INSIGHT.DATA.CLARITY.

August 2017

# THE STATISTICS OF RANDOM WALKS

## How low can you go and when should you panic?


### Executive summary

**Investors inevitably focus on the short term performance of the components of their portfolio and often attempt to make sense of recent bouts of consecutive negative returns or drawdowns in a given asset or strategy. These return streams can be modelled using random walks, a numerical technique that provides insight into the probability of observing a recent performance pattern given the assumptions about the quality of the investment. We demonstrate, given the modest Sharpe ratios that one should expect from an individual investment, that the depth and length of a cumulative period of negative performance can be surprisingly large. This observation leads us to the conjecture that investors systematically underestimate the length and depth of 'normal' drawdowns.**



**Contact details**

Call us +33 1 49 49 59 49

Email us cfm@cfm.fr

# Introduction

Everybody remembers their best trade! That one time when everything came together with returns that were so good that investing didn't seem to represent a challenge anymore. Markets, however, have a habit of humbling investors and, over time, those best trades blur into a mix of good and bad. Still, the bad trades are underweighted relative to the winners, leading to a false sense of one's ability to distinguish good investments from bad and in tough times it is very tempting to cut the losers and focus on the winners. Unfortunately, it is rare to find a genuine 'free lunch' and expectations of Sharpe ratios should be tapered to match the real world and account for the many biases that exist in investment. Having done so, deciding upon which strategy to invest in objectively becomes a real challenge.

The obvious question to ask in the presence of consecutive negative returns is whether the risk of the investment has been correctly estimated. This is often a question that our clients instinctively and immediately require an answer to when a given strategy heads south, wanting proof that risk is within the expected envelope. Of course, risk modelling is something that one should be concerned by, but generally, accumulated negative returns, or 'drawdowns', are more often simply the result of a random walk with a modest Sharpe ratio going about its business.

In fact, studying random walks provides the insight, demonstrated below, that the principal driver of drawdowns is the Sharpe ratio itself. The length of drawdowns scales inversely as the square of the Sharpe. For a Sharpe ratio of 1, one should expect typical drawdowns to last 1 year. Investing in a 0.3 Sharpe ratio strategy, a level one should expect from long term returns of equity indices for example, leads to typical drawdowns lasting $1/0.3^2$ or 10 years! Correspondingly, the depths of drawdowns are related to the inverse of the Sharpe ratio, meaning if one is invested in a strategy with a Sharpe ratio of 1 that produces a drawdown of $X\sigma$ (X representing a multiple of the volatility, $\sigma$, of the process) with a given probability then a strategy with a Sharpe ratio of 0.5 implies a drawdown of 1/0.5 or of $2X\sigma$ with the same probability. These relationships are investigated further below using numerical techniques.

One might be tempted to ask why focus so much energy on drawdowns rather than draw-ups? There is an obvious symmetry between the two but it remains the case that drawdowns are very well defined in being the difference in performance between the highest, most recent peak and lowest, most recent trough. A draw-up is not as easily defined with the price potentially never returning to its original point. In any case, investors are not overly concerned by draw-ups in studying the performance of their strategies. This short note is written as follows[1]: beginning with an introduction to numerical simulations we briefly give some intuition behind how the depth and length of drawdowns change with the Sharpe ratio of the random walk. We next look at the less interesting subject of the distribution of all drawdowns before dealing with the more instructive discussion of the worst drawdown for a given Sharpe ratio. We then also study the lengths of the longest drawdowns before using these numerical techniques to demonstrate the relationship between drawdown depths, lengths and the Sharpe ratio of the strategy. We conclude with some discussion regarding potential future work on using different models for processes that may better describe investment strategies.

# Numerical simulations – a pragmatic approach

The following introduction to numerical techniques may be considered technical by some readers and may be safely skipped in order to get to the key results. We first begin by introducing the basic tool of these simulations - the random walk. One can model the returns of an investment by constructing the price $p$ of the strategy as:

