

How Should a Robot Assess Risk?

Towards an Axiomatic Theory of Risk in Robotics

Anirudha Majumdar and Marco Pavone

Key words: Planning and decision making under uncertainty, risk metrics, safety.

Abstract Endowing robots with the capability of assessing risk and making risk-aware decisions is widely considered a key step toward ensuring safety for robots operating under uncertainty. But, how should a robot *quantify* risk? A natural and common approach is to consider the framework whereby costs are assigned to stochastic outcomes—an assignment captured by a cost random variable. Quantifying risk then corresponds to evaluating a *risk metric*, i.e., a mapping from the cost random variable to a real number. Yet, the question of what constitutes a “good” risk metric has received little attention within the robotics community. The goal of this paper is to explore and partially address this question by advocating *axioms* that risk metrics in robotics applications should satisfy in order to be employed as rational assessments of risk. We discuss general representation theorems that precisely characterize the class of metrics that satisfy these axioms (referred to as *distortion risk metrics*), and provide instantiations that can be used in applications. We further discuss pitfalls of commonly used risk metrics in robotics, and discuss additional properties that one must consider in sequential decision making tasks. Our hope is that the ideas presented here will lead to a foundational framework for quantifying risk (and hence safety) in robotics applications.

1 Introduction

Safe planning and decision making under uncertainty are widely regarded as central challenges in enabling robots to successfully operate in real-world environments. By far the most common conceptual framework for addressing these challenges is to assign costs to stochastic outcomes and then to use the *expected value* of the resulting cost distribution as a quantity that “summarizes” the value of a decision. Such a quantity can then be optimized, or bounded within a constrained formulation. However, in settings where risk has to be accounted for, this choice is rarely

Anirudha Majumdar and Marco Pavone

Department of Aeronautics and Astronautics, Stanford University, CA 94305, USA. Emails: {anirudha, pavone}@stanford.edu.

well justified beyond the mathematical convenience it affords. For example, imagine a safety-critical application such as autonomous driving; would a passenger riding in an autonomous car be happy to do so if she was told that the average behavior of the car is not to crash? While one can introduce some degree of risk sensitivity (i.e., sensitivity to the tails of the cost distribution) in the expected cost framework by simply shaping the cost function, this can quickly turn into an exercise in “cost function hacking”. Unless one is careful about the way one shapes the cost function, this can lead to the robot behaving in an irrational manner [3]. The common alternative approach, aimed at promoting risk sensitivity, is to consider a *worst-case assessment* of the distribution of stochastic outcomes. In practice, however, such an assessment can often be quite conservative: an autonomous car whose goal is to never crash would never leave the garage.

The expected value operator and the worst-case assessment are examples of *risk metrics*. Informally, a risk metric is a mapping from a random variable corresponding to costs to a real number. The expected cost corresponds to risk neutrality while the worst-case assessment corresponds to extreme risk aversion. For practical applications, we would like to explore risk metrics that lie in between these extremes. This raises the following question: is there a class of risk metrics that lie between these extremes while still ensuring that the robot quantifies risk (and hence safety) in a rational and trustworthy manner? This question is central to the problem of decision making under uncertainty since the choice of a risk metric is one that must be made in *any* framework that assigns costs to outcomes. Yet, despite the role of safe decision making under uncertainty as a core theme in practically all areas of robotics, this question has received very little attention within the robotics community. As a result, there is arguably no firm theoretical foundation for making an informed decision about what risk metric to use for a given robotics application.

The goal of this paper is to provide a first step towards such a principled framework. More precisely, we describe axioms (properties) that a risk metric employed by a robot should satisfy in order to be considered sensible. To our knowledge, this is the first attempt to provide such an axiomatic framework for evaluating risk in robotics applications. Our effort is inspired by a similar effort in the finance community that led to the identification of *coherent risk metrics* [4, 29] as a class of risk metrics that have desirable properties for assessing the risk associated with a financial asset (e.g., a portfolio of stocks). The influence that these ideas have had can be gauged by the fact that in 2014 the Basel Committee on Banking Supervision changed its guidelines for banks to replace the Value at Risk (VaR) (a *non*-coherent risk metric) with the Conditional Value at Risk (CVaR) (a coherent risk metric) for assessing market risk [22].

We believe that the question of what properties a risk metric should satisfy in robotics applications is a fundamental one: a robot’s inability to assess risks in a rational way could lead to behavior that is harmful both to itself and humans or other autonomous agents around it. Our hope is that the ideas presented in this paper can help the community converge upon a set of properties that risk metrics must satisfy in order for the robot’s decision-making system to be considered rational and trustworthy, paralleling a similar effort in the financial industry. Indeed, it is conceivable that in the not-so-distant future, robots such as autonomous cars or unmanned aerial

vehicles (UAVs) deployed in safety-critical scenarios will be subject to regulatory frameworks that mandate the use of “officially-approved” risk metrics.

Related Work: Our work here is inspired by the efforts over the last two decades towards the development of an axiomatic theory of risk in finance and operations research. These efforts have resulted in widespread acceptance of the class of coherent risk metrics as “rational” measures of risk for financial assets. Coherent risk metrics are defined by four axioms that any sensible assessment of financial risk must satisfy (see Section 4). Additional axioms beyond the four characterizing coherent risk metrics have also been studied and lead to more refined classes of risk metrics. We refer the reader to [12, 29] for an introduction to this vast literature. Our main contribution here is to parallel the efforts in the finance community on proposing and characterizing axioms that any sensible assessment of risk in a robotics application should satisfy. We propose these axioms in Section 3 and provide interpretations for them in a robotics context. As we will see, the class of risk metrics fulfilling these axioms corresponds precisely to the class of *distortion risk metrics*, which form a subset of coherent risk metrics and have previously been studied in the finance literature [32, 16, 12, 5].

