Risk Analytics

Machine Learning and Optimization
for Data-Driven Decision Making

Fernando S. Oliveira

Draft version — March 23, 2026

Chapter 3
The Architecture of Risk Preferences

Chapter 3

The Architecture of Risk Preferences

3.1 Introduction

The previous chapter examined the psychology of risk and uncertainty, focusing on how real decision makers perceive, frame, and respond to uncertain outcomes. Research in psychology and behavioral economics shows that judgment under risk is shaped by heuristics, emotions, framing effects, and other cognitive mechanisms that systematically influence choice [8]. These findings make clear that observed decisions often depart from the predictions of classical rational models.

This chapter turns from that descriptive perspective to the normative structure of choice under uncertainty. Instead of asking how people in fact behave, it asks how risky alternatives can be evaluated within a coherent formal framework. This is the central task of normative decision theory: to specify how uncertain prospects should be compared once outcomes, probabilities, and preferences are given.

That task is fundamental to risk analysis. Probabilistic models describe uncertainty, but probabilities alone do not determine what should be chosen. A decision rule is still required. In modern risk analysis, that evaluative role is typically played by preference-based criteria such as expected value or expected utility [5, 10]. These criteria are central not only in economics, finance, insurance, and operations research, but also in computational risk analytics, where forecasts and optimization models must ultimately be linked to explicit objectives.

Figure 3.1 summarizes this logic. An uncertain environment is represented probabilistically, evaluated through a preference structure, and translated into a decision criterion. The key insight is that risk analysis requires both a model of uncertainty and a model of preference.

The most influential normative framework for this purpose is expected utility theory. It provides a general way to evaluate risky prospects by

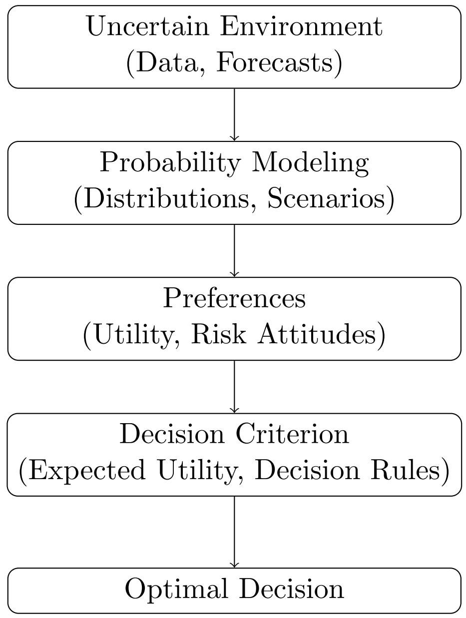

Figure 3.1: Conceptual pipeline of quantitative risk analytics. Uncertainty is represented probabilistically and then evaluated using preference-based decision criteria.

combining probabilities and utilities into a single criterion of choice. Within this framework, attitudes toward risk are represented through the curvature of the utility function, making it possible to distinguish formally between risk neutrality, risk aversion, and risk seeking.

The chapter proceeds as follows. It begins with expected value as the simplest benchmark and explains why it is not sufficient as a general criterion of rational choice. It then introduces expected utility theory, develops the relation between utility curvature and risk attitudes, and explains the concepts of certainty equivalent and risk premium. The chapter next turns to the modeling of risk aversion, including the role of standard utility specifications such as CARA and CRRA, then considers the representation of risk-seeking behavior, and finally discusses how risk attitudes are measured in practice.

3.2 Risk-Neutral Decision Making

The simplest normative model of decision making under uncertainty evaluates risky prospects according to their expected value. If a decision maker faces possible outcomes x_i with associated probabilities p_i, the expected value of the prospect is

$$E[x] = \sum_i p_i x_i.$$

Expected value therefore provides the most elementary criterion for evaluating uncertain outcomes and serves as a natural benchmark in many areas of economic analysis and decision theory [5, 11].

