OXFORD UNIVERSITY PRESS LTD JOURNAL **00** (0000), 1–15
doi:10.1093/OUP Journal/XXX000

Moment-free Sharpe ratio estimation from total drawdown durations

Damien Challet
Université Paris-Saclay

Received 07/10/2015

*Correspondence to: Damien Challet, Laboratoire MICS, CentraleSupélec, Université Paris-Saclay, 92290 Châtenay-Malabry, France, `damien.challet@centralesupelec.fr`, phone: 00 33 1 41 13 17 55

Abstract

The total duration of drawdowns is shown to be an efficient and robust estimator of Sharpe ratios. Its properties are distribution-dependent: the expected total drawdown duration is smaller for heavy-tailed returns than for Gaussian ones. As a consequence, in leptokurtic market conditions, the new estimator yields smaller Sharpe ratios than moment-based estimators, which implies that the standard estimator overestimates the information content of prices when the return distribution has heavy tails. Accordingly, using the standard estimator for taking trend-following decisions enhances large price fluctuations.

(*JEL*: C14, C58)

KEYWORDS: Keywords: Sharpe ratio, drawdown total duration, robust estimation, efficiency, price records

Introduction

Sharpe ratios (Sharpe, 1964) appear naturally in financial analysis and understandably so: they are nothing else than signal-to-noise ratios, a fundamental quantity in signal analysis. In a Gaussian world, they are also equivalent to the t-statistics. Finance does not live in an ideal world, however, and many problems arise in practice. Sharpe ratio's distribution (Lo, 2002), bias (Miller and Gehr, 1978; Jobson and Korkie, 1981) and corrections due to serial correlations (Lo, 2002; Mertens, 2002; Christie, 2005; Opdyke, 2007) have been characterized. Better estimating methods use the Generalized Moments Method (Lo, 2002; Christie, 2005) and block bootstraps (Ledoit and Wolf, 2008). Although Sharpe ratios only depend on the first and second moments of price returns, their variance depends on the third and fourth moments (Lo, 2002; Mertens, 2002; Christie, 2005; Opdyke, 2007). Given the definition of the Sharpe ratios, it is not surprising that all these methods rely on the computation of moments of price returns. But as noted e.g. in Opdyke (2007), this may be problematic as the fourth moment may not be defined (Dacorogna et al., 2001; Jondeau and Rockinger, 2003).

Here, I propose a new way to estimate Sharpe ratios that does not require the computation of any moment and that may be extended to measure the drift of time series with infinite variance. It is based on the fact that the total duration of all drawdowns in a price time series of a given length is a monotonic function of the Sharpe ratio; by symmetry, the same holds for the total duration of all drawups. As a consequence, one may estimate Sharpe ratios by computing the difference between the total durations of drawups and drawdowns. This quantity is bounded by definition and leads to an estimator that is both robust to outliers and more efficient than direct estimates of Sharpe ratios for heavy-tailed data.

Intuitively, the sum of all drawdown durations, i.e., the total drawdown duration of a time series of fixed length is linked to the number of upper price records since a new price return pushes the price either to an all time high (a new upper record) or to a drawdown (see Fig. 1). This implies that if n is the length of a time series and R_+ is the number of its upper price records, the total drawdown duration, denoted by T_-, is $T_- = n - R_+$. Because of this equivalence, total drawdown/up duration and the numbers of price records lead to two equivalent estimators; accordingly, we will use either wordings. Assuming that log prices are random walks, drawdown/up durations are determined by first-passage times, themselves derived from persistence (or survival) properties (Redner, 2001). The connection between persistence and price dynamics, especially in the context of market microstructure, is well known (Lo et al., 2002; Eisler et al., 2009).

Persistence is at the core of a noteworthy recent result about discrete-time unbiased random walks, derived in a different scientific field. In a financial context, it may be stated as follows: the distribution of the number of upper (or lower) records of a price time series with independent and identically distributed return (i.i.d.), of a fixed length, does not depend on the increment distribution provided that the latter is symmetric and continuous (Majumdar and Ziff, 2008). This universality is behind the robustness and power of the r-statistics, a family of statistics based on the number of records of a time series, which not only provides a powerful non-parametric location test (Challet, 2015) but also, as shown here, an efficient estimator of Sharpe ratios. Their robustness come from the fact that the influence of outliers is much dampened because sample values are transformed into an integer number with bounded admissible values.

Majumdar et al. (2012) show that the distribution of the number of records converges to a Gaussian distribution in the limit of infinitely long time series provided that the price return distribution has a finite variance. Even better, the support of the finite-size sample distribution of the new estimator is bounded, contrarily to that of Sharpe ratios (and t-statistics), and is accordingly more peaked than a normal distribution (Challet, 2015). When the true Sharpe ratio is different from zero, the expected number of records and its variance are distribution-dependent; exact expressions are only known for exponentially distributed increments, hence one has to resort to approximations and numerical simulations for other types of distribution in the limits of large and small Sharpe ratios.

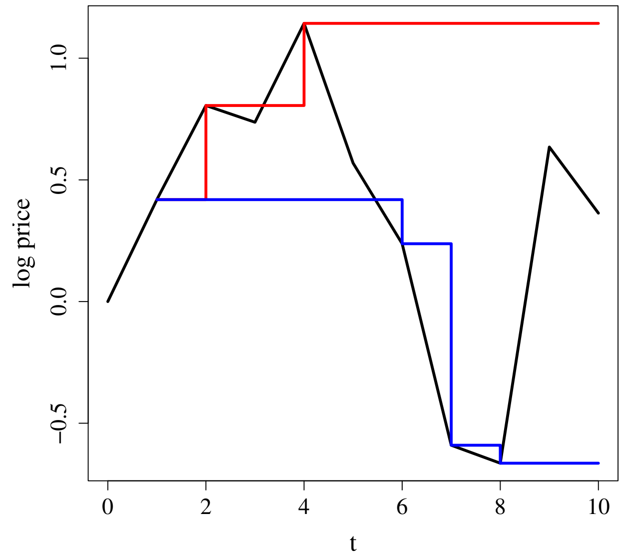

Figure 1: Example log price time series (black lines), its running maximum (red lines), and running minimum (blue lines). The number of upper (lower) records R_+ (R_-) is equal to the number of jumps of the running maximum (minimum) plus one since the first point counts as a record by convention: here $R_+ = 3$ and $R_- = 4$. The total drawdown duration is $T_- = 7$ and the total drawup duration is $T_+ = 6$. Clearly, $R_+ + T_- = R_- + T_+ = 10$, the number of returns.

