Finance 400

A. Penati - G. Pennacchi

Risk Aversion and Risk Premia

An individual is said to be risk averse if and only if the individual's utility function is *concave*. This aversion to risk implies that the individual would not accept a "fair" lottery (asset). A fair or "pure risk" lottery is defined as one that has an expected value of zero. To see why a risk averse individual would have lower expected utilility from accepting a fair lottery, consider the following example. Let there be a lottery that has a random payoff, $\widetilde{\varepsilon}$, where

$$\widetilde{\varepsilon} = \begin{cases} h_1 \text{with probability } p \\ h_2 \text{ with probability } 1-p \end{cases} \tag{1}$$

The requirement that it be a fair lottery restricts its expected value to equal zero:

$$E\left[\widetilde{\varepsilon}\right] = ph_1 + (1-p)h_2 = 0 \tag{2}$$

which implies $h_1/h_2 = -\left(1-p\right)/p$, or, solving for p, $p = -h_2/\left(h_1 - h_2\right)$.

Now suppose a von Neumann-Morgenstern expected utility maximizer whose current wealth equals W is offered the above lottery. Would this individual accept it, that is, would she place a positive value on this lottery?

If the lottery is accepted, expected utility is given by $V = E\left[U\left(W + \widetilde{\varepsilon}\right)\right]$. Instead, if it is not accepted, expected utility is given by $V = E\left[U\left(W\right)\right] = U\left(W\right)$. Thus, if the individual refuses to accept a fair lottery, it implies

$$U\left(W\right) > E\left[U\left(W + \widetilde{\varepsilon}\right)\right] = pU\left(W + h_1\right) + (1-p)U\left(W + h_2\right) \tag{3}$$

To show that this is equivalent to having a concave utility function, note that $U\left(W\right)$ can be re-written as

$$U(W) = U\left(W + ph_1 + (1-p)h_2\right) \tag{4}$$

since $ph_1 + (1-p)h_2 = 0$ by the assumption that the lottery is fair. Re-writing inequality (3),

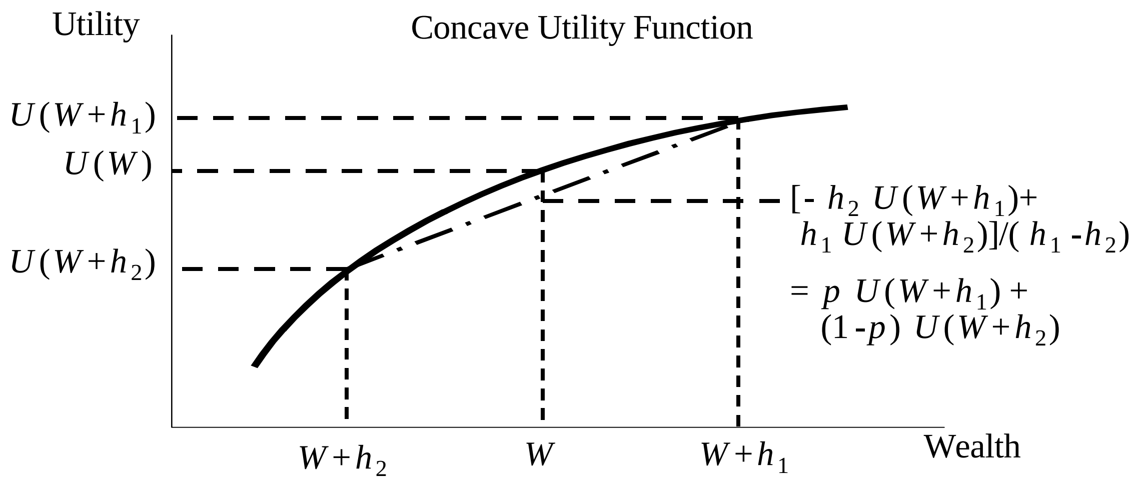

Figure 0-1: Fair Lotteries Lower Utility

we have

$$U\left(W + ph_1 + (1-p)h_2\right) > pU\left(W + h_1\right) + (1-p)U\left(W + h_2\right) \tag{5}$$

which is the definition of U being a concave function. A function is concave if a line joining any two points of the function lies entirely below the curve. When $U(W)$ is concave, a line connecting the points $U(W+h_2)$ to $U(W+h_1)$ lies below $U(W)$ for all W such that $W+h_2 < W < W+h_1$. $pU(W+h_1) + (1-p)U(W+h_2)$ is exactly the point on this line directly below $U(W)$. This is clear by substituting $p = -h_2/(h_1 - h_2)$. Note that when $U(W)$ is a continuous, second differentiable function, concavity implies that its second derivative, $U''(W)$, is less than zero.

