

The Impact of Skew on Performance and Bias

How Skew Distorts Short Term Performance, Triggers Bias, and Changes
Drawdowns

January 2019

Dr. Zachary Dugan, zack.dugan@isam.com, 561-990-2021, Corresponding Author, ISAM, 5100 Town Center Circle, Tower II Suite 430, Boca Raton FL, 33486. Zack Dugan is a Research Scientist at ISAM and a Visiting Senior Research Scientist at the Johns Hopkins University Institute for Data Intensive Science and Engineering and SciServer.

Dr. Alexander Greyserman, alex.greyserman@isam.com, 561-990-2021, ISAM. Alex Greyserman is the Chief Scientist of ISAM and an Adjunct Professor in the Columbia University Department of Mathematics.

The Impact of Skew on Performance and Bias

How Skew Distorts Short Term Performance, Triggers Bias, and Changes Drawdowns

January 2019

Abstract

In this paper, we explore the relationship between the statistic skew and known behavioral biases. We investigate the impact that skew has on the perception of performance as a function of time, and we show that negative skew artificially improves performance over the short term, while positive skew has the opposite effect. We quantify the relationship between skew and drawdown depth and length, and we show that negative skew increases drawdown depth and length, and that again positive skew does the opposite. Finally, we explore the relationship between skew, volatility, and drawdown, and we show that negative skew amplifies the increase that volatility causes in drawdown depth and length, while positive skew has a corresponding dampening effect.

Keywords: skew, volatility, drawdown, behavioral biases, loss aversion, recency bias

1 Introduction

Certain investment strategies pick up the proverbial pennies in front of a moving train. This approach results in a high probability of long sequences of small gains at the risk of a catastrophic loss. Other strategies, such as venture capital, depend on a small number of investments being homeruns at the cost of many small losers. Trend following, like venture capital, brings positive skew to investors' portfolios. This statistical attribute stems from the nature of the system, frequent but small losses in search of the rare but large wins. Dugan and Greyserman [2018] document how this aspect of the strategy can make it painful for investors. Dugan, Greyserman, and Friccione [2017] explore the relationship a strategys volatility and Sharpe have with drawdown depth and length. However, the simulations drew from normal distributions with no skew.

In this paper, we build on concepts from both papers by incorporating skew into the analysis. Through simulation, we demonstrate that skew can distort perception of performance over short periods of time, and we document the relationship between skew and behavioral biases. We quantify the impact that skew has on drawdown depth and length under various volatilities, returns, and time durations. Ultimately, we show that while negative skew appeals to inherent human biases, it increases both drawdown depth and length, and it amplies the increase in drawdown depth that volatility causes.

In Section 2, we review the behavioral biases relevant to this analysis. We demonstrate the role that skew plays in probability distributions of returns in Section 3. In Section 4, we present results from simulations of outperformance as a function of skew and time period. We show the impact that skew has on drawdown depth and length in Section 5. In Section 6, we conclude. We review our manufacture of skew through gamma distributions and skew normal distributions in Appendices A and B, and in Appendix C, we describe the full return construction.

2 Behavioral Biases

Dugan and Greyserman [2018] describe how behavioral biases impact the positively skewed quantitative investing strategy, trend following. Here, we quickly review two of the biases discussed in that paper: loss aversion and recency bias. In 1979, Amos Tversky and Daniel Kahneman introduced Loss Aversion to describe the theory that people dislike losing two times as much as they enjoy winning, and that people tend to magnify losses in their minds.[11, 17, 12, 3, 4, 8, 14, 15] As Dugan and Greyserman [2018] document, this bias negatively influences investor's perception of return series with frequent losses.

Recency bias is the over-emphasis of recent information. It is the financial version of a cognitive bias called the availability heuristic, in which people tend to make predictions based on information that is usually more recent, available, and easy to remember. In finance, this can be particularly problematic when the same calculation on different sets of time yield different results.[7, 10, 16] When investors place greater importance on recent data, they lose comprehensive perspective, which can lead to poor decisions.

As an example, 100% of the profits time series momentum strategies come from the top 10% of the months.[6] Positive skew dictates that if an investor analyzed only a small, more recent sample, and missed one of the big months, he or she would form a very different opinion on the strategy than if they analyzed the entire history. Thus, the nature of the return series with positive skew make those strategies seem less valuable to investors swayed by recency bias.

For a more detailed look at these behavioral biases and others as they pertain to quantitative time-series momentum strategies, see Dugan and Greyserman [2018] and references therein.

3 Probability Distributions

Probability Distribution Functions (PDFs) of returns show the probabilities of various returns randomly occurring and can help visualize how the probabilities of specific returns change with skew. More specifically, PDFs can be modified to generate specific returns, volatilities, and skews.

In Exhibit 1, we show gamma distributions of daily returns with various negative skews that each produce return series with annual compounding returns of 20% and volatilities of 20% over long periods of time. We describe the construction of gamma distributions in Appendix A. The average daily return necessary to produce this return annual is 0.0726%. As the skew decreases, the mode of the distribution grows increasingly larger than the average. However, the mean of the distribution remains the same because the probability of large negative returns increases simultaneously. Thus, for negatively skewed distributions, the most common returns may be small and positive, but the probability of a large negative return is also much greater.




