useful. As in the expected utility case, these preference functions can be empirically assessed and then used to predict the individual's behavior in other situations.

However, while such forms allow for the modelling and analysis of preferences that are more general than those allowed by the expected utility hypothesis, they possess two limitations. (i) Each requires a different set of conditions on its component functions $v(\cdot)$, $\pi(\cdot)$, $\tau(\cdot)$, $g(\cdot)$, . . . for the properties of stochastic dominance preference, risk aversion–risk preference, comparative risk aversion, and so on, so that expected utility theorems linking properties of the function $U(\cdot)$ to such aspects of behavior will typically not extend, for example, to the corresponding properties of the function $v(\cdot)$ in these models. (ii) Each replaces the independence axiom by some other (albeit more general) global restriction on preferences, possibly subject to similar types of systematic empirical violations.

An alternative approach to the study of non-expected utility preferences proceeds not by specifying a particular nonlinear form for the preference function, but rather by considering nonlinear functions in general and by using calculus to extend results from expected utility theory in the same manner in which it is typically used to extend results involving linear functions (*29, 30*). Specifically, taking the first order Taylor expansion of a differentiable ("smooth") preference function $V(\cdot)$ about the distribution P yields

$$V(P^*) - V(P) = \Sigma U(x_i;P)[p_i^* - p_i] + o(\|P^* - P\|) \qquad (4)$$

where P^* is any other distribution, $U(x_i;P) \equiv \partial V(P)/\partial \text{prob}(x_i) = \partial V(p_1, \ldots, p_n)/\partial p_i$, $\|P^* - P\| \equiv [\Sigma(p_i^* - p_i)^2]^{1/2}$ is the Euclidean norm, and $o(\cdot)$ denotes a function which is zero at zero and of higher order than its argument. In the case of a preference functional $V(\cdot)$ over cumulative distribution functions, this expansion can be shown to take the form

$$V(F^*) - V(F) = \int U(x;F)[dF^*(x) - dF(x)] + o(\|F^* - F\|) \qquad (5)$$

where $\int U(x;F)[dF^*(x) - dF(x)]$ is the classical variational derivative of the functional $V(\cdot)$ and $\|F^* - F\| \equiv \int |F^*(x) - F(x)|dx$ is the standard L^1 norm.

In each of these cases, it follows that the individual's evaluation of differential shifts from the distributions P or $F(\cdot)$ will be determined by the first order (that is, linear) terms $\Sigma U(x_i;P)[p_i^* - p_i]$ or $\int U(x;F)[dF^*(x) - dF(x)]$. However, since these may be written as $\Sigma U(x_i;P)p_i^* - \Sigma U(x_i;P)p_i$ or $\int U(x;F)dF^*(x) - \int U(x;F)dF(x)$, it follows that an individual with smooth preferences will evaluate alternative differential shifts from either P or $F(\cdot)$ precisely as would an expected utility maximizer with "local utility function" $U(x_i;P)$ or $U(x;F)$. Thus, for example, an individual would prefer all first order stochastically dominating differential shifts from P or $F(\cdot)$ if and only if $U(x_i;P)$ or $U(x;F)$ were increasing in x, and would be averse to (prefer) all differential increases in risk (*7*) if and only if $U(x_i;P)$ or $U(x;F)$ were concave (convex) in x. Intuitively, this result follows immediately from (multivariate or variational) calculus: since smooth preferences are "locally linear" in the probabilities and linearity in the probabilities is equivalent to expected utility maximization, smooth preferences will be "locally expected utility maximizing."

