Asymmetric loss utility: an analysis of decision under risk

Alex Strashny
astrashn@uci.edu
University of California, Irvine

29 January 2004 [Draft]
Original: 10 November 2003

Abstract

This paper develops a utility model for evaluating lotteries. In estimating utility, risk averse people use an asymmetric loss function. Expected utility is seen as a special case that is a good approximation in some cases. The model resolves several paradoxes and makes easily falsifiable predictions. When used in hypothesis testing, the model allows researchers to directly specify their attitudes toward risk.

Keywords: choice under uncertainty, non-expected utility theory, risk aversion, Allais paradox, Ellsberg paradox, St. Petersburg paradox, equity premium puzzle, decision theory.

1 Introduction

Suppose a lottery L has M outcomes numbered $m = 1, 2, \ldots M$, and outcome b_m occurs with probability p_m, $\sum p = 1$. The *elementary utility* of an outcome b is $u(b)$. The most common way of calculating the utility of the lottery is using the expected utility theory, which states that the utility of a lottery is the expected utility:

$$U_E(L) = Eu = \sum_{m=1}^{M} p_m u(b_m). \tag{1}$$

There are several problems with expected utility. Theoretical problems **I will describe here XXX**. Another problem is more practical: expected utility fails to explain several well known paradoxes, such as the Allais paradox, Ellsberg paradox, and St. Petersburg paradox.

In this paper, I take a different approach. The utility of a lottery, *after* the lottery's outcome is known, is the utility of the outcome. That is, if we know that the lottery's outcome is b_m, then $U(L) = u(b_m)$. *Before* the outcome is known, the utility of a lottery is a random variable. What people mean when they refer to the utility of a lottery is actually a point estimate of the random variable.

Expected value is the optimal estimator under quadratic loss, which is a symmetric loss function. In my view, a symmetric loss function reflects risk neutrality. A risk averse individual uses asymmetric loss function in which higher loss comes from overestimating utility. This causes his estimates of utility to be, in general, less than expected utility.

In section 2, I describe the model. In section 3, I look at several well-known paradoxes. Finally, in section 4, I discuss the model in the context of hypothesis testing.

2 Utility

Suppose a person plans to play some lottery a total of n times, and that outcome b_m occurs x_m times. The total utility from all n games is:

$$n\tilde{U}(L) = \sum_{m=1}^{M} x_m u\left(b_m\right). \tag{2}$$

Define the utility of the lottery as the average per-game utility after n games. That is,

$$\tilde{U}(L) = \sum_{m=1}^{M} \left(\frac{x_m}{n}\right) u\left(b_m\right). \tag{3}$$

In the above equations, since counts x_m are random variables, so is $\tilde{U}(L)$. When evaluating lotteries, people *estimate* the utilities of those lotteries. I denote the random variable by $\tilde{U}\,(L)$ and its estimate by by $U(L)$. *Risk neutral* individuals care about overestimating and underestimating utility equally: they have a symmetric loss function. For instance, people who minimize expected quadratic loss, estimate the utility of a lottery by the expected value over elementary utilities. *Risk averse* individuals would rather underestimate utility than overestimate it: they have an asymmetric loss function in which overestimation is penalized more than underestimation.

If the probabilities p_m are known, then counts $\{x_m\}$ have the Multinomial distribution as follows:

$$\{x_m\} \, | n, \{p_m\} \sim Mult\left(n, \{p_m\}\right). \quad (4)$$

In the case when there are $M = 2$ possible outcomes, this reduces to the Binomial distribution.

Sometimes, the probabilities p_m themselves are unknown. This is not a problem if their distribution is known. Suppose the probabilities $\{p_m\}$ are *a priori* Dirichlet distributed with parameter vector $\{\alpha_m\}$. (In case $M = 2$, this is the same as the Beta distribution.) We observe k outcomes. In these observations, outcome b_m occurs y_m times. Then, the posterior distribution of the probabilities is Dirichlet:

$$\{p_m\} \, | \, \{y_m\} \sim D\left(\{\alpha_m + y_m\}\right). \quad (5)$$

From this, when probabilities p_m are not known, but there is data, counts x_m have the Multinomial-Dirichlet distribution.

Note that, if the probabilities p_m are known, as the number of games n approaches infinity, by the frequentist definition of probability, the weights $\left(\frac{x}{n}\right)$ approach the probabilities p. Thus,

$$\lim_{n \to \infty} \tilde{U}\left(L\right) = U_E\left(L\right). \quad (6)$$

So, any reasonable estimator of $\tilde{U}(L)$ must also approach expected utility $U_E(L)$. In other words, expected utility theory can be thought of in two ways. First, it accurately reflects preferences of risk neutral individuals with symmetric quadratic loss. Second, it is a limiting case when the number of games n approaches infinity and there is no uncertainty.

