# Using Statistical Process Control to Monitor Active Managers

*Rapidly identify troubled portfolios with confidence.*

Thomas K. Philips, Emmanuel Yashchin, and David M. Stein



Most traditional performance measurement algorithms measure the performance of a portfolio over a fixed horizon, typically the three to five most recent years, and suffer consequently from a serious limitation. Good performance in some years can mask poor performance in others, making it difficult to estimate the portfolio's current performance, and harder still to identify transitions from good performance to bad.

Indeed, it is often claimed that it takes 40 years to determine whether an active manager has outperformed a benchmark by a statistically significant margin. While this is true if the manager's mean excess return is stationary and if one requires 95% confidence that it is positive, it is of little use to investors, who should not and do not wait 40 years to determine whether their portfolio's mean excess return is both stationary and positive.

In fact, time variation in the mean excess return is the norm in an efficient market. As inefficiencies appear and are then arbitraged away, all investment processes will at some times flourish, and at other times stagnate or stumble, necessitating a fundamentally different approach to performance measurement. Investors ought to continually estimate the *current* performance of their portfolios, and to rigorously reevaluate each manager's investment strategy as soon as they can determine that it no longer adds value.

Dynamic performance measurement and change point detection are closely related, and have long been addressed in the fields of sequential analysis and statisti-

**Thomas K. Philips** is chief investment officer at Paradigm Asset Management in New York (NY 10019).
**tkp@paradigmasset.com**

**Emmanuel Yashchin** is a research staff member at IBM's T.J. Watson Research Center in Yorktown Heights (NY 10547).
**yashchi@us.ibm.com**

**David M. Stein** is chief investment officer at Parametric Portfolio Associates in Seattle (WA 98109).
**dstein@paraport.com**

cal process control (SPC). In a seminal book, Wald [1947] shows how sequential analysis can greatly speed up the detection of change in a wide variety of systems. His insights have led to the development of many algorithms for change point detection.

We describe one such procedure, known as the CUSUM (which is a contraction of the words *cumulative sum*). We show that with an appropriate choice of parameters, it can detect a transition from outperformance to flat-to-the-benchmark performance in 40 months, and a transition to underperformance faster still. In fact, Moustakides [1986] proves that for any fixed rate of false alarms, the CUSUM is the fastest way there is to detect a change in performance.

If a portfolio's performance is stationary, i.e., if its expected excess return is constant, the power of the CUSUM test is similar to that of Wald's [1947] sequential probability ratio test (SPRT), which, being asymptotically optimal, is more efficient than the widely used t-test for a difference in means. The CUSUM offers the user the best of both worlds. When performance is stationary, it differentiates good performance from bad about as well as the SPRT, and better than the t-test. When performance is time-varying, it detects changes much more quickly than both.

While the CUSUM procedure is easily described, its mathematics are complex. We want here to describe the CUSUM in laypeople's language, and to provide sufficient information to allow its implementation. The mathematically sophisticated reader may want to consult Wald [1947], Moustakides [1986], Basseville and Nikiforov [1993], and Yashchin, Philips, and Stein [1997].

We first specify the measure of performance we think it best to monitor, and then touch briefly on the theoretical underpinnings of the CUSUM. We describe the CUSUM procedure algorithmically, and explain the mathematics underlying each step. Finally, we discuss its applications, along with its strengths and limitations.

## THE INFORMATION RATIO

A useful measure of performance must incorporate both risk and return. While there are many such measures of performance, the one we find most useful is the information ratio, defined as the ratio of a portfolio's excess return relative to an appropriate benchmark to its tracking error, or the standard deviation of its excess return relative to the same benchmark. A good discussion of the properties of the information ratio appears in Grinold and Kahn [1999].