Drawdown durations are by definition bounded integer numbers, which is not optimal to estimate a real number. The solution comes from random permutations. Assuming that the price returns are i.i.d., one can shuffle their order at will. As a consequence, the resulting (shuffled) price time series will be an equally valid representation of a given set of price returns and may lead to a different number of upper and lower records. Thus, to obtain a more precise estimate of the Sharpe ratio, one takes the average between the total drawdown and drawup durations over many such permutations (see Fig. 3 for a graphical explanation).

The structure of this paper is as follows: Section 1 introduces the necessary notations to define price record statistics and shows that when prices have a positive trend, heavy-tailed increments lead to more upper price records than Gaussian increments; a mathematical derivation of the expected number of price records for Student's t-distributed increments is reported in Appendix A, which focuses on the case of tail exponent equal to 4 (3 degrees of freedom) for the sake of analytical tractability. Section 2 investigates the efficiency of the number of price records as Sharpe ratio estimators relative to the vanilla estimator and shows that the new estimator is several times more efficient than moment-based methods for heavy-tailed variables and almost as efficient as the vanilla estimator in the case of Gaussian variables. Section 3 estimates the 252-day rolling Sharpe ratios of SPY and of 400 US equities with both methods. It turns out that in leptokurtic times, the estimates from both methods may differ very significantly because the vanilla Sharpe ratio estimator is not only more volatile, but also systematically overestimates the information content of price time series that have heavy-tailed returns. Finally, a naive trading strategy illustrates the advantages of the new estimator: taking a long or short position when the absolute value of the estimated Sharpe ratio is large enough leads to very different outcomes depending on which estimator one uses.

1 Record statistics of random prices

Financial data exist in discrete time, which will be the point of view adopted in this paper. Let us assume that the initial log price is $S_0 = 0$ and that its value at time $n > 0$ follows

$$S_n = S_{n-1} + r_n + c \tag{1}$$

where r_n is the increment at time n, assumed to be identically and independently drawn from a continuous distribution $P(r)$, and c is a constant trend. Let the running maximum $M_k = \max_{1 \leq t \leq k} S_t$ (see Fig. 1). The number of upper records of a time series of length n is the number of jumps of M_n, which by convention always includes M_1; it will be denoted by R_+ and its distribution by $P(R_+, n)$. In the same spirit, one defines R_-, the number of lower records, as the number of jumps of the running minimum.

Majumdar and Ziff (2008); Le Doussal and Wiese (2009); Majumdar et al. (2012) demonstrate that many quantities of interest are fully characterized by the persistence function $q_-(n)$ of the process, i.e., the probability that the price has never exceeded its starting value after n steps. It is advantageous to work with its characteristic function $\tilde{q}_-(z) = \sum_{n>0} z^n q_-(n)$.

For example, the characteristic function of $P(R_+, n)$ is $\tilde{P}(R_+, z) = [1 - (1-z)\tilde{q}_-(z)]^{R_+-1}\tilde{q}_-(z)$ (Majumdar and Ziff, 2008), while the characteristic function of the expected number of upper records $m_+(n) = E(R_+)(n)$ can be written as $\tilde{m}_+(z) = [(1-z)^2\tilde{q}_-(z)]^{-1}$ (Le Doussal and Wiese, 2009).

Generalized Sparre Andersen theorem (Andersen, 1953; Feller, 2008) provides a constructive way to compute $\tilde{q}_-$: for any continuous and symmetric $P(r)$,

$$\log\left(\tilde{q}_-(z)\right) = \sum_{n=1}^{\infty} \frac{z^n}{n} P(S_n < 0). \tag{2}$$

A direct consequence of this theorem is the universality of the unbiased case $c = 0$ since $P(S_n < 0) = \frac{1}{2}$ for all symmetric and continuous distributions. As a consequence, $\tilde{q}_\pm(z) = \tilde{q}(z)$ is the same for all such distributions and

$$P(R, n) = \binom{2n - R + 1}{R} 2^{-2n+R-1},$$

where R may either be R_+ or R_-, by symmetry (Majumdar and Ziff, 2008). This implies that the first two moments of this distribution are

$$E(R_\pm)(n) \simeq 2\sqrt{\frac{n}{\pi}}, \quad \text{and} \quad E[((R_\pm - E(R_\pm))^2](n) \simeq (2 - 4/\pi)n.$$

Analytical results are harder to obtain in the case of non-zero drift ($c \neq 0$) since Sparre Andersen theorem requires the full knowledge of all convolutions of the elementary increments. Denoting the standard deviation of the increments r_k by σ^2, good approximations of the expected number of upper records are known for Gaussian increments in the case of small Sharpe ratios, i.e., when $c/\sigma \ll 1$ and $n \gg 1$ while $cn/\sigma \ll 1$ (Majumdar et al., 2012):

$$E(R_+)(c/\sigma, n) \simeq 2\sqrt{\frac{n}{\pi}} + \frac{c\sqrt{2}}{\sigma\pi}\left[n \arctan(\sqrt{n}) - \sqrt{n}\right]. \tag{3}$$

The case of heavy-tailed increments of finite variance has not been thoroughly investigated. We will focus on Student's t-distributions because of their abilities to describe fat-tailed and Gaussian returns. They are known to describe the unconditional price return distribution (i.e., forgetting about volatility heteroskedasticity) (Bouchaud and Potters, 2000; Longin, 2005; Opdyke, 2007) and innovations (see e.g. Bollerslev (1987)). Let us therefore assume from now on that the price returns r_k are distributed according to a Student's t-distribution of variance σ with ν degrees of freedom (we use this wording only to parametrize the return distribution), denoted by $P(r)$. Sparre Andersen theorem requires the knowledge of the n-time convoluted return distribution, denoted by $P^{(n)}(r)$, of which no explicit expression exists for generic values of n and ν. In passing, $P^{(n)}(r)$ can be explicitly computed for any value of n provided that ν is odd but the expressions quickly become cumbersome as n grows (Nadarajah and Dey, 2005). This is why we shall resort to approximations.