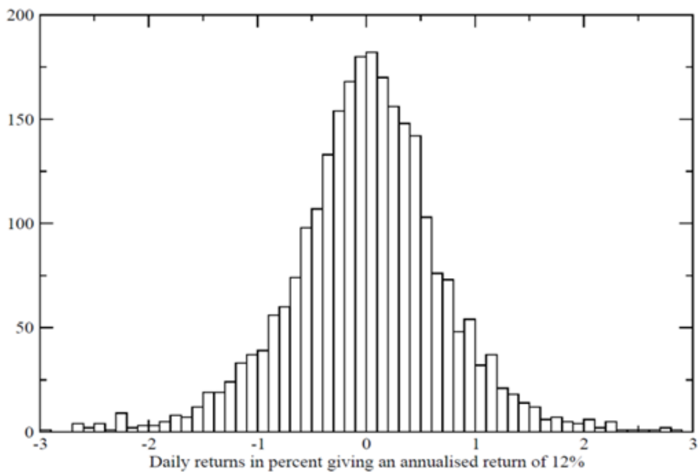


Figure 1: A histogram illustrating the bell shaped distribution of the random numbers used in the random walks. The random numbers are centered on zero and have tails that fit financial time-series well.

[1] This work has been conducted in parallel with our academic paper *You are in drawdown. When should you start worrying?* A Rej et al, CFM arXiv:1707.01457v2

$$p = \sum_{n=0}^{N} (\mu + \eta_n)$$

where n is the counter, say the days for a daily return and N is the total number of days in the time-series of returns[2]. The η term is simply a zero mean noise term or random number generator with a bell shaped distribution that best models the returns $\delta p = \mu + \eta$, of the investment strategy. A histogram of these random numbers can be seen in Figure 1 showing a distribution centered on zero with tails representative of financial returns[3]. The μ term is a constant added to the unpredictable 'noise' $\eta_n$ at every time step to generate a random walk with a 'drift' or a positive return. Figure 2 shows the results of generating random walks with Sharpe ratios of 0, 0.5 and 1 by varying the drift term to achieve the Sharpe ratio we require. Obviously, a Sharpe ratio of zero is generated by applying no drift term at all i.e. setting $\mu$ to zero and allowing the zero mean of the $\eta_n$ random numbers to generate a flat (on average) zero Sharpe random walk.

We now have a framework within which to simulate many random walks with any particular Sharpe ratio we desire, each realisation being different due to the existence of the $\eta_n$ term. The time-series in Figure 2 show how these random walks resemble the prices of for example investment indices or individual fund returns.

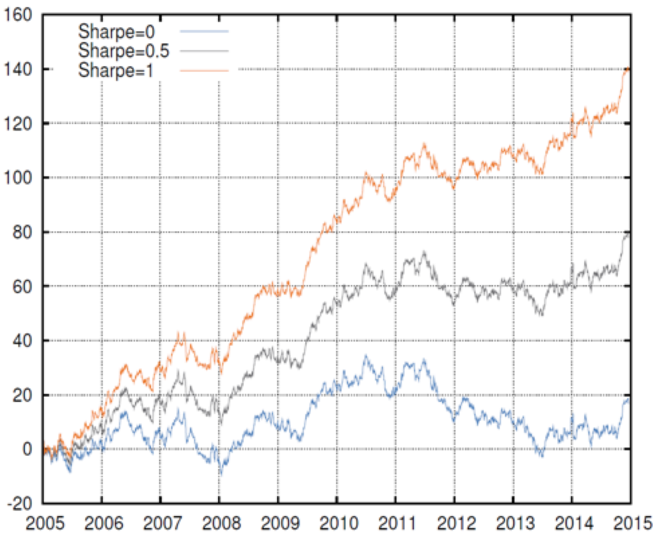


**Figure 2:** Random walks generated with 3 Sharpe ratios illustrating how varying the μ parameter allows us to easily change the drift and hence the Sharpe ratio.

# How long should a drawdown be? A rough guide

Given this 'random walk' model for a strategy it is instructive to consider the interplay between the two terms, the first driving the returns or the drift, $\mu$, and the second driving the volatility, $\eta$. A feature of the random component is that the envelope of expected outcomes is described by a square root function. We can write the price, P, at time, T, as the following:

$$P(T) = \mu T + \sigma\sqrt{T}$$

This can be seen visually in figure 3. We have generated many different realisations of zero Sharpe random walks and find that they are enveloped by a square root function of time. Adding in the drift term, $\mu T$, as in the above equation, we see that on short timescales the random fluctuations of the noise dominate and the drift is not strong enough to bring the random walk outside of the square root envelope. As time increases, however, we apply a larger and larger positive support to the random walk such that, at some point, the drift dominates and the positive Sharpe ratio expresses itself. This is intuitive as a zoom into the returns of any strategy, even those with high Sharpe ratios, reveals something that resembles white noise. A zoom out of any strategy, contrarily and even for low Sharpe ratios, looks misleadingly good, despite the low Sharpe ratio. A cursory look at, for example, two centuries of trend following[4], visually looks much better than one would expect from the measured Sharpe ratio of 0.8.

