Maximum drawdown

The maximum loss from a market peak to a market nadir, commonly called the maximum drawdown (MDD), measures how sustained one's losses can be. Malik Magdon-Ismail and Amir Atiya present analytical results relating the MDD to the mean return and the Sharpe ratio. The MDD factors into many risk-adjusted measures of performance, such as the Calmar ratio. Magdon-Ismail and Atiya propose new scaling laws for such ratios, analogous to the 'square-root-T scaling law' for the Sharpe ratio, which facilitates the comparison of funds with track records of different length. They also discuss the portfolio implications of their results

The maximum cumulative loss from a market peak to the following trough, often called the maximum drawdown (MDD), is a measure of how sustained one's losses can be. Large drawdowns usually lead to fund redemptions, and so the MDD is the risk measure of choice for many money management professionals – a reasonably low MDD is critical to the success of any fund. The MDD is related to the Calmar ratio[1], a risk-adjusted measure of performance that is given by the formula:

$$\text{Calmar}(T) = \frac{\text{Return over}[0,T]}{\text{MDD over}[0,T]}$$

The Sharpe ratio is similar in that it is also a risk-adjusted measure of performance. However, the MDD risk measure is replaced by the standard deviation of the returns over intervals of size T.

The 'square-root-T-law' is a well-known law describing how the unnormalised Sharpe ratio scales with time. This law allows one to scale the Sharpe ratio so that comparing different systems is possible even when their Sharpe ratios are calculated using different values of T. On the other hand, similar scaling laws for the Calmar ratio are not known. As a result, the common practice is to compare Calmar ratios for portfolios over equal length time intervals (the typical choice is three years). Such a constraint on the use of the Calmar ratio is artificial and, based upon the results that we will present, unnecessary.

Another important task for fund managers is the ability to construct portfolios that are optimal with respect to the risk-adjusted performance. When the performance measure used is the Sharpe ratio, this leads to mean-variance portfolio analysis. A similar approach to portfolio optimisation using the Calmar ratio as a criterion is not prevalent mainly because of a lack of an analytical understanding regarding how the MDD of a portfolio is related to performance characteristics of the individual instruments.

In this article, we present analytical results relating the expected MDD to the mean return and the standard deviation of the returns. The detailed mathematical derivations are given in Magdon-Ismail *et al* (2004). We also present formulas that relate the Calmar ratio to the Sharpe ratio. We introduce the normalised Calmar ratio, which can be immediately compared for two portfolios. We also present some plots illustrating some of the portfolio aspects of the MDD – in particular, how the correlation factors in. Among our findings is that an instrument with a negative return can be beneficial from the Calmar ratio point of view, if it is sufficiently uncorrelated.

The drawdown at time t (a related but analytically simpler measure than MDD) has been studied, and its distribution can be obtained analytically from the joint density of the maximum and the close of a Brownian motion (see, for example, Karatzas & Shreve, 1997). Most work on the maximum drawdown is empirical in nature (for example, Acar & James, 1997, Burghardt, Duncan & Liu, 2003, Harding, Nakou & Nejjar, 2003, and Sornette, 2002). The most relevant theoretical result is for the case of a Brownian motion with zero drift, in which case, the full distribution of the maximum drawdown is given in Douady, Shiryaev & Yor (2000). Since we wish to relate the MDD to the drift, we cannot assume that the drift is zero. Portfolio optimisation using the drawdown has also been considered in Chekhlov, Uryasev & Zabarankin (2003).

The expected maximum drawdown

Assume that the value of a portfolio follows a Brownian motion:

$$dx = \mu dt + \sigma dW \qquad 0 \le t \le T$$

where time is measured in years, and μ is the average return per unit time, σ is the standard deviation of the returns per unit time and dW is the usual Wiener increment. This model assumes that profits are not reinvested. If profits are reinvested, then a geometric Brownian motion is the appropriate model:

$$ds = \hat{\mu} s dt + \hat{\sigma} s dW \qquad 0 \le t \le T$$

For such a case, equivalent formulas can be obtained by taking a log transformation: if $x = \log s$, then x follows a Brownian motion with $\mu = \hat{\mu} - 1/2\,\hat{\sigma}^2$ and $\sigma = \hat{\sigma}$. (The MDD in this case is defined with respect to the percentage drawdown rather than the absolute drawdown.) If the portfolio value follows a more complicated process, then the results for the Brownian motion can be used as a benchmark.

