Long-Horizon Stock Returns Are Positively Skewed

Adam Farago Erik Hjalmarsson*

April 16, 2021

Abstract

At long horizons, multiplicative compounding induces strong-to-extreme positive skewness into stock returns; the magnitude of the effect is primarily determined by single-period volatility. Consequently, at horizons greater than five years, returns –individual or portfolio– will be positively skewed under reasonable parametrizations. From an investor perspective, the strong positive skewness implies that the mean compound return will serve as a poor guide for typical long-horizon outcomes. Moreover, the large effects of compounding on higher-order moments are shown to affect the validity of Taylor expansions used to approximate preferences for skewness, when applied to returns of annual or longer horizons.

JEL classification: C58, G11, G12, G17.
Keywords: Compound returns; Long-run returns; Portfolio choice; Skewness.

*Both authors are at the Department of Economics and the Centre for Finance, University of Gothenburg. Contact information: P.O. Box 640, SE 405 30 Gothenburg, Sweden. Email: adam.farago@cff.gu.se and erik.hjalmarsson@economics.gu.se. We have benfitted from comments by Hendrik Bessembinder, Peter Bossaerts, Dominice Goodwin, Amit Goyal, Andrew Harvey, Oliver Linton, Andrew Patton, Lu Zhang, Filip Zikes, as well as by seminar participants at University of Cambridge, City University of Hong Kong, Second Swedish National Pension Fund (AP2), 17th Paris December Finance Meeting, and the Workshop on Financial Econometrics at Örebro University. The authors gratefully acknowledge financial support from Marianne and Marcus Wallenberg Foundation (grant number 2019.0117), Jan Wallander's and Tom Hedelius' Foundation and Tore Browaldh's Foundation (grant number P19-0117), Nasdaq Nordic Foundation, as well as from the Foundation for Economic Research in West Sweden.

1 Introduction

Individual investors typically have long investment horizons. In the 2019 Survey of Consumer Finances, almost 70% of households report to have an investment horizon that spans at least a few years, and 63% of these report a horizon longer than 5 years.[1] For certain saving purposes like retirement savings, the relevant horizon can easily span over several decades. Therefore, understanding the properties of returns over such horizons is important.

A long-run investor collects the total compound returns over the investment horizon. The moments of such long-run returns are non-trivial functions of the moments of short-run period-by-period returns. This is particularly true for higher-order moments such as skewness. Even under the simplest assumption of independently and identically (iid) distributed period returns, the single-period skewness does not scale with the investment horizon in a straightforward manner. We show in this paper that skewness becomes a characterizing feature of long-run return distributions. Therefore, understanding the asymmetry implied by this skewness is key to understanding the properties of long-horizon returns.

Empirically, Bessembinder (2018) shows that long-run compound stock returns behave very differently from short-run (monthly or annual) returns. Through simulation exercises, he illustrates how compounding induces strong *positive* skewness into multiperiod returns – even if the single-period returns are symmetric. While Bessembinder (2018) primarily focuses on individual stocks, his simulation results suggest that the skew-inducing effect of compounding should also be present in aggregate returns, albeit to a much lesser extent. In stark contrast, Neuberger and Payne (2021) argue that long-run (aggregate) stock returns are substantially negatively skewed.

In this paper, we aim to shed more light on the properties of long-run compound returns. We start by deriving a theoretical formula for the skewness (and other higher-order moments) of compound returns in an iid setting, given the moments of one-period (i.e., short-run) returns. The clear conclusion from the theoretical results is that compounding inevitably leads to positively skewed long-run returns. The strength of the skew-inducing effect of compounding depends primarily on the level of volatility in the single-period return – the higher the volatility, the stronger the effect – and is *not* qualitatively affected by potential asymmetries in the single-period return distribution. The skewness of 5-year market returns lies in the range of 1 to 3, where skewness values around 1 correspond to low volatility markets (17% yearly volatility), and skewness values around 3 are representative of high

[1] The survey asks the following question: "In planning or budgeting your saving and spending, which of the time periods is most important to you?" with the choices being: "1. Next few months, 2. Next year, 3. Next few years, 4. Next 5-10 years, 5. Longer than 10 years." The data are available at https://www.federalreserve.gov/econres/scfindex.htm.

volatility markets (30% yearly volatility). If 30-year market returns are considered instead, the corresponding skewness values end up in a range from about 5 (for low-volatility markets) to 30 (for higher volatility markets). Individual stocks are more volatile than aggregate markets, and therefore the skew-inducing effect of compounding is considerably stronger in the case of individual stock returns. The skewness of 5-year stock returns can easily exceed 10, and the skewness of 30-year stock returns can be in the order of *millions*. As the last result shows, compounding leads to extreme positive skewness in the long-run returns for individual stocks.

The above results are derived in an iid setting. Deviations from the iid assumption do have some quantitative effects on the skewness of long-horizon returns, but do not change the overall conclusions. We consider the cases of serially correlated single-period returns and conditional heteroskedasticity. Starting with serial correlation, our analytical results show that even large degrees of mean reversion (or mean aversion) in returns cannot affect the qualitative conclusions from the iid setting. In the case of conditional heteroskedasticity, analytical results appear elusive, but simulation results provide clear guidance. Time-varying volatility alone has little effect on long-run skewness compared to the iid case. When time-varying volatility is coupled with the leverage effect – where low (negative) returns predict higher future volatility – there is an effect. For returns with volatilities and leverage similar to that observed in U.S. market returns, the leverage effect has a noticeable negative impact on the skewness of long-horizon returns. However, the effect is not strong enough to change the qualitative conclusions from the iid setting. Compound 5-year returns are still clearly positively skewed, and at the 30-year horizon, using parameter values relevant for the U.S. market, compound returns from the leverage model have a skewness of about 4, compared to the iid case that has a skewness of about 5. To summarize, deviations from the iid assumption do have an effect on the skewness of long-horizon returns, but it is not strong enough to change the general conclusions from the iid setting.

The skewness of long-horizon returns impacts the estimation of this skewness as well. Joanes and Gill (1998) point out that for non-normal distributions, the skewness estimator is typically biased. Since long-horizon returns can be highly non-normal, assessing the properties of the estimator in this context becomes an important issue. We rely on simulations to assess the small sample properties of the skewness estimator, and derive its analytical asymptotic distribution to study the large sample properties. The results show that direct empirical estimation of skewness is mostly limited to annual or shorter horizons. Due to available sample sizes, skewness estimates will be biased for longer horizons both for market returns (where only time-series data are available) and for individual stock returns (where cross-sectional data are also available). For horizons longer than one year, sufficient sample sizes are thus typically only available in simulation settings. However, even simulations have to be carried out with caution for long-horizon returns. For low-volatility returns (e.g., market returns), simulated sample sizes in the order of millions might be needed, which is considerably larger than what is used in typical simulation studies. For individual stocks, it is virtually impossible to obtain useful simulation evidence at horizons greater than 15 years. In these cases, one can only rely on theoretical results.

The likely positive skewness of long-horizon returns has been noted in previous papers. Arditti and Levy (1975) show that compounding induces positive skewness into multiperiod returns, but they only analyze the iid case and do not recognize the dramatic long-run effects, since they only consider horizons up to 20 months. Bessembinder (2018) and Fama and French (2018) rely on simulations to study the distribution of long-run returns. However, given our results on the difficulties of estimating skewness, simulation-based evidence can potentially be misleading, especially from relatively small simulated samples, and our theoretical results provide an important foundation for understanding long-run returns. We offer several additional insights relative to previous literature, including detailed discussion on (i) the determinants of long-run skewness (i.e., the first order importance of single-period volatility), (ii) the problems of the skewness estimator in the context of long-run returns, (iii) a comparison to other measures of asymmetry, and (iv) some important implications for portfolio choice that, to our knowledge, has not been previously discussed.

While we focus on the standard (moment-based) skewness of simple returns throughout the paper, other measures have also been used to describe the asymmetry of return distributions. Kim and White (2004) and Ghysels et al. (2016) advocate the use of quantile-based measures, Campbell and Hentschel (1992) and Neuberger (2012) focus on the skewness of log returns, and Neuberger and Payne (2021) develop a new measure of return asymmetry. We derive analytical formulas of these measures for long-run returns in a setting where single-period returns are iid, and we rely on simulations to assess the effect of certain deviations from the iid assumption. We show that, when interpreted correctly, these alternative measures all lead to conclusions consistent with those from the standard skewness measure regarding the asymmetry of long-horizon return distributions.

Specifically, a clear distinction has to be made between the distributions of simple returns and log returns over long horizons. Since the investor collects the simple return from the investment, the log return is only useful to the extent that it provides valuable information about the corresponding simple return.[2] But for long horizons, the distributions of

[2] For very short horizons (i.e., daily and intra-daily returns), the simple and log returns correspond to each other closely, so the distribution of log returns will be informative about the simple returns. E.g., if daily log returns have a significant negative skew, the corresponding daily simple returns will also be negatively skewed.

simple and corresponding log returns have very different shapes. In particular, if simple returns are log-normal, then they are positively skewed, but the skewness of the corresponding log returns is always zero. Heuristically, the skewness of log returns therefore provides an interpretation *relative to log-normality.* Negative skewness of long-horizon log returns indicates that the corresponding simple returns are less skewed than a comparable log-normal distribution. However, those corresponding simple returns can still feature a considerable positive skewness, since log-normality implies a high (potentially extreme) positive skewness for long-horizon simple returns. We show that this is indeed the case. Some model specifications (e.g., stochastic volatility with a strong leverage effect) can lead to considerable negative skewness in even 10- or 20-year log returns, but the corresponding simple returns will feature a high positive skewness.

This also helps to resolve the seeming contradiction with Neuberger and Payne (2021), who argue that long-run stock market returns are substantially negatively skewed. The skewness measure developed by Neuberger and Payne (2021) is different from the skewness of simple returns or the skewness of log returns, but it also takes the value of zero in the case of log-normal simple returns (just as the skewness of log returns). When long-run market returns are negatively skewed in the Neuberger and Payne (2021) sense, this therefore implies that the returns are less skewed than if they were log-normally distributed, but they can in fact still be (strongly) positively skewed when the standard measure of skewness for simple returns is considered.

Our analysis thus clearly points in one direction: For longer horizons, compound (simple) returns are positively skewed. From an investor perspective, the most important immediate implication of this strong positive skewness is that the long-run mean return (i.e., the compounded mean single-period return) serves as an increasingly poor guide for the typical outcome of a long-horizon investment. To illustrate this point, consider the frequently quoted rule-of-thumb that if an asset delivers a 7% annual expected return, it takes 10 years to double the initial investment. This statement is valid in expected terms, but it does not take into account the shape of the long-run return distribution. If the asset has a 17% annual volatility (like the U.S. market), there is a 50% chance that the initial investment is doubled only after 13 years, and there is a 30% chance that it takes at least 20 years. If the annual volatility is 30% instead (typical for emerging markets), there is a 50% chance that it takes more than 22 years to double the initial investment, and with a 30% probability it takes more than 59 years, despite the fact that in expectation the investment doubles in 10 years.[3]

An understanding of the likely long-run outcomes of an investment therefore requires considerably more mental effort than the corresponding short-run exercise. Typical investment

[3]The numbers in this illustration are based on the assumption that monthly returns are iid log-normal.

advertisements or prospectuses will mention average annual returns and volatility. Transforming these into a meaningful characterization of the distribution of 10- or 30-year returns is clearly well beyond the capabilities of most investors. Ensthaler et al. (2018) illustrate how these problems play out in an experimental setting, and show that the participants almost completely fail to anticipate the asymmetry that arises in the long-run compound returns.

From a more formal perspective, the skew-inducing effect of compounding can also easily result in misleading conclusions. We consider the standard practice of using third- or fourth-order Taylor expansions of utility functions to (approximately) take into account the skewness (and kurtosis) preferences of a risk-averse investor (see, among others, Kraus and Litzenberger, 1976; Dittmar, 2002; Guidolin and Timmermann, 2008; Martellini and Ziemann, 2010; or Neuberger and Payne, 2021). As is evident from the Taylor expansion, a standard investor prefers more (positive) skewness to less.

Suppose a long-run investor uses the approximated expected utility over terminal wealth to rank different investment opportunities. Specifically, suppose there are two assets that both have iid period returns and that differ only in their single-period volatilities. Intuitively, a risk-averse investor should prefer the lower-volatility asset, regardless of investment horizon. However, since the skewness of long-run returns is strongly increasing in volatility, using the third-order utility approximation, a risk-averse investor might end up ranking the high-volatility asset over the low-volatility asset. That is, the strong skew-inducing effect of compounding makes the high-volatility asset seem more attractive to a long-run investor using the third-order Taylor approximation. We emphasize that this is indeed the "wrong" conclusion, since the exact utility calculations show that the long-run investor actually prefers the low-volatility asset, a result we prove formally.

A natural reaction to this finding is to argue that one should instead use a fourth-order expansion, to achieve a better approximation. In this case, preferences over kurtosis are also accounted for. This may help in some cases, but it does not provide a general solution. Kurtosis also grows with the investment horizon, and its increase will be affected by the lower-order moments, including skewness. In an analogous manner to what happens in the third-order expansion case, the fourth-order expansion can therefore also lead to the wrong conclusion. Since all higher-order moments increase with horizon, there is generally little hope to solve the approximation problem by adding further terms in the expansion.[4]

[4]These findings are related to the analysis by Kane (1982), who makes the observation that the portfolio allocation decision between a risky and a risk-free asset, based on a third-order Taylor approximation of the utility function, lacks an interior solution when the skewness of the risky asset exceeds a certain threshold. In this case, the solution using only the mean and variance tradeoff will be superior, and Kane attributes the failure of the third-order approximation to a neglect of higher order moments. While this is clearly formally

Importantly, these "approximation errors" do not only occur for very long horizons. We show that they can already occur at an annual investment horizon, with the implication that using Taylor approximations for utility functions based on annual return statistics can be misleading.

2 Skewness of compound returns

2.1 The case of iid returns

We start by deriving a general formula for the higher-order moments of compound returns, given the short-run (one-period) moments, and assuming that the period-returns are independently and identically distributed (iid). While the result applies to all higher-order moments, we focus the discussion on skewness.

Let x represent the *one-period gross return* on a given asset or portfolio. Throughout the paper, we will denote the expected value, standard deviation, and skewness of the *one-period return* as

$$\mu \equiv E[x] \quad , \qquad \sigma \equiv \sqrt{E\left[(x-\mu)^2\right]} \quad , \qquad s \equiv \frac{E\left[(x-\mu)^3\right]}{\sigma^3} \ . \tag{1}$$

Define the product process X_T as

$$X_T = x_1 \times x_2 \times ... \times x_T \ . \tag{2}$$

That is, X_T represents compound returns over T periods. If the x_t-s are assumed to be iid and have the same distribution as x, then it is straightforward that the k-th order (non-central) moment of X_T is

$$E\left[X_T^k\right] = E\left[x_1^k\right] \times E\left[x_2^k\right] \times ... \times E\left[x_T^k\right] = E\left[x^k\right]^T \ . \tag{3}$$

The mean and variance of X_T can easily be computed using (3) as

$$E[X_T] = \mu^T \qquad \text{and} \qquad Var(X_T) = \left(\mu^2+\sigma^2\right)^T - \mu^{2T} \ . \tag{4}$$

Proposition 1 *Let x and x_t, $t = 1, ..., T$, be iid random variables. Define the compound process $X_T = \prod_{t=1}^{T} x_t$. Denoting*

$$\theta_j \equiv \frac{E\left[x^j\right]}{E[x]^j} \ , \tag{5}$$

true, our results here suggest that it is not enough to simply include additional terms in the expansion when dealing with long-run returns.

the k-th order standardized moment of the compound process, for $k > 2$, is given by

$$\frac{E\left[\left(X_T - E\left[X_T\right]\right)^k\right]}{Var\left(X_T\right)^{k/2}} = \frac{\sum_{j=0}^{k} \binom{k}{j} (-1)^j \theta_{k-j}^T}{\left(\theta_2^T - 1\right)^{k/2}} . \quad (6)$$

Proof. See the proof in Appendix A.1. ■

With the help of Proposition 1, all the higher-order standardized moments of X_T can be obtained. Consequently, the skewness and kurtosis of compound returns are

$$Skew\left(X_T\right) = \frac{\theta_3^T - 3\theta_2^T + 2}{\left(\theta_2^T - 1\right)^{3/2}} \quad \text{and} \quad Kurt\left(X_T\right) = \frac{\theta_4^T - 4\theta_3^T + 6\theta_2^T - 3}{\left(\theta_2^T - 1\right)^2} , \quad (7)$$

where

$$\theta_2 = \frac{\sigma^2}{\mu^2} + 1 , \qquad \theta_3 = s\frac{\sigma^3}{\mu^3} + 3\frac{\sigma^2}{\mu^2} + 1 , \qquad \theta_4 = Kurt(x)\frac{\sigma^4}{\mu^4} + 4s\frac{\sigma^3}{\mu^3} + 6\frac{\sigma^2}{\mu^2} + 1 . \quad (8)$$

The formulas in equations (6) and (7) are fairly straightforward to derive. However, they imply extreme higher order moments for long-horizon returns, which, to our knowledge, has not been discussed in previous literature.[5]

In order to assess the quantitative implications of the above formulas, representative values for the single-period return moments have to be picked. We use $\mu = 1.01$, i.e., a 1% expected monthly return throughout the paper. Changing the value of μ to 1.008 or 1.012 (i.e., to a 0.8% or 1.2% monthly expected return) only marginally changes the results. We use $\sigma = 0.05$ and $\sigma = 0.08$ (i.e., 5% and 8% monthly volatilities) to represent well-diversified portfolios, and $\sigma = 0.17$ (i.e., 17% monthly volatility) to represent individual stocks.[6] $\sigma = 0.11$ (11% monthly volatility) represents an intermediate case between fully diversified market portfolios and individual stocks. The value of the single-period skewness

[5] Arditti and Levy (1975) derive a related result on the third moment of compound returns, although they consider the non-standardized moment rather than the actual skewness. Proposition 1 generalizes their result to all higher order (standardized) moments. Arditti and Levy (1975) note that compounding induces skewness, but their focus is on portfolio choice and they do not examine the long-run implications of compounding as they do not consider periods longer than 20 months.

[6] Ghysels et al. (2016) present summary statistics of the market returns for various countries using data between 1996 and 2013. The mean monthly return is 0.8% in the U.S. and on a value-weighted developed markets portfolio, while it is 1.1% on a value-weighted emerging markets portfolio. Bessembinder (2018) reports a mean of 1.1% for the pooled sample of monthly returns on U.S. common stocks from 1926 to 2016, while Oh and Wachter (2019) report 1.3% for the 1945-2016 subsample. Ghysels et al. (2016) report the volatility of the monthly returns to be 5.7% in the U.S., 6.2% on the developed market portfolio, and 9.0% on the emerging market portfolio. Finally, Bessembinder (2018) reports a monthly 18.1% volatility for his pooled sample of individual U.S. stocks, while Oh and Wachter (2019) report 16.8% for the 1945-2016 subsample.

is varied between $s = -2$ and $s = 2$ to illustrate its effect on skewness in long-run returns.

Table 1 shows the skewness of X_T, calculated via equation (7), when the single-period returns correspond to monthly returns with $\mu = 1.01$ and volatility that varies across the columns of the table. The return horizon, T, varies between 1 month and 30 years. Panel A shows the skewness of compound returns, when the single-period returns are symmetric (zero-skew). Several results are worth noting. First, compound returns are positively skewed, and their skewness increases non-linearly with the horizon. That is, compounding induces skewness in long-horizon returns even if single-period returns are symmetric (as previously also noted by Arditti and Levy, 1975, and Bessembinder, 2018). Second, skewness increases dramatically and highly non-linearly in σ, for a given T. In other words, the single-period volatility has a huge effect on the degree of skewness induced by compounding. If the volatility of the monthly returns is $\sigma = 0.05$, which corresponds to a well-diversified portfolio, the effect of compounding is clear, although not extreme: The skewness of the 30-year returns is 5.19. On the other hand, for $\sigma \geq 0.11$ the skewness induced by compounding increases very rapidly with the horizon. This leads to our third observation: For large T and σ, the skewness values are extreme. For example, the skewness of 30-year returns when $\sigma \geq 0.17$ is in the order of millions. Finally, it is worth highlighting that the results in Panel A hold for any symmetric distribution. They are equally valid if one-period returns are, for example, normally or uniformly distributed.

