Network for Studies on Pensions, Aging and Retirement

NETSPAR ACADEMIC SERIES

Comparative risk aversion vs. threshold choice in the Omega ratio

Anne Balter, Ki Wai Chau, Nikolaus Schweizer

DP 10/2023-058

Omega 123 (2024) 102992

Contents lists available at ScienceDirect

Omega

journal homepage: www.elsevier.com/locate/omega

Comparative risk aversion vs. threshold choice in the Omega ratio ☆

Anne G. Balter [a], Ki Wai Chau [b], Nikolaus Schweizer [c,*]

[a] Department of Econometrics and Operations Research & Netspar, Tilburg University, The Netherlands
[b] Department of Economics, Econometrics and Finance, University of Groningen, The Netherlands
[c] Department of Econometrics and Operations Research, Tilburg University, P.O. Box 90153, 5000 LE Tilburg, The Netherlands

ARTICLE INFO

JEL classification:
D81
G11

Keywords:
Omega ratio
Risk preferences
Decision criteria
Background risk

ABSTRACT

We study conditions under which the threshold parameter in the Omega ratio represents risk aversion in the sense of monotonicity of risk premia. To this end, we derive asymptotic expansions for risk premia associated with taking a small additional risk on top of a background risk. These risk premia have the expected monotonicity behavior if, roughly speaking, the variance of the additional risk decreases with the background risk and if the density of the background risk is log-concave. When these conditions are violated, the threshold in the Omega ratio does not represent risk aversion in general. Finally, we compare our sufficient conditions for the Omega ratio to those that are needed to guarantee monotonicity of risk premia with an expected utility criterion under background risk. We argue that the conditions that are needed for the Omega threshold to represent risk aversion are comparable to those that are needed for expected utility with exponential utility functions.

1. Introduction

Ever since its proposal by Keating and Shadwick [12], ranking risky prospects by the Omega ratio has been a popular alternative to utility-based approaches on the one side and criteria like the Sharpe ratio on the other. The Omega ratio is a risk-return ratio which compares the expected gains above a given threshold to the expected losses below it. Plotting the Omega ratio Ω_K as a function of the threshold K gives a representation of the entire distribution and enables a user to quickly compare, e.g., different investment opportunities by their Omega curves. These plots have been especially popular with practitioners. In the academic literature, a surprisingly controversial literature has discussed whether the risk preferences captured by the Omega ratio are in line with risk-averse behavior, see e.g. [7,13,14] and, for a recent survey and summary, [8].[1] The central question in this literature is whether a ranking in the sense of second-order stochastic dominance (SSD), i.e., a ranking that all risk-averse expected utility maximizers can agree upon, implies the same ranking in the Omega ratio.[2] In a nutshell, this is the case if attention is restricted to risks X whose Omega ratio is greater than 1 or, equivalently, whose expected value $\mu = E[X]$ is above the threshold K.[3] Thinking in terms of a plot of Omega against the threshold, consistency with risk-averse behavior is thus ensured if we cut off the y-axis at $\Omega = 1$ and only keep the upper part.

In this paper, we address a direct follow-up question. Given that consistency with risk-averse behavior is guaranteed for a range of threshold values, how should we interpret an increase in the threshold? In particular, we study conditions for the validity of a claim that is sometimes made,[4] especially in the practitioner-oriented literature, providing one of the main rationales behind plots of the Omega ratio: An increase in the threshold K corresponds to an increase in risk appetite and thus, conversely, to a decrease in risk aversion. Thinking in terms of Omega plots, we thus ask what the x-axis means when it comes to risk preferences. In contrast, the previous literature on consistency with risk-averse behavior – as surveyed in [8] – has mostly focused on the ranking across the y-axis for a fixed point on the x-axis, i.e., for a fixed threshold K.

☆ Area: Decision Analysis and Preference Modeling. This manuscript was processed by Associate Editor Salvatore Corrente.
* Corresponding author.
E-mail address: n.f.f.schweizer@uvt.nl (N. Schweizer).

[1] There is also a sizeable and closely related literature on using the Omega ratio as a decision and optimization criterion in operations research and related fields, see e.g. [1–3], and [4].

[2] See, e.g., [5] for a more general discussion of decision criteria and their consistency with SSD. For background on SSD from the perspective of the theory of stochastic orders, see, e.g., [6] (where it corresponds to the "increasing concave order").

[3] We refer, e.g., to [7] or [8] for a more in-depth discussion and more precise statements.

[4] See, for instance, [9,10] or [11].

https://doi.org/10.1016/j.omega.2023.102992
Received 23 January 2023; Accepted 22 October 2023
Available online 29 October 2023

1.1. Overview of main results

Our analysis starts with two ideas that go back to Pratt [15] who addressed similar questions for expected utility. First, in order to rank different decision criteria in terms of their degree of risk aversion, we consider risk premia, i.e., the required compensations that make an agent indifferent whether to accept a certain risk or not. If the threshold K captures the degree of risk appetite, the opposite of risk aversion, then risk premia should be decreasing with K. Second, in order to make this question analytically tractable, we focus mostly on risk aversion "in the small". Thus, we consider Taylor expansions of the risk premium in the limit of vanishing risk, and study the monotonicity behavior of the leading order terms. Our main results are sufficient conditions for these leading order terms to be decreasing in K. These results guarantee that the interpretation of K as measuring the degree of risk appetite are valid at least for sufficiently small risks.

In line with the previous literature on consistency with risk-averse preferences, we restrict attention throughout to risky prospects X whose mean μ lies above the threshold, $\mu \geq K$. However, when discussing risk premia, some additional restrictions are natural. First, we mostly rule out the limiting case $\mu = K$: An agent whose preferences are captured by the Omega ratio with threshold K is indifferent between all risks X with mean $E[X] = K$, implying a risk premium of zero. For risks with mean K, the agent is thus effectively risk-neutral.[5] Second, the Omega of a non-risky prospect $x > K$ is infinite since the gain/loss ratio becomes $(x - K)/0$. Accordingly, an agent with Omega ratio preferences will want an infinitely high compensation for exchanging a certain prospect $x > K$ against *any* risky prospect with finite mean and unbounded support, no matter how favorable. Consequently, we limit attention to choices between risky prospects.

We thus arrive at a setting that is closely related to models from the classical literature on utility-based choice under background risk.[6] An agent decides between a risky prospect X and another risky prospect $X + hS + \pi$. Here, $h > 0$ is a positive constant while X and S are two possibly correlated random variables with $E[X] = \mu > K$ and $E[S|X] = E[S] = 0$. $X + hS$ can thus be seen as a mean-preserving spread of X. The risk premium π is chosen in such a way that $X + hS + \pi$ has the same Omega ratio as X itself, thus corresponding to the indifference price of the additional risk.

As noted above, the literature on consistency of the Omega ratio and risk aversion has shown that π is always positive in this setting. We want to know whether π is also decreasing in K, confirming the interpretation of K as a measure of risk appetite. As a first main result, we provide an explicit expansion of π around $h = 0$, similar to expansions for utility-based risk premia from [15] and the subsequent literature on background risk. We then show that risk premia are decreasing in K for sufficiently small h if two conditions are satisfied, (i) the probability density g of the background risk X is log-concave[7] in the lower tail, i.e., below K and (ii) the conditional variance of the additional risk $\sigma^2(x) = \mathrm{Var}(S|X = x)$ is weakly decreasing in x for $x \leq K$.[8]

Thus, we find that K can be interpreted as a measure of risk appetite under natural conditions. Nevertheless, these conditions are not *always* satisfied. In particular, while the log-concavity condition is somewhat stronger than needed,[9] we can easily create counterexamples by relaxing it. In particular, we find that when X follows a Student-t distribution and S is independent of X, risk premia are non-monotonic in K. Log-concavity can be interpreted as combining two types of conditions, a local regularity condition that e.g. rules out jumps in the density and a global regularity condition that rules out tail behavior which is heavier than exponential. The example of the Student-t distribution shows that heavy tails can indeed destroy the validity of interpreting K as a measure of risk aversion. This matters because the Omega ratio's ability to handle heavy-tailed risks has been cited as an important advantages over alternative criteria.[10]

As noted above, rankings based on the Omega ratio are consistent with risk aversion as long as means are finite and above the threshold. However, our result shows that plots of Omega against the threshold should be interpreted with caution because the interpretation of the threshold as capturing risk appetite breaks down when the lower tails are too heavy. In Section 3.2, we demonstrate some additional plots that can be made to get a sense of whether such violations are a concern in a particular application. Of course, taking an agnostic stand, one can plot the Omega against the threshold without claiming that the threshold has any interpretation beyond it being the threshold. Yet then, it is not easy to say what we learn precisely from a plot of the Omega against its threshold. What we do see in such a plot is a collection of rankings that are all consistent with risk averse behavior. To make a choice between these rankings, a decision-maker could, for instance, determine a single threshold based on economic criteria (such as using a risk-free return as a benchmark like in the Sharpe ratio), reducing the Omega curves to single points again. Alternatively, the decision-maker could use the Omega plot as a starting point to filter out a number of promising alternatives, investigating those alternatives that are ranked highest at different threshold levels in more detail with additional tools. Finally, as we briefly discuss in Section 3.2, our results on the connection between risk premia and threshold levels can also be used to translate risk premia (which are, arguably, more intuitive) into threshold levels.

1.2. Comparison with expected utility

For our second set of results, we compare our findings about Omega-based risk premia to the situation for risk premia based on classical parametric families of utility functions like the power (CRRA) and the exponential (CARA) family. What can we say about monotonicity of risk premia under background risk when we vary the risk aversion parameter in either of these families? How do the conditions on the distribution of the background risk X and the additional risk S compare? Put differently: Suppose you plot risk premia based on power utility or exponential utility against the "risk aversion parameter" of these families of utility functions. Will there be the same problems regarding the interpretation of the x-axis?

For this situation of background risk, the classical references are [19,20]. Leaving $\sigma^2(x) = \mathrm{Var}(S|X = x)$ largely unrestricted, Ross [19] shows that a rather rigid notion of comparative risk aversion between two utility functions needs to be satisfied to guarantee that differences in risk premia have the expected sign. In particular, two power utility or two exponential utility functions can generally not be ranked in this strong order, implying that under background risk

[5] This observation is closely related to the discussion of "non-strict dominance compatibility" in [8].

[6] See [16] for an introduction and [17] for a recent contribution with many references.

[7] A function is called log-concave if its logarithm is concave.

[8] Throughout the paper, we use the terms "decreasing" and "weakly decreasing" synonymously to indicate that for a function $f : \mathbb{R} \to \mathbb{R}$ the condition $x < y$ implies $f(x) \geq f(y)$ for all $x, y \in \mathbb{R}$. We add the "weakly" mostly to emphasize that we are not ruling out constant functions which are often among the most important special cases.

[9] To be precise, while log-concavity of the density function is a stronger condition than needed, we provide another log-concavity condition, which is basically a necessary and sufficient condition, see Lemma 3, Corollary 2 and the surrounding discussion.

[10] See, e.g., [18] who argue that, unlike the Sharpe ratio, the Omega can account for heavier than normal tail behavior as captured by higher moments like skewness and kurtosis. At the same time, the Omega can be estimated without actually estimating these higher moments — which may be a challenge with limited amounts of data.

their parameters do not represent a degree of risk aversion in general. In contrast, the main result of Kihlstrom et al. [20] implies that when $\sigma^2(x) \equiv \bar{\sigma}^2$ is constant and when one of the utility functions under consideration has decreasing absolute risk aversion, Pratt [15]'s notion of comparisons in absolute risk aversion carries over and risk premia have the desired monotonicity behavior. Because power and exponential utility functions exhibit (weakly) decreasing absolute risk aversion, it follows that changes in the "risk aversion parameters" go together with the desired monotonicity behavior of risk premia.

Thus, it seems that with constant σ^2, no further conditions need to be imposed on the distribution of X to guarantee that the parameters in the classical families of utility functions can be interpreted as degrees of risk aversion under background risk. In contrast, our results for the Omega ratio required an additional log-concavity condition. Yet this view is a bit misleading since the regularity conditions that need to be satisfied for the criteria to be well-defined are stronger for expected utility: Power utility cannot handle negative realizations while exponential utility requires existence of a moment-generating function which is closely related to log-concavity. In this sense, the requirements for the Omega case and for exponential utility are remarkably similar. The Omega ratio needs less integrability to be well-defined but some additional regularity to justify the interpretation of K as a degree of risk aversion. In many counterexamples where this interpretation breaks down, risk premia based on expected utility are not available as an alternative.

Comparing our results for the Omega ratio to those of Ross [19] and Kihlstrom et al. [20] for expected utility – in particular exponential and power utility – a gap between these two classical papers becomes apparent: While the negative result of Ross [19] leaves $\sigma^2(x)$ completely unrestricted, [20] limit attention to the extreme case where σ^2 is independent of x. In contrast, our positive results for the Omega ratio merely need that $\sigma^2(x)$ is weakly decreasing in x. We show that the main result of Kihlstrom et al. [20] can be extended to weakly decreasing $\sigma^2(x)$, thus closing part of this gap. In fact, we can replace their assumption of decreasing absolute risk aversion by assuming that the product of $\sigma^2(x)$ and the absolute risk aversion function is weakly decreasing.

1.3. A quartet of decision criteria

We close the introduction with some more discussion of how the Omega ratio fits into the broader decision-theoretic picture. Since the 1970s, various competitors to expected utility theory have been proposed for a variety of reasons. One departure from expected utility that is shared by (perhaps surprisingly) many of these competitors is the introduction of a reference or aspiration level to which possible outcomes are compared. Besides the Omega with its threshold parameter, this includes prominent models such as (cumulative) prospect theory [21,22], regret theory [23,24], habit formation [25] and the reference-dependent preferences of [26]. In fact, as a decision-criterion, the Omega forms part of a quartet with three other models of this type that base decisions, respectively, on success probabilities, quantiles and expectiles.