Using results on the first passage time of a reflected Brownian motion, we find that the expected MDD has drastically different behaviour according to whether the portfolio is profitable, breaking even or losing money. This 'phase shift' in the behaviour is highlighted by the asymptotic ($T \to \infty$) behaviour in the formulas below. The asymptotic behaviour is important because most trading desks are interested in long-term performance, that is, systems that can survive over the long run, with superior return and small drawdowns. The expression for the expected MDD is:

$$E(MDD) = \begin{cases} \frac{2\sigma^2}{\mu} Q_p\left(\frac{\mu^2 T}{2\sigma^2}\right) \xrightarrow{T\to\infty} \frac{\sigma^2}{\mu}\left(0.63519 + 0.5\log T + \log\frac{\mu}{\sigma}\right) & \text{if } \mu > 0 \\ 1.2533\sigma\sqrt{T} & \text{if } \mu = 0 \\ \frac{-2\sigma^2}{\mu} Q_n\left(\frac{\mu^2 T}{2\sigma^2}\right) \xrightarrow{T\to\infty} -\mu T - \frac{\sigma^2}{\mu} & \text{if } \mu < 0 \end{cases}$$

As can be noted, the scaling of the expected MDD with T undergoes a phase transition from T to $\sqrt{T}$ to $\log T$ as μ changes from negative to zero to positive. One immediate use of this behaviour is as a hypothesis test to determine if a portfolio is profitable, even or losing. The functions $Q_n(x)$ and $Q_p(x)$ are complicated integral expansions that do not have a convenient analytical form. They are independent of μ, σ and T, and so they are

[1] *Similar to the Calmar ratio is the Sterling ratio:*

$$\text{Sterling}(T) = \frac{\text{Return over}[0,T]}{\text{MDD over }[0,T] - 10\%}$$

and our discussion applies equally well to the Sterling ratio

1. *E*(*MDD*) regimes and its dependence on *Shrp*

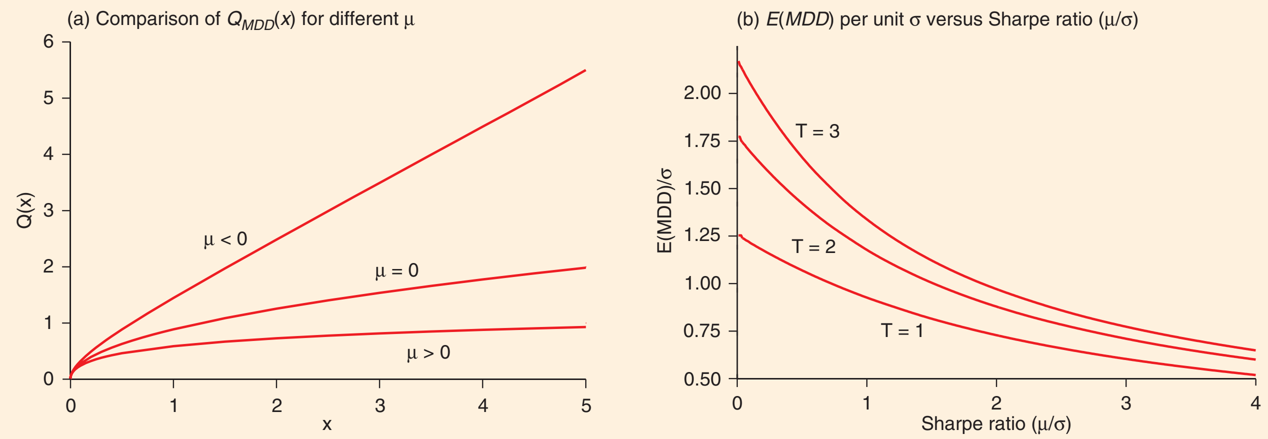

Note: (a) shows the behaviour of the $Q_p(x)$, $Q_n(x)$ and the equivalent function for μ = 0, illustrating the behaviour of these functions for different μ regimes; (b) shows how the expected *MDD* per unit variance depends on the Sharpe ratio for different values of *T*