Appendix A reports analytical results for the case $\nu = 3$, i.e., for a tail exponent of 4.[1] The resulting expected number of upper records becomes, in the same limit $c/\sigma \ll 1$ and $n \gg 1$ while $cn/\sigma \ll 1$,

$$E(R_+)(n, c/\sigma) \simeq \frac{2\sqrt{n}}{\sqrt{\pi}} + \frac{c\sqrt{2}}{\sigma\pi}\left[n \arctan(\sqrt{n}) - \sqrt{n}\right] + \frac{c}{\sigma}\frac{8}{\sqrt{3}\pi^{3/2}}\sqrt{n}\left(\operatorname{atanh}\sqrt{1 - \frac{1}{n}} - \sqrt{1 - \frac{1}{n}}\right). \tag{4}$$

Although a first order expansion, Eq. (4) is not very accurate even in the limit of small $n(c/\sigma)$, because the approximations needed to obtain explicit equations are quite rough (see Fig. 2). However, it was worth computing it for several reasons. First, it contains the correct dependence of $E(R_+)(n, c/\sigma)$ on n for small Sharpe ratios, which means that one may use this functional form to fit numerical simulations. Second, the presence of the third term, due to the difference between Gaussian and t-distributions at the origin, correctly implies that the prices with positive trends and heavy tails (and small Sharpe ratios) have a larger expected number of price records, which emphasises the importance of accounting for the tails of price return distributions when using price records to estimate Sharpe ratios. Appendix B contains the derivation of $E(R_+)(n, c/\sigma)$ in the large Sharpe ratio limit.

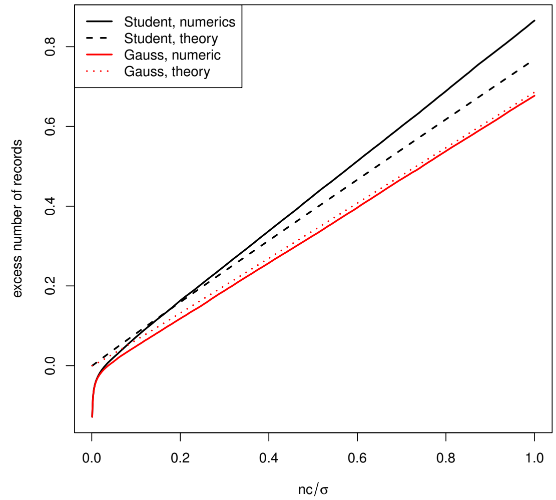

Figure 2: Excess number of records $E(R_+|c/\sigma, n) - E(R_+|0, n)$ for biased random walks with Student-t increments ($\nu = 3$). Interrupted lines are theoretical predictions and continuous lines are from numerical simulations. $c = 0.001$, $\sigma = 1$, averages over 10^7 samples.

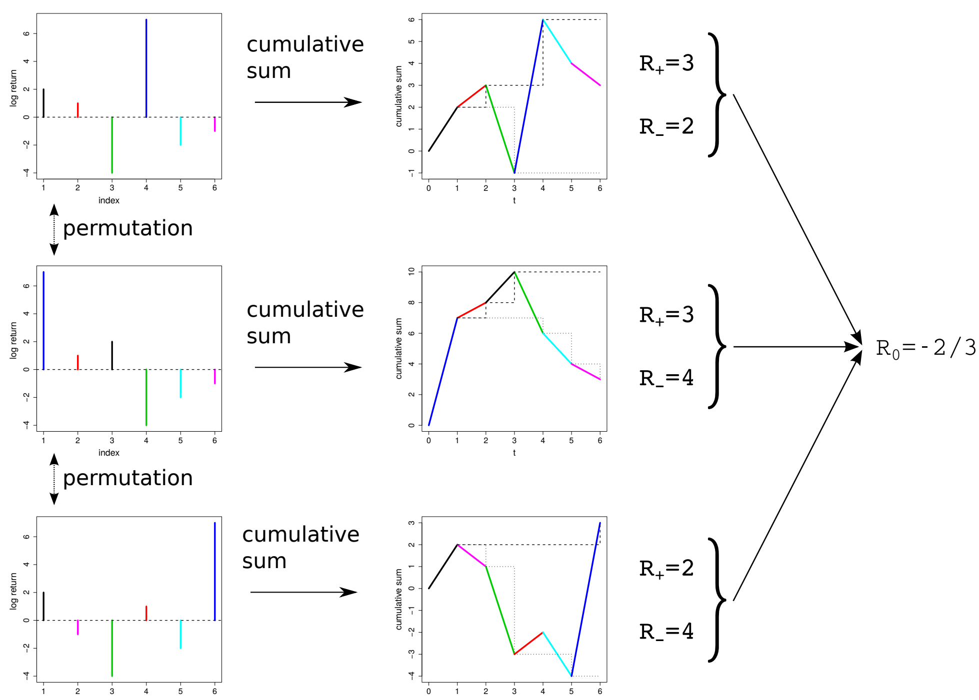

Figure 3: Schematic explanation of the idea behind the permutation estimator of Sharpe ratios: one computes the difference between total drawup and drawdown durations, or equivalently, the number of jumps of the running maximum (dashed lines) and the number of jumps of the running minimum (dotted lines) of the cumulated sums of the sample values, averaged over many random permutations. By convention, the first point counts as a first jump for both the running maximum and minimum.

2 Estimating Sharpe ratios with permutations

The fact that the expected number of price records is a monotonous function of the Sharpe ratio $\theta = c/\sigma$ means that one may estimate the latter from measures of the former. The main problem of a number of records is that it is an integer number by definition, which yields an estimator with unacceptable precision for short time series. The fundamental idea of the r-statistics (Challet, 2015), in this context, consists in assuming that its log returns are i.i.d.. In that case, one may build many other log price paths based on random permutations of the original