To show the reverse, that concavity of utility implies the unwillingness to accept a fair lottery, we can use a result from statistics known as Jensen's inequality. If $U(\cdot)$ is some concave function, and $\tilde{x}$ is a random variable, then Jensen's inequality says that

$$E[U(\tilde{x})] < U(E[\tilde{x}]) \tag{6}$$

Therefore, substituting $\tilde{x} = W + \tilde{\varepsilon}$, with $E[\tilde{\varepsilon}] = 0$, we have

$$E\left[U(W+\tilde{\varepsilon})\right] < U\left(E\left[W+\tilde{\varepsilon}\right]\right) = U(W) \tag{7}$$

which is the desired result.

We have defined risk aversion in terms of the individual's utility function. We now consider how this aversion to risk can be quantified. This is done by defining a *risk premium*, the amount that an individual is willing to pay to avoid a risk.

Let π denote the individual's risk premium for a particular lottery, $\tilde{\varepsilon}$. It can be likened to the maximum insurance payment an individual would pay to avoid a particular risk. Pratt (1964) defined the risk premium for lottery (asset) $\tilde{\varepsilon}$ as

$$U(W-\pi) = E\left[U(W+\tilde{\varepsilon})\right] \tag{8}$$

This is the definition of a risk premium that is commonly used in the insurance literature. In financial economics, a somewhat different concept is often used, namely, that an asset's risk premium is its expected rate of return in excess of the risk-free rate of return. This alternative concept will be considered later.

To analyze this Pratt (1964) risk premium, we continue to assume the individual is an expected utility maximizer and that $\tilde{\varepsilon}$ is a fair lottery, that is, its expected value equals zero. Further, let us consider the case of $\tilde{\varepsilon}$ being "small," so that we can study its effects by taking a Taylor series approximation of equation (8) around the point $\tilde{\varepsilon} = 0$ and $\pi = 0$.[1] Expanding the left hand side of (8) around $\pi = 0$ gives

$$U(W-\pi) \cong U(W) - \pi U'(W) \tag{9}$$

and expanding the right hand side of (8) around $\tilde{\varepsilon} = 0$ (and taking a three term expansion since $E\left[\tilde{\varepsilon}\right] = 0$ implies that a third term is necessary for a limiting approximation) gives

$$E\left[U(W+\tilde{\varepsilon})\right] \cong E\left[U(W) + \tilde{\varepsilon}U'(W) + \tfrac{1}{2}\tilde{\varepsilon}^2 U''(W)\right] \tag{10}$$

[1] By describing the random variable $\tilde{\varepsilon}$ as "small" we mean that its probability density is concentrated around its mean 0.

$$= U(W) + \tfrac{1}{2}\sigma^2 U''(W)$$

where $\sigma^2 \equiv E\left[\tilde{\varepsilon}^2\right]$ is the lottery's variance. Equating the results in (9) and (10), we have

$$\pi = -\tfrac{1}{2}\sigma^2 \frac{U''(W)}{U'(W)} \equiv \tfrac{1}{2}\sigma^2 R(W) \tag{11}$$

where $R(W) \equiv -U''(W)/U'(W)$ is the Pratt (1964) - Arrow (1970) measure of absolute risk aversion. Note that the risk premium, π, depends on the uncertainty of the risky asset, σ^2, and on the individual's coefficient of absolute risk aversion. Since σ^2 and $U'(W)$ are both greater than zero, concavity of the utility function ensures that π must be positive.

From (11) we see that the concavity of the utility function, $U''(W)$, is not sufficient to quantify the risk premium an individual is willing to pay, even though it is necessary and sufficient to indicate if an individual is risk-averse or not. We also need the first derivative, $U'(W)$, which tells us the marginal utility of wealth. An individual may be very risk averse ($-U''(W)$ is large), but he may be unwilling to pay a large risk premium if he is poor since his marginal utility is high ($U'(W)$ is large).