Of course, the above results will only hold exactly in a vanishingly small neighborhood of any probability distribution. However, we can exploit another result from standard calculus to show how "expected utility" theory may be applied to the exact global analysis of smooth preferences over probability distributions. Recall that in many instances, a differentiable function will exhibit a particular qualitative property if and only if this property is exhibited by its linear approximations at every point: for example, a differentiable function will be nondecreasing if and only if its linear approximations are nondecreasing (that is, its partial derivatives are nonnegative) at each point. Most of the fundamental aspects of attitudes toward risk and their expected utility characterizations turn out to be of this type. In particular, it can be shown that (i) a smooth preference functional $V(F)$ will prefer all local or global first order stochastically dominating shifts if and only if its local utility functions $U(x;F)$ are increasing in x for all $F(\cdot)$, (ii) a smooth preference functional $V(F)$ will be averse to (prefer) all local or global increases in risk (*7*) if and only if its local utility functions $U(x;F)$ are concave (convex) in x for all $F(\cdot)$, and (iii) the smooth preference functional $V_a(F)$ will be more risk averse (in the appropriately defined senses) than $V_b(F)$ if and only if the Arrow-Pratt curvature indices $-U_a''(x;F)/U_a'(x;F)$ of its local utility functions $U_a(x;F)$ are greater than $-U_b''(x;F)/U_b'(x;F)$ for all x and $F(\cdot)$. Analogous results hold for the preference function $V(P)$ and its local utility functions $U(x_i;P)$ in the case of a finite outcome set $\{x_1, \ldots, x_n\}$.

Figure 7 illustrates this second result for the outcome set $\{x_1,x_2,x_3\}$. The solid curves are the indifference curves of a smooth non-expected utility preference function $V(P)$. The solid parallel lines near the point P_o denote the expected utility preference field that approximates these preferences in the neighborhood of P_o, or in other words, the indifference curves generated by the local utility function $U(x_i;P_o)$. Since these lines are tangent to the actual indifference curve at P_o, a differential shift from this distribution will be preferred (lead to a higher indifference curve) if and only if it is preferred by the local expected utility approximation (leads to a higher tangent line). Since increases in risk consist of northeast movements along the dashed iso-expected value lines, global risk aversion is equivalent to the condition that the indifference curves be everywhere steeper than these lines. However, this is equivalent to the condition that all of the tangent approximations be steeper than these lines, which is in turn equivalent to the condition that all of the local expected utility approximations be risk averse, or in other words, that all of the local utility functions $U(x_i;P)$ be concave in x_i (*30*).

This approach may also be used to obtain a uniform characterization of the type of departures from linearity in the probabilities described in the previous section. In particular, it can be shown that an individual's indifference curves will fan out in the (p_1,p_3) triangle for every triple $\{x_1,x_2,x_3\}$ if and only if the local utility function becomes more concave (in the sense of Arrow and Pratt) when evaluated at stochastically dominating distributions. Although the refutable predictions and public policy implications of this characterization (*30*) are of course weaker than those of the expected utility model, they are at least more closely tied to what we have actually observed about preferences.