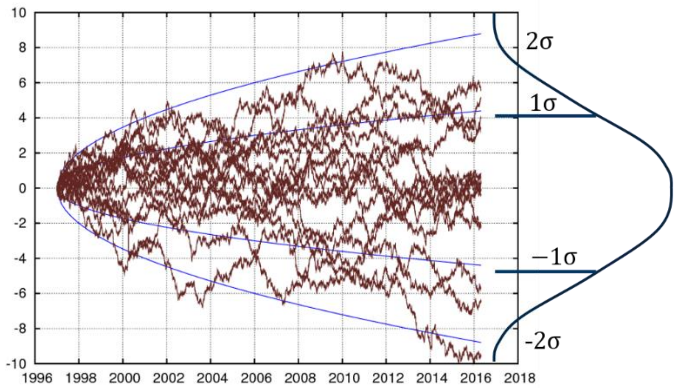


**Figure 3:** Zero Sharpe ratio random walks generated using the technique described in the text. Also shown are the 1 and 2 sigma

[2] One may enquire as to why we sum over returns rather than compounding them. Summed returns represent the performance of a constant risk investment, a more practical framework within which to simulate investments. However, all results can be reproduced using a random walk $p = \prod_{n=0}^{N}(\mu + \eta_n)$ with drawdowns defined as $(p_{peak} - p_{trough})/p_{peak}$

[3] The choice of the distribution of returns can change the results of the study. Here we use a Student's distribution with 4 degrees of freedom, a distribution which is naturally "fat tailed" and fits financial time-series well

[4] Two centuries of trend following, Lemperiere et al, CFM, published in the Journal of Investment Strategies

lines that envelope the evolving prices. In adding in a positive drift term, μ, the random walks slowly evolve away from zero, the drift term beginning to dominate at the time $T^*$ described in the text.

With this in mind we are now in a position to formulate this effect by setting the two terms of our random walk to be equal at a time $T^*$.

$$\mu T^* = \sigma\sqrt{T^*}$$

$T^*$ representing the point at which the random noise term finds it more and more difficult to stop the drift term from pulling the strategy upwards[5]. This can be rearranged to give the following rule of thumb

$$T^* = \frac{\sigma^2}{\mu^2} = 1/Sharpe^2$$

As the Sharpe ratio of the random walk changes, the timescale at which these two forces offset scales inversely to the square of the Sharpe ratio. One can now also take this characteristic timescale and plug it back into the random walk. At time $T^*$ therefore

$$-\sigma\sqrt{T^*} = -\sigma\sqrt{\frac{1}{Sharpe^2}} = -\sigma/Sharpe$$

One finds that the depth that the noise component gets us to in the random walk varies inversely to the Sharpe ratio. These two rules of thumb will be discussed further later on in the text.

## All drawdowns

Having now set up our framework to study drawdown probabilities numerically we can generate many realisations and average over them. The first obvious exercise is to look at the distribution of all drawdowns, described as the cumulative losses generated in periods of negative performance. Any one drawdown, $d$, as illustrated in figure 4, is defined as:

$$d = \sum_{last\ peak}^{now} \delta p_t$$

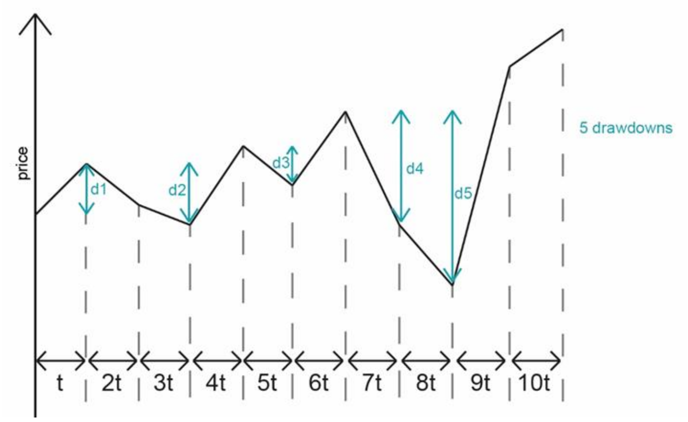


Figure 4: An illustration of how we calculate the drawdown for a random walk of 10 time periods. Each of the drawdowns is calculated from the highest recent peak to the trough. Each return that moves the price above the highest peak does not therefore count as a drawdown. Most of these moves down are small and correspond to the wiggles in the random walk.

