Middle out: only extreme deciles matter[1]

Ashwin Alankar,[2] Philip Maymin[3] and Myron Scholes[4]

Abstract

For many return distributions, both empirical and simulated, replacing the middle with zeroes and keeping only returns in the top decile or bottom decile still yields a 90% correlation with the original and exhibits 90% of its volatility. This heuristic holds broadly across various asset classes, including 20 years of daily returns on broad equity and fixed income indexes, and regardless of whether the breakpoints are calculated on an overall or a rolling historical basis. We show that for the thin-tailed Gaussian, middling out results in a lower correlation of 80%, while for more realistic distributions such as the Student T distribution with three degrees of freedom, the correlation is again 90%, and these results also hold in the presence of changing distributions over time. Further, we document two interesting theoretical properties of the 80% middle-out cutoff for the Gaussian distribution by showing it is the only cutoff that equals the correlation and also the only cutoff whose marginal correlation exactly equals one in magnitude. Middling out is stronger when kurtosis is higher or return horizons are shorter. In addition to these quantitative correlation and volatility results, we argue that tail returns should also have an even greater importance for investment managers and asset allocators because of their outsized effects on compound returns, the difficulty in hedging or reacting to tail conditions when they do occur, and the observation that tail returns are the ones that reveal actual asset information, while the middle returns are merely middling noise. Investment professionals should apply middle-out thinking and focus their risk management and portfolio decision policies on tail events.

Keywords: middles, tails, deciles, risk, compound returns.

1 The opinions and views expressed are those of the authors and are subject to change without notice. They do not necessarily reflect the views of others in Janus Henderson's organisation and no forecasts can be guaranteed.

2 Head of Global Asset Allocation and Portfolio Manager, Janus Henderson Investors, ashwin.alankar@janushenderson.com.

3 Director of Asset Allocation Strategies, Janus Henderson Investors, philip.maymin@janushenderson.com, corresponding author.

4 Chief Investment Strategist, Janus Henderson Investors, myron.scholes@janushenderson.com.

1. Introduction

Among the American Finance Association (2022)'s top 25 cited Journal of Finance articles of all time, 14 use information from financial market prices. Of those, 13 focus on the average returns (Carhart (1997), Fama and French (1992), Jegadeesh and Titman (1993), Daniel et al (1998), Glosten et al (1993), De Bondt and Thaler (1985), Fama and French (1996), Markowitz (1952), Baker and Wurgler (2006), Ang et al. (2006), Sharpe (1964), Bansal and Yaron (2004), and Hong and Stein (1999)); the only exception is Forbes and Rigobon (2002), who distinguish market co-movement from correlation by adjusting for a correlation bias due to market volatility.

This focus on average returns is rampant in the industry as well. Portfolio managers report and are ranked by average or risk-adjusted returns (eg Morningstar (2022)), even though there is little evidence that such ratings have any ability to forecast future performance (Blake and Morey (2009)). Risk managers are effectively required to use metrics that explicitly ignore tail risks. Even the latest Basel (2017) standards for bank capital requirements, which ultimately flow through to govern virtually all risk management with its centralised definition of the "standardised approach," explicitly ignores tail risk by, for example, routinely requiring risk to be measured to within some confidence interval. One implicit problem with confidence intervals is that they highlight a single point such as the 0.05 level. A common misinterpretation is to observe volatility of, say, 15% and erroneously presume this implies a 30% loss on their fund in extreme conditions, or a ratio of 2:1. The actual losses could be as great as 45%, or a ratio of 3:1. Graph 1 illustrates by plotting on a daily rolling basis the ratio of the forward-looking one-year maximum drawdown to the backward-looking one-year realised volatility, on the S&P 500 from 1928 to 2022. Indeed, ratios in excess of 3:1 occurred 10 times more frequently than would be expected by a Gaussian distribution.

