Journal Of Investment Management, Vol. 15, No. 4, (2017), pp. 20–29

www.joim.com

INVESTMENT HORIZON RISK AND VOLATILITY METRICS

Robert Korkie[a]

We re-examine the literatures' disparate conclusions that stock returns are more (less) volatile over longer investment horizons. We claim that the commonly employed variance ratio is incapable of generally determining whether investment risk increases with investment horizon. We demonstrate that the use of effective returns and standard deviation ratios have significantly different results compared to continuous return variance ratios. Using basic return generating processes and standard deviations, plus a recent well-specified study, we find stocks are less volatile over short to long horizons but are more volatile over very long horizons. The conclusions are consistent with some research for very long horizons but inconsistent in short to long horizons.

Conventional but old wisdom says that stock returns are less risky over longer investment horizons; however, some research contrarily finds that stocks are more risky over longer horizons when viewed from an investor's perspective.[1] The cause is imprecise conditional expected returns and their uncertain future path that is surely descriptive of financial markets. The issue is an important one for asset allocation, hedging, portfolio rebalancing, and glide paths, for example.[2]

Commonly, the literature's analysis of horizon risk uses the variance ratio computed from continuously compounded nominal returns, generally referred to as log returns. This variance ratio was initially intended for tests of the random walk hypothesis via autocorrelations in subperiod returns and not for testing the horizon risk hypothesis.[3] The basis of this metric in random walk analysis is that the variance of the sum of T uncorrelated returns equals T times the 1-period return variance. We demonstrate that the fundamental problem for horizon risk analysis is that, under quite reasonable measures of horizon return, the equality holds but the quantity is not equal to the T-period variance.

The first premise for our argument against the traditional variance ratio is that investors should use effective (simple) return and not nominal (log) return in measuring horizon performance

[a]Professor Emeritus, University of Alberta, Edmonton, Alberta, Canada. Postal: 733 2 St. Canmore AB T1W 2K3 Canada. E-mail: bkorkie@ualberta.ca

because of the former's linear correspondence with wealth.[4] Secondly, the standard deviation (volatility) of effective returns is the preferred horizon risk measure because it is in the same units as return and, jointly with expected return, linearly determines the confidence interval for return generated over a specified horizon. Thirdly, the average risk per period, or equivalently the per unit time risk from a T-period investment, is correctly obtained from division by the number of periods and not its square root, irrespective of the risk measure but particularly applicable to standard deviation. Failure to do so implicitly changes the length of the investment horizon because the product of T times that average risk is not equal to the horizon risk.

Rather than estimating a particular model, we define return generating conditions starting from pure random (iid) series in continuous nominal returns to heteroskedastic effective returns that are lognormally distributed and conclude with a recent study that is well estimated and has verifiable conditions. We demonstrate with very simple analysis that, irrespective of a particular model, the use of effective returns and standard deviation ratios can have significantly different horizon behavior compared to the traditional continuously compounded nominal return variance ratios. We conclude that, in short to long investment horizons, stocks are less volatile than in an annual horizon but are more volatile in very long horizons, for reasonable return processes and distributions.

The first section investigates the commonly used iid return assumption applied to continuous nominal return or effective return. Section 2 discusses the impact of changing variance and volatility. Section 3 compares the Pástor and Stambaugh (2012) predictive variance ratio results with our corresponding standard deviation ratios. Section 4 concludes.

1 Return process and the risk ratios

The iid assumption or the weaker assumption of uncorrelated returns, applied to continuously compounded nominal (log) returns, facilitates a nice financial interpretation if variance ratios are used to examine horizon risk. A probable explanation for the use of variance ratios in the conventional analyses of investment horizon and risk is that if sequential returns are iid or simply uncorrelated with finite variance, the sum of T returns has variance equal to T times the per period variance and going the other way, the per unit time variance equals the horizon variance divided by T. Additionally, the sum of the variances is not equal to the horizon variance unless the return is defined as the continuously compounded return. Furthermore, if continuously compounded (i.e., log) returns are normally distributed, limited liability is not violated.