Basing decisions on maximizing success probabilities like the probability of exceeding a certain target level is a very simple and intuitive criterion. In the literature, it has been discussed, e.g., in the context of the "security-potential/aspiration" theory of Lopes [27] or, under the name "utility mass model" by Manski [28].[11] In fact, Manski discusses this criterion as one of a pair of models, the other being the maximization of a given quantile of the outcome distribution. The two models are, of course, closely related due to the mathematical connection between the quantile function and the cumulative distribution function.[12] Besides their widespread use in risk management under the name Value-at-Risk (e.g., [30]), quantiles as decision criteria have been studied, among others, by [31], [32,33] and [34].

Expectiles were first proposed by Newey and Powell [35] who investigated an alternative to quantile regression that replaces the usual weighted-absolute-error criterion by a weighted least-squares criterion. Newey and Powell [35] called the minimizers of these criteria expectiles and noted that "expectiles have properties that are similar to quantiles". In fact, [36] showed that there is a non-linear transformation of the original distribution such that the quantiles of the transformed distribution coincide with the expectiles of the original distribution. Under the same transformation, the Omega curve of the original distribution matches the odds curve of the transformed distribution.[13] Since success probabilities p and odds of success $\frac{p}{1-p}$ are monotonically related, maximizing the odds of success is equivalent to maximizing the success probability. Consequently, the relation between Omega ratios and expectiles as decision criteria is completely analogous to the relation between success probabilities and quantiles.

In the past decade, expectiles have been studied quite actively[14] from a decision-theoretic perspective, starting with the observation of Bellini et al. [39] and Ziegel [40] that – when viewed as financial risk measures – certain expectiles are the only risk measures satisfying the axioms of coherence and elicitability. The study of the connection between Omega ratios and expectiles was initiated in Bellini et al. [41]. They show that when the entire curves of Omegas between two payoffs can be ranked, i.e. if the curves do not cross, then the same is true for the entire curves of expectiles – and vice versa. In these cases, the criteria are equivalent. However, the criteria are not equivalent in general – analogously to Manski's utility mass and quantile models which are not equivalent either.

1.4. Structure

The paper is organized as follows: Section 2 introduces the problem, the setting and some technical assumptions. Section 3.1 contains our technical main result, the asymptotic expansions for Omega-based risk premia. In Section 3.2, we illustrate these results in an example, and in Section 3.3, we contrast our results on risk premia for mean-preserving spreads to analogous results for location-scale families. Section 4 discusses related results for risk premia based on expected utility. Some technical proofs are found in Appendix A. Finally, Appendix B recalls some details on the connection between Omegas and expectiles while Appendix C provides some technical details on risk premia based on expected utility under background risk.

2. The setting

Let $(\Theta, \mathcal{F}, P)$ be a probability space on which all subsequent random variables are defined.[15] For a real-valued, integrable random payoff X and a real number K, the Omega ratio $\Omega_K(X)$ with threshold K is defined as

$$\Omega_K(X) = \frac{E[(X-K)^+]}{E[(K-X)^+]}$$

[11] See [29] for a recent empirical contribution with a broad overview of this literature.

[12] Manski also notes that these two models share one "disagreeable feature": They have large indifference classes, i.e., large sets of payoff distributions that are considered equivalent. As discussed above, this feature is also shared by the Omega ratio.

[13] For the reader's convenience, we summarize the mathematics behind these facts in Appendix B.

[14] For recent contributions in this literature, see, e.g., [37,38] and the references therein.

[15] A concrete example could be $\Theta = \mathbb{R}^n$ equipped with the Borel sigma algebra $\mathcal{F}$ where n is the number of real-valued random variables under consideration.

where $E[\cdot]$ denotes the expected value and $(\cdot)^+$ the positive part. The Omega ratio is thus the ratio between the expected gains and the expected loss in reference to the threshold K. The Omega ratio is non-negative and decreasing in K. It takes the values 0 or ∞ when $E[(X-K)^+]=0$ or $E[(K-X)^+]=0$. We extend the definition to the case where $X=K$ holds almost surely[16] by setting $\Omega_K(K)=1$. Introducing the notation $P_X(K)=E[(K-X)^+]$, $C_X(K)=E[(X-K)^+]$ and $\mu_X=E[X]$, we can rewrite the Omega as 1 plus the "Sharpe-Omega" of Kazemi et al. [42],

$$\Omega_K(X)=\frac{C_X(K)}{P_X(K)}=1+\frac{\mu_X-K}{P_X(K)}$$

using the "put–call-parity" for expected payoffs, $C_X(K)=P_X(K)+\mu_X-K$.

The literature on consistency of the Omega ratio and risk-averse behavior (e.g. [8]) has shown that consistency holds whenever $\mu_X\geq K$ or, equivalently, whenever $\Omega_K(X)\geq 1$.[17]

Admissible thresholds: We want to study whether an increase in the gain–loss threshold K makes risk preferences based on the Omega ratio less risk-averse. For this question to be meaningful, we limit attention to situations in which the Omega ratio is finite and in line with risk-averse behavior. We say that the threshold K is *admissible* for X if $K<\mu_X$ and if $P_K(X)=E[(K-X)^+]>0$. The latter condition is equivalent to $\text{Prob}(X<K)>0$. Thus, K is admissible for X if it lies in the support of X but below its mean. Formally, $\mathcal{R}(X)$, the range of admissible thresholds for X is given by the interval

$$\mathcal{R}(X)=(x_{\min},\mu_X)\ \text{ with }\ x_{\min}=\inf\{K\in\mathbb{R}|\text{Prob}(X<K)>0\}$$

where we allow $x_{\min}=-\infty$. When X is deterministic, $X=x$ a.s. for some real number x, then $\mathcal{R}(X)$ is empty. Conversely, when X is not deterministic, then $\mathcal{R}(X)$ is non-empty. When X has full support on $\mathbb{R}$ then $\mathcal{R}(X)=(-\infty,\mu_X)$.

By restricting attention to admissible thresholds we thus effectively rule out deterministic payoffs X. The Omega ratio was not designed for comparing those – and it is also not very good at it. An agent whose preferences are based on the Omega ratio with threshold K prefers a deterministic payoff $x>K$ over any payoff with finite Omega ratio since $\Omega_K(x)=\infty$. Similarly, the agent is indifferent between any pair x and y of deterministic payoffs that lie above K.

The condition $K\in\mathcal{R}(X)$ is equivalent to $1<\Omega_K(X)<\infty$. Thus, in line with the literature on consistency with risk-averse behavior, admissibility means that we are in the range of thresholds where consistency holds and, additionally, where Ω_K is finite. Finally, note that $\mathcal{R}(X)$ is an open interval. Thus, $K\in\mathcal{R}(X)$ implies $K\pm\varepsilon\in\mathcal{R}(X)$ for sufficiently small $\varepsilon>0$. We can thus always vary K a little bit without leaving $\mathcal{R}(X)$.

Risk premia and risk appetite: In order to investigate whether the threshold K in the Omega ratio is a measure of risk appetite, we study how premia for taking up additional risks vary with K. To this end, consider some non-deterministic, integrable baseline risk X and let $K\in\mathcal{R}(X)$ be a threshold. Denote by S another integrable, real-valued random variable with $E[S|X]=0$. Thus, $X+S$ is a mean-preserving spread of X. Since $K\in\mathcal{R}(X)$, $X+S$ has a weakly smaller Omega ratio than X, see also (1) below. The risk premium π_K is defined as the deterministic monetary amount that needs to be added to $X+S$ in order to compensate the additional risk and make the agent indifferent. The following lemma shows that risk premia are well-defined and gives some basic properties.

[16] This is the only case where $E[(X-K)^+]=0$ and $E[(K-X)^+]=0$ hold simultaneously.

[17] In a nutshell, increasing risk increases the expected "option payoff" $P_X(K)$. Whether this translates into an increase or decrease in risk depends on the sign of μ_X-K. In particular, a strict increase in $P_X(K)$ (while keeping μ_X fixed) will translate into a strict decrease of the Omega whenever $\mu_X>K$.

Lemma 1. *Assume that X and S are integrable payoffs with $K\in\mathcal{R}(X)$ and $E[S|X]=0$. Then, there exists a unique risk premium $\pi_K\in[0,\infty)$ which solves*

$$\Omega_K(X)=\Omega_K(X+S+\pi_K).$$

Moreover, $K\in\mathcal{R}(X+S+\pi_K)$.

Proof. By Jensen's inequality, $E[(K-X-S)^+]\geq E[(K-X)^+]$ which implies

$$\Omega_K(X+S)=1+\frac{E[X+S]-K}{E[(K-X-S)^+]}\leq 1+\frac{E[X]-K}{E[(K-X)^+]}=\Omega_K(X) \quad (1)$$

since $E[S]=0$. Next, note that $E[(K-X-S-\pi_K)^+]$ is decreasing in π_K so that

$$\Omega_K(X+S+\pi_K)=1+\frac{E[X]+\pi_K-K}{E[(K-X-S-\pi_K)^+]}$$

is strictly increasing in π_K and going to ∞ as π_K goes to ∞. Consequently, there exists a unique, finite $\pi_K\geq 0$ which solves $\Omega_K(X+S+\pi_K)=\Omega_K(X)$. Finally, $1<\Omega_K(X)<\infty$ implies $1<\Omega_K(X+S+\pi_K)<\infty$ which can only hold if $E[(K-X-S-\pi_K)^+]>0$ and $E[X]+\pi_K>K$. Thus $K\in\mathcal{R}(X+S+\pi_K)$. □

We are now ready to make our main research question precise: Does the threshold K represent risk appetite in the sense that an increase in K leads to a decrease in π_K?

Definition 1. Assume that X and S are integrable payoffs with $K\in\mathcal{R}(X)$ and $E[S|X]=0$. Denote by π_K the unique risk premium defined via

$$\Omega_K(X)=\Omega_K(X+S+\pi_K). \quad (2)$$

We say that K represents risk appetite at $(X,X+S)$ if π_K is decreasing in $K\in\mathcal{R}(X)$.

In the remainder of this paper, we study under which conditions K represents risk appetite in this sense of monotonicity of risk premia. This definition of comparative risk aversion is in line with a large literature starting e.g. with Pratt [15]'s famous work on risk premia under expected utility preferences.

Technical assumptions and notation: So far, we have made very few assumptions on the distributions of X and S, leaving it, e.g., open whether distributions are discrete or continuous. In particular, the notions of risk premia and increasing risk appetite apply both in the discrete and in the continuous case and can, e.g., be applied to Omegas from empirical distributions. For our theoretical main results in the following sections, we need to assume more however.

We assume that the baseline risk X is integrable and continuously distributed with bounded density function g with respect to the Lebesgue measure. We denote by G the associated cumulative distribution function. Moreover, we write $\mu=E[X]$, $P(K)=E[(K-X)^+]$ and $C(K)=E[(X-K)^+]$, dropping the subscript X from the previous notation. For the additional risk, we assume that S is square-integrable with $E[S|X]=0$. We also define the conditional variance function

$$\sigma^2(x):=\text{Var}(S|X=x). \quad (3)$$

For some of our results, in particular the asymptotic expansions of risk premia, we need additional regularity assumptions. Intuitively, we face a technical challenge that is well-known from the quantitative finance literature on computing sensitivities of standard options like puts and calls.[18] Since the "payoffs" $(x-K)^+$ and $(K-x)^+$ are weakly differentiable once but not twice, some structure is needed to ensure existence of higher order derivatives of their expectations. To this end, we introduce the following additional assumption on the joint distribution of X and S which is invoked where needed.

[18] See, e.g., Chapter 7 of Glasserman [43].

Assumption 1. Denote by $l_{X,S}$ the joint density of X and S with respect to the Lebesgue measure. We assume that $l_{X,S}$ satisfies the following two regularity conditions:

1. (Global convergence) There exists a bounded, integrable function $b : \mathbb{R} \to \mathbb{R}$ and a constant $\kappa > 3$ such that

$$l_{X,S}(x,s) \le b(x)\varphi(s) \tag{4}$$

where $\varphi : \mathbb{R} \to \mathbb{R}$ is defined by

$$\varphi(s) := \begin{cases} \frac{1}{|s|^{\kappa}} & \text{for } |s| \ge 1; \\ 1 & \text{for } |s| < 1. \end{cases} \tag{5}$$

2. (Local continuity) For all real numbers x, y and s, we have

$$|l_{X,S}(x,s) - l_{X,S}(y,s)| \le C \max(1,|s|)|x-y|, \tag{6}$$

where C is a constant independent of x, y and s.

To illustrate Assumption 1 and the families of distributions for which it holds, we first have a look at the case where X and S are independent:

Example 1. Suppose X and S are independent with respective density functions $g : \mathbb{R} \to \mathbb{R}$ and $p : \mathbb{R} \to \mathbb{R}$ so that $l_{X,S}(x,s) = g(x)p(s)$ and so that $\sigma^2(x) = \mathrm{Var}(S)$ is constant in x. Then Assumption 1 is satisfied under the following conditions on g and p:

(a) X is integrable and g is bounded from above and Lipschitz continuous on $\mathbb{R}$.
(b) S is square-integrable with $E[S] = 0$, p is bounded from above and there exists a constant $C > 0$ such that $p(s) \le \frac{C}{|s|^{\kappa}}$ for all $s \in \mathbb{R}$ with $|s| \ge 1$ and for some $\kappa > 3$.

Condition (a) is satisfied, e.g., for the normal distribution, the Laplace distribution, the Gamma distribution (with $\alpha \ge 2$),[19] the Beta distribution (with $\alpha \ge 2$ and $\beta \ge 2$), the Lognormal distribution, the Weibull distribution (with $\alpha \ge 2$), the skew-normal distribution, and the t-distribution with $\nu > 1$. Here, the parameter constraint for the t-distribution is needed to ensure integrability while the remaining parameter constraints ensure Lipschitz continuity (i.e. a finite slope in $x = 0$ and, in case of the Beta distribution, also in $x = 1$). Up to a possible shift to ensure $E[S] = 0$, condition (b) is satisfied for the same list of distributions, except that we need to tighten the parameter constraint for the t-distribution to $\nu > 2$ to ensure square-integrability and the density bound $p(s) \le C|s|^{-\kappa}$. For the Gamma distribution, we can relax the parameter constraint to $\alpha \ge 1$ because we only need boundedness but no longer Lipschitz continuity. Similar reasoning applies to the Beta and to the Weibull distribution. Finally, (b) is also satisfied by distributions like a uniform distribution on $[-w, w]$ for some $w > 0$ whose density is bounded from above but not Lipschitz continuous.