Forward-looking one-year maximum drawdown divided by the backward-looking one-year realised volatility on the S&P 500 — Graph 1

Daily rolling ratios

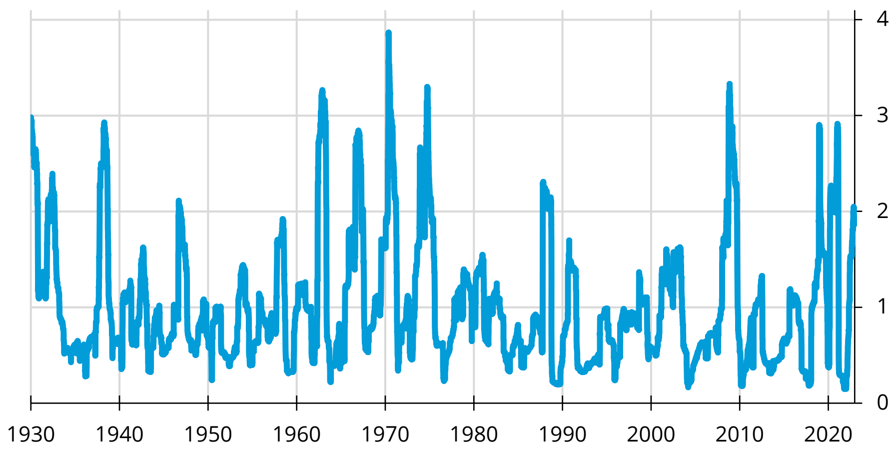

Source: Bloomberg, 1928–2022. As of Nov 2022.

We aim to argue here that this approach, while ubiquitous and conventional, is entirely upside-down. Rather than viewing tails and extreme events as unexplainable, unforeseen, unavoidable outliers that need to be ignored while focusing on the more common situations, we argue that the tails are where all the information and value is, and that the middles of the distribution are the actual noise. Even scenario analysis can't forecast the next tail event; it can only extrapolate from the prior observed tail events.

This paper proceeds as follows. Section 2 describes the simple method of middling out and shows its result on theoretical distributions. Section 3 evaluates middling out on empirical return distributions. Section 4 expands the discussion to implications for portfolio managers, asset allocators, risk managers and other investment professionals. Section 5 concludes.

2. Middle out

Middling out a distribution means transforming the distribution to keep only the values in the tails and replacing the middle values with a replacement value. Graph 2 illustrates an 80% middled-out normal distribution, meaning values in the top 10% and bottom 10% of the distribution remain unchanged while all the values in the middle 80% are replaced with zero; the probability density function of the standard normal distribution is superimposed with the middled-out distribution.

While in complete generality a middled-out distribution could use a replacement value other than zero (for example, the mean or the median), for the remainder of this paper we focus only on middle-outs with zero as the replacement value. From the perspective of calculating correlations, volatilities and similar measures, this is a conservative choice.

A standard normal and an 80% middled-out standard normal — Graph 2

Probability density functions

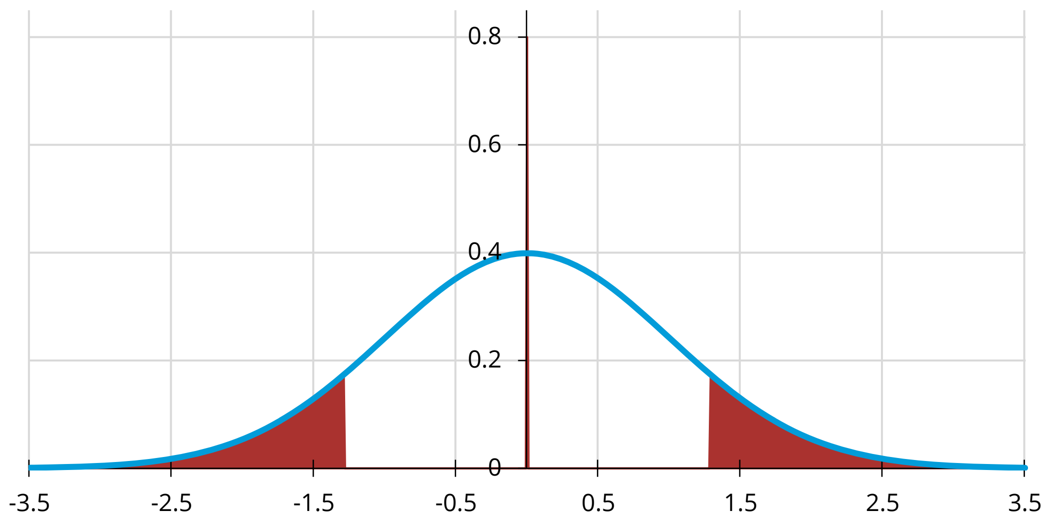

Source: Simulations.