2.1 Loss function

A loss function $C(\tilde{U}, U)$ expresses the loss that an individual experiences from estimating random utility $\tilde{U}$ by some estimate U. To find the optimal estimate, the individual minimizes expected loss. If $f(\tilde{U})$ is the probability mass function of $\tilde{U}$, then

$$E\left[C\left(\tilde{U},U\right)\right] = \sum_{\tilde{U}} C\left(\tilde{U},U\right) f(\tilde{U}) \tag{7}$$

$$U = \arg\min_{U} E\left[C\left(\tilde{U},U\right)\right] \tag{8}$$

Each individual may have his own unique loss function $C(\cdot)$. When discussing attitudes toward risk, we are interested in how loss function penalizes overestimation as opposed to underestimation of utility. Let $\Delta > 0$ be some constant. If a person is indifferent between overestimating and underestimating utility, that is, if $C(U, U+\Delta) = C(U, U-\Delta)$, I call him *risk neutral.* If a person is more afraid to overestimate utility, that is, if $C(U, U+\Delta) > C(U, U-\Delta)$, then he is *risk averse.* Finally, if he is more afraid of underestimating utility, $C(U, U+\Delta) < C(U, U-\Delta)$, then he is *risk loving.*

To perform calculations, we need to define a specific loss function. In this paper, I use the following *asymmetric loss*:

$$C_a\left(\tilde{U},U\right) = \begin{cases} \left(\tilde{U}-U\right)^a & \text{if } \tilde{U} \geq U \\ c_a\left|\tilde{U}-U\right|^a & \text{if } \tilde{U} < U \end{cases}, \tag{9}$$

where $a \geq 0$ and $c_a \geq 0$ are constants. $c_a > 1$ reflects risk aversion, while $c_a = 1$ reflects risk neutrality. When $a \to 0$ and $c_0 = 1$, the loss function is called the *all-or-none loss*; when $a = 1$, it is the *asymmetric linear loss*; when $a = 2$, this is the *asymmetric quadratic loss*; when $a = 2$ and $c_2 = 1$, it is the usual quadratic loss.

Because a defines the general shape of the loss function, I call it the *type of loss*; because c_a defines the degree of asymmetry, I call it *risk aversion.*

For linear (type 1) loss, there is an analytical solution for the best estimator U. Define $q = (1 + c_1)^{-1}$; the best estimator under type 1 loss is the q-th quantile of $\tilde{U}$. Thus, for symmetric linear loss, the best estimator is the median of $\tilde{U}$. If the loss is symmetric quadratic ($a = 2$, $c_2 = 1$), then the best estimator is the expected value of $\tilde{U}$. For arbitrary values of a and c_a, however, I don't know of a general analytical solution, and I find best estimates numerically.

The type of loss and risk aversion that people commonly have need to be determined experimentally. In this paper, I concentrate on three types of loss. I use the linear (type 1) loss because there is an analytical solution for the best estimator under this loss. A problem with this loss is that for discrete random variables, the estimator of utility $U(L)$ is not a continuous function of risk aversion c_1. Quadratic (type 2) loss is also attractive, because symmetric quadratic loss is used so often. But because there is no analytical solution for it, I don't use it for the more complicated problems, such as the Equity Premium Puzzle of section 3.4. Finally, sometimes, type 2 loss produces estimates that are too high to be of interest, while type 1 produces estimates that are two low to be of interest. This is the case with the St. Petersburg paradox of section 3.3. In that case, I use type 1.5 loss.

2.2 Risk aversion

Here, I suggest a couple of thought experiments which can help each person find his own value of risk aversion c_a. This is done by imagining lotteries the values of which the person knows for himself. For example, suppose the utility of a lottery has the Discrete Uniform distribution, as follows: it is -1 with probability 50% and 1 with probability 50%. In a person's view, the utility of this lottery is U. Then, for $a > 1$, his risk aversion c_a is

$$c_a = (1 - U)^{a-1} (1 + U)^{1-a}. \tag{10}$$

Selected values of risk aversion c_a are tabulated in table 1. Based on this thought experiment, for type 1.5 loss, risk aversion $c_{1.5}$ between 1.5 and 2 seems reasonable; for type 2 (quadratic) loss, risk aversion c_2 of 2.3 to 4 seems reasonable.

Let's perform a similar thought experiment with another lottery. Suppose the utility of a lottery has the Standard Normal distribution and that a person values this lottery at U. For type 1 (linear) loss, the corresponding risk aversion is

$$c_1 = \frac{1}{\Phi(U)} - 1, \tag{11}$$

where $\Phi(\cdot)$ is the cumulative density of the Standard Normal. For other types of loss function, I do not know of an analytical solution for risk aversion. I solve for risk aversion numerically and tabulate the values in table 2. Based

Utility U	Type of loss a	
	1.5	2.0
0.0	1.0	1.0
-0.1	1.1	1.2
-0.2	1.2	1.5
-0.3	1.4	1.9
-0.4	**1.5**	**2.3**
-0.5	**1.7**	**3.0**
-0.6	**2.0**	**4.0**
-0.7	2.4	5.7
-0.8	3.0	9.0
-0.9	4.4	19.0

Table 1: Risk aversion c_a corresponding to selected values of Discrete Uniform utility described in the text.

on this thought experiment, the following values of risk aversion c_a don't seem very high: for linear (type 1) loss, $c_1 = 6$; for type 1.5 loss, $c_{1.5} = 9$; for quadratic (type 2) loss, risk aversion $c_2 = 13$ does not seem very high.