Clearly, the independence assumption in the previous example is somewhat restrictive. The next lemma shows that Assumption 1 also holds under fairly mild conditions in the "heteroskedastic" case where $S = \Sigma(X)Z$ and Z is a mean-zero random variable that is independent of X, implying $\sigma^2(x) = \Sigma(x)^2\mathrm{Var}(Z)$ in this case.

Lemma 2. *Let S be given by $S = \Sigma(X)Z$ where $\Sigma(\cdot)$ is a non-negative, continuous function and Z is a continuously distributed, mean-zero random variable which is independent of X and has density h. Assume that there are constants $C > 0$ and $\kappa > 3$ such that $h(z) \le C\varphi(z)$ for $\varphi(z)$ defined in (5). Moreover, assume that g, h and $\frac{1}{\Sigma(\cdot)}$ are bounded and Lipschitz continuous and that the expected values $E[\Sigma(X)^{\kappa-1}]$ and $E[\frac{1}{\Sigma(X)}]$ are finite. Then Assumption 1 holds.*

[19] Regarding the parametrization of the Beta, Gamma and Weibull distributions, we follow [44].

The proof of the lemma is given in the appendix. Unlike in Example 1, we now also need Lipschitz continuity for the density h of Z. Thus, the conditions of the lemma are, e.g., satisfied if (up to a possible shift to ensure $E[Z] = 0$) both g and h come from the normal distribution, the Laplace distribution, the Gamma distribution (with $\alpha \ge 2$), the Beta distribution (with $\alpha \ge 2$ and $\beta \ge 2$), the Lognormal distribution, the Weibull distribution (with $\alpha \ge 2$), the skew-normal distribution, or the t-distribution with $\nu > 2$ (where $\nu > 1$ is sufficient for X). For the function $\Sigma(\cdot)$, we can, e.g., assume that Σ is Lipschitz continuous and there exist real numbers Σ_0 and Σ_1 such that $0 < \Sigma_0 \le \Sigma(x) \le \Sigma_1$ for all $x \in \mathbb{R}$.

3. Main results

3.1. Asymptotic expansions

Our goal is to understand how risk premia based on the Omega ratio depend on the threshold K. To make this question tractable, we study $\pi_K(h)$, the risk premium for taking on the additional risk $h\,S$ on top of the existing risk X where $h > 0$ is a real-valued parameter. Formally, in line with (2), for $K \in \mathcal{R}(X)$, $\pi_K(h)$ is defined as the unique solution to the equation

$$\Omega_K(X) = \Omega_K(X + hS + \pi_K(h))$$

which is well-defined by Lemma 1. The next proposition provides our asymptotic expansion of $\pi_K(h)$ around $\pi_K(0) = 0$.

Proposition 1. *Under Assumption 1, for $K \in \mathcal{R}(X)$ and $h > 0$, we have the expansion*

$$\pi_K(h) = \sigma^2(K)R(K)\frac{h^2}{2} + O(h^3) \tag{7}$$

where $\sigma^2(K)$ is defined in (3) and where

$$R(K) = \frac{(\mu - K)g(K)}{P(K) + (\mu - K)G(K)} \ge 0. \tag{8}$$

The proof of the proposition is given in the appendix. The proposition shows that the leading order term in $\pi_K(h)$ depends on K through the product of two non-negative factors, a factor $R(K)$ that only depends on the distribution of X and a second factor $\sigma^2(K)$ which only depends on the distribution of S conditional on X. If both factors are weakly decreasing in K, then risk premia are decreasing in K, in line with the interpretation of K as risk appetite. This is summarized in the following corollary:

Corollary 1. *In the setting of Proposition 1 and for sufficiently small $h > 0$, if $R(K)$ and $\sigma^2(K)$ are both weakly decreasing in $K \in \mathcal{R}(X)$ then risk premia are weakly decreasing in K. Thus, K represents risk appetite at $(X, X + hS)$ for all sufficiently small h.*

In the language of Pratt [15], Corollary 1 mixes elements of a result about risk aversion "in the small" (locally) and risk aversion "in the large" (globally). While the result is "in the small" in the sense that we only consider sufficiently small (additional) risks, it is "in the large" in the sense that we study a range of thresholds, in line with our goal of understanding how to interpret plots of Omegas against their thresholds.

In the corollary, we could, of course, also allow for one factor being increasing if the other decreases sufficiently strongly. What matters is how their product behaves. When X and S are independent, the function $\sigma^2(\cdot)$ is constant so R needs to be decreasing, reducing the two conditions to a single one. Moreover, recall that $\mathcal{R}(X)$ is the lower part of the support of X up to the mean μ. Thus, monotonicity only needs to be guaranteed in that range. For example, if $\sigma^2(x)$ is largest in the tails of the distribution of X in the sense that $\sigma^2(x) = \psi(|x - \mu|)$ for some positive and increasing function ψ, then $\sigma^2(x)$ is decreasing over $\mathcal{R}(X)$.

To understand why the monotonicity behavior of $\sigma^2(x)$ matters, suppose for a moment that $\sigma^2(\cdot)$ is increasing over $\mathcal{R}(X)$ so that the additional risk S has the largest conditional variance when X is realized in the center of its distribution. Recall that the risk measure $P(K)$ in the Omega ratio only considers the riskiness of the tail below K. Thus, a risk premium π_K that only accounts for the lower tail ($K \ll \mu$) may well be smaller than a risk premium that also looks at the center ($K \approx \mu$). The following example illustrates this.

Example 2. Suppose that $\sigma^2(K) = \bar{\sigma}^2 1_{\{K>t\}}$ for some real numbers $\bar{\sigma}^2 > 0$ and $t \in \mathcal{R}(X)$ where $1_{\{\cdot\}}$ denotes the indicator function. Then, $S = 0$ a.s. conditional on $X = K$ for any $K < t$. Accordingly, $E[(K - X - S)^+] = E[(K - X)^+]$ for $K < t$. Thus, $\pi_K = 0$ for $K < t$. In contrast, for $K \in (t, \mu)$ the risk premium will be strictly positive in general. Thus, the risk premium is not decreasing in K and K does not represent risk appetite. Moreover, since the risk premium vanishes in the limit $K \to \mu$, K does not represent risk aversion either.

The example shows that upwards jumps,[20] in the conditional variance function $\sigma^2(K)$ can threaten the interpretation of K as a measure of risk appetite. Inspecting the factor $R(K)$, we see that the same is true for the density function $g(K)$. An upwards jump in $g(K)$ would change the slope but not the level of $G(K)$ and $P(K)$, thus leading to an upwards jump in the risk premium around K. Beyond this, the dependence of R on the distribution is a bit more complex. To understand it better, we rewrite R in terms of another function L:

Lemma 3. *The function $R(K)$ from (8) can be written as*

$$R(K) = \frac{d}{dK} \log(L(K)) \tag{9}$$

where $L(K)$ is given by

$$L(K) = P(K) + (\mu - K)G(K) = \int_{x_{\min}}^{K} (\mu - k)g(k)dk.$$

Consequently, if L is log-concave in $K \in \mathcal{R}(X)$ then R is decreasing in $K \in \mathcal{R}(X)$.

Proof. Note that $L(K) = P(K) + (\mu - K)G(K)$ is the denominator in (8). Moreover, $\frac{d}{dK}L(K) = (\mu - K)g(K)$ since $\frac{d}{dK}P(K) = G(K)$. This shows (9), R is the derivative of $\log(L)$. Consequently, log-concavity of L implies monotonicity of R. Finally, the integral expression for L follows from $\frac{d}{dK}L(K) = (\mu - K)g(K)$ and $\lim_{K \downarrow x_{\min}} L(K) = 0$. □

Thus, in order to check whether R is decreasing, it suffices to check whether $\log(L)$ is concave over $\mathcal{R}(X)$. A plot of $\log(L)$ can easily be made using the same software that is used to plot the Omega ratio itself, relying on basically the same calculations. If the curve $\log(L)$ deviates clearly from concavity for some risks under consideration, there is reason to be concerned about the interpretation of K as a measure of risk appetite.

The next lemma relates log-concavity of L to a more standard condition, log-concavity of the density function g.

Lemma 4. *If the density function $g(K)$ is log-concave in $K \in \mathcal{R}(X)$, then the function $L(K)$ is log-concave and, accordingly, the function $R(K)$ from (8) is decreasing.*

Proof. Since $\frac{d}{dK}L(K) = (\mu - K)g(K) \geq 0$ for $K \in \mathcal{R}(X) = (x_{\min}, \mu)$, we know that L is increasing over $\mathcal{R}(X)$. Moreover, since $\lim_{K \downarrow x_{\min}} L(K) = 0$, $H(K) = L(K)/L(\mu)$ is a valid cumulative distribution function on $\mathcal{R}(X)$ with density function $h(K) = (\mu - K)g(K)/L(\mu)$. Being the product of the two log-concave functions, g and the linear function $(\mu - \cdot)/L(\mu)$, h is itself log-concave. Moreover, log-concavity is inherited by the cumulative distribution function H and thus by L.[21] □

Thus, log-concavity of the density of X is a sufficient condition for R being decreasing. For instance, when X is normally distributed, the log-density is a concave quadratic polynomial, implying that R is decreasing. In contrast, with heavier than exponential lower tail behavior, this sufficient condition breaks down and there is reason to be concerned that lack of log-concavity can lead to non-monotonic risk premia. In the example of a Student-t distribution below in Section 3.2, we will see a case where R is non-monotonic when X follows a smooth and unimodal distribution that is not log-concave, confirming that, indeed, heavy tails can be a problem for the interpretation of K as a measure of risk appetite.

In the above discussion, we have focused on Proposition 1 as a positive result, providing sufficient conditions for the threshold parameter in the Omega ratio to be aligned with risk appetite. However, arguably, the result is even stronger as a negative result, implying a necessary condition rather than a sufficient condition. This is made precise in the following direct corollary of the proposition.

Corollary 2. *Suppose that the density of an integrable payoff X is Lipschitz continuous and bounded. Moreover, suppose that for some interval $(k_0, k_1) \subseteq \mathcal{R}(X)$ the function $R(K)$ defined in (8) is strictly increasing in $K \in (k_0, k_1)$.[22] Let S be standard normally distributed and independent of X. Thus, (X, S) satisfy Assumption 1 with $\sigma^2(K) \equiv 1$. Then, for all sufficiently small h, the risk premium $\pi_K(h)$ for switching from X to $X + hS$ is increasing in $K \in (k_0, k_1)$.*

The corollary shows that as soon as $R(K)$ is not weakly decreasing at a given point, we can add small independent risks hS to X to generate alternative payoffs $Y = X + hS$ for which the threshold does not represent risk appetite when comparing X and Y.

3.2. Illustrations

If X follows a normal distribution or a Student-t distribution, we can compute R explicitly. The results are summarized in the following lemma:

Lemma 5.

(i) If X is normally distributed with mean μ and standard deviation Σ, then the function R is given by

$$R(K) = \frac{\mu - K}{\Sigma^2} \tag{10}$$

(ii) If X follows a Student-t distribution with $\nu > 1$ degrees of freedom and location and scale parameters[23] μ and Σ, then the function R is given by

$$R(K) = \frac{(\mu - K)(1 - \frac{1}{\nu})}{\Sigma^2 + \frac{1}{\nu}(\mu - K)^2}. \tag{11}$$

Proof. These are tedious but straightforward calculations so we just give the main stepping stones. In both cases, we first compute the

[20] While easy to grasp, discontinuous jumps in $\sigma^2(x)$ or $g(x)$ are, of course, outside the scope of Assumption 1. In order to construct similar examples that satisfy Assumption 1 one could replace the jump by a rapid but smooth local increase.

[21] See [45] for these elementary properties of log-concave functions.

[22] A function $f : \mathbb{R} \to \mathbb{R}$ is called strictly increasing if $x < y$ implies $f(x) < f(y)$ for all $x, y \in \mathbb{R}$.

[23] Often, the Student-t distribution is introduced as a standard t-distribution with only a single parameter ν. Our X is distributed as $\mu + \Sigma Z$ where Z follows a standard t-distribution with ν degrees of freedom. In particular, Σ does not coincide with the standard deviation of X. X only possesses a finite standard deviation for $\nu > 2$. In that case, its standard deviation is given by $\sqrt{\frac{\nu}{\nu-2}}\Sigma$.

function $R_{0,1}(K)$ for the case $\mu = 0$ and $\Sigma = 1$. For the normally distributed case, we find $R_{0,1}(K) = -K$ while for the Student-t case, we find

$$R_{0,1}(K) = \frac{-K(1-\frac{1}{\nu})}{1+\frac{1}{\nu}K^2}. \tag{12}$$

This result can be extended to general μ and Σ using the relation

$$R_{\mu,\Sigma}(K) = \frac{1}{\Sigma} R_{0,1}\left(\frac{K-\mu}{\Sigma}\right). \quad \square$$

For the case of a normal distribution, we find that $R(K)$ is decreasing in K which is expected because the normal distribution has a log-concave density. Note also that, in the admissible case $\mu > K$, the expression (10) (and thus the risk premium) goes to ∞ when Σ goes to zero, capturing that – with an admissible threshold – risk premia for deterministic baselines X are infinite in this case.

Comparing cases (i) and (ii), we recall that in the limit, as the "degrees of freedom" parameter ν goes to ∞, the Student-t distribution converges to a normal and, indeed, (11) converges to (10). However, for finite ν, the function R from (11) behaves quite differently from (10). In particular, as summarized in the first part of the following corollary, the function R is non-monotonic. The second part of the corollary summarizes the consequences

Corollary 3.