Exhibit 1: Probability Distribution Functions of Gamma Distributions of negative skew that each produce returns of 20% and volatilities of 20%. Returns are daily and compounding. Average return for all distributions is plotted in black.

In Exhibit 2, we show distributions of daily returns with various positive skews

that all produce return series with the same annual compounding returns of 20% and volatilities of 20% over long periods of time. As the skew increases, the mode of the distribution grows increasingly negative, which may seem difficult to reconcile with the distributions' positive annual return. This apparent contradiction is possible because the mean of each distribution remains identical. As the skew increases, the probability of increasingly large and positive returns also increases. Therefore, while positively skewed distributions may frequently deliver negative returns, the probability of a large positive return is also much greater.




Exhibit 2: Probability Distribution Functions of Gamma Distributions of positive skew that each produce returns of 20% and volatilities of 20%. Returns are daily and compounding. Average return for all distributions is plotted in black.

To help visualize the probability of outperformance, in Exhibit 3, we plot the cumulative probability distribution functions (CDFs) of return series with varying skews, which show the probabilities that a return series will be equal to or less than a given benchmark return. Exhibit 3 shows that series with negative skew have a greater likelihood of delivering large, negative returns of values equal to or less than -1.5%. Returns with positive skew are more likely to deliver negative values. In fact, series with skews of 1.6 are over 20% more likely to deliver negative returns than return series with skews of -1.6. This reality alone will push investors swayed by recency bias and loss aversion to prefer a strategy with negative skew.

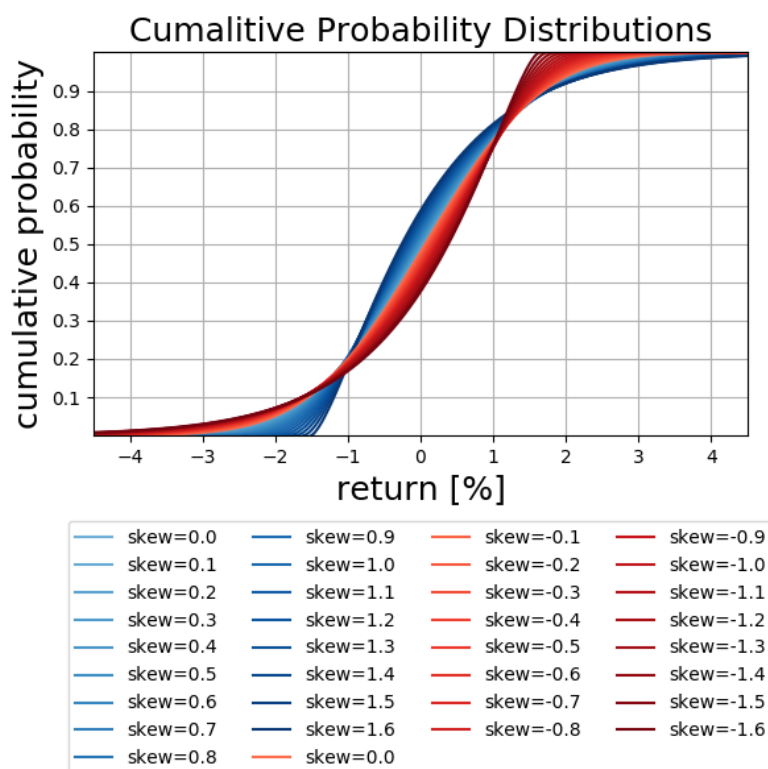


Exhibit 3: Cumulative Probability Distribution Functions of Gamma Distributions that each produce returns of 20% and volatilities of 20%. Returns are daily and compounding.

On the other hand, return series with positive skew have a greater probability than those of negative skew of delivering outsized positive returns, of values equal to or greater than 1.5%. Also interesting, we see the CDFs of positive and negative skews cross over each other twice. The asymmetric nature of the gamma distributions, also visible in Exhibits 1 and 2, causes these cross overs.

We also consider the relative probability distributions of return series with both different skews and different annual returns. Here, we focus on scenarios in which the negatively skewed return series have annual returns of 8% and the positively skewed return series have annual returns of 16%. In Exhibit 4, we plot the PDFs of series with 8% annual returns, 20% volatilities, and varying negative skews. In Exhibit 5, we plot the PDFs of series with 16% annual returns, 20% volatilities, and varying positive skews. If these distributions appear similar to those in Exhibits 1 and 2, it is because they are. The average return in the first scenario, for both the positively and negatively skewed series, is 0.0726%. In this scenario, the average return for the negatively skewed and positively skewed series respectively are 0.0307% and 0.0591%. In turn, the distributions in this scenario will be shifted slightly to the left on the x axis. This shift requires a close look to discern in Exhibits 4 and 5.




Exhibit 4: Probability Distribution Functions of Gamma Distributions of negative skew that each produce annual returns of 8% and annual volatilities of 20%. Returns are daily and compounding. Average return for all distributions is plotted in black.