2. How *Clmr* depends on *T* and *Shrp*

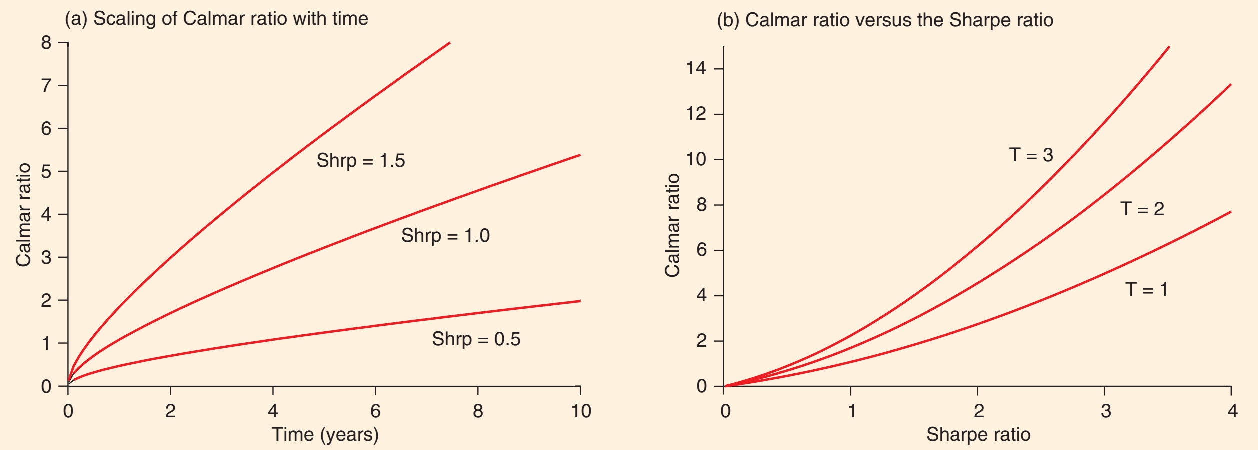

Note: (a) illustates how the *Clmr* scales with time for different portfolio characteristics; (b) shows how it scales with *Shrp* for different times

'universal functions' in the sense that they can be evaluated once and tabulated for future use. Such a table is given in Magdon-Ismail *et al* (2004) and can also be downloaded from Q-functions (see reference box). Figure 1(a) shows the functions $Q_p(x)$ and $Q_n(x)$. The exact functional form, including the distribution of the MDD, as well as a tabulation of values can be found in Magdon-Ismail *et al* (2003, 2004). From now on, we focus on the more interesting case of profitable (μ > 0) portfolios. The discussion can easily be extended to all three regimes of μ.

Define the $\sqrt{T}$-scaled Sharpe ratio of expected performance by $Shrp = \mu/\sigma$. The expected MDD normalised per unit of σ can be written entirely in terms of *Shrp*:

$$\frac{E(MDD)}{\sigma} = \frac{2Q_p\left(\frac{T}{2}Shrp^2\right)}{Shrp}$$

Figure 1(b) illustrates the dependence of *E*(*MDD*) normalised per unit of σ on the Sharpe ratio, *Shrp*.

The normalised Calmar ratio

First, we will deduce a relationship between the Sharpe ratio and the Calmar ratio. Consider the Calmar ratio of expected performance, *Clmr*, given by:

$$Clmr(T) = \frac{\mu T}{E(MDD)}$$

Substituting this definition into the expression for *E*(*MDD*), we obtain:

$$Clmr(T) = \frac{\frac{T}{2}Shrp^2}{Q_p\left(\frac{T}{2}Shrp^2\right)} \xrightarrow{T\to\infty} \frac{T\,Shrp^2}{0.63519 + 0.5\log T + \log Shrp} \tag{1}$$