A decision maker who evaluates alternatives solely according to expected value is said to be *risk neutral*. Risk neutrality implies that only the average payoff matters, while the dispersion of outcomes around that average is irrelevant. Any two prospects with the same expected payoff are therefore treated as equivalent.

To illustrate, consider a choice between receiving \$50 with certainty or accepting a gamble that yields \$100 with probability 0.5 and \$0 with probability 0.5. The expected value of the gamble is

$$E[x] = 0.5(100) + 0.5(0) = 50.$$

A risk-neutral decision maker is therefore indifferent between the two options.

Risk neutrality is analytically important because it provides a benchmark against which richer models can be compared. However, it is too restrictive to describe many economically important choices, because in practice individuals often care not only about the mean outcome but also about the uncertainty surrounding it. Equal expected values need not imply equal desirability once the decision maker cares about the shape of the risk rather than its average payoff alone.

3.3 Why Expected Value Is Not Enough

Expected value is useful as a benchmark, but it is too narrow to serve as a general theory of rational choice under uncertainty. It treats all prospects with the same mean as equivalent, regardless of how outcomes are distributed or how individuals value changes in wealth. As a result, it cannot explain why individuals may prefer certainty to a fair gamble, pay for insurance, or reject highly volatile prospects despite attractive averages.

One of the most influential objections to expected value reasoning is the *St. Petersburg paradox*, first analyzed by [4]. The paradox shows that expected value can generate implausible prescriptions and therefore motivates the search for a more satisfactory criterion.

The St. Petersburg Paradox

Consider a lottery in which a fair coin is tossed repeatedly until it lands heads. If the first head appears on the nth toss, the player receives 2^n dollars. The probability of this event is $1/2^n$.
The expected value of the lottery is

$$E[x] = \sum_{n=1}^{\infty} \frac{1}{2^n} 2^n = \sum_{n=1}^{\infty} 1 = \infty.$$

According to expected value reasoning, a rational person should be willing to pay any finite amount to enter the game.
In practice, however, most people would only pay a modest amount.
Bernoulli's resolution was that outcomes should not be evaluated by monetary amounts alone, but by their utility. If utility rises with wealth at a diminishing rate, the expected utility of the lottery becomes finite and the paradox disappears [4].

The paradox is important because it shows that rational evaluation under uncertainty requires more than averaging payoffs. It requires a model in which the subjective value of outcomes can differ from their monetary magnitude. This is precisely the role played by expected utility theory.

3.4 Expected Utility Theory and Its Foundations

Expected utility theory provides the standard normative framework for evaluating risky prospects. Its central idea is that uncertain alternatives should be assessed not by the expected value of outcomes themselves, but by the expected value of a utility function defined over those outcomes.

If x_i denotes outcomes, p_i their probabilities, and $u(x_i)$ the utility associated with each outcome, expected utility is defined as

$$EU = \sum_i p_i u(x_i).$$

A rational decision maker is assumed to prefer the alternative with the highest expected utility.

The modern axiomatic formulation of expected utility theory was developed by von Neumann and Morgenstern [15]. Its significance lies not only in the formula above, but in the fact that, under a set of consistency axioms on preferences over lotteries, those preferences can be represented as the maximization of expected utility. In this way, expected utility theory provides the normative foundation for utility-based analysis of risky choice.

This framework resolves the main limitation of expected value reasoning by allowing the value of outcomes to depend on preferences rather than money alone. It therefore makes it possible to represent systematically the preference for certainty, the acceptance of some risks, and the rejection of others. Once utility rather than money is the object of expectation, the shape of the utility function becomes the key determinant of how risk is evaluated.

Despite its elegance, standard expected utility theory does not perfectly describe observed behavior. Empirical evidence documents systematic violations of its axioms, such as the Allais paradox and preference reversals [1, 9]. Nevertheless, it remains the central normative benchmark against which alternative theories are compared.

