BA 513: Ph.D. Seminar on Choice Theory
Professor Robert Nau
Spring Semester 2008

Notes and readings for class #3:   utility functions, risk aversion, and state-preference theory (revised September January 24)

Primary readings:

1a. "Risk aversion in the small and in the large" by John Pratt (1964)
1b. "On the role of securities in the optimal allocation of risk-bearing" by Kenneth Arrow (1953/1963/1974)
1c. Excerpt on "unidimensional utility theory" from *Decisions With Multiple Objectives: Preferences and Value Tradeoffs"* by Ralph Keeney and Howard Raiffa (1976)
1d. Excerpt on "the fundamental theorem of risk bearing" from *The Analytics of Uncertainty and Information* by Jack Hirshleifer and John Riley (1992)
1e. Excerpts on "utility theory" and "arbitrage and pricing" from *Theory of Financial Decision Making* by Jonathan Ingersoll (1987)
1f. Excerpt on "risk sharing and group decisions" from *Decision Analysis: Introductory Lectures on Choice Under Uncertainty* by Howard Raiffa (1968)

Supplementary readings:

2a. "A generalization of Pratt-Arrow measure to non-expected-utility preferences and inseparable probability and utility" by Robert Nau, *Management Science*, August 2003
2b. "Evaluating income streams" by Jim Smith, *Management Science*, December 1998
2c. "Scoring rules, generalized entropy, and utility maximization" by Victor Richmond R. Jose, Robert Nau, and Robert Winkler, forthcoming in *Operations Research*

Other related readings:

3a. "The theory of syndicates" by Robert Wilson (1968)
3b. "Increasing risk: I. a definition" by Michael Rothschild and Joseph Stiglitz
3c. "One-switch utility functions" by David Bell
3d. "Marginal value and intrinsic risk aversion" by David Bell and Howard Raiffa
3e. "The effect of the timing of consumption decisions and the resolution of lotteries on the choice of lotteries" by Michael Spence and Richard Zeckhauser (1972)
3f. "An aggregation theorem for securities markets" by Mark Rubinstein (1974)
3g. Excerpt on "theoretical foundations of value and utility measurement" from *Decision Analysis and Behavioral Research* by Detlof von Winterfeldt and Ward Edwards (1986)

3.1 Utility for money, risk aversion, and risk tolerance

This week's readings are concerned with developments in the theory of expected utility in the decades following von Neumann and Morgenstern's and Savage's pioneering axiomatic work, up to and including some recent work. During this time, expected utility theory became the tool of choice for modeling rational behavior under conditions of risk and uncertainty in microeconomics and applied decision analysis. The term "von Neumann-Morgenstern utility function" refers generically to a utility function that is used to model preferences under conditions of risk (objective probabilities) or uncertainty (subjective probabilities), and which is linear in the probabilities attached to consequences. Von Neumann-Morgenstern utility functions can be defined on lotteries with arbitrary consequences, but *utility functions for money* are of particular interest for obvious reasons: a function of one variable is the simplest mathematical function, and money plays a distinguished role as a medium of exchange and pricing in a competitive economy.

A utility function for money is now commonly called a *Bernoulli utility function* in recognition of the first use of a logarithmic utility function by Daniel Bernoulli (1738). In the 1960's and 1970's, the mathematical properties of general Bernoulli utility functions were worked out in detail. It was recognized that there are important *parametric forms* of utility functions— exponential, logarithmic, power, etc.—just as there are important parametric forms of probability distributions—normal, lognormal, gamma, etc. Indeed, we will see that particular utility functions pair up naturally with particular probability distributions to yield convenient closed-form solutions in various applications. Different-shaped utility functions represent different qualitative attitudes toward risk, and questions naturally arise as to the conditions under which one type of utility function or another is appropriate for decision analysis or economic modeling.

Pratt (1964) and Arrow (1965) introduced the now-standard definitions and measures of *risk aversion* in terms of properties of Bernoulli utility functions, although their ideas were anticipated in a 1952 paper of de Finetti. Bernoulli utility functions for money are functions of final wealth (at the end of the decision problem), not gains and losses relative to a (possibly uncertain) initial position. Thus $u(x)$ stands for the utility assigned to a final wealth of x. Technically the quantity $u(x)$ does not represent the happiness or well-being that is "experienced" *ex post* when final wealth is x. Rather, it is just an index that is used *ex ante* to evaluate choices under risk or uncertainty, and as such it may reflect intrinsic attitudes toward risk as well as decreasing marginal value of money. But mathematically it behaves just like a cardinal measure of experienced utility, and this important distinction is often glossed over. (It is also possible to define *multiattribute* utility functions, which are utility functions of more than one variable. For example, it is possible to have a multiattribute utility function defined over amounts of consumption that take place at different points in time. But never mind those complications just now...)

Let $\tilde{x}$ denote a *random variable* representing your level of final wealth—perhaps you are the owner of some lottery tickets or shares of stock or other financial assets and liabilities whose value, at the end of the time horizon being modeled, is uncertain. Presumably there is some amount of money which, if you possessed it with certainty, would make you just as happy as having the uncertain wealth $\tilde{x}$. This is what Pratt calls the *cash equivalent* of $\tilde{x}$, more

commonly known nowadays as the **certainty equivalent** of $\tilde{x}$. We will write this quantity as CE($\tilde{x}$). If $u(x)$ is your Bernoulli utility, the *expected utility* of $\tilde{x}$ is denoted by E[$u(\tilde{x})$], where the expectation is taken with respect to your own "true" subjective probability distribution, and CE($\tilde{x}$) is then defined by the equation:

$$u(CE(\tilde{x})) = E[u(\tilde{x})]$$

or equivalently

$$CE(\tilde{x}) = u^{-1}(E[u(\tilde{x})])$$

where u^{-1} denotes the inverse of the function u. The *expected value* of $\tilde{x}$, denoted E[$\tilde{x}$], is what we would normally consider the "actuarially fair" value of the uncertain wealth level $\tilde{x}$ according to your own true probabilities. However, most persons would trade $\tilde{x}$ for something less than its actuarially fair value, in order to be rid of the risk associated with $\tilde{x}$. Such persons are said to be **risk averse.** The **risk premium** associated with $\tilde{x}$ is the amount of expected value you are willing to give up—i.e., the "premium" you are willing to pay to an insurance company—in order to get rid of all the risk that you currently face. Let the risk premium of $\tilde{x}$ be denoted by RP($\tilde{x}$). Then RP($\tilde{x}$) is defined by the equation

$$RP(\tilde{x}) = E[\tilde{x}] - CE(\tilde{x})$$

You are risk averse "in the large" if RP($\tilde{x}$) > 0 for any wealth $\tilde{x}$ that is genuinely risky (i.e., non-constant), which will be true if your utility function u is a concave function (i.e., a function that exhibits diminishing marginal utility for money). On the other hand, if RP($\tilde{x}$) = 0 for you for all $\tilde{x}$, then you are said to be **risk neutral**, which will be true if u is a linear function.

Measures of risk aversion and risk tolerance. The key concept that de Finetti, Pratt and Arrow independently discovered is that there is a measure of "risk aversion in the small" that underlies the phenenomon of risk aversion in the large. The function $r(x)$ defined by

$$r(x) = -u''(x)/u'(x)$$

measures your local aversion to risk in the vicinity of wealth level x. That's minus the second derivative of u at x divided by the first derivative of u at x: it is a measure of the relative "curvature" of the utility function u at wealth x. Since the first derivative of u is presumably always positive, $r(x)$ is always opposite in sign to the second derivative of u, so it is positive when u has a negative second derivative, i.e., when u exhibits decreasing marginal utility. This function is now called the **Pratt-Arrow measure of absolute risk aversion**. The reciprocal of $r(x)$, which I will denote here as $t(x)$, is called your local **risk tolerance** at x. That is,

$$t(x) = 1/r(x) = -u'(x)/u''(x)$$

The function $r(x)$—or equivalently $t(x)$—contains all the essential information about your risk preferences that is encoded in your utility function, and it eliminates the arbitrary scale parameters and imaginary utility units that complicate the interpretation of a utility function.

Recall that if a function $u(x)$ represents your attitude toward risk, then so does any other function $au(x)+b$ where a is a positive constant and b is an arbitrary constant. These arbitrary constants drop out when $r(x)$ and $t(x)$ are calculated. $t(x)$ is conveniently measured in units of *money* (hence it is uniquely defined relative to a given currency), and $r(x)$ is correspondingly measured in units of reciprocal money.

Pratt's risk premium formula: Suppose that your wealth is $x + \tilde{z}$, where $\tilde{z}$ is a very *small* (in the limit, infinitesimal) and *actuarially fair* gamble (i.e., E[$\tilde{z}$]=0). In other words, your wealth is almost constant except for the small risk $\tilde{z}$. Then:

$$CE(x + \tilde{z}) \approx x - \tfrac{1}{2} r(x)Var[\tilde{z}]$$

or equivalently:

$$RP(x+\tilde{z}) \approx \tfrac{1}{2} r(x)E[\tilde{z}^2] = \tfrac{1}{2} r(x)Var[\tilde{z}] \qquad (1)$$

Thus, your premium for risks is proportional to the *variance* of the payoff, where the constant of proportionality is your local risk aversion measure (times one-half). This means that $\tfrac{1}{2} r(x)$ is the "price per unit of variance" that you are willing to pay to get rid of small risks in the vicinity of an otherwise-constant wealth level x. Note that the units of measurement work out just right here: variance is measured in units of squared dollars, and risk aversion is measured in reciprocal dollars, so their product is measured in dollars. (Pratt also derives a nifty, similar-looking "probability premium" for binary gambles in which you stand to win or lose a fixed quantity of money with equal probabilities.)

Risk aversion in the small, as measured by $r(x)$, determines risk aversion in the large, as measured by RP($\tilde{x}$). For, let the risk aversion measures of two individuals be denoted $r_1(x)$ and $r_2(x)$. If $r_1(x) > r_2(x)$ for all x, this suggests that person 1 is more risk averse than person 2, and indeed, in this case, it can easily be proved that for any large risk $\tilde{x}$, $RP_1(\tilde{x}) > RP_2(\tilde{x})$.

