A New Example of a Closed Form Mean-Variance Representation, and Implications for the Equity Premium Puzzle

by

Keith R. McLaren
Department of Econometrics and Business Statistics
Faculty of Business and Economics
Monash University, Clayton, Vic, 3800, AUSTRALIA

First version May, 2009[1], corrected May 2011
Revised March 2012; 13 June 2013; corrected 13 December 2013

Abstract: The two sufficient conditions that allow derivation of mean-variance preferences are either quadratic (von-Neumann-Morgenstern) utility with an arbitrary probability distribution of stochastic assets, or arbitrary utility with Normally distributed assets. In the first case, the specific functional form of mean-variance preferences follows from simple algebra. The second case is a qualitative result, and the only specific functional form usually provided is the case of negative exponential utility. In this paper the specific functional form for mean-variance preferences is derived for the much more realistic example of lognormally distributed assets and constant relative risk aversion (CRRA) preferences. With this use of CRRA, the observed equity premium appears not to be a puzzle.

1. Introduction

Mean-variance preferences can be stated as a primitive, or can be derived as a special case of expected (von-Neumann-Morgenstern) utility theory. In general, the two sufficient conditions that allow this derivation are either a quadratic utility function with an arbitrary probability distribution of stochastic assets, or arbitrary utility functions with Normally distributed assets. In the first case, the specific functional form of mean-variance preferences follows constructively from simple algebra. In the second case, the general result is qualitative, following from the fact that any Normal distribution is completely characterized by its mean and variance. The only specific illustrative functional form usually provided is the case of negative exponential utility. In this paper the specific functional form for mean-variance preferences is derived for the much more realistic example of lognormally distributed assets and constant relative risk aversion (CRRA) preferences. This closed form solution does not seem to be well-known in the literature. As well as being an interesting illustrative example, it also provides empirical evidence on numerical values of the CRRA parameter. CRRA utility is the standard functional form used for the Consumption Capital Asset Pricing Model (CCAPM), for which simple empirical applications lead to the well-known Equity Premium

[1] The first version of this paper was released as a Working Paper, McLaren (2009a), and a summary version as an online Working Paper, McLaren (2009b). I thank Benjamin Fu for noting an error in equation (2) in 2011, and Diego Peñarrubia for noting an error in the equation prior to (2) in 2013.

Puzzle. It is noted that an analogous application using the closed form CRRA mean-preference function in conjunction with the standard CAPM leads to much more reasonable estimates of the implied coefficient of risk aversion.

2. Background

Mean-variance preferences are usually introduced and motivated by an argument something like the following. Let A_0 represent (non-random) initial period assets. Any investment in a portfolio of risky assets will generate a probability distribution for end of period assets, to be denoted A. Under certain axioms there exists a von-Neumann-Morgenstern utility function $\tilde{U}(A)$ such that, when evaluating uncertain prospects, the decision maker acts as if he/she maximizes $E\left[\tilde{U}(A)\right]$. See, for example, Arrow (1971) or Lengwiler (2004). The notation $\tilde{U}$ is used to emphasise that von-Neumann-Morgenstern utility is random, being a function of the random variable A, and to distinguish it from the special case of a closed form expression for expected utility as a function of specific arguments, to be denoted U .The standard properties of $\tilde{U}(\quad)$ are that it is: non-decreasing i.e. $\tilde{U}'(A) \geq 0$; and concave i.e. $\tilde{U}''(A) \leq 0$, so the decision maker prefers higher wealth to less, but is averse to risk.

Associated with $\tilde{U}$ (A) are the two standard measures of risk aversion (Arrow-Pratt):