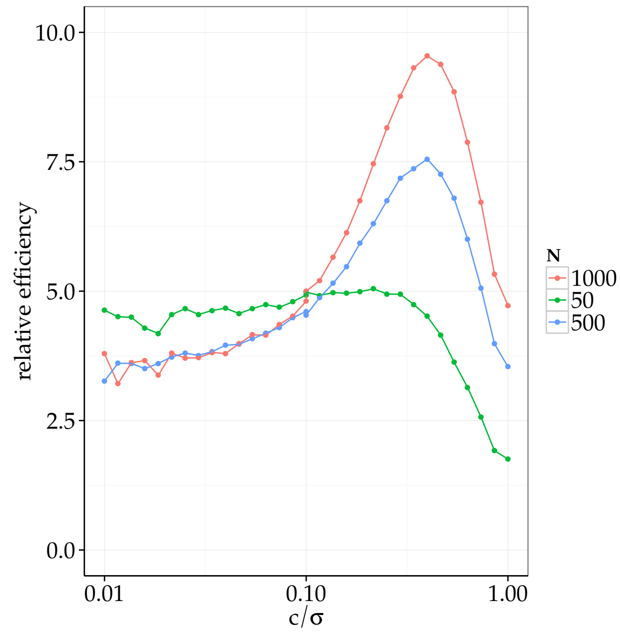

Figure 4: Efficiency of the record-based estimator θ_0 relative to that of the vanilla estimator, defined by the ratio of the variance of the new estimator θ_0 and the usual one θ_S as a function of the true Sharpe ratio c/σ of the synthetic data. Averages over 10^6 samples per point; record numbers have been averaged over 1000 permutations; Student-distributed increments with tail exponent set to 4.

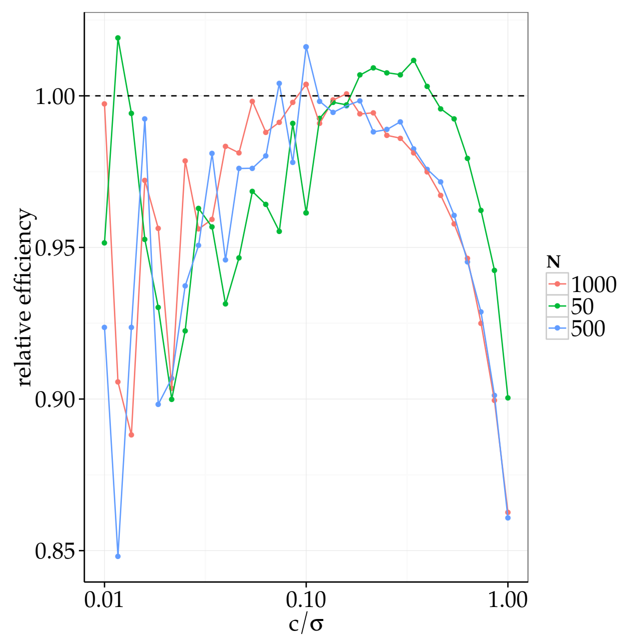

Figure 5: Efficiency of the record-based estimator θ_0 relative to that of the straightforward t-statistics, defined by the ratio of the variance of the new estimator θ_0 and the usual one θ_S as a function of the true Sharpe ratio c/σ of the synthetic data. Averages over 10^6 samples per point for $N = 500$ and $N = 1000$, and 10^7 samples per point for $N = 50$; record numbers have been averaged over 1000 permutations; Gaussian-distributed increments.

returns and thus measure the average number of records of the cumulated sums over many permutations (see Fig. 3). Mathematically, denoting the random permutation of index $i \in \{1, \cdots, n\}$ by $\pi(i)$ and the ensemble of all permutations by Π, the average number of records is $\bar{R}_+ = \frac{1}{|\Pi|}\sum_{\pi\in\Pi} R_{+,\pi}$ where $R_{+,\pi}$ is the number of upper records of $S_{n,\pi} = \sum_{m=1}^{n} r_{\pi(m)}$. In practice, one restricts computations to a subset of Π for speed reasons, which has little influence on the end result. The new Sharpe ratio estimator is then $R_0 = \bar{R}_+ - \bar{R}_-$.

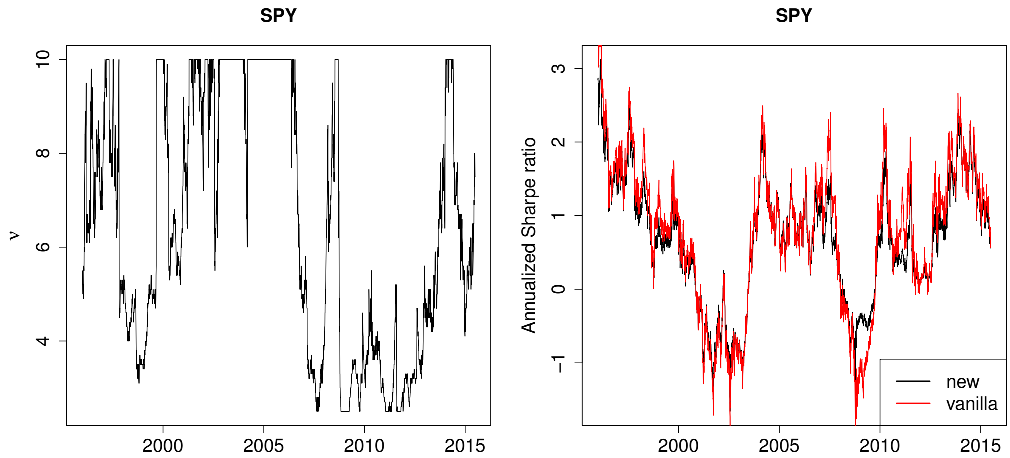

Figure 6: Left plot: parametric fit of the number of degrees of freedom of Student's t-distribtuion in a sliding window of 252 close-to-close returns of SPY. Right plot: estimated Sharpe ratios with the new estimator and from a vanilla estimation. 1000 permutations have been used to estimate R_0.

Equation (4) implies that R_0 is also a function of θ, hence that one may invert the exact relationship $E(R_0)(n,\theta)$ to infer θ. Note that $E(R_0)(n,\theta) \neq E(\bar{R}_+)(n,\theta) - E(\bar{R}_-)(n,\theta)$ because the number of upper and lower records of a single random walk are not independent. Let us denote by θ_0 the Sharpe ratio inferred from R_0. The standard deviation of θ_0, denoted by σ_θ, is obtained by the method of Deltas, i.e., from the relationship $\sigma_\theta = \sigma_R \frac{1}{\frac{dE(R_0|n,\theta)}{d\theta}}$ where σ_R is the standard deviation of R_0; the numeric derivative of $E(R_0|n,\theta)$ was computed numerically with splines. The relative efficiency of θ_0 with respect to the straightforward estimator $\theta_S = \hat{\mu}/\hat{\sigma}$ is then defined as $\rho = \sigma_S^2/\sigma_R^2$ where σ_S is the standard deviation of θ_S. Figure 2 reports the relative efficiency of θ_0 for various N for Student's t-distributed returns and $\nu = 4$. The new estimator is unambiguously more powerful than the vanilla estimator. This result holds as long as the returns are heavy tailed.