Exhibit 5: Probability Distribution Functions of Gamma Distributions of positive skew that each produce returns of 16% and volatilities of 20%. Returns are daily and compounding. Average return for all distributions is plotted in black.

Likewise, the CDFs of these distributions with returns of 8% and 16%, shown in Exhibit 6, are also shifted to the left along the x axis from their counterpart distributions in Exhibit 3. Surprisingly, return series with skews of 1.6 and annual returns of 16% have a 60% chance to be negative, while return series with annual returns of 8% and skews of -1.6 are less than a 40% chance of being negative. In spite of delivering annualized returns that are two times greater, the positively skewed return series are still 20% more likely to be negative. Skew is the sole reason for this paradox. Furthermore, negative skew will drive investors susceptible to recency bias and loss aversion to prefer a return series with fewer losses even though the risk adjusted returns are only half as good.

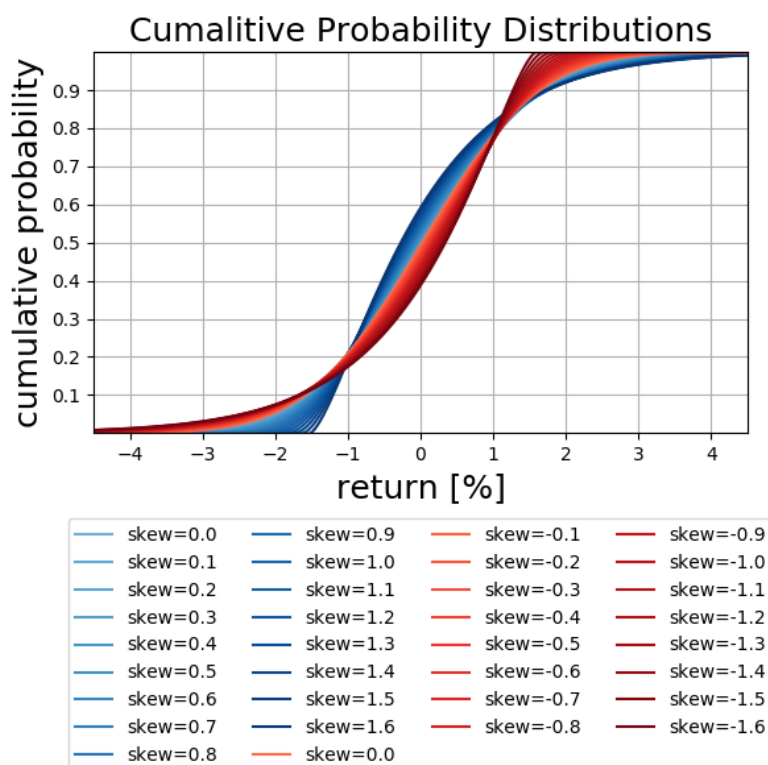


Exhibit 6: Cumulative Probability Distribution Functions of Gamma Distributions of negative skew that each produce annual returns of 8% and annual volatilities of 20% and Cumulative Probability Distribution Functions of positive skew that each produce returns of 16% and volatilities of 20%. Returns are daily and compounding.

4 Outperformance

Exhibits 1-6 each show how the nature of skew can influence our ability to judge the quality of performance. In particular, small time scales with fewer data points will magnify these misperceptions. As we alluded to in the previous section, negatively skewed returns series should outperform on smaller time scales. Conversely, positively skewed return series are more likely to deliver small losses over small time scales and not the outsized winners that create positive skew. In this section, we analyze large sets of simulated returns to quantify the relationship between outperformance, skew, and time horizon.

We first compare returns with zero skew to returns with varying skew over a range of time periods. We prescribe 20% volatilities for all returns. Employing gamma distributions described in Appendix A and the computational methods described in Appendix C, we fix the skew of the first series at zero, and we vary the skew of the second from -10 to 10. These values are fixed over a 20 year period. We note that a skew of positive 10 is nearly impossible for returns over a 20 year period and that large negative skews are rare but not without precedent. We explore this extreme range to more definitively demonstrate the impact that skew has on out-

performance in general. Over 20 year periods, we fix the Sharpe ratios of both the zero skew series and the varied skew series at 1.0. Therefore, we know that over a 20 year period, neither series should outperform the other. However, when we look at smaller subsets of these return series and analyze outperformance, the percentage of outperformance is no longer 50%/50%. We vary the length of the subsets from 5 days to 2 years and summarize the comparison in Exhibit 7. In Exhibit 8, we plot the percentage of time the return series with the varying skew outperforms the zero skew series.

Quantity	Fixed Skew Return Series	Varying Skew Return Series
20-year Volatility	0.20	0.20
20-year Sharpe	1.0	1.0
20-year Skew	0.0	-10.0 - 10.0

Exhibit 7: Prescribed characteristics of return series in the first case.




Exhibit 8: Frequency of short term outperformance between zero skew return series with 20% volatility and returns of 20% over 20 years and return series with the same volatility and returns but varying skews. Exhibit 7 lists the values of the comparison.