The Pratt-Arrow measure of risk aversion completely summarizes the local risk preferences of an individual whose current wealth is *constant*. For someone with *stochastic* prior wealth, the characterization of local risk preferences is (seemingly) more complicated. If you have large prior stakes in the outcomes of events, and if complete markets for contingent claims on those events do not exist, it may not be feasible to eliminate all the risk that you face, even if you wished to do so. Under these conditions, the Pratt-Arrow definition of risk premium (i.e., the price you would pay to eliminate all the risk you face) is not particularly useful. Moreover, when assessing the price at which you would be willing to buy or sell an additional small risk $\tilde{z}$, you must (seemingly) consider not only the mean and variance of $\tilde{z}$, but also its *covariance* with your prior wealth. Various authors (e.g., Ross) have studied generalizations of the Pratt-Arrow results to such conditions, and typically they have assumed that new risks were to some extent statistically independent of prior wealth. (Later I will discuss a generalization of the Pratt-Arrow measure that finesses away these difficulties.)

3.2 Important parametric families of utility functions.

It seems reasonable that you should become *less* risk averse as your wealth level rises. Intuitively, a person who can afford to absorb larger losses should be willing to undertake greater risks. The richer you are, the less you should worry about any given amount of variance in your wealth, and the smaller the premium you should pay to be rid of it. Utility functions can be classified according to whether—and in what manner—risk aversion decreases as a function of wealth. The most important and widely used family of utility functions is the generalized power family, which includes the linear, log, and exponential functions as limits or special cases. There are several alternative (but preferentially equivalent) functional forms for power utility functions. The one that is most commonly used in the finance literature is the following:

$$u_\gamma(x) = \frac{x^{1-\gamma} - 1}{1 - \gamma}$$

Sometimes the minus-1 in the numerator is dropped. It is just an additive constant, so it makes no difference except when the limit at $\gamma = 1$ is evaluated. With the minus-1 in the numerator, this function conveniently converges to $\ln(x)$ at $\gamma = 1$, otherwise it is necessary to define $u_\gamma(x) = \ln(x)$ at $\gamma = 1$ as a special case. (Some authors also drop the $1-\gamma$ in the denominator, and some express the exponent directly as a power c.) The exponential utility function, for which (absolute) risk aversion is constant, is also defined to be a special case of the generalized power function, although it is not actually a limiting case of the expression above.

The Pratt-Arrow measure of absolute risk aversion determined by $u_\gamma(x)$ is $r_\gamma(x) \equiv -u_\gamma''(x)/u_\gamma'(x) = \gamma/x$, which is a linear function of γ and a hyperbolic function of x, hence this family of functions is commonly called the **HARA** family, which stands for **Hyperbolic Absolute Risk Aversion**. This rather awkward term seems to have been inflicted on us by finance theorists in the 1970's. Ingersoll suggests the more sensible term **LRT**, for **Linear Risk Tolerance**, because the corresponding *risk tolerance measure*, which is just the reciprocal of the Pratt-Arrow risk aversion measure, is $t_\gamma(x) = 1/r_\gamma(x) = x/\gamma$, and I will henceforth use both terms HARA and LRT interchangeably. The **Pratt-Arrow measure of relative risk aversion** is defined as $-xu_\gamma''(x)/u_\gamma'(x) = xr(x)$, which is equal to the constant γ at all levels of wealth for HARA/LRT utility functions as parameterized above, so this family also has the property of **constant relative risk aversion (CRRA)**. Equivalently, it can be said to have the property of **constant relative risk tolerance**, because in this case risk tolerance is a characteristic fraction of wealth.

> In general, a utility function with the property of linear risk tolerance need not also exhibit CRRA. Linear risk tolerance means that risk tolerance is a linear (affine) function of wealth, i.e., $t(x) = a + bx$ for some constants a and b with $b > 0$. The CRRA property is satisfied only if $a = 0$, i.e., if risk tolerance equals zero when wealth equals zero, whatever it means to have zero wealth! Risk tolerance is observable: it is uniquely determined by preferences and can be elicited by asking simple questions regarding changes in wealth relative to the status quo. The zero point of someone's wealth, on the other hand, is not always observable, nor is it even necessarily a well-defined quantity outside of abstract models. What would you consider to be the zero point of *your* wealth? Is it the boundary between a positive and negative balance in your bank

account, or is the state you would be in if you gave away all your worldly possessions, or is it the deepest state of indebtedness you could incur before your loan shark would kill you? In applications, the fictional zero point of a person's wealth is never closely approached, and the CRRA property is used more to characterize cross-sectional phenomena, i.e., how local risk tolerance varies among agents with very different status quo levels of wealth. See the discussion of Howard's the "one-sixth rule" below.

The HARA/LRT utility function $u_\gamma(x)$, as parameterized above, is only defined for positive x, where x is interpreted to be the decision maker's total wealth, as just noted. The wealth level $x = 1$ (in whatever monetary units are used) serves as a reference point for choosing the origin and units of the utility scale: $u_\gamma(x)$ is scaled so that $u_\gamma(1) = 0$, $u_\gamma'(1) = 1$, and $r_\gamma(1) = \gamma$. Note that zero wealth does not mean zero utility under this parameterization: that is merely an artifact of the minus-1 in the numerator. The zero point of utility is irrelevant and can be chosen arbitrarily, because Bernoulli utility functions are only unique up to positive affine transformations.

The HARA/LRT class of utility functions includes the **linear** utility function (for which risk tolerance is infinite), the **exponential** utility function (for which risk tolerance is constant), the **logarithmic** utility (for which risk tolerance is a linear function of wealth with a slope of 1), as well as the **reciprocal** utility function, the **square root** utility function, and the **quadratic** utility function (for which $\gamma = -1$). These special cases are summarized in the following table and chart:

Table 3.1: Special cases of the HARA/LRT family of utility functions

Relative risk aversion	Functional form	$u_\gamma(x) =$
$\gamma = 0$	Linear	$x - 1$
$\gamma = \frac{1}{2}$	Square root	$2(\sqrt{x} - 1)$
$\gamma = 1$	Logarithmic	$\ln(x)$
$\gamma = 2,$	Reciprocal	$1 - 1/x$
$\gamma = \infty$	Exponential	$1 - \exp(1 - x)$ *

* $u_\gamma(x)$ actually converges to $1/\gamma$ as $\gamma \to \infty$, but it is defined to be an exponential utility function at $\gamma = \infty$. The particular exponential form $u(x) = 1 - \exp(1-x)$ is shown here because it has the same origin and scale properties $u(1) = 0$ and $u'(1) = 1$ that are satisfied by $u_\gamma(x)$ for $\gamma < \infty$. The conventional HARA/LRT functional form also has to be patched up for negative γ, but this case is no longer taken seriously in finance applications because it implies that risk aversion increases rather than decreases with increasing wealth. If the power utility function is parameterized in terms of $\beta = 1/\gamma$ instead, it can be affinely scaled so that it is valid for all real values of β as well as their limits, as shown in my recent paper "Scoring Rules, Generalized Entropy, and Utility Maximization" with Bob Winkler and Victor Jose. In this version, the exponential utility function falls in the center of the spectrum (at $\beta = 0$), and it forms the boundary between increasing-risk-averse utility and decreasingly-risk-averse utility.

Figure 3.1: Linear-Risk-Tolerance utility functions
(Gamma = risk aversion coefficient)
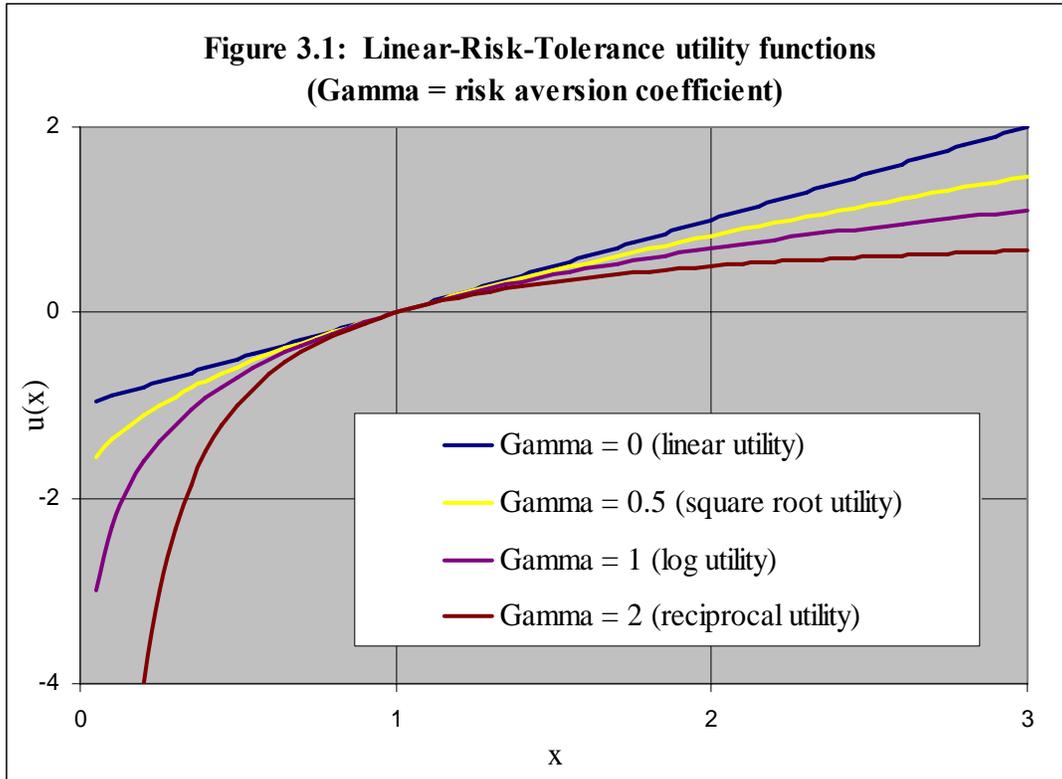
The general exponential utility function can be written as $u(x) = -\exp(-x/t)$ where t is a positive constant, and it has the property of **constant absolute risk aversion (CARA)**, or equivalently, **constant risk tolerance**, because risk tolerance is equal to t for all x under this function. If your utility function were really exponential, you would not change your portfolio of investments even if, through extraordinary good or bad luck, you suddenly became enormously rich or deeply in debt. (In the latter case you would borrow to finance the same investments as before—assuming anyone would loan you the money!) As such, the exponential utility function doesn't seem very suitable for representing risk attitudes over a potentially wide range of final wealth levels. But it is very suitable as a *local approximation* to anyone's utility function for evaluating small to moderate gambles, and for this reason it is widely used in applied decision analysis. On the assumption that you have an exponential utility function, your risk tolerance parameter can be elicited by merely asking the following question: "for what value of t would you be just willing to accept a gamble that offers a 50% chance of winning $\$t$ and a 50% chance of losing $\$ \frac{1}{2} t$?" The value of t that makes you just indifferent to accepting this gamble is approximately, but almost exactly, your risk tolerance. This rule of thumb is obtained as follows: suppose your utility function really is exponential with risk tolerance t, and suppose (very importantly) that your current wealth position is riskless. Then your current level of wealth, whatever it is, has no bearing on your preferences among gambles. So, without loss of generality, assume you currently have wealth 0, which yields a utility of $-\exp(0) = -1$. Now, a 50-50 gamble between winning t or losing $\frac{1}{2}t$ yields an expected utility of $-\frac{1}{2}(\exp(-1) + \exp(\frac{1}{2}))$ $= -\frac{1}{2}(0.368 + 1.649) = -1.008 \approx -1$, almost exactly the same as utility of not gambling.