The Allais Paradox

Consider the following two problems:

Choice 1
Option A: $1 million for certain.
Option B: 89% chance of $1 million, 10% chance of $5 million, and 1% chance of nothing.

Choice 2
Option C: 11% chance of $1 million and 89% chance of nothing.
Option D: 10% chance of $5 million and 90% chance of nothing.

The paradox arises because many individuals choose A over B because certainty is especially attractive, but choose D over C because, once certainty disappears, the larger prize becomes more appealing. That switch is what expected utility theory cannot accommodate. It shows that individuals often place a disproportionate value on certainty when making choices under risk [1, 9].

3.5 Risk Attitudes under Expected Utility

Once preferences over risky prospects are represented by expected utility, attitudes toward risk can be characterized by the curvature of the utility function. This is one of the key insights of normative decision theory: the same probability distribution may be evaluated differently because decision makers attach different utility to the same monetary outcomes.

Figure 3.2 illustrates the three canonical cases.

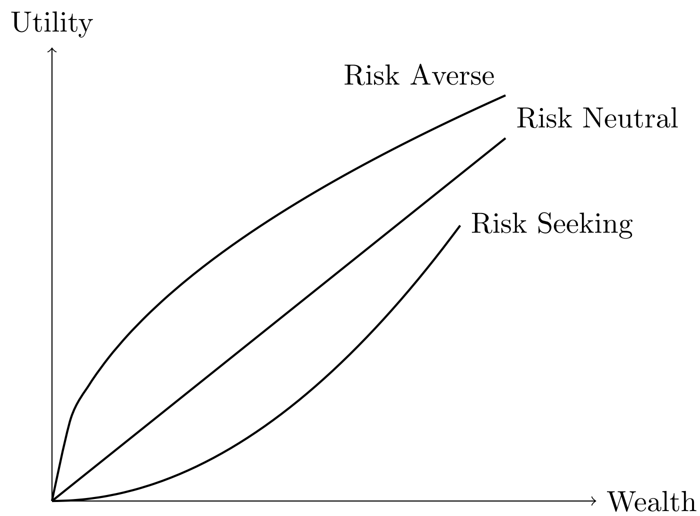

Figure 3.2: Schematic utility functions representing different attitudes toward risk.

A linear utility function corresponds to risk neutrality. A concave utility function corresponds to risk aversion, reflecting diminishing marginal utility of wealth and implying a preference for certainty over fair gambles. A convex utility function corresponds to risk seeking, where variability is valued positively.

These differences in curvature allow expected utility theory to represent a wide range of attitudes toward uncertainty within a unified framework. They also create a bridge to two important monetary concepts: the *certainty equivalent* and the *risk premium*. Once risk attitudes are expressed through curvature, they can be translated into monetary comparisons that are directly useful in practical decision settings.

3.6 Certainty Equivalents and Risk Premia

Expected utility theory provides a direct way to compare risky prospects with certain amounts through the concept of a *certainty equivalent*. The certainty equivalent of a risky prospect is the guaranteed amount that yields the same utility as the prospect itself.

Formally, for a prospect with outcomes x_i and probabilities p_i, the certainty equivalent CE is defined by

$$u(CE) = \sum_i p_i u(x_i).$$

The decision maker is therefore indifferent between receiving CE with certainty and facing the risky prospect itself.

The difference between the expected value of the prospect and its certainty equivalent is called the *risk premium*:

$$RP = E[x] - CE.$$

The risk premium measures the maximum amount the decision maker would be willing to give up in order to eliminate the risk. It is therefore a monetary measure of the cost of uncertainty.

The sign and size of the risk premium depend on the curvature of the utility function. For a risk-neutral decision maker, utility is linear and $CE = E[x]$, so the risk premium is zero. For a risk-averse decision maker with concave utility, the certainty equivalent lies below the expected value:

$$CE < E[x],$$

so the risk premium is positive. For a risk-seeking decision maker with convex utility, the opposite can hold.