The above thought experiments are done just to get the feel for the correct magnitude of risk aversion. In this paper, I use low values of risk aversion to illustrate that even a small deviation from the standard symmetric loss approach (which is implied by expected utility) can resolve a number of apparent problems. Thus, when applying quadratic loss to discrete utilities, I use $c_2 = 3.0$ (see table 1); when applying linear loss to the Normal distrib-

Utility U	Type of loss a		
	1.0	1.5	2.0
0.0	1.0	1.0	1.0
-0.5	2.3	2.9	3.6
-0.8	3.8	5.5	7.7
-1.0	**5.4**	**8.7**	**13.1**
-1.2	7.8	13.8	22.4
-1.5	14.1	28.4	52.5

Table 2: Risk aversion c_a corresponding to selected values of Standard Normal utility described in the text.

ution, I use $c_1 = 3.8$ (see table 2); when the intermediate, type 1.5, loss is needed, as for the St. Petersburg paradox, I use $c_{1.5} = 1.7$ (table 1). Actually, the paradoxes discussed below are resolved under a wide range of values of a and c_a. Using higher risk aversions c_a resolves them even more easily.

2.3 Buying and selling

When a person *buys* a lottery, he will receive one of its potential outcomes; when he *sells* a lottery, he will have to pay out one of its potential outcomes.[1] That is, from the seller's point of view, all the outcomes are negated. This means that, if a person is not risk neutral, his estimated utility of a lottery as a buyer is not equal to the (negative) estimated utility as a seller.

For example, consider a lottery with two equally likely outcomes: $b_1 = 0$, $b_2 = 1$; the lottery will be played once ($n = 1$). A person's elementary utility is $u(b) = b$; he has quadratic (type 2) loss and a risk aversion of $c_2 = 3.0$. If the person is considering buying the lottery, he evaluates it at 0.25; that is, he is willing to pay 0.25 or less for playing the lottery once. If, on the other hand, the person is considering selling the lottery, he evaluates it at -0.75. This means that he requires a payment of 0.75 or more if he will have to make the payments specified by the lottery.

2.4 Backward induction

If, by picking one of several lotteries, a person commits to playing it n times, everything is straightforward: the person calculates the utility $U(L)$ of playing each lottery n times, and picks the lottery with the greatest n-game utility.

Suppose, however, that the person knows that he will play n times, but he can decide which lottery to play before each game. Then, he does not choose a single lottery, but rather chooses a lottery path. This is easily done with backward induction.

The person knows that during the last game, game n, he will pick the lottery with the greatest one-game utility. Thus, when making the choice for game $n - 1$, he already knows his choice for game n. Similarly, when choosing the lottery for game $n - 2$, he already knows his choices for the last

[1] The terms *going long* a lottery and *going short* a lottery might be more accurate.

two games. In this way, the person can calculate the best path of lotteries for all n games. I use this logic in the Equity Premium puzzle of section 3.4.

3 Paradoxes

The "paradoxes" discussed below are situations in which the standard expected utility theory predicts one outcome, while we observe another outcome in experiments. This points to a failure of the theory to make correct predictions in some circumstances. I resolve the paradoxes by showing that if the model developed in this paper is used, the predicted and observed outcomes match.

3.1 Allais paradox

Here is a usual statement of the Allais paradox. A person is asked to choose between the following two gambles:

Gamble A Receive \$1M (one million dollars) with 100% probability.

Gamble B Receive \$5M with 10% probability, \$1M with 89% probability, or nothing with 1% probability.

He is also asked to choose between the following two gambles:

Gamble C Receive \$1M with 11% probability, and nothing with 89% probability.

Gamble D Receive \$5M with 10% probability, and nothing with 90% probability.

It is observed that most people choose A over B and choose D over C. However, according to expected utility theory, if a person prefers A over B, he must also prefer C over D.

Denote the possible outcomes as $b = \{0, 1, 5\}$. Preferring A to B means that the difference in utilities of these gambles must be greater than zero. In expected utility terms, that difference is

$$U_E(A) - U_E(B) = 0.11u(1) - 0.10u(5) - 0.01u(0). \qquad (12)$$

	$U(L)$	$U_E(L)$
A	0.001	0.001
B	-0.05	0.09
$A-B$	**0.05**	**$-$0.09**
C	-6.61	-6.06
D	-6.63	-6.15
$D-C$	**0.02**	**0.09**

Table 3: Allais utilities as estimated under asymmetric quadratic (type 2) loss, and by expected value. The paradox disappears under asymmetric loss.

Likewise, the difference between expected utilities of D and C is

$$U_E(D) - U_E(C) = 0.10u(5) + 0.01u(0) - 0.11u(1). \tag{13}$$

In other words, these expected utility differences have the opposite sign

$$U_E(A) - U_E(B) = -\left(U_E(D) - U_E(C)\right). \tag{14}$$

Thus, according to expected utility, if a person prefers A to B, it is *impossible* that he prefers D to C.