As discussed previously, the most commonly used risk metrics in robotics are the expected cost and worst-case metrics. While these risk metrics are justifiable in certain contexts (e.g., expected cost in applications where the distribution of costs is known to not have a long tail or worst-case assessments in low-level control tasks such as trajectory tracking), many robotics applications call for a more nuanced assessment of risk. *Chance constrained programming* [7] provides one avenue towards such assessments. In particular, a chance constraint specifies an upper bound on the probability of incurring a cost higher than a given threshold. Chance constraints have been widely studied in robotics for motion planning under uncertainty [6, 10, 23]. While chance constraints are suitable for capturing risk corresponding to *boolean* events (e.g., collisions with obstacles), they do not take into account variations in the tails of cost distributions (since they are not affected by changes to the value of the cost above the given threshold). Chance constraints are thus not suitable for capturing risk in settings where one must consider a range of cost outcomes (in contrast to boolean events).

Another popular way to quantify risk in control theory and decision making is through the notion of *distributional robustness* [9, 34, 30]. Distributional robustness captures the idea that the underlying distribution from which random outcomes of the world are generated may itself be uncertain. Such scenarios are referred to as *ambiguous* in the literature on human decision making [14]. In such scenarios, while the precise distribution may be unknown, one may know certain properties of the underlying distribution (e.g., its first few moments, or that it lies in a given set of possible distributions). One can then compute the worst-case expectation of the cost function over the *set of distributions* that satisfy the known properties (e.g., the set of all distributions that have the given moments). As we will see in Section 4, the set of risk metrics fulfilling the axioms advocated in this paper have an interpretation in terms of distributional robustness. However, not all distributionally robust risk metrics satisfy the proposed axioms.

Beyond robotics and finance, the notion of risk is of central concern to the theory of human decision making under uncertainty. A historically prominent theory

is *subjective expected utility (SEU)* theory [31], where humans make decisions that maximize the expected value of a utility function. This is analogous to the method for taking into account risk aversion in autonomous agents by “shaping” the cost function. However, SEU theory is inconsistent with a number of experimental observations [11, 17, 2]. In particular, SEU theory does not take into account the experimentally observed fact that humans are *ambiguity averse* [11, 2], i.e., averse to situations where there is uncertainty about the underlying probability distribution from which outcomes are drawn (previously discussed in the context of distributional robustness). Models of ambiguity aversion include *prospect theory*, which has risen to prominence over the last several decades, along with more recent and sophisticated theories [14] that model the human as minimizing a distributionally robust risk metric. These theories are thus closely related to the class of distortion risk metrics (which we advocate for in this paper for use in robotics) since they can also be interpreted as a special class of distributionally robust risk metrics.

Outline: The outline of this paper is as follows. Section 2 formally introduces the notion of a risk metric (Section 2.1) and proposes an interpretation of risk in robotics applications (Section 2.2). In Section 3 we advocate axioms that risk metrics in robotics applications should satisfy in order to be considered sensible (Sections 3.1), provide examples of metrics that satisfy them, and discuss pitfalls stemming from using risk metrics that do not satisfy these axioms (Section 3.2). Section 4 discusses representation theorems that precisely characterize the class of risk metrics fulfilling the axioms we propose. Section 5 proposes additional properties that one must consider in sequential decision making tasks. Section 6 concludes the paper with directions for future research. We note that we highlight a number of *points of discussion* throughout the paper. Our hope is that these will form the basis for discussion and debate at the Blue Sky session during the conference and will provoke new directions for future research.

2 Assessing Risk: Preliminaries

In this section, we formally introduce risk metrics and propose an intuitive interpretation of risk quantification in robotics contexts. This interpretation will form the basis for the axioms we advocate in Section 3.

2.1 Risk Metrics

We denote the set of possible outcomes that may occur when a robot operates in uncertain settings as Ω . In order to avoid heavy use of measure theoretic notions, we take Ω to be finite. Denote by \mathbb{P} a probability mass function that assigns probabilities $\mathbb{P}(\omega)$ to outcomes $\omega \in \Omega$. Consider a cost function $Z : \Omega \rightarrow \mathbb{R}$ that assigns costs $Z(\omega)$ to outcomes. The cost Z is then a random variable, namely the cost random variable. Let \mathcal{Z} denote the set of all random variables on Ω . A *risk metric* is a mapping $\rho : \mathcal{Z} \rightarrow \mathbb{R}$. In other words, a risk metric maps a cost random variable to a real number.

2.2 Interpretation of Risk in Robotics Applications

Imagine a fictional government agency known as the Robot Certification Agency (RCA) that is responsible for certifying if a given robot is safe to operate in the real

world. How should the RCA quantify the risk faced by this robot? As an example, consider an autonomous car driving from one city to another. While performing this task, the robot will incur random costs Z (e.g., due to fuel consumption, time, crashes, mechanical wear and tear, etc.). In order to provide clear interpretations of the axioms for risk metrics discussed below, we will take the following axiom as a key starting point.

A0. Monetary costs. The costs Z are expressed in monetary terms.

This axiom ensures that the costs assigned to outcomes have a tangible and interpretable value, which will be instrumental in defining a meaningful notion of risk below. Such an axiom may also provide a handle on reasoning about insurance policies for safety-critical robots (e.g., autonomous cars). We note that our starting point contrasts with one where one considers a more abstract and subjective notion of cost (e.g., quadratic state and control costs for Linear Quadratic Regulator problems).