Because the variance and higher order raw return moments are not in the same units as return, in contrast to the standard deviation, its use may cause some conflicting conclusions regarding horizon risk. We claim that the standard deviation applied to effective returns is the more reliable metric for analyzing horizon risk.

In order to address the horizon risk issue, we use the traditional variance ratio, $VarRatio(T) = \frac{Var(R_T)/T}{Var(r_1)}$, and its less common volatility ratio counterpart, $VolRatio(T) = \frac{StdDev(R_T)/T}{StdDev(r_1)}$, where R_T is the T-period return and r_1 is the period-1 return. These ratios measure the T-period average horizon risk relative to the period-1 horizon risk and are not restricted to continuous returns nor is the numerator defined as the simple sum of T single period variances.

1.1 iid Continuous return risk ratios

To provide a starting 1-year horizon variance and volatility (standard deviation) for our analysis, the S&P500 continuously compounded mean and volatility are parameterized from S&P500 total annual returns over the period Jan 31, 1994 to Dec. 31, 2014 resulting in the continuous mean of 9.01%/year and the volatility and variance of 19.08%/year and 364%^{2}/year, respectively. The intention is not to claim that these sample moments provide the true annual moments; rather, they are reasonable values for the first year of our horizon risk analysis.

The starting 1-year horizon values are not critical to our analysis. Henceforth, we simply postulate the 1-year volatility as being derived from a predictive model or as a conditional volatility depending on the context and employ our preceding values as an example. Then, specific processes and distribution assumptions are made describing the evolution of annual variance and volatility. As we demonstrate in subsequent sections, the interpretation of the variance and standard deviation (volatility) ratios conflict depending on the continuous nominal versus effective return convention, as well as on the properties of the distributions and return generating processes.

Figure 1 shows the variance and volatility ratios for our parameters. The assumption of iid continuously compounded returns results in the well-known conclusion that the variance ratio is constant with an increasing horizon and the volatility ratio declines at a decreasing rate. In this simplistic case, stocks are of constant risk in the long run, when risk is measured by the continuously compounded return variance and stocks are less risky in the long run, when risk is measured by the continuously compounded return volatility. Conclusions about the effective return ratios are not possible without a distributional assumption, as in Section 1.3.

1.2 *iid Effective return risk ratios*

For comparison, Figure 2 shows the effective variance and volatility ratios from the parameters with the assumption of iid effective annual returns. The effective return variance for horizon t is calculated recursively using the formula,

$$\begin{aligned} V(R_t) &= V((1+R_{t-1})(1+r_t)) \\ &= (E(1+r_t))^2 V(1+R_{t-1}) \\ &\quad + (E(1+R_{t-1}))^2 V(1+r_t) \\ &\quad + V(1+R_{t-1}) V(1+r_t), \\ &\qquad t = 1, 2, \ldots, T \end{aligned} \tag{1}$$

where R_t is the t-year horizon effective return, $R_0 = 0$, r_t is the iid 1-year effective return in year t, and T is the horizon limit in years.

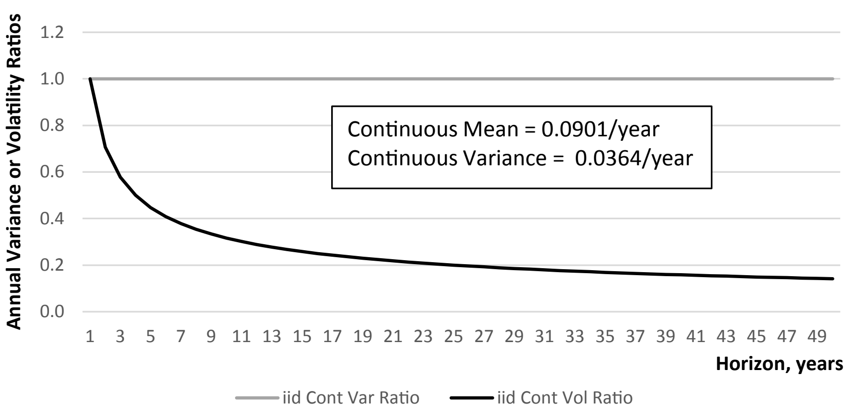

Figure 1 S&P500 annualized continuously compounded nominal variance & volatility ratios, assuming iid continuous nominal returns.