Here is a resolution of the paradox. Since the potential payoffs are so large compared to the resources of most players, suppose that the lotteries are only offered once: $n = 1$. To reflect decreasing marginal utility of money, let the utility of each outcome $b = \{0, 1, 5\}$ be logarithmic: $u(b) = \ln(v + b)$, where $v > 0$ is a constant that reflects player's resources. I set $v = 10^{-3}$, which corresponds to the player having about a thousand dollars.

Under quadratic loss with $c_2 = 3.0$, the paradox disappears. Refer to table 3 which shows that A is in fact preferred to B at the same time as D is preferred to C.

According to the model developed here, I predict that if people are explicitly told that they can play the lotteries a very large number of times, then they will make choices consistent with expected utility.

3.2 Ellsberg paradox

Here is a usual statement of the Ellsberg paradox. An urn contains 300 balls: 100 are red; of the rest, some are blue and some are green. A person draws a random ball from the urn and is asked to choose between the following gambles:

Gamble A Receive $1 if the ball is red.

Gamble B Receive $1 if the ball is blue.

He also has to choose between the following two gambles:

Gamble C Receive $1 if the ball is not red.

Gamble D Receive $1 if the ball is not blue.

People usually prefer A to B and C to D. However, if we use expected utility theory, it appears that a person who prefers A to B has to also prefer D to C. Let p_R be the probability of a red ball and p_{NR} be the probability that the ball is not red. Then, under expected utility theory, preferring A to B implies that, in the person's view, $p_R > p_B$. But preferring C to D implies that, in his view, $p_{NR} > p_{NB}$. Since $p_{NR} = 1 - p_R$, both of these inequalities cannot be true.

Now, consider these gambles with the model presented in this paper. A person plans to play his chosen gamble n times; x_R is the number of times he draws a red ball. Letting the utility of no payment be zero, $u(0) = 0$, the random utility of A is

$$\tilde{U}(A) = \frac{x_R}{n} u(1). \tag{15}$$

The estimate of this utility is the same thing but with x_R replaced by its estimate $\hat{x}_R$. Thus, utilities of A and B are

$$U(A) = \frac{\hat{x}_R}{n} u(1) \tag{16}$$

$$U(B) = \frac{\hat{x}_B}{n} u(1) \tag{17}$$

Since drawing any ball is equally likely, the person knows that the probability of a red ball is $p_R = \frac{100}{300} = \frac{1}{3}$. From this, he knows the distribution of the number of red draws x_R: it is Binomial with parameters n and $p_R = \frac{1}{3}$.

$$x_R \sim B(n, p_R). \tag{18}$$

	$U(L)$
A	0.29
B	0.24
$A-B$	**0.05**
C	0.62
D	0.57
$C-D$	**0.05**

Table 4: Ellsberg utilities.

On the other hand, he does not know the probability of a blue ball p_B. The person might believe that the probability p_B is distributed according to some probability function. Since probability p_B can be anything between 0 and $\frac{2}{3}$, the person might think that p_B is distributed Uniformly between 0 and $\frac{2}{3}$. If $f(p_B)$ is the probability density of p_B, then the distribution of the number of blue balls is

$$f(x_B) = \int_0^{\frac{2}{3}} B(n, p_B) f(p_B) \, dp_B. \tag{19}$$

If $f(p_B)$ is in fact Uniform, the expected number of red draws x_R is equal to the expected number of blue draws x_B: $E[x_R] = E[x_B]$. However, the variance of blue draws is greater than the variance of red draws: $Var(x_B) > Var(x_R)$. This is because the probability of red draws is certain, while the probability of blue draws is not. The added uncertainty adds to the variance. Under asymmetric loss, with equal expected values and unequal variances, the estimate $\hat{x}_B$ is less than $\hat{x}_R$. Because of this, the person prefers A over B.

The same logic applies to the second pair of gambles. Let x_{NR} be the number of non-red draws. Then, because of equal expected values and unequal variances, the estimate of the number of non-blue draws is less than the estimate of the number of non-red draws: $\hat{x}_{NB} < \hat{x}_{NR}$. And so, $U(A) > U(B)$ and $U(C) > U(D)$.

As a numerical example, let's say that $u(1) = 1$, $u(0) = 0$, and the number of games is $n = 20$. I use quadratic loss with risk aversion $c_2 = 3.0$. The lottery utilities are shown in table 4. As the table shows, A is preferred to B while C is preferred to D.

3.3 St. Petersburg paradox

Here is a discussion of the St. Petersburg paradox based on (Martin 2001). A fair coin is flipped until it comes up heads for the first time. Let k be the toss on which this happens. Then, the St. Petersburg gamble pays $\$2^k$. The question is, how much would someone be willing to pay for playing this gamble? The expected value of the gamble is infinite:

$$EV = \frac{1}{2}2 + \frac{1}{4}4 + \ldots = \infty. \tag{20}$$

From this, it might appear that people would be willing to pay an infinite amount of money for the gamble. This is obviously wrong as, in reality, people are only willing to pay much less than infinity. One flaw with the above presentation of the paradox is that it is made in terms of payoffs, not utilities. In fact, the first solution of the paradox, proposed by Bernoulli, is that people perceive payoffs in terms of utilities which are increasing, but at a decreasing rate.