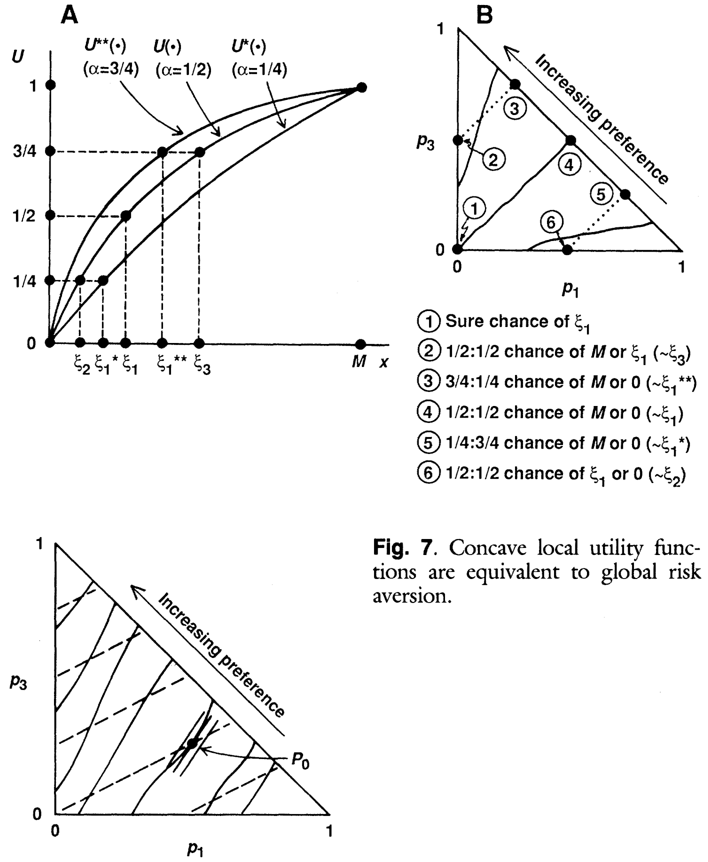

Fig. 7. Concave local utility functions are equivalent to global risk aversion.

My fellow researchers and I have shown how this and similar techniques can be applied to the analysis of both general and specific non-expected utility preference functionals in a manner that simultaneously exploits and demonstrates the robustness of the large body of theoretical results derived in the expected utility framework (*29–33*). It is hoped that the development and successful application of such an approach will tempt the mainstream of expected utility-trained researchers to examine the empirical nature and theoretical implications of behavior which departs from this standard model.

Unresolved Issues

Although the departures from linearity in the probabilities cited above constitute the most well-documented and systematic body of evidence on the validity of the expected utility hypothesis, investigators have uncovered additional empirical phenomena that call into question even more fundamental aspects of the model (*34*). One such class of examples is known as the preference reversal phenomenon. Here, individuals are presented with three questions, typically embedded in a random order among a group of several other questions. One asks their preferences over a pair of prospects of the form

$$P\text{-bet}\begin{cases} p \text{ chance of } \$X \\ 1-p \text{ chance of } \$x \end{cases} \qquad \text{versus} \qquad \$\text{-bet}\begin{cases} q \text{ chance of } \$Y \\ 1-q \text{ chance of } \$y \end{cases}$$

where X and Y are respectively greater than x and y, p is greater than q, and Y is greater than X (the names "P-bet" and "\$-bet" come from the greater probability of winning in the first gamble and greater possible gain in the second). In some cases, x and y took on small negative values. The second and third questions involve obtaining the individuals' sure monetary equivalents ξ_P and $\xi_\$$ of these two prospects.

When the probabilities and outcomes are such that the expected values of the gambles are similar, experimenters have found a tendency for the majority of subjects to prefer the P-bet to the \$-bet in a direct comparison. However, when asked for their sure monetary equivalents, a majority will assign a higher sure monetary equivalent to the \$-bet than to the P-bet (*35, 36*). On the assumption that individuals will always prefer greater sure amounts to smaller ones, this implies that the \$-bet is indifferent to the amount $\xi_\$$, which is preferred to the (smaller) amount ξ_P, which is indifferent to the P-bet, which is preferred to the \$-bet, so that preferences are cyclic or intransitive. Since such intransitivity is incompatible with the existence of any real-valued maximand or preference function (expected utility or otherwise), it constitutes evidence against both the expected utility model as well as each of the non-expected utility models described above. It is worth noting that this phenomenon has proven to be remarkably robust, having withstood testing by (initially) skeptical experimenters, the use of real-money gambles, and the use of subjects presumably well-versed in making decisions under uncertainty (professional gamblers) (*36*).

Another class of empirical phenomena, termed framing effects, consists of examples in which alternative, but probabilistically equivalent, statements of a decision problem yield different expressed choices. Early examples of this phenomenon involved different ways of generating the same probability distributions out of compound gambles or additive pairs of gambles, and resulted in reversals of choice (*37*). More recent studies have found, for example, that the preferred choice will depend on whether an otherwise identical problem is phrased as a decision whether or not to "gamble" versus whether or not to "insure," whether the effects of a preventive vaccine program are specified in terms of "lives saved" versus "lives lost," and so on (*38*).