**EXHIBIT 1**

**Probability of Outperformance versus Information Ratio**

| *Information Ratio* | *Probability of Outperforming Benchmark Over Given Horizon* | | |
|---|---|---|---|
| | **1 Year** | **3 Years** | **5 Years** |
| **0.00** | 0.50 | 0.50 | 0.50 |
| **0.25** | 0.60 | 0.67 | 0.71 |
| **0.50** | 0.69 | 0.81 | 0.87 |
| **0.75** | 0.77 | 0.90 | 0.95 |
| **1.00** | 0.84 | 0.96 | 0.99 |

Even though investors typically specify a portfolio's performance in terms of its excess return, there are many advantages to monitoring its information ratio instead. First, a portfolio's excess return is positive if and only if its information ratio is positive. Second, its excess return is the product of its information ratio and its tracking error, readily allowing a translation between one and the other. Third, while the achievable level of excess return varies significantly across asset classes, information ratios are far more stable, and monitoring the information ratio instead of the excess return allows the CUSUM procedure to be used across asset classes without modification. Finally, and most important, the information ratio is directly related to a manager's skill.

With an appropriate benchmark, the sequence of excess returns is uncorrelated and approximately normally distributed, and the probability that a portfolio will outperform its benchmark over any specified horizon is related simply to its information ratio. Exhibit 1 illustrates the relationship, and reveals two surprises. First, even a modest information ratio (0.25) leads to a reasonable probability of outperformance over a five-year horizon. Second, an information ratio of 1.00 virtually guarantees outperformance.

While the CUSUM is readily adapted to monitor other measures of performance such as the Sharpe ratio, the Treynor ratio, or the intercept and slope ($\alpha$ and $\beta$) in a single-factor model, the variance of these parameters tends to be significantly higher than that of the information ratio, slowing the detection of changes in performance.

## THE CUSUM PROCEDURE: THEORY

The CUSUM procedure was first proposed by Page [1954] as a method to rapidly detect changes in the mean of a noisy random process. While the mathematics of Page's solution are complex, we can appreciate his insight every time we compare a chart of monthly and cumulative excess returns for an active manager. The sequence of monthly excess returns is noisy, and it is extremely difficult to determine the average value, or even the sign, of the mean excess return. It is harder still to detect changes in the mean.

The cumulative excess return chart, however, tells a more convincing story, as the slope of the cumulative excess return plot equals the portfolio's average excess return. If the portfolio outperforms its benchmark, the cumulative excess return plot trends upward. Correspondingly, if the portfolio underperforms its benchmark, it trends downward. Finally, if the portfolio's performance is flat to its benchmark, the cumulative excess return plot is trendless.

Changes in slope are easily detected visually, and a quick look at the cumulative excess return plot provides a great deal of information about a portfolio's performance that is not easily divined from its monthly excess returns. The CUSUM procedure is a mathematical implementation of this insight. Formally, the CUSUM procedure is a *backward-looking sequential probability ratio test.* Each time a new return is recorded (typically every month), the model recomputes the optimum interval over which to measure the portfolio's performance relative to its benchmark. It then computes the relative probability that the observed sequence of excess returns over this interval is generated by good and bad investment processes, and raises an alarm when this ratio exceeds a user-defined threshold.

Exhibits 2 through 5 show the monthly excess return, the exponentially weighted tracking error, the cumulative excess return and the cumulative estimated information ratio of a large-value manager relative to the Russell 1000 Value Index. The protractors of annualized excess return (in percent per year) and annualized information ratio displayed on the left of the cumulative sum plots allow one to quickly estimate the performance over any period of interest by comparing the slope over the period to the slopes displayed in the protractors.

The visual power of the cumulative sum plots is striking. At a single glance, one can identify periods of underperformance and outperformance, and see when transitions from one regime to the next occurred.

## THE CUSUM PROCEDURE: DESCRIPTION

Following an initial phase when the CUSUM's current estimate of the information ratio is set to 0 and its estimate of the tracking error is set to its expected value, a three-step procedure is executed every time a new return is recorded. We explain each step after this simple algorithmic description.