The rest of Table 1 helps us understand the effect of single-period skewness. Panel B corresponds to the case where monthly returns have a skewness equal to that of a log-normal distribution.[7] Panels C and D represent cases with more greatly skewed one-period returns, with $s = Skew(x)$ being 2 and -2, respectively. Our main observation is that the effect of single-period skewness depends on the level of the single-period volatility. When σ is low, single-period skewness does not have a large effect on the skewness of long-horizon returns (up to a 30-year horizon). Take the column with $\sigma = 0.05$; the skewness of the 30-year returns is 4.50 when $s = -2$, and it is 5.95 for $s = 2$. That is, the difference in skewness at the 30-year horizon is much lower than at the monthly level. In fact, it is striking how quickly the negative skewness is reduced in magnitude through compounding. For $\sigma = 0.05$, the skewness goes from -2 at the monthly horizon to -0.13 at the annual horizon. When single-period volatility is high, single-period skewness can have a large effect on the skewness of compound returns, especially at long horizons. For example, if $\sigma = 0.17$, the skewness of 30-year returns is of the order of 10^4 when $s = -2$, and of the order of 10^7 for $s = 2$. However, large absolute differences between the corresponding cells of different Panels in Table 1 only

[7]The log-normal distribution does not have an explicit skewness parameter, but its skewness is a function of the mean and variance of the distribution. Specifically, $s = \frac{\sigma}{\mu}\left(\frac{\sigma^2}{\mu^2} + 3\right)$.

occur when the values in Panel A (non-skewed single-period returns) are already extreme. In these cases, it is hard to give an interpretation to the differences in the extreme skewness levels (i.e., it is difficult to interpret the difference between the skewness levels of 10^4 and 10^7).

Figure 1 provides a graphical illustration of the results in Table 1 by plotting the skewness of compound returns as a function of horizon. Single-period volatility, σ, is varied across the panels, while differing single-period skewness, s, is represented by different lines. Panel A clearly illustrates that for low single-period volatility, the skewness in long-horizon compound returns is almost identical regardless of inherent skewness in the single-period returns. As the volatility of the single-period returns increases (through Panels B-D), the skewness in compound returns can easily reach extreme values. However, for a given volatility, the cases with $s = 0$ and $s = \pm 2$ result in qualitatively similar patterns. To that extent, it is the volatility of the single-period returns, and not their skewness, which is of first order importance for the skewness of compound returns. In other words, the patterns in $Skew\left(X_T\right)$ are more similar within the panels of Figure 1 (where single-period skewness is varied), than they are across the panels (where single-period volatility is varied).[8]

2.2 Theory versus estimation

Bessembinder (2018) reports the skewness of the *90-year* compound return on a bootstrapped U.S. single-stock position to be around 100. At the same time, Table 1 shows that the theoretical skewness of the *30-year* compound return is already in the order of millions if the single-period volatility is $\sigma = 0.17$ (which is similar to the monthly volatility of the bootstrapped single-stock position). There is a large discrepancy between the estimation results and theory. As we show in this section, the discrepancy arises because the common skewness estimator is not adequate for measuring the skewness of long-horizon returns, especially in the case of individual stocks.

The commonly used estimator of skewness is

$$\hat{S} \equiv \frac{\frac{1}{n}\sum_{i=1}^{n}\left(z_i - \bar{z}\right)^3}{\left(\frac{1}{n}\sum_{i=1}^{n}\left(z_i - \bar{z}\right)^2\right)^{\frac{3}{2}}}\ , \tag{9}$$

where z_i, $i = 1, ..., n$ is a sample (of size n) corresponding to a general random variable

[8]Bessembinder et al. (2019) provide empirical support to these theoretical predictions. They show in a cross-country analysis that the standard deviation of monthly individual stock returns is negatively associated with the proportion of stocks that outperform the risk-free investment in the long run (which can be viewed as a measure of asymmetry), while the skewness of monthly individual stock returns does not have significant explanatory power.

Z, and $\bar{z}$ is the sample average. For non-normal distributions, $\hat{S}$ is typically biased, but theoretical expressions for the bias are generally not available (Joanes and Gill, 1998). Since long-horizon returns can be highly non-normal, evidenced by the high positive skewness shown in Figure 1, assessing the size of the bias becomes a relevant issue.

We start with a Monte Carlo simulation to show the finite sample distribution of the estimator, but later we also derive its asymptotic distribution. For the simulation, we assume that single-period returns are log-normal. For a given sample size n, we simulate n independent realizations of the long-run compound return X_T and estimate the skewness of the sample using the estimator $\hat{S}$. We repeat this step 100,000 times to get the distribution of the skewness estimates. Figure 2 summarizes these distributions for different monthly volatilities (σ) and compounding horizons (T), with the sample size n ranging from 50 to 250,000; the expected monthly return is always equal to 1% ($\mu = 1.01$). Specifically, each panel in the figure shows the range between the 5^{th} and the 95^{th} percentiles, the median (square marker), and the mean value (diamond marker) of the skewness estimates $\hat{S}$. The true skewness value, that $\hat{S}$ tries to estimate, is indicated by the dotted vertical line in each graph.

The graphs in the top row of Figure 2 use 8% for the single-period volatility, representing well-diversified portfolios and market indexes (the results are qualitatively similar if 5% single-period volatility is used). Panel A presents results for annual returns ($T = 12$). If one considers a portfolio or a market index, one might have somewhere between 50 and 100 years of time-series data, and therefore the sample sizes of $n = 50$ and $n = 100$ correspond to directly estimating the skewness from annual returns. As seen from Panel A of Figure 2, there is a clear downward bias in $\hat{S}$ when $n = 50$. For $n = 100$, $\hat{S}$ is closer to being unbiased, but the precision of the estimator is still not great. The precision improves quickly as n increases, but the larger sample sizes are only available in simulations. Moving to Panel B for 5-year returns ($T = 60$), $n = 50$ is already empirically out of reach for direct estimation of skewness, and $\hat{S}$ is a biased estimator for $n \leq 100$. However, unbiasedness and a high precision can be achieved in simulation settings.

The 20-year horizon reveals the real limitations of the skewness estimator. Wilkins (1944) shows that there is an upper limit to the absolute value of $\hat{S}$, which depends solely on the sample size:

$$\left|\hat{S}\right| \leq \frac{n-2}{\sqrt{n-1}} \approx \sqrt{n}\ . \tag{10}$$

This is an exact bound that holds in any sample: It is *numerically impossible* for the skewness estimator $\hat{S}$ to exceed $\sqrt{n}$. The upper bound is indicated by the cross marker in the graphs of Figure 2. At $n = 50$ and $n = 100$ in Panel C, the upper limit is lower than the true value of skewness (which equals 12 in Panel C), implying that all the sample estimates will be lower than the true parameter value. As the sample size increases, the upper bound from equation (10) is not binding anymore, but $\hat{S}$ remains downward biased even at $n = 25,000$. At $n = 250,000$, $\hat{S}$ is close to being unbiased, but its distribution is right-skewed and has a wide dispersion around the true skewness value. Therefore, an even higher sample size (e.g., $n \geq 10^6$) is needed to get a precise estimate of the skewness for 20-year returns in this example. Obviously, these issues become even worse as the horizon further increases.

The graphs in the bottom row of Figure 2 use a monthly volatility of 17% to represent the case of individual stocks. Start again with looking at annual returns ($T = 12$). In a purely empirical situation, using a single cross-section of annual returns on individual stocks, one might face a sample size around $n = 2,500$. As is seen in Panel D of Figure 2, the skewness estimator is almost unbiased and reasonably precise in this case. For longer horizons, much larger sample sizes are needed. In the case of 5-year returns, the skewness estimator is severely downward biased for $n = 2,500$, and it becomes unbiased only for $n > 250,000$. If one pools individual stock returns from many years, considerably larger sample sizes than $n = 2,500$ can be achieved. In the CRSP data set for U.S. stock returns, using observations from the entire 1926 to 2019 period, one can construct close to $300,000$ independent annual stock returns and $50,000$ independent 5-year returns. However, even if one ignores the issue that a sample spanning so many stocks and such a great time period is likely to suffer from large heterogeneity, and thus a less clear interpretation of the estimates, the available sample size at the 5-year horizon is still not enough to ensure a reliable skewness estimate. The required sample sizes are only available in simulation settings.

Estimating the skewness of 20-year (and longer) returns on individual stocks proves to be essentially impossible. The true skewness of the 20-year returns in Panel F of Figure 2 is 23,366. The upper bound on $\hat{S}$ is therefore highly restricting for all the presented sample sizes, even for $n = 250,000$ (where the upper bound is 500). The skewness estimates are highly misleading because of an extreme downward bias in the estimator. Equation (10) implies that a sample size above 5.5×10^8 would be needed just for the upper limit on $\hat{S}$ not to be binding in this case. However, simply exceeding the required minimum sample size from equation (10) is not enough, as suggested by Panel C of Figure 2, and a sample considerably larger than $n = 5.5 \times 10^8$ is needed to obtain useful estimates.

The above results imply that extremely large sample sizes are needed to avoid a biased estimator when estimating the skewness of long-horizon individual stock returns. However, even if the sample is big enough to avoid a substantial bias, the standard error of the estimator can still be so large that the skewness estimates are next to useless. To assess this concern, we derive the asymptotic distribution of the skewness estimator.

Proposition 2 *The asymptotic distribution of the skewness estimator $\hat{S}$, as $n \to \infty$, is given by*

$$\sqrt{n}\left(\hat{S} - Skew(Z)\right) \xrightarrow{d} N\left(0, \frac{\mu_6 - 6\mu_2\mu_4 + 9\mu_2^3 - \mu_3^2}{\mu_2^3} - \frac{3\mu_3\left(\mu_5 - 4\mu_2\mu_3\right)}{\mu_2^4} + \frac{9\mu_3^2\left(\mu_4 - \mu_2^2\right)}{4\mu_2^5}\right), \tag{11}$$

where $\mu_k \equiv E\left[(Z - E[Z])^k\right]$.

Proof. See the proof in Appendix A.2. ■

That is, $\hat{S}$ is consistent and asymptotically normally distributed, with an asymptotic variance that is a function of the central moments up to order six. In the case of compound returns, the μ_k values for $k \geq 3$ can be obtained from the single-period moments via Proposition 1. Using the asymptotic variance from equation (11), Figure 3 shows two-standard error bounds of the skewness estimator around the true parameter value. The monthly returns are log-normally distributed with a 17% volatility (representing individual stocks), and the horizon over which the returns are compounded changes along the horizontal axis.[9] The bounds are shown for two different sample sizes, $n = 10^{15}$ and $n = 10^{20}$, both large enough that the upper bound on $\hat{S}$ is not binding even for the 30-year ($T = 360$) returns. As is evident from Figure 3, the error bounds on the skewness estimator widen very quickly after a certain horizon, such that the estimates essentially become uninformative due to the large standard errors. Even with a simulated sample size of $n = 10^{20}$ observations, skewness estimates are uninformative for return horizons over 15 years.

To summarize, direct empirical estimation of skewness is essentially limited to annual or shorter horizons, both for market indexes (where time-series data are available), and for individual stocks (where cross-sectional data are also available); if one pushes the available data to the limit and pool individual stock returns over very long time periods (i.e., the last 100 years), estimation at somewhat longer horizons (e.g., 2 or 3 years) might sometimes be possible, but the 5-year horizon is already infeasible. For longer horizons, we have to rely on simulations to achieve adequately-sized samples. However, even simulations have to be carried out with caution when studying the skewness of long-horizon returns. For aggregate returns ($\sigma \leq 0.08$), simulated sample sizes in the order of millions might be needed, which is larger than what is used in typical simulation studies. For individual stocks ($\sigma \geq 0.17$), it is essentially impossible to obtain useful simulation evidence at horizons greater than 15

[9] Log-normal single-period returns are used for the values in Figure 3, but the asymptotic result can be applied to any distribution where the first six central moments are known. The conclusions are qualitatively the same if we use the normal distribution or a more skewed distribution for the single-period returns. The single-period expected return is $\mu = 1.01$.

years. In these cases, one can only rely on theoretical results.[10]

2.3 Relaxing the iid assumption

The result in Proposition 1 was derived under an iid assumption for returns. Relaxing the "identical" part is relatively straightforward. The factoring of the moments according to the first equality in equation (3) continues to hold, but the last equality is no longer valid since the moments for the individual returns change over time. That is, $E\left[X_T^k\right] = E\left[x_1^k\right] \times ... \times E\left[x_T^k\right] \neq E\left[x_1^k\right]^T$. Simple formulas such as those expressed in equation (7) are therefore no longer attainable, but for a given set of time-varying moments, the skewness or kurtosis of the compound returns can relatively easily be calculated. Clearly, the moments of the compound returns will reflect a form of average moments for the period returns, and the insights from unconditional heteroskedasticy are limited.

Of more interest is therefore the relaxation of the "independence" assumption, while maintaining that the unconditional distribution remains the same over time. Allowing for some form of dependence across the individual returns immediately implies that $E\left[X_T^k\right] \neq E\left[x_1^k\right] \times ... \times E\left[x_T^k\right]$ for at least some k, and one needs to take into consideration the joint distribution of $x_1, ..., x_T$. Such a setting quickly becomes non-tractable, given the difficulties of analyzing products of dependent random variables. Below, we first discuss the case where period returns are serially correlated, where a (heuristic) analytical result is presented. We then discuss the effects of other types of dependence.

2.3.1 Serial correlation

Since compounding involves multiplication rather than summation, the exact effects of serial dependence on the compound returns is extremely difficult to derive. Therefore, we rely on a heuristic approximation based on the log-normal case. Using the skewness result in equation (7) under the assumption that period returns are iid log-normal implies

$$Skew\left(X_T\right) = \left(\left(1+\frac{\sigma^2}{\mu^2}\right)^T + 2\right)\left(\left(1+\frac{\sigma^2}{\mu^2}\right)^T - 1\right)^{\frac{1}{2}}. \qquad (12)$$

[10]The limits of what can be achieved with simulations clearly depend on computing power. As seen from Figure 3, one would need simulated sample sizes exceeding 10^{20} to achieve reliable results for individual stocks at the 15-year horizon. In the simulation results reported below in Table 2, we use a sample size of 10^7. While such a sample size is many times larger than those typically used in simulation studies, it still restricts us to consider horizons of 10 years or shorter for $\sigma = 0.08$, and to horizons of 2 years or shorter for $\sigma = 0.17$; see Foonote 14 for additional details.

Our result for serially correlated returns is stated as a corresponding (approximate) expression of the (exact) iid result in equation (12):

Proposition 3 *Suppose x and x_t, $t = 1, ..., T$ are log-normally distributed random variables with mean μ and variance σ^2, and let $X_T = \prod_{t=1}^{T} x_t$ be the corresponding compound returns. Further, let the variance ratio VR denote the ratio between the long-run and short-run variance of x_t, such that*

$$VR \equiv \frac{LR.Var\,(x)}{Var\,(x)} = \frac{\sum_{j=-\infty}^{\infty} Cov\,(x_t, x_{t+j})}{Var\,(x_t)}. \tag{13}$$

The skewness of X_T in the case when the x_t-s are serially correlated can be approximated as

$$Skew\,(X_T) \approx \left(\left(1 + \frac{VR \times \sigma^2}{\mu^2}\right)^T + 2\right)\left(\left(1 + \frac{VR \times \sigma^2}{\mu^2}\right)^T - 1\right)^{\frac{1}{2}}. \tag{14}$$

Proof. See the proof in Appendix A.3. ■

As is evident from equation (14), the effect of serial correlation is captured by scaling the variance of the single-period returns with the long-run to short-run variance ratio. Specifically, Proposition 3 implies that if single-period returns have mean μ, volatility σ, and are *serially correlated*, the skewness of the resulting compound returns behaves as if the single-period returns were *iid* with mean μ and volatility $\sqrt{VR} \times \sigma$.[11] There is a large literature suggesting that returns are mean-reverting over longer horizons, which implies that $VR < 1$ (see, for instance, Fama and French, 1988, Poterba and Summers, 1988, Cecchetti et al., 1990, Cutler et al., 1991, Siegel, 2008, and Spierdijk et al., 2012).[12] Thus, for example, if the non-iid single period returns have $\sigma = 0.17$ and $VR = 0.8$, the resulting $Skew\,(X_t)$ can be well approximated by the iid formula with volatility parameter equal to $\sqrt{0.8} \times 0.17 \approx 0.152$.

Figure 4 illustrates the effects of serial dependence on the skewness of compound returns. The effects of $VR = 0.9$ and $VR = 0.8$ are compared to the benchmark iid case within each panel, and the volatility of the single-period returns is varied across the panels. The conclusions are similar to those obtained when looking at the effect of single-period skewness. When σ is low, the effect of serial dependence on long-horizon skewness is small. When σ is

[11] Alternatively, one can think of this as replacing the standard (short-run) variance, σ^2, with the long-run variance, $VR \times \sigma^2$, in the skewness expression in equation (12). That is, under serial correlation, the appropriate measure of the variance in the process is the long-run variance, and it is this quantity that should be used when inferring the skewness in long-run compound returns.

[12] The presence of mean reversion in stock returns is not universally accepted, however, and other studies argue against it; for instance, Richardson and Stock (1989), Kim, Nelson, and Startz (1991), and Richardson (1993).

high, the effect of serial dependence can be sizable, but only in the range of extreme skewness levels, where interpretation of the different skewness values is not straightforward any more. To that extent, the effect of serial dependence is of second order importance compared to the effect of single-period return volatility.

2.3.2 Non-linear dependencies

The other common model deviation from independent returns is some form of conditional heteroskedasticity, where the first moment of returns are serially uncorrelated, but where there is higher-order (non-linear) dependence. GARCH and stochastic volatility processes, including those with so-called "leverage" effects, are typical examples of such models. Given their non-linear nature, such processes immediately become very difficult to analyze in a multiplicative compounding framework. In this section we present simulation results for such processes, and show that there are quantitative effects of introducing stochastic volatility. But qualitatively, long-run compound returns behave very similarly to the iid case.

Neuberger and Payne (2021) stress the importance of the "leverage" effect on the skewness of long-run returns. The leverage effect refers to the often observed negative correlation between past returns and current volatility. The extent of this effect varies across assets and appears stronger in index returns than in individual stock returns (e.g. Kim and Kon, 1994, and Andersen et al., 2001).

We consider a simulation experiment where data are generated by a continuous-time stochastic volatility (SV) model. Specifically, we use the Heston (1993) model of stochastic volatility, which is the same as the one used by Neuberger and Payne (2021) in their simulations.[13] As discussed at length in Section 2.2, simulation results can be misleading at long horizons and for large volatilities. Therefore, we restrict the results to the "safer" regions of the parameter space, i.e., to shorter horizons for higher monthly (unconditional) volatilities.[14]

[13]Jumps in returns are a common feature in continuous-time models and jumps that have a negative mean contribute to negative skewness in returns. However, since jumps are typically assumed independent they have little effect on long-run compound returns. In non-reported simulations, we find that at the five-year horizon the specifications with and without jumps deliver almost identical results. We therefore focus the analysis on specifications without jumps.

[14]Specifically, we first rule out all cases where the upper bound on the skewness estimator in equation (10) is (close to) binding. Second, we rule out cases where the asymptotic standard errors (under an iid log-normal assumption) would be "large". There is always a return horizon after which there is a substantial discrete jump in the standard error, and after which the skewness estimator thus becomes very imprecise; in all cases we report, the asymptotic standard errors are less than 0.15. And third, we compare the differences between the theoretical skewness values and the estimated skewness values from simulated data in the iid log-normal case. In this case, the theoretical skewness values are known (and we report these in Table 2), but as a final check we verify that the estimated values are close to the theoretical ones. For the reported return horizons in Table 2, this difference never exceeds 0.15 in absolute value.