(i) For the Student-t distribution, the function R as defined in (11) is increasing in K for $K < \mu - \Sigma\sqrt{\nu}$ and decreasing for $K \in (\mu - \Sigma\sqrt{\nu}, \mu)$.

(ii) In the setting of Proposition 1, suppose that X follows a Student-t distribution with $\nu > 1$ degrees of freedom and location and scale parameters μ and σ. Let S be independent of X so that $\sigma^2(K) = \mathrm{Var}(S|X = K)$ is constant in K. Then, for sufficiently small h, the risk premium from taking on the additional risk hS is increasing in K for $K < \mu - \Sigma\sqrt{\nu}$ and decreasing in K for $K \in (\mu - \Sigma\sqrt{\nu}, \mu)$. In particular, K does not represent risk appetite for $K < \mu - \Sigma\sqrt{\nu}$.

Proof. For (i), like in the proof of Lemma 5, we consider first the case $\mu = 0$, $\Sigma = 1$ and thus $\mathcal{R}(X) = (-\infty, 0)$. By (12), the derivative of $R_{0,1}$ is given by

$$\frac{d}{dK} R_{0,1}(K) = \frac{(K^2-\nu)(\nu-1)}{(K^2+\nu)^2}.$$

Investigating the sign of the derivative, we see that $R_{0,1}$ is increasing in K for $K < -\sqrt{\nu}$ and decreasing for $K \in (-\sqrt{\nu}, 0)$. Now, due to the relation $R_{\mu,\Sigma}(K) = \frac{1}{\Sigma} R_{0,1}(\frac{K-\mu}{\Sigma})$, we find the result of (i) by replacing K with $\frac{K-\mu}{\Sigma}$ in these conditions. (ii) follows directly from combining (i) with Proposition 1. $\square$

Corollary 3 thus shows that when the distribution of X violates log-concavity in the sense of too heavy tails, the connection between the threshold and risk appetite can break down.

In Figs. 1–3, we further illustrate these differences between the normally distributed and the t-distributed case based on simulations from the two examples. This discussion can also be thought of as a model for how one might investigate the behavior of risk premia in empirical applications, relying on the insights from Proposition 1. In the example, X follows either a normal distribution with parameters $\mu = 2$ and $\Sigma = 0.04$ or a Student-t distribution with $\nu = 3$ degrees of freedom and the same μ and Σ. The resulting two possible density functions for X are plotted in Fig. 1, highlighting that the differences between the two cases are visible but not extreme at first sight. The additional risk is hS where $h = 0.02$ and S is standard normal so that $\sigma^2(x) \equiv 1$. The black curves in Fig. 2 are empirical risk premia computed by solving the sample equivalents of $\Omega_K(X) = \Omega_K(X+hS+\pi_K)$ for different K and $m =$

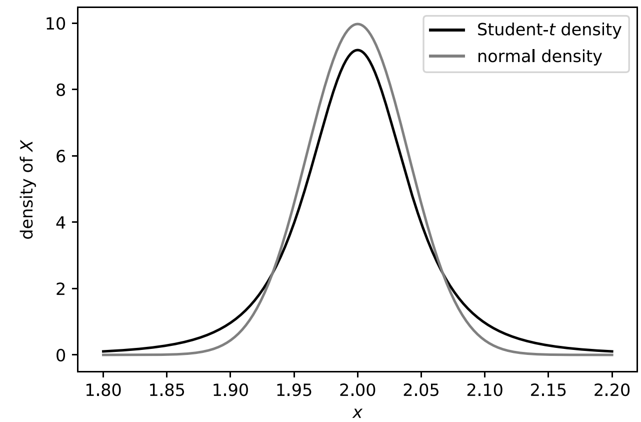

Fig. 1. The two candidate probability density functions for X, a normal distribution (gray curve) or a Student-t distribution X (black curve).

10^6 simulations.[24] The gray curves are our premium approximations $\pi_K \approx R(K)h^2/2$ using the closed-form expressions for $R(K)$ from, respectively, (10) and (11) and the true parameters. Comparing the two curves in the two cases, we see that approximation is fairly accurate and, in particular, captures the monotonicity behavior of π_K rather well. While risk premia in the normal case are decreasing in K, in line with the interpretation of K as risk appetite, this interpretation breaks down in the Student-t case where premia are clearly non-monotonic. Finally, the dashed, black curve is a non-parametric finite difference approximation of the gray curve,

$$R(K) = \frac{d}{dK}\log(L(K)) \approx \frac{\log(\widehat{L}(K+\Delta/2)) - \log(\widehat{L}(K-\Delta/2))}{\Delta} = \widehat{R}(K)$$

where $\widehat{L}(K) = \widehat{P}(K) + (\widehat{\mu} - K)\widehat{G}(K)$ and $\Delta = 0.02$. Here, $\widehat{P}(K)$, $\widehat{G}(K)$ and $\widehat{\mu}$ denote sample equivalents of $P(K)$, $G(K)$ and μ. Computing and plotting $\widehat{R}(K)$ is thus of similar complexity as plotting the empirical Omega ratio. The resulting approximation is close to the other two curves. Overall, the figure highlights that a small change in the distribution of X can have dramatic consequences for the monotonicity behavior of risk premia and, thus, for the interpretation of the threshold parameter K.

In Fig. 3, we illustrate the situation of Fig. 2 in a more traditional plot of the (log)-Omega ratios[25] of X and Y where $Y = X + hS + p$. Here, as the only change, we pick h slightly larger than before, $h = 0.08$, to increase the visual impact of the additional risk. The premium $p = 0.04$ is fixed throughout the plot.[26] Again, the plots are based on 10^6 simulations to reduce the impact of noise. In the normally distributed case, we find a single intersection between the two Omega curves, highlighting that our fixed risk premium is high enough for agents with a high threshold K but not for agents with a low threshold. In contrast, for the case where X follows a t-distribution, we see two intersections and thus a preference reversal that should not be there if K would capture risk appetite in this example.[27]

[24] In particular, we choose large numbers of samples to get a good sense of what is going on. For the normal distribution, we have chosen a somewhat more narrow plot range on the x-axis since we observe fewer extreme observations with this more light-tailed distribution.

[25] The idea of plotting the logarithm of Omega rather than the Omega itself to get clearer pictures goes back already to [12].

[26] The motivation for $p = 0.04$ is as follows: In Fig. 2, a premium of 0.0025 is within the range of observed premia for both distributions, giving in particular two intersections for the Student-t distribution. In light of Proposition 1, increasing h by a factor 4 should give an increase by a factor 4^2 in risk premia. We thus arrive at $4^2 \times 0.0025 = 0.04$.

[27] Even when all distributions are normal, there can be multiple intersections in an Omega-plot, compare e.g. Fig. 2 in [7]. The key distinction is that in the

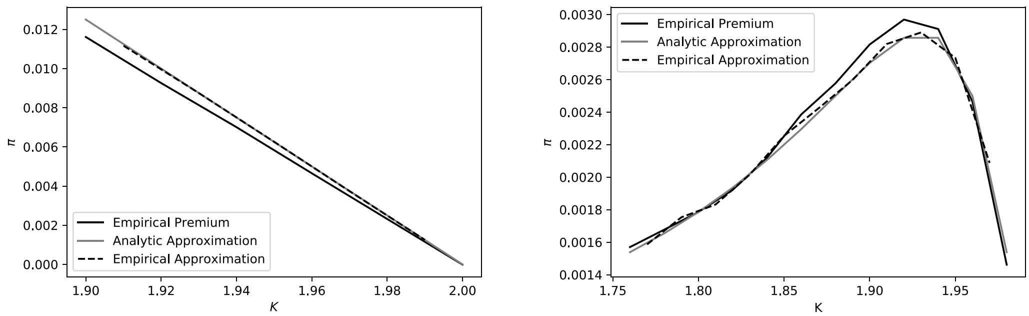

Fig. 2. Risk premia and approximations as a function of the threshold K for normally distributed X (left panel) and Student-t distributed X (right panel).

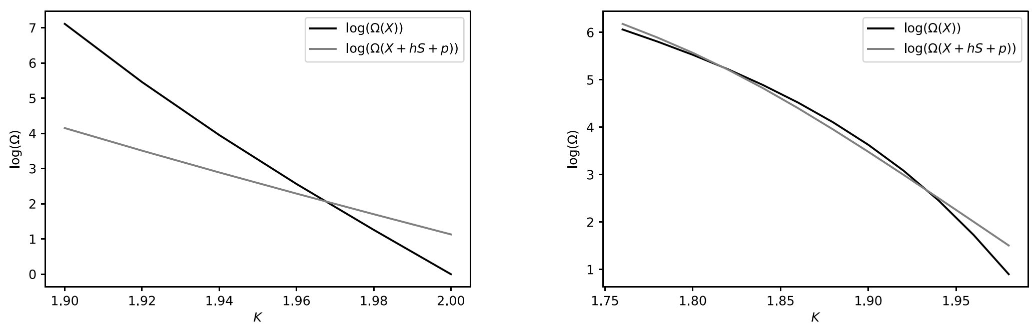

Fig. 3. Logarithmized Omega ratios of X and $X + hS + p$ for normally distributed X (left panel) and Student-t distributed X (right panel).

Some authors such as [10] have proposed the slope of the Omega curve as a measure of risk preferences with a preference for flatter curves indicating higher risk appetite. The Omega plot for the normally distributed case nicely illustrates how this logic is related to thresholds as measures of risk appetite. In this example, it is unambiguous which of the curves is flatter. Moreover, since Omega curves are decreasing, it is the flatter one which dominates to the right of the intersection — and thus at higher thresholds. In contrast, in the Student-t case, it is not even clear which of the two curves should be regarded as the flatter one: At the first intersection, it is the black curve, at the second one it is the gray one.

On a somewhat higher level, we observe that Omega ratios behave quite differently for the two types of distribution. This can be interpreted as a confirmation of the claim[28] that the Omega ratio takes into account all higher moments of the distribution and not just the first two.[29]

3.3. Location-scale families

When the two risks that are being compared come from the same location-scale family, the situation becomes a lot easier than the setting of mean-preserving spreads we have studied so far. In this case, risk premia have distribution-independent closed-form expressions — in the small as well as in the large. Moreover, risk premia are linearly decreasing in K, implying that K captures risk appetite. The fact that everything is easy and well-behaved in the location-scale setting is, of course, closely related to the well-known observation that the choice of performance criterion is largely irrelevant when attention is limited to a single location-scale family, see, e.g., [46]. If the world was two-dimensional in such a way that any comparison of two risks could be reduced to comparing mean and variance, there would be no need for the Omega ratio and many other concepts from decision theory. The following proposition summarizes the situation.

Proposition 2. *Suppose that two payoffs are given by $X = \mu + \sigma_X Z$ and $Y = \mu + \sigma_Y Z$ where Z is an integrable payoff with $E[Z] = 0$ and μ, σ_X and σ_Y are three real numbers with $\sigma_Y > \sigma_X > 0$. Then, for any threshold $K \in \mathbb{R}$ the risk premium π_K that solves $\Omega_K(X) = \Omega_K(Y + \pi_K)$ is given by*

$$\pi_K = (\mu - K)\left(\frac{\sigma_Y}{\sigma_X} - 1\right). \tag{13}$$

Thus K represents risk appetite at (X, Y).

Proof. We need to find π_K such that $\Omega_K(X) = \Omega_K(Y + \pi_K)$ or, equivalently,

$$\frac{\mu - K}{E[(K - \mu - \sigma_X Z)^+]} = \frac{\mu + \pi_K - K}{E[(K - \mu - \pi_K - \sigma_Y Z)^+]} \tag{14}$$

We can rewrite the right hand side into

$$\frac{(\mu + \pi_K - K)\frac{\sigma_X}{\sigma_Y}}{E\left[\left((K - \mu - \pi_K)\frac{\sigma_X}{\sigma_Y} - \sigma_X Z\right)^+\right]},$$

showing that (14) holds if we choose π_K such that

$$(\mu + \pi_K - K)\frac{\sigma_X}{\sigma_Y} = \mu - K.$$

Solving this for π_K gives the result. □

The mathematics behind this result is not new, going back to Theorem 8 in [14], even though the connection to risk premia is not spelled out there. While we only need the case of equal means, the extension to $\mu_X \neq \mu_Y$ is straightforward.

Remark 1. An advantage of the closed-form relation between risk premia and threshold levels in Eq. (13) is that it can easily be inverted to infer threshold levels from risk premia, thus suggesting an

right panel of our Fig. 3 we have multiple intersections within the domain where Omega is consistent with risk-averse behavior, $\Omega_K > 1$.

[28] See, for instance, already [12].

[29] Note though that, strictly speaking the t-distribution with $\nu = 3$ does not possess any finite moments beyond the first two. Note also that, as discussed in Footnote 23, the two distributions do not have exactly the same second moment. However, the picture for the normal distribution would look qualitatively the same if we would rescale that distribution to match the second moment of the Student-t distribution.

method for inferring the level K that matches a decision-maker's risk preferences – provided that those are captured by the Omega. Consider a decision-maker who has to choose between two stochastic returns. Under the first return distribution, the return is 1% or 3% with equal probability while under second distribution the return is π% or $4+\pi$% with equal probability. Which value of π makes a risk-averse decision-maker indifferent between the two lotteries? Clearly, for $\pi \leq 0$, the first lottery has lower risk and a higher return, while for $\pi \geq 1$ the second lottery dominates the first in the sense of first-order stochastic dominance. Thus, there should be an indifference point for some $\pi \in [0,1]$. For example, let us assume that our decision-maker is indifferent for $\pi = 0.5$. This situation fits into the setting of Proposition 2 with $\mu = 2$, $\sigma_X = 1$, $\sigma_Y = 2$ and Z taking values 1 and -1 with equal probability. Solving (13) for the threshold,

$$K = \mu - \frac{\pi}{\frac{\sigma_Y}{\sigma_X} - 1}$$

and plugging in shows that the decision-maker's indifference at $\pi = 0.5$ corresponds to a threshold of $K = 1.5$, i.e., at 1.5%.[30]

One case where the approximation from Proposition 1 and the closed-form solution for location-scale families from Proposition 2 can easily be compared is when everything is normally distributed. Let X be normally distributed with mean μ and standard deviation Σ, let S be standard normal and independent from X and let $Y = X + hS$ for some $h > 0$. Thus, Y is normally distributed with mean μ and standard deviation $\sqrt{\Sigma^2 + h^2}$. Moreover, $\sigma^2(K) \equiv 1$. Thus, using (10) and Proposition 1, our approximation of the risk premium for switching from X to $Y = X + hS$ is given by

$$\pi_K \approx (\mu - K)\frac{h^2}{2\Sigma^2}$$

while the closed-form expression from Proposition 2 is given by

$$\pi_K = (\mu - K)\left(\sqrt{1 + \frac{h^2}{\Sigma^2}} - 1\right).$$

The two approximations coincide up to $\sqrt{1+x^2} - 1 \approx \frac{x^2}{2}$, a second-order Taylor approximation around $x = 0$.