A histogram of these drawdowns is shown in figure 5 where one observes a distribution which is heavily biased towards small drawdowns. It seems on the face of it that deep drawdowns are in fact highly improbable. This is however misleading in the sense that the majority of negative periods are very small with only very few that persist. Most drawdowns are in fact very short for a random walk and these negative 'wiggles' are mostly irrelevant. As we shall see, it is these deepest or longest drawdowns that are the ones we should be looking at. Nonetheless, one can use this distribution to now construct a probability density function that gives us the probability of seeing a given level of drawdown. It is easiest to construct this function in its cumulative form[6] which can be seen in figure 5. If we now experience a drawdown of say 2 standard deviations then we can immediately read off that 87% of all drawdowns were less than this level for a Sharpe of 0.5 or, stated differently, 1 drawdown in 1/(1-0.87)=8 is deeper than this level – a 2σ drawdown will therefore occur 1 drawdown out of 8.

[5] at least to 1 sigma ie $T^*$ represents the time at which 1/3 of realisations will be negative overall and 2/3 will be positive while 2T^* would have only 2.5% of negative realisations

[6] The trick to determining the cumulative distribution is to sort the observed data in order of their observed values, for example in increasing order. The cumulative probability of a given value of the data at the sorted value, k, is then found to be k/(N+1), where N corresponds to the total amount of points in the data set

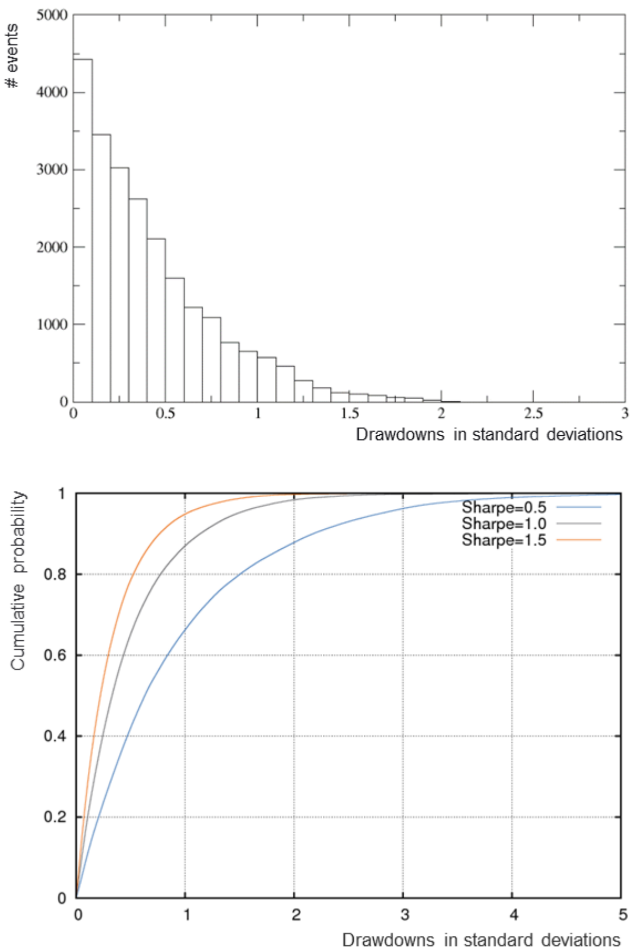


Figure 5: The top histogram shows the lengths of all drawdowns, in standard deviations σ, that have ever occurred in a time-series of fictitious price returns generated as described in the text. The bottom plot shows the cumulative probability density function for three values of Sharpe ratio – 0.5, 1 and 1.5. The y-axis represents the probability of seeing a drawdown of up to that depth, seen on the x-axis. For example, for the Sharpe=1.5 line with a probability of 0.5 one would see drawdowns up to ~0.3σ while with a probability of 0.95 one would see drawdowns up to ~1σ. 1 drawdown in 2 is therefore at least ~0.3σ deep while 1 drawdown in 20 is at least ~1σ deep for a Sharpe ratio of 1.5.