Some interesting points to note are that the Calmar depends on μ and σ only through the scaled Sharpe ratio (the dependence of *Clmr* on *T* and on the normalised Sharpe ratio are illustrated in figure 2); for fixed μ, σ, *Clmr* increases with *T*. Thus, knowing the Calmar ratio of a portfolio without knowing T is useless. If fund X has a Calmar of five and fund Y has a Calmar of six, it is not clear which is a better fund. In fact it is possible that fund X is better! To make a better comparison, it is necessary to know the time intervals over which each Calmar ratio was calculated, and scale appropriately. However, perhaps we can remove this dependence on T by standardising the way the Calmar ratio is quoted. This can be accomplished by normalising the Calmar ratio. More specifically, whenever a Calmar ratio is quoted, one should automatically incorporate the appropriate scaling so that the comparison becomes seamless. Despite how prevalent the MDD is as a measure of risk, such a systematic approach is not usually used, because the appropriate scaling behaviour was not known. Our results provide exactly the necessary scaling behaviour.

Fix a reference time frame τ (for example, $\tau = 1$). If all Calmar ratios were quoted on this time frame, then comparing portfolios would be easy. For a given portfolio, suppose we have calculated *Shrp*. In this case, from (1), for the time interval τ, we know that *Clmr* is expected to be $Clmr(\tau) = (\tau/2)\, Shrp^2/Q_p(\tau/2\, Shrp^2)$. Similarly, at time T, we know that $Clmr(T) = (T/2)\, Shrp^2/Q_p(T/2\, Shrp^2)$, and so to get the τ-normalised Calmar ratio, we need to scale by a normalising factor:

$$\gamma_\tau(T, Shrp) = \frac{\frac{1}{T} Q_p\left(\frac{T}{2} Shrp^2\right)}{\frac{1}{\tau} Q_p\left(\frac{\tau}{2} Shrp^2\right)}$$

More specifically, if everyone agrees on the base time scale τ, then having calculated the Calmar ratio, and μ, σ for a portfolio over the interval $[0, T]$, the τ-normalised Calmar ratio $\overline{\text{Calmar}}(\tau)$ is given by:

$$\overline{\text{Calmar}}(\tau) = \gamma_\tau(T, Shrp) \times \text{Calmar ratio}$$

Following the convention applied to quoting the Sharpe ratio, we suggest fixing the base time scale τ to one year.

□ **Example.** The idea is best illustrated by an example. Suppose that three portfolios Π_1, Π_2, Π_3 have the profit and loss statistics over their respective time intervals as illustrated in table A. How do we compare these portfolios if our criterion is the Calmar ratio? First, let us illustrate some of the intuition. If we calculate *Clmr* for Π_1, we get roughly 3.8. Since its actual Calmar is higher, Π_1 seems to have negative autocorrelation for its returns, that is, it seems to be outperforming. Similarly, $Clmr(\Pi_2) = 6.76$ and $Clmr(\Pi_3) = 4.55$. It seems that Π_2 is under-performing and is the worst, but it is not clear how to compare Π_1 with Π_3 at this point. By calculating the normalised (to $\tau = 1$) Calmar ratios, we will be in a better position. Specifically, the Calmar ratio of Π_1 is already normalised, that is, $\overline{\text{Calmar}}_1 = 5$. If we calculate the normalising factors for portfolios Π_2 and Π_3, we get $\gamma(\Pi_2) = 0.74$ and $\gamma(\Pi_3) = 0.60$, from which we get the normalised Calmar ratios: $\overline{\text{Calmar}}_2 = 4.41$ and $\overline{\text{Calmar}}_3 = 3.62$. It is now clear that $\Pi_1 > \Pi_2 > \Pi_3$, if we normalise to $\tau = 1$.

The normalised Calmar ratio may depend on the choice of τ, the normalising time. We can remove the τ-dependence by defining the relative strength $\beta(\Pi_1|\Pi_2)$ of portfolio Π_1 with respect to some other benchmark portfolio, Π_2. Π_2 could be (for example) the S&P 500. For normalising time τ, define the τ-relative strength $\beta_\tau(\Pi_1|\Pi_2)$ of Π_1 with respect to Π_2:

$$\beta_\tau(\Pi_1|\Pi_2) = \frac{\overline{\text{Calmar}_1}(\tau)}{\overline{\text{Calmar}_2}(\tau)}$$