Typically, individuals (whether acting on their own behalf or as corporate managers) tend to display a risk tolerance that is a characteristic fraction of the wealth at their disposal—i.e., it depends on what they can "afford to lose." Informal studies by Professor Ronald Howard at Stanford University—one of the pioneers of Bayesian methods of decision analysis—suggest that many subjects use a "one-sixth" rule: their risk tolerance is roughly one-sixth of their personal or corporate equity, which is consistent with the hypothesis that the CRRA property applies cross-sectionally (with $\gamma \approx 6$ in this case) even though individual decision makers may have CARA preferences within their own frames of action. (For more details, see the paper "Decision Analysis: Practice and Promise," *Management Science* 1988).

The attraction of the exponential utility function is that it is *extremely* convenient for purposes of mathematical modeling. For one thing, it is the only utility function that satisfies the so-called "delta property": if all the outcomes of a gamble are increased in value by the same amount Δ (which need not be small), then your certainty equivalent for the gamble is also increased by exactly the amount Δ. That is:

$$CE(\tilde{x} + \Delta) = CE(\tilde{x}) + \Delta$$

This is very handy when modeling multistage decision problems.

The exponential utility function is also the natural "conjugate" utility function for the normal probability distribution. If the agents in your economic model can be assumed to have normal probability distributions and exponential utility functions, your life as a theorist is vastly simplified. For one thing, under these conditions the formula for evaluating the risk premium of a gamble holds in-the-large as well as in-the-small. That is, under exponential utility, for any *normally distributed* wealth level $\tilde{x}$ with mean μ and variance σ^2:

$$CE(\tilde{x}) = \mu - \tfrac{1}{2} r \sigma^2 .$$

So, with normal probability and exponential utility, you are firmly in the realm of "mean-variance preferences." There is a long history in the finance literature of trying to model preferences for risky assets terms of *only* the means and variances of asset returns. But it is now known that preferences that satisfy the von Neumann-Morgenstern or Savage axioms are determined by means and variances alone only under fairly restrictive conditions: normal probability distributions or quadratic utility functions. (No one seriously advocates quadratic utility functions any more—they have the counterintuitive property of increasing risk aversion and they satiate at a finite level of wealth.) If normal probability distributions *and* exponential utility functions are assumed, then preferences are determined by a simple linear tradeoff between mean and variance as the above equation shows.

A more profound manifestation of the conjugacy between normal distributions and exponential utility functions is that the **risk neutral probability distribution** of an agent who has a normal probability distribution over states of the world and who has an exponential utility function, and whose wealth is a linear or quadratic function of the state variable, is again a normal distribution. (An agent's risk neutral probability distribution is the product of her probabilities and her state-

dependent relative marginal utilities for money. Much more will be said about this subject in section 3.4 below and in subsequent classes.) Hence, if the agents in your economic model face normally distributed risks, it is very convenient (and perhaps even normatively appropriate) to endow them all with exponential utilities.

Finally, it is only under the assumption of exponential utility that, in a complete market for contingent claims, a group of agents with heterogeneous beliefs and risk preferences behaves like a single "representative agent" whose probabilities and utilities are uniquely-determined composites of those of the individuals. Hence, only under the assumption of exponential utility can we *aggregate* the diverse beliefs and risk preferences of many individuals in a convenient and unambiguous way.   In particular, the risk tolerance of the representative agent is just the sum of the risk tolerances of the individual agents, and the probability distribution of the representative agent is a geometric weighted average of those of the individuals, with weights proportional to their risk tolerances. The aggregation problem is discussed in papers by Wilson (1968) and Rubinstein (1974).   Aggregation of multivariate normal probability distributions under exponential utility, with applications to the Capital Asset Pricing Model, is discussed in "Arbitrage, Rationality, and Equilibrium" by Nau and McCardle (1991).   The intuitive explanation of this aggregation property is that because risk tolerance does not depend on wealth under exponential utility, *the aggregate behavior of a group of agents with exponential utility, as seen from the perspective of an outside observer, does not depend on how they internally divide the surplus from trade.*

Recall from our discussion of consumer theory that the problem of determining the equilibrium prices that emerge from trade generally does not have a unique solution: different equilibrium prices will be observed at different Pareto optimal solutions, and the selection of one among many Pareto optimal solutions that dominate a given initial allocation typically depends on the amounts of surplus from trade that are captured by different agents—or by middlemen or arbitrageurs.   The Walrasian equilibrium concept, which attempts to force a unique solution through a fixed price schedule, is a high-handed way of side-stepping this fundamental indeterminacy. However, under exponential utility, the indeterminacy of the division of surplus doesn't result in indeterminacy of the marginal prices that the group is willing to pay for risky assets

The aggregation problem can also be solved with more general LRT utility functions, but only under strong restrictions on beliefs, tastes, and/or wealth. First of all, the agents must all have essentially the "same" utility function within this class—i.e., they must all have the same relative risk aversion coefficient (γ).   This condition ensures that, although the internal division of the surplus from trade may change the risk tolerances of individual agents, it does not change the *aggregate* risk tolerance of the group. Additionally, under all LRT utility functions other than logarithmic utility, they must also have identical beliefs—which is nonsense from a subjective point of view—while under log utility they must have the same stochastic prior wealth (at least up to a multiplicative constant), as if they all hold the same market portfolio of risky assets, which is also rather limiting. The aggregation property is desirable from a modeling perspective because it is much easier to model an economy consisting of one representative agent than one populated by a million idiosyncratic agents. But alas, the homogeneity conditions needed to

obtain a unique aggregation formula are seen to be very severe. Maybe nature is trying to tell us something by this.

The logarithmic utility function and other LRT utility functions with positive but finite γ nevertheless have their own compelling logic. Within the LRT family, they are the only functions that are *decreasingly risk averse,* a sensible property for modeling decision problems involving large variations in wealth. And the CRRA property conveniently implies that an individual always makes the same *proportional* investments in risky assets—i.e., she always takes the same "relative risks" regardless of her base level of wealth.

Whereas the exponential utility function jibes smoothly with normal probability distributions, the logarithmic and power utility functions jibe smoothly with *lognormal* probability distributions, which are often used to model movements in stock prices. (If the short-term returns on stocks are independently and identically distributed, then the longer-term returns, which are the cumulative product of the short-term returns, are approximately lognormally distributed—this is the multiplicative version of the central limit theorem.) Indeed, the logarithmic and power utility functions are naturally conjugate to the lognormal distribution in the same way that the exponential utility function is conjugate to the normal distribution: the resulting risk neutral distributions are still lognormal. The Black-Scholes options pricing mode, which we shall discuss in a later class, is implicitly based on an assumption of lognormal risk-neutral probabilities. Alas, the lognormal family of distributions is not closed under the formation of portfolios—a portfolio of lognormally distributed stocks generally does *not* have a lognormal distribution of returns—which is an inconvenience for finance theory. According to the central limit theorem, the return on a large portfolio of different stocks might be expected to have a normal distribution rather than a lognormal distribution. Empirically it does seem to be true that returns on individual stocks are more positively skewed than returns on large mutual funds or the market as a whole.

Exponential, logarithmic, and more general LRT utility functions are all widely used because of their convenient properties, although they are also widely criticized as being insufficient to explain various stylized facts about markets under uncertainty, particularly where preferences for both risk and intertemporal substitution are involved. Different authors have presented normative or empirical arguments in favor of different values for the relative risk aversion coefficient (γ) of a representative agent. In a paper entitled "The strong case for the generalized logarithmic utility function as the premier model of financial markets," Rubinstein (1976) argues in favor of the logarithmic utility function ($\gamma=1$) because it at least permits the aggregation problem to be solved with heterogeneous beliefs (though not heterogeneous prior wealth). The logarithmic utility function also crops up in other fields. In quantitative psychology, it is related to "Fechner's law," which states that sensitivity to changes in a stimulus are typically inversely proportional to the current level of the stimulus. In statistical decision theory, it is supported by the "Kelly criterion," according to which the optimal way to increase your wealth over time when facing a series of repeated gambles is to maximize the expected value of the logarithm of your wealth at every stage. The logarithmic utility function also has connections with Shannon's negative-entropy measure that is central to information theory (but more general LRT utility functions are associated with more general entropy measures that are also used on other fields). Thus, Daniel Bernoulli has finally found some measure of vindication after two centuries of

criticism by economists for his seemingly arbitrary adoption of a logarithmic utility function. Arrow (1965) has argued on theoretical grounds that agents should exhibit decreasing absolute risk aversion but *increasing* relative risk aversion (IRRA), rather than the CRRA proper of HARA/LRT utility functions. And Prescott and Mehra (1985, 2003) have shown that the "equity premium puzzle" in financial markets (the surprisingly large gap between average rates of return on stocks and bonds over long periods) is more suggestive of a coefficient of relative risk aversion that is greater 30, approaching an exponential utility function. The latter finding is not necessarily taken as an argument for using LRT utility functions with pathologically large relative risk aversion coefficients to model the behavior of agents in real markets. Rather, it suggests that simple Bernoulli utility functions do not suffice to explain the risk and time preferences of those agents. Extensions of the expected utility model to include features such as *non-expected-utility preferences*, *time-inseparability*, and *habit formation* have been widely explored as alternatives.