## Worst drawdowns

Unfortunately the distribution of all drawdowns above gives us a false sense of security and makes us think that deep drawdowns are unlikely. In reality, as we all know, drawdowns do exist and can be deep suggesting that they are not that improbable. So why does the above analysis suggest otherwise? The issue is that it is only the worst drawdowns in a given window of time, rather than all drawdowns, that we are concerned by. We can instead plot out the depth of the worst peak to valley drawdown seen in every calendar year, including the cases where that drawdown began in prior years. We do this again for very many years, with one worst drawdown per year, and see that this distribution is now very different. The distribution is shown in figure 6 where we see that the worst drawdowns are obviously rarely close to zero and can be significant. We once again convert this into a cumulative probability density function as shown in figure 6 where one observes, for example, for a Sharpe of 0.5 the most probable worst drawdown 1 year in 2 (the median) is at least 1.4σ while 5% of the worst drawdowns are at least 3.75σ. We can put this into context by setting $\sigma = 10\%$ (close to the volatility of the stock market in normal, calm periods) to see that a 'normal' year for a Sharpe=0.5 process would see a loss of at least 14% while 1 year in 20 the process loses at least 37.5% of its assets! These numbers may sound alarmingly high but are purely the result of the statistics of a random walk. Considering that many strategies exhibit a Sharpe of 0.5, it is safe to say that a firm existing based on only one such strategy has a high probability of going out of business based on a quite probable fluctuation!

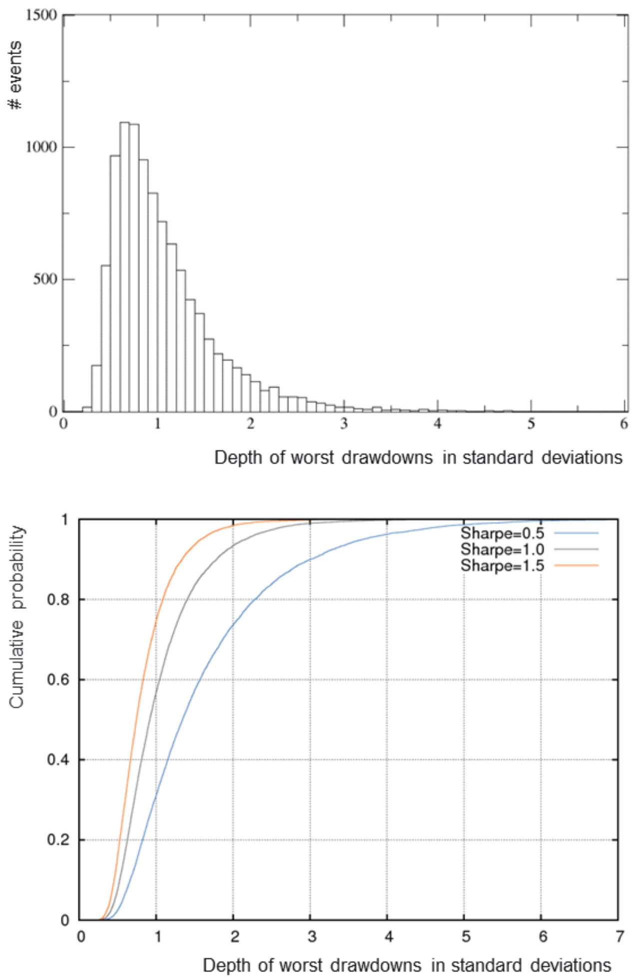


Figure 6: The top histogram shows the lengths of the deepest drawdowns per year, in standard deviations σ, that have ever occurred in a time-series of fictitious price returns generated as described in the text. Each data point therefore corresponds to

the worst drawdown seen in that year, even if that drawdown began in a prior year. The bottom chart shows the cumulative probability density function of these worst drawdowns for three values of Sharpe ratio – 0.5, 1 and 1.5. Taking for example the Sharpe=0.5 line one sees that 50% of worst drawdowns were at least 1.4 σ deep while 95% of the worst drawdowns were at least 3.75 σ or correspondingly 5% of years experienced a drawdown of 3.75 σ or more. Stated differently, 1 year in 2 will see a drawdown of at least 1.4 σ while 1 year in 20 a drawdown of at least 3.75 σ should be observed.