1. Estimate the portfolio's current information ratio using the newly recorded return.
2. Compute both the optimum performance measurement interval and the log-likelihood ratio over

## EXHIBIT 2

**Monthly Excess Returns: Large-Value Mutual Fund versus Russell 1000 Value Index**

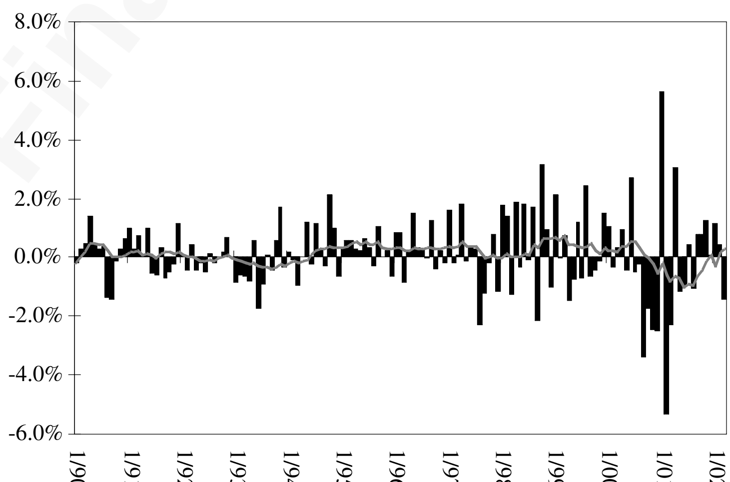

## EXHIBIT 3

**Exponentially Weighted Tracking Error: Large-Value Mutual Fund versus Russell 1000 Value Index**

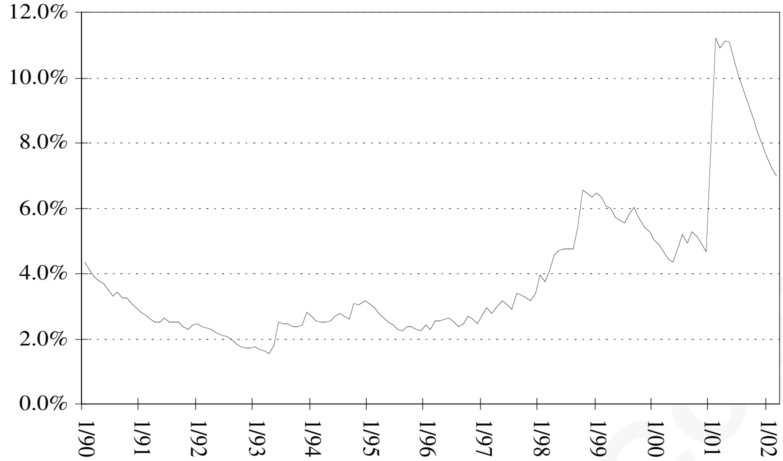


## EXHIBIT 4

**Cumulative Excess Return: Large-Value Mutual Fund versus Russell 1000 Value Index**

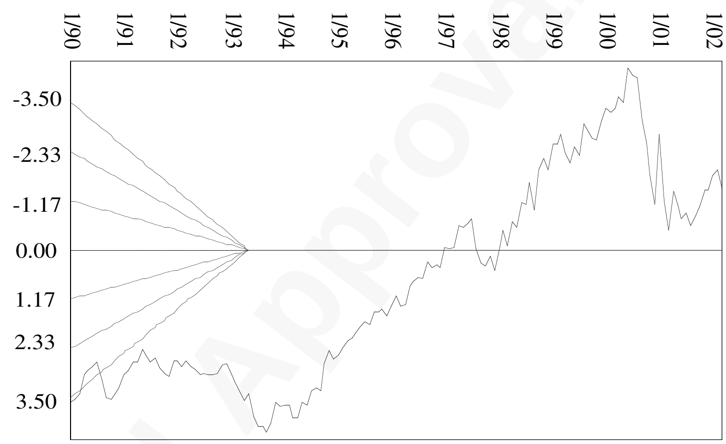


## EXHIBIT 5

**Cumulative Estimated Information Ratio: Large-Value Mutual Fund versus Russell 1000 Value Index**

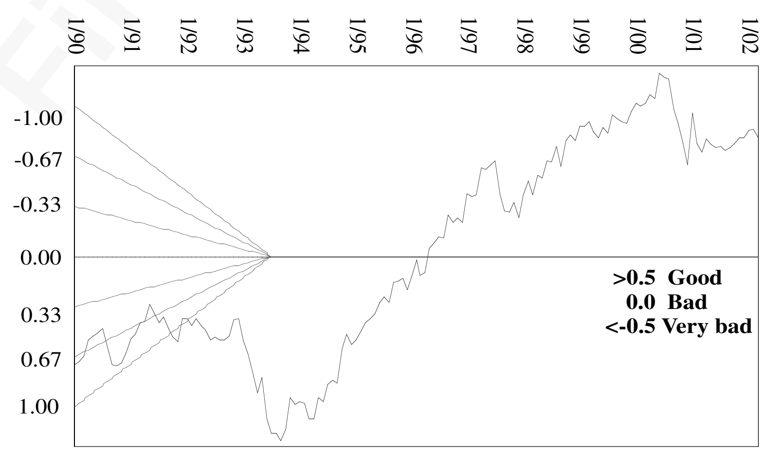

this interval. The log-likelihood ratio is the natural logarithm of the ratio of the probability that the observed sequence of returns was generated by a bad manager to the probability that it was generated by a good manager.

3. If the log-likelihood ratio does not exceed a user-defined threshold, do nothing. If it does, raise an alarm; there is sufficient statistical evidence to suggest that the manager's performance has changed from good to bad. Stop monitoring and conduct a thorough investigation to determine the root cause of the underperformance.

### Step 0: Initialize the CUSUM

Set the initial estimate of the manager's information ratio to 0, and then initialize the tracking error to its expected value. The expected tracking error can be derived from a time series of the portfolio's historical excess returns relative to its benchmark and an examination of the investment process.

