Coherent Risk Measures

A "Coherent" risk measure satisfies a number of seemingly desirable properties.

Preliminaries
Technical discussions of coherent risk measures impose a number of 'pre-cursor' conditions for the "probability space" of the random variables being measured. This space is usually specified as a 'triple' (Ω, $\mathscr{F}$, $\mathcal{P}$) consisting of a *sample space* (Ω), a family of subsets of the sample space ($\mathscr{F}$) representing *observable random events* associated with the sample space, and a *probability measure* ($\mathcal{P}$) which maps from the set of observable events to the [0,1] portion of the real line – effectively assigning to each observable random event a probability that is between 0 and 1 (inclusive).

'Probability theory' does not attempt to estimate the probability that any event will occur but instead assumes that 'somehow' the probabilities have already been determined and are available in advance –> this works well in gambling and coin flipping applications but it is less clear for situations where the random events are related to 'real world' business applications.

Key Question: Can probabilities be objectively measured or are probabilities inherently subjective?

Is it possible to fully define the sample space, the set of possible outcomes? Uncertainty is associated with events that do not appear in the sample space or, less rigidly, the difficulties of accurately defining a sample space.

Financial Risk Measures and Risk Metrics
Risk is a fundamental concept in both financial economics and actuarial science. It is conventional to distinguish *what is being measured* – the risk metric – and *how the risk metric is being measured* – the risk measure. Examples of risk metrics include: credit risk, liquidity risk, market risk, operational risk, general business risk. Examples of risk measures include: standard deviation, value-at-risk, expected shortfall, tail conditional probability.

Coherent Risk Measures
A *coherent risk measure* imposes specific, seemingly innocuous, technical requirements for a risk measure (ρ):

1) $\rho\,[0] = 0$

The risk of nothing is zero. This is referred to as *normalization*.

2) Consider two random outcomes, Z_1 and Z_2, e.g., the returns from two portfolios or capital assets. If $Z_2 > Z_1$ in all feasible future states then $\rho[Z_2] < \rho[Z_1]$.

This property is referred to as *monotonicity*. For financial applications, this implies that a security that always has higher return in all future states has less risk of loss.

3) Consider two random outcomes, Z_1 and Z_2, e.g., the returns from two portfolios or capital assets, then $\rho[Z_2] + \rho[Z_1] \geq \rho[Z_2 + Z_1]$.

This is referred to as *sub-additivity* or, in financial applications, the *diversification principle*. In effect, diversification is risk reducing.

4) For any random outcome Z_i and constant value $\gamma \geq 0$ then $\rho[\gamma Z_i] = \gamma\, \rho[Z_i]$.

This is referred to as *positive homogeneity*. In effect, if a portfolio or capital asset is, say, doubled ($\gamma = 2$), then the risk will also be doubled.

5) For any random outcome Z_i the addition of an additional outcome with a certain positive return (κ) will reduce the risk by that amount: $\rho[Z_i + \kappa] = \rho[Z_i] - \kappa$.

This is referred to as *translation invariance*. In effect, if an amount of cash κ (or risk free asset) is added to a portfolio, then the risk is reduced by that amount.

Any sensible risk measure needs to obey normalization, monotonicity and translation invariance.