The details of the simulation design are described in Appendix B. In brief, our parametrization are similar to those usually found in the literature, with the exception that we set the average volatility equal to 5%, 8%, 11%, or 17% on a monthly frequency, and the expected monthly return to 1%, so that the results are comparable to those in previous sections. The volatility of volatility, as well as the mean reversion parameter for the SV process are specified in the Appendix and are close to the typical values reported empirically. We consider three different values for the correlation between the return and volatility innovations: $0, -0.5$, and -0.75. We refer to these as the no-leverage (SV-NL), standard-leverage (SV), and high-leverage (SV-HL) cases, respectively. A value around -0.5 (SV) is often seen for U.S. stock index data (e.g., Andersen et al., 2002,. Eraker et al., 2003, Eraker, 2004), and is likely the empirically most relevant value for the low-volatility specifications that aim to capture the behavior of index returns. The high-leverage specification (SV-HL) is counterfactual to most empirical evidence and is intended to show the limits of the effects of leverage in our analysis. The no-leverage case (SV-NL) serves two purposes. First, the leverage effect is quite limited in individual stock returns (e.g., Andersen et al., 2001), so this is the empirically most relevant value for the high-volatility specifications. Second, by cutting off the dependence between the return and volatility innovations, one can examine the impact of time-varying volatility alone, without a leverage effect. The simulation results are based on 10^7 simulated returns for any given horizon.

The first four columns of Table 2 show estimates of the skewness in the simulated data, along with theoretical values for the iid log-normal case (labeled LN), which can be viewed as a special case of the SV model with constant volatility. The volatility σ now represents the average volatility in the SV specifications. Starting with the results for $\sigma = 0.05$, it is clear that stochastic volatility with leverage induces negative skewness at shorter horizons. At the monthly horizon ($T = 1$) the iid log-normal returns have a positive skewness of 0.15, whereas the SV returns with standard leverage have a skewness of -0.30. In the high-leverage case, the skewness increases in absolute value to -0.57. When the leverage effect is turned off (SV-NL), there is no negative skewness effect from time-varying volatility.

At the annual horizon ($T = 12$), there is still some negative skewness left in the specifications with leverage, but it has decreased in magnitude relative to the monthly horizon. At horizons longer than two years, the skewness is positive in all cases. As the average volatility is increased, the positive-skewness effect from compounding starts to dominate quicker: For $\sigma = 0.08$, all specifications have positive skewness already at the annual horizon. For $\sigma = 0.11$ and 0.17, we only report skewness estimates for shorter horizons, given the discussion in Section 2.2 on the bias and variance in the skewness estimator. However, even at shorter horizons, it is clear that the positive skewness from compounding dominates the negative skewness induced by the leverage effect.

The results in Table 2 strongly suggest that stochastic volatility and leverage effects do not overturn the main conclusions drawn from the iid analysis presented in Section 2. Time-varying volatility has very little impact if the leverage channel is turned off. Time-varying volatility coupled with leverage does have a negative effect on skewness. At short horizons, and for low levels of (average) volatility, this results in an actual negative skewness for compound returns. However, at horizons longer than 2 years, the positive skewness induced by compounding seems to inevitably dominate the negative skewness from leverage. Thus, while stochastic volatility with leverage can reduce the skewness in long-run compound returns, compounding still has a very similar effect to that in the iid case.

3 Other measures of asymmetry

The moment-based $Skew(X_T)$ is not the only measure that has been used in the literature to describe the asymmetry of return distributions. Kim and White (2004) and Ghysels et al. (2016) advocate the use of quantile-based measures of asymmetry because they are more robust to the presence of outliers. Hinkley's (1975) quantile-based coefficient of asymmetry is defined, for any $\alpha \in (0.5, 1)$, as

$$QSkew_\alpha (X_T) \equiv \frac{Q_\alpha (X_T) + Q_{1-\alpha} (X_T) - 2Q_{0.5} (X_T)}{Q_\alpha (X_T) - Q_{1-\alpha} (X_T)} , \tag{15}$$

where Q_α, $Q_{0.5}$, and $Q_{1-\alpha}$ denote the α quantile, median, and $1-\alpha$ quantile of the return distribution, respectively. For any symmetric distribution, $QSkew_\alpha (X_T) = 0$, and the measure is bounded with the maximum value of 1 representing extreme right asymmetry (or, right-tailed distribution) and the minimum value of -1 indicating extreme left asymmetry. Bowley's (1920) measure of asymmetry is obtained as a special case with $\alpha = 0.75$ (i.e., using quartiles). $QSkew_\alpha$ depends on the value of α, and it is not clear what value should be used. Therefore, Groeneveld and Meeden (1984) constructed an alternative by integrating over $0.5 < \alpha < 1$:

$$QSkew (X_T) \equiv \frac{\int_{0.5}^{1} \{Q_\alpha (X_T) + Q_{1-\alpha} (X_T) - 2Q_{0.5} (X_T)\} \, d\alpha}{\int_{0.5}^{1} \{Q_\alpha (X_T) - Q_{1-\alpha} (X_T)\} \, d\alpha} . \tag{16}$$

This measure is also zero for any symmetric distribution and is bounded by -1 and 1.

Other studies (e.g., Campbell and Hentschel, 1992; and Neuberger, 2012) are concerned with the skewness of *log* returns, $Skew(\log(X_T))$, which we will subsequently refer to as *logSkew,* for short. As we will highlight in this section, it is crucial to clearly distinguish between $Skew(X_T)$ and $Skew(\log(X_T))$ when discussing the "skewness of returns," especially for long horizons. More recently, Neuberger and Payne (2021) have also proposed a new measure of return asymmetry, which is closely related to the skewness of log returns:

$$NPSkew\left(X_T\right) \equiv \frac{E\left[6\left(\left(e^{\tilde{X}_T}+1\right)\tilde{X}_T - 2\left(e^{\tilde{X}_T}-1\right)\right)\right]}{E\left[2\left(e^{\tilde{X}_T}-1-\tilde{X}_T\right)\right]^{3/2}}, \tag{17}$$

with $\tilde{X}_T \equiv \log(X_T) - \tilde{m}T$, where $\tilde{m}$ is an appropriately defined constant (see the definition in Appendix C). $NPSkew$ is constructed such that it equals zero for the log-normal distribution (Neuberger and Payne, 2021).

In this section, we study the asymmetry of the distribution of long-horizon returns using *QSkew*, *logSkew*, and *NPSkew*.[15] In order to analytically calculate these measures, we have to specify the return distribution. Similar to previous results in the paper, our benchmark is the case of log-normally distributed single-period (monthly) returns. Appendix C.1 provides analytical formulas for all the asymmetry measures discussed above, in the case when single-period returns are iid log-normal. However, we also require a more flexible distribution to study how more asymmetric single-period returns affect the shape of the long-horizon return distribution.

The Normal Inverse Gaussian (NIG) distribution, introduced by Barndorff-Nielsen (1997), is a four-parameter distribution allowing for non-zero skewness and fat tails. In our alternative parametrization, we assume that the single-period returns, x_t, have a log-NIG distribution (i.e., that $\log(x_t)$ follow a NIG distribution) and are iid across time. The log-NIG distribution has several useful features:

- the log-normal distribution is nested as a special case,
- the parameters of the distribution can be set to match a certain combination of the first four moments (mean, variance, skewness, and kurtosis),
- if the single-period returns are iid log-NIG, the corresponding compound returns are also log-NIG, and
- we are able to derive analytical formulas for all the measures of asymmetry discussed above.

[15]Unreported results show that looking at $QSkew_\alpha$, with different α levels, leads to similar conclusions to those obtained from the integrated $QSkew$ measure. In order to save space, we therefore omit the results using $QSkew_\alpha$.

We provide the formulas and the details of the above statements in Appendix C.2.

Figure 5 shows *QSkew*, *logSkew*, and *NPSkew* as a function of horizon, when single-period returns are iid log-NIG. Three levels of monthly volatility (5%, 8%, and 17% across the panels) and three monthly skewness levels (-2, 2, and the skewness implied by log-normality, within each graph) are used; note that these skewness levels refer to the skewness of the simple (non-log) returns.[16] Let us start with the quantile-based measure shown in the panels in the first column. First, $QSkew$ increases with the horizon in all cases. As a consequence, long-run returns display positive asymmetry: For horizons longer than 2 years, all $QSkew$ values are positive, even if the monthly returns are negatively skewed. Second, single-period volatility has a much larger effect than single-period skewness on $QSkew$ in the long-run, i.e., $QSkew$ values for long-horizons vary much more across the graphs than within them. Third, long-run individual stock returns display extreme asymmetry: the value of $QSkew$ gets very close to its upper bound of one in Panel G for horizons longer than 5 years. All in all, studying $QSkew$ leads to the exact same qualitative conclusions as the standard measure of skewness, without providing much additional insight.

Let us now turn to *logSkew* and *NPSkew*. As is clear from Figure 5, the two measures give almost identical values. Neuberger and Payne (2021) also point out, based on their empirical and simulation results, that *NPSkew* generally lies rather close to *logSkew*. Our analytical results confirm this within the iid log-NIG setup. Both measures take the value of zero for the log-normal distribution and, heuristically, the signs of *logSkew* and *NPSkew* therefore provide an interpretation *relative to log-normality*; negative values indicate that the distribution features less right asymmetry than the log-normal distribution.[17] However, because the reference point – the skewness of compound (non-log) returns – is changing with the return horizon, it is hard to properly interpret specific values of these measures. If we consider the case where the single-period returns are negatively skewed with $Skew(x) = -2$, both the $Skew(X_T)$ values from Figure 1 and the *QSkew* values from Figure 5 show that for horizons longer than 2 years, the compound return distribution features right asymmetry (i.e., a longer right tail), despite the *logSkew* and *NPSkew* values being negative for all return

[16] To pin down all the parameters of the log-NIG distribution, the kurtosis of the simple returns also needs to be specified. When skewness is implied by log-normality, the corresponding kurtosis of the matched log-normal distribution is used. In the case with the skewness set to -2 or 2, the kurtosis is set to 15. While the log-NIG distribution allows us to match the first four moments, not all moment combinations are feasible and some combinations will result in fairly extreme parameter combinations. We pick a kurtosis of 15 because it leads to reasonable parameters values for all different levels of volatility considered. The results for horizons longer than two years in Figure 5 are barely affected by the choice of kurtosis.

[17] This is an heuristic interpretation. There may exist distributions with skewness less (greater) than the log-normal, but for which *logSkew* is positive (negative). However, in terms of trying to relate skewness of simple returns to skewness of log returns, benchmarking the latter against the log-normal distribution is likely a reasonable guide in most cases.

horizons. Finally, Figure 5 also shows that both *logSkew* and *NPSkew* converge to zero as the return-horizon increases, which is expected from central limit convergence.

The results in Figure 5 correspond to iid returns. Neuberger and Payne (2021) emphasize the importance of the effect of time-varying volatility with leverage on the skewness of long-run returns. Therefore, in Table 2 we also present simulation based estimates of the asymmetry measures corresponding to the stochastic volatility specifications detailed in Section 2.3.2. Our earlier conclusions are reinforced. $QSkew$ values show that stochastic volatility with leverage leads to somewhat less asymmetry in long-horizon returns compared to the iid case, but the differences are small, and long-run compound returns are positively skewed.

The *logSkew* and *NPSkew* values in Table 2 are very close to each other, and both measures are negative for the stochastic volatility with leverage specifications; i.e, they lead to less pronounced right asymmetry than a comparable iid log-normal specification. However, it needs to be emphasized again that negative *logSkew* and *NPSkew* values do not imply negatively skewed simple (non-log) compound returns. For example, 5-year returns from the SV specification with $\sigma = 0.05$ feature $Skew(\log(X_T)) = -0.4$ and $NPSkew(X_T) = -0.35$, but the distribution of these returns clearly features right asymmetry (positive skewness) as shown by the corresponding $Skew$ and $QSkew$ values (and by the actual density shown in Section 4 below). Again, the relevant interpretation of *logSkew* and *NPSkew* is relative to the skewness of log-normal returns, which are highly positively skewed for long horizons.

4 The density of long-run returns

We have looked at various measures of asymmetry in the previous sections, and concluded that long-horizon returns are positively skewed. In this section, we illustrate what the (often extreme) skewness numbers imply for the shape of the distribution of long-run compound returns. We start with our benchmark case, where monthly returns are iid log-normal. Figure 6 shows the densities of compound returns at different horizons, for different levels of single-period volatility. The monthly mean is $\mu = 1.01$, and the monthly volatility is either 5%, 8%, or 17%. Since log-normality is preserved under multiplication, the compound returns are also log-normal with the parameters scaled by the horizon. Figure 6 simply plots the resulting (analytical) densities for the the log-normal compound returns. The mean of the compound returns, which is identical across the three specifications in each panel, is indicated with a solid vertical line.

At the annual horizon ($T = 12$), the distribution associated with the lowest monthly volatility ($\sigma = 0.05$) appears fairly symmetric (skewness equal to 0.48), whereas for the more volatile cases clear signs of asymmetry are already noticeable. This is especially true for $\sigma = 0.17$, which is very far from symmetric (skewness equal to 1.86). As the compounding horizon increases, all three distributions become more positively skewed. This is most clearly evident for $\sigma = 0.17$, but also the less volatile cases result in high degrees of asymmetry as T grows large. For $T = 120$ and $T = 360$, the densities for $\sigma = 0.17$ become extremely peaked at small outcomes and the graphs are truncated along the y-axis to maintain useful information for the other two densities.

Figure 6 provides a graphical interpretation of the skewness numbers discussed in the previous sections. As is seen, even relatively small deviations from zero skewness result in markedly non-symmetric distributions (e.g., for $\sigma = 0.08$ and $T = 120$, the skewness is 4.24). Thus, Figure 6 clearly illustrates that there is a strong skew-inducing effect of compounding also for low-volatility assets. We previously highlighted the extreme skewness in long-run individual stock returns, or other volatile assets. In comparison, well diversified portfolios with low volatility appear reasonably well-behaved (compare the skewnesses for $\sigma = 0.05$ and $\sigma = 0.17$ in Table 1). However, this is a relative statement, and in absolute terms the long-run market returns are also very skewed and far from symmetric, as is evident from Figure 6. A portfolio with a monthly volatility of 5%, has a skewness of about 5 in the 30-year compound returns. A monthly volatility of 8% results in a skewness over 30 at the 30-year horizon. As seen in Panel D of Figure 6, the distributions corresponding to these skewness numbers are highly asymmetric.

To further emphasize how compounding induces asymmetry, it is useful to calculate how much mass in the distribution that is to the right of the mean outcome; i.e., the probability of a given outcome exceeding the expected value. Table 3 shows the mean, median, and the probability of the returns exceeding the mean for different combinations of σ and T, still assuming log-normally distributed returns. For $T = 360$, the expected compound return is equal to $\mu^{360} = 35.9$ when $\mu = 1.01$. If the monthly volatility is equal to $\sigma = 0.05$ (0.08), the *median* 30-year compound return is equal to 23.1 (11.7). For $\sigma = 0.05$ (0.08), there is only a 32% (23%) chance of the portfolio beating its 30-year expected return. Naturally, for $\sigma = 0.17$, these figures are even more extreme: The median 30-year compound return is equal to 0.24 (i.e., a net return of -76%) and the probability of beating the mean is less than 6%.

Figure 6 and Table 3 correspond to the iid log-normal case, but we also assess the effect of deviations from iid log-normality. Figure 7 shows the density of 5-year returns for two previously discussed alternatives. First, when single-period returns are iid but not log-normal. In particular, monthly returns are assumed to have a log-NIG distribution with

$Skew(x) = -1$.[18] The second deviation is when returns are non-iid and generated by a stochastic volatility model. In particular, the parameterization with standard leverage is used, which is labeled as SV in Table 2. The iid log-normal density is also presented as a benchmark. Note that the log-normal and log-NIG densities are analytical, while the SV densities are based on simulations. The overall conclusion from Figure 7 is that the different assumptions on the single-period returns result in quite similar densities for 5-year compound returns. The densities are not identical, but it is clear that the long-run distributions are dominated by the compounding effects and not by the short-run properties of the return generating process. Unreported graphs at the 10 and 30-year compounding horizons illustrate how the densities become even more similar as the return horizon increases, as one would expect from central limit convergence.

5 Implications for portfolio choice

We end our analysis with a discussion on implications for portfolio choice. The majority of the literature studying the effect of return skewness on portfolio choice and asset pricing relies on a theoretical framework that uses a third- or fourth-order Taylor expansion of some utility function.[19] As we demonstrate below, this approach can be misleading when applied to longer investment horizons.

The investor's end of horizon wealth is

$$W_T = W_0 X_T \ , \tag{18}$$

where W_0 denotes the initial wealth. Assume that the investor has power utility over terminal wealth with risk aversion $\gamma > 0$, i.e., $u(W_T) = \frac{W_T^{1-\gamma}}{1-\gamma}$. Without loss of generality we set $W_0 = 1$, implying $W_T = X_T$. The investor's utility function can be approximated by a Taylor expansion around the expected wealth. Using a fourth order expansion, the investor's

[18]Ghysels et al. (2016) report the monthly skewness of market returns for 43 countries using data from 1996 to 2013. Except for one country, all reported monthly skewness values are larger than -1, so the $Skew(x) = -1$ we use in this illustration can be considered a strong negative skewness in the case of index returns. Also note that we set the kurtosis of monthly returns to 5, but the densities would be very similar if the monthly kurtosis was set to a higher value, e.g., 10.

[19]Studies concerned with portfolio choice include Scott and Horvath (1980), Conine and Tamarkin (1981), Kane (1982), Jondeau and Rockinger (2006), Guidolin and Timmermann (2008), Martellini and Ziemann (2010), or Ghysels et al. (2016). For papers considering asset pricing implications, see Rubinstein (1973), Kraus and Litzenberger (1976), Dittmar (2002), Chabi-Yo (2012), Langlois (2013), and Neuberger and Payne (2021).

expected utility becomes

$$
\begin{aligned}
E\left[u\left(X_T\right)\right] \approx & \frac{E\left[X_T\right]^{1-\gamma}}{1-\gamma}-\frac{\gamma}{2} \frac{Var\left(X_T\right)}{E\left[X_T\right]^{\gamma+1}}+\frac{\gamma(\gamma+1)}{6} \frac{Var\left(X_T\right)^{3/2}}{E\left[X_T\right]^{\gamma+2}} Skew\left(X_T\right) \\
& -\frac{\gamma(\gamma+1)(\gamma+2)}{24} \frac{Var\left(X_T\right)^2}{E\left[X_T\right]^{\gamma+3}} Kurt\left(X_T\right) \ .
\end{aligned} \tag{19}
$$

If the expansion is only up to the third order, the kurtosis-related term is omitted. Moments of the long-horizon return, X_T, can be calculated via Proposition 1, and therefore calculating the expected utility from equation (19) for various investment horizons is straightforward. In this section, we demonstrate some implications of using the approximated utility, with the below examples clearly illustrating the problems that can arise from using the finite order Taylor expansions.

5.1 Assets with different volatility

Consider an investor who tries to rank two assets (or strategies) based on expected utility. Label the two assets by A and B, and assume that they have the *same* single-period expected return, skewness, and kurtosis, but they have *different* single-period volatilities with $\sigma_A < \sigma_B$. Both assets have iid returns across time. If the investor is risk-averse, it is clear that the low-volatility asset should be preferred for any investment horizon (we show this formally at the end of this section). However, as the following proposition shows, using the *third order* expansion of the expected utility leads to a different conclusion.

Proposition 4 *Let A and B be two assets with single-period returns x_A and x_B. The single-period return moments are $\mu_A = \mu_B$, $\sigma_A < \sigma_B$, and $s_A = s_B$. Both assets have iid returns across time. An investor whose expected utility is described by the third order version of* (19) *prefers asset B over the investment horizon T if*

$$
\frac{\theta_{2B}^T - \theta_{2A}^T}{\theta_{3B}^T - \theta_{3A}^T} < \frac{\gamma+1}{3(\gamma+2)} \ , \tag{20}
$$

where θ_{2j} and θ_{3j} for the two assets $j \in \{A, B\}$ are defined in (8). *Assuming that $s_B > \max\left(-2\frac{\mu_B}{\sigma_B}, -3\mu_B \frac{\sigma_B^2-\sigma_A^2}{\sigma_B^3-\sigma_A^3}\right)$, this also implies that*

(i) the investor prefers asset B as $T \to \infty$, and

(ii) there exists a horizon T^* such that the investor prefers asset B for all $T > T^*$.