When both skews are equal to zero, the percentage of outperformance is 50% as one would expect. However, on smaller time scales, Exhibit 8 shows that the negatively skewed return series outperform zero skew series by up to 56% of the time. Likewise, the most positively skewed return series underperform the zero skew series by up to 44% on short scales. Furthermore, the magnitude of these outperformances can be striking. Exhibit 8 shows that an investor susceptible to loss

aversion and recency bias will inherently lean toward negative skew despite the reality that all the return series in this scenario have the same long term performance.

Next, we compare return series with skews of -1.6 to return series with skews ranging from -10 to 10. Again, we prescribe 20% volatilities for all series. Over 20 year periods, we fix the Sharpe ratios of the -1.6 skew series at 0.4, and we fix the sharpe ratios of the series with varying skew at 0.8. Therefor, we know that over a 20 year period, the second return series will always outperform the first. However, when we look at smaller subsets of these return series and analyze outperformance, the percentage of outperformance is no longer 100%. As before, we vary the length of the subsets from 5 days months to 2 years. We summarize the comparison in Exhibit 9. In Exhibit 10, we plot the percentage of time the return series with 0.8 Sharpe and vary skew outperforms the return series with -1.6 skew and 0.4 Sharpe.

Quantity	Fixed Skew Return Series	Varying Skew Return Series
20-year Volatility	0.20	0.20
20-year Sharpe	0.4	0.8
20-year Skew	-1.6	-10.0 - 10.0

Exhibit 9: Prescribed characteristics of return series in the second case.

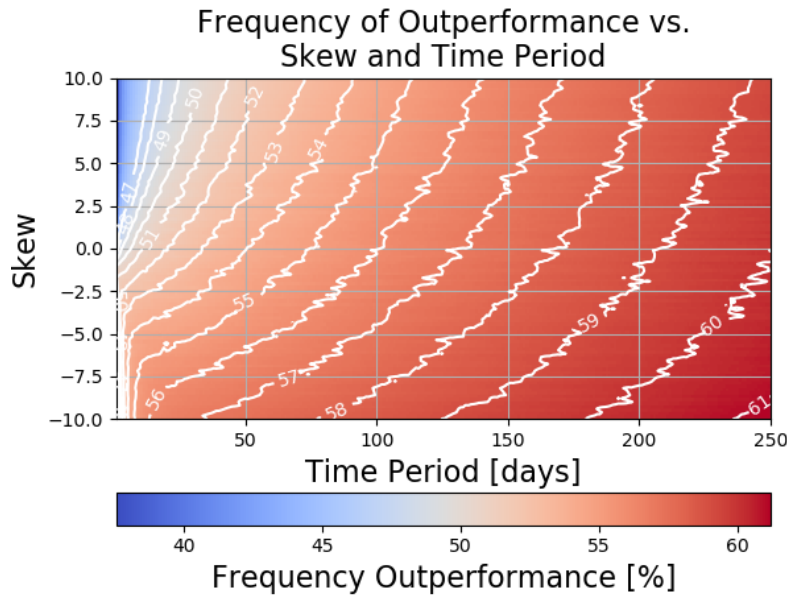


Exhibit 10: Frequency of short term outperformance between return series with 20% volatility, 8% returns, and skews of -1.6 over 20 years and return series with the same volatility, 16% returns, and varying skews. Exhibit 9 lists the values of the comparison.

Exhibit 10 shows two general trends. First, as the skew becomes more negative,

the outperformance becomes more frequent. These results are broadly in line with those from the previous experiment shown in Exhibit 8. Second, as the time period of observation increases, so does the outperformance. Again, we see that the presence of skew in returns necessitates longer windows to more accurately evaluate performance.

The results from Exhibit 10 also provide two specific points for discussion. First, it is noteworthy that even after 2 years, the returns with long term Sharpe ratios of 0.8 and varying skews are only outperforming returns with long term Sharpe ratios of 0.4 and skews of -1.6 ~60% of the time. Skew and volatility combine to cause this low frequency of outperformance. Second, and more remarkable, on the short time scale, return series with long term Sharpe ratios of 0.8 and positive skew actually underperform those with long term Sharpe ratios of 0.4 and skews of -1.6. Positive skew makes returns with twice the risk adjusted returns seem worse on short time scales.

5 Drawdowns

Greyserman, Dugan, and Friccione [2017] investigate the impact of Sharpe ratio and volatility on drawdown depth over 20-year periods. They show that volatility increases drawdown depth approximately linearly and that Sharpe ratio does not decrease drawdowns linearly. However, those simulations used returns with zero skew. In this section, we investigate the impact that skew, volatility, and time period have on drawdown depth and length.

First, we employ large sets of simulated returns to examine the effect that skew and time period have on drawdowns. For all returns, we prescribe a volatility of 20% and a Sharpe of 0.8. We vary the prescribed skew from -3 to 3, and we vary the time period we observe from 6 months to 20 years. For each combination of skew and duration, we simulate 100,000 return series and find the deepest and longest drawdowns for each. In Exhibit 11, we show the average maximum drawdown as a function of skew and time period.