When X is Student-t distributed, there is a stark contrast between the non-monotonicity result for mean-preserving spreads $X + S$ that follows from Proposition 1 and the linearly decreasing expression for location-scale families from Proposition 2. This shows that the intuitive equivalence between adding an additional risk versus scaling up an existing one is "Gaussian thinking" that does not hold up under heavy tails.

4. Comparison with expected utility

So far, we have established sufficient conditions such that Omega-based risk premia for mean-preserving spreads $X + hS$ have the expected monotonicity behavior in the threshold K for small $h > 0$ and, consequently, the parameter K represents risk appetite. A sufficient condition is that X has a log-concave probability density and that the conditional variance σ^2 is weakly decreasing.

The goal of this section is to compare these sufficient conditions to the conditions that are needed to ensure that risk premia based on expected utility have the intended monotonicity behavior in the risk aversion parameters of classical families of utility functions. Is the threshold parameter K a more or a less robust representer of risk appetite or its converse, risk aversion, than, e.g. the parameter of a CRRA utility function? In particular, how does a plot of Omega ratios as a function of K compare to a more traditional plot of risk premia based on CRRA or CARA utility functions against the respective risk aversion parameters of those functions?

For a smooth, strictly increasing and strictly concave function u, we define the utility-based risk premium $\pi_u(h)$ via

$$E[u(X)] = E[u(X + hS + \pi_u(h))]. \tag{15}$$

As before, we assume $E[S|X] = 0$ and define $\sigma^2(x) = \mathrm{Var}(S|X = x)$. For the utility function u, we are especially interested in the cases where it is either a constant absolute risk aversion (CARA) utility function with risk aversion parameter $\alpha > 0$ or a constant relative risk aversion (CRRA) utility function with risk aversion parameter $\gamma > 0$[31]

$$u(x) = -\exp(-\alpha x) \quad \text{or} \quad u(x) = \frac{x^{1-\gamma} - 1}{1 - \gamma}.$$

Our main question in this section is now as follows: Under which conditions are risk premia (for small $h > 0$) increasing in the respective risk aversion parameters α and γ of the CRRA and CARA classes of utility functions? In principle, this is a classical question from the literature on background risk [16]. However, most of that literature has assumed that X and S have a compact support, thus abstracting away from possible regularity problems and the technical conditions that rule them out. In Appendix C we thus present concise statements of the relevant results, while keeping the discussion in this section somewhat less technical.

4.1. Scope

Let us first discuss the scope of the various criteria: Under which conditions are the base quantities $E[u(X)]$ and $\Omega_K(X)$ mathematically well-defined?[32] For the Omega ratio $\Omega_K(X)$ to be well-defined, the main necessary condition is integrability, $E[|X|]$ needs to be finite. This is a relatively mild condition that treats the upper and lower tails symmetrically. In contrast, utility functions like the CARA and CRRA can handle the St. Petersburg paradox, i.e., they still give a finite expected utility for risks with a very heavy upper tail which possibly violates integrability. For example, a risk X with the (power law) cumulative distribution function $F(x) = 1 - x^{-1}$ on $[1, \infty)$ does not have a finite mean or a finite Omega ratio but it does have finite expected utility $E[u(X)]$ for any CARA or CRRA utility function.[33]

When it comes to heavy lower tails, expected utility is far more limited in scope. For CRRA utility to be well-defined for any $\gamma > 0$, we need to assume that X is non-negative and that $E[X^{1-\gamma}] < \infty$ for all $\gamma > 0$. For CARA utility, we are more flexible and can allow for negative realizations of X as long as $E[\exp(-\alpha X)]$ is finite for all $\alpha > 0$. Thus, CARA utility is well-defined for normally distributed X, but cannot handle heavier than exponential tails in the density of X. For instance, when X follows a Student-t distribution, CARA utility is negative infinity.

We thus see that expected-utility criteria tend to have a broader scope than the Omega regarding the heaviness of the upper tail of the distribution of X but a narrower scope when it comes to the lower tail. By using the Omega rather than CARA or CRRA utility, we can no longer resolve the St. Petersburg paradox but need not worry as much about distributions with heavy lower tails.

---

[30] We leave it to further research to study this elicitation method in greater detail, investigating among others its practical performance, its performance beyond location-scale families and the portability of the resulting threshold estimates from one setting to the next.

[31] In the CRRA case, we assume, as usual, that $u(x) = \log(x)$ for $\gamma = 1$ and that the support of X and $X + hS$ is contained in $\mathbb{R}^+$.

[32] Of course, as discussed further below, we need to make additional assumptions to ensure that risk premia are finite and satisfy our asymptotic expansions, see Assumption 1 for the Omega ratio and Assumption 2 in Appendix C for expected utility.

[33] Some authors, starting with Karl Menger in the 1930s, have argued that boundedness from above of the utility function is needed for a genuine resolution of the paradox. This is satisfied by CRRA utility with $\gamma > 1$ and by CARA for all $\alpha \geq 0$. See [47] for details, references and more discussion of expected utility and the St. Petersburg Paradox.

4.2. Asymptotic expansions

We now look into asymptotic expansions for the utility-based risk premia from (15), providing the counterparts of our Proposition 1. Classical arguments from the literature on background risk [16] show that, provided all expected values are well-defined, $\pi_u(0) = \pi_u'(0) = 0$ and thus

$$\pi_u(h) = \pi_u''(0)\frac{h^2}{2} + O(h^3)$$

where

$$\pi_u''(0) = -\frac{E[\sigma^2(X)u''(X)]}{E[u'(X)]}. \tag{16}$$

We refer to Proposition 4 in the appendix for a precise statement of the technical conditions behind this and to Lemmas 6 and 7 for their specialization to, respectively, CARA and CRRA utility.

Absence of background risk: If there is no background risk, i.e., if $X = x$ almost surely, (16) simplifies to

$$\pi_u''(0) = -\sigma^2(x)\frac{u''(x)}{u'(x)}. \tag{17}$$

This special case goes back already to Pratt [15]. In particular, unlike for the Omega ratio, utility-based risk premia can easily be defined in the absence of background risk.

Here, if two utility functions u and v can be ranked in terms of their absolute risk aversion functions, $-u''(x)/u'(x) \geq -v''(x)/v'(x)$, condition (17) immediately implies a ranking of the risk premia $\pi_u \geq \pi_v$. In particular, such a ranking holds when u and v are both CARA or both CRRA functions with different values of the risk aversion parameters α or γ since, respectively, $-u''(x)/u'(x) = \alpha$ in the CARA case and $-u''(x)/u'(x) = \gamma/x$, $x > 0$, in the CRRA case. Thus, in the absence of background risk, the parameters α and γ represent risk aversion within their respective families of utility functions.

Comparing risk premia under background risk: Ross [19] and Kihlstrom et al. [20] showed that in general when X is stochastic a ranking of the absolute risk aversion functions is not sufficient for guaranteeing a ranking of π_u and π_v. One main observation of [19] is the following[34]:

Fact 1. *For an interval $W \subseteq \mathbb{R}$, let $u : W \to \mathbb{R}$ and $v : W \to \mathbb{R}$ be two strictly increasing utility functions and let X be a random variable with values in W such that $E[u'(X)]$, $E[v'(X)]$, $E[\sigma^2(X)u''(X)]$ and $E[\sigma^2(X)v''(X)]$ are well-defined. Then, the condition*

$$-u''(x)/u'(y) \geq -v''(x)/v'(y) \tag{18}$$

for all $x, y \in W$ implies that

$$-\frac{E[\sigma^2(X)u''(X)]}{E[u'(X)]} \geq -\frac{E[\sigma^2(X)v''(X)]}{E[v'(X)]} \tag{19}$$

and thus $\pi_u''(0) \geq \pi_v''(0)$ in the setting of (16) regardless of the shape of $\sigma^2(X)$ and the distribution of X.

Proof. Note first that by strict monotonicity of u and v, $E[u'(X)]$ and $E[v'(X)]$ are positive. Introduce an independent copy Y of X and replace X by Y in the denominators of (16), $E[u'(X)] = E[u'(Y)]$ and $E[v'(X)] = E[v'(Y)]$. Then, using independence, we can rewrite $\pi_u''(0) \geq \pi_v''(0)$ into

$$E[\sigma^2(X)(u''(X)v'(Y) - v''(X)u'(Y))] \leq 0.$$

Condition (18) implies that expression inside the expected value is always negative, thus giving the result. □

[34] See also [48].

Ross also showed that condition (18) is in a sense tight. Unfortunately, (18) is a fairly rigid condition that is generally violated even when u and v are two CARA or two CRRA functions. Thus, additional assumptions are needed to guarantee that risk premia are monotonic in the parameters of those families of utility functions. The following example of CARA utility in connection with an increasing function σ^2 makes this claim concrete. The logic behind this counterexample is the same as in Example 2 for the Omega ratio, showing that some problems are shared between the different criteria.

Example 3. Suppose that X is exponentially distributed with parameter λ and that u is a CARA utility function with parameter α. Moreover, let $\sigma^2(x) = \bar{\sigma}^2 1_{\{x>t\}}$ for some positive constants $\bar{\sigma}^2$ and t, i.e., noise is added only to sufficiently high realizations of X. In that case,

$$\pi''(0) = -\frac{E[\sigma^2(X)u''(X)]}{E[u'(X)]} = \bar{\sigma}^2\alpha\frac{E[1_{X>t}\exp(-\alpha X)]}{E[\exp(-\alpha X)]} = \bar{\sigma}^2\alpha\exp(-t(\lambda+\alpha))$$

which is *decreasing* in α for sufficiently large α except in the boundary case $t = 0$ where $\sigma^2(x)$ is constant.

In contrast to the case of general $\sigma^2(x)$ treated in [19], the results of Kihlstrom et al. [20] apply to the constant case $\sigma^2(x) \equiv \bar{\sigma}^2$. In that case, (16) becomes[35]

$$\pi_u''(0) = -\bar{\sigma}^2\frac{E[u''(X)]}{E[u'(X)]}. \tag{20}$$

The main sufficient condition of [20] can be summarized as follows:

Fact 2. *For an interval $W \subseteq \mathbb{R}$, let $u : W \to \mathbb{R}$ and $v : W \to \mathbb{R}$ be two strictly increasing utility functions and let X be a random variable with values in W such that $E[u'(X)]$, $E[v'(X)]$, $E[u''(X)]$ and $E[v''(X)]$ are well-defined. Then, the conditions that (i) $-u''(x)/u'(x) \geq -v''(x)/v'(x)$ for all $x \in W$ and that (ii) at least one of the functions $-u''/u'$ or $-v''/v'$ is (weakly) decreasing imply*

$$-\frac{E[u''(X)]}{E[u'(X)]} \geq -\frac{E[v''(X)]}{E[v'(X)]} \tag{21}$$

and thus $\pi_u''(0) \geq \pi_v''(0)$ in the setting of (20).

The proof of this fact is a special case of the proof of Proposition 3 below. Conditions (i) and (ii) are satisfied e.g. when u and v are both CRRA utility functions or both CARA utility functions. Thus, for the case where σ^2 is constant, the results of Kihlstrom et al. [20] imply that $\pi_u''(0)$ must be increasing in α or, respectively, γ. For example, for the CARA case, we find $\pi''(0) = \alpha\mathrm{Var}(S)$ which indeed increases in α.

4.3. A third fact

Comparing Facts 1 and 2 and their implications for CARA and CRRA parameters as measures of risk aversion, we thus have a negative result from [19] for general σ^2, and a positive result from [20] for constant σ^2. In Examples 2 and 3, we saw that problems with increasing functions σ^2 appear both for Omega-based risk premia and for utility-based ones. It is thus natural to ask whether the positive result for decreasing functions σ^2 under the Omega ratio from Corollary 1 also holds under expected utility.

In Proposition 3, we show that the main result of Kihlstrom et al. [20] is indeed preserved if σ^2 is decreasing rather than constant. Moreover, a sufficiently strong decrease in σ^2 can counteract increases in $-u''/u'$ and vice versa. What matters is that the product of the two functions is decreasing. The proposition is an extension of their Theorem (p. 916) with the proof adapted from their line of argument.

[35] The setting actually studied by Kihlstrom et al. [20] is a slightly different one: They do not impose $E[S] = 0$ but make the stronger assumption of independence between X and S. Yet it is easy to see that their analysis also applies in our setting.

Proposition 3. *For an interval $W \subseteq \mathbb{R}$, let $u : W \to \mathbb{R}$ and $v : W \to \mathbb{R}$ be two twice continuously differentiable, strictly increasing, strictly concave functions. Let X be a random variable with values in W. Let $\sigma^2 : W \to \mathbb{R}$ be a non-negative, measurable function. Assume that (i) $-u''(x)/u'(x) \geq -v''(x)/v'(x)$ for all $x \in W$ and that (ii) at least one of the two functions $S_u(x) = -\sigma^2(x)u''(x)/u'(x)$ and $S_v(x) = -\sigma^2(x)v''(x)/v'(x)$ is weakly decreasing over W. Then,*

$$\frac{E[\sigma^2(X)u''(X)]}{E[u'(X)]} \geq \frac{E[\sigma^2(X)v''(X)]}{E[v'(X)]}. \tag{22}$$

provided all expected values in (22) exist.