T^* is the value of T that solves (20) *as an equality.*

Proof. See the proof in Appendix A.4. ■

As the horizon increases, the investor starts to prefer the high-volatility asset at a certain point.[20] Why does this happen? The skewness of the high-volatility asset grows (much) quicker with the horizon (see Panel A of Table 1), and the skewness-related term in the third-order expansion therefore becomes so dominant as the horizon increases that the expected utility from the high-volatility asset outgrows the utility from the low-volatility asset. That is, the third order approximation leads to the erroneous conclusion that the high-volatility asset is preferred for long investment horizons.

Panel A of Figure 8 illustrates this by showing the expected utilities, based on the third order approximation, from the two assets as a function of the investment horizon in a specific example ($\mu_A = \mu_B = 1.01$, $\sigma_A = 0.05$, $\sigma_B = 0.1$, $s_A = s_B = 0$, and $\gamma = 5$). For short horizons, the investor prefers the low-volatility asset, but for horizons longer than 13 months, the investor apparently prefers the high-volatility asset. Panel B shows the value of T^* as a function of σ_B and γ (while the other parameters are unchanged from Panel A). T^* is the threshold horizon at which the investor starts to prefer the high-volatility asset (according to the third order approximation). This threshold decreases as (i) the investor's risk aversion increases or (ii) the volatility difference between the two assets increases. Also note that the switch in preference occurs at rather short horizons (within a year) if the investor is highly risk averse, or the volatility difference between the assets is high.

One could argue that a fourth order expansion of the utility should be used instead. Indeed, the investor prefers the low-volatility asset for any horizon if the fourth order approximation is used, because the kurtosis of the high-volatility asset increases quicker with the horizon than that of the low-volatility asset (which is another direct implication of Proposition 1), and this compensates for the effect of skewness in (19). However, as shown in the next example, there are also cases when the fourth order expansion leads to analogous errors as those just illustrated for the third order expansion.

5.2 Assets with different skewness

Consider next two assets that have the *same* single-period expected return, volatility, and kurtosis, but their single-period return skewness is *different* with $s_A > s_B$. Both assets have iid returns across time. Intuitively, since the assets are identical apart from the greater skewness in asset A, the investor ought to prefer A to B, but this is not always the answer

[20] The condition $s_B > \max\left(-2\frac{\mu_B}{\sigma_B}, -3\mu_B\frac{\sigma_B^2-\sigma_A^2}{\sigma_B^3-\sigma_A^3}\right)$ in Proposition 4 rules out *extreme negative* values for the common single-period skewness. It is only a technical condition that will be satisfied for any reasonable parameterizations at the monthly level. For the example presented in Panel A of Figure 8, this condition implies $s_B > -20.2$, i.e., the common monthly return skewness has to be above -20.

delivered from the fourth order approximation.

Proposition 5 *Let A and B be two assets with single-period returns x_A and x_B. The single-period return moments are $\mu_A = \mu_B$, $\sigma_A = \sigma_B$, $s_A > s_B$, and $Kurt(x_A) = Kurt(x_B)$. Both assets have iid returns across time. An investor whose expected utility is described by the fourth order expansion in* (19) *prefers asset B over the investment horizon T if*

$$\frac{\theta_{3B}^T - \theta_{3A}^T}{\theta_{4B}^T - \theta_{4A}^T} < \frac{\gamma + 2}{4(\gamma + 3)} , \tag{21}$$

where θ_{3j} and θ_{4j} for the two assets $j \in \{A, B\}$ are defined in (8). *Assuming that $s_B > -\left(\frac{Kurt(x_B)}{3}\frac{\sigma_B}{\mu_B} + \frac{\mu_B}{\sigma_B}\right)$, this also implies that*

(i) the investor prefers asset B as $T \to \infty$, and

(ii) there exists a horizon T^{**} such that the investor prefers asset B for all $T > T^{**}$. T^{**} is the value of T that solves (21) *as an equality.*

Proof. See the proof in Appendix A.5. ■

As the horizon increases, the investor starts to prefer the asset with the lower (e.g., more negative) skewness at a certain point.[21] Using Proposition 1, it can be shown that the kurtosis of compound returns grows much quicker with the horizon if the single-period return skewness is higher. For longer horizons, the kurtosis-related term dominates the expression in (19), and the lower long-horizon kurtosis implies that the approximated expected utility is higher for the less skewed (or more negatively skewed) asset.

Panel C of Figure 8 illustrates this by showing the expected utilities, based on the fourth order approximation, from the two assets as a function of horizon in a specific example ($\mu_A = \mu_B = 1.01$, $\sigma_A = \sigma_B = 0.08$, $s_A = 0$, $s_B = -1$, $Kurt(x_A) = Kurt(x_B) = 5$, and $\gamma = 5$). For short horizons, the investor prefers the non-skewed asset, but for horizons longer than 8 months, the investor seems to prefer the asset with the negative skewness. Panel D shows the value of T^{**} as a function of the common volatility σ and γ (while the other parameters are unchanged from Panel C). The threshold horizon at which the investor starts to prefer the negatively skewed asset (according to the fourth order approximation) decreases as (i) the investor's risk aversion increases or (ii) the common volatility of the assets increases.[22] Also note that the switch in preference occurs at rather short horizons

[21] The condition $s_B > -\left(\frac{Kurt(x_B)}{3}\frac{\sigma_B}{\mu_B} + \frac{\mu_B}{\sigma_B}\right)$ in Proposition 5 rules out *extreme negative* values for the single-period skewness of asset B. It is only a technical condition that will be satisfied for any reasonable parameterizations at the monthly level. For the example presented in Panel C of Figure 8, this condition implies $s_B > -12.8$, i.e., the monthly return skewness of asset B has to be above -12.

[22] Note that the graph in Panel D of Figure 8 does not change considerably if we change the skewness of asset B (e.g., to $s_B = -2$ or $s_B = -0.5$) or the common kurtosis (e.g., to $Kurt(x_A) = Kurt(x_B) = 10$).

(within a year) if the investor is highly risk averse, or the common volatility of the assets is high.

5.3 An exact solution

The approach using approximated utility functions can be misleading because moments scale differently with the horizon. Therefore, omitting higher order moments by truncating the Taylor series at an arbitrarily chosen term can lead to erroneous conclusions. By "erroneous," we mean that an exact solution of the problem would lead to different conclusions. We show this by considering a setting where the single-period returns from the assets are iid log-NIG (the case of log-normal period returns is nested). This assumption allows for an exact analytical formula for the expected utility from holding the assets of the previous examples at an arbitrary horizon T.

Proposition 6 *Let A and B be two assets with single-period returns x_A and x_B. Both x_A and x_B have a log-NIG distribution, and the assets have iid returns across time. If an investor with power utility prefers asset A over a single period, then the investor also prefers asset A over any investment horizon T, i.e.,*

$$E\left[u(x_A)\right] > E\left[u(x_B)\right] \quad \Leftrightarrow \quad E\left[u(X_{TA})\right] > E\left[u(X_{TB})\right] \ . \tag{22}$$

Proof. See the proof in Appendix A.6. ■

The investor will prefer the low-volatility asset for any horizon in the example illustrated in Panel A of Figure 8, and will prefer the non-skewed asset for any horizon in the example corresponding to Panel C. That is, the misleading conclusions of Figure 8 are indeed the result of the Taylor approximations.

6 Conclusion

Multiplicative compounding of simple gross returns almost inevitably induces positive skewness. This effect is primarily a function of the volatility of the single-period returns and is relatively little affected by the specific single-period distribution. For volatility levels associated with market-wide index returns, the skew-inducing effect of compounding is strong, whereas for individual stock returns it is extreme. Skewness for individual stock returns at a 30-year horizon can easily be of the order of *millions*, and most of the mass of the distribution is concentrated to outcomes close to zero, despite a very large mean. These results provide a theoretical foundation to Bessembinder's (2018) study on the empirical behavior of long-run individual stock returns. Our results also question recent evidence suggesting that long-run (e.g., 5-year) aggregate returns have negative skewness. At long horizons, simple returns and log returns are fundamentally different, and while log returns can easily remain negatively skewed at long horizons, this is almost impossible for compound simple returns.

The skewness-inducing effect (and the effects on other higher-order moments) are important to investors. We point to two immediate implications. First, large positive skewness makes the mean a fairly uninformative statistic of typical return outcomes at long horizons. With increasing positive skewness, it becomes more and more difficult to beat the mean outcome as the investment horizon increases. Second, standard third- and fourth-order Taylor expansions of utility functions, which attempt to account for skewness and kurtosis, can lead to erroneous conclusions when applied to longer-horizon returns. In fact, these effects can manifest already at annual horizons.

A Proofs

A.1 Proof of Proposition 1

Define the compound process $X_T = x_1 \times ... \times x_T$. Since x_t are iid, we have $E\left[X_T^j\right] = E\left[x^j\right]^T$, which also implies

$$\frac{E\left[X_T^j\right]}{E\left[X_T\right]^j} = \left(\frac{E\left[x^j\right]}{E\left[x\right]^j}\right)^T = \theta_j^T \ . \tag{A1}$$

Using the binomial expansion of the k-th central moment of X_T (for $k > 2$),

$$\begin{aligned}\frac{E\left[\left(X_T - E\left[X_T\right]\right)^k\right]}{Var\left(X_T\right)^{k/2}} &= \frac{E\left[\sum_{j=0}^k \binom{k}{j}(-1)^j X_T^{k-j} E\left[X_T\right]^j\right]}{Var\left(X_T\right)^{k/2}} = \frac{\sum_{j=0}^k \binom{k}{j}(-1)^j E\left[X_T^{k-j}\right] E\left[X_T\right]^j}{Var\left(X_T\right)^{k/2}} \\ &= \frac{\sum_{j=0}^k \binom{k}{j}(-1)^j \frac{E\left[X_T^{k-j}\right]}{E\left[X_T\right]^{k-j}}}{\left(\frac{Var(X_T)}{E[X_T]^2}\right)^{k/2}} = \frac{\sum_{j=0}^k \binom{k}{j}(-1)^j \theta_{k-j}^T}{\left(\theta_2^T - 1\right)^{k/2}} \ .\end{aligned} \tag{A2}$$

To get to the second line above, divide both the numerator and the denominator by $E\left[X_T\right]^k$. We also provide a simple recursive formula to calculate the θ_k values. It is straightforward that

$$\theta_1 = \frac{E\left[x\right]}{E\left[x\right]} = 1 \qquad \text{and} \qquad \theta_2 = \frac{E\left[x^2\right]}{E\left[x\right]^2} = 1 + \left(\frac{E\left[x^2\right]}{E\left[x\right]^2} - 1\right) = 1 + \frac{\sigma^2}{\mu^2} \ . \tag{A3}$$

To determine θ_k for $k > 2$, start again with the binomial expansion of the k-th central moment, and then divide both the numerator and denominator by $E\left[x\right]^k$:

$$\begin{aligned}\frac{E\left[\left(x - E\left[x\right]\right)^k\right]}{Var\left(x\right)^{k/2}} &= \frac{E\left[\sum_{j=0}^k \binom{k}{j}(-1)^j x^{k-j} E\left[x\right]^j\right]}{Var\left(x\right)^{k/2}} = \frac{\sum_{j=0}^k \binom{k}{j}(-1)^j E\left[x^{k-j}\right] E\left[x\right]^j}{Var\left(x\right)^{k/2}} \\ &= \frac{\sum_{j=0}^k \binom{k}{j}(-1)^j \frac{E\left[x^{k-j}\right]}{E[x]^{k-j}}}{\left(\frac{Var(x)}{E[x]^2}\right)^{k/2}} = \frac{\sum_{j=0}^k \binom{k}{j}(-1)^j \theta_{k-j}}{\left(\theta_2 - 1\right)^{k/2}} \\ &= \frac{\theta_k + \left(\sum_{j=1}^{k-2} \binom{k}{j}(-1)^j \theta_{k-j}\right) + (-1)^k\left(1 - k\right)}{\left(\theta_2 - 1\right)^{k/2}} \ .\end{aligned} \tag{A4}$$

Rearranging equation (A4) yields

$$\theta_k = (-1)^k\left(k - 1\right) - \left(\sum_{j=1}^{k-2} \binom{k}{j}(-1)^j \theta_{k-j}\right) + \left(\theta_2 - 1\right)^{k/2} \frac{E\left[\left(x - \mu\right)^k\right]}{\sigma^k} \ . \tag{A5}$$

A.2 Proof of Proposition 2

Let the k-th central moment of the variable Z be denoted by μ_k, and its sample analogue by m_k, i.e.,

$$\mu_k \equiv E\left[(Z - E[Z])^k\right] \qquad \text{and} \qquad m_k \equiv \frac{1}{n}\sum_{i=1}^{n}(z_i - \bar{z})^k \ . \tag{A6}$$

Then the skewness of Z, and its estimator from equation (9), $\hat{S}$, are defined as

$$Skew(z) = \frac{\mu_3}{\mu_2^{3/2}} \qquad \text{and} \qquad \hat{S} = \frac{m_3}{m_2^{3/2}} \ . \tag{A7}$$

Provided the third moment of Z exists, m_2 and m_3 trivially converge to μ_2 and μ_3, respectively, by a law of large numbers. Further, from Serfling (1980, page 72), as $n \to \infty$,

$$\sqrt{n}\begin{bmatrix} m_2 - \mu_2 \\ m_3 - \mu_3 \end{bmatrix} \xrightarrow{d} N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{array}{cc} \mu_4 - \mu_2^2 & \mu_5 - 4\mu_2\mu_3 \\ \mu_5 - 4\mu_2\mu_3 & \mu_6 - \mu_3^2 - 6\mu_2\mu_4 + 9\mu_2^3 \end{array}\right). \tag{A8}$$

By the delta method, and provided the sixth moment of Z exists, $\hat{S}$ satisfies the result in equation (11). The asymptotic normality result in Serfling (1980) is derived for the iid case, although the result should extend to more general cases as long as sufficient conditions for a central limit theorem apply.

A.3 Proof of Proposition 3

Let x_t be log-normally distributed with parameters ψ and η. That is, the log-returns $y_t \equiv \log(x_t)$ are normally distributed with mean ψ and volatility η. Assume further that y_t follows a linear (infinite moving average) process, such that

$$y_t = \psi + u_t, \tag{A9}$$

and

$$u_t = C(L)\,\epsilon_t = \sum_{j=0}^{\infty} c_j \epsilon_{t-j}. \tag{A10}$$

The innovations ϵ_t are assumed to be iid standard normal, i.e., $\epsilon_t \sim N(0,1)$.

The compound return over T periods is given by $X_T = \prod_{t=1}^{T} x_t$ and the log-compound returns satisfy,

$$Y_T = \log(X_T) = \sum_{t=1}^{T} y_t = \psi T + \sum_{t=1}^{T} u_t. \tag{A11}$$

Using the BN decomposition (Beveridge and Nelson, 1981), we can write

$$u_t = C(L)\,\epsilon_t = C(1)\,\epsilon_t + \tilde{\epsilon}_{t-1} - \tilde{\epsilon}_t, \tag{A12}$$

where

$$\tilde{\epsilon}_t = \tilde{C}(L)\,\epsilon_t = \sum_{j=0}^{\infty} \tilde{c}_j \epsilon_{t-j} \qquad \text{and} \qquad \tilde{c}_j = \sum_{s=j+1}^{\infty} c_s \ . \tag{A13}$$

$C(1) = \sum_{j=0}^{\infty} c_j$ denotes the so-called long-run moving average coefficient. The process Y_T can therefore be written as,

$$Y_T = \psi T + C(1) \sum_{t=1}^{T} \epsilon_t + \sum_{t=1}^{T} (\tilde{\epsilon}_{t-1} - \tilde{\epsilon}_t) = \psi T + C(1) \sum_{t=1}^{T} \epsilon_t - \tilde{\epsilon}_T \ , \tag{A14}$$

using the fact that $\sum_{t=1}^{T} (\tilde{\epsilon}_{t-1} - \tilde{\epsilon}_t) = \tilde{\epsilon}_0 - \tilde{\epsilon}_T$ and imposing $\tilde{\epsilon}_0 = 0$.

The BN decomposition decomposes the process into a drift component (ψT), a martingale component $\left(C(1) \sum_{t=1}^{T} \epsilon_t\right)$, and a transitory component $(\tilde{\epsilon}_T)$. For large T, the permanent (martingale) component has a variation that is of an order of magnitude greater than the transitory component, and will therefore dominate the stochastic properties of Y_T. We can therefore write the "long-run" part of Y_T as

$$Y_T^{LR} \equiv \psi T + C(1) \sum_{i=t}^{T} \epsilon_t \approx Y_T. \tag{A15}$$

Since $\epsilon_t \overset{iid}{\sim} N(0,1)$, it follows that $\sum_{t=1}^{T} \epsilon_t \sim N(0,T)$ and $Y_T^{LR} \sim N\left(\psi T, C(1)^2 T\right)$. Thus, from the definition of the log-normal distribution, $e^{Y_T^{LR}} \sim LN\left(\psi T, C(1)^2 T\right)$. That is, since $\log(X_T) = Y_T \approx Y_T^{LR}$, $X_T \approx LN\left(\psi T, C(1)^2 T\right)$. The parameters ψT and $C(1)^2 T$ pin down the distribution, and therefore also the skewness, of the compound returns as discussed previously.

In order to assess the effects of serial dependence on compound returns vis-à-vis the iid setting, consider the case where u_t is iid. In this case, $\log(X_T) = Y_T^{LR} \sim N(\psi T, \eta^2 T)$, where the equality between $\log(X_T)$ and Y_T^{LR} is now exact. Compared to the serially correlated case, the mean parameter of the distribution of Y_T^{LR} is the same, but the variance is different. The effect of the serial dependence is therefore summarized by the differences between the variance of Y_T^{LR} in the serially correlated case, and the variance of Y_T^{LR} in the iid case.

Note that the (short-run) variance of y_t is given by,

$$\eta^2 = Var\left(y_t\right) = Var\left(u_t\right) = \sum_{j=0}^{\infty} c_j^2, \tag{A16}$$

and the so-called long-run variance of y_t is given by

$$LR.Var\left(y_t\right) \equiv \sum_{j=-\infty}^{\infty} Cov\left(y_t, y_{t+j}\right) = C\left(1\right)^2 = \left(\sum_{j=0}^{\infty} c_j\right)^2. \tag{A17}$$

The variance of Y_T^{LR} in the iid case is thus equal to $T \times Var\left(y_t\right)$, whereas the variance of Y_T^{LR} in the serially correlated case is equal to $T \times LR.Var\left(y_t\right)$. Define the variance ratio, of the long-run variance of y_t over its short-run variance,

$$VR \equiv \frac{LR.Var\left(y_t\right)}{Var\left(y_t\right)} = \frac{\left(\sum_{j=0}^{\infty} c_j\right)^2}{\sum_{j=0}^{\infty} c_j^2}. \tag{A18}$$

A given serial correlation structure $\{c_j\}_{j=0}^{\infty}$ reduces or increases the long-run variance of y_t, relative to the iid case, by a factor given in the expression above. The impact of serial correlation on skewness in compound returns is therefore evaluated by comparing the skewness implied for compound returns when using the short-run variance and the skewness implied when using the long-run variance.

The variances in equation (A18) correspond to log-returns, y_t, whereas the inputs into the skewness formula in Corollary 1 of the main text correspond to variances of simple returns, x_t. Since $x_t \sim LN\left(\psi, \eta^2\right)$, the relationship between the parameter η^2 and $Var\left(x\right) = \sigma^2$ is given by $\sigma^2 = \mu^2\left(e^{\eta^2} - 1\right)$, where $\mu = E\left[x\right]$. Defining $\sigma_{VR}^2 \equiv \mu^2\left(e^{VR\eta^2} - 1\right)$ and using the first-order Taylor approximation $e^{VR\eta^2} - 1 \approx VR\eta^2$ (which is a good approximation, since the typical values of η^2 in our context are close to zero), it can be shown that

$$\frac{\sigma_{VR}^2}{\sigma^2} = \frac{e^{VR\eta^2} - 1}{e^{\eta^2} - 1} \approx \frac{VR\eta^2}{\eta^2} = VR\,.$$

That is, $\sigma_{VR}^2 \approx VR\sigma^2$. As the log-variance shifts by a factor VR, so does the variance of simple returns, up to a first order approximation. In order to assess the effect of serial correlation on skewness, one can therefore also equally well calculate the variance ratio of the simple returns, rather than for the log-returns.