Absolute Risk Aversion $$ARA(A) = -\frac{\tilde{U}''(A)}{\tilde{U}'(A)}$$

Relative Risk Aversion $$RRA(A) = -\frac{\tilde{U}''(A)A}{\tilde{U}'(A)}.$$

Behaviour toward risk is preserved under linear transformations of $\tilde{U}(\quad)$ (because of the linearity of the expectation operator). In general, it is argued that *ARA* should be decreasing with the level of assets, but *RRA* could be close to constant, with reasonable values satisfying say $0 \leq RRA \leq 6$. Although these definitions are quite standard, they are in fact a little strange. For example, if end of period assets are nonstochastic, then the preference ordering is based on $E\left[\tilde{U}\left(A\right)\right] = \tilde{U}\left(A\right)$ and a measure of risk aversion attaches to a risk-free choice. Conversely, if end of period assets A are stochastic, in principle the distribution of A is arbitrary. However, it is well known that for an arbitrary probability distribution there is no unique scalar measure of riskiness of a random variable. And yet both *ARA* and *RRA* provide a scalar measure of aversion to "risk", somehow defined. What is clear is that both *ARA* and *RRA* are measures of curvature of the von-Neumann-Morgenstern utility function, and will hence enter into any calculation of aversion to risk.

There are three simple but popular examples that are used to illustrate preferences based on von-Neumann-Morgenstern utility functions:

(a) $\tilde{U}(\quad)$ a quadratic function:

$$\tilde{U}(A) = \alpha A - \frac{\beta}{2}A^2 \qquad \alpha, \beta > 0, A < \frac{\alpha}{\beta}.$$

In this case $ARA = \dfrac{\beta}{\alpha - \beta A}$ and $RRA = \dfrac{\beta A}{\alpha - \beta A}$ and it is argued that increasing *ARA* would seem to be counterintuitive.

(b) $\tilde{U}(\quad)$ a negative exponential function:

$$\tilde{U}(A) = \gamma - \delta e^{-\eta A} \qquad \gamma, \delta, \eta, > 0\,.$$

In this case $ARA = \eta$ and $RRA = \eta A$, and constant *ARA* would appear counterintuitive. Note that the parameters δ and γ are in fact redundant, being the parameters of a positive linear transformation.

(c) $\tilde{U}(\quad)$ a Constant Relative Risk Aversion (CRRA) function:

$$\tilde{U}(A) = \frac{A^{1-\gamma}}{1-\gamma} \qquad \gamma > 0.$$

In this case $ARA = \gamma / A$ and $RRA = \gamma$, and decreasing *ARA*, but constant *RRA*, are appealing as an illustrative example. The CRRA function is the workhorse of macrofinance, in models based on the CCAPM, but is not usually used in a mean-variance context.

It is then commonly argued that in at least two special cases the expected value of the von-Neumann-Morgenstern utility function can be replaced by a function of mean and variance alone, the mean-variance function,

i.e. replace the criterion $E[\tilde{U}(A)]$ by $U(\mu_A, \sigma_A^2)$, $\qquad U_1 > 0,\ U_2 < 0.$

U_1 denotes the partial derivative with respect to μ_A, and U_2 the partial derivative with respect to σ_A^2. These two special cases are:

(i) For arbitrary probability distributions on *A*, let $\tilde{U}(\quad)$ be quadratic, example (a) above.

Then $E\left[\tilde{U}(A)\right] = \alpha E(A) - \dfrac{\beta}{2}E(A^2) = \alpha\mu_A - \dfrac{\beta}{2}(\mu_A^2 + \sigma_A^2) = U(\mu_A, \sigma_A^2)$

which is clearly a mean-variance function, quadratic in mean and linear in variance. In this case the existence of a mean-variance function follows by construction. However, its relation to measures of risk aversion is instructive. The function *U* can be thought of as representing a preference ordering over two characteristics of random assets, with isoquants representing isoutility contours. To map out an indifference curve in mean-standard deviation space (standard deviation rather than variance to keep units comparable), set *U* equal to a constant

and take total differentials: $\alpha d\mu_A - \beta\mu_A d\mu_A = \beta\sigma_A d\sigma_A$ which allows the derivation of the marginal rate of substitution $MRS = \frac{d\mu_A}{d\sigma_A} = \frac{\beta\sigma_A}{(\alpha - \beta\mu_A)} = ARA(\mu_A)\times\sigma_A$ (ignoring the usual negative sign used in defining MRS, since here isoquants are positively sloped). But as noted above, the underlying quadratic von-Neumann-Morgenstern utility function is unattractive, and thus this mean-variance function is also unattractive.