Given this assumption, suppose that the RCA demands that the robot's owner must deposit an amount of money $\rho(Z)$ before the robot is deployed such that the RCA is satisfied that the owner will be able to cover the potential costs incurred during operations (e.g., making repairs to the robot due to an accident) with the amount ρ . We define the amount $\rho(Z)$ as the perceived *risk* from operating the robot. The particular risk metric $\rho : \mathcal{Z} \rightarrow \mathbb{R}$ the RCA uses will depend on its attitude towards risk and may depend on the application under consideration. For example, the RCA may ask for a deposit $\rho(Z) = \mathbb{E}[Z]$ if it is risk neutral. If it wants to be highly conservative, the RCA may demand a deposit equal to the worst-case cost outcome. The question we will pursue in Section 3 is the following: what properties must ρ satisfy in order for it to be considered sensible?

We note that we are using the RCA here as a pedagogical tool to provide an interpretation of risk in robotics applications and to motivate the axioms described in Section 3. In reality, the robot's decision-making system will assess risks and make decisions based on those assessments.

3 An Axiomatization of Risk Metrics for Robotics Applications

3.1 Axioms and their Interpretations

We now parallel results in finance [12, Chapter 4] [29] and propose six axioms for risk metrics. Specifically, we make the case that these axioms should be fulfilled by any risk metric used in a robotics application in order for it to be considered a sensible assessment of risk. For each axiom, we first provide a formal statement and then an intuitive interpretation based on the interpretation of risk from Section 2.2.

A1. Monotonicity. Let $Z, Z' \in \mathcal{Z}$ be two cost random variables. Suppose $Z(\omega) \leq Z'(\omega)$ for all $\omega \in \Omega$. Then $\rho(Z) \leq \rho(Z')$.

Interpretation: If a random cost Z' is guaranteed to be greater than or equal to a random cost Z *no matter what random outcome occurs*, then Z' must be deemed at least as risky as Z . One can think of the random costs as corresponding to two different robots, or the same robot performing different tasks, or executing different controllers. Given our interpretation of risk, this axiom states that the RCA must

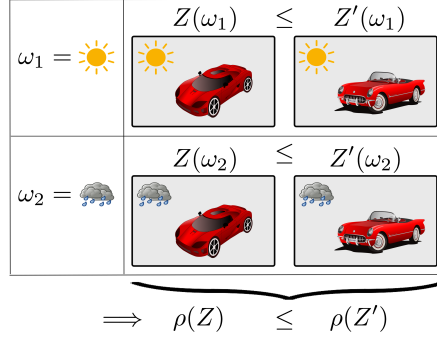


Fig. 1 Monotonicity. Consider two robot cars that incur random costs Z and Z' respectively. The situation consists of two outcomes ω_1 and ω_2 corresponding to sunny and rainy weather, respectively. If the car that incurs costs Z (first column) incurs a lower cost than the other car (second column) no matter what the weather is, monotonicity states that it should be considered less risky.

demand at least as large a deposit for covering costs for the robot (or task) corresponding to Z' as for the robot (or task) corresponding to Z . This is a sensible requirement since we are guaranteed to incur at least as high a cost in the second scenario as the first no matter which outcome $\omega \in \Omega$ is realized. An example is illustrated in Figure 1, where we have two robot cars corresponding to Z and Z' and two outcomes ω_1 and ω_2 corresponding to sunny and rainy weather, respectively. If the car corresponding to Z incurs a lower cost no matter what the weather is, monotonicity states that it should be considered less risky.

A2. Translation invariance. Let $Z \in \mathcal{Z}$ and $c \in \mathbb{R}$. Then $\rho(Z + c) = \rho(Z) + c$.

Interpretation: If one is charged a deterministic cost c (in addition to the random costs incurred when the robot is operated), then the RCA should demand that this amount c be set aside in addition to money for covering the other costs from operating the robot:

$$\rho(\underbrace{\text{🚗} + \text{💰}}_{Z + c}) = \rho(\underbrace{\text{🚗}}_Z) + \underbrace{\text{💰}}_c$$

Note that this axiom also implies that $\rho(Z - \rho(Z)) = 0$. Thus, $\rho(Z)$ is the smallest amount that must be deducted from the costs in order to make the task risk-free.

A3. Positive homogeneity. Let $Z \in \mathcal{Z}$ and $\beta \geq 0$ be a scalar. Then $\rho(\beta Z) = \beta \rho(Z)$.

Interpretation: If all the costs incurred by the robot (regardless of the random outcome) are scaled by β , the RCA demands that the deposit is scaled commensurately. This is reasonable since this corresponds to simply changing the units of cost (recall that we assumed that the costs are expressed in monetary terms).

A4. Subadditivity. Let $Z, Z' \in \mathcal{Z}$. Then $\rho(Z + Z') \leq \rho(Z) + \rho(Z')$.

Interpretation: This axiom encourages diversification of risk. For example, imagine a system with two robots. Suppose that Z and Z' are costs incurred by Robot 1

and Robot 2 respectively. The left-hand side (LHS) of the inequality corresponds to the deposit that the RCA demands when both robots are run simultaneously, while the right-hand side (RHS) corresponds to the sum of the deposits when the robots are deployed separately. Axiom A4 then states that deploying both robots simultaneously is at most as risky as deploying them separately:

$$\rho(\underbrace{\text{🤖🤖}}_{Z+Z'}) \leq \rho(\underbrace{\text{🤖}}_Z) + \rho(\underbrace{\text{🤖}}_{Z'})$$

This captures the intuition that one robot acts as a hedge against the other robot failing (i.e., if one of the robot fails in some way, then the other will make up for this loss). Another interpretation is that this axiom promotes redundancy in the system. The exact interpretation of the LHS and RHS of the inequality in A4 will depend on the particular application under consideration. For example, the two sub-costs Z and Z' may correspond to two separate sub-tasks that the robot must perform. In this case, the LHS corresponds to the robot performing both tasks simultaneously while the RHS corresponds to performing them independently. A4 encodes the intuition that performing both tasks simultaneously is less risky since one sub-task can act as a hedge against the other.