Worked Example: Certainty Equivalent and Risk Premium in a Contract Decision

The following example is constructed for illustration. It is not taken from a specific source, but applies the standard definitions of certainty equivalent and risk premium used in expected utility theory [12, 6].

Suppose a consulting firm is considering a performance-based contract. The contract pays \$100,000 if the project succeeds and \$0 if it fails, with each outcome occurring with probability 0.5.

The expected monetary value of the contract is

$$E[x] = 0.5(100{,}000) + 0.5(0) = 50{,}000.$$

If the firm is indifferent between accepting this risky contract and receiving \$35,000 for certain, then \$35,000 is the *certainty equivalent*:

$$CE = 35{,}000.$$

The *risk premium* is the difference between the expected value and the certainty equivalent:

$$RP = E[x] - CE = 50{,}000 - 35{,}000 = 15{,}000.$$

This means that, although the contract has an expected value of \$50,000, the firm values it as equivalent to only \$35,000 with certainty. The \$15,000 difference is the monetary cost of bearing the risk.

In formal models, the certainty equivalent is derived from an explicit utility function. At this stage, however, the key idea is conceptual: certainty equivalents and risk premia translate risky prospects into monetary terms that can be compared directly. This makes expected utility operational in practical settings such as insurance, investment, contracting, and project evaluation, and it prepares the ground for more specific models of risk aversion in which utility curvature is given an explicit functional form.

3.7 Modeling Risk Aversion

Risk aversion is the most extensively studied attitude toward uncertainty in economics. A decision maker is risk averse if they prefer a certain outcome to a risky prospect with the same expected value. In expected utility theory, this corresponds to concavity of the utility function:

$$u''(x) < 0.$$

This property reflects diminishing marginal utility of wealth: an additional unit of wealth matters less at higher wealth levels than at lower ones. As a result, unfavorable outcomes weigh more heavily than equally sized favorable outcomes when risky prospects are evaluated in utility terms.

A useful quantitative representation of risk aversion is given by the Arrow–Pratt coefficients [12, 3]:

$$A(x) = -\frac{u''(x)}{u'(x)}, \qquad R(x) = -x\frac{u''(x)}{u'(x)}.$$

The first is the coefficient of absolute risk aversion and the second the coefficient of relative risk aversion. They measure the local curvature of utility and therefore the intensity of risk aversion. These measures are directly linked to practical evaluation: a more concave utility function implies a lower certainty equivalent and a larger risk premium for a given gamble.

In applied work, this curvature must usually be represented by a specific functional form. The most common benchmark specifications are linear utility, CARA utility, and CRRA utility [6, 10].

Common Utility Specifications for Modeling Risk Attitudes

Linear utility (risk neutral) [10]

$$u(x) = x$$

CARA utility (constant absolute risk aversion) [3, 6]

$$u(x) = -e^{-ax}, \qquad a > 0$$

CRRA utility (constant relative risk aversion) [6, 10]

$$u(x) = \frac{x^{1-\gamma}}{1-\gamma}, \qquad \gamma > 0,\ \gamma \neq 1$$

The Arrow–Pratt measures of absolute and relative risk aversion are defined by [12, 3]

$$A(x) = -\frac{u''(x)}{u'(x)}, \qquad R(x) = -x\frac{u''(x)}{u'(x)}.$$

Applying these definitions yields [12, 3, 6]:

$$\text{Linear: } A(x) = 0,\ R(x) = 0$$

$$\text{CARA: } A(x) = a,\ R(x) = ax$$

$$\text{CRRA: } A(x) = \frac{\gamma}{x},\ R(x) = \gamma$$

These expressions summarize the economic meaning of the three forms. Linear utility implies risk neutrality. CARA implies that sensitivity to a fixed absolute gamble does not depend on wealth, because absolute risk aversion is constant. CRRA implies that sensitivity to a proportional gamble remains constant as wealth changes, because relative risk aversion is constant. The choice between these forms is therefore substantive rather than merely technical: it determines how changes in wealth affect the evaluation of risk.

