this discussion, we assume the wealth levels are continuous and that the customer has preferences for wealth that satisfy the conditions of Theorem E.4. Also, assume the lotteries are now continuous distributions F on $\Re$.[3]

Consider now any given lottery F (a distribution on possible wealth outcomes) and μ_F denote the mean of the distribution. A customer is said to have *risk-averse preferences* if he prefers the certain wealth μ_F to the lottery F itself for all possible lotteries F. That is, the customer always prefers the certainty of receiving the expected wealth rather than a gamble with the same mean. The customer is said to have *risk-seeking preferences* if he prefers the gamble F to the certain outcome μ_F for all F. Finally, he has *risk-neutral preferences* if he is indifferent between the lottery F and the certain reward μ_F.[4] We then have the following result:

THEOREM E.5 *A customer's preference $\succ$ for lotteries exhibits risk-aversion (risk-seeking) behavior if and only if their von Neumann-Morgenstern utility function $u(w)$ is concave (convex). Their preference is risk-neutral if and only if $u(w)$ is affine.*

Thus, risk preferences are linked directly to concavity or convexity of the customer's utility function. The reason is quite intuitive; with a concave utility function for wealth, a customer gains less utility from a given increase in wealth than he loses in utility from the same decrease in wealth. Hence, the upside gains produced by the volatility in outcomes do not offset the downside losses, and customers therefore prefer the certain average to the uncertain outcomes of the lottery. Since most customers have a decreasing marginal utility for wealth, risk aversion is a good assumption in modeling customer behavior.

Still, the concept of risk aversion has to be addressed with care in operational modeling. While it is true that most customers are risk-averse when it comes to *large* swings in their wealth, often the gambles we face as consumers have a relatively small range of possible outcomes relative to our wealth. For example, a customer may face a price risk in buying a CD or book online. However, the differences in prices for such items are extremely small compared to his total wealth. In such cases, the utility function is "almost linear" in the range of outcomes affecting the decision and the customer tends to behave "as if" he were risk-neutral.[5] Similar statements apply to firms. Generally, they are risk-averse too, but for decisions and gambles that involve "small" outcomes relative to their total wealth and income, they tend to be approximately risk-neutral. Hence, risk-neutrality is a reasonable assumption in operational models and, indeed, is the standard assumption in RM practice.

[3]The extension of Theorem E.4 to the continuous case requires some additional technical conditions that are beyond the scope of this chapter. See Kreps [313].

[4]Note that a customer's preferences may not fall into any of these three categories. For example, many consumers take out fire insurance, preferring a certain loss in premium payments every year to the gamble between making no payments but potentially loosing their house, yet simultaneously play their local state lottery, which has an expected loss but provides a small probability of a large wealth pay-off. Such behavior violates a strict risk preference.

[5]Formally, one can see this by taking a Taylor series approximation of the utility function about the customer's current wealth to; the first-order approximation is affine, corresponding to risk-neutrality.