Expected Utility, Mean-Variance, and Stochastic Dominance

by James R. Garven*

September 12, 2016

Updated: September 18, 2019

1 The Expected Utility Model

The expected utility model is the *foundation* for decision-making under risk and uncertainty. It requires making the following calculation:

$$E(U(W)) = \sum_{s=1}^{n} p_s U(W_s), \tag{1}$$

where $U(W)$ = utility of wealth, p_s= the probability that state s occurs, and W_s= state-contingent wealth.

The "good" news is that the expected utility model always "works", in the sense that it provides a logically consistent way to evaluate risk which lines up nicely with one's preferences. We can model all sorts of different risk attitudes using the expected utility framework; e.g., we can assume that people are risk averse, risk neutral, or risk loving. However, we typically assume that people are risk averse, but that they differ with respect to the *degree* to which they are risk averse.

So much for the good news. The "bad" news about the expected utility model is that one must know the *form* of the utility function in order to carry out these calculations. One way that the form of the utility function can be determined is by inference from a Risk Tolerance Quiz. Another way around this problem is to formulate a set of decision rules (specifically, the mean-variance and stochastic dominance models) which make it possible to determine how *arbitrarily risk averse* individuals compare different risks.

2 The Mean-Variance Model

Commonly in finance, we assume that risk = variance. Indeed, most of your other finance courses characterize the "risk-return tradeoff" as the tradeoff between accepting higher (lower) variance in

*James R. Garven is the Frank S. Groner Memorial Chair in Finance and Professor of Finance & Insurance at Baylor University (Address: Foster 320.39, One Bear Place #98004, Waco, TX 76798, telephone: 254-307-1317, e-mail: James_Garven@baylor.edu).

order to possibly earn higher (lower) expected returns. This is also commonly referred to as the mean-variance model. But where does this mean-variance model come from, and can we trust that it will always enable us to make sound decisions about risk and reward? It turns out that the answer to this question, like many others in finance, is that it "all depends".

First, let's consider the logical connection between the expected utility and the mean-variance decision models. On pp. 9–13 of the Decision Making under Risk and Uncertainty, part 3, I show that expected utility is positively related to expected wealth and skewness, and negatively related to variance and to how "fat-tailed", or kurtotic the risk is; i.e.,

$$E(U(W)) = f(\underbrace{\text{mean } (E(W))}_{+}, \underbrace{\text{variance } (\sigma_W^2)}_{-}, \underbrace{\text{skewness } (Sk_W)}_{+}, \underbrace{\text{kurtosis } (K_W)}_{-}). \tag{2}$$

The mean-variance model implicitly assumes that the risks being considered are neither skewed nor fat-tailed, which means that the last two parameters on the right hand side of equation (2) are zero. Thus we can rewrite equation (2) as follows:

$$E(U(W)) = f(\underbrace{\text{mean } (E(W))}_{+}, \underbrace{\text{variance } (\sigma_W^2)}_{-}). \tag{3}$$

Now suppose you are comparing two symmetric (i.e., not skewed) and thin-tailed (i.e., not kurtotic) risks called x and y (certainly if x and y are both normally distributed random variables, then they have these properties). Equation (3) implies the following set of "dominance" conditions:

$$\begin{aligned} E(x) &> E(y) \text{ and } \sigma_x^2 < \sigma_y^2; \\ E(x) &> E(y) \text{ and } \sigma_x^2 = \sigma_y^2; \text{ and} \\ E(x) &= E(y) \text{ and } \sigma_x^2 < \sigma_y^2. \end{aligned} \tag{4}$$

Furthermore, if x and y have the same mean and variance, then this is an "indifference" condition. Finally, if x has a higher mean *and* a higher variance, this is an "ambiguity" condition. In this latter case, the mean-variance model cannot provide us with any guidance as to which risk is preferred, and we must resort to either the expected utility model or the stochastic dominance model, which we discuss next.

3 The Stochastic Dominance Model

The math behind stochastic dominance is summarized in my optional reading entitled "Technical Note on Stochastic Dominance and Expected Utility". What I show there is that if there are two risks under consideration and the expected utility of risk x is greater than the expected utility of risk y, then it *must* be the case that

$$\sum_{s=1}^{n} (F(y_s) - F(x_s)) > 0, \tag{5}$$

where $F(y_s)$ = the cumulative distribution for y_s and $F(x_s)$ = the cumulative distribution for x_s. This is commonly referred to as "second order stochastic dominance". First order stochastic dominance says that risk x first order stochastically dominates risk y if $F(y_s) \geq F(x_s)$ for all s. Obviously, if this condition holds, then so must inequality (5) above.

The good news in the case of the stochastic dominance model is that it "works" in most cases, and it is less restrictive than the mean-variance model because it doesn't matter how skewed or fat-tailed the risk might be. The bad news with stochastic dominance is that this model requires being able to map out the entire probability distribution. Finally, there are cases where there is neither first nor second order stochastic dominance. If and when this happens, one has no choice but to resort to using the expected utility model.