Another class of (non-LRT) utility functions that has gained some popularity—more in decision analysis than in economics—is the **sumex/linex** class, consisting of utility functions that are the sum of two exponential functions or (as a limiting case) the sum of an exponential function and a linear function. These functions have the important property of decreasing absolute risk aversion while retaining some of the convenience of the exponential form. The key property of the sum of two exponential functions is that the one with the lower risk tolerance determines risk preferences at very low levels of wealth while the one with the higher risk tolerance determines risk preferences at very high levels of wealth. The decision maker thus exhibits decreasing absolute risk aversion by virtue of gradually shifting between a lower bound and an upper bound on risk tolerance as wealth increases. David Bell's paper on "One-switch utility functions" shows that the linex (linear+exponential) utility function is the only decreasingly-risk-averse utility function satisfying an intuitively pleasing "one-switch" property, namely that your direction of preference between two fixed gambles should not switch more than once as your wealth level increases. The sumex (sum of two exponentials) utility function has often been used in decision analysis applications. Some applications of this function are described in the excerpt from Keeney and Raiffa's book, which reports that programs for fitting sumex utility functions were commonly used at Harvard Business School in the 1960's and 1970's—a poignant reminder of the vision that mathematical methods would soon revolutionize all of decision making. (Actually, this has eventually come to pass, but more through the impact of the electronic spreadsheet than the widespread adoption of utility theory. We still do teach utility theory to MBA's, but a single exponential utility function is typically used when risk attitudes are formally modeled at all.) Keeney and Raiffa's book also contains (on page 202) a table of handy "exponential transforms" that can be used for performing expected utility calculations with exponential utility functions (including sumex and linex as well as single exponential) in the presence of various probability distributions from the exponential family— not only the normal distribution, but also the exponential, gamma, Cauchy, binomial, geometric, and Poisson distributions. (Nowadays we would just do the expected-utility calculations by performing Monte Carlo simulation in Excel.)

3.3 Markets under uncertainty and state-preference theory

In the 1950's, Arrow and Debreu generalized the concept of a competitive equilibrium in an exchange economy to include preferences for risk and time as well as preferences for physical commodities. For this and their work on proving the existence of general competitive equilibria, they eventually received Nobel prizes. Arrow and Debreu extended the concept of a "commodity" to include not only the physical description of a good but also the time and/or the state of the world in which it is received. Thus, a unit of a commodity might not be just a bushel of corn, but rather a bushel of corn delivered on a particular date and/or contingent on the value of an uncertain economic variable (say, the price of corn) assuming a specified value. By trading time- and state-contingent claims to money and other commodities, agents are able to hedge their risks and smooth their consumption over time and also profit from differences in *beliefs* and differences in *risk tolerance* as well as differences in tastes for commodities. The term "Arrow-Debreu economy" generically refers to an exchange economy in which goods are time-stamped and/or state-contingent, and an **Arrow-Debreu security** is a (sometimes hypothetical) security that pays \$1 when a particular state of nature occurs on a specified date. An Arrow-Debreu security is essentially the same as a unit lottery ticket in de Finetti's operational definition of subjective probability, except that the payoff may occur at a future date, so that time preference also plays a role in determining how much you are willing to pay for the ticket today.

The Arrow-Debreu approach to modeling uncertainty via preferences for state-contingent commodities is known as **state-preference theory**, and it is consistent with, but potentially more general than, Savage's subjective expected utility theory. Consider the special case of a economy in which there are n mutually exclusive, collectively exhaustive, states of nature and a single divisible commodity (money) for which state-contingent claims can be traded. The wealth distribution of an individual can then be represented by an n-vector $\boldsymbol{w}$, whose j^{th} element w_j denotes the quantity of money received in state j. (Henceforth I will use boldface vector notation rather than "tilde" notation to denote a stochastic wealth position or a risky asset.) If the individual's preferences among wealth distributions satisfy the standard axioms of consumer theory, then they are represented by a plain old ordinal utility function $U(\boldsymbol{w})$. If no additional restrictions are placed on preferences, the individual is rational according to the usual standards of consumer theory, but she may have "non-expected utility preferences" in the sense that her valuation of a risky asset may not have an additively separable form in which utilities for consequences are multiplied by the probabilities of states in which they are received. For example, she may behave as if amounts of money received in different states are substitutes or complements for each other, which is forbidden under expected utility theory.

To cover these more general cases, Yaari (1969) introduced a broader definition of risk aversion, namely that **a decision maker is [strictly] risk averse if her preferences for state-contingent wealth are [strictly] convex.** Convexity of preference means that $\boldsymbol{x} \succcurlyeq \boldsymbol{z}$ and $\boldsymbol{y} \succcurlyeq \boldsymbol{z}$ implies $\alpha\boldsymbol{x} + (1-\alpha)\boldsymbol{y} \succcurlyeq \boldsymbol{z}$ for any α between 0 and 1, and strict convexity means that the last preference is strict if α is strictly between 0 and 1. This requires the decision maker's ordinal utility function to be [strictly] quasi-concave,. Recall that this is the same assumption that is commonly used in modern consumer theory to generalize the concept of diminishing marginal utility for commodities. When the "commodities" are state-contingent amounts of money, convexity implies that she has diminishing marginal utility for all risky assets. Thus, the basic concept of

aversion to risk does not presuppose that the individual is an expected-utility maximizer, nor does it require any references to "true" probabilities or riskless wealth states—about which more will be said below.

If state-preferences are additionally assumed to satisfy the independence condition (Savage's P2), then $U(\mathbf{w})$ has an additively separable representation:

$$U(\mathbf{w}) = v_1(w_1) + v_2(w_2) + \ldots + v_n(w_n).$$

(The marginalist economists assumed that preferences among commodities such as apples and bananas had an additively separable representation like this, which rules out complementarity or substitutability.)   This representation of preferences can be interpreted as **state-dependent expected utility without unique determination of the probabilities**, because we can choose any old probabilities $(p_1, \ldots, p_n)$ and then define correspondingly scaled state-dependent utility functions $u_i(w_i) = v_i(w_i)/p_i$ to obtain the equivalent representation:

$$U(\mathbf{w}) = p_1 u_1(w_1) + p_2 u_2(w_2) + \ldots + p_n u_n(w_n).$$

The idea that utilities might be state-dependent is a natural one and is useful for many applications involving potentially life-changing events. However, it leads to the difficulty just noted, namely that *scale factors for the state-dependent utilities get all mixed up with the probabilities*, so that there is no way to behaviorally distinguish differences in subjective probabilities between states from differences in utility scale factors between states. Many authors have made heroic efforts over the years to try to uniquely separate probabilities from utilities when the utilities are state-dependent, and this topic remains controversial. We'll see later that the problem of separating probabilities from state-dependent utilities can be finessed by not worrying about "true" probabilities and instead doing the analysis in terms of observable "risk-neutral" probabilities.

If preferences are further assumed to be conditionally state-independent (the equivalent of Savage's P3), then $U(\mathbf{w})$ has a **state-independent expected utility** representation:

$$U(\mathbf{w}) = p_1 u(w_1) + p_2 u(w_2) + \ldots + p_n u(w_n),$$

where $u(x)$ is a univariate utility function for money, as in Pratt's model of risk aversion.
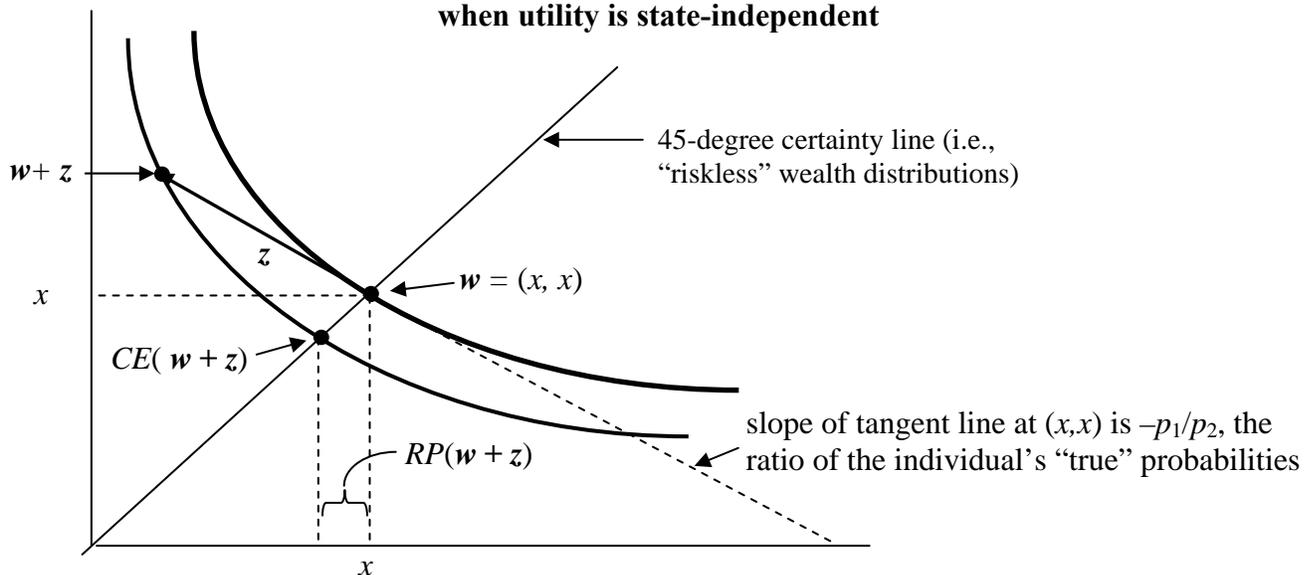
State-preference theory, like the consumer theory that it generalizes, lends itself to nice geometrical arguments and illustrations in the case of two dimensions. If the decision maker is assumed to have convex state-preferences (i.e., if she is risk averse), then her indifference curves for state-contingent wealth are convex. If she is assumed to be a state-independent expected utility maximizer (i.e., to also satisfy P2 and P3), then her indifference curves have other special properties.   Recall that, in 2-dimensional consumer theory, the slope of a tangent line to an indifference curve is (minus) the ratio of marginal utilities of the two commodities. By differentiating the preceding formula for $U(\mathbf{w})$, the marginal utilities of a state-independent expected-utility maximizer are seen to be

$$\partial U(\boldsymbol{w})/\partial w_j = p_j\, u'(w_j),$$

where $u'(w_j)$ is the derivative of the utility function for money, i.e., the marginal utility of money, at wealth level w_j.   Hence, **the marginal utility of wealth in state j is proportional to the product of probability and relative marginal utility for money in that state.**   The slope of the tangent line passing through the indifference curve at $(w_1,\ w_2)$ is therefore equal to $-p_1 u'(w_1)/p_2 u'(w_2)$.