### Longest drawdowns

The previous section described the deepest drawdowns in any year which is of obvious concern to any investor. Of a lesser concern but still relevant in selecting strategies and making investment decisions is to consider the length of a drawdown. A shallow drawdown that continues for too long may make an investor question his assumptions behind the Sharpe ratio of an investment decision whereas deep drawdowns should statistically not be quick. The deepest drawdown and the longest drawdown in any one given period of time do not necessarily correspond to the same event and we can instead focus on the distribution of the lengths of the longest drawdowns per year, again even if they had started prior to the beginning of the year, which as the Sharpe ratio decreases becomes ever more probable! In figure 7 we show a histogram of the lengths of the longest drawdowns per year which again is converted to a cumulative probability distribution function in figure 7. We see for example that for a Sharpe of 0.5 that a typical, median longest drawdown lasts for 350 days (approximately 17 months) while 1 year in 20 sees a longest drawdown lasting 2500 days (approximately 10 years!). If one is dealing with such low Sharpe ratio strategies it becomes clear that it is almost impossible to make an objective investment decision based on disappointing, flat performance over a 10 year period.

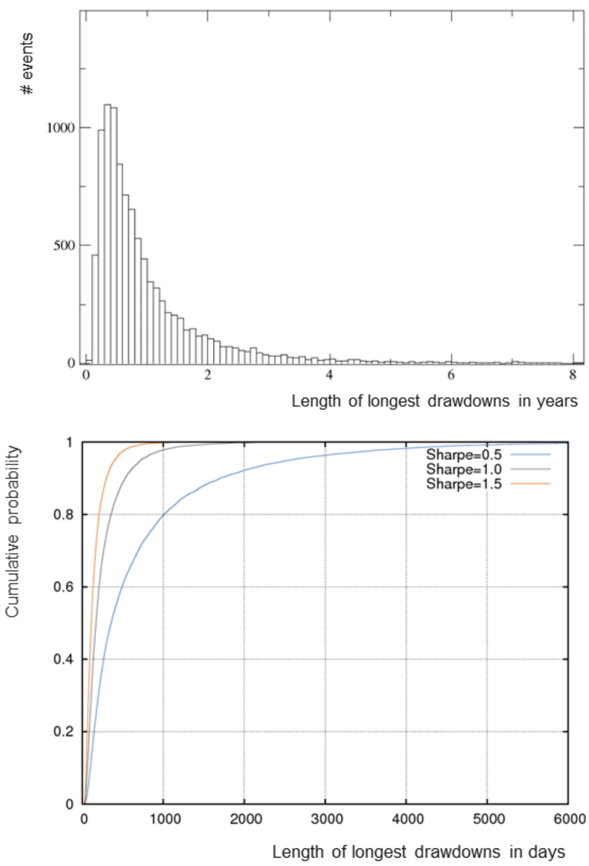


Figure 7: The top histogram shows the lengths of the longest drawdowns per year, in years, that have ever occurred in a time-series of fictitious price returns generated as described in the text. Each data point therefore corresponds to the longest drawdown seen in that year, even if that drawdown began in a prior year. The bottom graph shows the cumulative probability density function of these longest drawdowns (in days) for three values of Sharpe ratio – 0.5, 1 and 1.5. Taking for example the Sharpe=0.5 line one sees that 50% of the longest drawdowns were at least 350 days long while 95% of the longest drawdowns were at least 2500 days long, or correspondingly 5% of years experienced a drawdown that lasted for 2500 days or more. Stated differently, 1 year in 2 will see a drawdown that lasts at least 350 days while 1 year in 20 a drawdown lasting at least 2500 days should be observed.

## How do drawdown depths and lengths depend on the Sharpe ratio?

We are now in position to investigate the relationship between drawdown depths, lengths and the Sharpe ratio of the process. Figure 8 shows the 1 year in 20 drawdown depths for a range of Sharpe ratios with the function $1/Sharpe$ overlaid. Also shown is the 1 year in 20 drawdown lengths for a range of Sharpe ratios, this time with the function $1/Sharpe^2$ overlaid. These relationships describe well the way that expectations of drawdown depths and lengths should scale with Sharpe ratio and confirm the intuition derived above from the construction of the random walk.