However, we can easily circumvent this solution by making the payoffs not 2^k, but higher. If $u^{-1}(\cdot)$ is the inverse of the elementary utility function, let's make the payoff $u^{-1}\left(2^k\right)$. In this case, the *expected utility* of the gamble is infinite, and so it appears, once again, that people should be willing to pay an infinite amount for the gamble.

Now, consider this lottery in the framework presented here. For simplicity, let the utility of an outcome be equal to the outcome: $u(b) = b$. Under quadratic (type 2) loss, the value of utility diverges to infinity, regardless of what the risk aversion c_2 is. Under linear (type 1) loss with either risk neutrality or risk aversion ($c_1 \geq 1$), the utility $U(L)$ is always 2. While this is one possible answer, it is rather uninteresting. It could be argued that the answer is simply an artifact of the fact that the linear loss produces discontinuous estimates when applied to discrete distributions.

Consider now a type of loss that is between the two types discussed above, namely, the type 1.5 loss. Under this loss function, the utility of the lottery converges to a value greater than 2. For example, if the number of games is $n = 1$ and the risk aversion is $c_{1.5} = 1.7$, the utility of the lottery converges to $U(L) \approx 3.85$.

3.4 Equity premium puzzle

This section provides some insight into the equity premium puzzle. The difference between the return on equities and the return on almost riskless bonds is called the *equity premium*. Because equities are risky while bonds are not, expected equity premium is positive. In the United States and some other countries, when viewed through the lens of various asset pricing models, the premium seems excessive. According to asset pricing models, risk aversion required to sustain such a large premium is unrealistically high (Obstfeld & Rogoff 1996, p. 310). There are two possible explanations for this paradox: either the asset pricing models do not accurately describe human behavior, or people are, in fact, extremely risk averse, at least in some situations.

The classic paper on the subject is (Mehra & Prescott 1985). The paper examines real returns on stocks and almost riskless bonds between the years 1889 and 1978. The average real return on stocks was 6.98% per year, while on bonds, it was 0.8% per year. Thus, the equity premium is $6.98\% - 0.8\% = 6.18\%$ per annum. But, according to asset pricing models, under reasonable values of risk aversion, equity premium should not be greater than about 1% or 2%.

Similarly large equity premiums are present in other data as well. (Shiller 2000) gives annual data on Stocks, Bonds, and the Consumer Price Index between the years 1871 and 1997.[2] Stocks data consists of January values of the Standard and Poor's Composite index and yearly dividend data for those stocks. Bonds data is the total nominal return from investing in January and then reinvesting in July at the six month prime commercial paper rate. Based on the (Shiller 2000) data, I calculate real returns for holding Stocks and Bonds. Rate r is the real annual return; *Logarithm Return* (LR) is $\ln(1+r)$. The calculated statistics are in table 5. According to this data, the equity premium is $8.60\% - 2.96\% = 5.63\%$, still much greater than 2%.

I use the model developed in this paper, with a relatively low value of risk aversion, to predict investor behavior in two cases: when the returns are as described by the data; and when returns are such that the equity premium is 2%. I find that the predicted investor behavior if returns are as observed seems reasonable, while behavior when equity premium is set to 2% seems very unreasonable.

Suppose only two investment vehicles are available: Stocks and Bonds:

[2] Thanks to John Nuttall of University of Western Ontario for providing the data.

	Return r		Log return $\ln(1+r)$	
	Mean	Std dev	Mean μ	Std dev σ
Stocks	0.0860	0.1750	0.0689	0.1669
Bonds	0.0296	0.0681	0.0270	0.0658
	Correlation coefficient			
Stocks-Bonds	0.107		0.129	

Table 5: Returns on Stocks and Bonds from 1871 to 1997.

$V = \{S, B\}$. In the minds of investors, the Logarithm Rates are Normally distributed with statistics as shown in table 5, and zero correlations across time. Each investor knows the number of years n that he will invest. For example, this could be the number of years to retirement. To simplify computations, I use the linear (type 1) loss function. Each investor knows his risk aversion c_1, which could be related to personality, family situation, and so on. I assume the relatively low risk aversion of $c_1 = 3.8$ (see table 2 in section 2.2).

For convenience, let t be the number of years until the investor stops investing (such as until retirement). Thus, the first year of investing is $t = n$ while the last year is $t = 1$. Define the utility from investing is the logarithm of the total return. In other words,

$$\tilde{U} = \ln \prod_{t=1}^{n} (1 + r_t) = \sum_{t=1}^{n} \ln(1 + r_t), \tag{21}$$

where r_t is the real return that the investor receives in year t.