Graph 3 computes the correlation between 10,000 simulated values from a standard normal distribution and a Student T distribution with three degrees of freedom with the middled-out distributions for a variety of middle-out percentages. For example, the 80% middled-out normal distribution has a correlation of 80.74% with its original simulated values and the 80% middled-out Student T distribution with three degrees of freedom has a correlation of 90.96%.

Correlation between middled-out and raw simulation values — Graph 3

For Student T distribution with 3 degrees of freedom, and Normal

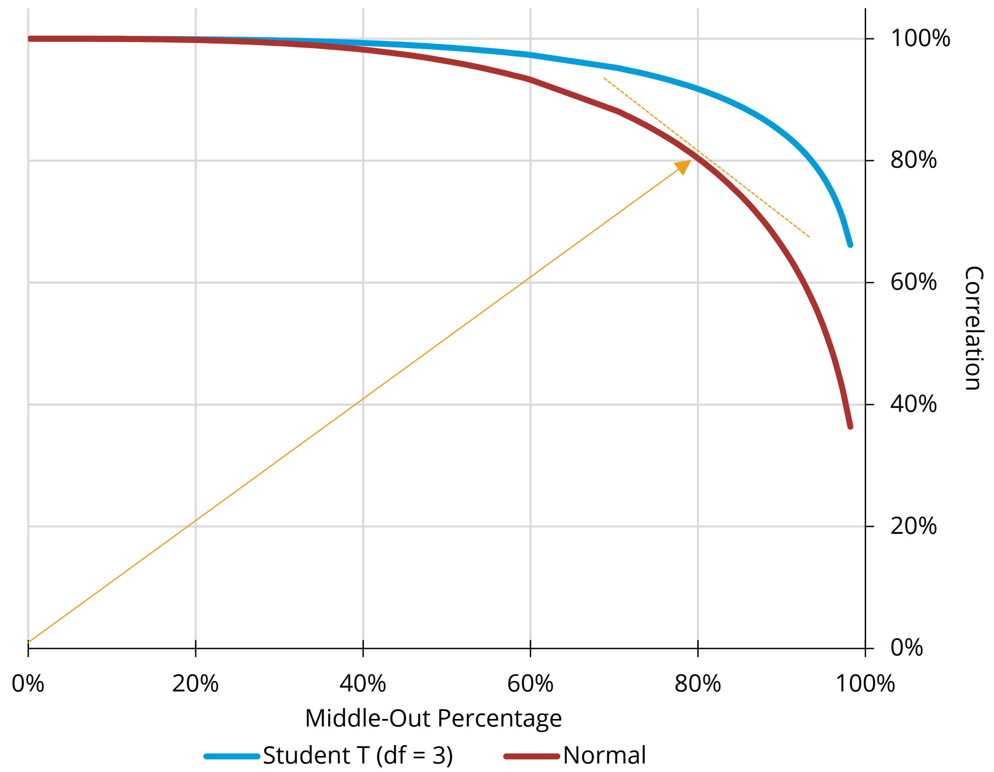

Source: Simulations.

We can derive a closed-form formula for the correlation ρ between a standard normal and a middled-out version that replaces the middle m percent of the distribution with zeros:

$$\rho = \sigma = \sqrt{1 - m + \frac{2e^{-\mathrm{erfc}^{-1}(1-m)^2}\,\mathrm{erfc}^{-1}(1-m)}{\sqrt{\pi}}}$$

where erfc^{-1} is the inverse complementary error function. Note that the standard deviation σ of the middled-out distribution on a standard normal will also equal the correlation ρ. Hence all conclusions about the correlation in this case also apply to the middled-out volatility.