The proof is deferred to the appendix. By (16), it follows that for CARA and for CRRA utility functions, the parameters capture the degree of risk aversion as long as σ^2 is weakly decreasing because then condition (ii) in Proposition 3 is satisfied. More precisely, since $-u''/u'$ is constant for CARA utility, the condition on σ^2 in that case is that $\sigma^2(x)$ is weakly decreasing. For CRRA utility, we have $-u''/u' = \gamma/x$ and thus only need to impose the weaker condition that $\sigma^2(x)/x$ is weakly decreasing. No additional conditions on the distribution of X like the log-concavity in the Omega case is needed. However, as discussed in the beginning of this section, differences in the scopes of the various criteria need to be taken into account.[36] For example, when X has a heavy lower tail, like for a Student-t distribution, only the Omega-based risk premium is well-defined — even though the interpretation of the threshold in terms of risk appetite may no longer be valid.

5. Conclusion

By and large, we conclude that the situations where risk premia are well-defined and the parameter properly captures risk appetite (or, conversely, risk aversion) are fairly similar for the Omega ratio and expected utility with an exponential utility function. An advantage of the Omega ratio is that it remains well-defined in situation with a heavy lower tail – even though the interpretation of the threshold in terms of risk appetite may break down. Conversely, a considerable advantage of expected utility is that its risk premia remain well-defined in the important boundary case of vanishing background risk, and in situations with a non-integrable upper tail as in the St. Petersburg paradox. One might argue that these differences partly reflect the origins of the two theories. Arguably, expected utility was developed by Daniel Bernoulli as a response to the St. Petersburg paradox, see, e.g., [47]. In contrast, the Omega ratio was developed as a decision criterion in a world where lower tails can be heavy, where background risk never vanishes and where St. Petersburg gambles are rarely on offer.

Nevertheless, besides these encouraging observations, we have also seen that sometimes the threshold in the Omega ratio can be disconnected from its intended interpretation as a measure of risk appetite. Our paper thus stands in the line of papers surveyed in [8], showing that some caution is needed when working with the Omega ratio.

Declaration of competing interest

Herby, Anne G. Balter, Ki Wai Chau and Nikolaus Schweizer certify that they have NO affiliations with or involvement in any organization or entity with any financial interest (such as honoraria; educational grants; participation in speakers bureaus; membership, employment, consultancies, stock ownership, or other equity interest; and expert testimony or patent-licensing arrangements), or non- financial interest (such as personal or professional relationships, affiliations, knowledge or beliefs) in the subject matter or materials discussed in this manuscript.

[36] For details, see also the precise statements formulated and proved in Appendix C.

Data availability

No data was used for the research described in the article.

Appendix A. Technical proofs

A.1. Proof of Lemma 2

A direct calculation shows that

$$l_{X,S}(x,s) = g(x)h\left(\frac{s}{\Sigma(x)}\right)\frac{1}{\Sigma(x)}.$$

Noting that we can write $\varphi(z)$ as $\varphi(z) = \min(1, |z|^{-\kappa})$ and using that $h(z) \leq C\varphi(z)$, we can bound $l_{X,S}$ from above by

$$l_{X,S}(x,s) \leq Cg(x)\varphi\left(\frac{s}{\Sigma(x)}\right)\frac{1}{\Sigma(x)} = Cg(x)\min\left(1, |s|^{-\kappa}\Sigma(x)^{\kappa}\right)\frac{1}{\Sigma(x)}$$

Next, using the elementary inequality $\min(1, ab) \leq \min(1, a)\max(1, b)$ for non-negative real numbers a and b, we obtain

$$l_{X,S}(x,s) \leq Cg(x)\max\left(\frac{1}{\Sigma(x)}, \Sigma(x)^{\kappa-1}\right)\varphi(s).$$

This shows the boundedness part of the global convergence condition. Integrability now also follows, using that $E[\Sigma(X)^{-1}]$ and $E[\Sigma(X)^{\kappa-1}]$ are finite. For the local continuity, denote by C a generic, positive constant whose value may change from line to line and note that

$$\begin{aligned}
&|l_{X,S}(x,s) - l_{X,S}(y,s)| \\
&= \left| g(x)h\left(\frac{s}{\Sigma(x)}\right)\frac{1}{\Sigma(x)} - g(y)h\left(\frac{s}{\Sigma(y)}\right)\frac{1}{\Sigma(y)} \right| \\
&\leq \left| g(x)h\left(\frac{s}{\Sigma(x)}\right)\frac{1}{\Sigma(x)} - g(y)h\left(\frac{s}{\Sigma(x)}\right)\frac{1}{\Sigma(x)} \right| \\
&\quad + \left| g(y)h\left(\frac{s}{\Sigma(x)}\right)\frac{1}{\Sigma(x)} - g(y)h\left(\frac{s}{\Sigma(y)}\right)\frac{1}{\Sigma(x)} \right| \\
&\quad + \left| g(y)h\left(\frac{s}{\Sigma(y)}\right)\frac{1}{\Sigma(x)} - g(y)h\left(\frac{s}{\Sigma(y)}\right)\frac{1}{\Sigma(y)} \right| \\
&\leq C\left(|g(x) - g(y)| + \left| h\left(\frac{s}{\Sigma(x)}\right) - h\left(\frac{s}{\Sigma(y)}\right) \right| + \left| \frac{1}{\Sigma(x)} - \frac{1}{\Sigma(y)} \right| \right) \\
&\leq C\left(|x-y| + (|s|+1)\left| \frac{1}{\Sigma(x)} - \frac{1}{\Sigma(y)} \right| \right) \leq C\max(1,|s|)|x-y|.
\end{aligned}$$

A.2. Proof of Proposition 1

Throughout this proof, C denotes a generic, positive constant whose value may change from line to line. The constant C is always independent of ϵ and h but may depend on some of the other parameters. Formally applying a second order Taylor expansion to the function π_K around 0 for small enough h, we obtain

$$\pi_K(h) = \pi_K(0) + \pi_K'(0)h + \pi_K''(0)\frac{h^2}{2} + O(h^3).$$

To prove the proposition, we thus need to justify the Taylor expansion and make sure that the coefficients match those given in the proposition. After acknowledging the obvious fact that $\pi_K(0) = 0$, the remainder of the proof thus consists of proving the following three points:

1. π_K is twice differentiable at 0,
2. $\pi_K'(0) = 0$, and
3. $\pi_K''(0) = \sigma^2(K)R(K)$.

To begin, we spell out the defining Eq. (2) of the risk premium $\pi_K(h)$

$$1 + \frac{\mu - K}{E[(K-X)^+]} = 1 + \frac{\mu + \pi_K(h) - K}{E[(K - X - hS - \pi_K(h))^+]}, \tag{23}$$

and rearrange it into

$$(\mu - K)E[(K - X - hS - \pi_K(h))^+] = (\mu + \pi_K(h) - K)E[(K - X)^+]. \quad (24)$$

Note that the condition $K \in \mathcal{R}(X)$ and Lemma 1 ensure that both denominators in (23) are non-zero so that (24) is meaningful. Thus, we can establish that $\pi'(h)$ exists and is bounded by the implicit function theorem.

Taking a derivative with respect to h on both sides of (24) yields

$$(K - \mu)E[1_{\{K-X-hS-\pi_K(h)>0\}}(S + \pi'_K(h))] = \pi'_K(h)E[(K - X)^+]. \quad (25)$$

In the limiting case $h = 0$, the expected value on the left hand side becomes

$$E[S1_{\{K-X>0\}}] + \text{Prob}(\{K - X > 0\})\pi'_K(0) = \text{Prob}(\{K - X > 0\})\pi'_K(0),$$

where we use that $E[S|X] = 0$. Evaluating (25) at $h = 0$ thus gives

$$0 = \pi'_K(0)\left\{E[(K - X)^+] + (\mu - K)\text{Prob}(\{K - X > 0\})\right\}. \quad (26)$$

By our assumption that $K \in \mathcal{R}(X)$, $E[(K - X)^+]$ is strictly positive and $\mu \geq K$. On the right hand side of (26), we thus have the product of $\pi'_K(0)$ and a strictly positive number. It follows that $\pi'_K(0) = 0$. This concludes the proof for point 2 from our list.

For the next steps of the proof, we need to take one more derivative with respect to h. To this end, we rewrite (25) as

$$(K - \mu)f_1(h) + (K - \mu)f_2(h)\pi'_K(h) = \pi'_K(h)E[(K - X)^+], \quad (27)$$

where

$$f_1(h) = E[S1_{\{K-X-hS-\pi_K(h)>0\}}] \text{ and } f_2(h) = \text{Prob}(\{K - X - hS - \pi_K(h) > 0\}).$$

It is clear from (27) and the fact that the factor in curly brackets from (26) is positive, that π'_K is differentiable if f_1 and f_2 are differentiable, implying point 1 on our list. Moreover, if we can compute the derivatives of f_1 and f_2, we can directly calculate $\pi''_K(0)$ and arrive at point 3. In the remainder of this proof, we thus verify that f_1 and f_2 are differentiable and calculate their derivatives.

We first focus on f_1. For $\epsilon \in (0, 1)$, denote by Φ_ϵ the cumulative distribution function of a normal distribution with mean 0 and variance ϵ^2 and by ϕ_ϵ the associated density function. In the limit $\epsilon \to 0$, the function

$$f_{1,\epsilon}(h) = E[S\Phi_\epsilon(K - X - hS - \pi(h))]$$

converges to f_1 because $\Phi_\epsilon(\cdot)$ converges to $1_{\{\cdot>0\}}$ almost everywhere and the expression inside the expectation is bounded by the integrable random variable $|S|$ and since the boundary event $\{K - X - hS - \pi(h) = 0\}$ has probability 0 due to the assumption of continuous distributions. Integrability of $|S|$ follows from Assumption 1. The goal is to derive the derivative of f_1 by studying the derivative of $f_{1,\epsilon}$, exploiting that $\Phi_\epsilon(\cdot)$ is smooth even though $1_{\{\cdot>0\}}$ is not. To establish the differentiability of $f_{1,\epsilon}$, we first note the differentiability of the function within the expectation and compute its derivative,

$$\begin{aligned}\frac{d}{dh}s\Phi_\epsilon(K - x - hs - \pi(h)) =& s\frac{1}{\epsilon\sqrt{2\pi}}e^{-\frac{(K-x-hs-\pi(h))^2}{2\epsilon^2}}(-s - \pi'(h))\\ =& -s^2\phi_\epsilon(K - x - hs - \pi(h))\\ &- s\pi'(h)\phi_\epsilon(K - x - hs - \pi(h)).\end{aligned}$$

For any fixed $\epsilon > 0$, the absolute value of the derivative is bounded by

$$\left|\frac{d}{dh}s\Phi_\epsilon(K - x - hs - \pi(h))\right| \leq \frac{1}{\epsilon\sqrt{2\pi}}s^2 + |\pi'(h)|\frac{1}{\epsilon\sqrt{2\pi}}|s|. \quad (28)$$

Since $\pi'(h)$ is bounded, the right hand side of (28) is integrable since Assumption 1 implies the integrability of $|S|$ and S^2. Taking the derivative of $f_{1,\epsilon}$ with respect to h now gives

$$f'_{1,\epsilon}(h) = -E[S^2\phi_\epsilon(K - X - hS - \pi(h))] - \pi'(h)E[S\phi_\epsilon(K - X - hS - \pi(h))].$$

To complete our study of $f_{1,\epsilon}$, it thus suffices to study the expressions $E[S^k\phi_\epsilon(K - X - hS - \pi(h))]$, for $k = 0, 1, 2$.[37] We thus write

$$\begin{aligned}E[S^k\phi_\epsilon(K - X - hS - \pi(h))] &= \int_{\mathbb{R}}\int_{\mathbb{R}} s^k l_{X,S}(x, s)\frac{1}{\epsilon\sqrt{2\pi}}e^{-\frac{(K-x-hs-\pi(h))^2}{2\epsilon^2}}dsdx\\ &= \int_{\mathbb{R}}\int_{\mathbb{R}} s^k l_{X,S}(K - \epsilon y - hs - \pi(h), s)\phi(y)dsdy\end{aligned} \quad (29)$$

where $\phi \equiv \phi_1$ is the standard normal density and note that by Assumption 1

$$\begin{aligned}&\int_{\mathbb{R}}\int_{\mathbb{R}}\left|s^k l_{X,S}(K - \epsilon y - hs - \pi(h), s)\phi(y)\right|dsdy\\ \leq& C\int_{\mathbb{R}}\int_{\mathbb{R}}|s|^k b(K - \epsilon y - hs - \pi(h))\varphi(s)\phi(y)dsdy\\ \leq& C\int_{\mathbb{R}}\int_{\mathbb{R}}|s|^k\varphi(s)\phi(y)dsdy < \infty\end{aligned}$$

where the last step uses that b is bounded. We can therefore apply the dominated convergence theorem for the limit $\epsilon \to 0$. It follows that

$$\begin{aligned}&\lim_{\epsilon\to0}E[S^k\phi_\epsilon(K - X - hS - \pi(h))]\\ =& \int_{\mathbb{R}}\int_{\mathbb{R}} s^k\lim_{\epsilon\to0}l_{X,S}(K - \epsilon y - hs - \pi(h), s)\phi(y)dsdy\\ =& \int_{\mathbb{R}}\int_{\mathbb{R}} s^k l_{X,S}(K - hs - \pi(h), s)\phi(y)dsdy\\ =& \int_{\mathbb{R}} s^k l_{X,S}(K - hs - \pi(h), s)ds.\end{aligned}$$