On the shorter time scales, particularly less than 2 years, change in skew makes almost no difference to the maximum drawdown, while time is the main determining factor. However, as the time period increases, so does impact of skew until it becomes the dominant factor. For samples of approximately 5 years, return series with skews of 3 and -3 have maximum drawdowns of 14% and 30% respectively. The difference in skew amplifies the depth of drawdown by a factor of 2. Looking at the very long term of 20 years, return series with skews of 3 and -3 have maximum drawdowns of 22% and 45% respectively. Again, the difference in skew doubles the average maximum drawdown. The longer timescales highlight the hidden risks of negative skew and the benefits of positive skew.

We can also think about the average maximum drawdown as a metrics for underlying risk. While negative skew will probably not make a difference to drawdowns on short time scales, the possibility for greater loss still exists, but with a greater probability of being realized on longer time periods.

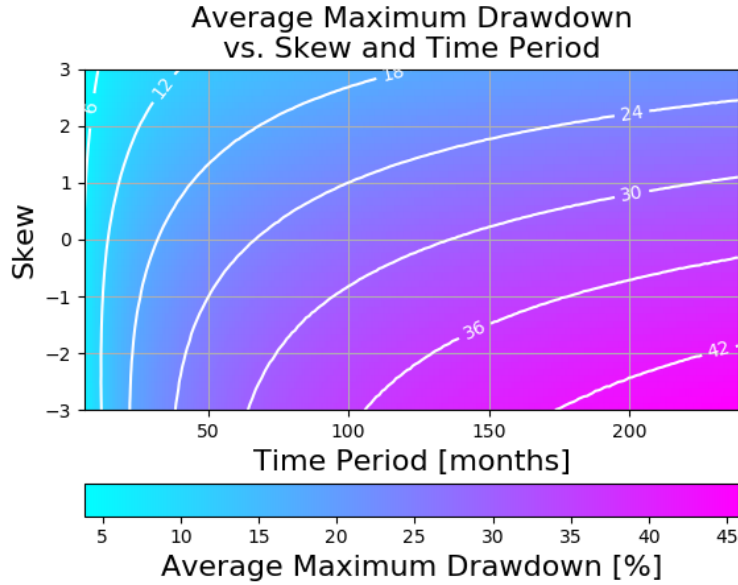


Exhibit 11: Gamma Distributions: Average maximum drawdown as a function of skew and time period for return series with volatilities of 20% and Sharpe ratios of 0.8.

To corroborate our results, we also employ a skew normal distribution, which we document in Appendix B, in place of the gamma distribution, which we document in Appendix A. In Exhibit 12, we show the average maximum drawdown as a function of time period and skew, as generated from the skew normal distributions. Because of limitations in the construction of the skew normal distribution, we examine a range of skews from -0.75 to 0.75. Again we see that on short time scales, skew has little impact on drawdown depth, but its impact grows to become the most important determinant of drawdown depth as time period increases.

The drawdown depths in Exhibit 12 agree well with those in the same range of skews in Exhibit 11. For samples of approximately 75 months, return series with skews of 0.75 and -0.75 have maximum drawdowns of 24% and 28% respectively, a difference of roughly 17% of the drawdown. Over the very long term of 20 years, return series with skews of 0.75 and -0.75 have maximum drawdowns of 32% and 40% respectively, an increase of roughly 25% of the drawdown. The magnitude of difference in drawdown depth displayed in Exhibit 11 as opposed to Exhibit 12 results from the greater range of skews, but the similarity in distribution and relationship corroborates the conclusion that negative skew increases drawdown depth.

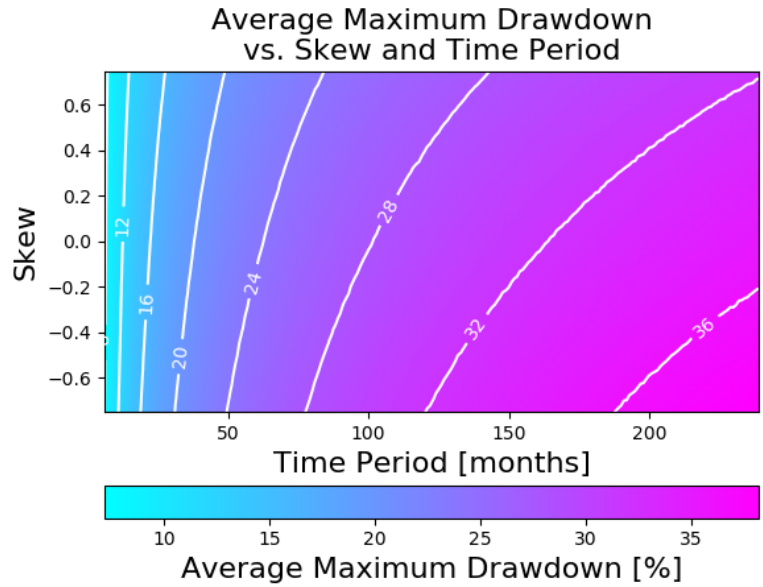


Exhibit 12: Skew Normal Distributions: Average maximum drawdown as a function of skew and time period for return series with volatilities of 20% and Sharpe ratios of 0.8.