The correlation between the original distribution and the "middled-in" distribution, where only values in the middle are kept while those in the extreme deciles are replaced with zeroes, can be computed similarly. The middled-in correlation begins to exceed the middled-out correlation when ρ = 0.50, which occurs when m = 0.88. In other words, for a standard normal distribution, the middled-out and middled-in correlations are the same when the top and bottom 6% of the distribution are considered, thus providing an upper bound of m = 0.88 for the middle-out percentage. However, the (rounded) value of m = 0.80 displays two striking theoretical facts, graphically shown with the orange lines in Graph 3:

(1) It is the only value for which $\rho(m) = m$. (Exact solution: m = 80.3%.)

(2) It is also the only value for which $\rho'(m) = -1$. (Exact solution: m = 79.7%.)

The above discussion shows that the top and bottom deciles are the primary drivers of risk in terms of volatility and correlation with the original distribution. In other words, if the choice is between the traditional approach of ignoring the extremes or the suggested approach here of ignoring the middles instead, the middle-out approach is superior, as the tails contain the most important risk information. However, this begs the question: if all the returns are available, why not use all of the returns? Granted that ignoring the tails is far worse than ignoring the middles, but why ignore anything at all?

One way to address this question theoretically is to consider compound returns. If we simulate a year's worth of daily returns on a normal distribution with illustrative mean μ = 5% / 252 and standard deviation σ = 15% / √252, we can split the year into the top decile, bottom decile and middle, and compute the compound return of each group. If we do 100,000 simulations, we can report both the average and the standard deviation of the compound returns.

The first column of Table 1 reports the results. First, it is important to note that the compound return of the middle of the distribution, as well as the total, falls short of the expected average annualised return of 5%. This drag is due to the variance of the returns: a risk-free sequence would have had a compound return of exactly 5%. Second, the compounding effect appears to be stronger for the upside than for the downside, but of course the cumulative contributions of both tails combined must be near zero because the distribution had no skew: (1 – 34%) * (1 + 51%) – 1 ≈ 0%. Third, the total compound return exceeds the middle-only compound return, even for the skew-less normal distribution, because the tails incorporate some of the expected returns.

The second column of Table 1 reports similar results for a skew-normal distribution whose expectation and variance match the normal distribution above but with a modest +0.75 positive skewness. Now, the middle of the distribution has a very negative compound return, because it is by construction missing out on the positive skewness. By construction, the average total compound return and the standard deviation of the total compound return are the same across the normal and skew-normal distributions, but the locations of the volatility differ: the downside volatility is much higher absent skew. Incorporating a skew substantially heightens the importance of the tails. Now, if one were able to ignore the middle, the return wouldn't be break-even around zero, but substantially positive: (1 – 28%) * (1 + 61%) – 1 ≈ 16%. In addition, rather than missing out on the 4.4% positive compound of the normal middle, now we are avoiding a 9.3% loss in the skew-normal middle.

Compound returns of 252 normal returns and skew-normal — Table 1

Simulated 100,000 times

	Normal		Skew-Normal	
	Average	Standard deviation	Average	Standard deviation
Left (bottom decile)	–34%	2%	–28%	1%
Middle	4%	13%	–9%	11%
Right (top decile)	51%	4%	61%	6%
Total	5%	16%	5%	16%

Source: Simulations.

The real world does not represent repeated draws from a single static distribution, of course. The next section explores empirical results, but we can first explore what happens if the distribution changes over time. We simulate 100 years of returns where each year we choose an annualised mean uniformly randomly between –30% and +30% and a volatility uniformly randomly between 5% and 45% to draw 252 daily returns from a normal distribution with that year's mean and volatility. We then compute the correlation between the resulting 25,200 returns and the 80% middled-out returns. Graph 4 plots the histogram of results across 10,000 simulations. The minimum correlation was about 85% and the maximum 90%. In other words, the middle-out perspective continues to hold even for dynamic distributions that vary wildly over time, even for thin-tailed distributions such as the normal. The results would be even more extreme for fat-tailed distributions with varying parameters.