To illustrate this point, consider the following utility function:

$$U(W) = -e^{-bW}, b > 0 \tag{12}$$

Note that $U'(W) = be^{-bW} > 0$ and $U''(W) = -b^2 e^{-bW} < 0$. Consider the behavior of a very wealthy individual, that is, one whose wealth approaches infinity:

$$\lim_{W\to\infty} U'(W) = \lim_{W\to\infty} U''(W) = 0 \tag{13}$$

As $W \to \infty$, the utility function is a flat line. Concavity disappears, which might imply that this very rich individual would be willing to pay very little for insurance against a random event, $\tilde{\varepsilon}$, certainly less than a poor person with the same utility function. However, this is not true because the marginal utility of wealth is also very small. This neutralizes the effect of smaller concavity. Indeed:

$$R(W) = \frac{b^2 e^{-bW}}{be^{-bW}} = b \tag{14}$$

which is a constant. Thus, we can see why this utility function is sometimes referred to as a constant absolute risk aversion utility function.

If we want to assume that absolute risk aversion is declining in wealth, a necessary, though not sufficient, condition for this is that the utility function have a positive third derivative, since

$$\frac{\partial R(W)}{\partial W} = -\frac{U'''(W)U'(W) - [U''(W)]^2}{[U'(W)]^2} \tag{15}$$

Also, it can be shown that the coefficient of risk aversion contains all relevant information about the individual's risk preferences. Note that

$$R(W) = -\frac{U''(W)}{U'(W)} = -\frac{\partial \left(\ln\left[U'(W)\right]\right)}{\partial W} \tag{16}$$

Integrating both sides of (16), we have

$$-\int R(W)dW = \ln[U'(W)] + c \tag{17}$$

Taking the exponential function of (17)

$$e^{-\int R(W)dW} = U'(W)e^c \tag{18}$$

Integrating once again gives

$$\int e^{-\int R(W)dW} dW = e^c U(W) + d \quad \sim \quad U(W) \tag{19}$$

Because expected utility functions are unique up to a linear transformation, $e^c U(W) + d$ reflects the same risk preferences as $U(W)$.

Relative risk aversion is another frequently used measure of risk aversion and is defined simply as

$$R_r(W) = WR(W) \tag{20}$$

Some useful utility functions:

a. *Negative exponential* (constant absolute risk aversion)

$$U(W) = -e^{-bW}, b > 0$$

As we saw earlier, $R(W) = b$, so that $R_r(W) = bW$.

b. *Power* (constant relative risk aversion)

$$U(W) = \frac{1}{\gamma}W^{\gamma}, \gamma < 1$$

implying that $R(W) = -\frac{\gamma(\gamma-1)W^{\gamma-2}}{\gamma W^{\gamma-1}} = \frac{(1-\gamma)}{W}$ and, therefore, $R_r(W) = 1 - \gamma$.

c. *Logarithmic*

This is a limiting case of power utility. To see this, write the power utility function as $\frac{1}{\gamma}W^{\gamma} - \frac{1}{\gamma} = \frac{W^{\gamma}-1}{\gamma}$. (Recall that we can do this because utility functions are unique up to a linear transformation.) Next take the limit of this utility function as $\gamma \to 0$. Note that the numerator and denominator both go to zero, so that the limit is not obvious. However, we can re-write the numerator in terms of an exponential and natural log function and apply L'Hospital's rule to obtain:

$$\lim_{\gamma\to 0} \frac{W^{\gamma}-1}{\gamma} = \lim_{\gamma\to 0} \frac{e^{\gamma \ln(W)}-1}{\gamma} = \lim_{\gamma\to 0} \frac{\ln(W)W^{\gamma}}{1} = \ln(W)$$

Thus, logarithmic utility is equivalent to power utility with $\gamma = 0$, or a coefficient of relative risk aversion of unity:

$R(W) = -\frac{W^{-2}}{W^{-1}} = \frac{1}{W}$ and $R_r(W) = 1$.