We also analyze the impact skew has on drawdown length. In Exhibit 13, we plot the average longest drawdown as a function of skew and time period. The results again show no discernable difference between various skews on short time scales. However, on longer time scales, skew makes a substantial difference in drawdown length as well. After approximately 18 years, the average longest drawdown for series with skew of 3 we 36 months, while the mean for series with skews of -3 was 42 months, an increase of approximately 18%.

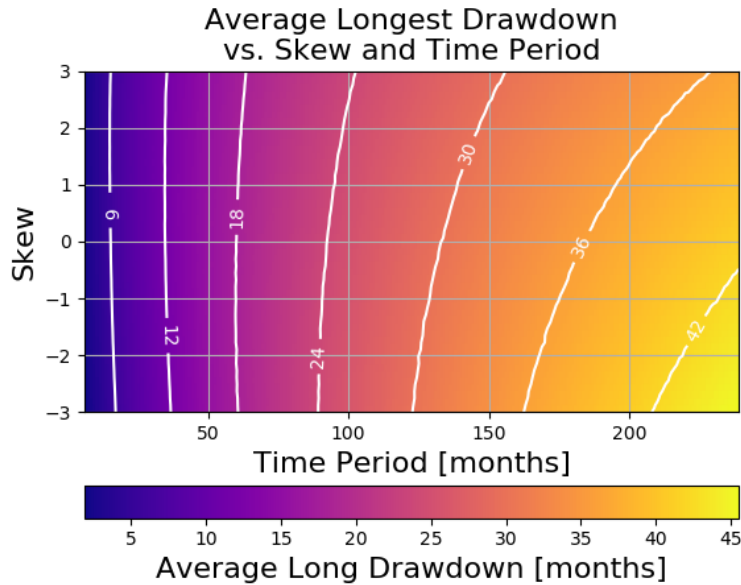


Exhibit 13: Gamma Distributions: Average longest drawdown as a function of skew and time period for return series with volatilities of 20% and Sharpe ratios of 0.8.

Together, Exhibits 11, 12, and 13 show that negative skew increases both drawdown depth and length, while positive skew decreases both. Each of the figures also shows that longer time scales are necessary to observe skew’s effects on drawdown.

We also investigate the combined impact of skew and volatility on drawdown depth. We simulate large sets of return series with Sharpe ratios of 0.8 and time periods of 20 years. We vary the skew from -3 to 3 and the volatility from 5% to 30%. In Exhibit 14, we show the average maximum drawdown as a function of skew and volatility. The results show that for a given volatility, positive skew reduces the maximum drawdowns while negative skew increases them. At 20% volatility, return series with skew of 3 have maximum drawdowns of 22% while those with skew of -3 have maximum drawdowns of 46%, more than a factor of 2 greater.

Greyserman, Dugan, and Friccione [2017] show that for normal distributions, volatility increases drawdown depth approximately linearly. Exhibit 14 also shows that positive skew dampens the increase in drawdown depth that volatility causes. For a skew of 2, the increase in volatility from 5 to 30% increases maximum drawdown by 28%. For a skew of -2, the same increase in volatility increases maximum drawdown by 46%.

In Exhibit 15, we show the average longest drawdown as a function of skew and volatility. Similarly, we find that for a given volatility, positive skew reduces the longest drawdown, while negative skew increases it. Exhibit 15 also demonstrates that positive skew dampens the increase in drawdown length that volatility causes. For a skew of 2, the increase in volatility from from 5 to 30% barely increases drawdown length. But for a skew of -2, the increase from 5-30% increases average longest drawdown length by 6 months.

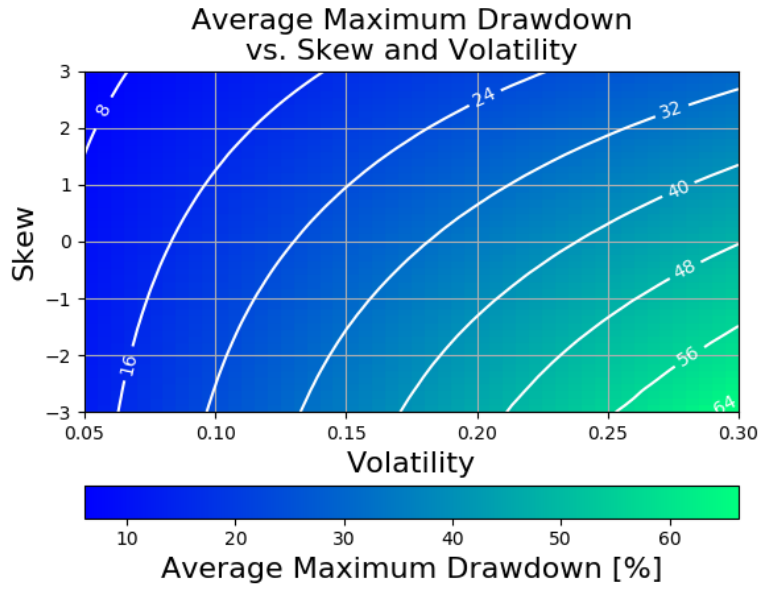


Exhibit 14: Average maximum drawdown as a function of skew and volatility for return series with Sharpe ratios of 0.8 over a 20 year period.