### Step 1: Estimate the Current Information Ratio

Define $r_i$ and $b_i$ to be the return of the portfolio and the benchmark, respectively, in month $i$. The logarithmic excess return in month $i$, $e_i$, is defined by

$$e_i = \ln\left(\frac{1 + r_i}{1 + b_i}\right) \tag{1}$$

The use of logarithmic excess returns ensures that the managers' excess returns are correctly compounded. Furthermore, if $e_i$ is normally distributed, it is a maximum-likelihood estimate of the portfolio's current compound excess return.

Next, define $\sigma_i$ and $IR_i$ to be the annualized tracking error and information ratio, respectively, of the portfolio in month $i$. We estimate the tracking error using an exponentially weighted variant of the von Neumann et al. [1941] estimator of variance. This estimator exploits the fact that for two uncorrelated random variables, $x$ and $y$, with the same mean, $12E[(x - y)^2] = \sigma_x^2 + \sigma_y^2$, where $E[.]$ denotes the expectation. It adapts to a slowly changing mean and variance, and is given by

$$\hat{\sigma}_0^2 = \hat{\sigma}_1^2 = \sigma_0^2 \tag{2}$$

$$\hat{\sigma}_i^2 = \gamma\hat{\sigma}_{i-1}^2 + (1 - \gamma)\frac{12 \times \left(e_i - e_{i-1}\right)^2}{2}$$
$$0 < \gamma < 1 \text{ for } i = 2, 3, \ldots, n \tag{3}$$

$$\hat{\sigma}_i = \sqrt{\hat{\sigma}_i^2} \text{ for } i = 0, 1, 2, 3, \ldots, n \tag{4}$$

Setting $\gamma = 1$ is equivalent to assuming that the tracking error equals $\sigma_0$ at all times. If, on the other hand, $\gamma < 1$, the estimator adapts to changes in tracking error. Clearly, there is a trade-off to be made; lowering $\gamma$ shortens the time taken to respond to changes in tracking error, but at the expense of increasing the noise in the estimate. In practice, setting $\gamma = 0.9$ provides a good compromise between detection speed and noise.