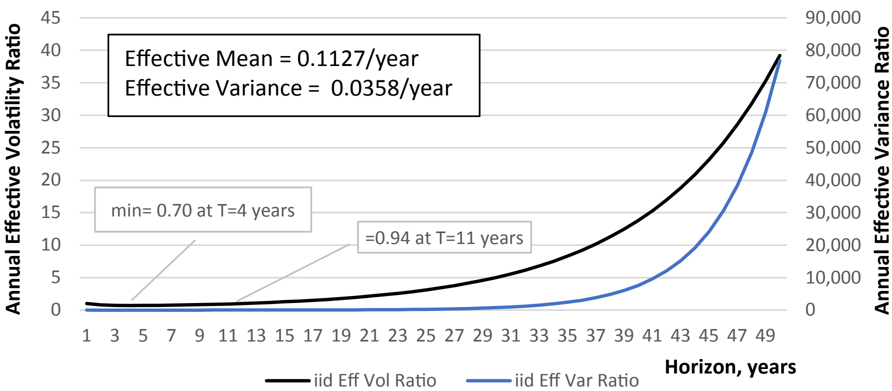

Figure 2 S&P500 annualized effective return variance & volatility ratios, assuming iid effective returns.

In contrast to the iid continuously compounded case of a constant variance ratio, the iid effective return variance ratio increases monotonically and is extremely large at $T = 50$ years. In contrast to Figure 1's monotonically declining continuous return volatility ratio, the effective return volatility ratio decreases until the minimum of 0.7 at $t = 4$ years, remaining below 1 at $t = 11$, and increasing monotonically thereafter to 39.2 at $T = 50$. Note that because it is the effective returns that are iid, their variance sum over T years is not equal to the T-year effective return variance. For example, the sum of the first three years' variances is 0.107 whereas the 3-year horizon variance is 0.169.

If effective returns are iid and the literatures' conventional variance ratio was calculated from continuous lognormal returns, the 4-year horizon result in this example would be a variance ratio of 1.0 with the conclusion that the 4-year horizon is of equal risk to the 1-year horizon. In contrast, the effective return volatility ratio is 0.70 with the conclusion that the 4-year horizon is of lower risk. Not only is there a conflict with the traditional continuous variance ratio but from the effective return variance ratio, stocks are more risky at all horizons; whereas, from the effective return volatility ratio, stocks are less risky over short to medium horizons and more risky at long horizons.

1.3 *iid Continuous returns and effective return risk ratios from lognormal distributions*

In this subsection, we add a lognormal effective return distribution assumption with iid continuous nominal returns and use effective return risk ratios, for comparisons with the previous subsection's iid effective return, effective return risk ratios, and no distribution assumption. Figure 3 shows a similar behavior to the risk ratios in Figure 2.

Figure 3 lognormal effective return moments are calculated from the horizon continuous nominal return moments as

$$E(R_t) = e^{(E(R_t^*)+0.5V(R_t^*))} - 1,$$
$$V(R_t) = (e^{V(R_t^*)} - 1)(e^{2E(R_t^*)+V(R_t^*)}), \qquad (2)$$
$$t = 1, 2, \ldots, T$$

where R_t is the t-year horizon effective return, R_t^* is the t-year horizon continuous nominal return, and T is the horizon limit in years.