On a 2-dimensional plot of w_2 versus w_1, the 45-degree line through the origin represents the set of "certain" wealth distributions, hence it is commonly called the **45-degree certainty line**. Where an indifference curve intersects this line, the marginal utilities of money are the same in both states, so its slope must be equal to $-p_1/p_2$, the (negative) ratio of the decision maker's subjective probabilities.   The Pratt-Arrow risk premium is easily illustrated on such a plot: suppose the decision maker initially has constant wealth $\boldsymbol{w} = (x,\ x)$ and acquires a small risk $\boldsymbol{z} = (z_1,\ z_2)$ which is actuarially neutral   (i.e., has zero expected value).   Because $\boldsymbol{z}$ is actuarially neutral, its direction is parallel to the tangent to the indifference curve at $(x,\ x)$.   If the decision maker is risk averse (i.e., has convex preferences), the new wealth position $\boldsymbol{w}+\boldsymbol{z}$ must lie on a lower indifference curve.   The point at which this lower indifference curve intersects the 45-degree certainty line is the certainty equivalent (CE) of $\boldsymbol{w}+\boldsymbol{z}$, and the absolute difference between the latter quantity and the original wealth x is the risk premium (RP) associated with $\boldsymbol{z}$.

Figure 3.2
Construction of the certainty equivalent and risk premium
when utility is state-independent

45-degree certainty line (i.e., "riskless" wealth distributions)

$\boldsymbol{w}+\boldsymbol{z}$

$\boldsymbol{z}$

x

$\boldsymbol{w}=(x,\ x)$

$CE(\boldsymbol{w}+\boldsymbol{z})$

$RP(\boldsymbol{w}+\boldsymbol{z})$

slope of tangent line at (x,x) is $-p_1/p_2$, the ratio of the individual's "true" probabilities

x

3.4 Risk neutral probabilities

The decision maker's relative marginal utilities, $\{\partial U(\boldsymbol{w})/\partial w_j\}$, play a central role in state-preference theory, just as they do in consumer theory. It is convenient to normalize the relative marginal utilities so that they sum to 1, yielding a probability distribution $\boldsymbol{\pi}(\boldsymbol{w})$ defined by:

$$\pi_j(\boldsymbol{w}) = \frac{\partial U(\boldsymbol{w})/\partial w_j}{\displaystyle\sum_{k=1}^{n}\left(\partial U(\boldsymbol{w})/\partial w_k\right)}$$

$\boldsymbol{\pi}(\boldsymbol{w})$ is called the decision maker's **risk neutral probability distribution** at wealth $\boldsymbol{w}$ because she prices very small assets in a seemingly risk-neutral manner with respect to it: a small gamble which has zero expectation with respect to the probability distribution $\boldsymbol{\pi}(\boldsymbol{w})$ yields zero change in marginal utility by construction, and hence it should be just barely acceptable. (Some authors refer to risk neutral probabilities as "risk-adjusted probabilities" or "martingale probabilities" or "utilprobs.") In the two-dimensional case, the slope of the tangent line to the indifference curve passing through the decision maker's current wealth position $\boldsymbol{w}$ is equal to $-\pi_1(\boldsymbol{w})/\pi_2(\boldsymbol{w})$, the (negative) ratio of the decision maker's risk neutral probabilities of the two states. For a decision maker who is a state-independent expected-utility maximizer, risk neutral probabilities are proportional to the product of her true subjective probabilities and her relative marginal utilities for money at her current wealth position, i.e.,

$$\pi_j(\boldsymbol{w}) \propto p_j u'(w_j).$$

If an attempt is made to elicit the decision maker's subjective probabilities by de Finetti's method—i.e., by asking which gambles she is marginally willing to accept—the probabilities that are observed will be her risk neutral probabilities rather than her true probabilities. The two distributions will differ if the decision maker has sufficiently large prior stakes in the outcomes of events to affect her marginal utilities for money. Risk neutral distributions play a central role in models of markets under uncertainty—either implicitly or explicitly—and as we proceed through the course, I will argue that they should play a central role in rational choice theory more generally.

If the decision maker is assumed to be embedded in a complete market for contingent claims—i.e., a market where prices exist for all Arrow-Debreu securities—the fundamental theorems of consumer theory can be immediately re-expressed in terms of risk neutral probabilities. For example, recall that in the solution to the consumer's optimization problem under a budget constraint, the decision maker's marginal rates of substitution between commodities must equal the corresponding ratios of prices in any optimal solution. Similarly, in a market under uncertainty, the decision maker should trade risky assets to arrive at a position in which the ratios of her risk neutral probabilities equal the corresponding ratios of prices for Arrow-Debreu securities. That is, in any optimal solution, it must be true that
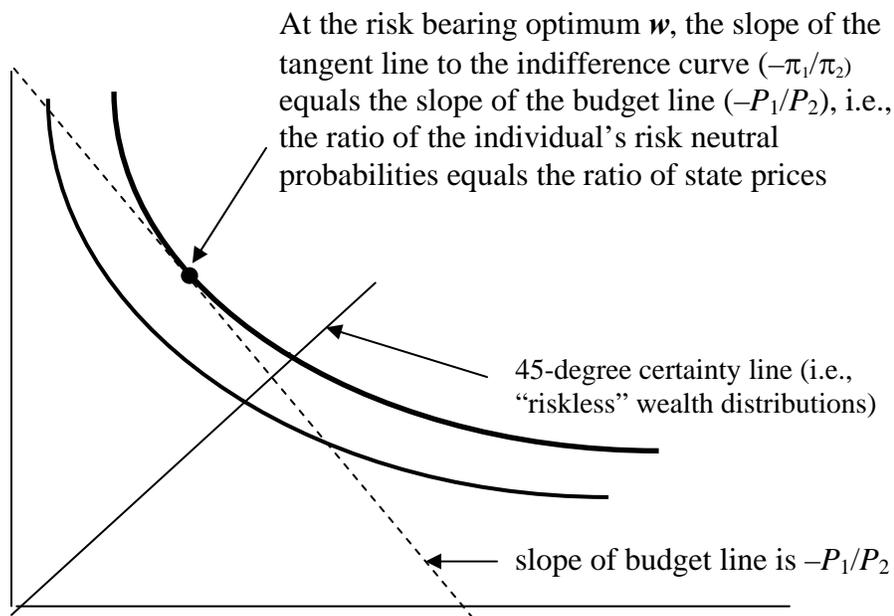
$$\pi_j(\boldsymbol{w})/\pi_k(\boldsymbol{w}) = P_j/P_k$$

for all states j and k, or equivalently

$$\pi_1(\boldsymbol{w})/P_1 = \pi_2(\boldsymbol{w})/P_2 = \ldots = \pi_n(\boldsymbol{w})/P_n,$$

where P_j is the price of a security that pays \$1 in state j. This result, which follows from the principle of equi-marginal utility, was christened the **fundamental theorem of risk bearing** by Arrow. In two dimensions, the picture looks like the one in Figure 3.3, which is essentially the same as Figure 1.3 in the notes from class 1, which was used to illustrate optimization under a budget constraint:

Figure 3.3
The fundamental theorem of risk bearing

At the risk bearing optimum $\boldsymbol{w}$, the slope of the tangent line to the indifference curve ($-\pi_1/\pi_2$) equals the slope of the budget line ($-P_1/P_2$), i.e., the ratio of the individual's risk neutral probabilities equals the ratio of state prices

45-degree certainty line (i.e., "riskless" wealth distributions)

slope of budget line is $-P_1/P_2$

In a market under uncertainty with many heterogeneous agents who are all subjective expected utility maximizers, the risk neutral probabilities of all agents must be the same in equilibrium, even though their true probabilities may differ. Indeed, their true probabilities will be disguised to some extent by the distortions caused by their heterogeneous state-dependent marginal utilities, hence they may "agree to disagree" about their true probabilities even as they agree about prices of contingent claims. This phenomenon is discussed in more detail in a 1995 paper of mine entitled "The Incoherence of Agreeing to Disagree."

3.5 Generalization of the Pratt-Arrow measure to state preferences without separation of probability from utility

It was pointed out above that state-preference theory encompasses all the standard results of subjective expected utility theory as they apply to decisions with monetary outcomes, provided the decision maker is assumed to satisfy Savage's P2 (independence of common consequences, a.k.a. the sure-thing principle) and P3 (state-independence of preferences for consequences). The beauty of state-preference theory is that is works even when P2 and P3 are not assumed, and as such it applies to situations in which (i) utility may be state-dependent, (ii) subjective probabilities may not be uniquely separable from utilities, (iii) the decision maker's prior wealth may be unobservable (in which case the location of the 45-degree certainty line is unknown) and/or (iv) the decision maker may have non-expected utility preferences. In the coming weeks, we will discuss these issues in more detail. In this section, I will show that it is possible to use the state-preference approach to recast the Pratt-Arrow measure of risk aversion in terms that do not refer to the "true" probabilities of events nor to any "riskless" states of wealth.