This estimator has two significant advantages over the standard sum of squared deviations from the sample mean estimator. First, subtracting adjacent returns effects a cancellation between their means, making the estimator insensitive to slow changes in the process mean. Second, an abrupt change in the mean distorts only two point estimates, and its effect is rapidly extinguished through exponential smoothing.

The current information ratio is estimated as

$$\hat{IR}_i = \frac{12e_i}{\hat{\sigma}_{i-1}} \tag{5}$$

Using $\hat{\sigma}_{i-1}$ instead of $\hat{\sigma}_i$ prevents $e_i$ from appearing simultaneously in both the numerator and the denominator, and ensures that our estimates of the information ratio remain approximately unbiased and uncorrelated.

### Step 2: Determine Optimum Estimation Interval and Compute Log-Likelihood Ratio

Each time a new return is recorded, compute both the optimum performance measurement interval and the log-likelihood ratio over this interval. The log-likelihood ratio is the logarithm of the ratio of the probability that the observed sequence of information ratios was generated by a bad manager to the probability that the observed sequence of information ratios was generated by a good manager. The optimum performance measurement interval is the interval that maximizes the log-likelihood ratio.

Our collective investment experience suggests that *an information ratio of 0.5 or better constitutes good performance.*

As an active portfolio that cannot outperform its benchmark is easily replaced by an index fund, *an information ratio of 0 or worse constitutes bad performance*. Correspondingly, we define a good manager as one whose information ratio is 0.5 or better, and a bad manager as one whose information ratio is 0 or worse.

If we have $N$ observations since the portfolio's inception, the log-likelihood ratio based on the $k$ most recent observations is given by

$$L_N(k) = \log\left[\frac{\text{Probability } [k \text{ most recent observations} \mid \text{Information Ratio} = 0]}{\text{Probability } [k \text{ most recent observations} \mid \text{Information Ratio} = 0.5]}\right] \quad (6)$$

After computing $L_N(1)$, $L_N(2)$, …, $L_N(N)$, the CUSUM determines the optimum performance measurement interval by identifying $k^*$, the value of $k$ for which $L_N(k)$ is maximized, thus maximizing its ability to differentiate good and bad active managers. In practice, it proves convenient to work with just the non-negative part of the log-likelihood ratio, which we define by $L_N = \max[0, L_N(k^*)]$.

If $L_N(k^*) < 0$, the information ratio is more likely to be positive than negative, and no information is discarded by setting $L_N$ to 0. If $L_N(k^*) \geq 0$, however, $L_N = L_N(k^*)$, and $L_N$ parsimoniously encapsulates all the evidence available at time $N$ to reject the hypothesis that the portfolio's performance is satisfactory. More important, if the excess returns of the portfolio are independent normal random variables, $L_N$ can be computed using a simple recursion:

$$\begin{aligned} L_0 &= 0 \\ L_N &= \max\left[0, L_{N-1} + L_N(1)\right] \quad \text{for } N = 1, 2, \ldots \\ &= \max\left[0, L_{N-1} - \hat{IR}_N + 0.25\right] \end{aligned} \quad (7)$$

This conscious maximization of the log-likelihood ratio over the optimal estimation interval differentiates the CUSUM from all other performance measurement methods. It is the root source of both its ability to detect a change in performance in the shortest possible time and its robustness to the distribution of portfolio returns. It can be shown that the distribution of $L_N$ is only weakly dependent on the distribution of the sequence of excess returns, allowing the user of the CUSUM to monitor active portfolios in almost any publicly traded asset class without modification.

### Step 3: Compare Log-Likelihood Ratio to Threshold and Raise an Alarm if Necessary

If $L_N$ is positive, we ask *how likely it is that the portfolio's underperformance is caused by a decline in the ability of the investment process to add value*. Clearly, if the log-likelihood ratio is low, the underperformance is not statistically significant; all active managers are afflicted by an occasional bout of underperformance. If, on the other hand, it is high, it is likely that the underperformance is caused by a decline in the ability of the investment process to add value. The ability to discriminate between a fundamentally flawed investment process and random noise in a viable investment process is an increasing function of the log-likelihood ratio.