In Figure 3, the effective return variance ratio increases monotonically and is large at $T = 50$ years. The effective return volatility ratio decreases until the minimum of 0.71 at $t = 4$ years, remaining below 1 at $t = 11$, and increasing monotonically thereafter to 47.5 at $T = 50$, nearly identical to Figure 2.

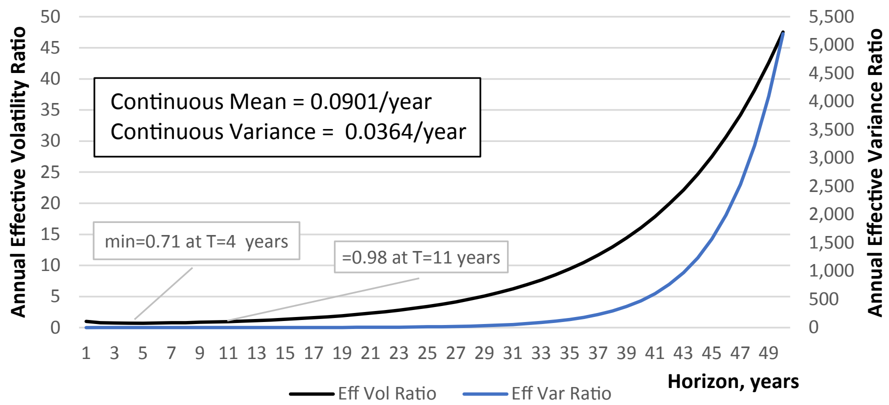

Figure 3 S&P500 annualized effective return variance & volatility ratios, assuming lognormal distributions and iid continuous nominal returns.

From an investor's perspective of the effective return conditional volatility ratio, stocks are less risky over short to medium horizons and more risky at long horizons; whereas from the perspective of the effective return variance ratio, stocks are more risky at all horizons. Whether continuous or effective returns are iid, does not affect this conclusion.

2 Heteroskedastic continuous returns, effective return risk ratios from lognormal distributions

In this section, we extend the previous Section 1.3 lognormal distribution and continuous returns iid assumptions but make the iid assumption conditional on a fixed multiplicative variance factor, F, and a constant mean.[5] A factor value exceeding 1 increases the conditional continuous variance per year. For example, the previous sections' continuous volatility is 0.1908/year and the continuous variance is 0.0364/year for every year including the starting year 1 but in the following figures, a factor value of $F = 1.3$ increases the continuous variance to 0.0473/year resulting in a continuous volatility of 0.2175/year.

In Figures 4a and 4b, each line is the horizon effective variance ratio or volatility ratio conditional on a specific factor value ranging from 1 to 1.5. In Figures 5a and 5b, the independent variable axis

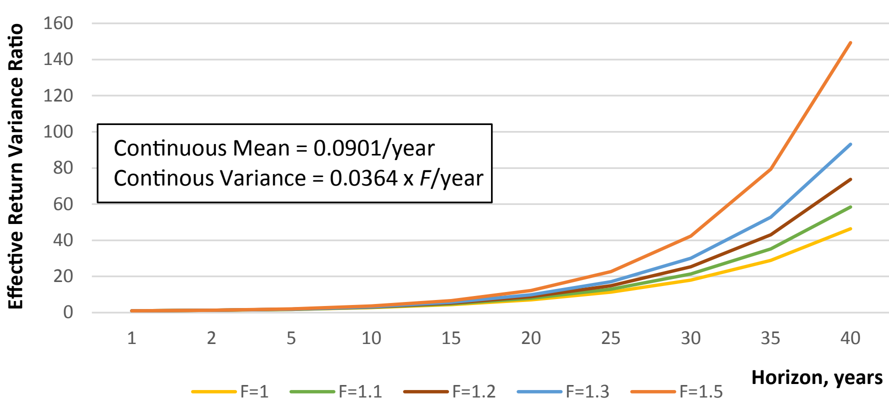

Figure 4a S&P500 effective return variance ratios, from lognormal distributions, iid conditional continuous return with fixed annual mean and annual variance that changes with a multiplicative factor F.