We note that A3 and A4 together imply convexity:

$$\rho(\lambda Z + (1 - \lambda)Z') \leq \lambda \rho(Z) + (1 - \lambda)\rho(Z'), \text{ for all } \lambda \in [0, 1].$$

A5. Comonotone additivity. Suppose Z and Z' are comonotone, i.e., $(Z(\omega) - Z(\omega'))(Z'(\omega) - Z'(\omega')) \geq 0$, for all $(\omega, \omega') \in \Omega \times \Omega$. Then $\rho(Z + Z') = \rho(Z) + \rho(Z')$.

Interpretation: This axiom supplements A4. In particular, A5 states that if two costs rise and fall together, then there is no benefit from diversifying (e.g., if one robot always performs poorly at a task when the other does or when a robot performs poorly at a sub-task when it also performs poorly at another one).

A6. Law invariance. Suppose that Z and Z' are identically distributed. Then $\rho(Z) = \rho(Z')$.

Interpretation: If two tasks have the same distribution of costs, then the RCA demands an equal deposit in both cases. For example, suppose $\Omega = \{\omega, \omega'\}$ with both outcomes having probability 0.5. Further, suppose $Z(\omega) = 1$, $Z(\omega') = 10$, $Z'(\omega) = 10$, $Z'(\omega') = 1$. The two situations must be considered equally risky even though the assignment of costs to events is different.

Taken together, Axioms A1 – A6 capture a fairly exhaustive set of essential properties that we believe any reasonable quantification of risk in robotics should obey given the interpretation of risk we proposed in Section 2.2. A hypothetical RCA that quantifies risk in a manner that is consistent with these axioms would be considered a sensible one. Moreover, robots that assess risks according to risk metrics that fail to satisfy some of these axioms can behave in a manner that would be con-

sidered extremely unreasonable and arguably very unsafe, as illustrated in Section 3.2. We thus advocate risk metrics that satisfy Axioms A1 – A6 for use in robotics applications.

3.2 Examples and Pitfalls of Commonly Used Risk Metrics

In this section, we first discuss examples of existing risk metrics that fulfill Axioms A1 – A6. We then discuss commonly used risk metrics that do not fulfill some of these axioms, along with pitfalls stemming from their use. Collectively, the discussion provided here motivates the formal introduction in the next section of distortion risk metrics (equivalently, risk metrics satisfying A1 – A6) as a general class of risk metrics for robotic applications.

An important risk metric that satisfies Axioms A1 – A6 is the Conditional Value at Risk (CVaR) [24]. The CVaR_α for a random cost Z at level α is defined as:

$$\text{CVaR}_\alpha(Z) := \frac{1}{\alpha} \int_{1-\alpha}^1 \text{VaR}_{1-\tau}(Z) d\tau, \quad (1)$$

where $\text{VaR}_\alpha(Z)$ is the *Value at Risk (VaR)* at level α , i.e., simply the $(1 - \alpha)$ -quantile of the cost random variable Z :

$$\text{VaR}_\alpha(Z) := \min\{z \mid \mathbb{P}[Z > z] \leq \alpha\}. \quad (2)$$

Intuitively, CVaR_α is the expected value of Z in the conditional distribution of Z 's upper $(1 - \alpha)$ -tail. It can thus be interpreted as a risk metric that quantifies “how bad is bad.” We note that the expected cost and worst-case assessment also satisfy A1 – A6. Figure 2 provides a visualization of the expected cost, worst case, VaR, and CVaR. Axioms A1 – A6 thus define a broad class of risk metrics that capture a wide spectrum of risk assessments from risk-neutral to worst-case. In Section 4, we will discuss theorems that allow us to precisely characterize *all* risk metrics that satisfy A1 – A6 and easily generate new examples of such metrics.

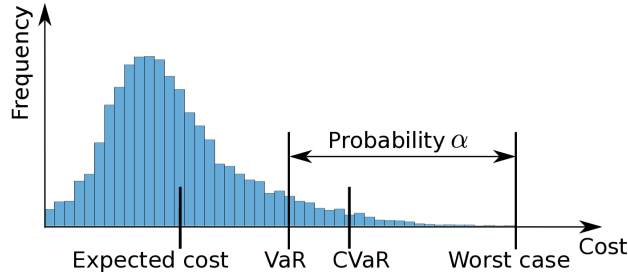


Fig. 2 An illustration of four important risk metrics, namely: expected cost, worst case, Value at Risk (VaR), and Conditional Value at Risk (CVaR). Intuitively, VaR is the $(1 - \alpha)$ -quantile of the cost distribution. CVaR is the expected value of costs in the conditional distribution of the cost distribution's upper $(1 - \alpha)$ -tail and is thus a metric of “how bad is bad.” The CVaR, expected cost, and worst case metrics satisfy Axioms A1 – A6, but VaR does not.

However, there are many examples of popular risk metrics that *do not* fulfill the axioms we advocate. For example, a very popular metric to quantify risk in robotics applications is the mean-variance risk metric: $\mathbb{E}[Z] + \beta \text{Variance}[Z]$ (see, e.g., [18, 13, 21]). The mean-variance metric satisfies A6 but fails to satisfy the other axioms. This can lead to a robot that utilizes the mean-variance metric making decisions that would be considered unreasonable. Consider the setup in Table 1 (based on [19]), where $\omega_1, \omega_2, \omega_3, \omega_4$ are disturbance outcomes, and Z and Z' are the costs resulting from executing two different controllers π and π' . Which controller should the robot execute? Controller π results in lower costs *no matter what the disturbance outcome is* and hence should be preferred by any sensible decision maker. However, computing the mean-variance risk metric with $\beta = 1$, we see that:

$$\mathbb{E}[Z] + \text{Variance}[Z] = 3.75 > \mathbb{E}[Z'] + \text{Variance}[Z'] = 3.4375.$$

The robot would hence strictly prefer π' . This unreasonable behavior is a result of the mean-variance risk metric failing to satisfy Axiom A1 (monotonicity).