If $Shrp_1 \neq Shrp_2$, then the τ-relative strength depends on τ. The limiting (that is, $\tau \to \infty$) long-term behaviour of the relative strength is well defined, and so we define the relative strength $\beta(\Pi_1|\Pi_2) = \lim_{\tau\to\infty} \beta_\tau(\Pi_1|\Pi_2)$. One can show that:

$$relative\ strength = \beta(\Pi_1|\Pi_2) = \frac{\text{Calmar}_1}{\text{Calmar}_2} \times \frac{\frac{1}{T_1} Q_p\left(\frac{T_1}{2} Shrp_1^2\right)}{\frac{1}{T_2} Q_p\left(\frac{T_2}{2} Shrp_2^2\right)}$$

which is independent of τ. If the relative strength is greater than or equal to one, then Π_1 is 'better' than Π_2, written $\Pi_1 \succeq \Pi_2$. Since $\beta(\Pi_1|\Pi_3) = \beta(\Pi_1|\Pi_2)\beta(\Pi_2|\Pi_3)$, the relative strength index is transitive ($\Pi_1 \succeq \Pi_2$ and $\Pi_2 \succeq \Pi_3$ implies $\Pi_1 \succeq \Pi_3$), which is certainly a desirable consistency condition for any such strength index. It is complete and anti-symmetric, because $\beta(\Pi_1|\Pi_2) = 1/\beta(\Pi_2|\Pi_1)$ (so either $\Pi_1 \succeq \Pi_2$ or $\Pi_2 \succeq \Pi_1$ and $\Pi_1 \succeq \Pi_2 \Rightarrow \Pi_2 \preceq \Pi_1$). Thus $\succeq$ is a total order. Further, the choice of the reference instrument does not affect the total ordering, because $\beta(\Pi_1|\Pi_2) = \beta(\Pi_1|\Pi_3)/\beta(\Pi_2|\Pi_3)$ (so $\beta(\Pi_1|\Pi_3) \geq \beta(\Pi_2|\Pi_3) \Rightarrow \Pi_1 \succeq \Pi_2$). The relative strengths of the portfolios in the example, with Π_1 as benchmark, are given in table A.

□ **Real data.** In table B, we give the MDD-related statistics for some indexes and funds. The data (in non-bolded font) was obtained from the International Advisory Services Group (see reference box). Notice that the expected MDD is generally slightly lower than predicted. One reason for this is the discretisation bias (the data is built from monthly statistics, but the model is continuous). Notice that the time periods over which the funds are quoted are quite different, since the funds have been in existence for different periods of time. Some have not been around for three years, and some have been around significantly longer. Thus, it is not clear how to compare the funds using Calmar ratios for some standardised time period, three years being the norm in the industry. If a fund has been around less than three years, then it is not possible, and choosing (say) the most recent three-year period for a well-established fund ignores valuable data. However, the normalised Calmar ratios and the relative strengths facilitate seamless comparison among the funds using all the available data.

■ **Summary.** We now have a systematic way to quote Calmar ratios so that systems can be easily compared. Further, there is a direct (monotonic) relationship between the Calmar ratio and the Sharpe ratio. A deviation observed from historical data indicates a non-Brownian phenomenon at work, which could for example be due to the presence or absence of excessive correlation between successive loss periods, or the presence or absence of fat-tailed behaviour for the returns (note, however, that it has been empirically found that higher moments have a negligible impact on the Calmar ratio (Burghardt, Duncan & Liu, 2003)). Such features may depend on the nature of the trading system, the types of markets (for example, trending or mean-reverting) and the degree of diversification. For example, for a passive buy-and-hold strategy, if the Calmar ratio is lower than indicated by the theory, that could be due to positive autocorrelation for the returns, indicating the need for more risk control measures such as diversification or hedging. Alternatively, if a trend-following sys-

A. Some example portfolios

Portfolio	μ(%)	σ(%)	Calmar	Time interval (yrs)	Relative strength
Π_1	25	10	5	[0, 1]	1.00
Π_2	30	10	6	[0.5, 2]	0.97
Π_3	25	12.5	6	[0, 2]	0.64