(ii) For arbitrary preferences, assume $A \sim N(\mu_A, \sigma_A^2)$ which, being Normally distributed, must be characterised solely by its mean and variance. Hence it must be the case that $E[\tilde{U}(A)] = U(\mu_A, \sigma_A^2)$, an implication of Normality. However, this is a qualitative result, and the explicit functional form of $U(\mu_A, \sigma_A^2)$ will not in general be known. An exception is the popular illustrative example of function (b) above. Assume $\tilde{U}(\quad)$ is negative exponential, $\tilde{U}(A) = -e^{-\eta A}$, and A is Normally distributed, $A \sim N(\mu_A, \sigma_A^2)$. Appealing to results on the lognormal distribution, if $X \sim N(\mu_X, \sigma_X^2)$ and $Y = e^X$, then Y is lognormally distributed (i.e. the natural logarithm of Y is Normal) and $\mu_Y = e^{\mu_X + \frac{1}{2}\sigma_X^2}$. In the negative exponential example $-\eta A \sim N(-\eta\mu_A, \eta^2\sigma_A^2)$ and hence by this basic result

$$E\left[\tilde{U}(A)\right] = -E(e^{-\eta A}) == -e^{[-\eta\mu_A + \frac{1}{2}\eta^2\sigma_A^2]} = -e^{-\eta[\mu_A - \frac{1}{2}\eta\sigma_A^2]}$$

$$= -e^{[-\eta\mu_A + \frac{1}{2}\eta^2\sigma_A^2]} = -e^{-\eta[\mu_A - \frac{1}{2}\eta\sigma_A^2]}.$$

Hence, by monotonicity, maximizing $E\left[\tilde{U}(A)\right]$ is equivalent to maximizing the function

$$U(\mu_A, \sigma_A^2) = \mu_A - \frac{\eta}{2}\sigma_A^2$$

with respect to mean and variance. Note that this function is linear in mean <u>and</u> variance. Again, its relation to measures of risk aversion is instructive. The indifference curves in mean-standard deviation space satisfy $d\mu_A = \eta\sigma_A d\sigma_A$ which gives $MRS = \eta\sigma_A = ARA\times\sigma_A$. But again, as noted above, the underlying negative exponential von-Neumann-Morgenstern utility function is unattractive, and thus this mean-variance function is also an unattractive illustrative example. In fact the qualitative result that Normality of assets generates mean-variance expected utility extends to the *linear distribution class* of probability distributions for assets (see Meyer(1987), who also shows for example that under CRRA the corresponding preferences in mean-standard deviation space are homothetic). Anticipating the results of the next section, note that the lognormal distribution is not a member of the linear distribution class. In passing, we note two other ways to understand the relations between risk aversion and variances. In derivations of the CAPM based on mean-variance

preferences, one step involves equating $\frac{-2U_2A_0}{U_1}$ with the observed market price of risk per unit of nondiversifiable risk. Another illuminating derivation is to use a Taylor series approximation of expected utility around a small risk to generate the *Arrow-Pratt approximations*. See, for example, Lengwiler (2004), pp. 98-99.

In the above, preferences over risky outcomes have been motivated in terms of preferences over the uncertain levels of *assets*. In portfolio analysis preferences are often expressed in terms of the uncertain levels of *returns* on the portfolio. Given the fixed initial level of assets A_0, the two ideas are equivalent, provided that returns are expressed as ordinary compound returns. Assume that the uncertainty is over the return R on a given portfolio. Then $A = A_0(1+R)$ for A_0 fixed, and hence

$$\mu_A = A_0(1+\mu_R) = A_0 + A_0\mu_R \quad ; \quad \sigma_A^2 = A_0^2\sigma_R^2$$

and so means and variances of assets are positively linearly related to the means and variances of returns, using only the level of initial assets A_0, and it is easy to translate mean-variance preferences over assets to mean-variance preferences over compound returns (plus the initial level of assets A_0), and vice-versa. In fact

$$U\left(\mu_A, \sigma_A^2\right) = U\left(A_0 + A_0\mu_R, A_0^2\sigma_R^2\right) = U^R\left(\mu_R, \sigma_R^2, A_0\right)$$

where the symbol $U^R(.)$ indicates mean-variance utility over means and variances of returns *R*. This is not the case for continuously compounded returns, which are used below.