Worked Example: Insurance Demand under CARA Utility

The following example is constructed for illustration using standard expected-utility definitions and the CARA specification discussed in the literature [3, 12, 6]. It is not copied from a single source.
Consider an individual with initial wealth of \$100,000 who faces a 1% probability of a loss of \$10,000. Final wealth without insurance is therefore

$$W = \begin{cases} 100{,}000 & \text{with probability } 0.99, \\ 90{,}000 & \text{with probability } 0.01. \end{cases}$$

The expected wealth without insurance is

$$E[W] = 0.99(100{,}000) + 0.01(90{,}000) = 99{,}900.$$

Suppose preferences are represented by CARA utility

$$u(W) = -e^{-aW}, \qquad a > 0.$$

Expected utility without insurance is

$$EU = -\left[0.99e^{-a\cdot 100{,}000} + 0.01e^{-a\cdot 90{,}000}\right].$$

The certainty equivalent wealth level CE_W is defined by

$$-e^{-a\,CE_W} = -\left[0.99e^{-a\cdot 100{,}000} + 0.01e^{-a\cdot 90{,}000}\right].$$

Solving for CE_W gives

$$CE_W = -\frac{1}{a}\ln\left(0.99e^{-a\cdot 100{,}000} + 0.01e^{-a\cdot 90{,}000}\right).$$

Because CARA utility is concave for $a > 0$, this certainty equivalent is below expected wealth:

$$CE_W < 99{,}900.$$

The difference

$$RP = E[W] - CE_W$$

is the risk premium associated with the uninsured position.
Now suppose full insurance is available. If the individual pays a premium π, final wealth becomes certain at

$$100{,}000 - \pi.$$

The individual is indifferent between remaining uninsured and buying insurance when

$$100{,}000 - \pi = CE_W.$$

Hence the maximum premium the individual is willing to pay for full insurance is

$$\pi^* = 100{,}000 - CE_W.$$

This reservation premium exceeds the actuarially fair premium of \$100 whenever the risk premium embedded in the uninsured position is positive.

This example shows how insurance demand can be derived formally from the same concepts developed earlier in the chapter. The decision does not rest on a vague dislike of uncertainty, but on the comparison between expected utility, certainty-equivalent wealth, and the payment required to eliminate the risk. More broadly, it illustrates how explicit utility specifications turn abstract ideas about risk aversion into operational models that can be used in insurance, finance, and decision analysis.

3.8 Modeling Risk-Seeking Behavior

Although economic analysis often emphasizes risk aversion, individuals and organizations may also exhibit risk-seeking behavior under specific conditions. A decision maker is risk seeking if they prefer a risky prospect to a certain outcome with the same expected value. Within expected utility theory, this corresponds to local convexity of the utility function:

$$u''(x) > 0.$$

Risk-seeking behavior can be approached from both normative and descriptive perspectives. In descriptive theories of choice, it is often associated with losses, reference points, or aspiration levels, as emphasized in prospect theory [9, 14]. In organizational and financial settings, it may also arise when incentives are convex, when agents face severe downside pressure, or when survival or promotion depends on achieving a threshold outcome rather than on average performance. In such cases, volatility itself may become attractive because it increases the chance of reaching the relevant target.

Within expected utility theory, risk-seeking behavior can be represented by any utility function that is locally convex over the relevant range. One way to do this is to use the convex counterpart of the CARA family. Standard CARA utility for risk aversion is written as

$$u(x) = -e^{-ax}, \qquad a > 0,$$

which is strictly concave. If the sign is reversed, however, the exponential form becomes convex:

$$u(x) = e^{bx}, \qquad b > 0.$$

This is no longer the standard CARA specification used for risk aversion. Instead, the positive parameter b measures the intensity of local risk-seeking behavior: larger values of b imply greater curvature and therefore a stronger preference for upside volatility over the relevant range.