Histogram of correlations — Graph 4

Between middled-out and dynamic simulated random returns

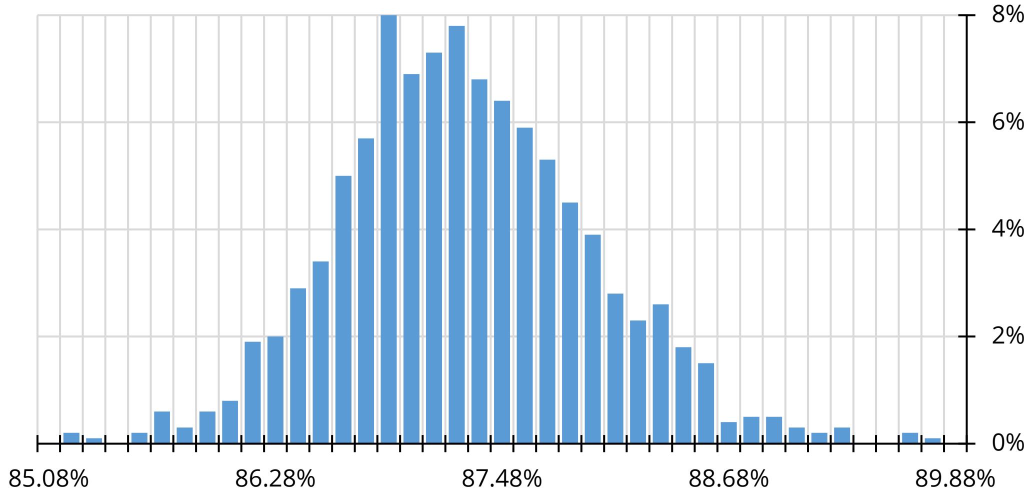

Source: Simulations.

3. Empirical results

To evaluate the empirical results of middle-out, we focus on 20 years of history of global equities and global bonds. For global equities, we use the MSCI ACWI Net Total Return Local Index (Bloomberg ticker: NDLEACWF Index). For global bonds, we use the Bloomberg Global Aggregate Index, described as a flagship measure of global investment grade debt from 24 local currency markets (Bloomberg ticker: LEGATRUH Index). The daily returns were calculated from Bloomberg from 26 February 1999 through 14 April 2023. Similar results were obtained for equities and debt comprising only US assets, using S&P 500 total returns and the returns of the Bloomberg US Government/Credit Bond Index (LUGCTRUU Index).

Graph 5 shows the correlation curve as a function of the middle percentage for both of these global indexes. The correlation between 80% middled-out equities and the original equities is 91%. The correlation between 80% middled-out bonds and the original bonds is 87%. Not shown but virtually identical are the ratio of volatilities between the middled-out and the original distributions: these numbers were 91% and 87%, respectively, as well.