Before the beginning of each year, investors decide what percentage of their money to put into each of the investment vehicles. They determine the optimal investment path by backward induction, as discussed in section 2.4. The utility is Normally distributed as follows:

$$\tilde{U} \sim N\left(\sum_{t=1}^{n} \mu_t, \sum_{t=1}^{n} \sigma_t^2\right), \tag{22}$$

where μ_t and σ_t are mean and standard deviation of the Logarithm Return of the investment mix chosen for year t. If $\pi_{V,t}$ is the fraction invested in V in year t, then

$$\mu_t = \sum_V \pi_{V,t}\mu_V \tag{23}$$

$$\sigma_t^2 = \left(\sum_V \pi_{V,t}^2\sigma_V^2\right) + 2\pi_{S,t}\pi_{B,t}\sigma_S\sigma_B\rho_{SB} \tag{24}$$

The right hand side variables, μ_V, σ_V, and ρ_{SB}, are taken from table 5.

Figure 1 shows the proportion of money invested in Stocks, as a function of time remaining t, for a person with risk aversion $c_1 = 3.8$; the rest of the money is held in Bonds. Until there are $t = 21$ years left, the person holds all his money in Stocks. Starting from $t = 20$, he gradually begins to shift from Stocks into Bonds. In $t = 8$, proportion invested in Stocks falls below 50%. During the last year of investing, when $t = 1$, the proportion is $\pi_{S,1} = 22\%$.

Now, suppose that, following asset pricing models, equity premium was 2%. That is, all the data is as is, except that the returns of Stocks r_S are shifted by $-E[r_S]+E[r_B]+0.02$. This changes the expected return on Stocks without changing their risk (standard deviation of return) or correlation to Bonds. The estimated mean of Logarithm Return for Stocks becomes $\mu_S = E[\ln(1+r_S)] = 0.0339$. Figure 2 shows the proportion of money invested in Stocks, as a function of t, under this scenario. Now, the person starts shifting into Bonds at $t = 786$; he starts investing less than 50% in Stocks at $t = 312$; he starts investing less than 20% in Stocks at $t = 27$. Such an investment path is very unrealistic.

4 Hypothesis testing

Hypothesis testing allows us to tell whether sufficient evidence exists for some proposition of interest. For example, based on data related to some quantity β, we might want to know if there is sufficient evidence that $\beta > 0$. The *alternate hypothesis*, H_A, is the proposition for which we would like to know whether sufficient support exists; the *null hypothesis*, H_0, is the complement of the alternate. In this example, $H_0 : \beta \leq 0; H_A : \beta > 0$.

In conventional hypothesis testing, the *size of the test* is usually set to 5%; sometimes, it is also set to either 10% or 1%. The researcher sets the size of the test somewhat arbitrarily, without any direct reference to his attitude toward risk. Many times a null can be rejected at one common level, such

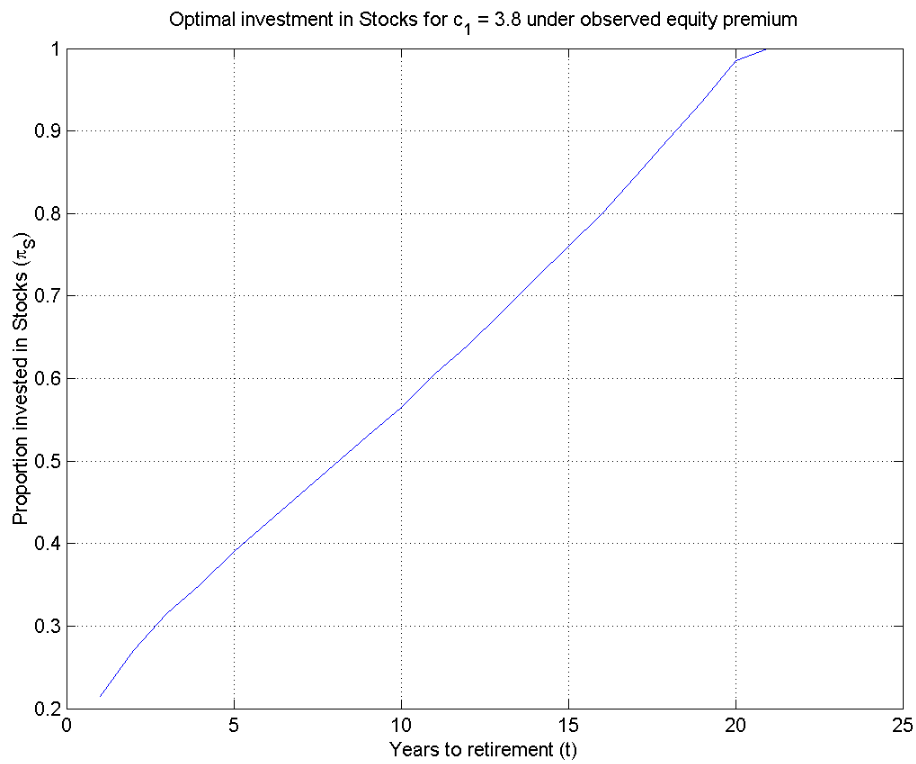

Figure 1: Optimal investment path under observed equity premium. Proportion invested in stocks versus years to retirement.