B. MDD-related statistics of some indexes and funds available through the IASG

Fund	μ(%)	σ(%)	T(yrs)	MDD	Calmar	E[MDD]	$\overline{\text{Calmar}}$	β
S&P 500	10.04	15.48	24.25	46.28	5.261	**44.56**	**0.6104**	**1**
FTSE 100	7.01	16.66	19.83	48.52	2.865	**55.54**	**0.4395**	**0.5003**
Nasdaq	11.20	24.38	19.42	75.04	2.899	**77.87**	**0.4402**	**0.5407**
DCM	15.65	5.78	3.08	3.11	15.50	**4.770**	**6.541**	**27.76**
NLT	3.35	16.03	3.08	25.40	0.4062	**31.35**	**0.2202**	**0.1331**
OIC	17.19	4.52	1.16	0.42	47.48	**2.493**	**42.31**	**212.0**
TGF	8.48	9.83	4.58	8.11	4.789	**15.84**	**1.752**	**3.589**

Note: DCM = Diamond Capital Management; NLT = Non-Linear Technologies; OIC = Olsen Investment Corporation; TGF = Tradewinds Global Fund. The normalised Calmar ratio, $\overline{\text{Calmar}}$, is normalised to $\tau = 1$ yr. The relative strength index is calculated with respect to the S&P 500 as benchmark

3. Impact of correlation on Calmar ratio

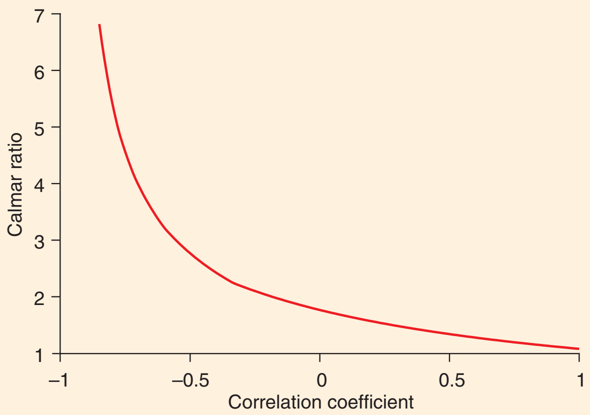

Note: the Calmar ratio for a portfolio of two equally weighted trading systems/markets (with $\mu_1 = \mu_2 = 0.2$ and $\sigma_1 = \sigma_2 = 0.2$) against the correlation coefficient of the two systems

tem were to pick the trends accurately, then it could significantly improve the Calmar ratio.

Portfolio aspects of MDD

Mean variance analysis exploits the correlation structure between assets to build a portfolio with good Sharpe ratio characteristics. This ability is facilitated by the fact that the variance and return of a portfolio can be calculated given these properties of the individual assets. As we have shown in the previous results, these parameters are also sufficient to obtain the $E(MDD)$ of the resulting portfolio, so we should be able to perform such a similar analysis to optimise the MDD. Further, since the Calmar ratio is monotonic in the Sharpe ratio, we can directly transfer portfolio optimisation methods for the Sharpe ratio over to the Calmar ratio. We briefly illustrate some of these issues here. Assume throughout that Calmar ratios are normalised to one year.

☐ **The impact of correlation.** Consider for simplicity a portfolio of two instruments. If the correlation of the returns of the two instruments is low, then we should be able to construct a portfolio better than either asset, from the risk-adjusted-return point of view. We want to quantify this effect using the previous analysis, and the Calmar ratio as a performance measure. For illustration, consider a portfolio in which the mean return of each instrument is 20%, and the standard deviation of the returns of each instrument is 20% (all annualised). Assume the portfolio is equal-weighted. Figure 3 shows the Calmar ratio as a function of the correlation coefficient of the returns of the two instruments. While the fact that the Calmar ratio decreases with increasing correlation is not surprising, the extent of the change is higher than expected. We should mention, however, that it is quite difficult to find trading systems/markets both with positive returns and highly negative correlation. Highly negative correlations are typically achieved by a long-type system versus a short- type system, in which case their mean returns would typically be of opposite signs. So the part of the curve deep into the negative correlation portion is probably difficult to attain.