The CUSUM procedure explicitly recognizes this trade-off between detection speed and the ability to discriminate between good and bad managers. It requires us to specify a threshold for the log-likelihood ratio below which no decision is made about the state of the investment process. When this threshold is crossed, sufficient evidence has accrued to conclude that its ability to add value has declined.

The threshold determines both the average time taken to detect underperformance by a bad manager and the rate of false alarms (incorrect identifications of a good manager as bad). If the threshold is set low, underperformance is rapidly detected, but we experience a correspondingly high rate of false alarms. If the threshold is set high, it takes longer to detect underperformance, but the rate of false alarms drops.

Every statistical decision procedure must make this trade-off between detection speed and the rate of false alarms. Moustakides [1986] proves that under some relatively general conditions, the trade-off made by the CUSUM procedure is optimal. For *any* distribution of returns that is drawn from the exponential family, and for *any* fixed rate of false alarms, no other procedure can detect underperformance faster. His article represents the culmination of a 25-year effort to prove this remarkable result.*

In practice, a threshold setting that detects flat-to-the-benchmark performance in three and a half years and allows one false alarm in seven years proves satisfactory.

## EXHIBIT 6
**Detection Speed versus Threshold**

| *Threshold for $L_N$* | *Expected Time To Cross Threshold (months)* | | |
|---|---|---|---|
| | IR = 0.5: $T_0$ | IR = 0: $T_1$ | IR = –0.5: $T_2$ |
| 11.81 | 24 | 16 | 11 |
| 15.00 | 36 | 22 | 15 |
| 17.60 | 48 | 27 | 18 |
| 19.81 | 60 | 32 | 21 |
| 21.79 | 72 | 37 | 23 |
| 23.59 | 84 | 41 | 25 |

This setting was determined by examining the actual investment results of a large number of domestic and international equity and fixed-income managers. It affords a tenfold improvement in detection speed over the traditional t-test for a difference in means while still maintaining a reasonably low rate of false alarms.

Exhibit 6 shows the thresholds required for various expected times between false alarms. These thresholds are computed using the approximate computational approach described in Yashchin, Philips, and Stein [1997] and verified using a simulator. Appropriate thresholds can be computed for any desired level of good and bad performance—see Woodall [1983], Vance [1986], and Yashchin, Philips, and Stein [1997] for a detailed description of the numerical procedure.

### THE CUSUM PROCEDURE: IMPLEMENTATION

The CUSUM procedure is implemented using the following sequence of steps:

1. Initialize the tracking error $\sigma_0$ to its expected value and set $L_0 = 0$.
2. Each time a new return is recorded:
   a. Compute the logarithmic excess return using Equation (1).
   b. Update the estimate of tracking error using Equations (2), (3), and (4).
   c. Estimate the current information ratio using Equation (5).
   d. Update the likelihood ratio $L_N$ using Equation (7).
3. Compare the updated likelihood ratio to the thresholds in Exhibit 6.
   a. If it does not exceed the user-chosen threshold, do nothing.
   b. If it does exceed the user-chosen threshold, raise an alarm and investigate the manager's process to determine the cause of underperformance:
      - If the investment process is found satisfactory, consider this to have been a false alarm. Go back to Step 1, reset the CUSUM, and continue monitoring the portfolio.
      - If the investment process is found deficient, stop the CUSUM, and decide on a course of action.

Exhibit 7 displays the log-likelihood ratio plot (also known as Page's plot in honor of E.S. Page) for the manager whose performance is shown in Exhibits 2-5. When the portfolio underperforms its benchmark, the log-likelihood ratio increases (i.e., moves downward). When the portfolio outperforms its benchmark, it declines (i.e., moves upward). We have inverted the direction of the y-axis so that it corresponds to our established notion that good performance points upward while poor performance points down.