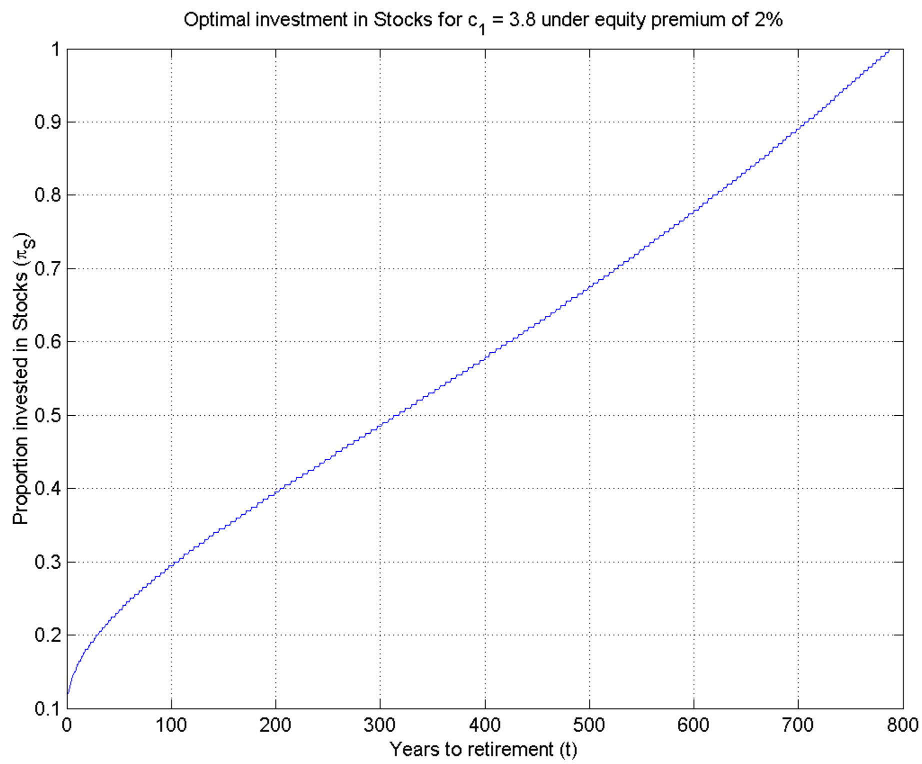

Figure 2: Optimal investment path if equity premium was 2%. Proportion invested in stocks versus years to retirement.

as 5%, but not at another common level, such as 1%. In this situation in particular, the researcher has to have his preferences well-quantified.

4.1 The basic utility approach

Let B be a random variable, such as the response to some treatment. $u(B) > 0$ means that the response is desirable. B can take on M possible values, subscripted with $m = 1, 2, \ldots, M$. A researches has already observed some values of this random variable. He wants to know that, if the treatment is applied in the future, the average utility of the treatment will be positive.

The possible values of the random variable $\{b_m\}$ are known, but their probabilities p_m are not. In observations, outcome b_m occurs y_m times. The researcher plans to apply this treatment n times. The utility of the treatment, $\tilde{U}(L)$, is the average per-application utility. If the researcher's estimate of this utility is positive, he concludes that, on average, the treatment produces desirable results; if the estimate is zero or negative, then the researcher concludes that, on average, the treatment produces undesirable results. The researcher explicitly accounts for his attitudes toward risk in the estimation process by choosing an appropriate loss function.

For simplicity, assume that the treatment, if chosen, will be applied a lot of times (n is very large). Then, the ratios $\frac{x}{n}$ reduce to the probabilities. But keep in mind that the probabilities themselves are not known but rather follow the Dirichlet distribution. The random variable utility becomes

$$\tilde{U}(L) = \sum_{m=1}^{M} p_m u(b_m) . \tag{25}$$

4.2 Sample calculation

Suppose $u(B)$ can be any integer between -10 and 10, all *a priori* equally likely. A researcher observes the utilities of $k = 20$ treatments, shown in table XXX. He want to know whether, if he applies the treatment a very large number of times, the average per-application utility will be positive.

Let μ be the expected value of $u(B)$. In the conventional hypothesis test framework, test the following hypothesis: $H_0 : \mu \leq 0; H_A : \mu > 0$. Let $\hat{\mu}$ be the estimator of expected value μ, and let $\hat{\sigma}$ be its (estimated) standard deviation. *A priori*, if $\mu = 0$, the ratio $\dfrac{\hat{\mu}}{\hat{\sigma}}$ has the t distribution with $20 - 1 = 19$ degrees of freedom. From this, if the size of the test is 5%, the critical value of the test statistic is 1.73; if the size is 1%, the critical value is 2.54.

The observed $\hat{\mu}$ is 2.55; the observed $\hat{\sigma} = 1.30$. The obtained test statistic is 1.97. Thus, the researcher accepts the hypothesis that $\mu > 0$ at the 5% level, but not at the 1% level. It's not clear what decision the researcher should make since there is no direct correspondence between the size of the test and the researcher's attitudes toward risk.