In addition to this evidence that probabilistically equivalent restatements of a problem may affect decisions, there is also evidence that when experimenters do not explicitly state probabilities, subjects may not formulate uncertain choice problems in any probabilistically coherent manner. In the simplest of a class of examples due to Daniel Ellsberg (*14, 18, 39*), subjects were presented with a pair of urns, the first containing 50 red balls and 50 black balls and the second also containing 100 red and black balls but in an unknown proportion. When faced with the choice of staking a prize on drawing a red ball from either urn, a majority of subjects strictly preferred to draw from the first urn. However, a majority also strictly preferred drawing from the first urn when the prize was staked on drawing a black ball. It is clear that there exists no subjectively assigned pair of probabilities $p:(1-p)$ of drawing a red versus a black ball from the second urn (even 1/2:1/2) which can simultaneously generate both of these strict preferences.

Researchers have attempted to extend the non-expected utility models of the previous section to address each of these types of empirical phenomena. Models of nontransitive preference rankings have been developed that are able to accommodate the preference reversal phenomenon as well as many of the observed violations of linearity in the probabilities (*40*). Although there is of yet no single uniform principle (analogous to fanning out) known to underlie the diverse group of framing examples, some progress has nonetheless been made in this area as well (*41*). Finally, Ellsberg-type phenomena have inspired generalizations of expected utility theory where individuals' probabilistic beliefs are represented by "nonadditive" probability measures which do not satisfy the standard laws of probability theory (*42*). As with the non-expected utility models described earlier, the long-run success of these endeavors will depend on the extent to which they can usefully address the tremendous number of issues in the theory of individual and group choice under uncertainty to which the expected utility model has been applied.