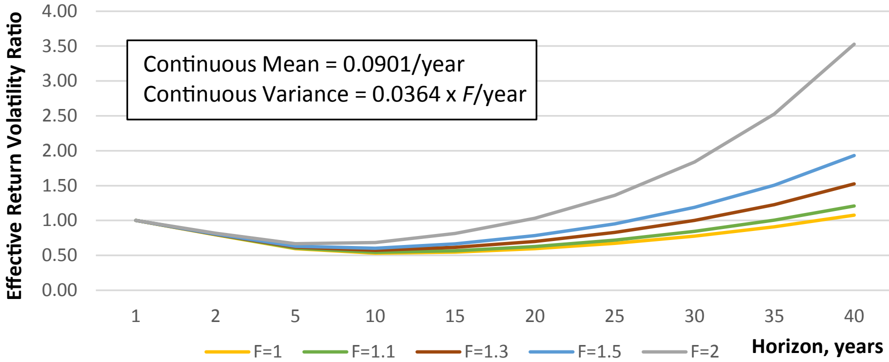

Figure 4b S&P500 effective return volatility ratios, from lognormal distributions, iid conditional continuous return with fixed annual mean and annual variance that changes with a multiplicative factor *F*.

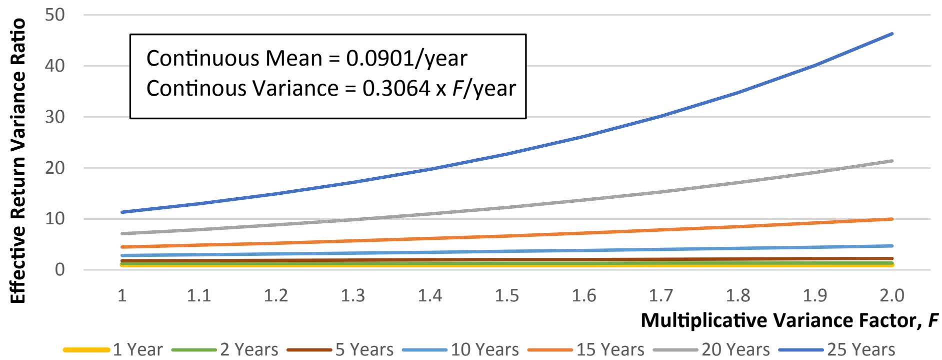

Figure 5a S&P500 fixed horizon effective return variance ratios, from lognormal distributions with fixed annual continuous mean and heteroskedastic continuous annual variance that changes with a multiplicative factor *F*.

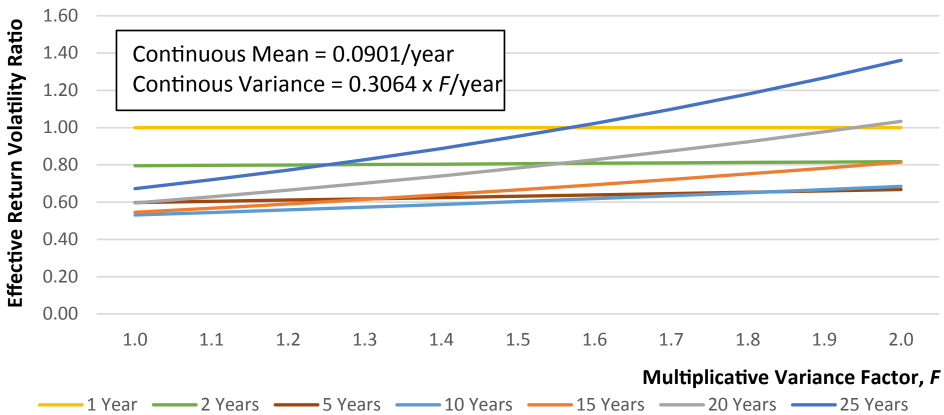

Figure 5b S&P500 fixed horizon effective return volatility ratios, from lognormal distributions with a fixed annual continuous mean and heteroskedastic continuous annual variance that changes with a multiplicative factor *F*.

is the factor value and each line represents a fixed horizon's effective return variance ratio or volatility ratio from lognormal distributions with a fixed annual nominal continuous mean and changing (heteroskedastic along lines) continuous annual nominal variance.