Correlation between middled-out and historical index returns — Graph 5

Feb 1999–Apr 2023 for NDLEACWF and LEGATRUH indexes

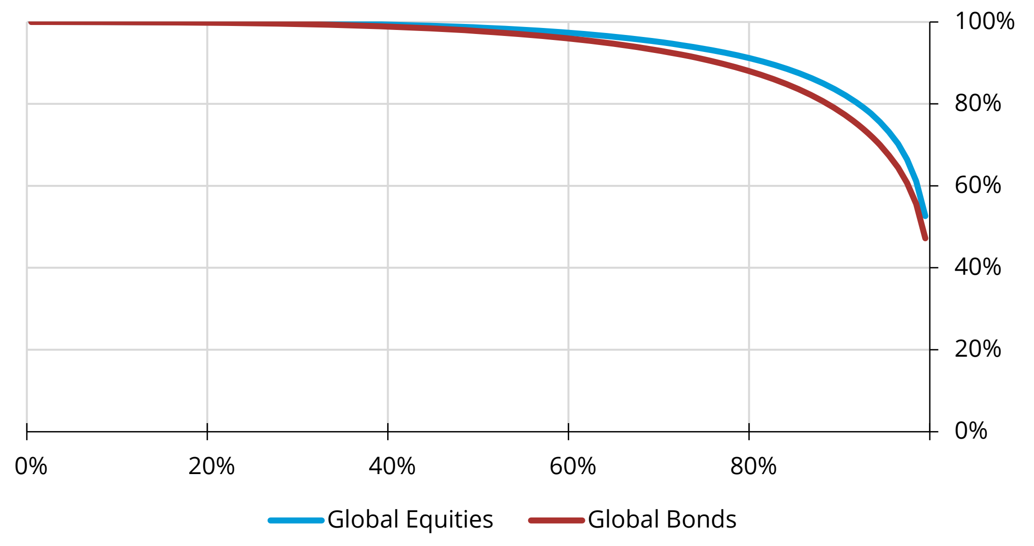

Source: Simulations.

One potential concern is that the middling-out thresholds are chosen using the entire history, which could potentially hide a form of lookahead bias. To address that concern, we run a rolling middle-out where the thresholds are chosen only from the returns of the past 60 trading days. Graph 6 shows those results, which are broadly the same. The correlation between 80% rolling-middled-out equities and the original equities is 85%. The correlation between 80% middled-out bonds and the original bonds is 83%. Again, the ratios of volatilities were also almost identical at all points.

Correlation between rolling middled-out and historical index returns | Graph 6

60-day rolling, Feb 1999–Apr 2023 for NDLEACWF and LEGATRUH indexes

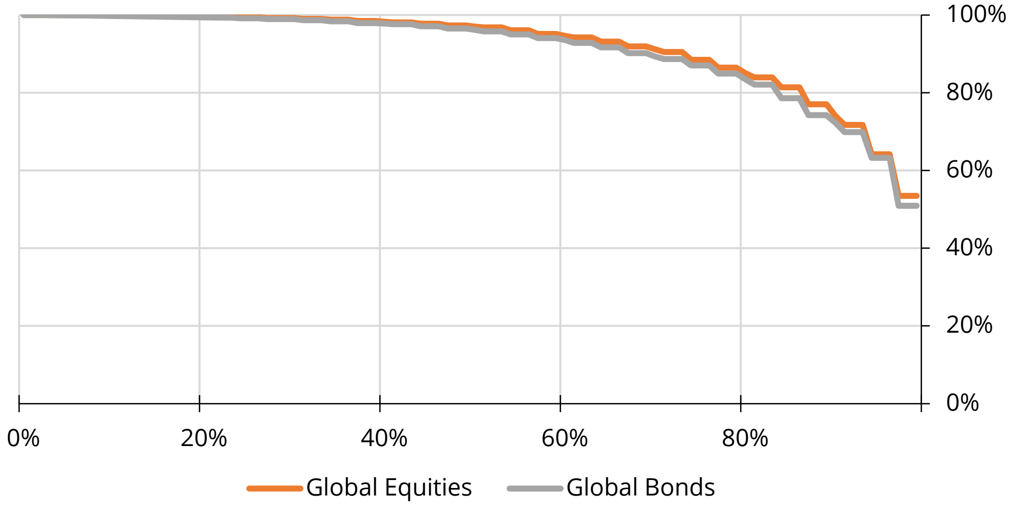

Source: Simulations.

As with the theoretical results, one may note that the above has convinced us, first, that the tails are more important than the middles, and, second, that the optimal threshold for middling out appears to be around 80%. In other words, rather than ignoring the extremes, it is much better to ignore the middles instead. Yet we are faced with the same question as before: why ignore anything at all? The earlier simulation results showed that, for skewed distributions, ignoring the middle can lead to substantially better performance. We now explore the same question in the context of empirical results.