REFERENCES AND NOTES

1. The infinite expectation in this example is not really necessary to make the point: agreeing to limit the game to at most 1 million tosses will still lead to a striking discrepancy between most individuals' valuations of the modified gamble and the expected payoff of \$500,000.
2. G. Cramer, correspondence to N. Bernoulli (1728), and D. Bernoulli, *Comment. Acad. Sci. Imp. Petropolitanae* **5**, 175 (1738) [English translations *Econometrica* **22**, 23 (1954)]. For historical discussions of the St. Petersburg Paradox and its impact, see I. Todhunter, *A History of the Mathematical Theory of Probability* (Chelsea, New York, 1949); P. Samuelson, *J. Econ. Lit.* **15**, 24 (1977).
3. F. Ramsey, *Foundations of Mathematics and Other Logical Essays* (Harcourt Brace, New York, 1931); J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior*, (Princeton Univ. Press, Princeton, NJ, ed. 3, 1953); L. Savage, *Foundations of Statistics* (Wiley, New York, 1954).
4. J. Hey, *Uncertainty in Microeconomics* (Martin Robinson, Oxford, 1979); F. Hillier and G. Lieberman, *Introduction to Operations Research*, (Holden-Day, Oakland, CA, ed. 4, 1986); E. Ells, *Rational Decision and Causality* (Cambridge Univ. Press, Cambridge, 1982); M. DeGroot, *Optimal Statistical Decisions* (McGraw-Hill, New York, 1970).
5. The Riemann-Stieltjes integral $\int U(x)dF(x)$ denotes the integral of $U(\cdot)$ with respect to the probability measure associated with the cumulative distribution function $F(\cdot)$ [R. Bartle, *Elements of Real Analysis* (Wiley, New York, ed. 2, 1976), sec. 29].
6. This result from probability theory states that the expectations $\Sigma U(x_i)p_i$, $\int U(x)f(x)dx$ or $\int U(x)dF(x)$ will be no greater (no less) than $U(\bar{x})$ whenever $U(\cdot)$ is a concave (convex) function [W. Feller, *An Introduction to Probability Theory and Its Applications* (Wiley, New York, ed. 2, 1971), p. 153].
7. More generally, a concave (convex) utility function will be averse to (prefer) all increases in risk, where a pair of shifts of probability mass is said to constitute an increase in risk if it moves mass away from the center of a probability distribution out toward its tails in a manner that preserves the expected value of the distribution [M. Rothschild and J. Stiglitz, *J. Econ. Theory* **2**, 225 (1970)].
8. K. Arrow, *Essays in the Theory of Risk-Bearing* (North-Holland, Amsterdam, 1974); J. Pratt, *Econometrica* **32**, 122 (1964).
9. The $\alpha:(1-\alpha)$ probability mixtures of the density functions $f^*(\cdot)$ and $f(\cdot)$ and the cumulative distribution functions $F^*(\cdot)$ and $F(\cdot)$ are given by $\alpha f^*(\cdot) + (1-\alpha)f(\cdot)$ and $\alpha F^*(\cdot) + (1-\alpha)F(\cdot)$, and possess similar interpretations.
10. This follows since $V(P^*) > (=)\ V(P)$ implies $V(\alpha P^* + (1-\alpha)P^{**}) = \alpha V(P^*) + (1-\alpha)V(P^{**}) > (=)\ \alpha V(P) + (1-\alpha)V(P^{**}) = V(\alpha P + (1-\alpha)P^{**})$. Early statements of this axiom may be found in J. Marschak [*Econometrica* **18**, 111 (1950)] and P. Samuelson [*ibid.* **20**, 670 (1952)].
11. M. Allais, *ibid.* **21**, 503 (1953). See also M. Allais, in *Expected Utility Hypotheses and the Allais Paradox*, M. Allais and O. Hagen, Eds. (Reidel, Dordrecht, Netherlands, 1979), pp. 27–145 and 437–681.
12. Algebraically, these cases are equivalent to the cases where $0.10U(5{,}000{,}000) - 0.11U(1{,}000{,}000) + 0.01U(0)$ is negative or positive, respectively.
13. D. Morrison, *Behav. Sci.* **12**, 373 (1967); H. Raiffa, *Decision Analysis: Introductory Lectures on Choices Under Uncertainty* (Addison-Wesley, Reading, MA, 1968), p. 80.
14. P. Slovic and A. Tversky, *Behav. Sci.* **19**, 368 (1974).
15. H. Moskowitz, *Decision Sci.* **5**, 225 (1974).
16. The Allais Paradox choices a_1, a_2, a_3, and a_4 correspond to b_1, b_2, b_4, and b_3, where $\alpha = 11/100$, $x = \$1{,}000{,}000$, P is a (10/11):(1/11) chance of \$5,000,000 or \$0, $P^* = \delta_{\$0}$, and $P^{**} = \delta_{\$1{,}000{,}000}$.
17. K. MacCrimmon, in *Risk and Uncertainty*, K. Borch and J. Mossin, Eds. (Macmillan, London, 1968), pp. 3–23.
18. ——— and S. Larsson, in *Expected Utility Hypotheses and the Allais Paradox*, M. Allais and O. Hagen, Eds. (Reidel, Dordrecht, Netherlands, 1979), pp. 333–409.
19. D. Kahneman and A. Tversky, *Econometrica* **47**, 263 (1979).
20. S. Chew and W. Waller, *J. Math. Psychol.* **30**, 55 (1986).
21. O. Hagen, in *Expected Utility Hypotheses and the Allais Paradox*, M. Allais and O. Hagen, Eds. (Reidel, Dordrecht, Netherlands, 1979), pp. 271–302.
22. Thus, for example, Kahneman and Tversky (*19*) found that 80% of their subjects preferred a sure gain of 3000 Israeli pounds to an 80% chance of winning 4000, but 65% preferred a 20% chance of winning 4000 to a 25% chance of winning 3000.
23. A. Tversky, *Erkenntnis* **9**, 163 (1975).
24. U. Karmarkar, *Organizational Behav. Hum. Performance* **21**, 61 (1978).
25. ———, "Effect of probabilities on the subjective evaluation of lotteries" (Sloan School of Business working paper 689-74, Cambridge, MA, 1974); M. McCord and R. de Neufville, in *Foundations of Utility and Risk Theory with Applications*, B. Stigum and F. Wenstøp, Eds. (Reidel, Dordrecht, Netherlands, 1983), pp. 181–199; M. McCord and R. de Neufville, *Large Scale Sys.* **6**, 91 (1984).
26. Set $\{x_1, x_2, x_3\} = \{0, \xi_1, M\}$. Then the ordering of the values ξ_1, ξ_2, ξ_3, ξ_1^* and ξ_1^{**} in Fig. 6A is derived from the individual's preference ordering over the five distributions in Fig. 6B for which they are the respective sure monetary equivalents.
27. R. Battalio, J. Kagel, D. MacDonald, *Am. Econ. Rev.* **75**, 597 (1985).
28. W. Edwards, *J. Exp. Psychol.* **50**, 201 (1955); J. Quiggin, *J. Econ. Behav. Organ.* **3**, 323 (1982); P. Fishburn, *J. Econ. Theory* **31**, 293 (1983); M. Yaari, *Econometrica* **55**, 95 (1987).