First, it is necessary to appropriately generalize the definition of a risk premium. The usual (Pratt) definition of a risk premium is that of a "selling" (a.k.a. "equivalent") premium. That is, it is the premium you would pay, in expected-value terms, to dispose of an existing risk. This is a useful definition in situations where it is possible to imagine selling off all the risk that you currently bear. Under more general conditions, where it may not be possible to unload all the risk (and where it may not be obvious even what a riskless state of wealth would look like), it is more natural to use the notion of a "buying" (a.k.a. "compensating") risk premium. The buying risk premium is the premium *you would have to be paid*, in expected-value terms, to take on a new risk in addition to any (possibly-correlated) background risk that you currently face.

More precisely, the decision maker's *buying price* for an asset z at wealth w, denoted $B(z; w)$, is determined by $U(w+z-B(z; w)) = U(w)$, while her *selling price* $C(z; w)$, otherwise known as her *certainty equivalent* for z, satisfies $U(w+C(z; w)) = U(w+z)$. The buying and selling prices are generally similar, but not identical, and they are related by $C(z; w) = -B(-z; w+z)$. The *buying risk premium* associated with z at wealth w, here denoted $b(z; w)$, is defined as the difference between the asset's risk neutral expected value and its buying price: $b(z; w) = E_{\pi(w)}[z] - B(z; w)$. The *selling risk premium* is similarly defined by $c(z; w) = E_{\pi(w)}[z] - C(z; w)$. (See Figure 3.4 below.) Note that the risk neutral probability distribution rather than the "true" probability distribution has been used in these definitions, for generality. In the special case where the decision maker is an expected utility maximizer with state-independent utility and where w is a riskless (i.e., non-stochastic) wealth position, risk neutral probabilities and "true" probabilities are identical, and Pratt's risk premium is the special case of the selling risk premium that obtains under these conditions. But under more general conditions, true probabilities may be nonexistent or unobservable, and in any case, what matters for purposes of marginal analysis are the decision maker's local betting rates for money.

The advantage of focusing on the buying rather than selling risk premium is that the buying risk premium is more well-behaved and has a direct connection with definition of risk aversion in terms of convex preferences, namely:

Theorem 1: *The decision maker is risk averse (according to Yaari's definition) if and only if her buying risk premium is non-negative for every asset regardless of her prior wealth distribution.*

We are now in a position to generalize Pratt's formula for the risk premium of a small neutral asset z. Intuitively, the risk premium depends on "how risk averse" the decision maker is in the vicinity of her current wealth, which in turn depends on "how curvy" her indifference curves are. So, we should expect the risk premium to depend on some measure of the local curvature of the indifference curves. (Recall that the familiar Pratt-Arrow measure, $r(x)$, is a measure of the relative curvature of the utility function for money at wealth x.) In the general state-preference framework, the slope of the difference curve at wealth w is measured by the risk neutral probability distribution $\pi(w)$. Now, geometrically, the curvature of the indifference curves must have something to do with the rate of change of their slopes, which are measured by the *derivatives of the risk neutral probabilities*, and indeed, this turns out to be exactly the measure of curvature that is needed to generalize Pratt's local approximation formula.

The derivatives of the risk neutral probabilities constitute a matrix, which will be denoted $D\pi(w)$, whose jk^{th} element is

$$D\pi_{jk}(w) = \partial\pi_j(w)/\partial w_k$$

The risk premium formula is then generalized as follows:

Theorem 2: *The buying risk premium of a small neutral asset z at wealth w satisfies*

$$b(z; w) \approx -\tfrac{1}{2} z \cdot D\pi(w) z.$$

Here a "neutral" asset means an asset whose risk-neutral expectation is zero. Thus, in the n-dimensional state-preference framework, the expression $r(x)E[z^2]$ in Pratt's risk premium formula for a neutral asset is replaced by the more general matrix expression $- z \cdot D\pi(w) z$. (This is a quadratic form in which the matrix $D\pi(w)$ is pre- and post-multiplied by the vector z.) Evidently $D\pi(w)$ encodes both the decision maker's beliefs and local risk preferences, and indeed it can be factored into a product of two matrices, one of which contains the decision maker's risk neutral probabilities and the other of which is constructed from ratios of second and first derivatives of the ordinal utility function, generalizing the Pratt-Arrow measure. To show this, define the local *risk aversion matrix* as the matrix $R(w)$ whose jk^{th} element is the following ratio of second to first derivatives:

$$r_{jk}(w) = -(\partial^2 U(w)/\partial w_j w_k)/(\partial U(w)/\partial w_j).$$

Under expected-utility preferences, $R(w)$ would be an observable diagonal matrix, and with state-independent utility and riskless prior wealth, its diagonal elements would be $r_{jj}(w) = r(x)$

for every j, as will be discussed in more detail below. But under general preferences, $R(w)$ is neither a diagonal matrix nor is it observable, since it is not invariant to monotonic transformations of U. In particular, if $\hat{U}(w) = f(U(w))$, where f is monotonic and twice differentiable, then $\hat{U}$ represents the same preferences as U, but the corresponding risk aversion matrix $\hat{R}(w)$ differs from $R(w)$ by an additive constant in each column. To eliminate the arbitrary constants, let the *normalized risk aversion matrix* $\overline{R}(w)$ be defined as the matrix whose jk^{th} element is

$$\overline{r}_{jk}(w) = r_{jk}(w) - E_{\pi(w)}[r_k(w)],$$

where $r_k(w)$ denotes the k^{th} column of $R(w)$. In other words, $\overline{R}(w)$ is obtained from $R(w)$ by subtracting a constant from each column that is equal to the risk neutral expected value of the original column, so that every column of $\overline{R}(w)$ has a risk neutral expectation of zero. The normalized risk aversion matrix is invariant to monotonic transformations of U and is observable. It turns out that $D\pi(w)$ and $\overline{R}(w)$ are related by

$$D\pi(w) = -\Pi(w)\overline{R}(w),$$

where $\Pi(w) = \text{diag}(\pi)$. In other words, the jk^{th} element of $\overline{R}(w)$ is $(\partial\pi_j(w)/\partial w_k)/\pi_j(w)$, which is the *relative* rate of change of the risk neutral probability of state j as wealth increases in state k. In these terms, the risk premium formula in the preceding theorem can be rewritten as follows:

$$b(z; w) \approx \tfrac{1}{2} z \cdot \Pi(w)\overline{R}(w)z = \tfrac{1}{2} z \cdot \Pi(w) R(w)z$$

Comparison with Pratt's original formula (1) shows that, under general conditions, the "true" distribution in the risk premium formula is replaced by the local risk neutral distribution $\pi(w)$, while the scalar Pratt-Arrow measure $r(x)$ is replaced by the **matrix risk aversion measure** $R(w)$, or equivalently by its normalized, observable form $\overline{R}(w)$.

In the special case where the decision maker satisfies the independence condition, U is additively separable and its cross-derivatives are zero. Hence $R(w)$ becomes an observable diagonal matrix, namely $R(w) = \text{diag}(r(w))$, where $r(w)$ is **a vector-valued Pratt-Arrow measure of risk aversion** whose j^{th} element is

$$r_j(w) = -(\partial^2 U(w)/\partial w_j^2)/(\partial U(w)/\partial w_j) = -u_j''(w_j)/u_j'(w_j),$$

and u_j is the utility function for money in state j in an arbitrary state-dependent expected-utility representation. (The true probabilities and utility-scale factors conveniently drop out when r is computed.) The risk premium approximation formula can then be simplified:

$$b(z; w) \approx \tfrac{1}{2} E_{\pi(w)}[r(w) z^2]$$

19

Note that this is almost the same as Pratt's original formula, except that (i) the risk aversion measure is state-dependent and multiplies z^2 *inside* the expectation formula, and (ii) the risk neutral distribution rather than the "true" distribution is used in the expectation. This is the appropriate generalization of the risk premium formula to use when utility is state-dependent and/or prior wealth is stochastic. By allowing the risk aversion measure to vary with the state and by using the risk neutral distribution to take the expectation, we finesse both of those complications at once.

**Figure 3.4   Construction of buying and selling prices and risk premia**

The vector $\pi(w)$ is the risk neutral distribution, i.e., the normalized gradient of U, at wealth w. The dashed line is the tangent hyperplane at w, whose normal vector is $\pi(w)$. The marginal price of any asset at wealth w is its inner product with $\pi(w)$. Asset z is neutral at wealth w—it has a marginal price of zero because it lies in the tangent hyperplane—while z' is not neutral because it does not lie in the tangent hyperplane. $w + z$ and $w + z'$ lie on the same indifference curve, below that of w. The small dashed triangles are 45-degree right triangles with altitudes a, b, c, d. The prices and risk premia are:

Marginal price	Buying price	Selling price	Buying risk premium	Selling risk premium
$P(z; w) = 0$	$B(z; w) = -a$	$C(z; w) = -b$	$b(z; w) = a$	$c(z; w) = b$
$P(z'; w) = -c$	$B(z'; w) = -(c+d)$	$C(z'; w) = -b$	$b(z'; w) = d$	$c(z'; w) = b - c$