A.4 Proof of Proposition 4

Let x_A and x_B be single-period gross returns on assets A and B, respectively, with $\mu_A = \mu_B$, $\sigma_A < \sigma_B$, and $s_A = s_B$. Define $X_{Tj} = \prod_{t=1}^{T} x_{tj}$ where the x_{tj}-s are iid and have the same distribution as x_j for both assets $j \in \{A, B\}$. Using the third order version of the utility expansion (19), asset B is preferred over horizon T, i.e., $E[u(X_{TB})] > E[u(X_{TA})]$, if

$$\begin{aligned}&\frac{E\left[X_{TB}\right]^{1-\gamma}}{1-\gamma} - \frac{\gamma}{2}\frac{Var\left(X_{TB}\right)}{E\left[X_{TB}\right]^{\gamma+1}} + \frac{\gamma\left(\gamma+1\right)}{6}\frac{Var\left(X_{TB}\right)^{3/2}}{E\left[X_{TB}\right]^{\gamma+2}}Skew\left(X_{TB}\right) \\ &\quad > \frac{E\left[X_{TA}\right]^{1-\gamma}}{1-\gamma} - \frac{\gamma}{2}\frac{Var\left(X_{TA}\right)}{E\left[X_{TA}\right]^{\gamma+1}} + \frac{\gamma\left(\gamma+1\right)}{6}\frac{Var\left(X_{TA}\right)^{3/2}}{E\left[X_{TA}\right]^{\gamma+2}}Skew\left(X_{TA}\right) \ .\end{aligned} \tag{A19}$$

Using that $E[X_{TB}] = E[X_{TA}]$ (since $\mu_A = \mu_B$) and the formulas for $Var(X_{Tj})$ and $Skew(X_{Tj})$ in equations (4) and (7), straightforward algebraic manipulation shows that the inequality in (A19) is equivalent to

$$\frac{\gamma+1}{3\left(\gamma+2\right)} > \frac{\theta_{2B}^T - \theta_{2A}^T}{\theta_{3B}^T - \theta_{3A}^T} \ , \tag{A20}$$

where for $j \in \{A, B\}$,

$$\theta_{2j} = \frac{\sigma_j^2}{\mu_j^2} + 1 \qquad \text{and} \qquad \theta_{3j} = s_j\frac{\sigma_j^3}{\mu_j^3} + 3\frac{\sigma_j^2}{\mu_j^2} + 1 \ . \tag{A21}$$

Using the definitions in (A21), the following inequalities can be established:

- $\theta_{3B} > \theta_{2B}$ and $\theta_{3A} > \theta_{2A}$ if $s_B > -2\frac{\mu_B}{\sigma_B}$,
- $\theta_{2B} > \theta_{2A}$ if $\sigma_B > \sigma_A$ (which is given by assumption), and
- $\theta_{3B} > \theta_{3A}$ if $s_B > -3\mu_B\frac{\sigma_B^2-\sigma_A^2}{\sigma_B^3-\sigma_A^3}$.

Given that $s_B > \max\left(-2\frac{\mu_B}{\sigma_B}, -3\mu_B\frac{\sigma_B^2-\sigma_A^2}{\sigma_B^3-\sigma_A^3}\right)$, the above inequalities imply

$$\lim_{T\to\infty}\frac{\theta_{2B}^T - \theta_{2A}^T}{\theta_{3B}^T - \theta_{3A}^T} = \lim_{T\to\infty}\left(\frac{\theta_{2B}}{\theta_{3B}}\right)^T\frac{\left(1-\left(\frac{\theta_{2A}}{\theta_{2B}}\right)^T\right)}{\left(1-\left(\frac{\theta_{3A}}{\theta_{3B}}\right)^T\right)} = 0 \ . \tag{A22}$$

Let us define the function $g_\theta(T) \equiv \frac{\theta_{2B}^T-\theta_{2A}^T}{\theta_{3B}^T-\theta_{3A}^T}$. It is easy to show that $g_\theta(2) < g_\theta(1)$, as

$$\frac{\theta_{2B}^2 - \theta_{2A}^2}{\theta_{3B}^2 - \theta_{3A}^2} = \frac{\left(\theta_{2B}-\theta_{2A}\right)\left(\theta_{2B}+\theta_{2A}\right)}{\left(\theta_{3B}-\theta_{3A}\right)\left(\theta_{3B}+\theta_{3A}\right)} < \frac{\left(\theta_{2B}-\theta_{2A}\right)\left(\theta_{3B}+\theta_{3A}\right)}{\left(\theta_{3B}-\theta_{3A}\right)\left(\theta_{3B}+\theta_{3A}\right)} = \frac{\theta_{2B}-\theta_{2A}}{\theta_{3B}-\theta_{3A}} \ , \tag{A23}$$

where the inequality is true because we increase the numerator ($\theta_{3B} > \theta_{2B}$ and $\theta_{3A} > \theta_{2A}$). That is, there exists $1 < T' < 2$ such that $g'_\theta(T') < 0$ (i.e., where the derivative is negative). After somewhat tedious and unenlightening algebra it can be shown that if $g'_\theta(T') < 0$, then $g'_\theta(T) < 0$ for all $T > T'$. Taken all the above results together, the value of $\frac{\theta_{2B}^T - \theta_{2A}^T}{\theta_{3B}^T - \theta_{3A}^T}$ decreases, and approaches zero as T increases. Since $\lim_{\gamma \to 0^+} \frac{\gamma+1}{3(\gamma+2)} = \frac{1}{6}$ and $\frac{\gamma+1}{3(\gamma+2)}$ is strictly increasing in γ (its first order derivative is always positive), the lower bound of $\frac{\gamma+1}{3(\gamma+2)}$ is $\frac{1}{6}$ when $\gamma > 0$. That is, the left hand side of (A20) is greater than $\frac{1}{6}$, and the right hand side decreases and approaches zero as T increases, so (A20) will be satisfied for some $T^* < \infty$ value. If it is satisfied for T^*, then it is also satisfied for all $T > T^*$.

A.5 Proof of Proposition 5

Let x_A and x_B be single-period gross returns on assets A and B, respectively, with $\mu_A = \mu_B$, $\sigma_A = \sigma_B$, $s_A > s_B$, and $Kurt(x_A) = Kurt(x_B)$. Define $X_{Tj} = \prod_{t=1}^{T} x_{tj}$ where the x_{tj}-s are iid and have the same distribution as x_j for both assets $j \in \{A, B\}$. Using the utility expansion (19) and that $E[X_{TA}] = E[X_{TB}]$ and $Var(X_{TA}) = Var(X_{TB})$ (since $\mu_A = \mu_B$ and $\sigma_A = \sigma_B$), asset B is preferred over horizon T, i.e., $E[u(X_{TB})] > E[u(X_{TA})]$, if

$$\begin{aligned} &\frac{\gamma(\gamma+1)}{6} \frac{Var(X_{TB})^{3/2}}{E[X_{TB}]^{\gamma+2}} Skew(X_{TB}) - \frac{\gamma(\gamma+1)(\gamma+2)}{24} \frac{Var(X_{TB})^2}{E[X_{TB}]^{\gamma+3}} Kurt(X_{TB}) \\ &> \frac{\gamma(\gamma+1)}{6} \frac{Var(X_{TA})^{3/2}}{E[X_{TA}]^{\gamma+2}} Skew(X_{TA}) - \frac{\gamma(\gamma+1)(\gamma+2)}{24} \frac{Var(X_{TA})^2}{E[X_{TA}]^{\gamma+3}} Kurt(X_{TA}) \ . \end{aligned} \tag{A24}$$

Using again that $E[X_{TA}] = E[X_{TB}]$ and $Var(X_{TA}) = Var(X_{TB})$, and the formulas for $Skew(X_{Tj})$ and $Kurt(X_{Tj})$ in equation (7), straightforward algebraic manipulation shows that the inequality in (A24) is equivalent to

$$\frac{\gamma+2}{4(\gamma+3)} > \frac{\theta_{3B}^T - \theta_{3A}^T}{\theta_{4B}^T - \theta_{4A}^T} \ , \tag{A25}$$

where for $j \in \{A, B\}$,

$$\theta_{3j} = s_j \frac{\sigma_j^3}{\mu_j^3} + 3\frac{\sigma_j^2}{\mu_j^2} + 1 \qquad \text{and} \qquad \theta_{4j} = Kurt(x_j)\frac{\sigma_j^4}{\mu_j^4} + 4 s_j \frac{\sigma_j^3}{\mu_j^3} + 6\frac{\sigma_j^2}{\mu_j^2} + 1 \tag{A26}$$

Using the definitions in (A26), the following inequalities can be established:

- $\theta_{4B} > \theta_{3B}$ and $\theta_{4A} > \theta_{3A}$ if $s_B > -\left(\frac{Kurt(x_B)}{3}\frac{\sigma_B}{\mu_B} + \frac{\mu_B}{\sigma_B}\right)$,
- $\theta_{3A} > \theta_{3B}$ if $s_A > s_B$ (which is given by assumption), and

- $\theta_{4A} > \theta_{4B}$ if $s_A > s_B$ (which is given by assumption).

Given that $s_B > -\left(\frac{Kurt(x_B)}{3}\frac{\sigma_B}{\mu_B} + \frac{\mu_B}{\sigma_B}\right)$, the above inequalities imply

$$\lim_{T\to\infty} \frac{\theta_{3B}^T - \theta_{3A}^T}{\theta_{4B}^T - \theta_{4A}^T} = \lim_{T\to\infty} \left(\frac{\theta_{3A}}{\theta_{4A}}\right)^T \frac{\left(\left(\frac{\theta_{3B}}{\theta_{3A}}\right)^T - 1\right)}{\left(\left(\frac{\theta_{4B}}{\theta_{4A}}\right)^T - 1\right)} = 0 \ . \tag{A27}$$

Let us define the function $h_\theta(T) \equiv \frac{\theta_{3B}^T - \theta_{3A}^T}{\theta_{4B}^T - \theta_{4A}^T}$. It is easy to show that $h_\theta(2) < h_\theta(1)$, as

$$\frac{\theta_{3B}^2 - \theta_{3A}^2}{\theta_{4B}^2 - \theta_{4A}^2} = \frac{(\theta_{3B} - \theta_{3A})(\theta_{3B} + \theta_{3A})}{(\theta_{4B} - \theta_{4A})(\theta_{4B} + \theta_{4A})} < \frac{(\theta_{3B} - \theta_{3A})(\theta_{4B} + \theta_{4A})}{(\theta_{4B} - \theta_{4A})(\theta_{4B} + \theta_{4A})} = \frac{\theta_{3B} - \theta_{3A}}{\theta_{4B} - \theta_{4A}} \ , \tag{A28}$$

where the inequality is true because we increase the numerator ($\theta_{4B} > \theta_{3B}$ and $\theta_{4A} > \theta_{3A}$). That is, there exists $1 < T'' < 2$ such that $h'_\theta(T'') < 0$ (i.e., where the derivative is negative). After somewhat tedious and unenlightening algebra it can be shown that if $h'_\theta(T'') < 0$, then $h'_\theta(T) < 0$ for all $T > T''$. Taken all the above results together, the value of $\frac{\theta_{3B}^T - \theta_{3A}^T}{\theta_{4B}^T - \theta_{4A}^T}$ decreases, and approaches zero as T increases. Since $\lim_{\gamma\to 0^+} \frac{\gamma+2}{4(\gamma+3)} = \frac{1}{6}$ and $\frac{\gamma+2}{4(\gamma+3)}$ is strictly increasing in γ (its first order derivative is always positive), the lower bound of $\frac{\gamma+2}{4(\gamma+3)}$ is $\frac{1}{6}$ when $\gamma > 0$. That is, the left hand side of (A25) is greater than $\frac{1}{6}$, and the right hand side decreases and approaches zero as T increases, so (A25) will be satisfied for some $T^{**} < \infty$ value. If it is satisfied for T^{**}, then it is also satisfied for all $T > T^{**}$.

A.6 Proof of Proposition 6

Assume that $x_A \sim LNIG(\varphi_A, \beta_A, \nu_A, \delta_A)$ and $x_B \sim LNIG(\varphi_B, \beta_B, \nu_B, \delta_B)$. For further details about the log-NIG distribution see Appendix C.2. Define $X_{Tj} = \prod_{t=1}^T x_{tj}$ where the x_{tj}-s are iid and have the same distribution as x_j for both assets $j \in \{A, B\}$. Then, similar to (A53),

$$X_{Tj} \sim LNIG\left(\varphi_j, \beta_j, T\nu_j, T\delta_j\right) \ . \tag{A29}$$

Since the NIG distribution is closed under affine transformations (see (A47)), this implies

$$X_{Tj}^{1-\gamma} \sim LNIG\left(\frac{\varphi_j}{|1-\gamma|}, \frac{\beta_j}{1-\gamma}, (1-\gamma)\, T\nu_j, |1-\gamma|\, T\delta_j\right) \ . \tag{A30}$$

Using (A48),

$$E\left[X_{Tj}^{1-\gamma}\right] = e^T \times \exp\left((1-\gamma)\,\nu_j + \delta_j\left(\sqrt{\varphi_j^2 - \beta_j^2} - \sqrt{\varphi_j^2 - \left(\beta_j + 1 - \gamma\right)^2}\right)\right) \ . \tag{A31}$$

Applying the above for $T = 1$ as well, reveals that $E\left[X_{Tj}^{1-\gamma}\right] = e^T E\left[x_j^{1-\gamma}\right]$. Therefore, for any horizon T, and any $\gamma > 0$ ($\gamma \neq 1$) we have

$$E\left[x_{tB}^{1-\gamma}\right] > E\left[x_{tA}^{1-\gamma}\right] \quad \Leftrightarrow \quad E\left[X_{TB}^{1-\gamma}\right] > E\left[X_{TA}^{1-\gamma}\right] \ , \tag{A32}$$

which also implies the result in (22).

B Details of the simulation specification

The simulation results in Section 2.3.2 are based on simulated paths of continuous-time stochastic-volatility (SV) models. Specifically, we use a specification along the lines of that originally proposed by Heston (1993). Let $X(t)$ denote the continuous-time price process, standardized such that $X(0) = 1$. For any t, $X(t)/X(0) = X(t)$ now represents the compound gross returns from time 0 to time t. Prices are assumed to be generated according to the following diffusion model,

$$\frac{dX(t)}{X(t)} = \alpha dt + \sqrt{V(t)} dW_X(t), \tag{A33}$$

$$dV(t) = \phi(\theta - V(t))\, dt + \sigma_V \sqrt{V(t)} dW_V(t). \tag{A34}$$

Here, $W_X(t)$ and $W_V(t)$ are standard Brownian motions with correlation ρ, and $V(t)$ represents the stochastic variance process. Expected net returns per time unit are given by α. ϕ is the mean reversion parameter for the variance process, θ is the long-run (average) level of the variance processes, and σ_V represents the "volatility of volatility". The process is simulated by a quadratic exponential discretization scheme, using a step size equal to 1/10th of a trading day. That is, each simulated month of data represents 210 trading intervals, using a 21-day trading month. The simulations are initialized with $X(0) = 1$ and $V(t) = \theta$. For each simulated path, the equivalent of a 40-year price series is generated. The first 10 years are discarded, such that the starting point for the volatility process is effectively a draw from its stationary distribution. The remaining 30-years are used to create returns at the 1, 12, 24, 60, 120, 240, and 360-month horizons. The simulation results are based on 10^7 simulated sample paths.

The parameters ϕ and σ_V are set mostly in line with empirical studies on aggregate U.S. stock returns (e.g., Andersen et al., 2002,. Eraker et al., 2003, Eraker, 2004). Specifically, on a daily frequency, the mean reversion parameter ϕ is set 0.02 and the volatility-of-volatility parameter is set to 0.002 (i.e., 0.2%). The value of θ is set to correspond to the target volatility level σ in each specification. Thus, in the daily parametrization of the model,

$\theta = \sigma^2/\,21$, where σ represents the monthly volatility (set to $0.05, 0.08, 0.11$, and 0.17). Estimates of the correlation, or "leverage" parameter, ρ, varies somewhat across studies and data sets, with estimates typically ranging between -0.3 and -0.5. In our main specification, we set $\rho = -0.5$, but we also simulate alternatives with $\rho = 0$, and $\rho = -0.75$. The latter value is likely empirically too large in absolute magnitude, but serves to illustrate the effects of strengthening the leverage channel. Setting $\rho = 0$ implies that the Brownian motions driving prices and volatility are independent, turning off the leverage effect. While the leverage effect is quite strong in index returns, it appears much weaker in individual stock returns (e.g. Kim and Kon, 1994, and Andersen et al., 2001), and for large σ, $\rho = 0$ might be the most realistic parameter value. Finally, we set $\alpha = 0.01/\,21$, representing a monthly gross return of $\mu = 1.01$.

Note that the returns from an iid log-normal process can be viewed as discretely sampled returns from a special case of the above model, where $V\left(t\right) \equiv \sigma^2$ for all t. In this case, the SV model collapses to a standard geometric Brownian motion.

C Alternative measures of asymmetry

C.1 The log-normal distribution

Let x_t, $t = 1, ..., T$, be iid log-normal variables with parameters ψ and η, i.e., $x_t \sim LN(\psi, \eta^2)$. Note that the parameters of the distribution can be set to match the first two moments of x_t via

$$\psi = \log\left(\frac{\mu^2}{\sqrt{\mu^2 + \sigma^2}}\right) \qquad \text{and} \qquad \eta = \sqrt{\log\left(\frac{\sigma^2}{\mu^2} + 1\right)}\;, \tag{A35}$$

where μ and σ are the expected value and standard deviation of x_t, respectively. The T-period compound return, $X_T = \prod_{t=1}^{T} x_t$, is then also log-normal:

$$X_T \sim LN\left(T\psi, T\eta^2\right)\;. \tag{A36}$$

It is obvious that $Skew\left(\log(X_T)\right) = 0$, since $\log(X_T)$ is normally distributed. The asymmetry measure of Neuberger and Payne (2021), $NPSkew$, is defined as in equation (17), which can be rewritten as

$$NPSkew\left(X_T\right) \equiv \frac{E\left[6\left(\tilde{X}_T e^{\tilde{X}_T} + \tilde{X}_T - 2e^{\tilde{X}_T} + 2\right)\right]}{E\left[2\left(e^{\tilde{X}_T} - 1 - \tilde{X}_T\right)\right]^{3/2}}\;, \tag{A37}$$

where $\tilde{X}_T \equiv \log(X_T) - \tilde{m}T$. The constant $\tilde{m}$ is defined such that $E[X_T] = e^{\tilde{m}T}$. Note that this implies $E\left[e^{\tilde{X}_T}\right] = e^{-\tilde{m}T}E[X_T] = 1$, such that equation (A37) can be rewritten as

$$NPSkew(X_T) \equiv \frac{6E\left[\tilde{X}_T e^{\tilde{X}_T}\right] + 6E\left[\tilde{X}_T\right]}{\left(-2E\left[\tilde{X}_T\right]\right)^{3/2}} \ . \tag{A38}$$

Also note that $E[X_T] = e^{\tilde{m}T}$ implies

$$\tilde{m} \equiv \psi + \frac{\eta^2}{2} \tag{A39}$$

in the log-normal case. The distributional assumption in (A36) thereby implies

$$E\left[\tilde{X}_T\right] = E\left[\log(X_T) - \tilde{m}T\right] = -\frac{T\eta^2}{2} \qquad \text{and} \qquad E\left[\tilde{X}_T e^{\tilde{X}_T}\right] = \frac{T\eta^2}{2} \ . \tag{A40}$$

Substituting the above into equation (A38) implies $NPSkew(X_T) = 0$.

Hinkley's (1975) measure of asymmetry, $QSkew_\alpha$, is defined in equation (15). The distributional assumption in (A36) implies that for any $0 < a < 1$,

$$Q_a(X_T) = \exp\left(T\psi + \sqrt{T}\eta\Phi^{-1}(a)\right) \ , \tag{A41}$$

where $\Phi^{-1}(\cdot)$ denotes the inverse cdf of the standard normal distribution. Using the above in equation (15) gives an analytical formula for $QSkew_\alpha$.