3. A New Result

A third, much more realistic, example of a closed form expression for a mean-variance function, that does not seem to appear in standard text books or journal articles, is the following. Assume that the von-Neumann-Morgenstern utility function is example (c) above, CRRA, and that it is the continuously compounding rate of return (often called, confusingly, "log returns"), *r*, that are Normally distributed. (CRRA preferences are used for von-Neumann-Morgenstern preferences in Courakis (1989), but the corresponding mean-variance function is derived as an approximation using a quadratic approximation to the underlying utility function.) Normality of continuously compounded returns, and hence LogNormality of assets, is a far more attractive example than Normality of compound returns and assets. For example, Normality of log returns implies that asset prices are LogNormally distributed, and hence constrained to be nonnegative, a restriction consistent with limited liability. Further, Normality is preserved under temporal aggregation of log returns. For more details, see Fama (1976). In the new example , more general results from the LogNormal distribution will be used. These results are that, if $X \sim N(\mu_X, \sigma_X^2)$ and $Y = e^X$, then

$$\mu_Y = e^{\mu_X + \frac{1}{2}\sigma_X^2}; \sigma_Y^2 = \left(e^{\sigma_X^2} - 1\right) e^{2\mu_X + \sigma_X^2}$$

and, conversely, that

$$\sigma_X^2 = \ln\left(1 + \frac{\sigma_Y^2}{\mu_Y^2}\right) \; ; \; \mu_X = \ln \mu_Y - \frac{1}{2}\ln\left(1 + \frac{\sigma_Y^2}{\mu_Y^2}\right).$$

These results will be used repeatedly, with X representing either r or $(1-\gamma)r$, with $Y = (1+R)$, and with $A = A_0 Y$.

Now consider the CRRA $\tilde{U}(A) = \dfrac{A^{1-\gamma}}{1-\gamma}$ and set $A = A_0(1+R) = A_0 e^r$ with $r \sim N\left(\mu_r, \sigma_r^2\right)$.

Then

$$\begin{aligned}
\tilde{U}(A) &= \frac{A^{1-\gamma}}{1-\gamma} = \frac{1}{1-\gamma} e^{(1-\gamma)\ln A} \\
&= \frac{1}{1-\gamma} e^{(1-\gamma)\ln A_0 + (1-\gamma)r} \\
&= K e^{(1-\gamma)r} \quad \text{with } K = \frac{1}{1-\gamma} e^{(1-\gamma)\ln A_0}.
\end{aligned}$$

Since $(1-\gamma)r$ is also Normal, $\tilde{U}(A)$ is scaled lognormal, and hence $E\left[\tilde{U}(A)\right]$ has the closed form expression

$$(1) \qquad E\left[\tilde{U}(A)\right] = K e^{(1-\gamma)\mu_r + \frac{1}{2}(1-\gamma)^2 \sigma_r^2}$$

a function of the parameter of the CRRA utility function, the mean and variance of r, and of the initial level of assets A_0. Interestingly, in this case, the slope of expected utility with respect to mean return is positive, but the slope with respect to variance of return is positive or negative according to whether $\gamma < 1$ or $\gamma > 1$. In particular, in the case of log utility ($\gamma = 1$ and writing (1) in the extended Box-Cox form to derive the limiting expression) the variance of log returns doesn't even enter into expression (1). This is because the variance of r directly affects both the mean <u>and</u> the variance of assets A, (which exactly compensate each other in the log case) and illustrates the fact noted above that mean-variance analysis should strictly be expressed in terms of the parameters of the distribution of A, and this can be also expressed monotonically in terms of returns only in the ordinary compound return case. Thus the above result (1) is *not* a true mean-variance representation. However, a simple transformation allows the application of the extended lognormal results to assets. Define the

gross rate of return $Y = (1+R) = e^r$, and since $A = A_0(1+R) = A_0 Y = A_0 e^r$ then $\mu_A = A_0\mu_Y; \sigma_A^2 = A_0^2\sigma_Y^2$, and hence by substitution

$$\sigma_r^2 = \ln\left(1+\frac{\sigma_Y^2}{\mu_Y^2}\right) = \ln\left(1+\frac{\sigma_A^2/A_0^2}{\mu_A^2/A_0^2}\right) = \ln\left(1+\frac{\sigma_A^2}{\mu_A^2}\right)$$

$$\mu_r = \ln\mu_A - \ln A_0 - \frac{1}{2}\ln\left(1+\frac{\sigma_A^2}{\mu_A^2}\right).$$