The plot has six labeled horizontal levels. Each corresponds to a row in Exhibit 6, and is labeled with three numbers. The first number is the expected time (in months) to cross this level if the manager is good, i.e., if the information ratio is 0.5. The second number is the expected time to cross this level if the manager is bad, or if the information ratio is 0. The third number is the expected time to cross this level if the manager is exceptionally bad, or if the information ratio is –0.5. On average, a good manager's log-likelihood ratio will cross the first level once in 24 months, while a bad manager's log-likelihood ratio will cross it once in 16 months, and an exceptionally bad manager's log-likelihood ratio will cross it once in 11 months.

Notice that the average time to cross a level increases monotonically as we go down levels, and declines monotonically from left to right at any given level. This is to be expected. The more confident we wish to be of our ability to discriminate between good and bad managers, the longer we must wait. The worse a portfolio's information ratio, the sooner its log-likelihood ratio will cross any given level.

When the log-likelihood ratio crosses the lowest level, an alarm is raised. At this point, sufficient evidence has accrued to warrant an investigation of the manager's process. It takes only 41 months to detect flat-to-the-benchmark performance. Once in 84 months, a good manager's log-likelihood ratio will cross this level by chance, and a false alarm will be raised. Following the

## EXHIBIT 7

**Page's Plot of Log-Likelihood Ratios: Large-Value Mutual Fund versus Russell 1000 Value Index**

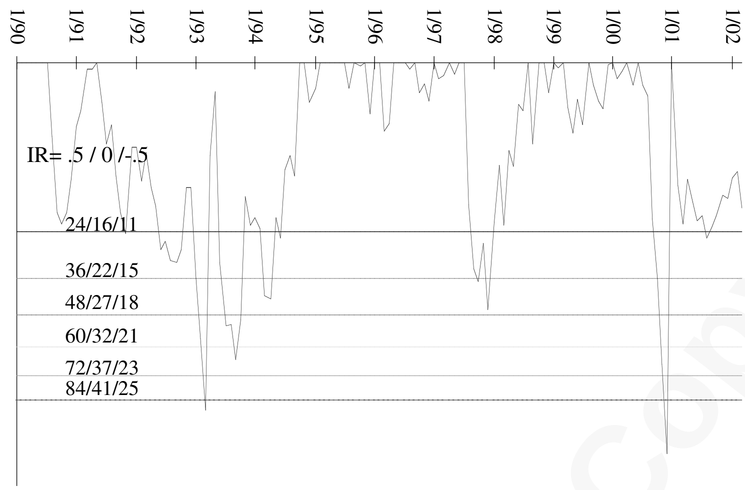


alarm, an investigation must be launched. If, after the investigation, it is concluded that the manager's investment process is satisfactory and that the alarm is false, the CUSUM is reset, and the monitoring process is restarted.

In our example, two alarms were raised: one in March 1993, and the other in December 2000. The first alarm was, with the benefit of hindsight, a false alarm—the manager's performance recovered after the alarm was raised. The second one was valid—the manager's performance has become far more volatile, and in spite of some sharp reversals, its performance continues to deteriorate.

This raises a very important question. What ought one to do when an alarm is raised? The CUSUM procedure detects underperformance, but offers no causal explanation for it. It should not, therefore, be used as a tool to engage and terminate asset managers solely upon the basis of recent performance. It is incumbent upon, and indeed imperative for, the user to launch a thorough investigation into the manager's investment process and to decide if it is likely to result in satisfactory performance going forward.

If the user concludes that the investment process is likely to generate satisfactory returns in the future, the alarm is likely to have been false, and the manager ought to be retained and the CUSUM reinitialized. If, on the other hand, the user feels that the process is unlikely to perform well in the future, a number of courses of action, up to and including termination, are available, and an appropriate decision must be made. In the example presented, an examination of the portfolio and a discussion with the portfolio manager resulted in an understanding of why its risk had increased so sharply in recent years.