d. *Quadratic*

$$U(W) = W - \frac{b}{2}W^2, b > 0$$

Note that $U'(W) = 1 - bW$, which is > 0 if and only if $b < \frac{1}{W}$. Thus, this utility function only makes sense when $W < \frac{1}{b}$, which is know as the "bliss point." We have $R(W) = \frac{b}{1-bW}$ and $R_r(W) = \frac{bW}{1-bW}$.

e. *HARA* (hyperbolic absolute risk aversion)

$$U(W) = \frac{1-\gamma}{\gamma}\left(\frac{\alpha W}{1-\gamma} + \beta\right)^{\gamma}$$

subject to the restrictions $\gamma \neq 1$, $\alpha > 0$, $\frac{\alpha W}{1-\gamma} + \beta > 0$, and $\beta = 1$ if $\gamma = -\infty$. Thus, $R(W) = \left(\frac{W}{1-\gamma} + \frac{\beta}{\alpha}\right)^{-1}$. Since $R(W)$ must be > 0, it implies $\beta > 0$ when $\gamma > 1$. $R_r(W) = W\left(\frac{W}{1-\gamma} + \frac{\beta}{\alpha}\right)^{-1}$. HARA utility nests constant absolute risk aversion ($\gamma = -\infty, \beta = 1$), constant relative risk aversion ($\gamma < 1$, $\beta = 0$), and quadratic ($\gamma = 2$) utility functions. Thus, depending on the parameters, it is able to display constant absolute risk aversion or relative risk aversion that is increasing, decreasing, or constant.

Kenneth Arrow (1970) independently derived a coefficient of risk aversion that is identical to Pratt's measure, but Arrow's derivation is based on a concept of a risk premium that is commonly used in financial markets. Suppose that an asset (lottery), $\widetilde{\varepsilon}$, has the following payoffs and probabilities (this could be generalized to other types of fair payoffs):

$$\widetilde{\varepsilon} = \begin{cases} +h \text{ with probability } \frac{1}{2} \\ -h \text{ with probability } \frac{1}{2} \end{cases} \tag{21}$$

Note that, as before, $E[\widetilde{\varepsilon}] = 0$. Now consider the following question. By how much should we change the expected value (return) of the asset, by changing the probability of winning, in order to make the individual indifferent between taking and not taking the risk? If p is the probability of winning, we can define the risk premium as

$$\theta = prob\left(\widetilde{\varepsilon} = +h\right) - prob\left(\widetilde{\varepsilon} = -h\right) = p - (1-p) = 2p - 1 \tag{22}$$

Therefore, from (22) we have

$$\begin{array}{c} prob\left(\widetilde{\varepsilon} = +h\right) \equiv p = \frac{1}{2}(1+\theta) \\ prob\left(\widetilde{\varepsilon} = -h\right) \equiv 1-p = \frac{1}{2}(1-\theta) \end{array} \tag{23}$$

These new probabilities of winning and losing are equal to the old probabilities, $\frac{1}{2}$, plus half of the increment, θ. Thus, the premium, θ, that makes the individual indifferent between accepting and refusing the asset is

$$U(W) = \frac{1}{2}(1+\theta)U(W+h) + \frac{1}{2}(1-\theta)U(W-h) \tag{24}$$

Taking a Taylor series approximation around $h = 0$, gives

$$U(W) = \frac{1}{2}(1+\theta)\left[U(W) + hU'(W) + \tfrac{1}{2}h^2U''(W)\right] \tag{25}$$

$$+\frac{1}{2}(1-\theta)\left[U(W) - hU'(W) + \tfrac{1}{2}h^2U''(W)\right]$$

$$= U(W) + h\theta U'(W) + \tfrac{1}{2}h^2U''(W)$$

Re-arranging (25) implies

$$\theta = \tfrac{1}{2}hR(W) \tag{26}$$

which, as before, is a function of the coefficient of absolute risk aversion. Note that the Arrow premium, θ, is in terms of a probability, while the Pratt measure, π, is in units of a monetary payment. If we multiply θ by the monetary payment received, h, then (26) becomes

$$h\theta = \tfrac{1}{2}h^2R(W) \tag{27}$$

Since h^2 is the variance of the random payoff, $\widetilde{\varepsilon}$, equation (27) shows that the Pratt and Arrow measures of risk premia are equivalent. Both were obtained as a linearization of the true function around $\widetilde{\varepsilon} = 0$.