Substituting these in the expression for expected utility then gives the fundamental result

(2) $$E\left[\tilde{U}(A)\right] = \frac{1}{1-\gamma}e^{(1-\gamma)\ln\mu_A - \frac{1}{2}\gamma(1-\gamma)\ln\left(1+\frac{\sigma_A^2}{\mu_A^2}\right)} = U(\mu_A, \sigma_A^2)$$

which is a closed form, and a much more realistic, illustrative (or even empirically applicable) example of a mean-variance utility function. It can be seen that the slope of this function with respect to μ_A is positive, and the slope with respect to σ_A^2 is negative, consistent with the usual assumptions of mean-variance preferences. (Although this is not necessary – see Bigelow (1993), Hadar and Russell (1969), Meyer (1987).) In terms of the ordinary compound return *R* the corresponding expression is

(3) $$E\left[\tilde{U}(A)\right] = \frac{1}{1-\gamma}e^{(1-\gamma)\ln(A_0+A_0\mu_R) - \frac{1}{2}\gamma(1-\gamma)\ln\left(1+\frac{\sigma_R^2}{1+2\mu_R+\mu_R^2}\right)} = U^R(\mu_R, \sigma_R^2, A_0)$$

4. Implications for the Equity Premium Puzzle

The underlying von-Neumann-Morgenstern utility function used to derive (3), the CRRA form, is the standard function used in empirical applications of the Consumption Capital Asset Pricing Model (CCAPM), in particular in the discussion of the Equity Premium Puzzle. For a survey, see Kocherlakota (1996). Using the unconditional expected value form of the CCAPM implications, the Euler equations, and sample data from 1889 to 1978, Kocherlakota shows that for reasonable values of other parameters in the model, equating sample means to unconditional expectations results in an unreasonably high estimate of the parameter corresponding to CRRA. In a similar way, the implications of (3) can also be used to derive an implied measure of risk aversion from sample averages. In the standard derivation of the CAPM, a preliminary result allows the measure of risk aversion to be replaced by the market price of risk. This result takes the form

(4) $$\frac{-2U_2^R}{U_1^R} = \frac{\mu_M - R_F}{\sigma_M^2}$$

and is usually used to derive the CAPM relations that express excess expected return of an individual asset as being proportional to that asset's beta, with the factor of proportionality being the market price of risk $\mu_M - R_F$, without any reference to a specific functional form for $U^R(\mu_R, \sigma_R^2, A_0)$. However, with the specific functional form (3), this intermediate result gives the possibility of using observed data on market risk and excess return to estimate the coefficient of relative risk aversion, analogous to what is done in the CCAPM literature. Using sample means and variances from Dimson, Marsh and Staunton (2008), covering the period 1900-2005 for the United States, to estimate the numerator and denominator of the Expected Excess Return to Variance Ratio on the rhs of (4), and evaluating the lhs of (4) for the functional form in (3), allows a corresponding calibration of the CRRA risk aversion parameter of $\hat{\gamma} = 6.5$. While this is toward the upper bound, it is perfectly consistent with the usual a priori values considered reasonable for such a measure of risk aversion.[2]

5. Conclusion

The closed form mean-variance expression (2) is based on a reasonably realistic CRRA von-Neumann Morgenstern utility function, and exploits the properties of the Normal distribution by associating Normality more realistically with the continuously compounded rate of return, rather than with the actual distribution of assets. Thus it is a far more appealing, though admittedly a little more complex, specific example of a mean-variance preference function derived from maximizing behaviour than the examples based on either quadratic preferences or negative exponential preferences. In the same way that the CCAPM with CRRA utility allows estimation of a measure of risk aversion from first order conditions, the Euler equations, this closed form expression allows an analogous estimate based on the CAPM, and the corresponding estimate is far more realistic.