From Figure 4a, the effective return variance ratio increases monotonically for a fixed factor, as in Figure 3. Similar to Figure 3, Figure 4b shows that the effective return volatility ratio for a fixed F decreases until the minimum and remains below 1 for a relatively long time, increasing monotonically thereafter. This is true for all factor values less than 10 (not shown) but the minimum is reached at shorter horizons for larger factor values. For factor values from 1 to approximately 1.8, a minimum volatility ratio occurs at approximately the same 10-year horizon. Factors near 1 have relatively small effects on either risk ratio at short horizons. For example, with $F = 1$ the minimum effective return volatility ratio of 0.53 occurs at $t = 10$ years and exceeds 1 at $t = 40$ years; with $F = 1.8$, the minimum effective return volatility ratio of 0.65 occurs at $t = 10$ years and exceeds 1 between $t = 20$ and 25 years.

In Figures 5a and 5b, a factor greater than 1 increases the continuous variance per year. Figure 5a shows that variance ratios at any horizon are all monotonic, upward sloping, and are nonlinear in F, which is more pronounced at longer horizons. At short horizons, the factor value has little effect on the variance ratio. Figure 5b shows the effective return volatility ratios increase with the factor value for horizons exceeding 1. The factor value has little effect on the effective return volatility ratio, at shorter horizons.

From an investor's perspective of the effective return conditional volatility ratio, stocks are less risky up to quite long horizons and more risky at very long horizons; whereas from the perspective of the conditional variance ratio, stocks are more risky at all horizons, particularly for long horizons. However if the factor value, a surrogate for iid uncertainty, uncertainty about future expected returns, uncertainty about current expected return, and estimation risk is sufficiently large, the effective return volatility ratio is monotonically increasing in the horizon like the effective variance ratio. With representative S&P500 annual returns, continuous nominal return variance per year would have to be approximately 10 times its unconditional representative value, while maintaining the representative continuous mean, for effective volatility ratios to be monotonically increasing. A factor of 10 implies a huge effective return 1-year S&P500 volatility of 87%/year under lognormality. Therefore, in lognormal "normal market conditions", stocks are less volatile up to quite long horizons.

3 Pástor and Stambaugh comparative variance and standard deviation ratios

Pástor and Stambaugh (2012) (PS) re-examine the conventional wisdom that stock returns are less risky over longer investment horizons and contrarily find that stocks are more risky over long horizons, when viewed from an investor's perspective. This is because an investor's information set does not reveal the true return generating parameters and their predictor variables resulting in an imperfect proxy for the time dependent conditional expected return.[6] PS specify and estimate predictive systems over 206 years of S&P500 return data to analyze various horizons' predictive variances and their components. Conclusions are based on the traditional variance ratio metric calculated from continuously compounded return nominal variances. They conclude that the 50-year predictive variance per year is "at least 1.3 times higher than the variance at a 1-year horizon" and predictive variance "at a 10-year horizon is (often) higher than at a 1-year horizon".