Growth of $1 in excess of cash invested in US large-cap equities | Graph 7

Hypothetical, Ibboton 1857–1925 and S&P 500 1926–2016

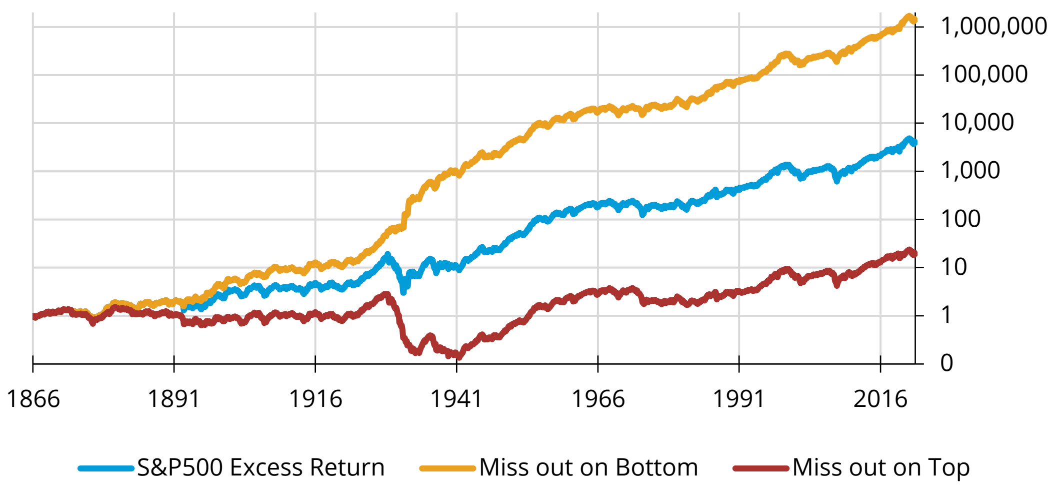

Source: Ibbotson and Bloomberg.

Using Ibbotson, for the period 1/1/1857–12/31/1925, individual security returns were gathered from US financial periodicals on a monthly basis, beginning with the official list of the New York Stock Exchange during that time period; from the period 1/1/1926–12/31/2022 returns were represented by the S&P 500 Index.

Graph 7 plots the hypothetical growth of a single $1 investment in US large-cap equities at the beginning of the period for three scenarios: always invested, missing the extreme tail losses, and missing the extreme tail gains. Here, an extreme tail gain or loss is any monthly period whose performance is two standard deviations above or below the average monthly return for the entire period.

Another way of exploring the importance of tails might be to calculate the marginal improvement due to a 50% improvement in either the left tail, the right tail, or the middle. To do that, we use gross daily returns of the S&P 500 since January 1928 through April 2023 from Bloomberg (SPX Index).

Table 2 shows the results.

The left and right tails are almost 40x larger in magnitude than the middle, about 200 bp each versus about 5.5 bp for the middle. The overall compound annualised return is 9.5%. That increases to 15.3% if the middles improve by 50%, ie if we multiply the middle return by 1.5. But it increases to more than 40% if we improve either the left or the right tails by 50% (reduce the downside by 50% or increase the upside by 50%), and to nearly 85% if both tails are improved and the middle remains untouched. It is hard to imagine a more dramatic difference.

Compound annualised returns of the S&P 500
1928–2022

Table 2

	Average
Average of bottom 10% of daily returns	–209 bp
Average of middle 80% of daily returns	5.5 bp
Average of top 10% of daily returns	209 bp

	Compound Annualised Returns
Overall	9.5%
Multiply middle returns by 1.5	15.3%
Multiply downside tail returns by 0.5	43.2%
Multiply upside tail returns by 1.5	41.0%
Multiple upside by 1.5 and downside by 0.5	84.5%

Source: Bloomberg.

4. Discussion and extensions

The results shown here are generalisable to other distributions and other measures of co-movement. For example, the fatter-tailed a distribution is, the higher the correlation would be: a Cauchy distribution is nearly 100% correlated with its middled-out version for essentially any middle-out percentage. On the flip side, using rank correlation generates distribution-free results that are similar to the Normal case.

In addition to these correlation and volatility results, tail returns should also have an even greater importance for investment managers and asset allocators because of their outsize effects on compound returns. The compound return of an investment, or its geometric average, underperforms the simple average return, and this amount of underperformance increases with larger volatility. And as we have seen, virtually all of the volatility of a financial asset comes from its tails.