Worked Example: Risk Seeking in a Distressed Firm's Turnaround Choice

Consider a firm under severe financial pressure. Management must choose between two restructuring strategies for the coming year:

- *Safe strategy:* a guaranteed cash-flow improvement of $2 million;
- *Turnaround strategy:* a 50% chance of a $6 million improvement and a 50% chance of a $2 million loss.

Let c denote the firm's net cash-flow outcome, measured in millions of dollars.
Both strategies have the same expected monetary value:

$$E[c]_{\text{safe}} = 2, \qquad E[c]_{\text{turnaround}} = 0.5(6) + 0.5(-2) = 2.$$

Suppose that, over the relevant range, management evaluates outcomes with a convex exponential utility function

$$u(c) = e^{bc}, \qquad b > 0,$$

where larger values of b imply stronger local risk-seeking behavior.
The utility of the safe strategy is

$$u(2) = e^{2b}.$$

The expected utility of the turnaround strategy is

$$EU_{\text{turnaround}} = 0.5e^{6b} + 0.5e^{-2b}.$$

Because the exponential function is convex, the expected utility of the risky turnaround strategy exceeds the utility of the certain option with the same expected value:

$$0.5e^{6b} + 0.5e^{-2b} > e^{0.5(6b)+0.5(-2b)} = e^{2b}.$$

Hence

$$EU_{\text{turnaround}} > u(2),$$

so the risky turnaround strategy is preferred to the safe strategy even though both have the same expected monetary value.
The certainty equivalent CE of the turnaround strategy satisfies

$$e^{bCE} = 0.5e^{6b} + 0.5e^{-2b},$$

so

$$CE = \frac{1}{b} \ln\left(0.5e^{6b} + 0.5e^{-2b}\right).$$

Since utility is convex, this certainty equivalent exceeds the expected value of $2 million:

$$CE > 2.$$

The implied risk premium is therefore negative:

$$RP = E[c] - CE < 0.$$

This example shows that the same framework used to model risk aversion can also represent risk-seeking behavior. The difference lies not in the underlying logic of evaluation, but in the curvature of utility over the relevant domain.

In the present case, the negative risk premium has a clear managerial interpretation. Because the firm values the upside potential of a successful turnaround more than it dislikes the downside risk of failure, the risky strategy is worth more to management than its expected cash-flow value alone would suggest. The certainty equivalent therefore exceeds the expected value of $2 million, which means that management would require a guaranteed payment greater than $2 million to be willing to give up the turnaround option. The negative risk premium is thus the monetary expression of a preference for volatility: rather than paying to avoid risk, the firm would forgo part of a certain payoff in order to retain exposure to a high-upside opportunity.

This helps explain why distressed firms, highly leveraged investors, or managers facing threshold-based incentives may rationally choose volatile strategies when safer alternatives do not offer a realistic path to recovery, survival, or exceptional performance. In such settings, risk is not merely tolerated; it can become actively desirable because it preserves the possibility of reaching outcomes that matter disproportionately.

The formal models developed so far show how risk attitudes can be represented. The next question is how such attitudes can be identified or estimated in actual decision settings.

3.9 Measuring Risk Attitudes in Practice

The formal models developed in this chapter show how risky alternatives can be evaluated once a utility function or preference structure has been specified. In applied settings, however, these risk attitudes are often not directly known and must instead be inferred from observed or experimental choices.

Common empirical approaches include lottery-choice tasks, certainty-equivalent elicitation, and multiple price list designs such as those used by Holt and Laury [7]. These methods attempt to recover the parameter values or preference patterns that best rationalize observed choices, including the implied curvature of utility and the associated certainty equivalents and risk premia. In this way, the measurement of risk attitudes provides an empirical counterpart to the normative models of expected utility, certainty equivalents, and risk premia.

At the same time, measured attitudes toward risk are often sensitive to the decision environment. Experimental evidence shows that inferred parameters can vary across elicitation methods, payoff scales, framing conditions, and decision domains [2, 13]. A person may appear strongly risk averse in one setting and less so in another, not necessarily because the theory is wrong, but because the empirical expression of risk preference is context dependent.