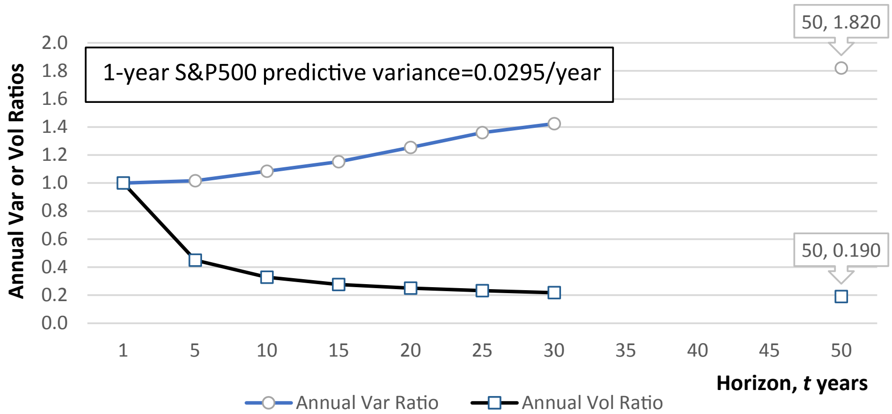

Figure 6 S&P500 annualized continuously compounded nominal variance (Var) & volatility ratios (Vol), from Pástor and Stambaugh (2012).

In this section, we utilize the results from the PS Predictive System 1. Figure 6 charts the PS variance ratio results obtained from their Figure 6, Panel A and Table 1, and Panel A and B Benchmark case. Our Figure 6 also shows the volatility (standard deviation) ratios derived from the PS annualized variances, as follows. First, a horizon volatility is computed as the square root of the quantity, horizon times the horizon annualized variance. The horizon volatility is divided by the horizon to produce the annual volatility, which is used to compute the continuously compounded nominal volatility ratios relative to year 1 volatility.[7]

The behavior of the continuous variance and volatility ratios differ substantially over the horizons. The variance ratio increases at an increasing rate until a horizon of approximately $t = 20$ years and then seems to increase at a constant rate until a value of 1.82 at $T = 50$ years. Contrarily, the volatility ratio decreases at a decreasing rate reaching a value of 0.19 at $T = 50$ years. Note that the PS predictive continuous variance ratio is somewhat larger than our Section 1 iid continuous variance ratio, particularly at long horizons, suggesting that the underlying return processes are not the same. However, the PS predictive continuous volatility ratio is only marginally larger than the iid volatility ratio, with values of 0.19 and 0.14 at $T = 50$ years and 0.33 and 0.32 at $t = 10$ years, respectively.

At a 10-year horizon, the PS continuous variance ratio is 1.085 or approximately 1.1 and the continuous volatility ratio is 0.33. Our Section 2 factor values are defined as multipliers of the 1-year continuous variance, under the assumptions of independent lognormal returns. Although, the assumptions differ, we can interpret a factor value as a continuous variance ratio for rough comparisons with the PS results. For a factor value of 1.1 and a 10-year horizon, the Figure 4a or 5a *effective return* variance ratio is 2.97 indicating a much riskier stock than indicated by the PS continuous variance ratio. However, the Figure 4b or 5b *effective return* volatility ratio is 0.545, indicative of a less risky stock at a 10-year horizon, somewhat consistent with the implied PS continuous return volatility ratio (0.33).

At a 50-year horizon, the PS continuous variance ratio is 1.82 implying, from an extrapolated Figure 4a, the very large risk and *effective return variance ratio* of 1400 and from an extrapolated Figure 4b, implying the *effective return volatility ratio* of 5.37. The latter conflicts with the PS implied continuous volatility ratio of 0.19 and is indicative of a much riskier stock at a 50-year horizon, consistent with the PS continuous variance ratio indicator.

In summary, the PS model empirical estimates show conflicting interpretations of the horizon risk when it is measured by the continuously compounded nominal return variance versus the continuously compounded nominal return volatility ratio and our preferred effective return volatility ratio. Rather than increasing risk with the horizon, our implied and approximate effective return volatility ratio shows declining risk to a minimum at approximately 10 years and increasing thereafter to 1 at approximately 28 years. From an investor's perspective of the continuous return predictive volatility ratio, stocks are less risky over all horizons and from the perspective of the effective return predictive volatility ratio, stocks are less risky at quite long horizons but not at very long horizons.