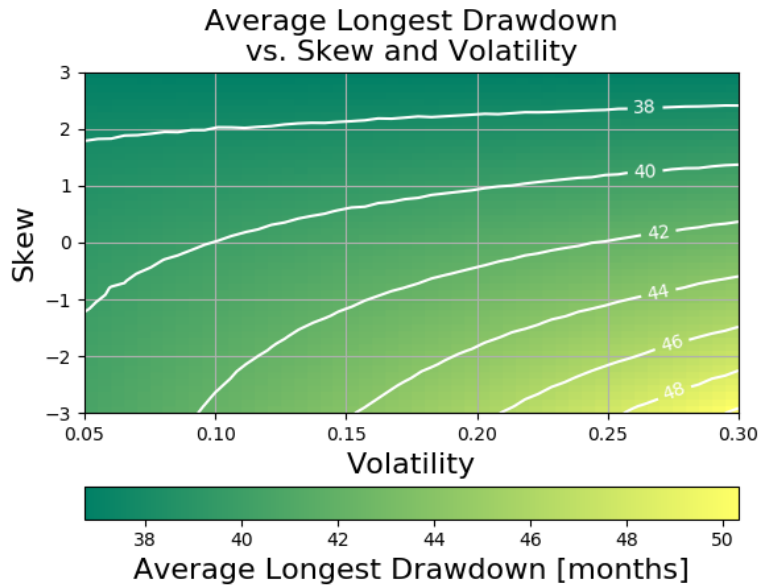


Exhibit 15: Average longest drawdown as a function of skew and volatility for return series with Sharpe ratios of 0.8 over a 20 year period.

6 Conclusion

In this paper, we show that returns with negative skew appeal to classic behavioral biases, namely recency bias and loss aversion, because of the frequent winners and

rare losers that inherently come with the distributions. We show that over short time scales, these returns frequently outperform their zero skewed and positively skewed counterparts, making them more attractive to those investors swayed by bias.

However, while returns with negative skew deliver wins more frequently over the short term, they also incur deeper and longer drawdowns over the long term. We also show that negative skew amplifies the increase that volatility causes on drawdown depth and length, while positive skew mitigates that same impact. From this perspective, positive skew allows investors to capture the outsized returns of higher volatility while shielding them from some of the requisite pain. Investors must be aware of biases toward strategies that seem attractive over the short term, but that eventually meet deeper, longer drawdowns, and in the worst cases, ruin.

Appendices

A Gamma Distributions

Gamma distributions depend on two input parameters: the shape k and the scale θ . The probability distribution function (PDF) is:[9]

$$\phi(k, \theta, x) = \frac{1}{\Gamma(k)\theta^k} x^{k-1} e^{-x/\theta} \quad (1)$$

in which $\Gamma(k)$ is the gamma function of k . In turn this produces the cumulative distribution function (CDF):

$$\Phi(k, \theta, x) = \frac{1}{\Gamma(k)} \gamma(k, x/\theta) \quad (2)$$

in which γ is the lower incomplete gamma function. In turn, these distributions produce a skewness, a mean, and a variance of:

$$s = 2/\sqrt{k} \quad (3)$$

$$\bar{x} = k\theta \quad (4)$$

$$\sigma^2(x) = k\theta^2 \quad (5)$$

B Skew Normal Distributions

The skew normal distribution is an extension of the classical normal distribution that allows for the prescription of skew. The standard normal distribution is defined as:[9]

$$\phi(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2} \quad (6)$$

for which the cumulative distribution function is:

$$\Phi(x) = \int_{-\inf}^x \phi(t)dt \quad (7)$$

Ohagna and Leonard [1976] introduce the skew normal distribution, and Azzalini [2014] documents the parametrization with shape parameter α , scale parameter ω , and location parameter ξ . The skew normal distribution has the following probability distribution:[1]

$$\phi_s(x) = \frac{2}{\omega} \phi\left(\frac{x-\xi}{\omega}\right) \Phi\left(\alpha\left(\frac{x-\xi}{\omega}\right)\right) \quad (8)$$

This distribution has the resulting skew:

$$\gamma = \frac{4-\pi}{2} \frac{(\delta\sqrt{2/\pi})^3}{(1-2\delta^2/\pi)^{3/2}} \quad (9)$$

in which:

$$\delta = \frac{\alpha}{\sqrt{1+\alpha^2}} \quad (10)$$

Setting $\delta = 1$ generates the maximum skewness of 1.