On a practical level, when markets are in the middle region of a distribution, hedging and reacting to events can appear deceptively smooth and easy. When tail conditions do occur, hedging or reacting can become inordinately difficult, and may come as a surprise to those accustomed to middling returns.

On a philosophical level, tail returns are likely the ones that reveal actual asset information, while the middle returns are merely middling noise. The classic temptation is to treat tail returns as outliers, noise, or acts of God and effectively ignore them in standard portfolio construction methodologies such as mean-variance optimisation and standard risk management metrics such as value-at-risk, but this is exactly backwards. Earnings announcements, news and one-off events drive tail events and change asset prices substantially; without these drivers, prices essentially randomly fluctuate without any signal. Investment professionals should instead apply middle-out thinking and focus their risk management and portfolio decision policies on tail events.

Finally, if we take the view that the returns on a portfolio have a systematic component and a residual component, with tail events, the correlations among assets increase, thus reducing the diversification effect of individual securities. The volatility of the portfolio thus increases as a result. Even more likely, the total volatility is also related to tail events.

Conclusion

Broadly speaking, the two extreme deciles of financial returns represent about 90% of the volatility. Yet typically, extreme events are the ones that are ignored by conventional risk models that focus on estimating risk primarily from the middle of the distribution; tails are incorrectly viewed as "noise."

This conventional approach is exactly backwards: indeed, it is the extremes where the information is, and the middle of the distribution is the noise. Investors, asset allocators, risk managers and other finance professionals should stop focusing on the middling contribution of the middle and focus instead on the tales of the tails.

References

American Finance Association (2022): "Top 50 cited articles of all time", 21 April, https://afajof.org/most-cited-articles/.

Ang, A, R Hodrick, Y Xing and Z Zhang (2006): "The cross-section of volatility and expected returns", *Journal of Finance*, vol 61, no 1.

Baker, M and J Wurgler (2006): "Investor sentiment and the cross-section of stock returns", *Journal of Finance*, vol 61, no 4.

Bansal, R and A Yaron (2004): "Risks for the long run: a potential resolution of asset pricing puzzles", *Journal of Finance*, vol 59, no 4.

Basel Committee on Banking Supervision (2017): *Basel III: Finalising post-crisis reforms*, 21 April.

Blake, C and M Morey (2009): "Morningstar ratings and mutual fund performance", *Journal of Financial and Quantitative Analysis*, vol 35, no 3, pp 451–83.

Carhart, M (1997): "On persistence in mutual fund performance", *Journal of Finance*, vol 52, no 1.

Daniel, K, D Hirshleifer and A Subrahmanyam (1998): "Investor psychology and security market under- and overreactions", *Journal of Finance*, vol 53, no 6.

De Bondt, W and R Thaler (1985): "Does the stock market overreact?", *Journal of Finance*, vol 40, no 3.

Fama, E and K French (1992): "The cross-section of expected returns", *Journal of Finance*, vol 47, no 2.

——— (1996): "Multifactor explanations of asset pricing anomalies", *Journal of Finance*, vol 51, no 1.

Forbes, K and R Rigobon (2002): "No contagion, only interdependence: measuring stock market comovements", *Journal of Finance*, vol 57, no 5.

Glosten, L, R Jagannathan and D Runkle (1993): "On the relation between the expected value and the volatility of the nominal excess return on stocks", *Journal of Finance*, vol 48, no 5.

Hong, H and J Stein (1999): "A unified theory of underreaction, momentum trading, and overreaction in asset markets", *Journal of Finance*, vol 54, no 6.

Jegadeesh, N and S Titman (1993): "Returns to buying winners and selling losers: implications for stock market efficiency", *Journal of Finance*, vol 48, no 1.

Markowitz, H (1952): "Portfolio selection", *Journal of Finance*, vol 7, no 1.

Sharpe, W (1964): "Capital asset prices: a theory of market equilibrium under conditions of risk", *Journal of Finance*, vol 19, no 3.