For the sake of completeness, we computed the efficiency of record statistics for log prices with Gaussian increments. Since the vanilla estimator is asymptotically optimal in this case (Neyman and Pearson, 1933), any other estimator is bound to be less efficient for large N. Figure 5 plots the relative efficiency of θ_0 for Gaussian increments, which depends on c/σ. Remarkably, θ_0 may be slightly more efficient than the t-statistics itself for small N.

3 Application to real data

The i.i.d. assumption is totally unrealistic regarding asset price returns, if only because of volatility heteroskedasticity. Applying straightforwardly the above estimator would therefore make little sense on long time series. The approach followed here is to consider smaller time windows and to assume that stationarity approximately holds in each time window. The second current limitation of the proposed estimator to keep in mind here is that it does not account explicitly for skewness. At any rate, this section is meant to provide a clear illustration of how different the estimates of both methods may be.

The method is as follows: we first perform extensive numerical simulations to establish the relationships $E(R_0|n,\theta,\nu)$ for $\nu = \{2.5, \cdots, 10\}$ with increments of 0.1 and $n = 252$ (a year of daily data); we take 31 values of $\theta \in [0.001, 1]$ that grow geometrically. Since there is little difference between a Gaussian distribution and a Student's t-distribution when $\nu = 10$, this range of values of ν is able to account for a variety of market conditions. Then, in each time window, we measure R_0 averaged over 1000 permutations. In order to find the corresponding Sharpe ratio, we assume that price returns are conditionally leptokurtic (Bollerslev, 1987): in each time window, we fit the returns with Student's t-distribution by maximum likelihood and obtain an estimate $\hat{\nu}$. Finally, Sharpe ratios are determined from the pre-computed relationship between R_0 and θ with $\nu \simeq \tilde{\nu} = \texttt{round}(10\hat{\nu})/10$, bounded by 2.5 and 10.

Figure 6 shows the difference between annualized Sharpe ratios of SPY estimated with the new and vanilla estimator. When $\tilde{\nu}$ is close to 10, both estimators yield the same Sharpe ratio, as expected. On the other hand, when tails are heavier, i.e. when $\tilde{\nu} < 10$, the two estimates significantly differ. Indeed, the new term in Eq. (4) with respect to Eq. (3) implies that vanilla estimates are too large in absolute values. This is confirmed in Fig.

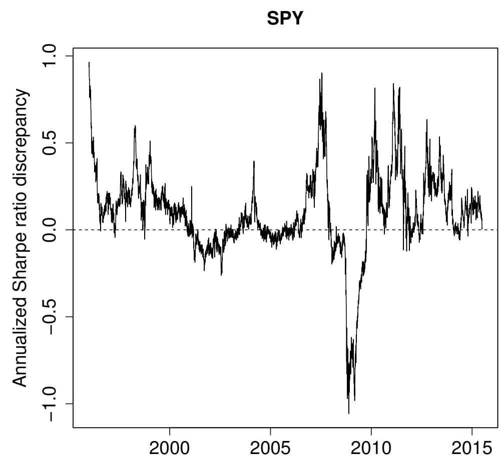

Figure 7: Discrepancy of the estimates of the annualized Sharpe ratio of SPY with moving time windows of 252 days between the new and the vanilla estimators. 1000 permutations have been used to estimate R_0.

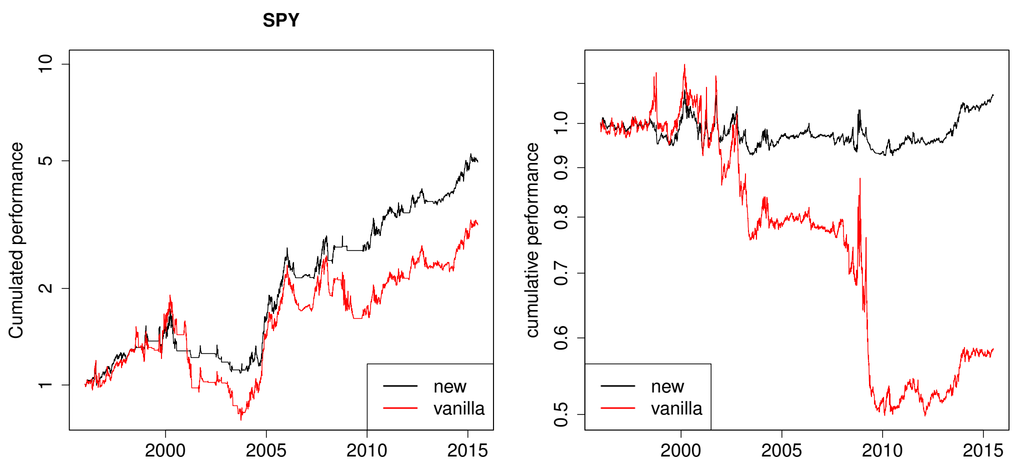

Figure 8: Cumulated performance of a trading strategy consisting in investing when the estimated annualized Sharpe ratio is larger than 1 in absolute value; short positions are allowed in case of negative Sharpe ratio; estimates over rolling windows of 252 trading days (close-to-close price returns). Left plot: SPY. Right plot: equivally weighted portfolio built with same strategy applied to 400 equities whose symbols start with A; unbiased historical database of liquid assets (price larger than \$20, rolling median daily volume larger than 250000 shares, computed over 60 trading days rolling windows); no transaction costs. 1000 permutations have been used to estimate R_0.

7. The difference between both estimates is very large in leptokurtic times, e.g. in 2008 and 2009; in addition, in these difficult times, the new estimator is clearly less volatile.

The fact that the moment-based method overestimates the Sharpe ratio in leptokurtic times means that anybody using it for trading purposes would be lead to take wrong trading decisions more often (the power of the r-statistics is indeed much larger than that of t-statistics for heavy-tailed data (Challet, 2015)). Let us try the following naive trading strategy (without transaction costs): whenever the estimated annualized Sharpe ratio is larger than 1 in absolute value in the last 252 close-to-close price returns, one takes a long or short single-day position, depending on the sign of the Sharpe ratio. Figure 8 reports the cumulated performance of this strategy when applied to SPY and to 400 US equities for the period 1995-01-01 to 2015-06-30. The difference is unambiguous and stems from periods either with a marked change of trend or high volatility.