Outcome	ω_1	ω_2	ω_3	ω_4
Probability	0.25	0.25	0.25	0.25
Z	1	2	3	4
Z'	2	2	3	4

Table 1 Issues with the mean-variance risk metric. Any rational agent would choose controller π (with associated costs Z) since it results in lower costs no matter what the disturbance outcome is. But, using the mean-variance risk metric results in choosing π' (with associated costs Z').

Outcome	ω_1	ω_2	ω_3
Probability	0.4	0.4	0.2
Z	1	2	3
Z'	1	1.99	10^{10}

Table 2 Issues with the Value at Risk (VaR) metric. Any reasonable agent would prefer costs Z to Z' . However, utilizing VaR results in choosing Z' .

The Value at Risk (VaR) (defined in Equation (2)) is another example of a risk metric that does not satisfy all the axioms (it is easily verified that VaR does not satisfy A4 (subadditivity), but satisfies the other axioms). VaR is closely related to chance constraints (see Section 1) since the constraint $\text{VaR}_\alpha(Z) \leq 0$ corresponds to the chance constraint $\mathbb{P}[Z > 0] \leq \alpha$. Using the VaR metric can also lead to behavior that is arguably very unreasonable and unsafe. For example, consider the costs Z and Z' in Table 3.2. Any reasonable agent would prefer costs Z to Z' due to the extremely large costs associated with ω_3 . However, we see that $\text{VaR}_{0.3}(Z) = 2$, while $\text{VaR}_{0.3}(Z') = 1.99$. Thus, utilizing VaR results in a strict preference for Z' . We note that using CVaR instead results in a preference for Z .

Table 3 lists the axioms satisfied by popular risk metrics in the literature (we refer the reader to [29] for definitions). We note that the standard semi-deviation is widely used in finance [29], while the entropic risk metric has been popular in control theory for risk-averse control [33, 15].

4 Distortion Risk Metrics

Risk metrics satisfying Axioms A1 – A6 have been studied in the context of portfolio optimization in finance and are known as *distortion risk metrics* [32, 16, 12, 5].

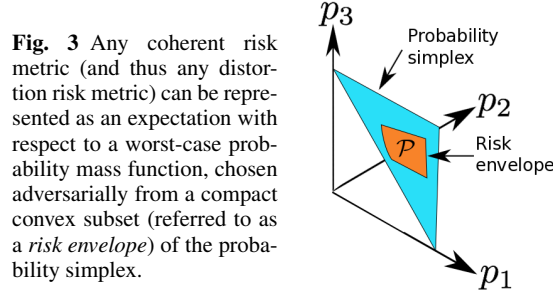
	Axioms
Conditional Value at Risk (CVaR)	A1 – A6
Expected Cost	A1 – A6
Worst case	A1 – A6
Mean–Variance	A6
Entropic risk	A1, A2, A6
Value at Risk (VaR)	A1 – A3, A5, A6
Standard semi-deviation	A1 – A4, A6

Table 3 Axioms satisfied by popular risk metrics.

These risk metrics are also equivalent to the class of *spectral risk measures* [1]. They enjoy an elegant characterization, which we discuss below, in terms of the CVaR metric. Before doing so, we first discuss representations for risk metrics satisfying subsets of the above axioms. In particular, Axioms A1 – A4 correspond to those of *coherent risk metrics* (CRMs). CRMs enjoy a universal representation theorem:

$$\rho(Z) = \max_{p \in \mathcal{P}} \mathbb{E}_p[Z], \quad (3)$$

where \mathcal{P} is a compact convex set of probability mass functions. In other words, any coherent risk metric can be represented as an expectation with respect to a worst case probability mass function, chosen adversarially from a given compact convex set (referred to as a *risk envelope*). An example of a risk envelope is visualized in Figure 3. Coherent risk metrics thus capture, as a by-product, the notion of *distributional robustness* (see Section 1), i.e., robustness to uncertainty over the underlying distribution itself.



Risk metrics satisfying Axioms A1 – A5 are known as *comonotonic risk metrics* [12, Chapter 4.7]. These risk metrics have a characterization in terms of Choquet integrals, which we now discuss with the help of additional terminology.

Definition 1. A set function $g : 2^\Omega \rightarrow [0, 1]$ is called *monotone* if $g(A) \leq g(B)$ for all $A \subseteq B \subseteq \Omega$, and *normalized* if $g(\emptyset) = 0$ and $g(\Omega) = 1$. If, in addition, g satisfies

$$g(A \cup B) + g(A \cap B) \leq g(A) + g(B),$$

g is called *submodular*.

Recall that Ω corresponds to the set of outcomes that may occur when the robot operates (see Section 2.1). The set 2^Ω is the power set (i.e., set of all subsets) of Ω . We now define the Choquet Integral [8] and state a representation theorem for comonotonic risk metrics in terms of these integrals.

Definition 2. The *Choquet Integral* of a random variable $Z \in \mathcal{Z}$ with respect to a monotone, normalized set function $g : 2^\Omega \rightarrow [0, 1]$ is

$$\int Z dg := \int_{-\infty}^0 (g(Z > z) - 1) dz + \int_0^\infty g(Z > z) dz.$$

Here, the integrals in the RHS are Riemann integrals. Informally, the Choquet integral is a generalization of the Lebesgue integral that allows the integration operation to have a *nonlinear* dependence on Z (in contrast to the Lebesgue integral, which is a linear operator).

Theorem 1 (Representation of comonotonic risk metrics [27]). *A coherent risk metric $\rho : \mathcal{Z} \rightarrow \mathbb{R}$ is comonotonic if and only if it can be written as a Choquet Integral $\int Z dg$, where $g : 2^\Omega \rightarrow [0, 1]$ is a monotone, normalized, and submodular set function.*

Risk metrics satisfying Axioms A1 – A6 are known as distortion risk metrics. These metrics inherit the representation theorems for coherent and comonotonic risk metrics and enjoy a further elegant characterization in terms of the CVaR metric.

