Finance Research Letters 000 (2016) 1–7

Contents lists available at ScienceDirect

Finance Research Letters

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

Risk aversion vs. the Omega ratio: Consistency results

Sven Balder[a], Nikolaus Schweizer[b,*]

[a] Federal Agency for Financial Market Stabilisation, Germany
[b] Department of Econometrics and OR, Tilburg University, Netherlands

ARTICLE INFO

Article history:
Received 27 April 2016
Revised 23 December 2016
Accepted 25 December 2016
Available online xxx

JEL Classification:
G11
D81

Keywords:
Performance measurement
Stochastic dominance
Omega ratio
Sharpe ratio

ABSTRACT

This paper clarifies when the Omega ratio and related performance measures are consistent with second order stochastic dominance and when they are not. To avoid consistency problems, the threshold parameter in the ratio should be chosen as the expected return of some benchmark – as is commonly done in the Sharpe ratio. When the ratio is below one, its value should be discarded – just like a negative Sharpe ratio. Finally, we show that a class of closely related performance measures has both better consistency properties and greater flexibility.

1. Introduction

In the past decade, the Omega ratio of Keating and Shadwick (2002) has received considerable attention as an alternative to the classical Sharpe ratio when measuring the performance of different investment opportunities.[1] The Omega ratio can be interpreted as a return-risk ratio. The return is measured as the expected gain above some threshold whereas risk is measured as the expected loss beyond the same threshold. Since only adverse deviations from the threshold contribute to the risk component of the Omega ratio, this measure is particularly attractive for situations where the return distribution is strongly asymmetric. This is in contrast to the Sharpe ratio which measures risk by the standard deviation and thus does not differentiate between upward and downward deviations from the benchmark, implicitly assuming that they are of the same magnitude.

In the literature, it is sometimes stated that another advantage of the Omega ratio over the Sharpe ratio is that it is consistent with Second Order Stochastic Dominance (SSD). The most important consequence of this property is that when one potential payoff is preferred over another by any risk-averse investor, then this payoff also has a greater Omega ratio. In a recent article, Caporin et al. (2015) point out that this property does not always hold, and provide various sufficient and necessary conditions.[2]

* Corresponding author.
E-mail addresses: sven.balder@fmsa.de (S. Balder), n.f.f.schweizer@uvt.nl (N. Schweizer).

[1] See Caporin et al. (2015) for a recent overview of many papers which apply the Omega ratio in a variety of contexts and also for a general overview of the topic.

[2] Caporin et al. (2015) also provide references to various authors who claim that the Omega ratio, or its special case the Gain-Loss-Ratio, is consistent with SSD. In some cases, these claims are indeed false (e.g. Prigent, 2007) while in others (e.g. Cherny and Madan, 2009) the claim is correct since attention is restricted to payoffs whose mean lies above the threshold.

http://dx.doi.org/10.1016/j.frl.2016.12.012
1544-6123/© 2017 Published by Elsevier Inc.

We revisit the main question of Caporin et al. (2015) and determine sufficient conditions for the ranking implied by the Omega ratio to be in line with the ranking implied by SSD. We show that their conditions and arguments can be simplified. In addition, we provide an elementary sufficient condition for SSD and the Omega ratio to be consistent: When choosing between more than two payoffs, a payoff which is dominant in SSD also has the highest Omega ratio, provided that at least one payoff in the choice set has an expected value above the threshold parameter chosen in the ratio. This observation applies not only to the Omega ratio but also to related performance measures such as the Sortino ratio and the Kappa of Kaplan and Knowles (2004).

Our work also extends and clarifies earlier results by Kazemi et al. (2004) and Darsinos and Satchell (2004). Kazemi et al. (2004) point out that the question whether Omega is smaller or greater than one is analogous to the question whether the Sharpe ratio is negative. However, they do not address consistency with stochastic dominance – which is one main aspect in which Omega has better properties than the Sharpe ratio. In contrast, Darsinos and Satchell (2004) derive consistency results which are similar to ours but fail to notice the pathological cases where these results are reversed.[3] Finally, there is a sizeable literature in the Operation Research community studying consistency with SSD for various deviations from the classical mean-variance framework, see e.g. Branda and Kopa (2014) and Ogryczak and Ruszczyński (2001).