4 Discussion

The proposed Sharpe ratio estimator is robust, efficient, and well-behaved as it does not rely on moment estimation. Large returns are not regarded as outliers, but contribute to record statistics in a smooth way. In addition, a real outlier (due e.g. to a data error, or a corporate action) may only create one spurious additional price record, while two outliers of the same magnitude and opposite signs have only a mild influence on R_0. Finally, the robustness of the estimator lies in the fact that the latter is based only on the duration of drawdowns, not on their amplitudes. This is to be contrasted with other quantities related to drawdowns. For example the expectation of the maximum drawdown of a Brownian motion is a known function of the Sharpe ratio (Magdon-Ismail et al., 2003), but is very sensitive to outliers, by definition.

Because of the lack of exact results, using this estimator requires to compute once the relationship between the number of upper records and Sharpe ratios for a given time series length numerically, which takes a few hours with current computers (full source code is available).

Estimating Sharpe ratios with price record/drawdown statistics is not limited to Student's t-distributed returns, as indeed one may calibrate their relationships for any return distribution with finite variance. In addition, the method introduced in this paper provides a generic way to build many types of estimators with record statistics as long as the relationship between price record statistics and the measurable to estimate is monotonic. For example, it may be used to estimate the drift of a Lévy process.

The main limitations of the proposed estimator are the assumptions of i.i.d and symmetric increments. Both must be accounted for numerically for the time being. An interesting challenge is to incorporate serial correlations into the analytical computation of record statistics: numerical results point to simple corrections in the case of AR(1) and GARCH(1,1) models (Wergen, 2014). Practically, a way to respect return auto-correlation and volatility heteroskedasticity is to use block-bootstraps, as in Ledoit and Wolf (2008).

Full source code (R and C++) and pre-calibrated estimators for time series of length 252 are available at `https://github.com/damienchallet/moment-free_sharperatio` .

An interactive webpage producings the plots of Section 3 for any symbol and time period may be found at `https://brillant.shinyapps.io/moment-free_Sharpe_ratio` .

A Expected number of records in the small Sharpe ratio limit

In this limit, one may use a first order expansion of the reciprocal cumulative function

$$P(S_n > 0) = \frac{1}{2} + P^{(n)}(0)cn + O([cn]^2). \tag{5}$$

One therefore needs to compute $P^{(n)}(0)$. Since the increments are assumed to be independent,

$$P^{(n)}(0) = \frac{1}{2\pi}\int_{-\infty}^{\infty}\phi^{(n)}(t)dt = \frac{1}{2\pi}\int_{-\infty}^{\infty}[\phi(t)]^n dt,$$

where $\phi^{(n)}(t)$ is the characteristic function of $P^{(n)}(x)$, and $\phi(t)$ that of $P^{(1)}(r) = P(r)$. Even if Eq. (5) only requires the computation of $P^{(n)}(0)$, which is nevertheless impossible for any n and ν. However the $\nu = 3$ case leads to workable expressions. One finds $P^{(n)}(0) = \frac{e^n}{\sigma\pi}E_{-n}(n)$, where $E_n(z)$ is the exponential integral function. The specific form $n = -z$ of the exponential integral function is easy to compute in a recursive way by integration by parts:

$$\begin{aligned} E_{-k-n}(k) &= \frac{e^{-k}}{k} + \frac{k-n}{k}E_{-(k-n-1)}(k) \\ E_0(k) &= \frac{e^{-k}}{k}. \end{aligned}$$

Therefore, after some elementary computations, $E_{-n}(n) = \frac{e^{-n}}{n}\frac{n!}{n^n}\left(\sum_{s=0}^{n}\frac{n^s}{s!}\right)$ and

$$P^{(n)}(0) = \frac{1}{\sigma\pi}\frac{1}{n}\frac{n!}{n^n}\sum_{s=0}^{n}\frac{n^s}{s!}. \tag{6}$$

Using the asymptotic expansion $K_n = \sum_{s=0}^{n}\frac{n^s}{s!} = e^n[\frac{1}{2} + \sqrt{\frac{2}{3\pi}}\frac{1}{\sqrt{n}} + O(n^{-1})]$ and the usual Stirling expansion, Eq. (6) then becomes $P^{(n)}(0) = \frac{1}{\sigma}\frac{1}{\sqrt{2\pi n}} + \frac{2}{\sigma\pi\sqrt{3n}} + O(n^{-3/2})$ and thus, in the case of small drifts, Eq. (5) reads

$$P(S_n > 0) = \frac{1}{2} + \frac{c}{\sigma}\frac{2}{\pi\sqrt{3}} + \frac{c}{\sigma}\sqrt{\frac{n}{2\pi}} + O(n^{-1/2}).$$

Higher-order expansions of K_n and $n!$ contribute terms of order $n^{-1/2}$ that are negligible. It is noteworthy that the additional correction for Student increments does not depend on n; accordingly, it is relevant for any value of n and has a larger relative weight for smaller n; this is consistent with the fact that convolutions of Student's t-distributions with $\nu = 3$ converge to a Gaussian distribution. Sparre Andersen theorem yields

$$\tilde{q}_-(z) = \frac{1}{\sqrt{1-z}}\left(1+\sum_{n=1}^{\infty}\frac{c}{\sigma}\frac{z^n}{\sqrt{2\pi n}}\right) - \frac{c}{\sigma}\frac{2}{\pi\sqrt{3}}\frac{\log(1-z)}{\sqrt{1-z}} + O[(c/\sigma)^2]. \tag{7}$$

The generating function of the number of records is then (Le Doussal and Wiese, 2009)

$$\tilde{m}_+(z) \simeq \frac{1}{(1-z)^{3/2}}\left[1+\frac{c}{\sqrt{2\pi}\sigma}\sum_{n=1}^{\infty}\frac{z^n}{\sqrt{n}} - \frac{2c}{\sigma\pi\sqrt{3}}\log(1-z)\right].$$