Appendix

Well-known results are that, if $X \sim N(\mu_X, \sigma_X^2)$ and $Y = e^X$, then

(5) $$\mu_Y = e^{\mu_X + \frac{1}{2}\sigma_X^2}$$

[2] As with the calculations of risk aversion for the CCAPM, this is very much a "back of the envelope" calculation based on simplifying assumptions. The rhs of equation (4) relates to the "Market" portfolio. In the CAPM, the market portfolio consists of both equities and bonds, so adjusting the figures from DMS makes use of an estimate of the relative sizes of the bond and equity markets. The capitalization of the US Stock market, according to August 2008 figures from Wilshire Associates, is $15.59 trillion. The capitalization of the US Bond market, according to 2007 figures released by FINRA, is in approximately $27 trillion.

(6) $$\sigma_Y^2 = \left(e^{\sigma_X^2} - 1\right) e^{2\mu_X + \sigma_X^2}.$$

(See, for example, Aitchison, J. and J.A.C. Brown, (1957).) From (5), (6) can be written as

$$\sigma_Y^2 = \left(e^{\sigma_X^2} - 1\right) e^{2\mu_X + \sigma_X^2} = \left(e^{\sigma_X^2} - 1\right) e^{2(\mu_X + \frac{1}{2}\sigma_X^2)} = \left(e^{\sigma_X^2} - 1\right) \mu_Y^2$$

Rearranging

$$\sigma_Y^2 = \left(e^{\sigma_X^2} - 1\right) e^{2\mu_X + \sigma_X^2} = \left(e^{\sigma_X^2} - 1\right) e^{2(\mu_X + \frac{1}{2}\sigma_X^2)} = \left(e^{\sigma_X^2} - 1\right) \mu_Y^2$$

$$e^{\sigma_X^2} = 1 + \frac{\sigma_Y^2}{\mu_Y^2} \quad \text{or} \quad \sigma_X^2 = \ln\left(1 + \frac{\sigma_Y^2}{\mu_Y^2}\right).$$

From (5),

$$\mu_X = \ln\left(\mu_Y\right) - \frac{1}{2}\sigma_X^2 = \ln\left(\mu_Y\right) - \frac{1}{2}\ln\left(1 + \frac{\sigma_Y^2}{\mu_Y^2}\right)$$

References

Aitchison, J. and J.A.C. Brown, *The Lognormal Distribution: with special reference to its uses in economics*, Cambridge University Press, Cambridge, 1957.

Arrow, Kenneth, *Essays in the Theory of Risk Bearing*, Markham Publishing Company, Chicago, 1971.

Bigelow, John P., "Consistency of Mean-Variance Analysis and Expected Utility Analysis", *Economics Letters* 43 (1993) pp. 187-192.

Courakis, Anthony S., "Does Constant Relative Risk Aversion Imply Asset Demands that are linear in Expected Returns*?", Oxford Economic Papers*, 41, 3, (1989) pp. 553-566.

Dimson, Elroy, Paul Marsh, and Mike Staunton, "The Worldwide Equity Premium : A Smaller Puzzle." In *Handbook of the Equity Risk Premium*, ed. Rajnish Mehra, 467–514. Elsevier, Amsterdam, 2008.

Fama, Eugene, *Foundations of Finance*, Basil Blackwell, Oxford, 1976.

Hadar, Josef and William R. Russell, "Rules for Ordering Uncertain Prospects", *American Economic Review*, 59, 1, (1969) pp. 25-34.

Kocherlakota, NR. , "The Equity Premium: It's Still a Puzzle." *Journal of Economic Literature* 34, 1, (1996) pp.42–71.

Lengwiler, Y., *Microfoundations of Financial Economics*, Princeton University Press, Princeton, 2004.

McLaren, Keith R., "A New Example of a Closed Form Mean-Variance Representation" Working Paper 1/2009, Dept of Econometrics and Business Statistics, Monash University, January 2009a.

McLaren, Keith R, "Closed Form Mean-Variance Representations", *Medium for Econometric Applications*, Volume 17, Edition 2, 2009b, pp. 14-18

Meyer, Jack, "Two-moment decision Models and Expected Utility Maximization", *American Economic Review*, 77, 3, (1987) pp. 421-430.