The skew normal distribution has a mean of:

$$\bar{x} = \xi + \omega\delta\sqrt{\frac{2}{\pi}} \quad (11)$$

and a mode of:

$$m = \xi + \omega m_o(\alpha) \quad (12)$$

in which:

$$m_o(\alpha) \sim \mu_z - \frac{\gamma\sigma_z}{2} - \frac{\text{sign}(\alpha)}{2} e^{-\frac{2\pi}{|\alpha|}} \quad (13)$$

in which:

$$\mu_z = \sqrt{\frac{2}{\pi}}\delta \quad (14)$$

and:

$$\sigma_z = \sqrt{1-\mu_z^2} \quad (15)$$

The skew normal distribution has a variance of:

$$\sigma^2 = \omega^2 \left(1 - \frac{2\delta^2}{\pi}\right) \quad (16)$$

C Prescribing Array Characteristics

To create a random array with a prescribed skew γ and scale ω , we first calculate δ from γ using Equation 5. We then create random normal arrays u_0 and v , which we combine in the following equation:[1]

$$u_1 = \omega(\delta u_0 + v\sqrt{1-\delta^2}) \quad (17)$$

We then multiply all negative values by -1 so that they are positive. To prescribe the location parameter ξ , we add ξ to the array:

$$u_1 = u_1 + \xi \quad (18)$$

We then prescribe the volatility of the array with the following equation:

$$u_1 = u_1 \times \sigma_p / (\sigma_{u1} \times \sqrt{n}) \quad (19)$$

in which σ_p is the prescribed volatility, σ_{u1} is the standard deviation of the array as it was, and n is the number of compounding periods per year. Next, we prescribe the compounding returns of the array using what the average returns per period should be:

$$r_p = (1 + r_a)^{1/n} - 1 \quad (20)$$

$$u_1 = u_1 - \bar{u}_1 + r_p \quad (21)$$

in which r_a is the prescribed annualized return. However, this formulation leads to realized annual returns that are systemically less than the prescribed value, though the standard deviation of realized returns is very low. The percentage error of returns varies with prescribed returns, volatility, and period of time. After making attempts at various analytical solutions, we settled on the computational solution of adding incrementally to the array until the annualized returns are within 3% of the prescribed value. The error on the prescribed skew and prescribed volatility are substantially less than 1%.

References

- [1] Azzalini, A. "The Skew Normal and Related Families." Cambridge University Press (2014), 24.
- [2] Chen M.K, Lakshminarayanan V., Santos L.R. . "How basic are behavioral biases? Evidence from capuchin monkey trading behavior." Journal of Political Economy (2006), 114, 517.
- [3] Camerer, C. "Three Cheers, Psychological, Theoretical, Empirical, for Loss Aversion." Journal of Marketing Research (2005), 42, 2, 129-133.
- [4] Chua, E. & Camerer, C., Experiments on Intertemporal Consumption with Habit Formation and Social Learning, 2003. Division of Humanities and Social Sciences, Caltech, [<https://authors.library.caltech.edu/22010/>].
- [5] Dugan, Z., Greyserman, A. "Trend and Skew Aversion." Journal of Alternative Investment (2018). [under review]
- [6] Greyserman, A., Dugan, Z., Fricchione, D. "The Nature of Drawdowns." ISAM Research Series (2017).
- [7] Healy, A.F.; Havas, D.A.; Parkour, J.T. "Comparing serial position effects in semantic and episodic memory using reconstruction of order tasks." Journal of Memory and Language (2000), 42, 147.
- [8] Ho, Teck and Juanjuan Zhang (2004), Experimental Tests of Solutions to the Double-Marginalization Problem: A ReferenceDependent Approach, Haas School of Business, University of California, Berkeley, [<http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.203.2342&rep=rep1&type=pdf>]
- [9] Hogg, R. & Craig, A. "Introduction to Mathematical Statistics." Macmillan Publishing (1958), 104-11.

- [10] Howard, M. W.; Kahana, M. "Contextual Variability and Serial Position Effects in Free Recall." *Journal of Experimental Psychology: Learning, Memory & Cognition* (1999), 25, 923.
- [11] Kahneman, D.; Tversky, A. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* (1979), 47, 263.
- [12] Novemsky, N. & Kahneman, D. The Boundaries of Loss Aversion. *Journal of Marketing Research* (2005), 42 (May), 11928.
- [13] O'Hagan, A. & Leonard, T. "Bayes Estimation Subject to Uncertainty about Parameter Constraints." *Biometrika* (1976), 63, 201-202.
- [14] Tovar, P. (2004), *The Effects of Loss Aversion on Trade Policy: Theory and Evidence*, 2004. Department of Economics, University of Maryland.
- [15] Tu, Q. "Reference Points and Loss Aversion in Intertemporal Choice," Institute of World Economics and Politics, Chinese Academy of Social Sciences; Wageningen University (December 21, 2004). [<https://ssrn.com/abstract=644142> or <http://dx.doi.org/10.2139/ssrn.644142>]
- [16] Tversky, A., Kahneman, D. "Availability: A heuristic for judging frequency and probability." *Cognitive Psychology* (1973), 5, 207.
- [17] Tversky, A.; Kahneman, D. "Loss Aversion in Riskless Choice: A Reference Dependent Model." *Quarterly Journal of Economics* (1991), 106.