Nevertheless, these results are not widely-known and appreciated in the applied literature. We provide a comprehensive and self-contained overview of the consistency properties of Omega and related ratios with an emphasis on potential problems and a view towards applications. In the latter regard, the two main messages of the paper are as follows:

(i) When ranking investment opportunities, the Omega ratio should be applied with care.[4] There is a simple rule which characterizes careful behavior: Omega ratios which are smaller than one should be discarded before the ranking.
(ii) One of the major appeals of the Omega ratio over the Sharpe ratio is that it allows to plot rankings for varying levels of risk aversion through varying the threshold. In Section 4, we argue that there are alternative performance measures which can be plotted in a similar way but which respect risk aversion for all assets with a positive excess return. The basic idea behind these measures is to separate the threshold which defines excess returns from the threshold which characterizes undesirable outcomes.

Even though the Sharpe ratio is generally inconsistent with SSD, it faces a similar “problem” that can illustrate the main idea. Consider some payoff X with mean μ and standard deviation σ. For any real number K, we can define a generalized Sharpe ratio with threshold K, $S_K(X)$, by

$$S_K(X) = \frac{\mu - K}{\sigma}.$$

Suppose that the mean μ lies below the threshold, $\mu < K$, so that S_K is negative. In this case, increasing σ pushes the ratio closer to zero – and thus increases it. Consequently, payoffs with the same expected return and higher risk have a larger Sharpe ratio. There is a simple reason why *this* paradox of the Sharpe ratio is largely irrelevant in practice. Typically, the threshold K is chosen equal to the risk-free rate or as the expected return of some other reference asset. This guarantees that payoffs with a return below the threshold are never optimal since their Sharpe ratio is negative. Thus, while the Sharpe ratio implies an absurd ranking of payoffs whose means lie below that of the reference asset, the relevant part of the ranking, i.e., the ranking of payoffs which are at least as good as the reference, is unaffected by this problem.

The same reasoning can eliminate potential consistency problems of the Omega ratio: The threshold parameter should be chosen at (or below) the mean of some available benchmark. The implied internal ranking of payoffs whose means lie below the threshold should not be taken at face value. Whereas this fact is widely accepted for the Sharpe ratio, it appears to be far less established for the Omega ratio where the threshold is often viewed as a free parameter, see for instance Mausser et al. (2006) who also report that the regime where the threshold lies below the mean ($\Omega < 1$) is particularly challenging from a computational point of view.

Some authors (e.g. Bernard and Ghossoub, 2010) have viewed the shortcomings of the Omega ratio as strengths, relating potential reversals of risk-averse behavior to the risk-lovingness in the loss domain pointed out by Kahneman and Tversky (1979). The present paper is closer in spirit to the normative (rather than behavioral) intentions of Keating and Shadwick (2002), identifying suitable decision criteria for risk-averse agents which capture variations in the degree of risk aversion while retaining the simplicity of the Sharpe ratio. While we do not follow Keating and Shadwick (2002) to the point of viewing Omega as a suitable replacement for utility theory, we do agree that often the impact of risk aversion on decisions is simple enough to be modeled by a single parameter.

The paper is organized as follows. Section 2 introduces the setting and the Omega ratio. Section 3 discusses consistency of the Omega ratio with second order stochastic dominance. Section 4 generalizes these consistency results to related performance measures including the Sortino ratio and the Kappa. We also propose an alternative class of performance measures which has better consistency properties than these measures while retaining their flexibility. All proofs are short and stated

[3] There is one difference in terminology between our paper and theirs. We define SSD via the increasing concave order while they define it via the concave order. Our SSD property is called MSSD in their paper. Their SSD property is equivalent to our SSD property together with the assumption that the distributions have equal means, the setting of our Proposition 1.

[4] Caporin et al. (2015) provide many references to studies which vary the threshold parameter without worrying about consistency. For a typical example, see Table 5.2 in Darbyshire and Hampton (2012).