To finish this proof, we need to show that the functions $E[S^k\phi_\epsilon(K - X - hS - \pi(h))]$ converge uniformly with respect to h when ϵ goes to 0. This is necessary to establish that the limit of the sequence is the derivative of the target function. For any given ϵ, we define $M = \epsilon^{-\frac{1}{2(k+2)}}$ and note that $M > 1$ by our assumption that $\epsilon \in (0, 1)$. We have

$$\begin{aligned}&\left|E[S^k\phi_\epsilon(K - X - hS - \pi(h))] - \int_{\mathbb{R}} s^k l_{X,S}(K - hs - \pi(h), s)ds\right|\\ \overset{(29)}{=}& \left|\int_{\mathbb{R}}\int_{\mathbb{R}} s^k(l_{X,S}(K - \epsilon y - hs - \pi(h), s) - l_{X,S}(K - hs - \pi(h), s))\phi(y)dsdy\right|\\ \leq& \left|\int_{\mathbb{R}}\int_{[-1,1]} s^k(l_{X,S}(K - \epsilon y - hs - \pi(h), s) - l_{X,S}(K - hs - \pi(h), s))\phi(y)dsdy\right|\\ &+ \left|\int_{\mathbb{R}}\int_{[-M,M]\setminus[-1,1]} s^k(l_{X,S}(K - \epsilon y - hs - \pi(h), s) - l_{X,S}(K - hs - \pi(h), s))\phi(y)dsdy\right|\\ &+ \left|\int_{\mathbb{R}}\int_{\mathbb{R}\setminus[-M,M]} s^k(l_{X,S}(K - \epsilon y - hs - \pi(h), s) - l_{X,S}(K - hs - \pi(h), s))\phi(y)dsdy\right|\\ \leq& \int_{\mathbb{R}}\int_{[-1,1]}|s|^k C|\epsilon y|\frac{1}{\sqrt{2\pi}}e^{-\frac{y^2}{2}}dsdy + \int_{\mathbb{R}}\int_{[-M,M]\setminus[-1,1]}|s|^{k+1}C|\epsilon y|\phi(y)dsdy\\ &+ \int_{\mathbb{R}}\int_{\mathbb{R}\setminus[-M,M]}|s|^k C(b(K - \epsilon y - hs - \pi(h)) + b(K - hs - \pi(h)))\frac{1}{|s|^\kappa}\phi(y)dsdy\\ \leq& C\epsilon\int_{\mathbb{R}}|y|\phi(y)dy + C\epsilon M^{k+2}\int_{\mathbb{R}}|y|\phi(y)dy + CM^{k+1-\kappa}\\ \leq& C(\epsilon + \epsilon^{\frac{1}{2}} + \epsilon^{\frac{\kappa-k-1}{2(k+2)}}).\end{aligned}$$

Since the last term is independent of h and converges to 0 as ϵ tends to 0, we have established the uniform convergence of $E[S^k\phi_\epsilon(K - X - hS - \pi(h))]$. Combining this result with the fact that $\pi'(h)$ is bounded, $f'_{1,\epsilon}$ also converges uniformly with respect to h. We can conclude that

$$\begin{aligned}f'_1(h) =& \lim_{\epsilon\to0}f'_{1,\epsilon}(h)\\ =& -\int_{\mathbb{R}} s^2 l_{X,S}(K - hs - \pi(h), s)ds - \pi'(h)\int_{\mathbb{R}} s l_{X,S}(K - hs - \pi(h), s)ds.\end{aligned}$$

Specifically for $h = 0$, we find that

$$f'_1(0) = -\int_{\mathbb{R}} s^2 l_{X,S}(K, s)ds - \pi'(0)\int_{\mathbb{R}} s l_{X,S}(K, s)ds = -\text{Var}(S|X = K)g(K)$$

[37] The case $k = 0$ is not needed right now. It is included because it appears later when we analyze $f_{2,\epsilon}$ in an analogous fashion.

where we use that $E[S|X] = 0$ and that $l_{X,S}(K,s)$ is the product of $g(K)$ and of the density of S conditional on $X = K$. Applying the same line of reasoning to the second function f_2 from (27), we find that

$$f_2'(h) = -\int_{\mathbb{R}} s l_{X,S}(K - hs - \pi(h), s)ds - \pi'(h)\int_{\mathbb{R}} l_{X,S}(K - hs - \pi(h), s)ds,$$

and

$$f_2'(0) = -\int_{\mathbb{R}} s l_{X,S}(K, s)ds - \pi'(0)\int_{\mathbb{R}} l_{X,S}(K, s)ds = 0.$$

Finally, we can now take one more derivative on both sides of (27) and plug in $h = 0$ to obtain the condition

$$(K-\mu)f_1'(0) + (K-\mu)f_2'(0)\pi_K'(0) + (K-\mu)f_2(0)\pi_K''(0) = \pi_K''(0)E[(K-X)^+]$$

which only depends on $\pi_K''(0)$, $f_2(0) = G(K)$ and terms we have already computed. Solving for $\pi_K''(0)$, we find that

$$\pi_K''(0) = \frac{(\mu - K)\mathrm{Var}(S|X = K)g(K)}{E[(K-X)^+] + (\mu - K)G(K)}.$$

which is well-defined because the denominator is again the factor in curly brackets from (26). This concludes the proof of point 3.

A.3. Proof of Proposition 3

The main intermediate step in the proof is to show that for all $x_1, x_2 \in W$ with $x_1 \geq x_2$ we have that

$$r\left\{(S_u(x_1) - S_v(x_1)) + (S_u(x_2) - S_v(x_2))\right\} + (1-r)(S_u(x_2) - S_v(x_1)) \geq 0 \quad (30)$$

where $r = (u'(x_1)/v'(x_1))/(u'(x_2)/v'(x_2))$. We first prove (30) and then show how to complete the proof from there. Note that the ratio $u'(x)/v'(x)$ is weakly decreasing in x since the associated first order condition follows from (i). Thus, $x_1 \geq x_2$ and the fact that u' and v' are positive imply $r \in [0,1]$. From (i), it follows that $S_u(x) \geq S_v(x)$ for all $x \in W$ so the term in curly brackets is non-negative. To show that $S_u(x_2) - S_v(x_1) \geq 0$, consider first the case in (ii) that S_u is weakly decreasing. In that case, we can write

$$S_u(x_2) \geq S_u(x_1) \geq S_v(x_1)$$

where the second step uses (i). In the other case, S_v is weakly decreasing so

$$S_u(x_2) \geq S_v(x_2) \geq S_v(x_1).$$

This shows (30). To derive the claim of the proposition from (30), we first rewrite (30)

$$\begin{aligned} &-u''(x_1)\sigma^2(x_1)v'(x_2) - u''(x_2)\sigma^2(x_2)v'(x_1) \\ &\geq -v''(x_1)\sigma^2(x_1)u'(x_2) - v''(x_2)\sigma^2(x_2)u'(x_1). \end{aligned} \quad (31)$$

Since each side of (31) is symmetric in x_1 and x_2, (31) must hold for all $x_1, x_2 \in W$ if it holds for $x_1 \geq x_2$. We can thus drop the assumption that $x_1 \geq x_2$. Now, consider two mutually independent copies X_1 and X_2 of X. Plugging in X_1 for x_1, X_2 for x_2, taking expectations and using independence turns (31) into

$$\begin{aligned} &-E[u''(X_1)\sigma^2(X_1)]E[v'(X_2)] - E[u''(X_2)\sigma^2(X_2)]E[v'(X_1)] \\ &\geq -E[v''(X_1)\sigma^2(X_1)]E[u'(X_2)] - E[v''(X_2)\sigma^2(X_2)]E[u'(X_1)] \end{aligned}$$

which implies

$$-2E[u''(X)\sigma^2(X)]E[v'(X)] \geq -2E[v''(X)\sigma^2(X)]E[u'(X)]$$

and thus (22).

Appendix B. Omegas, odds and expectiles

In this appendix, we informally recall how Omega curves and expectiles can be written as odds curves and quantile functions of the same transformed distribution, mostly translating observations from [36,41] into our notation. To this end, let X be a real-valued random payoff as introduced in Section 2 with cumulative distribution function G. For simplicity, assume that G is strictly increasing and thus invertible. Recalling the two decision criteria discussed in [28], we can associate to X a success probability $1 - G(K)$ at level $K \in \mathbb{R}$, and, for $\alpha \in (0,1)$, the α-quantile $q_\alpha(X)$ which is the solution to $G(q_\alpha(X)) = \alpha$. The odds at level K are given by

$$\frac{1 - G(K)}{G(K)}.$$

Since the odds are a monotonically increasing transformation of the success probabilities, $x \mapsto \frac{x}{1-x}$, it does not matter whether success probabilities or odds are used as a decision criterion. Now, using the notation $C_K(X) = E[(X-K)^+]$ and $P_K(X) = E[(K-X)^+]$, define a transformed distribution function[38]

$$F(K) = \frac{P_K(X)}{C_K(X) + P_K(X)}$$

The odds associated with the transformed distribution F coincide with the Omega ratios of the original distribution G

$$\frac{1 - F(K)}{F(K)} = \frac{C_K(X)}{P_K(X)} = \Omega_K(X).$$

Moreover, the condition for an α-quantile of F, $F(e_\alpha(X)) = \alpha$ is equivalent to

$$\Omega_{e_\alpha(X)}(X) = \frac{1-\alpha}{\alpha}, \quad (32)$$

i.e., the expectile can be interpreted as the value of the threshold parameter that fixes the Omega ratio at a level determined by α. Eq. (32) is the well-known connection between Omegas and expectiles (compare formula (5) in [41]), confirming that the expectiles of G are the quantiles of F.

Appendix C. Arrow–Pratt approximations for expected utility under background risk

In this technical appendix, we provide Arrow–Pratt approximations for expected utility under background risk. In principle, this material is classical, see, e.g., [19,20] or the textbook of Gollier [16]. However, we are not aware of a previous treatment of this material that explicitly spells out the regularity conditions that are needed when considering risks with an unbounded support. Yet, for a fair comparison with the Omega ratio, this additional generality is needed. We begin with a general result and then specialize it to CARA and CRRA utility functions. Our result relies on the following technical assumption which is the counterpart of Assumption 1 in our analysis of the Omega ratio.

Assumption 2. Let $u : W \to \mathbb{R}$ be a twice continuously differentiable, strictly increasing and concave function with support $W \subseteq \mathbb{R}$ of the form $W = (w, \infty)$ for some $w \in [-\infty, \infty)$. Fix some $\delta > 0$ and assume that the real-valued random variables X and S satisfy that $X + \delta S \in W$ a.s., that S is square-integrable and that $E[S|X] = 0$. Moreover, we assume the following:

(i) **Well-definedness.** $E[u(X + hS + C)] < \infty$ for all $h \in [0, \delta]$ and $C \geq 0$.
(ii) **Regularity.**
   1. There exists an integrable function $\theta_1(X, S)$ independent of h such that for all $h \in [0, \delta]$, $C \geq 0$, and $k \in \{0, 1\}$, $|S^k u'(X + hS + C)| \leq \theta_1(X, S)$.

[38] To see that this indeed a cumulative distribution function, one can check that the associated density is $G(K)/(C_K(X) + P_K(X)) \geq 0$ or verify that the expression is in line with [36].

2. There exists an integrable function $\theta_2(X,S)$ independent of h such that for all $h \in [0,\delta]$, $C \geq 0$, and $k \in \{0,1,2\}$, $|S^k u''(X+hS+C)| \leq \theta_2(X,S)$.

Notice that the strict monotonicity of u together with the regularity part of Assumption 2 implies the following condition to which we will refer as "non-degeneracy",

$$E[u'(X+hS+C)] > 0 \quad \text{for all} \quad h \in [0,\delta] \quad \text{and} \quad C \geq 0. \tag{33}$$

Moreover, as before, we denote by $\sigma^2(X)$ the conditional variance of S, $\sigma^2(X) = \mathrm{Var}(S|X)$. Then we have the following expansion for utility-based risk premia:

Proposition 4. *In the setting of Assumption 2, the utility-based risk premia $\pi_u(h)$, $h \in [0,\delta]$, that are defined via*

$$E[u(X+hS+\pi_u(h))] = E[u(X)] \tag{34}$$

satisfy

$$\pi_u(h) = \pi''(0)\frac{h^2}{2} + O(h^3)$$

where

$$\pi_u''(0) = -\frac{E[\sigma^2(X)u''(X)]}{E[u'(X)]}.$$

Proof. Formally applying a second order Taylor expansion to the function π_u around 0 for small enough h, we obtain

$$\pi_u(h) = \pi_u(0) + \pi_u'(0)h + \pi_u''(0)\frac{h^2}{2} + O(h^3).$$

To prove the proposition, we thus need to justify the Taylor expansion and make sure that the coefficients match those given in the proposition. After noting that, obviously, $\pi_u(0) = 0$, the remainder of the proof thus consists of proving the following three points:

1. π_u is twice differentiable at 0,
2. $\pi_u'(0) = 0$, and
3. $\pi_u''(0) = -\frac{E[u''(X)\sigma^2(X)]}{E[u'(X)]}$.

Consider the defining condition (34) of the risk premium. Given the first regularity assumption, the expression $E[u(X+hS+\pi)]$ is continuously differentiable with respect to π, and we have $\frac{\partial}{\partial\pi}E[u(X+hS+\pi)] > 0$ from the non-degeneracy condition (33). Thus, we can establish that $\pi'(h)$ exists and is bounded by the implicit function theorem.

Taking a derivative with respect to h on both sides of (34) yields

$$E[u'(X+hS+\pi_u(h))(S+\pi'(h))] = 0. \tag{35}$$

Note that we need regularity condition 1 to ensure that the derivative takes the desired form. In the limiting case $h = 0$, we have

$$0 = E[u'(X)(S+\pi_u'(0))] = E[u'(X)E[(S+\pi_u'(0))|X]] = \pi_u'(0)E[u'(X)]$$

where we use that $E[S|X] = 0$. By the non-degeneracy condition, we can conclude that $\pi_u'(0) = 0$. This concludes the proof for point 2 on our list.

For the next steps of the proof, we need to take one more derivative with respect to h. We can rewrite (35) as

$$\pi_u'(h) = -\frac{E[Su'(X+hS+\pi_u(h))]}{E[u'(X+hS+\pi_u(h))]} \tag{36}$$

It is clear that π_u' is differentiable if the numerator and denominator are differentiable and the denominator is non-zero. This is guaranteed by the non-degeneracy condition and by the regularity condition.