4 Conclusions

In general, there are two important variables in the measurement of relative horizon risk, the risk measure and the return measure. We examine the horizon volatility (standard deviation) and variance return risks both measured in continuous nominal or effective returns. The behavior of the return horizon risk with horizon length depends critically on the definition of the average per period risk over the horizon and the process(es) generating the returns. From an investor's perspective, we conclude that the effective return volatility ratio is a much preferred metric of horizon risk in comparison to the continuous nominal return variance ratio.

From simple continuous versus effective iid return generating examples and third party empirical results from a well-known and well-designed study, we show conflicting and opposite results for the continuous return volatility and variance risk ratios. Using these empirical results, we also compare the study's continuous variance ratios with our extrapolated effective return volatility ratios. In comparisons with 1-year risks, we conclude that, from an investor's perspective of the effective return volatility ratio, stocks are less risky at short to long horizons but more risky at very long horizons.

Acknowledgements

My gratitude to Richard Roll and Harry Turtle for their helpful comments and to Edward Elyn for programming assistance.

Notes

1 See Kritzman (1994, 2015) for a summary and an early contrarian discussion, Bennyhoff (2004) for a practitioner summary of significant conflicting opinions, Chung *et al.* (2009) for a utility based treatment of the subject, and Pástor and Stambaugh (2012) for recent work.

2 See Barberis (2000) for an application to asset allocation.

3 See Lo and MacKinlay (1988).

4 Linton and Smetanina (2016) employ a mean ratio also using effective returns to test the martingale hypothesis.

5 The constant continuous mean assumption makes the effective return variance smaller than if the mean increased with the factor value. For example, if the mean increases by multiplication by the factor value 1.3, the effective return volatility ratio at a 10-year horizon increases from 0.573 to 0.652, which is also the minimum value. Therefore, the basic conclusion is not affected in this example.

6 To examine the predictive variance and the conditional expected returns, PS insightfully partition the predictive variance into five components; iid uncertainty, mean reversion, uncertainty about future expected returns, uncertainty about current expected return, and estimation risk. They conclude that the last three components, particularly uncertainty about future expected returns, more than offset the second component resulting in increased long horizon predictive variance.

7 If continuous nominal returns were iid and we had divided the continuous nominal volatility by $\sqrt{T}$ rather

than T, the resulting number would be the 1-year horizon nominal volatility and not the per year average volatility from a T-year hold. If continuous nominal returns are not iid, the interpretation of the resulting number is not clear.

References

Barberis, N. (2000). "Investing for the Long Run when Returns are Predictable," *Journal of Finance*, **55**, 225–264.

Bennyhoff, D. G. (2008). "Time Diversification and Horizon-Based Asset Allocations," *Vanguard Investment Counseling and Research*, 1–10.

Campbell, J. Y., Lo, A. W., and MacKinlay, A. C. (1997). *The Econometrics of Financial Markets*, Princeton University Press.

Chung, K. H., Smith, W. T., and Wu, T. L. (2009). "Time Diversification; Definitions and Some Closed-Form Solutions," *Journal of Banking and Finance* **33**, 1101–1111.

Kritzman, M. (1994). "What Practitioners Need to Know about Time Diversification," *Financial Analysts Journal* **50**, 14–18, reprinted 2015, *Financial Analysts Journal*, Jan/Feb Special Issue.

Linton, O. and Smetanina E. (2016). "Testing the Martingale Hypothesis for Gross Returns," *Journal of Empirical Finance* **38B**, 664–689.

Lo, A. W. and Mackinlay, A. C. (1988). "Stock Market Prices Do Not Follow Random Walks," *Review of Financial Studies* **1**, 41–66.

Pástor, Ľ. and Stambaugh, R. F. (2012). "Are Stocks Really Less Volatile in the Long Run?" *Journal of Finance* **67**, 431–477.

Keywords: Investment horizon risk; variance ratio; standard deviation ratio