The two first terms in the brackets are the same ones as those of Gaussian biased random walks (Majumdar et al., 2012). The third term is new and due to the difference between a Gaussian and a t-distribution at the origin. The only way to make progress is to approximate sums by integrals, which yields

$$-\frac{1}{(1-z)^{3/2}}\log(1-z) \simeq \frac{2}{\sqrt{\pi}}\sum_{n\geq 1}\left[2\sqrt{n}\left(\operatorname{atanh}\sqrt{1-\frac{1}{n}} - \sqrt{1-\frac{1}{n}}\right)\right]z^n, \tag{8}$$

which is not a very good approximation even for large n but gives the correct asymptotic $\sqrt{n}$ dependence, with an additional logarithmic correction brought by $\operatorname{atanh}\sqrt{1-\frac{1}{n}} - \sqrt{1-\frac{1}{n}}$. Finally, approximating n by $n-1$ as in Wergen (2014) and identifying each term of the generating function with the value one is interested in gives

$$E(R_+)(c/\sigma, n) \simeq \frac{2\sqrt{n}}{\sqrt{\pi}} + \frac{c\sqrt{2}}{\sigma\pi}\left[n\arctan(\sqrt{n}) - \sqrt{n}\right] + \frac{c}{\sigma}\frac{8}{\sqrt{3}\pi^{3/2}}\sqrt{n}\left(\operatorname{atanh}\sqrt{1-\frac{1}{n}} - \sqrt{1-\frac{1}{n}}\right). \tag{9}$$

Given its derivation, this formula is relevant in the limit $cn \ll \sigma$ and large n.

B Expected number of price records in the large Sharpe ratio limit

Although quite rare in a financial context, the large Sharpe ratio limit also makes it possible to derive some analytical insights. Majumdar et al. (2012) give results for large, but not too large, cn/σ. Indeed, the central limit theorem states that the convergence of the distribution of convoluted variables to a Gaussian distribution occurs from the center of the distribution. This implies that the tails of any non-Gaussian distribution are non-Gaussian. Thus, intuitively, when cn/σ is large enough (whose meaning will be discussed below), $P(x_n < cn)$ comes from the non-Gaussian tails. This will lead to markedly different results for Student's t-distributions since the tails of convoluted t-distributions keep their power-law nature. Bouchaud and Potters (2000) give an intuitive argument to compute the n-time convoluted return $r_0^{(n)}$ at which the Student and Gaussian parts of the distribution have equal importance and find that $r_0^{(n)} \simeq \sigma\sqrt{n\log n}$ for $\nu = 3$. This means that the value of n_0 at which the power-law tail starts to prevail is such that $cn_0 \simeq \sigma\sqrt{n_0 \log n_0}$, i.e.,

$$\frac{c}{\sigma}\sqrt{n_0} \simeq \sqrt{\log n_0}. \tag{10}$$

Since the convoluted distribution has a continuous first derivative, there is no sharp transition between the Gaussian and power-law regimes, hence n_0 only approximately indicates where the Gaussian approximation begins to break down. Figure 9 plots $n_0(c/\sigma)$ and shows these two regions. In the region well below the line, a Gaussian approximation holds for Student convolutions. Reversely, when $n \gg n_0(c/\sigma)$,

$$P(S_n < 0) \simeq \int_{cn}^{\infty}\frac{2\sigma^3 n}{\pi x^4}dx = \left(\frac{\sigma}{c}\right)^3\frac{2}{3\pi n^2},$$

hence $\log\left(\tilde{q}_-(z)\right) = \left(\frac{\sigma}{c}\right)^3\frac{6\sqrt{3}}{\pi}\sum_{n=1}^{\infty}\frac{z^n}{n^3}$, thus

$$\tilde{m}_+(z) = \frac{1}{(1-z)^2}\exp\left[-\left(\frac{\sigma}{c}\right)^3\frac{2}{3\pi}\sum_{n=1}^{\infty}\frac{z^n}{n^3}\right] \simeq \frac{1}{(1-z)^2}\left[1-\left(\frac{\sigma}{c}\right)^3\frac{2}{3\pi}\sum_{n=1}^{\infty}\frac{z^n}{n^3} + O\left[\left(\frac{c}{\sigma}\right)^6\right]\right].$$

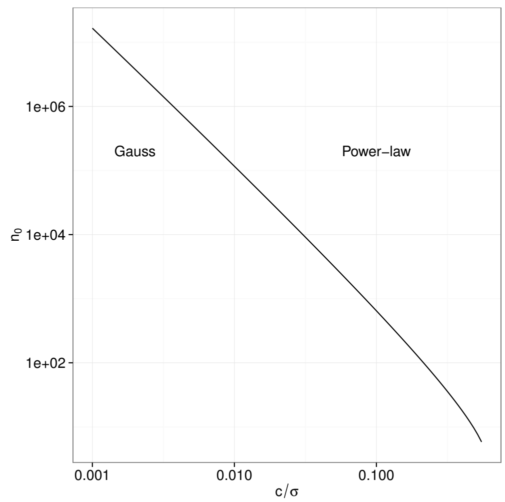

Figure 9: Limiting n_0 as a function of c/σ, from Eq. (10). Convoluted Student's t-distributions may be approximated by a Gaussian distribution below the continuous line, and by a power-law above this line.

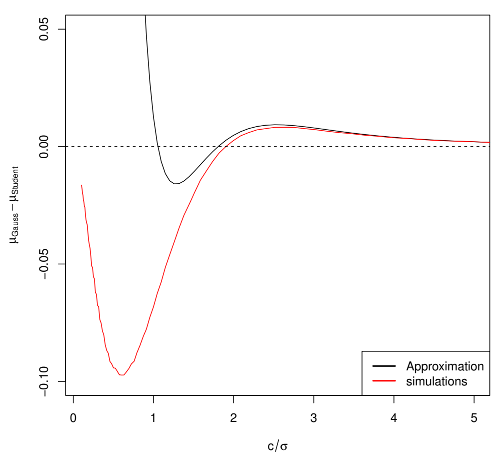

Figure 10: Large signal-to-noise ratio limit: difference of record rate between random walks with Gaussian increments and Student's t-distributed increments. The rate was determined as the average slope $[m(200) - m(100)]/100$ (averages over 10^5 samples).