Taking a derivative on both sides of (35), we obtain

$$\begin{aligned} 0 =& \pi_u''(h)E[u'(X+hS+\pi_u(h))] + E[u''(X+hS+\pi_u(h))(S^2+S\pi_u'(h))] \\ &+ \pi_u'(h)E[u''(X+hS+\pi_u(h))(S+\pi_u(h))] \end{aligned}$$

Plugging in $h = 0$, we find that

$$0 = \pi_u''(0)E[u'(X)] + E[u''(X)S^2].$$

Solving for $\pi_u''(0)$ and using that $E[S^2|X] = \sigma^2(X)$ yields

$$\pi_u''(0) = -\frac{E[u''(X)\sigma^2(X)]}{E[u'(X)]}.$$

This concludes the proof of point 3. □

The next lemma provides a sufficient condition on the distribution of X and S that ensures that Assumption 2 is satisfied for CARA utility. Here, we consider a range A for the risk aversion parameter α. This range should be thought of as the range of risk aversion level which are being compared.

Lemma 6. *Consider constant absolute risk aversion (CARA) utility functions u with parameter α,*

$$u(x) = -\exp(-\alpha x),$$

for α taken from some interval $A \subseteq (0,\infty)$ and support $W = \mathbb{R}$. Assume that S is square-integrable and that $E[S|X] = 0$. Then Assumption 2 is satisfied for all $\alpha \in A$ and some $\delta > 0$ when the joint Laplace transform of X and S exists, i.e.,

$$E[\exp(-aX - bS)] < \infty \tag{37}$$

for all $(a,b) \in \mathbb{R}^2$ with $a \in A$ and $b \in [-2, 2+a\delta]$.

Proof. Clearly, u is concave, twice differentiable and strictly increasing. The utility $u(X+hS)$ is integrable by assumption. We also have

$$u'(x) = \alpha\exp(-\alpha x) \text{ and } u''(x) = -\alpha^2\exp(-\alpha x).$$

Next, we need to verify regularity. It is easy to see that for $|S| > 1$, $1 \leq |S| \leq |S|^2 \leq \exp(2|S|)$ and for $|S| \leq 1$, $1 \geq |S| \geq |S|^2$. Therefore,

$$\begin{aligned} &|S^k u'(X+hS+C)| \\ =&\mathbf{1}_{\{|S|>1\}}\alpha|S|^k\exp(-\alpha(X+hS+C)) \\ &+ \mathbf{1}_{\{|S|\leq 1\}}\alpha|S|^k\exp(-\alpha(X+hS+C)) \\ \leq&\mathbf{1}_{\{|S|>1\}}\alpha\exp(-\alpha X)\exp(-\alpha hS + 2|S|)\exp(-\alpha C) \\ &+ \mathbf{1}_{\{|S|\leq 1\}}\alpha\exp(-\alpha X)\exp(-\alpha hS)\exp(-\alpha C) \\ \leq&\alpha\exp(-\alpha X)\exp(-\alpha hS)(\exp(2|S|)+1) \\ \leq&\alpha\exp(-\alpha X)(1+\exp(-\alpha\delta S))(1+\exp(-2S)+\exp(2S)), \end{aligned}$$

for $k = 0,1$. We use the fact the exponential function is positive and monotonic, and $a, b \leq \max\{a,b\} \leq a+b$ for $a, b \geq 0$. Similarly,

$$\begin{aligned} &|S^k u''(X+hS+C)| \\ &\leq \alpha^2\exp(-\alpha X)(1+\exp(-\alpha\delta S))(1+\exp(-2S)+\exp(2S)), \end{aligned}$$

for $k = 0,1,2$. We choose

$$\theta_1(X,S) = \alpha\exp(-\alpha X)(1+\exp(-\alpha\delta S))(1+\exp(-2S)+\exp(2S))$$

and

$$\theta_2(X,S) = \alpha^2\exp(-\alpha X)(1+\exp(-\alpha\delta S))(1+\exp(-2S)+\exp(2S)).$$

This concludes the proof. □

Remark 2. From Lemma 6, we see that the main condition is existence of the Laplace transform (37). This is always satisfied for X and S taken from a bounded support or, e.g., drawn from a Gaussian distribution. With heavier tails, however, (37) is not guaranteed. For instance, suppose that $-X$ is drawn from an exponential distribution with rate λ so that X is distributed on $(-\infty, 0)$ with density $\lambda\exp(\lambda x)$. In that case, (37) can only hold for $a < \lambda$.

We next turn to CRRA utility.

Lemma 7. *Let $A \subseteq (0,\infty)$ be an interval. Assume that S is square-integrable and that $E[S|X] = 0$. For $\gamma \in A$, consider constant relative risk aversion (CRRA) utility functions u with parameter γ,*

$$u(x) = \frac{x^{1-\gamma}-1}{1-\gamma} \quad \text{for} \quad \gamma \neq 1 \quad \text{and} \quad u(x) = \log(x) \quad \text{for} \quad \gamma = 1$$

and support $W = (0,\infty)$. Then Assumption 2 is satisfied for all $\gamma \in A$ and some $\delta > 0$ if for all $h \in [0,\delta]$, all $C \geq 0$ and all $\gamma \in A$, $u(X + hS + C)$ is integrable and if $X + \delta S \geq \epsilon > 0$ a.s. for some $\epsilon > 0$.

Proof. Integrability is given. We also have (both for $\gamma \neq 1$ and $\gamma = 1$)

$$u'(x) = x^{-\gamma} \text{ and } u''(x) = -\gamma x^{-\gamma-1} \tag{38}$$

Since the utility is strictly increasing, the non-degeneracy condition (33) is satisfied. Moreover, we have

$$|S^k u'(X + hS + C)| = \frac{|S|^k}{(X + hS + C)^\gamma} \leq \frac{1}{\epsilon^\gamma}(1 + |S|) =: \theta_1(S),$$

almost surely. Similarly,

$$|S^k u''(X + hS + C)| \leq \gamma \frac{|S|^k}{\epsilon^{(1+\gamma)}} \leq \frac{\gamma}{\epsilon^{(1+\gamma)}}(1 + |S| + |S|^2) =: \theta_2(S)$$

almost surely. This concludes the proof since S was assumed to be square-integrable. □

Remark 3.

1. If $A \subseteq (1,\infty)$, we do not need to assume integrability of $u(X + hS + C)$ separately since in that case u is bounded over $[\epsilon,\infty)$, guaranteeing integrability.
2. The assumption that $X + \delta S \geq \epsilon > 0$ can in principle be relaxed to $X + \delta S > 0$ a.s. for $\gamma > 0$ and even to $X + \delta S \geq 0$ a.s. for $\gamma \in (0,1)$ when suitable technical conditions on $\sigma^2(\cdot)$ and on the distribution of X near 0 are put in place.[39] We leave out further details here but note that – in one way or other – assumptions need to be in place that keep $X + hS$ away from the negative numbers and guarantee sufficient regularity near 0.

References

[1] Mausser Helmut, Saunders David, Seco Luis. Optimizing omega. Risk 2006;19(11):88–92.
[2] Guastaroba Gianfranco, Mansini Renata, Ogryczak Włodzimierz, Speranza Maria Grazia. Linear programming models based on Omega ratio for the enhanced index tracking problem. European J Oper Res 2016;251(3):938–56.
[3] Balbás Alejandro, Balbás Beatriz, Balbás Raquel. Omega ratio optimization with actuarial and financial applications. European J Oper Res 2021;292(1):376–87.
[4] Sehgal Ruchika, Sharma Amita, Mansini Renata. Worst-case analysis of Omega-VaR ratio optimization model. Omega 2023;114:102730.
[5] Ogryczak Włodzimierz, Ruszczyński Andrzej. Dual stochastic dominance and related mean-risk models. SIAM J Optim 2002;13(1):60–78.
[6] Shaked Moshe, Shanthikumar J George. Stochastic orders. 2nd ed.. Springer; 2007.
[7] Balder Sven, Schweizer Nikolaus. Risk aversion vs. the Omega ratio: Consistency results. Finance Res Lett 2017;21:78–84.
[8] Bernard Carole, Caporin Massimiliano, Maillet Bertrand, Zhang Xiang. Omega compatibility: A meta-analysis. Comput Econ 2022.
[9] Darbyshire Paul, Hampton David. Hedge fund modelling and analysis using Excel and VBA. John Wiley & Sons; 2012.
[10] Snyder Alan, Parrish Joel, Lam Tuling, Zou Yunfei. An investor's guide to the risk versus return conundrum. Working Paper, 2013, Shinnecock Partners.
[11] Vilkancas Renaldas. Characteristics of Omega-optimized portfolios at different levels of threshold returns. Bus Manag Educ 2014;12(2):245–65.
[12] Keating Con, Shadwick William F. A universal performance measure. J Perform Meas 2002b;6(3):59–84.
[13] Caporin Massimiliano, Costola Michele, Jannin Gregory, Maillet Bertrand. On the (ab)use of Omega? J Empir Finance 2018;46:11–33.
[14] Klar Bernhard, Müller Alfred. On consistency of the Omega ratio with stochastic dominance rules. In: Innovations in insurance, risk- and asset management. World Scientific; 2018, p. 367–80.
[15] Pratt John W. Risk aversion in the small and in the large. Econometrica 1964;32(1–2):122–36.
[16] Gollier Christian. The economics of risk and time. MIT Press; 2001.
[17] Mu Xiaosheng, Pomatto Luciano, Strack Philipp, Tamuz Omer. Background risk and small-stakes risk aversion. arXiv preprint 2010.08033, 2021.
[18] Keating Con, Shadwick William F. An introduction to Omega. In: AIMA newsletter. 2002a.
[19] Ross Stephen A. Some stronger measures of risk aversion in the small and the large with applications. Econometrica 1981;49(3):621–38.
[20] Kihlstrom Richard E, Romer David, Williams Steve. Risk aversion with random initial wealth. Econometrica 1981;49(4):911–20.
[21] Kahneman Daniel, Tversky Amos. Prospect theory: An analysis of decision under risk. Econometrica 1979;47(2):263–91.
[22] Tversky Amos, Kahneman Daniel. Advances in prospect theory: Cumulative representation of uncertainty. J Risk Uncertain 1992;5(4):297–323.
[23] Bell David E. Regret in decision making under uncertainty. Oper Res 1982;30(5):961–81.
[24] Loomes Graham, Sugden Robert. Regret theory: An alternative theory of rational choice under uncertainty. Econ J 1982;92(368):805–24.
[25] Abel Andrew B. Asset prices under habit formation and catching up with the Joneses. Amer Econ Rev 1990;80(2):38–42.
[26] Kőszegi Botond, Rabin Matthew. A model of reference-dependent preferences. Q J Econ 2006;121(4):1133–65.
[27] Lopes Lola L. When time is of the essence: Averaging, aspiration, and the short run. Organ Behav Hum Decis Process 1996;65(3):179–89.
[28] Manski Charles F. Ordinal utility models of decision making under uncertainty. Theory and Decision 1988;25:79–104.
[29] Zeisberger Stefan. Do people care about loss probabilities? J Risk Uncertain 2022;65(2):185–213.
[30] Jorion Philippe. Value at risk: the new benchmark for managing financial risk. McGraw-Hill; 1997.
[31] Mendelson Haim. Quantile-preserving spread. J Econom Theory 1987;42(2):334–51.
[32] Chambers Christopher P. An axiomatization of quantiles on the domain of distribution functions. Math Finance 2009;19(2):335–42.
[33] Rostek Marzena. Quantile maximization in decision theory. Rev Econom Stud 2010;77(1):339–71.
[34] de Castro Luciano, Galvao Antonio F, Montes-Rojas Gabriel, Olmo Jose. Portfolio selection in quantile decision models. Ann Finance 2022;18(2):133–81.
[35] Newey Whitney K, Powell James L. Asymmetric least squares estimation and testing. Econometrica 1987;55(4):819–47.
[36] Jones M Chris. Expectiles and M-quantiles are quantiles. Statist Probab Lett 1994;20(2):149–53.
[37] Bellini Fabio, Cesarone Francesco, Colombo Christian, Tardella Fabio. Risk parity with expectiles. European J Oper Res 2021;291(3):1149–63.
[38] Bellini Fabio, Fadina Tolulope, Wang Ruodu, Wei Yunran. Parametric measures of variability induced by risk measures. Insurance Math Econom 2022;106:270–84.
[39] Bellini Fabio, Klar Bernhard, Müller Alfred, Gianin Emanuela Rosazza. Generalized quantiles as risk measures. Insurance Math Econom 2014;54:41–8.
[40] Ziegel Johanna F. Coherence and elicitability. Math Finance 2016;26(4):901–18.
[41] Bellini Fabio, Klar Bernhard, Müller Alfred. Expectiles, Omega ratios and stochastic ordering. Methodol Comput Appl Probab 2018;20:855–73.
[42] Kazemi Hossein, Schneeweis Thomas, Gupta Bhaswar. Omega as a performance measure. J Perform Meas 2004;8:16–25.
[43] Glasserman Paul. Monte Carlo methods in financial engineering. Springer; 2004.
[44] Ross Sheldon M. Introduction to probability models. 10th ed.. Academic Press; 2010.
[45] Boyd Stephen P, Vandenberghe Lieven. Convex optimization. Cambridge University Press; 2004.
[46] Schuhmacher Frank, Eling Martin. A decision-theoretic foundation for reward-to-risk performance measures. J Bank Finance 2012;36(7):2077–82.
[47] Bassett Gilbert W. The St. Petersburg paradox and bounded utility. Hist Polit Econ 1987;19(4):517–23.
[48] Machina Mark J, Neilson William S. The Ross characterization of risk aversion: Strengthening and extension. Econometrica 1987;55(5):1139–49.

[39] For instance, suppose that the support of X goes down to 0. Then, since $E[S|X] = 0$, we need that $\sigma^2(x)$ vanishes for $x \downarrow 0$ to ensure that $X + hS$ cannot get negative.