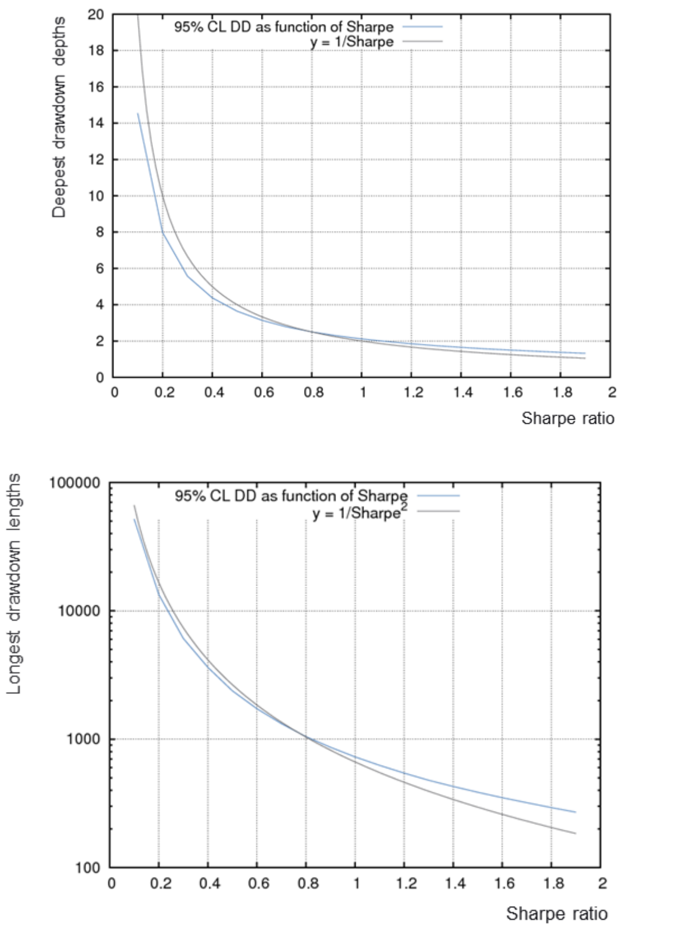


Figure 8: The depths and lengths of the 5% deepest and longest drawdowns experienced with a range of Sharpe ratios. Also shown are the relationships as described in the text that the worst drawdown depth is proportional to $1/Sharpe$ while the longest drawdown length is proportional to $1/Sharpe^2$

# Conclusions

In this short note we have studied a model of strategy returns that allows us to build up probability distributions of drawdowns for a given level of Sharpe ratio. This model was used to look at the most probable worst drawdowns and most probable longest drawdowns in any one calendar year even when that said drawdown began in a previous year. The results give surprisingly large numbers for what is considered typical levels of Sharpe ratios that investors would have in a portfolio, typically basing investment decisions on prior returns. For example, with a Sharpe of 0.5 one should expect:

- A deepest drawdown of at least 3.75σ 1 year out of 20, or a 37.5% loss for a 10% volatility
- A longest drawdown of at least 10 years, 1 year out of 20

The depth of the drawdown increases proportionally to the inverse of the Sharpe ratio while the length of the drawdown increases proportionally to the inverse of the square of the Sharpe ratio. This makes sense, considering a Sharpe ratio of zero implies strategy drawdowns that can in principle go on forever albeit with a low probability. These statistics are counterintuitively high and, along with our experience of interacting with allocators and advisors, lead us to conclude that investors tend to significantly underestimate the depth and length of statistically normal drawdowns.

In conclusion the only way to reduce the depth and length of drawdowns is to improve the Sharpe ratio of a portfolio and, generally speaking, the only way to improve Sharpe ratios is through diversifying among many well implemented strategies. The statistics presented in this paper are, of course, sensitive to the process used. It is the subject of further work in this field to study random walks with time fluctuating Sharpe ratios and/or autocorrelation in the returns and volatility. We will report on this further work in subsequent papers.

# Disclaimer

# CFM has pioneered and applied an academic and scientific approach to financial markets, creating award winning strategies and a market leading investment management firm.

## Contact us

**Capital Fund Management S.A.**
23, rue de l'Université, 75007
Paris, France
T +33 1 49 49 59 49
E cfm@cfm.fr

**CFM International Inc.**
The Chrysler Building, 405 Lexington Avenue - 55th Fl.,
New York, NY, 10174, U.S.A
T +1 646 957 8018
E cfm@cfm.fr

**CFM Asia KK**
9F Marunouchi Building, 2-4-1, Marunouchi, Chiyoda-Ku,
100-6309 Tokyo, Japan
T +81 3 5219 6180
E cfm@cfm.fr

**Capital Fund Management LLP**
64 St James's Street, London
SW1A 1NF, UK
T +44 207 659 9750
E cfm@cfm.fr