Theorem 2 (Representation of distortion risk metrics [12]). *A risk metric $\rho : \mathcal{Z} \rightarrow \mathbb{R}$ is a distortion risk metric if and only if there exists a function $v : [0, 1] \rightarrow [0, 1]$, satisfying $\int_{\alpha=0}^1 v(d\alpha) = 1$ (i.e., v defines a probability measure on the set $[0, 1]$), such that:*

$$\rho(Z) = \int_{\alpha=0}^1 \text{CVaR}_\alpha(Z) v(d\alpha). \quad (4)$$

This theorem provides us with a precise mathematical characterization of *all* distortion risk metrics and allows us to generate examples of such metrics by choosing functions v satisfying the assumptions of the theorem.

To summarize our discussion so far, based on our arguments in Section 3.1 we advocate the use of distortion risk metrics (i.e., risk metrics satisfying A1 – A6, or equivalently risk metrics of the form (4)) for evaluating risk in robotics applications. This is in contrast to popular risk metrics used in the robotics literature (e.g., mean-variance, or VaR) and other popular classes of risk metrics used in finance (e.g., coherent risk metrics, or comonotonic risk metrics). However, we end this section by noting the possibility that for certain applications distortion risk metrics may constitute too restrictive a class of risk metrics. We leave the following as a point for discussion and future work.

Discussion 1 (Axioms A1 – A6). The axioms of monotonicity (A1), translation invariance (A2), positive homogeneity (A3), and law invariance (A6) should arguably be applicable in any robotics application. Subadditivity (A4) and comonotone additivity (A5) are also intuitively appealing, particularly for applications that involve

some degree of high-level decision making (since a high-level decision making system should ideally diversify risks). However, the interpretation of diversification is somewhat unclear in certain applications. For example, imagine a humanoid robot whose goal is to minimize a cost function that is a sum of two components Z and Z' , where Z penalizes one aspect of the robot's motion (e.g., deviations of the robot's torso from the vertical orientation) while Z' penalizes another aspect (e.g., deviation of the robot's gaze from a target). For such low-level control tasks, the interpretation of $\rho(Z) + \rho(Z')$ is not entirely clear since it is not possible to perform the different subtasks corresponding to Z and Z' independently of each other. The following is thus a question for discussion and future work: should A4 and A5 be abandoned (or replaced by other axioms) for such tasks?

5 Sequential Decision Making and Time Consistency

The sequential nature of many decision-making tasks in robotics gives rise to additional important considerations beyond the ones discussed in Section 3. Again paralleling results in finance [28, 26, 29], we discuss properties that ensure the temporal consistency of risk assessments in sequential decision-making tasks.

Local property. At every state of the system, the optimal decision taken by the robot should *not* depend on scenarios that the robot knows cannot occur in the future (see [28, 26] for a formal discussion¹). This concept is illustrated in Figure 4. The optimal decision taken by the robot at state x should *not* depend on the scenarios shaded in red.

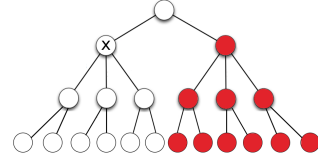


Fig. 4 Local property. The optimal decision taken by the robot at state x should *not* depend on scenarios that the robot knows cannot occur in the future (shaded in red).

Time consistency of risk assessments. Intuitively, time consistency stipulates that if a certain situation is considered less risky than another situation in all states of the world at time-step $k+1$, then it should also be considered less risky at time-step k . Before providing a formal definition, we note that failure to satisfy time consistency or the local property can lead to “irrational” behavior, including: (1) intentionally seeking to incur losses [21], or (2) deeming states to be dangerous when in fact they are favorable under any realization of the underlying uncertainty (we discuss an example of this below), or (3) declaring a decision-making problem to be feasible (e.g., satisfying a certain risk threshold) when in fact it is infeasible under *any* possible realization of the uncertainties [25].

As an example, consider the following planning problem with a CVaR cost. Given a Markov Decision Process (MDP) with initial state x_0 and time horizon $N \geq 1$, solve:

$$\tau^* := \min_{\pi} \text{CVaR}_{\alpha}(c_N(x_N)) \quad (5)$$

where $\alpha = 2/3$ and x_N is the state at time-step N . Consider the scenario tree (based on [4]) in Figure 5. Suppose we consider the solution of (5) acceptable if $\tau^* \leq 0$.

¹ We note that the local property is sometimes referred to as “time consistency”. However, here we use terminology that is consistent with the dynamic risk measurement literature [26].

One can then show that the optimization problem (over a single policy) results in an *unacceptable* solution since τ^* is positive. However, $\text{CVaR}_\alpha(c_N(x_N)) \leq 0$ is satisfied in every state of the world from the perspective of time-step $k = 1$. In other words, the decision maker would deem the solution of (5) unacceptable even though the solution appears acceptable from the perspective of the second stage ($k = 1$) under *any realization of the uncertainties*.

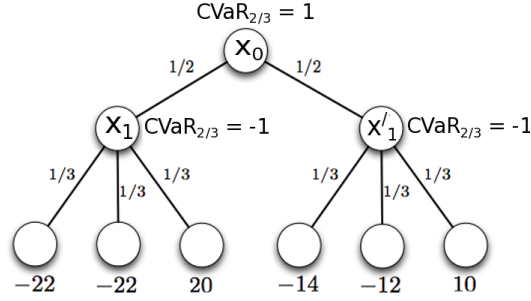


Fig. 5 Failure to satisfy time consistency can lead to “irrational” behavior. The numbers along the edges represent transition probabilities in the MDP, while the numbers below the terminal nodes represent the terminal costs $c_N(x_N)$. The problem involves a single control policy and there is thus a unique decision tree. The optimal cost appears acceptable at states x_1 and x'_1 , but unacceptable from the perspective of time-step 0 at state x_0 . In other words, the decision maker would deem the solution unacceptable even though the solution appears acceptable from the perspective of the second stage ($k = 1$) under *any realization of the uncertainties*.