According to Theorem 2 , the buying risk premium satisfies:

$$a \approx -\tfrac{1}{2} z \bullet D\pi(w)z \;=\; \tfrac{1}{2} z \bullet \Pi(w)\, R(w)z \;=\; \tfrac{1}{2} z \bullet \Pi(w)\, \overline{R}(w)z,$$

where $D\pi(w)$ is the matrix of derivatives of the risk neutral probabilities, $\Pi(w)$ is the diagonal matrix whose diagonal elements are the (observable) risk neutral probabilities, $R(w)$ is the risk aversion matrix, and $\overline{R}(w)$ is the (observable) normalized risk aversion matrix.
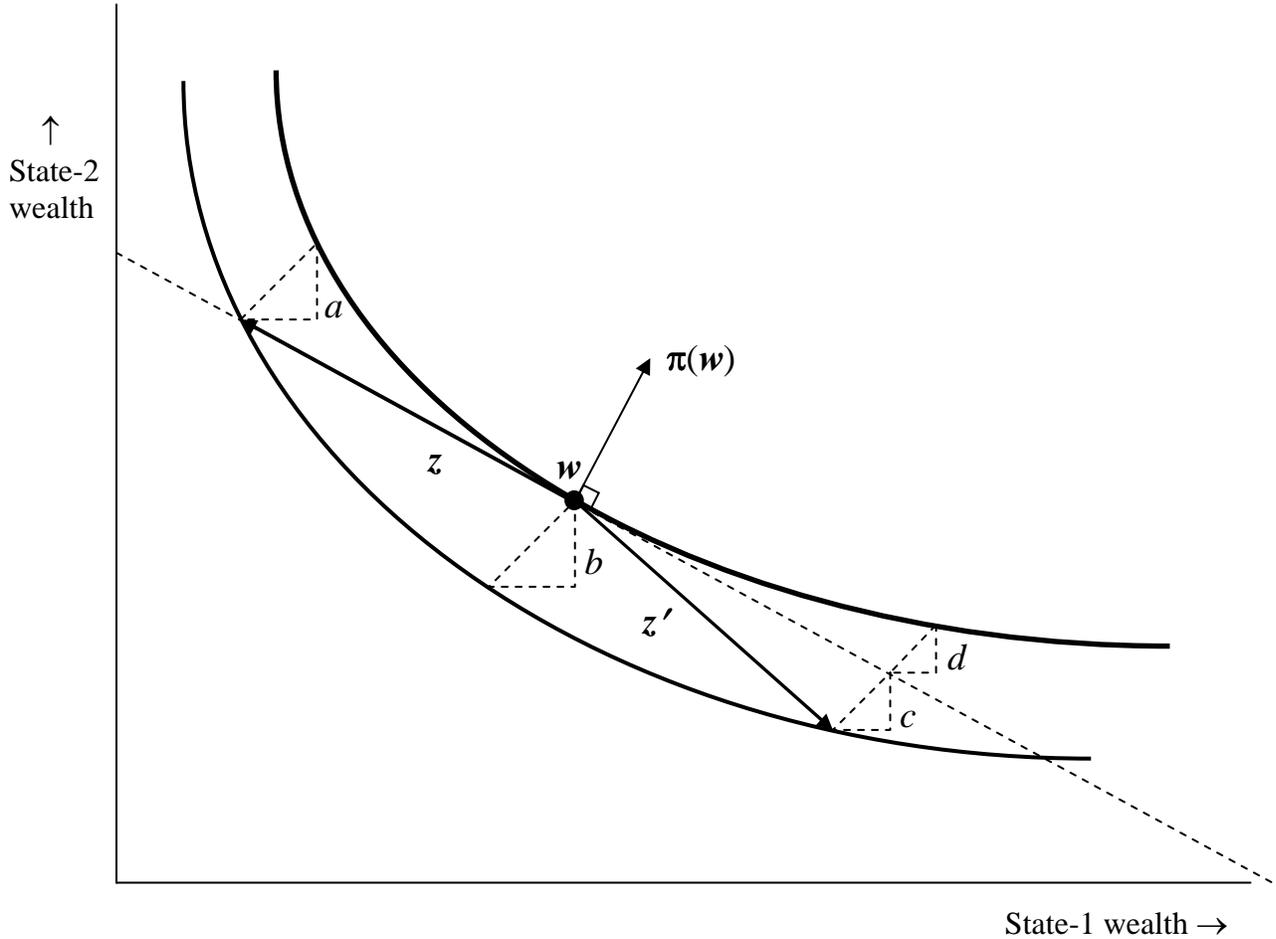
3.6 Summary of the various forms of the risk aversion measure and risk premium formula

1. When utilities are state-independent and prior wealth (x) is riskless, risk neutral probabilities are identical with "true" subjective probabilities, the risk aversion measure ($r(x)$) is *scalar-valued*, and Pratt's formula applies:

$$b(z; x) \approx \tfrac{1}{2} r(x) E[z^2]$$

2. When utilities are state-dependent and/or prior wealth (w) is stochastic, the risk aversion measure ($r(w)$) is *vector-valued*, risk neutral probabilities replace "true" probabilities in the expectation, and the risk premium formula becomes:

$$b(z; w) \approx \tfrac{1}{2} E_{\pi(w)}[r(w) z^2]$$

3. When the decision maker does not satisfy the independence condition (i.e., she has non-expected utility preferences), the risk aversion measure ($R(w)$ or its normalized, observable form $\overline{R}(w)$) is *matrix valued*, and the risk premium formula becomes:

$$b(z; w) \approx \tfrac{1}{2} z \bullet \Pi(w) \overline{R}(w)z = \tfrac{1}{2} z \bullet \Pi(w) R(w)z$$

Thus, regardless of whether the decision maker has expected-utility or non-expected utility preferences, and regardless of whether her utility is state-dependent or state-independent, and regardless of whether her prior wealth is stochastic or non-stochastic, her local beliefs and risk preferences can be conveniently summarized by an *observable risk-neutral probability distribution* and an *observable measure of local risk aversion*.

GUIDE TO THE READINGS

1a. "Risk aversion in the small and in the large" by John Pratt:

This classic paper—still very readable—covers the basic concepts of risk aversion.

1b. "On the role of securities in the optimal allocation of risk-bearing" by Kenneth Arrow (1953/1963/1974)

In an Arrow-Debreu economy in which commodities are time- and date-stamped, it seems at first glance as though an optimal allocation of resources would require trading of state-contingent claims on all commodities. For example, suppose there are two states (hot and cold) and three commodities (lemonade, tea, cocoa). We can imagine separate markets for lemonade-if-it's-hot, lemonade-if-it's-cold, tea-if-it's-hot, and so on. Different agents may have different preferences for commodities as a function of the state. (For example, some agents would rather have tea in cold weather rather than hot weather, and vice versa for others.) In general, if there are C different commodities and S different states, it seems we would need at total of $S{\times}C$ different markets in order to determine prices and make the trades necessary to achieve a competitive equilibrium allocation. This famous and oft-cited paper shows that it suffices to have a single "spot" market for each commodity plus a securities market in which state-contingent claims to money (i.e., Arrow-Debreu securities) are traded, for a total of $S{+}C$ markets. The key idea is that if agents can correctly anticipate the prices that will prevail on the spot markets as a function of the state that occurs, they can then execute trades in the securities market so as to yield exactly the amount of wealth in each state that is needed to buy their optimal allocation of commodities in that state.

On one level, this is a very profound result that seems to explain why we indeed have very well-organized securities markets and but relatively few markets for state-contingent commodities other than money. On another level, there is less to it than meets the eye. If you study it carefully, you will find that the proof depends on an extrapolation of the dubious assumption that Walrasian equilibrium prices can somehow be determined in the absence of trade. Arrow first describes how things would work if markets for state-contingent commodities existed: he assumes that an optimal allocation has been achieved, and he then invokes the second theorem of welfare economics to infer the existence of equilibrium prices supporting that allocation. (What this means, precisely, is that if those prices were posted, and if every agent were given a budget equal to the value of her allocation at those prices, then it would be optimal for every agent to buy that very allocation.) We have already seen that there is a very big question concerning how the equilibrium prices are actually determined. The Walrasian idea is that somehow equilibrium prices will be determined prior to trade, perhaps through the intervention of an auctioneer to whom everyone reports their preferences truthfully, but (as Kreps points out in his microeconomics text) this isn't very realistic. More plausibly, there will be a lot of haggling and trading at non-equilibrium prices on various markets en route to the final optimal allocation and equilibrium prices, and the final equilibrium allocation will not be uniquely determined by the initial allocations.

When Arrow moves on to consider the situation in which there is only a single securities market in addition to the spot commodity markets, he assumes that the prices that will eventually prevail on the spot markets, as a function of the state, will be the *same* prices that would have prevailed in the state-contingent commodities markets of the previous example. He then shows that it is possible to determine prices for securities such that every agent will be able to purchase the same, optimal, allocation in every state as she would have otherwise. But again, the price formation question is left unanswered, and now there are fewer markets on which to haggle. How will the agents form expectations about the prices that will prevail on the spot markets? (Recall Arrow's own observation from his 1987 paper: "Since no one has market power, no one sets prices; yet they are set and changed. There are no good answers to these questions, and I do not pursue them…") It is true that if the agents (somehow) have the same, correct expectations of equilibrium prices, then the securities market and the spot markets will enable them to make all the *trades* they need to make to allocate resources efficiently. But this way of framing the problem completely ignores the *informational* value of markets. You will often hear it said that this or that market is unnecessary because it merely replicates transactions that could be achieved by trading in several other existing markets. For example, it is sometimes argued that markets for stock options are "redundant" because you can achieve the same state-contingent payoff that a stock option would provide by continuous-time trading of the stock and a risk-free bond. But this overlooks the fact that an options market helps you discover the price that *other people* are willing to pay for the asset whose payoffs are generated in this way.

1c. Excerpt on "unidimensional utility theory" from *Decisions With Multiple Objectives: Preferences and Value Tradeoffs" by Keeney and Raiffa*

These pages provide more material on risk aversion and properties of utility functions—including a nice summary of properties of the commonly-used utility functions. The sections at the end provide examples of how utility functions are used in practice in decision analysis applications.

1d. Excerpt on "the fundamental theorem of risk bearing" from *The Analytics of Uncertainty and Information* by Jack Hirshleifer and John Riley (1992)

This chapter analyzes the problem of optimal risk-bearing by an individual in a market under uncertainty. It is noteworthy for including a concise statement of the fundamental theorem of risk bearing, which is just the generalization of the law of equi-marginal utility to the case of markets under uncertainty. This book is a good general reference on information economics.

1e. Excerpts on "utility theory" and "arbitrage and pricing" from *Theory of Financial Decision Making* by Jonathan Ingersoll (1987)

Ingersoll's book provides an accessible introduction to financial economics, and these excerpts indicate how decisions under uncertainty are treated in modern finance theory. The chapter on utility theory emphasizes the HARA/LRT class of utility functions and mentions the problem of modeling intertemporal consumption. This is followed by a chapter on "arbitrage and pricing" that makes almost no mention of utility whatever. This is typical of how subjective expected utility theory and arbitrage theory traditionally have been viewed as somewhat distinct subjects.

One of the themes of this course is that the two are deeply connected, and the former is in fact derivable from the latter. Hence, the chapter on arbitrage logically should come first! The chapter on arbitrage and pricing gives a short explanation of asset pricing in terms of risk neutral probabilities. The connection with subjective expected utility theory, which is not made explicit in this chapter, is that the market's risk neutral probabilities must coincide with those of the investors, where the latter are defined as probabilities of subjective probabilities and relative marginal utilities for money.