Finally, one finds without major difficulty

$$\tilde{m}_+(z) \simeq \sum_{n\geq 0} \left[(n+1) \left(1 - \left(\frac{\sigma}{c}\right)^3 \frac{2}{3\pi} K[1 + O(n^{-1})\right) \right] z^n$$

and

$$m_+(n) \simeq n \left[1 - \left(\frac{\sigma}{c}\right)^3 \frac{2}{3\pi} K \right].$$

Numerically, $K \simeq 1.202 \simeq \frac{6}{5}$ for large n; approximating sums with integrals yields the very different $K = 1/2$. Thus the number of records increases linearly for large n $m_+(n) \simeq n\mu_{Student}$ with an asymptotic rate given by $\mu_{Student} \simeq 1 - \left(\frac{\sigma}{c}\right)^3 \frac{4}{5\pi}$, to be compared with $\mu_{Gauss} \simeq 1 - \frac{\sigma}{c} \frac{1}{\sqrt{2\pi}} e^{-\frac{c^2}{2\sigma^2}}$. Figure 10 plots the difference of the record rate between Gaussian- and Student's t-distributed ($\nu = 3$) increments as a function of c/σ. Whereas the number of records of random walks with Student increments are larger than those with Gaussian ones for small Sharpe ratios, Fig. 10 shows, somewhat surprisingly, that Gaussian increments lead to a larger rate of records for very large Sharpe ratios.

List of Figures

References

Andersen, E. S. 1953. On the Fluctuations of Sums of Random Variables. *Mathematica Scandinavica* 1:263–285.

Bollerslev, T. 1987. A Conditionally Heteroskedastic Time Series Model for Speculative Prices and Rates of Return. *The Review of Economics and Statistics* 69: 542–547.

Bouchaud, J.-P., and M. Potters. 2000. *Theory of Financial Risks.* Cambridge: Cambridge University Press.

Challet, D. 2015. "One- and two-Sample Nonparametric Location Tests Based on Record Statistics." Working paper, available at `http://arxiv.org/abs/1502.05367`.

Christie, S. 2005. "Is the Sharpe Ratio Useful in Asset Allocation?" Technical report, Macquarie Applied Finance Centre Research Paper.

Dacorogna, M. M., R. Gencay, U. A. Müller, R. B. Olsen, and O. V. Pictet. 2001. *An Introduction to High-Frequency Finance.* London: Academic Press.

Eisler, Z., J. Kertesz, F. Lillo, and R. N. Mantegna. 2009. Diffusive Behavior and the Modeling of Characteristic Times in Limit Order Executions. *Quantitative Finance* 9:547–563.

Feller, W. 2008 *An Introduction to Probability Theory and its Applications* vol. 2. Hoboken, NJ: John Wiley & Sons.

Jansen, D. W., and C. G De Vries. 1991. On the Frequency of Large Stock Returns: Putting Booms and Busts Into Perspective. *The Review of Economics and Statistics* 73: 18–24.

Jobson, J. D., and B. M. Korkie. 1981. Performance Hypothesis Testing With the Sharpe and Treynor Measures. *The Journal of Finance* 36:889–908.

Jondeau, E., and M. Rockinger. 2003. Conditional Volatility, Skewness, and Kurtosis: Existence, Persistence, and Comovements. *Journal of Economic Dynamics and Control* 27(10):1699–1737.

Le Doussal, P., and K. J. Wiese. 2009. Driven Particle in a Random Landscape: Disorder Correlator, Avalanche Distribution, and Extreme Value Statistics of Records. *Physical Review E* 79:051105.

Ledoit, O., and M. Wolf. 2008. Robust Performance Hypothesis Testing With the Sharpe Ratio. *Journal of Empirical Finance* 1:850–859.

Lo, A. W. 2002. The Statistics of Sharpe Ratios. *Financial Analysts Journal* 58:36–52.

Lo, A. W., A. C. MacKinlay, and J. Zhang. 2002. Econometric Models of Limit-Order Executions. *Journal of Financial Economics* 65:31–71.

Longin, F. 2005. The Choice of the Distribution of Asset Returns: How Extreme Value Theory Can Help? *Journal of Banking & Finance* 29:1017–1035.

Magdon-Ismail, M., A. Atiya, A. Pratap, and Y. Abu-Mostafa. The Maximum Drawdown of the Brownian Motion. 2003 In *Computational Intelligence for Financial Engineering, 2003. Proceedings. 2003 IEEE International Conference on*, 243–247.

Majumdar, S. N., and R. M. Ziff. 2008. Universal Record Statistics of Random Walks and Lévy Flights. *Physical Review Letters* 101:050601.

Majumdar, S. N., G. Schehr, and G. Wergen. 2012. Record Statistics and Persistence for a Random Walk With a Drift. *Journal of Physics A: Mathematical and Theoretical* 45(35):355002.

Mertens, O. 'Comments on the Correct Variance of Estimated Sharpe Ratios in Lo (2002, FAJ) When Returns are iid." Research Note, available at `http://www.elmarmertens.com/research/discussion/soprano01.pdf`, accessed 2015-06-04.

Miller, R. E., and A. K. Gehr. 1978. Sample Size Bias and Sharpe's Performance Measure: a Note. *Journal of Financial and Quantitative Analysis* 13:943–946.

Nadarajah, S., and D. Dey. 2005. Convolutions of the t Distribution. *Computers & Mathematics with Applications* 49:715–721.

Neyman, J., and E. Pearson. 1933. On the Problems of the Most Efficient Tests of Statistical Hypotheses. *Philosophical Transactions of the Royal Society of London* 231: 289–337.

Opdyke, J. D. J. 2007. Comparing Sharpe Ratios: so Where are the p-Values? *Journal of Asset Management* 8:308–336.

Plerou, V., P. Gopikrishnan, L. A. N. Amaral, M. Meyer, and H. E Stanley. 1999. Scaling of the Distribution of Price Fluctuations of Individual Companies. *Physical Review E* 60:6519.

Redner, S. 2001 *A Guide to First-Passage Processes.* Cambridge: Cambridge University Press.

Sharpe, W. 1964. Capital Asset Prices: A theory of Market Equilibrium Under Conditions of Risk. *The Journal of Finance* 19:425–442.

Wergen, G. 2014. Modeling Record-Breaking Stock Prices. *Physica A: Statistical Mechanics and its Applications* 396:114–133.

Notes

[1]This precise value is the only one for which analytical computations seem workable. It also happens to be in line with the average tail exponent of US equities daily and intraday price returns (Jansen and De Vries, 1991; Plerou et al., 1999; Bouchaud and Potters, 2000; Longin, 2005).