For this reason, practical measurement should be interpreted with care. Risk attitudes are not observed directly; they are inferred through a model, a task, and a particular domain of choice. In applied risk analysis, it is therefore important to specify clearly both the empirical setting and the preference framework being used.

The broader lesson is that formal modeling and empirical measurement are complementary. Theory provides the structure for representing risky choice, while empirical methods attempt to identify which representation is most relevant for the decision makers and contexts under study.

3.10 Summary and Key Takeaways

This chapter developed the normative foundations of choice under risk. Its central message is simple: probabilistic information alone is not enough to guide decisions. Uncertainty must also be evaluated through a preference structure. Expected value provides a useful benchmark, but expected utility offers the more general framework because it allows risky prospects to be assessed in light of risk attitudes.

Within this framework, utility curvature determines whether the decision maker is risk neutral, risk averse, or risk seeking. Certainty equivalents and risk premia translate these attitudes into monetary terms: for risk-averse decision makers, the certainty equivalent lies below expected value and the risk premium is positive, whereas for risk-seeking decision makers, the certainty equivalent can exceed expected value and the risk premium becomes negative. Utility specifications such as CARA and CRRA make these ideas operational in applied settings, while empirical estimates of risk attitudes remain context dependent, so formal modeling and practical measurement must be treated as complementary.

More broadly, the chapter shows how preference-based criteria connect probabilistic models to actual decisions. In that sense, the architecture of risk preferences is what turns uncertainty analysis into risk evaluation.

Key Takeaways

- Do not rely on expected value alone when the dispersion of outcomes matters.

- Use expected utility when decisions depend on attitudes toward risk rather than average payoff alone.
- Use certainty equivalents and risk premia to compare risky and certain options: positive risk premia indicate aversion to risk, while negative risk premia indicate a preference for risk.
- Choose the utility specification carefully: CARA and CRRA imply different ways that risk attitudes change with wealth.
- Treat measured risk attitudes as context dependent, and always state clearly the decision setting and preference model being used.

Bibliography

[1] Maurice Allais. Le comportement de l'homme rationnel devant le risque. *Econometrica*, 21:503–546, 1953.

[2] Gosse A. G. Alserda. Measuring latent risk preferences: minimizing measurement biases. *Journal of Risk*, 21(5):83–103, 2019.

[3] Kenneth J. Arrow. Aspects of the theory of risk-bearing. *Yrjö Jahnsson Foundation*, 1965.

[4] Daniel Bernoulli. Specimen theoriae novae de mensura sortis. *Commentarii Academiae Scientiarum Imperialis Petropolitanae*, 5:175–192, 1738.

[5] Peter Fishburn. *Utility Theory for Decision Making*. Wiley, 1970.

[6] Christian Gollier. *The Economics of Risk and Time*. MIT Press, Cambridge, MA, 2001.

[7] Charles A. Holt and Susan K. Laury. Risk aversion and incentive effects. *American Economic Review*, 92(5):1644–1655, 2002.

[8] Daniel Kahneman. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, 2011.

[9] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, 1979.

[10] Andreu Mas-Colell, Michael Whinston, and Jerry Green. *Microeconomic Theory*. Oxford University Press, 1995.

[11] Philippe Mongin. Expected utility theory. In John B. Davis, D. Wade Hands, and Uskali Mäki, editors, *Handbook of Economic Methodology*, pages 342–350. Edward Elgar, Cheltenham, UK, 1997.

[12] John W. Pratt. Risk aversion in the small and in the large. *Econometrica*, 32(1/2):122–136, 1964.

[13] Philip J. Thomas. Measuring risk-aversion: The challenge. *Measurement*, 79:285–301, 2016.

[14] Amos Tversky and Daniel Kahneman. Advances in prospect theory. *Journal of Risk and Uncertainty*, 1992.

[15] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ, 1944.