The asymmetry measure of Groeneveld and Meeden (1984), $QSkew$, is calculated according the formula in equation (16). The formula can be rewritten as

$$\begin{aligned} QSkew(X_T) &= \frac{E[X_T] - Q_{0.5}(X_T)}{E\left[|X_T - Q_{0.5}(X_T)|\right]} \\ &= \frac{E[X_T \mid X_T > Q_{0.5}(X_T)] + E[X_T \mid X_T \le Q_{0.5}(X_T)] - 2Q_{0.5}(X_T)}{E[X_T \mid X_T > Q_{0.5}(X_T)] - E[X_T \mid X_T \le Q_{0.5}(X_T)]} \ , \end{aligned} \tag{A42}$$

where the first equality is shown by Groeneveld and Meeden (1984), and the second equality is a straightforward reformulation. The distributional assumption in (A36) implies

$$\begin{aligned} Q_{0.5}(X_T) &= \exp(T\psi) \ , \\ E[X_T \mid X_T > Q_{0.5}(X_T)] &= 2\exp\left(T\psi + \frac{T\eta^2}{2}\right)\Phi\left(\sqrt{T}\eta\right) \ , \\ E[X_T \mid X_T \le Q_{0.5}(X_T)] &= 2\exp\left(T\psi + \frac{T\eta^2}{2}\right)\Phi\left(-\sqrt{T}\eta\right) \ , \end{aligned} \tag{A43}$$

where $\Phi(\cdot)$ denotes the standard normal cdf. Substituting the above into equation (A42) gives an analytical formula for $QSkew$.

C.2 The log-NIG distribution

The Normal Inverse Gaussian (NIG) distribution, introduced by Barndorff-Nielsen (1997), is a four-parameter distribution. Let the random variable z follow a NIG distribution, i.e.

$$z \sim NIG(\varphi, \beta, \nu, \delta) \ , \tag{A44}$$

where $\varphi \geq 0$, β, ν, and $\delta > 0$ are the parameters with $\varphi^2 - \beta^2 > 0$. The pdf is

$$\xi(y; \varphi, \beta, \nu, \delta) \equiv \frac{\varphi\delta \exp\left(\delta\sqrt{\varphi^2 - \beta^2} + \beta(y-\nu)\right)}{\pi\sqrt{\delta^2 + (y-\nu)^2}} K_1\left(\varphi\sqrt{\delta^2 + (y-\nu)^2}\right) \ , \tag{A45}$$

where K_1 is the modified Bessel function of the third kind. The corresponding cdf is

$$\Psi(y; \varphi, \beta, \nu, \delta) \equiv \int_{-\infty}^{y} \xi(v; \varphi, \beta, \nu, \delta)\, dv \ . \tag{A46}$$

The distribution is closed under affine transformations, and in particular

$$a + bz \sim NIG\left(\frac{\varphi}{|b|}, \frac{\beta}{b}, a + b\nu, |b|\,\delta\right) \ , \tag{A47}$$

The moment-generating function of the distribution is

$$E[\exp(kz)] = \exp\left(k\nu + \delta\left(\sqrt{\varphi^2 - \beta^2} - \sqrt{\varphi^2 - (\beta+k)^2}\right)\right) \ , \tag{A48}$$

and correspondingly, its first four moments are

$$\begin{array}{ll} E[z] = \nu + \delta\frac{\beta}{\sqrt{\varphi^2-\beta^2}} & Var(z) = \delta\frac{\varphi^2}{\sqrt{\varphi^2-\beta^2}^3} \\ Skew(z) = 3\frac{\beta}{\varphi\sqrt{\delta\sqrt{\varphi^2-\beta^2}}} & Kurt(z) = 3 + \frac{3}{\delta\sqrt{\varphi^2-\beta^2}}\left(1 + 4\frac{\beta^2}{\varphi^2}\right) \end{array} \tag{A49}$$

The normal distribution (with mean μ and volatility σ) is nested in the NIG distribution with $\nu = \mu$, $\beta = 0$, $\delta = \sigma^2\varphi$, and letting $\varphi \to \infty$. Another convenient feature of the distribution is that it is closed under convolution in the following sense: if $z_1, z_2, ..., z_T$ are iid $NIG(\varphi, \beta, \nu, \delta)$ variables with parameters, then

$$\sum_{t=1}^{T} z_t \sim NIG(\varphi, \beta, T\nu, T\delta) \ . \tag{A50}$$

Assume that $z_t \sim NIG(\varphi, \beta, \nu, \delta)$ are iid for $t = 1, ..., T$, and the single-period gross return on an asset or portfolio is $x_t \equiv \exp(z_t)$. That is, the single-period gross return x_t follows a log-NIG distribution, which we denote

$$x_t \equiv \exp(z_t) \sim LNIG(\varphi, \beta, \nu, \delta) \ . \tag{A51}$$

The pdf of x_t can be easily derived from the pdf of z_t, similarly to any exponentially transformed random variable:

$$f_{x_t}(y) = \frac{1}{y}\xi(\log(y); \varphi, \beta, \nu, \delta) \ . \tag{A52}$$

We now provide some properties of the log-NIG distribution. First, since the normal distribution is a special case of the NIG distribution, the log-normal distribution is a special case of the log-NIG distribution. Second, the first four non-central moments of x_t are given by equation (A48), since $E\left[x_t^k\right] = E[\exp(kz_t)]$. Therefore, the parameters φ, β, ν, and δ can be set by matching the first four (non-central) moments. Third, the compound return over T periods, $X_T = \prod_{t=1}^{T} x_t = \exp\left(\sum_{t=1}^{T} z_t\right)$, also follows a log-NIG distribution as implied by equation (A50):

$$X_T \sim LNIG(\varphi, \beta, T\nu, T\delta) \tag{A53}$$

Finally, we provide the alternative measures of asymmetry discussed in the paper, when X_T is distributed as in (A53). $\log(X_T)$ in this case has a NIG distribution shown in (A50), and therefore its skewness (see in (A49)) is

$$Skew(\log(X_T)) = \frac{3\beta}{\varphi\sqrt{T\delta\sqrt{\varphi^2 - \beta^2}}} \ . \tag{A54}$$

For *NPSkew*, let us find the value of $\tilde{m}$ first. $E[X_T] = e^{\tilde{m}T}$ implies

$$\tilde{m} \equiv \nu + \delta\left(\sqrt{\varphi^2 - \beta^2} - \sqrt{\varphi^2 - (\beta+1)^2}\right). \tag{A55}$$

It can be shown that

$$
\begin{aligned}
E\left[\tilde{X_T}\right] &= T\delta\left(\frac{\beta}{\sqrt{\varphi^2-\beta^2}} - \sqrt{\varphi^2-\beta^2} + \sqrt{\varphi^2-(\beta+1)^2}\right), \\
E\left[\tilde{X}_T e^{\tilde{X}_T}\right] &= T\delta\left(\frac{(\beta+1)}{\sqrt{\varphi^2-(\beta+1)^2}} - \sqrt{\varphi^2-\beta^2} + \sqrt{\varphi^2-(\beta+1)^2}\right) \ .
\end{aligned} \tag{A56}
$$

Substituting the above into equation (A38) gives an analytical formula for $NPSkew\left(X_T\right)$.

The quantiles of X_T can be calculated via

$$
Q_a(X_T) = \exp\left(\Psi^{-1}\left(y;\varphi,\beta,T\nu,T\delta\right)\right) \ , \tag{A57}
$$

where Ψ^{-1} is the inverse of the cdf in (A46). It can also be shown that

$$
\begin{aligned}
&E\left[X_T \mid X_T > Q_{0.5}\left(X_T\right)\right] = \\
&= 2\exp\left(T\nu + T\delta\left(\sqrt{\varphi^2-\beta^2} - \sqrt{\varphi^2-(\beta+1)^2}\right)\right) \times \left(1-\Psi\left(\log\left(Q_{0.5}\left(X_T\right)\right);\varphi,\beta+1,T\nu,T\delta\right)\right) \\
&E\left[X_T \mid X_T \leq Q_{0.5}\left(X_T\right)\right] = \\
&= 2\exp\left(T\nu + T\delta\left(\sqrt{\varphi^2-\beta^2} - \sqrt{\varphi^2-(\beta+1)^2}\right)\right) \times \Psi\left(\log\left(Q_{0.5}\left(X_T\right)\right);\varphi,\beta+1,T\nu,T\delta\right) \ ,
\end{aligned} \tag{A58}
$$

Substituting the (A57) and (A58) into (A42) leads to a formula for $QSkew(X_T)$. Note that numerical integration has to be used to calculate the functions Ψ and Ψ^{-1}, but these are analytical formulas in the same sense as any formula using the functions Φ or Φ^{-1} (i.e., the cdf and inverse cdf of the standard normal distribution).

References

Andersen, T. G., T. Bollerslev, F.X. Diebold, and H. Ebens, 2001. The Distribution of Realized Stock Return Volatility. *Journal of Financial Economics* 61, 43-76.

Andersen, T.G., L. Benzoni, and J. Lund, 2002. An Empirical Investigation of Continuous-Time Equity Return Models. *Journal of Finance* 57, 1239-1284.

Arditti, F., and H. Levy, 1975. Portfolio Efficiency Analysis in Three Moments: The Multiperiod Case. *Journal of Finance* 30, 797-809.

Barndorff-Nielsen, O.E., 1997. Normal Inverse Gaussian Distributions and Stochastic Volatility Modelling, *Scandinavian Journal of Statistics* 24, 1-13.

Bessembinder, H., 2018. Do Stocks Outperform Treasury Bills? *Journal of Financial Economics* 129, 440-457.

Bessembinder, H., T.F. Chen , G. Choi, and K.C. Wei, 2019. Do Global Stocks Outperform Treasury Bills? Working paper, available at https://ssrn.com/abstract=3415739

Bowley, A.L., 1920. *Elements of Statistics*, Scribner's, New York, USA.

Campbell, J.Y., and L. Hentschel, 1992. No News Is Good News: An Asymmetric Model of Changing Volatility in Stock Returns, *Journal of Financial Economics* 31, 281-318.

Cecchetti, S.G., P.S. Lam, and N.C. Mark, 1990. Mean Reversion in Equilibrium Asset Prices, *American Economic Review* 80, 398-418.

Chabi-Yo, F., 2012. Pricing kernels with Stochastic Skewness and Volatility Risk. *Management Science* 58, 624–640.

Conine, T.E., and M.J. Tamarkin. 1981. On Diversification Given Asymmetry in Returns. *Journal of Finance* 36, 1143–1155.

Cutler, D.M., J.M. Poterba, and L.H. Summers, 1991. Speculative Dynamics. *Review of Economic Studies* 58, 529-546.

Dittmar, R.F., 2002. Nonlinear Pricing kernels, Kurtosis Preference, and Evidence from the Cross Section of Equity Returns. *Journal of Finance* 57, 369–403.

Ensthaler, L., O. Nottmeyer, G. Weizsäcker, C. Zankiewicz, 2018. Hidden Skewness: On the Difficulty of Multiplicative Compounding Under Random Shocks. *Management Science* 64, 1693-1706.

Eraker, B., 2004. Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices. *Journal of Finance* 59, 1367-1403.

Eraker, B., J. Johannes, and N. Polson, 2003. The Impact of Jumps in Volatility and Returns. *Journal of Finance* 58, 1269-1300.

Fama, E., and K. French, 1988. Permanent and Temporary Components of Stock Prices. *Journal of Political Economy* 96, 246-273.

Fama, E., and K. French, 2018. Long-Horizon Returns. *The Review of Asset Pricing Studies* 8, 232-252.

Ghysels, E., A. Plazzi, and R. Valkanov, 2016. Why Invest in Emerging Markets? The Role of Conditional Return Asymmetry. *Journal of Finance* 71, 2145-2191.

Groeneveld, R.A., and G. Meeden, 1984. Measuring skewness and kurtosis. *The Statistician* 33, 391-399.

Guidolin, M., and A. Timmermann, 2008. International Asset Allocation under Regime Switching, Skew, and Kurtosis Preferences. *Review of Financial Studies* 21, 889–935.

Heston, S.L., 1993. A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options. *Review of Financial Studies* 6, 327-343.

Hinkley, D.V., 1975. On Power Transformations to Symmetry. *Biometrika* 62, 101-111.

Joanes, D.N., and C.A. Gill, 1998. Comparing Measures of Sample Skewness and Kurtosis. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 47, 183-189.

Jondeau, E., and M. Rockinger, 2006. Optimal portfolio allocation under higher moments. *European Financial Management*, 12, 29–55.

Kane, A., 1982. Skewness Preference and Portfolio Choice. *Journal of Financial and Quantitative Analysis* 17, 15-25.

Kim, D., and S.J. Kon, 1994. Alternative Models for the Conditional Heteroscedasticity of Stock Returns. *Journal of Business* 67, 563-598.

Kim, M.J., C.R. Nelson, and R. Startz, 1991. Mean Reversion in Stock Prices? A Reappraisal of the Empirical Evidence, *Review of Economic Studies* 58, 515-528.

Kim, T.-H., H. White, 2004. On More Robust Estimation of Skewness and Kurtosis, *Finance Research Letters* 1, 56-73.

Kraus, A., and R. Litzenberger, 1976. Skewness Preference and the Valuation of Risky Assets. *Journal of Finance* 31, 1085-1100.

Langlois, H., 2013. Asset Pricing with Return Asymmetries: Theory and Tests. Working Paper, available at https://ssrn.com/abstract=2362944

Martellini, L., and V. Ziemann, 2010. Improved Estimates of Higher-Order Comoments and Implications for Portfolio Selection. *Review of Financial Studies* 23, 1467–1502.

Neuberger, A., 2012. Realized Skewness. *Review of Financial Studies* 25, 3423-3455.

Neuberger, A., and R. Payne, 2021. The Skewness of the Stock Market at Long Horizons. *Review of Financial Studies*, 34, 1572-1616.

Oh, S., and J.A. Wachter, 2019. Cross-Sectional Skewness. Working paper, The Wharton School, University of Pennsylvania.

Poterba, J., and L.H. Summers, 1988. Mean Reversion in Stock Prices: Evidence and Implications, *Journal of Financial Economics* 22, 27-59.

Richardson, M., 1993. Temporary Components of Stock Prices: A Skeptic's View, *Journal of Business and Economics Statistics* 11, 199-207.

Richardson, M., and J. Stock, 1989. Drawing Inferences from Statistics Based on Multi-year Asset Returns, *Journal of Financial Economics* 25, 323-348.

Rubinstein, M.E., 1973. The Fundamental Theorem of Parameter-Preference Security Valuation, *Journal of Financial and Quantitative Analysis* 8, 61-69.

Scott, R.C., and P.A. Horvath, 1980. On the Direction of Preference for Moments of Higher Order than the Variance. *Journal of Finance* 35, 915-919.

Serfling, R.J., 1980. *Approximation Theorems of Mathematical Statistics*. John Wiley & Sons, New York, USA.

Siegel, J.J., 2008. Stocks for the Long Run. McGraw Hill, New York.

Spierdijk, L., J.A. Bikker, and P. van den Hoek, 2012. Mean Reversion in International Stock Markets: An Empirical Analysis of the 20th Century, *Journal of International Money and Finance* 31, 228-249.

Wilkins, J.E., 1944. A Note on Skewness and Kurtosis. *Annals of Mathematical Statistics* 15, 333-335.

Figure 1: Skewness of compound returns

The graphs show the skewness of compound returns, $Skew\,(X_T)$, as a function of the compounding horizon, T, when single-period returns are iid. The expected value of the single-period gross return is $\mu = 1.01$, the volatility of the single-period return, σ, varies across the panels (see above each panel), while different single-period skewness values are represented by the different lines (see legends). The second skewness value in each panel corresponds to the log-normal case.

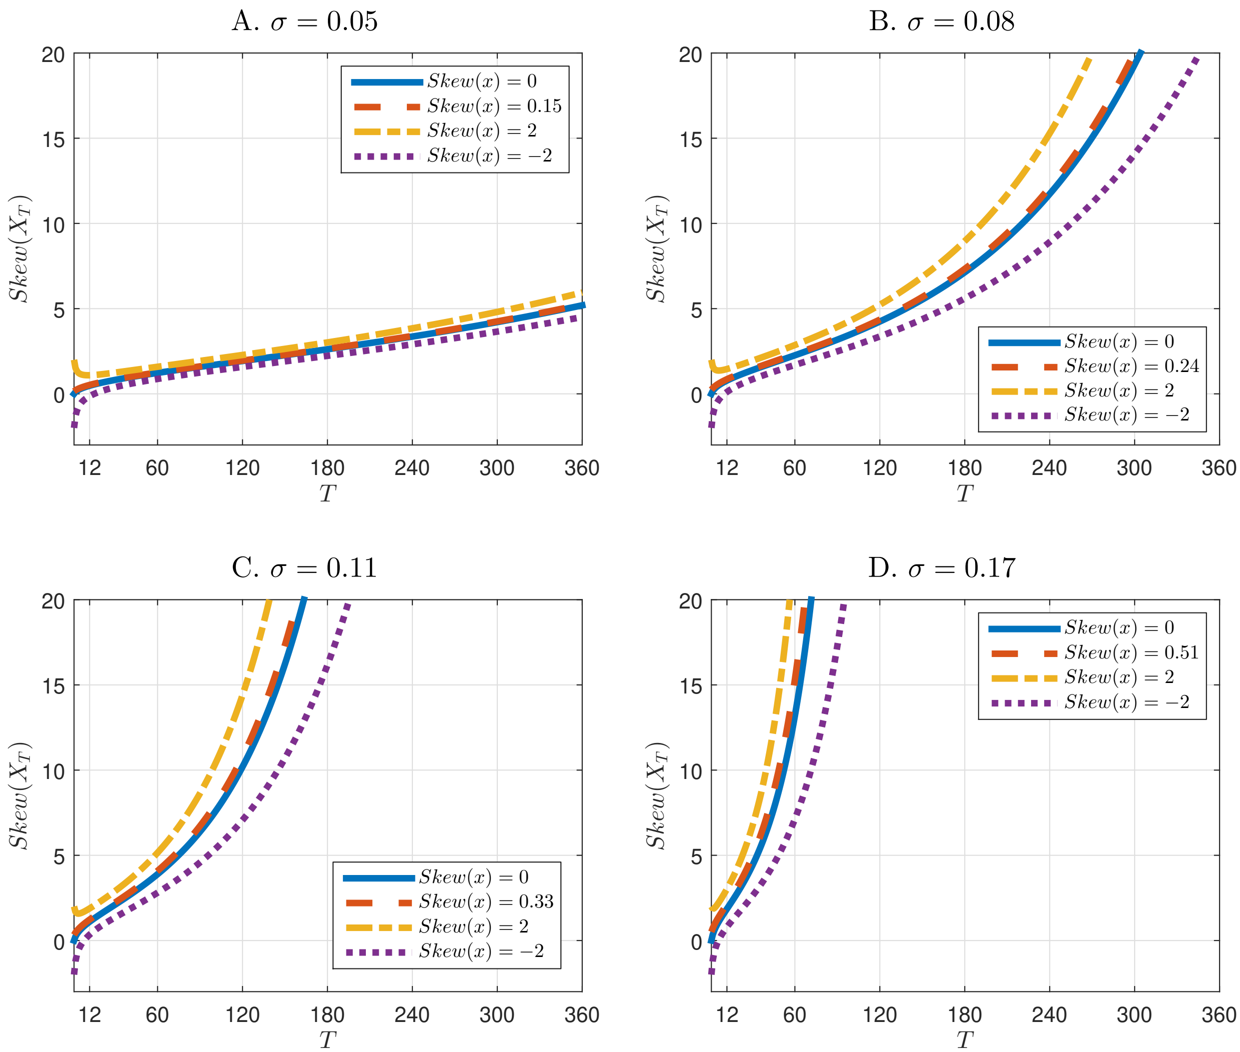

Figure 2: Distribution of the skewness estimator

The graphs display summary statistics of the distribution of the skewness estimator $\hat{S}$, when estimating the skewness of compound returns. The monthly returns are log-normal with an expected monthly return of 1% ($\mu = 1.01$) and monthly volatility, σ, that varies across the panels (see above each panel). Monthly returns are iid, and the compounding horizon is T (see above each panel). The sample size used for estimating the skewness, n, varies within each graph. The summary statistics shown are the range between the 5$^{\text{th}}$ and the 95$^{\text{th}}$ percentiles, the median (square marker), and the mean value (diamond marker) of the distribution of $\hat{S}$. The cross marker shows the upper bound of the possible values of $\hat{S}$. The true skewness value is indicated by the dotted vertical line in each graph. The summary statistics are obtained via Monte Carlo simulation using 100,000 simulated samples for each (n, σ, T) combination.