We note that there is nothing special about this example. In general, simply applying a risk metric to the sum of all costs incurred at each time-step does not generally lead to time consistency. In order to obtain time-consistent measures of risk, we need to construct a *sequence* of risk metrics $\{\rho_{k,N}\}_{k=0}^N$, each mapping a future stream of random costs into a risk assessment at time-step k (visualized in Figure 6). Such risk metrics are known as *dynamic* risk metrics [26] since they assess risk at multiple points in time (in contrast to static risk metrics, which only assess risk from the perspective of the initial stage as in the example above). A dynamic risk metric $\{\rho_{k,N}\}_{k=0}^N$ is called time-consistent if, for all time-steps $0 \leq l \leq k \leq N$ and all sequences of stage costs $\{Z_i\}_{i=l}^N$ and $\{Z'_i\}_{i=l}^N$, the conditions

$$Z_i(\omega) = Z'_i(\omega), \forall \omega \in \Omega, \forall i = l, \dots, k-1, \quad \text{and} \quad \rho_{k,N}(Z_k, \dots, Z_N) \leq \rho_{k,N}(Z'_k, \dots, Z'_N),$$

imply that

$$\rho_{l,N}(Z_l, \dots, Z_N) \leq \rho_{l,N}(Z'_l, \dots, Z'_N).$$

Remarkably, it can be shown that one can construct a time-consistent risk metric by *compounding* one-step risk metrics [26]:

$$\rho_{k,N} = Z_k + \rho_k(Z_{k+1} + \rho_{k+1}(Z_{k+2} + \dots + \rho_{N-2}(Z_{N-1} + \rho_{N-1}(Z_N)) \dots)), \quad (6)$$

where the $\rho_k(\cdot)$ functions are a set of single-period risk metrics (satisfying mild technical conditions which are satisfied by distortion risk metrics). These single period metrics assess the risk of a random cost incurred at time-step $k + 1$ from the perspective of time-step k (we refer the reader to [26] for a more thorough introduction to time consistency of risk metrics). Moreover, under certain mild conditions, *any* time consistent risk metric is of the form (6) [26]. It can also easily be shown that in addition to time-consistency, compounding distortion risk metrics also leads to the local property being satisfied [29, Chapter 6.8.5].

Discussion 2 (Time consistency). If one adopts a distortion risk metric as the one-step risk metric that is compounded over time as in Equation (6), one inherits the “rationality” of the one-step assessments while also ensuring time consistency. This is the form of risk metric we advocate for sequential decision making tasks. However, imposing time consistency still comes at a conceptual cost. In particular, while a risk metric of the form $\rho(\sum_{k=0}^N Z_k)$ is easy to interpret, the composition of one-step metrics in Equation (6) stipulated by time consistency is more difficult to interpret. A direction for future work is thus to establish if this tension between time consistency and interpretability is an avoidable one.

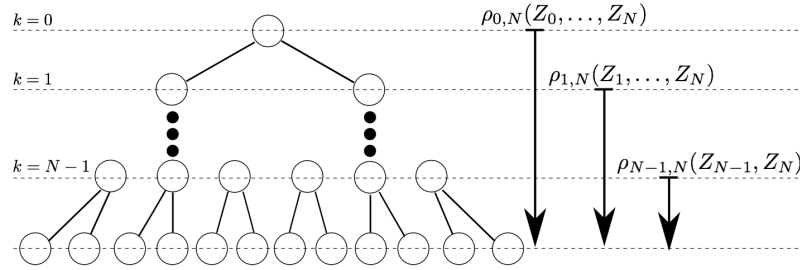


Fig. 6 Dynamic risk metrics assess risk at multiple points in time and lead to time-consistent risk assessments. Each $\rho_{k,N}$ maps a future stream of random costs to a risk assessment from the perspective of time-step k .

6 Discussion and Conclusions

Our goal in this paper has been to provide preliminary directions towards an axiomatic theory of risk for robotics applications. We have advocated properties that risk metrics employed by robots should satisfy in order for them to be considered sensible. These axioms define a class of risk metrics, known as distortion risk metrics, which have been previously used in finance. We further discussed properties that ensure the temporal consistency of risk assessments in sequential decision making tasks. We end with some questions that highlight areas for future thought in addition to the discussion points highlighted in Discussions 1 and 2 above.

Discussion 3 (Further axioms). While we have highlighted a number of axioms that we believe are particularly important, the identification of other axioms is an

important direction for future work. These may depend on the particular domain of application. Moreover, for certain applications it may not be necessary to impose all the axioms described here. For example, A4 and A5 (concerning diversification of risks) will generally be relevant to high-level decision making tasks where it is possible to diversify risks and may not be relevant for low-level control tasks where diversification may not be possible.

Discussion 4 (Choosing a particular risk metric). For a given application, we may wish to choose a particular risk metric from the class of metrics described here. How should such a metric be chosen? One possibility is to learn a distortion risk metric that explains how humans evaluate risk in the given application domain and then employ the learned risk metric. We describe first steps towards this in [20], where we have introduced a framework for *risk-sensitive inverse reinforcement learning* for learning humans’ risk preferences from the class of coherent risk metrics.