29. M. Machina, *Econometrica* **50**, 277 (1982).
30. ——, in *Foundations of Utility and Risk Theory with Applications*, B. Stigum and F. Wenstøp, Eds. (Reidel, Dordrecht, Netherlands, 1983), pp. 263–293.
31. S. Chew, *Econometrica* **51**, 1065 (1983).
32. P. Fishburn, *Math. Soc. Sci.* **8**, 63 (1984).
33. M. Machina, *J. Econ. Theory* **33**, 199 (1984); E. Dekel, *ibid.* **40**, 304 (1986); B. Allen, *ibid.*, in press; L. Epstein, *Econometrica* **53**, 945 (1985); E. Karni and Z. Safra, *ibid.*, in press.
34. For a more complete discussion of the following evidence and how researchers have responded to it, see M. Machina, *J. Econ. Perspect.*, in press.
35. H. Lindman, *J. Exp. Psychol.* **89**, 390 (1971); S. Lichtenstein and P. Slovic, *ibid.*, p. 46.
36. S. Lichtenstein and P. Slovic, *ibid.* **101**, 16 (1973); D. Grether and C. Plott, *Am. Econ. Rev.* **69**, 623 (1979); W. Pommerehne, F. Schneider, P. Zweifel, *ibid.* **72**, 569 (1982); R. Reilly, *ibid.*, p. 576; J. Mowen and J. Gentry, *J. Appl. Psychol.* **65**, 715 (1980).
37. P. Slovic, *J. Exp. Psychol.* **79**, 139 (1969); J. Payne and M. Braunstein, *ibid.* **87**, 13 (1971).
38. P. Shoemaker and H. Kunreuther, *J. Risk Insur.* **46**, 603 (1979); J. Hershey and P. Schoemaker, *ibid.* **47**, 111 (1980); A. Tversky and D. Kahneman, *Science* **211**, 453 (1981); D. Kahneman and A. Tversky, *Sci. Am.* **246**, 160 (January 1982); J. Hershey, H. Kunreuther, P. Schoemaker, *Manage. Sci.* **28**, 936 (1982).
39. D. Ellsberg, *Q. J. Econ.* **75**, 643 (1961); K. MacCrimmon, thesis, University of California, Los Angeles (1965).
40. D. Bell, *Oper. Res.* **30**, 961 (1982); P. Fishburn, *J. Math. Psychol.* **26**, 31 (1982); G. Loomes and R. Sugden, *Econ. J.* **92**, 805 (1982).
41. B. Fischhoff, *J. Exp. Psychol.* **9**, 103 (1983); D. Kahneman, A. Tversky, *Am. Psychol.* **39**, 341 (1984); A. Tversky, D. Kahneman, *J. Business* **59**, S251 (1986).
42. D. Schmeidler, unpublished manuscript.