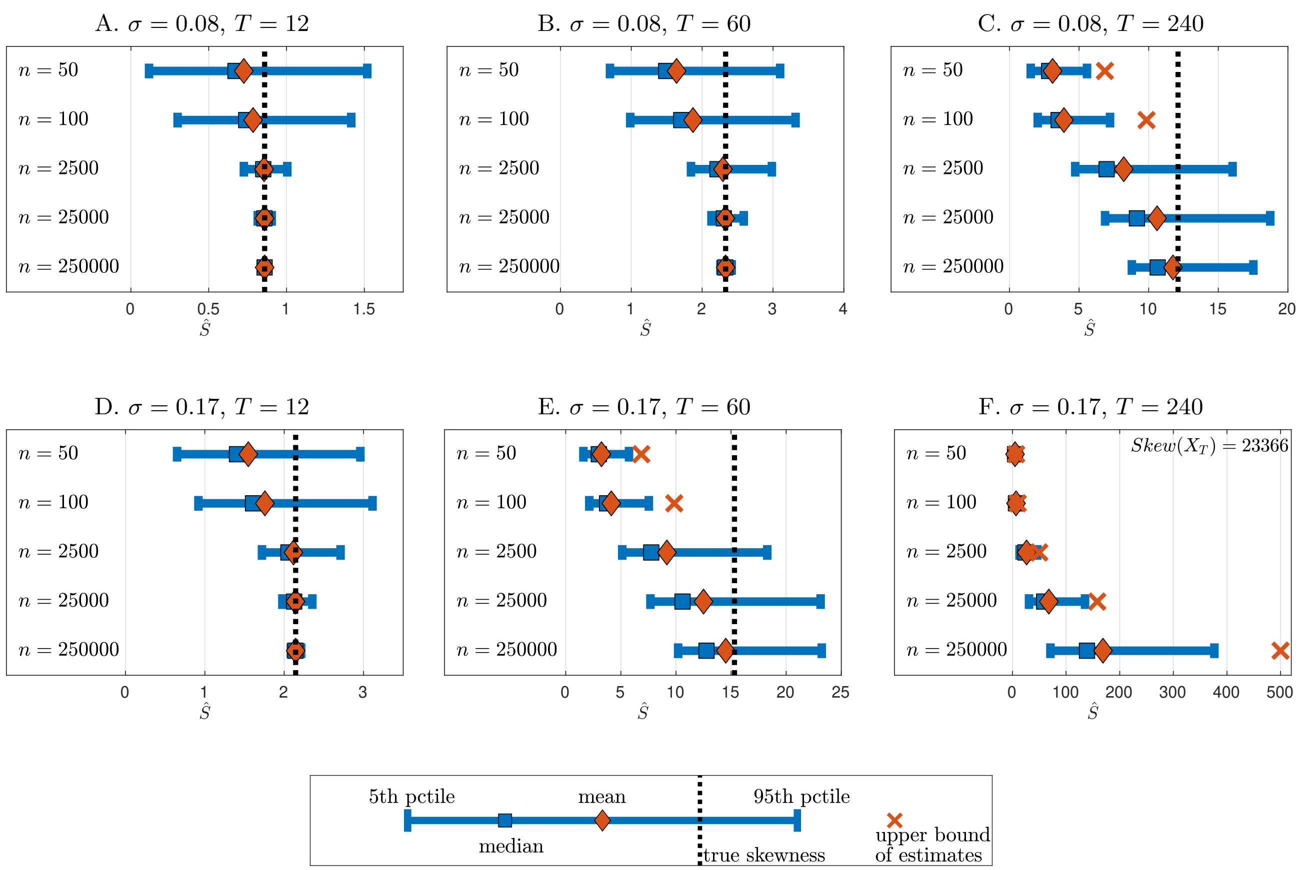

Figure 3: Asymptotic standard errors of the skewness estimator

The graph shows two-standard error bounds of the skewness estimator, $\hat{S}$, around the true parameter value, $Skew(X_T)$ (the true value is indicated by the solid line). The monthly returns are log-normal with a 1% monthly mean ($\mu = 1.01$) and 17% monthly volatility ($\sigma = 0.17$), and the compounding horizon changes along the horizontal axis. The two-standard error bounds are based on the asymptotic variance of $\hat{S}$ from equation (11). The bounds are shown for two different sample sizes, $n = 10^{15}$ and $n = 10^{20}$.

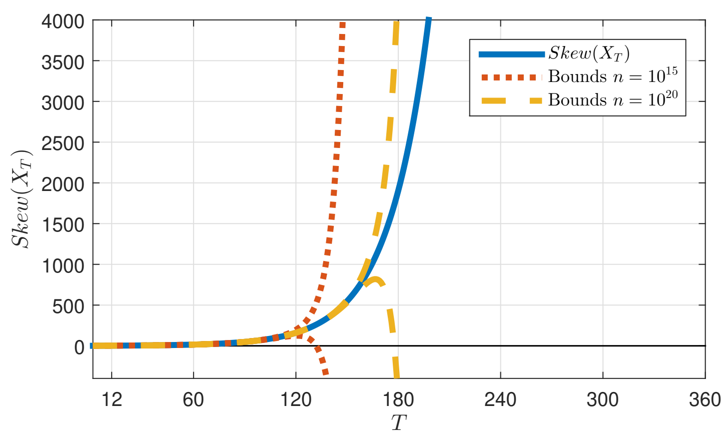

Figure 4: Skewness of compound returns - the effect of serial correlation

The graphs show the skewness of compound returns as a function of the compounding horizon, T, when single-period returns might be serially correlated. The values are calculated using equation (14). The single-period gross return is assumed to be log-normal with mean $\mu = 1.01$ and standard deviation σ that varies across the panels (see above each panel). The lines correspond to cases where single-period returns are independent across time (iid), or are serially correlated ($VR = 0.9$ or 0.8). VR is the ratio between the long-run and short-run variance of the single-period return, defined in equation (13).

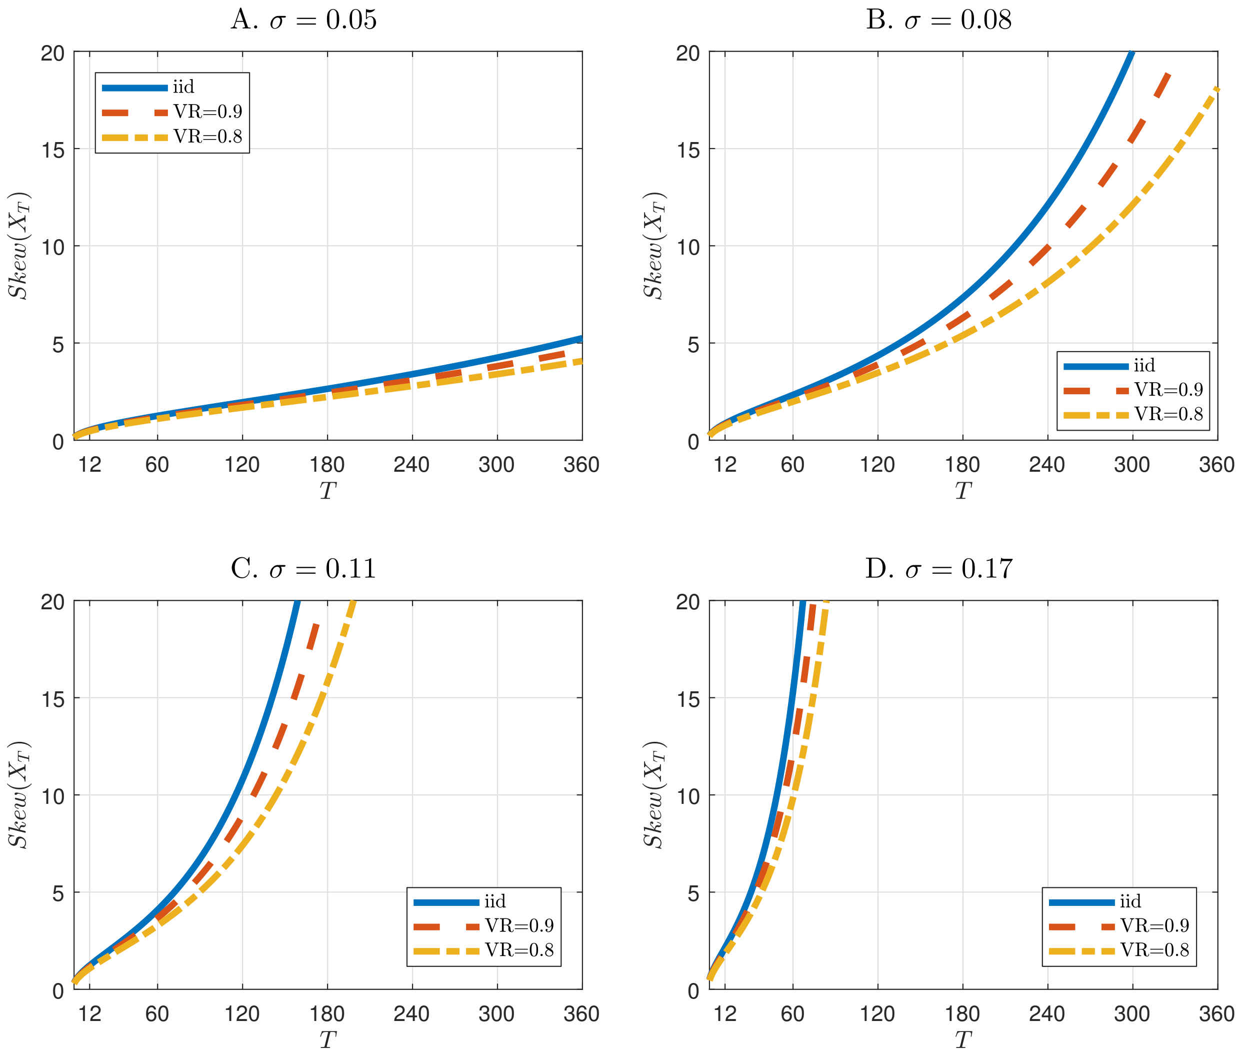

Figure 5: Alternative measures of asymmetry

The graphs show three measures of asymmetry for compound returns: the quantile-based measure of Groeneveld and Meeden (1984) (*QSkew*), the skewness of log returns (*logSkew*), and the asymmetry measure of Neuberger and Payne (2021) (*NPSkew*). Monthly returns are iid and the compounding horizon is T months (see the x-axis). The monthly expected return is 1% and the monthly volatility, σ, varies across the graphs (see above each panel). Three distributions for monthly returns are considered in each graph: (i) log-normal (dashed line), (ii) log-NIG with $Skew(x) = -2$ and $Kurt(x) = 15$ (solid line), and (iii) log-NIG with $Skew(x) = 2$ and $Kurt(x) = 15$ (dash-dotted line).

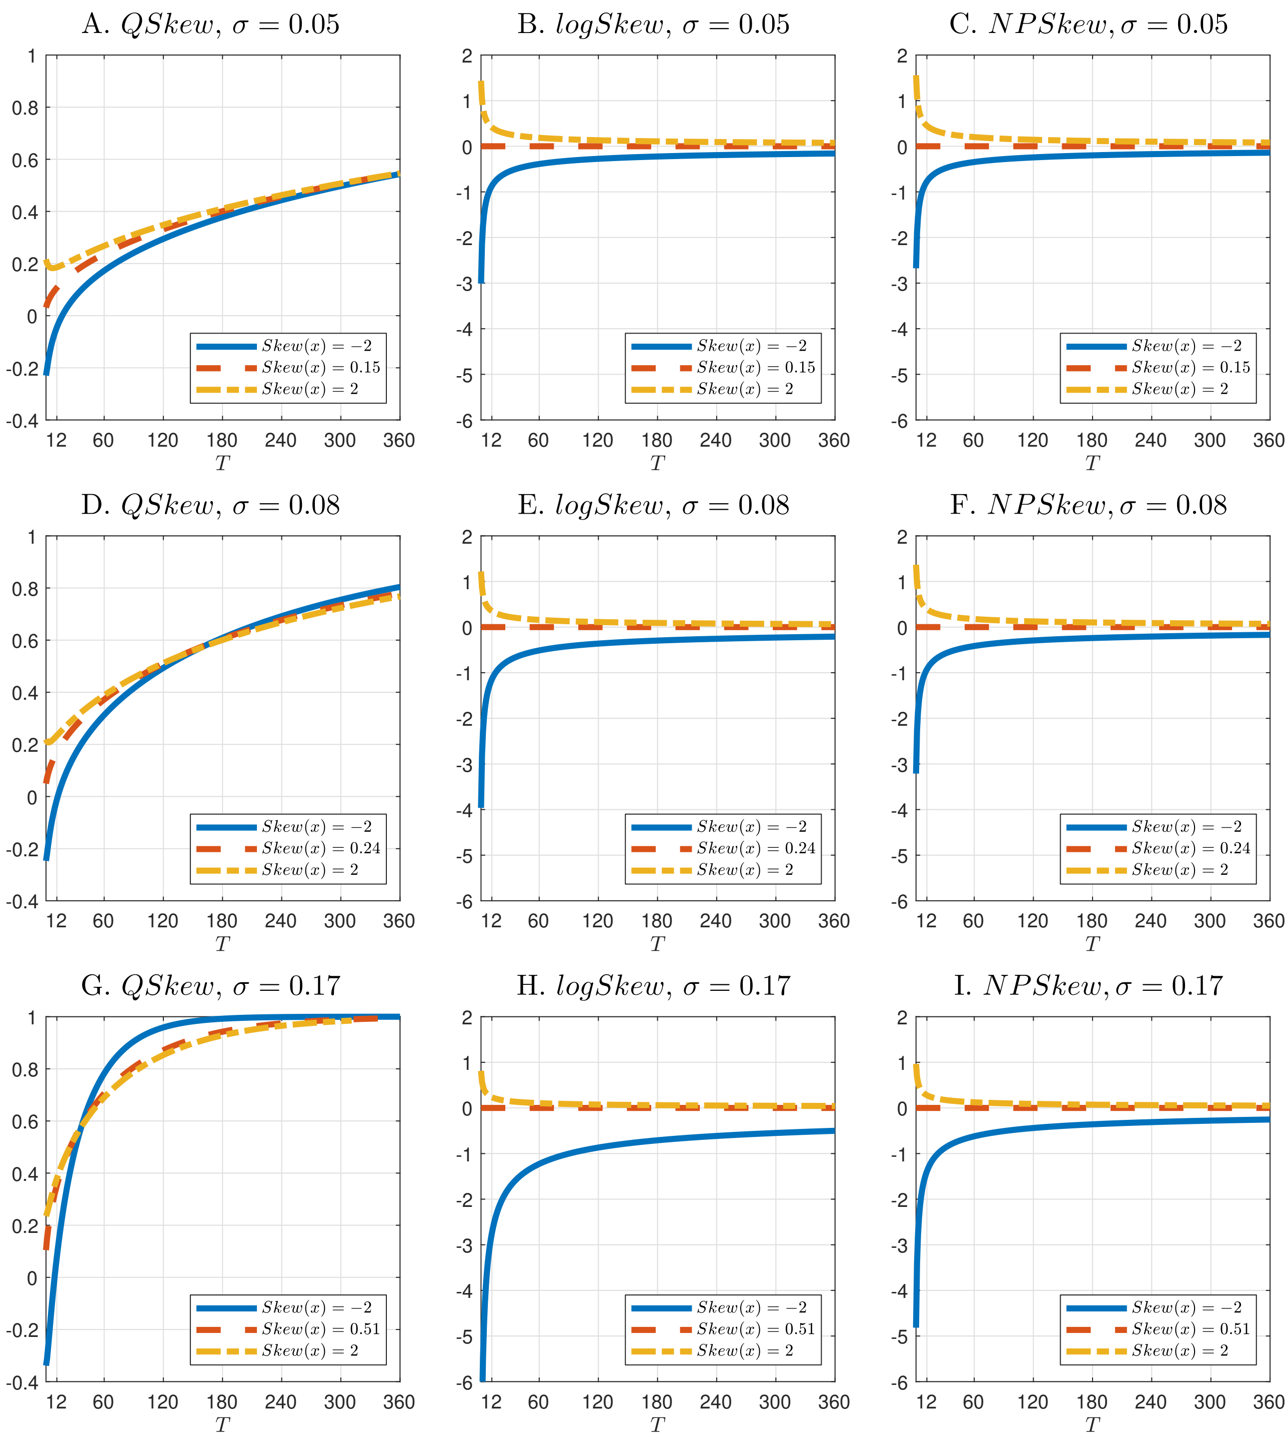

Figure 6: Analytical densities of log-normally distributed compound returns

The graphs show the density of compound returns for different compounding horizons, T, when monthly returns are iid log-normally distributed. The expected value of the single-period (monthly) gross returns are $\mu = 1.01$. Each line in each panel corresponds to different single-period volatilities, σ (see legends). The compounding horizon, T, is varied across the panels. The mean of the compound returns is indicated with a solid vertical line.

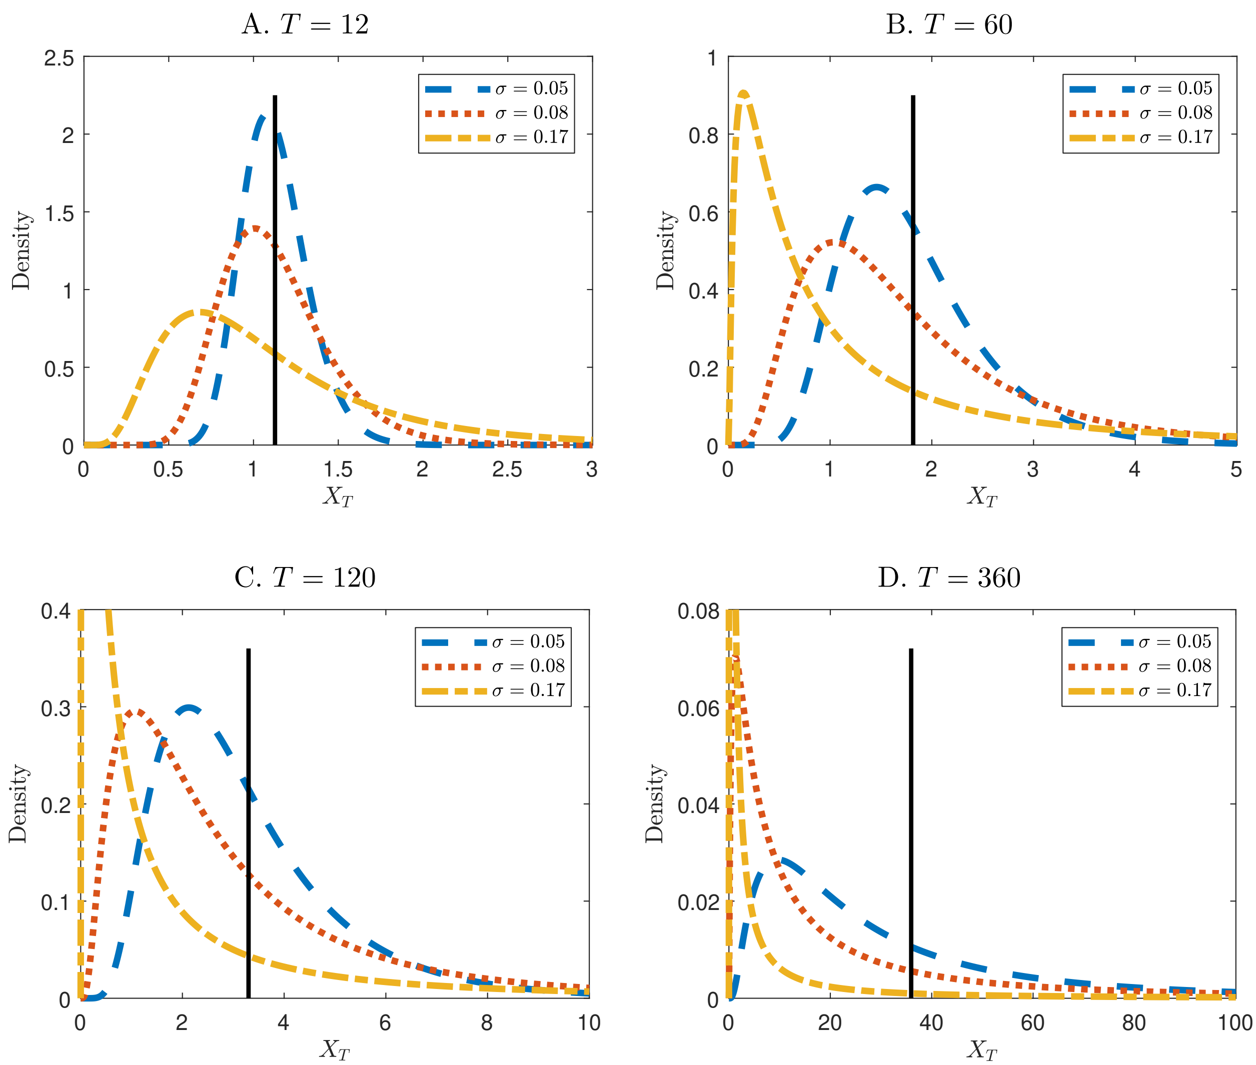

Figure 7: Densities of 5-year compound returns for various monthly return processes

The graphs show the density of compound returns for a compounding horizon of 5 years ($T = 60$ months) for different data generating processes and different single-period volatilities. The expected value of the single-period gross return is $\mu = 1.01$ in all cases, the volatility of the single-period return, σ, is varied across the panels (see above each panel). In the "log-normal" case, single-period returns are iid log-normal; the corresponding density is analytical. In the "log-NIG" case, single-period returns are iid and the higher order moments are set to $Skew(x) = -1$ and $Kurt(x) = 5$; the corresponding density is analytical. "SV" refers to the stochastic volatility with standard leverage model described in Section 2.3.2; the "SV" densities are obtained through kernel estimates of simulated data using 10^7 observations. The mean of the compound returns is indicated with a solid vertical line.