Discussion 5 (Legal frameworks). The question of safety for AI systems has received significant attention recently (see, e.g., [3] for a recent review). An important component of this discussion has been the consideration of legal frameworks and guidelines that must be placed on AI systems to ensure that they do not pose a threat to our safety. Such considerations for robots such as unmanned aerial vehicles are already extremely pressing for government agencies such as the Federal Aviation Administration (FAA). It is not difficult to imagine a future where the Robot Certification Agency (RCA) is in fact a real entity that certifies the safety of new robotic systems. How can we effectively engage lawmakers and government officials in discussions on how to evaluate risks in robotic applications?

Our hope is that the ideas presented in this paper will spur further work on this topic and eventually lead to a convergence upon a particular class of risk metrics that form the standard for assessing risk in robotics.

Acknowledgements The authors were partially supported by the Office of Naval Research, Science of Autonomy Program, under Contract N00014-15-1-2673.

References

- [1] C. Acerbi. Spectral measures of risk: A coherent representation of subjective risk aversion. *Journal of Banking & Finance*, 26(7):1505–1518, 2002.
- [2] M. Allais. Allais paradox. In *Utility and Probability*, chapter 2. Palgrave Macmillan UK, first edition, 1990.
- [3] D. Amodei, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, and D. Mané. Concrete problems in ai safety. *arXiv preprint arXiv:1606.06565*, 2016.
- [4] P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath. Coherent measures of risk. *Mathematical Finance*, 9(3):203–228, 1999.
- [5] D. Bertsimas and D. Brown. Constructing uncertainty sets for robust linear optimization. *Operations Research*, 57(6):1483–1495, 2009.
- [6] L. Blackmore, M. Ono, and B. C. Williams. Chance-constrained optimal path planning with obstacles. *IEEE Transactions on Robotics*, 27(6):1080–1094, 2011.
- [7] A. Charnes and W. W. Cooper. Chance-constrained programming. *Management Science*, 6(1):73–79, 1959.
- [8] G. Choquet. Theory of capacities. In *Annales de l’institut Fourier*, 1954.

- [9] E. Delage and Y. Ye. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations Research*, 58(3):595–612, 2010.
- [10] N. E. Du Toit and J. W. Burdick. Probabilistic collision checking with chance constraints. *IEEE Transactions on Robotics*, 27(4):809–815, 2011.
- [11] D. Ellsberg. Risk, ambiguity, and the savage axioms. *The Quarterly Journal of Economics*, 75(4):643–669, 1961.
- [12] H. Föllmer and A. Schied. *Stochastic Finance: An Introduction in Discrete Time*. Walter de Gruyter, 2011.
- [13] J. García and F. Fernández. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research*, 16(1):1437–1480, 2015.
- [14] I. Gilboa and M. Marinacci. Ambiguity and the Bayesian paradigm. In *Readings in Formal Epistemology*, chapter 21. First edition, 2016.
- [15] K. Glover and J. C. Doyle. Relations between H_∞ and risk sensitive controllers. In *Analysis and Optimization of Systems*. Springer-Verlag, 1987.
- [16] D. A. Iancu, M. Petrik, and D. Subramanian. Tight approximations of dynamic risk measures. *Mathematics of Operations Research*, 40(3):655–682, 2015.
- [17] D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, pages 263–291, 1979.
- [18] S. Kuindersma, R. Grupen, and A. Barto. Variable risk control via stochastic optimization. *Int. Journal of Robotics Research*, 32(7):806–825, 2013.
- [19] F. Maccheroni, M. Marinacci, A. Rustichini, and M. Taboga. Portfolio selection with monotone mean-variance preferences. *Mathematical Finance*, 19(3):487–521, 2009.
- [20] A. Majumdar, S. Singh, A. Mandlekar, and M. Pavone. Risk-sensitive inverse reinforcement learning via coherent risk models. In *Robotics: Science and Systems*, 2017. In press.
- [21] S. Mannor and J. N. Tsitsiklis. Mean-variance optimization in Markov decision processes. In *Int. Conf. on Machine Learning*, 2011.
- [22] B. C. on Banking Supervision. *Fundamental Review of the Trading Book: A Revised Market Risk Framework*. Bank for International Settlements, 2014.
- [23] M. Ono, M. Pavone, Y. Kuwata, and J. Balaram. Chance-constrained dynamic programming with application to risk-aware robotic space exploration. *Autonomous Robots*, 39(4):555–571, 2015.
- [24] R. T. Rockafellar and S. Uryasev. Optimization of conditional value-at-risk. *Journal of Risk*, 2:21–41, 2000.
- [25] B. Roorda, J. M. Schumacher, and J. Engwerda. Coherent acceptability measures in multi-period models. *Mathematical Finance*, 15(4):589–612, 2005.
- [26] A. Ruszczyński. Risk-averse dynamic programming for Markov decision process. *Mathematical Programming*, 125(2):235–261, 2010.
- [27] D. Schmeidler. Integral representation without additivity. *Proceedings of the American Mathematical Society*, 97(2):255–261, 1986.
- [28] A. Shapiro. On a time consistency concept in risk averse multi-stage stochastic programming. *Operations Research Letters*, 37(3):143–147, 2009.
- [29] A. Shapiro, D. Dentcheva, and A. Ruszczyński. *Lectures on stochastic programming: Modeling and theory*. SIAM, second edition, 2014.
- [30] T. Summers, J. Warrington, M. Morari, and J. Lygeros. Stochastic optimal power flow based on conditional value at risk and distributional robustness. *International Journal of Electrical Power & Energy Systems*, 72:116 – 125, 2015.
- [31] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.
- [32] S. Wang. A class of distortion operators for pricing financial and insurance risks. *Journal of Risk and Insurance*, 67(1):15–36, 2000.
- [33] P. Whittle. Risk-sensitive linear/quadratic/gaussian control. *Advances in Applied Probability*, 13(4):764–777, 1981.
- [34] H. Xu and S. Mannor. Distributionally robust Markov decision processes. In *Advances in Neural Information Processing Systems*, 2010.