Now, let's apply the approach developed above. Since *a priori* all outcomes are equally likely, set the prior parameter $\alpha_m = 1$ for all m. Figure 3 shows the distribution of utility of treatment, $\tilde{U}(L)$. The researcher calculates the point estimate of the utility, $U(L)$, by specifying the parameters a and c_a of the loss function that best reflect his preferences. In this example, the utility is positive under a wide range of very reasonable parameters. For instance, if the researcher has quadratic (type 2) loss with risk aversion $c_2 = 13.1$, then utility is $U(L) = 0.3$; if risk aversion is $c_2 = 22.4$, utility is $U(L) = 0.1$. A very high risk aversion of $c_2 = 52.5$ does produce a negative utility of $U(L) = -0.2$ though. Whether the researcher decides that the treatment has a desirable effect depends directly on his easily quantifiable degree of risk aversion.

4.3 Extensions

Elementary utility $u(\cdot)$ should be defined such that if b is desirable, $u(b) > 0$, whereas if b is undesirable, $u(b) < 0$. For example, if for some constant b^*, we want to show that, on average, $b > b^*$, we can set $u(b) = b - b^*$. If we want to show that, on average, $b < b^*$, one possibility is to set $u(b) = b^* - b$. The precise definition of elementary utility should reflect the relative desirability of various possible outcomes. If the estimated utility $U(L)$ is positive, there is sufficient evidence for the desired conclusion.

If we wish to show that $b \neq b^*$, first try to show that $b > b^*$ and then that $b < b^*$. If there is sufficient evidence that $b > b^*$ but not that $b < b^*$, or vice versa, that is sufficient evidence that $b \neq b^*$. If there is not sufficient evidence for either $b > b^*$ or $b < b^*$, then we cannot conclude that $b \neq b^*$.

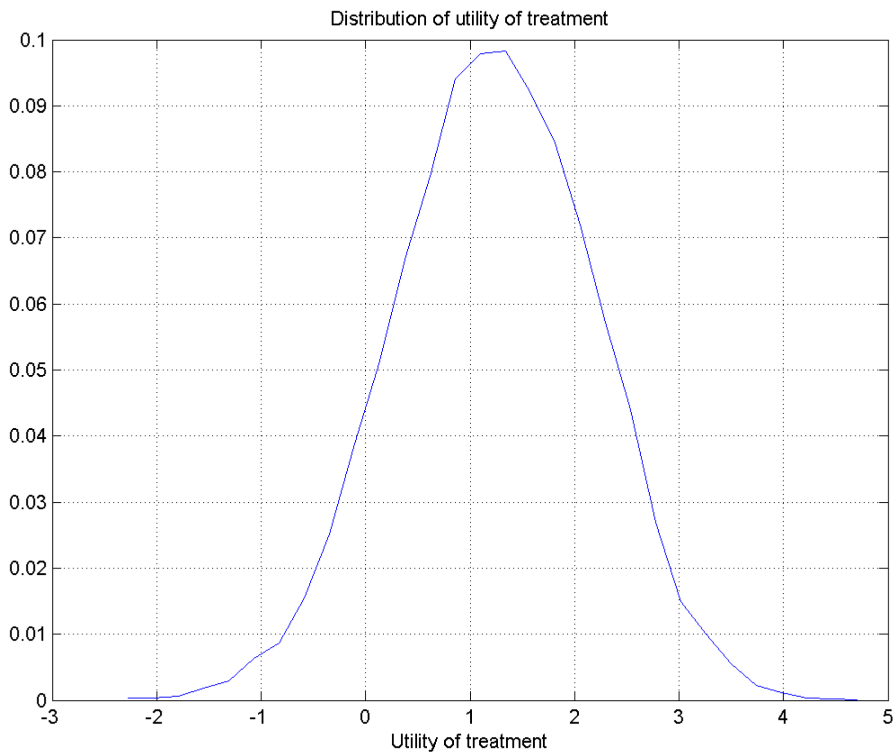

Figure 3: Distribution of utility for the sample calculation.

5 Conclusion

The utility model presented here is attractive theoretically since it is built upon solid statistical principles. The model sheds light on several well-known paradoxes. The model can also be put to good use in hypothesis testing.

References

Martin, R. (2001), The St. Petersburg paradox, *in* E. N. Zalta, ed., 'The Stanford Encyclopedia of Philosophy (Fall 2001 Edition)'.
URL: *http://plato.stanford.edu/archives/fall2001/entries/paradox-stpetersburg/*

Mehra, R. & Prescott, E. C. (1985), 'The equity premium: A puzzle', *Journal of Monetary Economics* **15**, 145–162.

Obstfeld, M. & Rogoff, K. (1996), *Foundations of International Macroeconomics*, third edn, MIT Press.

Shiller, R. J. (2000), *Irrational Exuberance*, Princeton University Press.
URL: *http://aida.econ.yale.edu/char' shiller/books.htm*