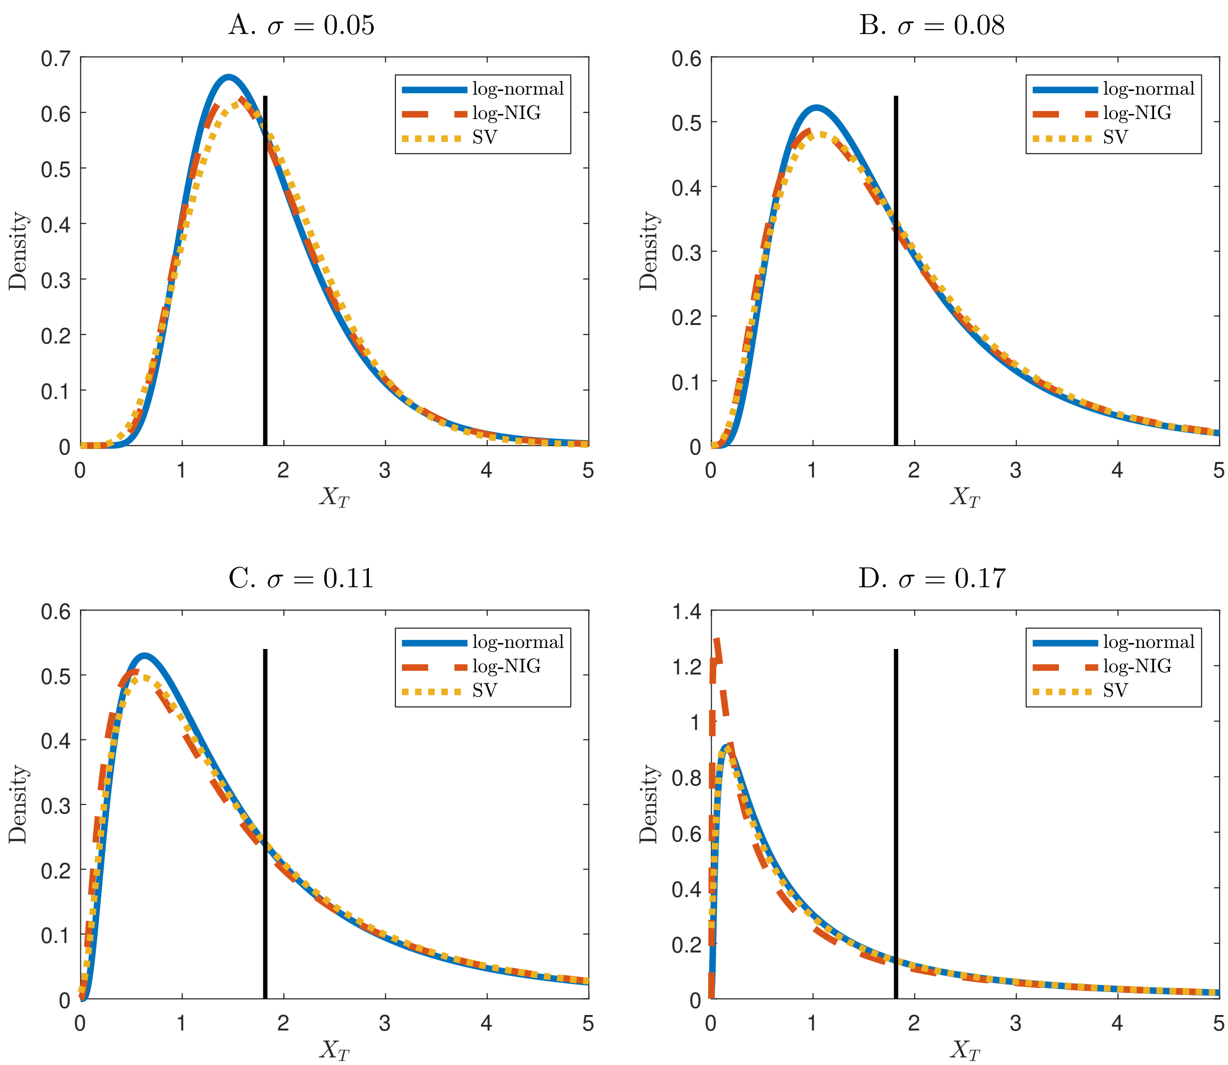

Figure 8: Ranking two assets using Taylor approximations of expected utility

Panel A shows the expected utility calculated via the third order version of (19) with $\gamma = 5$ for two assets as a function of the investment horizon T. The monthly return moments on the assets are $\mu_A = \mu_B = 1.01$, $\sigma_A = 0.05$, $\sigma_B = 0.1$, and $s_A = s_B = 0$. Panel B shows the threshold horizon, T^*; the investor prefers asset B (the high-volatility asset) over horizons longer than T^*. The values of γ and σ_B are varied in Panel B, while the other parameters are the same as in Panel A. Panel C shows the expected utility calculated via (19) with $\gamma = 5$ for two assets with monthly return moments $\mu_A = \mu_B = 1.01$, $\sigma_A = \sigma_B = 0.08$, $s_A = 0$, $s_B = -1$, and $Kurt(x_A) = Kurt(x_B) = 5$. Panel D shows the threshold horizon, T^{**}; the investor prefers asset B (the negatively skewed asset) over horizons longer than T^{**}. The values of γ and the common volatility $\sigma = \sigma_A = \sigma_B$ are varied in Panel D, while the other parameters are the same as in Panel C. The returns on the assets are iid across time in all cases.

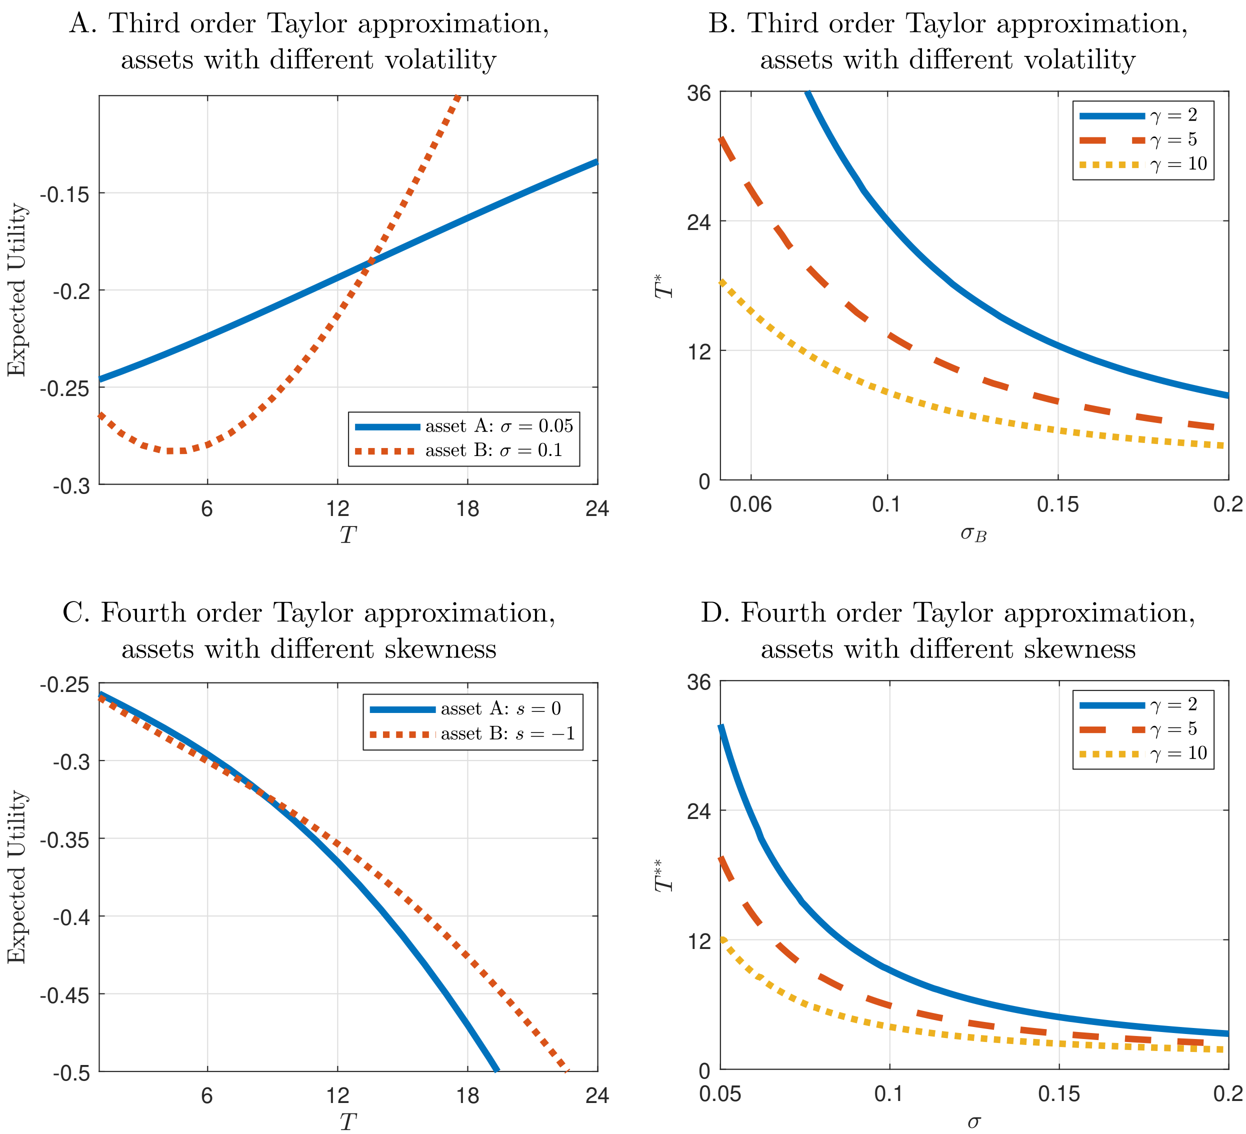

Table 1: Skewness of compound returns

The table shows the skewness of compound returns, $Skew\left(X_T\right)$, for various compounding horizons T (in different rows), when single-period returns are iid. The values are calculated using equation (7). The expected value of the single-period gross return is 1% ($\mu = 1.01$), the volatility of the single-period return, σ, varies across the columns, while the skewness of the single-period return, s, varies across the panels.

T	σ =0.05	σ =0.08	σ =0.11	σ =0.17
A. $s = 0$				
1	0.00	0.00	0.00	0.00
12	0.48	0.78	1.10	1.86
24	0.72	1.21	1.78	3.45
60	1.23	2.26	3.89	13.2
120	1.93	4.24	10.2	123
240	3.36	11.7	68.5	1.4×10^4
360	5.19	32.1	513	1.6×10^6
B. The log-normal case: $s = \frac{\sigma}{\mu}\left(\frac{\sigma^2}{\mu^2} + 3\right)$				
1	0.15	0.24	0.33	0.51
12	0.52	0.86	1.23	2.14
24	0.75	1.27	1.90	3.87
60	1.26	2.33	4.09	15.3
120	1.95	4.35	10.8	161
240	3.40	12.1	75.7	2.3×10^4
360	5.25	33.5	595	3.6×10^6
C. $s = 2$				
1	2.00	2.00	2.00	2.00
12	1.09	1.46	1.88	3.03
24	1.18	1.77	2.54	5.22
60	1.59	2.86	5.15	23.4
120	2.28	5.23	14.4	351
240	3.85	15.3	125	1.1×10^5
360	5.95	46.1	1.3×10^3	3.7×10^7
D. $s = -2$				
1	-2.00	-2.00	-2.00	-2.00
12	-0.13	0.11	0.34	0.79
24	0.26	0.65	1.06	2.00
60	0.88	1.70	2.81	7.13
120	1.58	3.37	7.06	42.1
240	2.91	8.91	37.3	1.6×10^3
360	4.50	22.3	209	6.6×10^4

Table 2: Asymmetry measures of compound returns from non-iid data.

The table shows four measures of asymmetry for compound returns: the standard skewness (*Skew*), the quantile-based measure of Groeneveld and Meeden (1984) (*QSkew*), the skewness of log returns (*logSkew*), and the asymmetry measure of Neuberger and Payne (2021) (*NPSkew*). The expected value of the monthly gross return is $\mu = 1.01$, and the (average) monthly return volatility, σ, is varied across Panels A to D (see in the title of each panel). Four data generating processes are considered: "LN" refers to the case where single-period returns are iid log-normal (the corresponding results are analytical), while the others are the stochastic volatility models described in Section 2.3.2 with standard leverage ("SV"), no leverage ("SV-NL"), and high leverage ("SV-HL"). Results for the stochastic volatility models are obtained from simulated data using 10^7 observations.

	Skew				*QSkew*				*logSkew*				*NPSkew*			
T	LN	SV	SV-NL	SV-HL	LN	SV	SV-NL	SV-HL	LN	SV	SV-NL	SV-HL	LN	SV	SV-NL	SV-HL
A. $\sigma = 0.05$																
1	0.15	−0.30	0.26	−0.57	0.03	−0.09	0.02	−0.15	0	−0.60	−0.06	−0.87	0	−0.55	0.00	−0.82
12	0.52	−0.06	0.69	−0.40	0.11	−0.03	0.10	−0.10	0	−0.75	−0.07	−1.08	0	−0.67	0.00	−0.99
24	0.75	0.25	0.90	−0.03	0.15	0.04	0.15	−0.02	0	−0.60	−0.05	−0.85	0	−0.53	0.00	−0.78
60	1.26	0.82	1.38	0.59	0.24	0.16	0.23	0.13	0	−0.40	−0.04	−0.57	0	−0.35	0.00	−0.52
120	1.95	1.49	2.11	1.25	0.33	0.28	0.33	0.26	0	−0.29	−0.03	−0.41	0	−0.26	0.00	−0.37
240	3.40	2.70	3.58	2.37	0.46	0.43	0.46	0.41	0	−0.21	−0.02	−0.30	0	−0.18	0.00	−0.27
360	5.25	4.11	5.56	3.58	0.55	0.53	0.55	0.51	0	−0.17	−0.02	−0.24	0	−0.15	0.00	−0.22
B. $\sigma = 0.08$																
1	0.24	−0.04	0.31	−0.22	0.05	−0.03	0.04	−0.06	0	−0.38	−0.03	−0.55	0	−0.34	0.00	−0.51
12	0.86	0.44	0.98	0.21	0.17	0.08	0.17	0.04	0	−0.47	−0.04	−0.67	0	−0.42	0.00	−0.62
24	1.27	0.87	1.40	0.64	0.24	0.17	0.24	0.14	0	−0.37	−0.03	−0.54	0	−0.33	0.00	−0.49
60	2.33	1.83	2.49	1.58	0.37	0.33	0.37	0.31	0	−0.25	−0.02	−0.36	0	−0.22	0.00	−0.33
120	4.35	3.45	4.68	3.04	0.51	0.49	0.51	0.47	0	−0.18	−0.02	−0.26	0	−0.16	0.00	−0.24
240	12.1				0.68	0.67	0.68	0.66	0	−0.13	−0.01	−0.19	0	−0.11	0.00	−0.17
360	33.5				0.78	0.78	0.78	0.77	0	−0.11	−0.01	−0.15	0	−0.09	0.00	−0.14
C. $\sigma = 0.11$																
1	0.33	0.12	0.38	−0.01	0.07	0.01	0.06	−0.01	0	−0.27	−0.03	−0.39	0	−0.25	0.00	−0.37
12	1.23	0.86	1.34	0.66	0.23	0.17	0.23	0.14	0	−0.34	−0.03	−0.49	0	−0.30	0.00	−0.45
24	1.90	1.48	2.05	1.26	0.33	0.28	0.33	0.26	0	−0.27	−0.03	−0.39	0	−0.24	0.00	−0.35
60	4.09	3.28	4.59	2.87	0.50	0.47	0.50	0.46	0	−0.18	−0.02	−0.26	0	−0.16	0.00	−0.24
120	10.8				0.66	0.65	0.67	0.64	0	−0.13	−0.01	−0.19	0	−0.12	0.00	−0.17
240	75.7				0.83	0.83	0.84	0.83	0	−0.09	−0.01	−0.13	0	−0.08	0.00	−0.12
360	595				0.92	0.92	0.92	0.92	0	−0.08	−0.01	−0.11	0	−0.08	0.00	−0.10
D. $\sigma = 0.17$																
1	0.51	0.38	0.55	0.29	0.10	0.07	0.10	0.05	0	−0.18	−0.02	−0.26	0	−0.16	0.00	−0.24
12	2.14	1.76	2.29	1.54	0.35	0.32	0.36	0.30	0	−0.22	−0.02	−0.32	0	−0.20	0.00	−0.29
24	3.87	3.19	4.14	2.81	0.48	0.46	0.49	0.45	0	−0.17	−0.02	−0.25	0	−0.16	0.00	−0.23
60	15.3				0.71	0.70	0.71	0.69	0	−0.12	−0.01	−0.17	0	−0.10	0.00	−0.15
120	161				0.87	0.87	0.88	0.87	0	−0.09	−0.01	−0.12	0	−0.08	0.00	−0.11
240	2.3×10^4				0.97	0.98	0.98	0.98	0	−0.06	−0.01	−0.09	0	−0.04	0.00	−0.07
360	3.6×10^6				0.99	1.00	1.00	1.00	0	−0.05	0.00	−0.07	0	−0.04	0.00	−0.07

Table 3: Mean, median, and probability of beating the mean for log-normal returns.

The table shows the mean, the median, and the probability of the returns exceeding the mean for different combinations of single-period volatility, σ, and compounding horizon T. The single-period returns are iid log-normal, with a 1% expected return ($\mu = 1.01$).

	Mean	Median				Probability of returns>mean			
T		$\sigma = 0.05$	$\sigma = 0.08$	$\sigma = 0.11$	$\sigma = 0.17$	$\sigma = 0.05$	$\sigma = 0.08$	$\sigma = 0.11$	$\sigma = 0.17$
1	1.01	1.01	1.01	1.00	1.00	49.0%	48.4%	47.8%	46.7%
12	1.13	1.11	1.09	1.05	0.95	46.6%	44.6%	42.5%	38.6%
24	1.27	1.23	1.18	1.10	0.91	45.2%	42.3%	39.5%	34.1%
60	1.82	1.69	1.51	1.28	0.79	42.4%	38.0%	33.7%	25.9%
120	3.30	2.85	2.27	1.63	0.62	39.3%	33.2%	27.6%	18.0%
240	10.9	8.12	5.14	2.65	0.38	35.1%	27.0%	20.0%	9.8%
360	35.9	23.14	11.66	4.30	0.24	31.9%	22.7%	15.1%	5.6%