43. I am grateful to the National Science Foundation (grant SES 83-08165) and the Alfred P. Sloan Foundation for financial support and to J. Sobel and M. Rothschild for helpful remarks on this material.

Accelerator Mass Spectrometry for Measurement of Long-Lived Radioisotopes

DAVID ELMORE AND FRED M. PHILLIPS

Particle accelerators, such as those built for research in nuclear physics, can also be used together with magnetic and electrostatic mass analyzers to measure rare isotopes at very low abundance ratios. All molecular ions can be eliminated when accelerated to energies of millions of electron volts. Some atomic isobars can be eliminated with the use of negative ions; others can be separated at high energies by measuring their rate of energy loss in a detector. The long-lived radioisotopes ^{10}Be, ^{14}C, ^{26}Al, ^{36}Cl, and ^{129}I can now be measured in small natural samples having isotopic abundances in the range 10^{-12} to 10^{-15} and as few as 10^5 atoms. In the past few years, research applications of accelerator mass spectrometry have been concentrated in the earth sciences (climatology, cosmochemistry, environmental chemistry, geochronology, glaciology, hydrology, igneous petrogenesis, minerals exploration, sedimentology, and volcanology), in anthropology and archeology (radiocarbon dating), and in physics (searches for exotic particles and measurement of half-lives). In addition, accelerator mass spectrometry may become an important tool for the materials and biological sciences.

RADIOISOTOPES HAVE LONG BEEN USED AS AN IMPORTANT (and sometimes the only) source of information regarding the chronology of geological processes, the history of meteorites and cosmic rays, human evolution, and the dynamics of biological systems. Short-lived radioisotopes, with half-lives less than about 1 year, can usually be measured with high sensitivity by conventional techniques, in which decay products are counted efficiently. Primordial radioisotopes with half-lives greater than about 10^9 years are relatively abundant naturally (since they have not completely decayed over the life of the solar system) and are used for dating by measuring the buildup of stable decay products. Radioisotopes with half-lives in the intermediate range of 10^3 to 10^8 years are difficult to measure with decay counting since only a small fraction of the atoms decay over a reasonable counting period of a few months or less. Many of the interesting processes that occur on the earth and in the solar system have time scales that fall in this interval. Over 30 elements have radioisotopes with half-lives in this range; five of these (Table 1) have now been measured with the new technique of accelerator mass spectrometry (AMS) with enough sensitivity for detection at natural levels. Although in some cases decay counting of large natural samples has been possible, sample sizes have been reduced by several orders of magnitude. For example, ^{36}Cl can be measured by conventional decay counting with 40 g of chloride (*1*) extracted from thousands of liters of contemporary water, whereas AMS requires only 1 mg of chloride and has backgrounds 10 to 100 times lower than those of conventional techniques.

Accelerator mass spectrometry can also be used to measure trace elements directly in unprocessed materials with backgrounds far lower than is possible with conventional techniques. Microprobe ion sources have been developed with 1-μm spatial resolution; these should be very useful in conjunction with AMS (*2*). Stable-isotope ratios of trace elements can also be measured with high sensitivity (*3*).

Although AMS had its beginnings in 1939 with the measurement of ^{3}He in helium at natural abundances by Alvarez and Cornog (*4*),

D. Elmore is a physicist with the Environmental Research Division, Argonne National Laboratory, Argonne, IL 60439. He remains temporarily at the Nuclear Structure Research Laboratory, University of Rochester, Rochester, NY 14627, where much of the research reported has been carried out. F. M. Phillips is a faculty member in the Geoscience Department and Geophysical Research Center, New Mexico Institute of Mining and Technology, Socorro, NM 87801.