1f. Excerpt on "risk sharing and group decisions" by Howard Raiffa (from *Decision Analysis: Introductory Lectures on Choice Under Uncertainty*, 1968)

This chapter provides a much more readable treatment of Wilson's results on decision making by syndicates.

Secondary readings:

2a. "A generalization of Pratt-Arrow measure to non-expected-utility preferences and inseparable probability and utility" by Robert Nau, *Management* Scienc, August 2003. This paper contains more details of the generalized Pratt-Arrow measure discussed in the last section of the notes above, including an application to decision analysis. It is available on-line at http://faculty.fuqua.duke.edu/~rnau/genprattarrow.pdf

2b. "Evaluating income streams" by Jim Smith, *Management Science*, December 1998.
In intertemporal models of decision making, it is natural to define the primitive preferences of the decision maker with respect to streams of *consumption* over time. However, decision alternatives are usually defined in terms of *income* streams. Preferences over income streams typically do not satisfy the axioms of expected utility, because utility for income is an "indirect" utility based on the solution of an optimal consumption problem in the presence of opportunities to save and invest. To evaluate an income stream, the decision maker must solve a stochastic dynamic programming problem to determine the optimal consumption plan. This paper shows that with time-additive exponential utility, the DM can solve the stochastic dynamic programming problem to obtain certainty equivalents at each node in the tree. As the DM moves through the tree following her optimal consumption plan, she experiences positive or negative "windfalls"—i.e., her certainty equivalent turns out to be larger or smaller than its expected value when she reaches a node. When she experiences a windfall, the DM "shares" it with her "future selves" in proportion to their risk tolerances, through investments in the market. It is available on-line at http://faculty.fuqua.duke.edu/~jes9/bio/Evaluating_Income_Streams.pdf.

Other related reading (available on request)

3a. "The theory of syndicates" by Robert Wilson (1968)

This is an important paper concerning the implications of expected utility theory for group decision making, but alas it is very tough sledding. The author uses a lot of idiosyncratic tools and terminology, and the narrative has been pared down to the point where it is hard to follow his tracks. (This is a more typical *Econometrica* paper—not for the uninitiated!) The problem

considered in the paper is this: suppose that a group of individuals forms a syndicate which must make a single decision under uncertainty on behalf of all its members, and the proceeds of the decision will somehow be shared. What decision should be made and how should the proceeds be shared? And under what conditions will the syndicate behave as though it were a single rational individual? The paper shows that, in general, this is a messy problem. The group behaves like an individual only if the agents either have identical probability distributions or else have nearly identical utility functions. (More precisely, the latter condition is that the derivative of the Arrow-Pratt risk aversion function, which is called the "cautiousness" function, must be identical for all agents.) Moreover, the rule for sharing the proceeds is linear and "determinate" (dependent only on the amount to be divided, not how it was obtained, and not otherwise randomized) only if the agents all have exponential utility functions. (A much more readable treatment of this material appears in Howard Raiffa's 1968 book *Decision Analysis*.)

3b. "Increasing risk: I. a definition" by Michael Rothschild and Joseph Stiglitz

Another classic paper. It is well known to be impossible to define an unambiguous measure of the "risk" inherent in a monetary random variable with a specified probability distribution. For example, we cannot in general equate "risk" with "variance," since an agent's risk premium for a gamble may depend on higher moments of the probability distribution, depending on the form of the utility function. This paper shows that it is nevertheless possible to develop a precise characterization of the sense in which one variable may be "more risky" than another. In particular, this paper shows that there are three *equivalent* ways in which to characterize increasing risk, namely that $\tilde{x}$ is more risky than $\tilde{y}$ if:

(i) $\tilde{x}$ can be generated from $\tilde{y}$ by the addition of an independent "noise" variable with mean zero.

(ii) $\tilde{x}$ yields lower expected utility than $\tilde{y}$ under *every* strictly risk averse (strictly concave) utility function

(iii) $\tilde{x}$ has more "weight in the tails" than $\tilde{y}$, in the sense that $\tilde{x}$ can be obtained from $\tilde{y}$ by a sequence of "mean preserving spreads" in probability.

These results have since been generalized in a recent (1997) paper by Mark Machina and John Pratt in *Journal of Risk and Uncertainty*.

3c. "One-switch utility functions" by David Bell

This paper discusses in detail the interesting properties of the linear+exponential ("linex") utility function, which is decreasingly risk averse and also has the property that the direction of preference between any two gambles does not switch more than once as wealth increases. Also, for an agent with linex utility, it *is* possible to derive an unambiguous measure $R(\tilde{x})$ of the risk of a gamble which has the property that $\tilde{x}$ is preferred to another gamble $\tilde{z}$ if and only if $E(\tilde{x}) - R(\tilde{x}) > E(\tilde{z}) - R(\tilde{z})$, where $E(.)$ denotes expected value. Thus, for an agent with linex utility, we can say that one gamble is preferred to another if its excess of return over risk is greater. (Alas, the normative importance of the one-switch property and the risk measure are debatable.)

A sequel to this paper called "Risk, Return, and Utility," published in *Management Science* in 1995, discusses the question of the existence of a measure of risk in more detail.

3d. "Marginal value and intrinsic risk aversion" by David Bell and Howard Raiffa

As we discussed last week, a vNMn utility function need not represent "strength of preference" among levels of wealth (or other outcomes) under conditions of certainty. It is therefore not the same concept of utility that was imagined by utilitarian and marginalist economists. For example, if $u(x) - u(y) = u(y) - u(z)$, this does not mean that the quantitative increase in happiness or satisfaction or well-being or pleasure-minus-pain that is obtained by moving from a certain wealth level of z to a certain wealth level of y is the same as that obtained by moving from y to x. This relation means *only* that you should be indifferent between having y for sure or having a 50-50 gamble between x and z. *The relative utility differences between wealth levels have no meaning except in a context where expected values over uncertain outcomes are calculated (or in a multiattribute context where the preferential independence condition can otherwise be applied).* Yet the idea of a quantitative measure of strength of preference is intuitively appealing, and indeed it can be axiomatized, along the lines pioneered by Frisch in the 1920's, by starting from the assumption that there exists a qualitative, quaternary relation for comparing strengths of preference. Let $v^*(x)$ denote a measure of strength of preference under this interpretation—that is, $v^*(x) - v^*(y)$ *does* represent the increase in happiness (or whatever) that is obtained in moving from y to x—and let $u(x)$ denote a von Neumann-Morgenstern utility function. Then, as a purely mathematical exercise, we can always write

$$u(x) = u^*(v^*(x))$$

for some increasing function u^*. Here, u^* evidently represents the individual's *intrinsic risk aversion*, distinct from the question of her "diminishing marginal utility" under conditions of certainty. In principle, the function u^* might be concave or convex—for example, a person might have strongly diminishing marginal utility for money (as measured by v^*) and yet have an attraction to risk (as measured by u^*), with the result that $u(x)$ shows only slight risk aversion. Or, on the other hand, the person might have linear utility for money (i.e., v^* might be linear), but might nonetheless be averse to risk (as measured by u^*). Various authors (most notably Dyer and Sarin in the late '70s') have explored the idea that it might be possible to separately measure an individual's strength-of-preference function (a.k.a. "measurable value function") and her intrinsic attitude toward risk. The paper by Bell and Raiffa explores this issue further, suggesting a method for calibrating the strength-of-preference function and arguing that the function u^*, if one exists, should be of the exponential form—i.e., *individuals ought to be constantly risk averse when payoffs are measured in preference units rather than monetary units.* In practice, though, it has proved nearly impossible to separately measure strength-of-preference and intrinsic risk aversion. Hence, vNM utility functions are often treated as though they represent riskless cardinal utility, even though this is theoretically not correct. If you are interested in more details about the relation between strength-of-preference functions and utility functions, see the supplementary reading from the book by von Winterfeldt and Edwards.

3e. "The effect of the timing of consumption decisions and the resolution of lotteries on the choice of lotteries" by Michael Spence and Richard Zeckhauser (1972)

This very short paper indicates some of the complications that arise when we try to model preferences for sequences of events that occur over time—e.g., the amounts of "consumption" that occur at different dates. Normally when we work with utility functions for money, we implicitly assume that an individual desires money because it can be used to purchase consumption goods that are the real "consequences" over which preferences are defined. In a single-period problem, it is often unnecessary to model the problem of how the individual chooses to spend her money, so there is no need to distinguish between "wealth" and "consumption." The individual must somehow consume all her wealth, and hence it suffices to use a univariate utility function for wealth to model her preferences. But in decision problems that unfold over time, the individual presumably is able to *borrow and lend* so as to accelerate or delay consumption. Hence, we need to distinguish between the individual's "direct" utility for consumption streams and her "indirect" utility for income streams, and we need to worry about the timing of the resolution of the uncertainty as well as the timing of the receipt of cash flows. In general, the individual's indirect utility for income need not behave like a von Neumann-Morgenstern utility function, as this paper shows. (Preferences for consumption over time are usually modeled under strong simplifying assumptions about the form of the multiattribute utility function for consumption streams. Usually an assumption of "utility independence" between consumption at different dates is made, so that the total utility is either a sum or product of univariate utility functions for consumption at each of the dates.)

3f. "An aggregation theorem for securities markets" by Mark Rubinstein (1974)

This paper treats a problem related to the one examined by Wilson: under what conditions do a group of agents in a market under uncertainty behave like a single, expected-utility-maximizing composite agent? The conclusions are essentially the same: the aggregation problem can be solved only if the agents all have exponential utility functions, or if they all have nearly-identical utility functions (with the same "cautiousness" or "taste" parameter), or if they have the same beliefs and satisfy additional homogeneity conditions. Rubinstein also focuses on the HARA class of utility functions.

3g. Excerpt on "theoretical foundations of value and utility measurement" from *Decision Analysis and Behavioral Research* by Detlof von Winterfeldt and Ward Edwards (1986)

This chapter explores the connections among axiomatic theories of expected utility, strength-of-preference measurement, multiattribute utility, and additive conjoint measurement. If you are interested in problem of separating values from risk preferences, or the broader connections between the literatures of mathematical decision theory and mathematical psychology, this is a good reference.