☐ **Can a losing system be beneficial?** It should be possible to explore the negative correlation region by combining a losing system with a profitable system. To illustrate, let us perform the following curious experiment: consider two instruments with annualised returns $\mu_1 = \mu_2 = 20\%$ and standard deviations $\sigma_1 = \sigma_2 = 20\%$, with correlation coefficient $\rho_{12} = 0.8$ for the returns of the two instruments. Applying the formulas presented earlier, we find that the best Calmar ratio that can be achieved is 1.154, using a portfolio that weights each instrument equally. Assume now that we add to the portfolio a losing instrument with $\mu_3 = –10\%$, $\sigma_3 = 30\%$, and with negative correlations $\rho_{13} = \rho_{23} = –0.8$. Let the new weightings for the three instruments in the portfolio be 45%, 45% and 10%. The Calmar ratio for the augmented portfolio is now 1.308. This unexpected result shows how a losing trading system, which might initially be regarded as useless, is actually beneficial and leads to improved performance. The benefit of the negative correlation outweighs its lack of profit performance. It is as if this trading system/instrument provides 'cohesion' to the portfolio. This instrument could, for example, be a shorted group of stocks or indexes, thus providing the negative correlation with the rest of the portfolio of long stock positions. This result sheds some light into long-short portfolios. Not only do they serve as diversification vehicles by producing returns over different cycles from traditional long-only portfolios, but they can also produce better risk-adjusted returns.

Even though correlation is currently considered by the industry to be an important factor when deciding whether to add a trading system/instrument to a portfolio, it is usually second to the average return. With respect to risk-adjusted returns, the correlation is almost on a par with average returns, and deserves to be given a higher weight (when evaluating a trading strategy).

Conclusion

The MDD is one of the most important risk measures. To be able to use it more effectively, its analytical properties have to be understood. As a step in this direction, we have presented a review of some analytic results that we have developed as well as some applications of the analysis. In particular, we highlight the introduction of the normalised Calmar ratio as a way to compare quantitatively the Calmar ratios of portfolios over different time horizons. We also indicate the possibly underrated role of correlations in the performance of portfolios, and these correlations can be systematically incorporated in optimising the Calmar ratio of a portfolio. We hope this study will spur more analysis of this important risk measure. ■

Malik Magdon-Ismail is an assistant professor in the department of computer science at the Rensselaer Polytechnic Institute in Troy, New York. He also consults in computational finance. Amir Atiya is an associate professor in the department of computer engineering at Cairo University, and a consultant in the financial industry. Email: magdon@cs.rpi.edu, amir@alumni.caltech.edu

REFERENCES

Acar E and S James, 1997
Maximum loss and maximum drawdown in financial markets
In Proceedings of International Conference on Forecasting Financial Markets, London

Burghardt G, R Duncan and L Liu, 2003
Deciphering drawdown
Risk September, pages S16–S20

Chekhlov A, S Uryasev and M Zabarankin, 2003
Drawdown measure in portfolio optimization
Technical report, ISE Department, University of Florida, September

Douady R, A Shiryaev and M Yor, 2000
On probability characteristics of downfalls in a standard Brownian motion
Siam, Theory Probability Appl 44(1), pages 29–38

Harding D, G Nakou and A Nejjar, 2003
The pros and cons of drawdown as a statistical measure for risk in investments
AIMA Journal, April 2003, pages 16–17

International Advisory Services Group
http://iasg.pertrac2000.com/main-frame.asp

Karatzas I and S Shreve, 1997
Brownian motion and stochastic calculus
Springer

Magdon-Ismail M, A Atiya, A Pratap and Y Abu-Mostafa, 2003
The maximum drawdown of the Brownian motion
In Proceedings of the IEEE Conference on Computational Intelligence in Financial Engineering 2003, Hong Kong, March, IEEE Press

Magdon-Ismail M, A Atiya, A Pratap and Y Abu-Mostafa, 2004
On the maximum drawdown of a Brownian motion
Journal of Applied Probability 41(1), March

Medova E, 2000
Measuring risk by extreme values
Risk November, pages S20–S26

Q-functions
http://www.cs.rpi.edu/_magdon/data/Qfunctions.html

Sornette D, 2002
Why do stock markets crash
Princeton University Press