Like any performance measurement technique, the CUSUM has both strengths and limitations, and users must be aware of both to maximize its utility. Its strengths include:

1. A transition from outperformance to underperformance can be detected much sooner than with traditional methods.
2. The CUSUM equations are simple, and the algorithm is easily implemented in a spreadsheet, a statistical package such as S-Plus, SAS, or Matlab, or a specialized process control package such as SPC/PI+ or SPC-PC.
3. The CUSUM is robust to the distribution of excess returns. It can therefore be used to monitor portfolios in all publicly traded asset classes without changing the threshold at which an alarm is raised.
4. It adapts to changes in tracking error, regardless of their cause.
5. Investigations launched after an alarm is raised almost always result in the identification of the underlying cause of the performance problem.

Its primary limitation is its sensitivity to correlation in the sequence of excess returns, as Equation (7) is valid only when returns are independent. In practice, however, serial correlation does not pose a major problem for two reasons:

1. The choice of an appropriate benchmark will ensure that the sequence of excess returns is essentially

uncorrelated. This is easily accomplished in practice; large-value managers should be benchmarked against large-value indexes, and fixed-income managers should be benchmarked against bond indexes of similar duration and credit quality

2. Yashchin [1993] shows that if the correlation in the series of estimated information ratios is reasonably low (i.e., if $|\rho| \leq 0.5$), simply modifying the threshold at which an alarm is raised is almost as effective as an exact computation of the log-likelihood ratio. Positive correlation requires the threshold to be raised, while negative correlation requires it to be lowered.

A decade of experience with a wide range of domestic and international equity and fixed-income managers has shown that if a portfolio is correctly benchmarked, its excess returns and estimated information ratios are almost always nearly uncorrelated. With this elementary precaution (which should be an integral part of any manager evaluation), even unsophisticated users can successfully use the CUSUM, for prior to an alarm being raised, there is no need for intervention by a skilled person.

## SUMMARY

The CUSUM procedure is an exceptionally powerful tool for monitoring the performance of actively managed portfolios. It detects flat-to-the-benchmark performance in 40 months, is robust to the distribution of excess returns, and works well with equity, fixed-income, and currency portfolios, both domestic and international, as well as hedge funds. It is currently being used by plan sponsors, consultants, and money managers on three continents to monitor over $500 billion in actively managed assets, and their experience with it has been uniformly positive.

Like all statistical methods that use only returns, the CUSUM does not provide a causal explanation for a manager's performance. It is therefore imperative to view an alarm as a call for an investigation into the root cause of a manager's underperformance, and not as an automatic signal to terminate the manager.

## ENDNOTES

The authors thank Richard Michaud for his many thoughtful comments, which led to substantial improvements in our article.

*Formally, the exponential family of distributions consists of all distributions whose density function can be written in the form $f(x,\theta) = h(x)g(\theta)\exp\left(\sum w_i(\theta)t_i(x)\right)$, where $h(x)$ and $t_i(x)$ are functions only of $x$, while $g(\theta)$ and $w_i(\theta)$ are functions only of $\theta$. It is a large family, and includes most distributions that are used to model portfolio returns, including the normal, lognormal, and gamma distributions.

## REFERENCES


Basseville, Michele, and Igor Nikiforov. *Detection of Abrupt Changes—Theory and Application*. Englewood Cliffs, NJ: Prentice-Hall, 1993.

Grinold, Richard, and Ronald Kahn. *Active Portfolio Management,* 2nd ed. New York: McGraw-Hill, 1999.

Moustakides, George. "Optimal Stopping Times for Detecting Changes in Distributions." *Annals of Statistics*, v. 14, no. 2 (1986), pp. 1379-1387.

Page, E.S. "Continuous Inspection Schemes." *Biometrika*, v. 41, no. 1 (1954), pp. 100-115.

Vance, Lonnie C. "Average Run Lengths of Cumulative Sum Control Charts for Controlling Normal Means." *Journal of Quality Technology,* v. 18, no. 3 (1986), pp. 189-193.

von Neumann, John, R.H. Kent, H.R. Bellinson, and B.I. Hart. "The Mean Square Successive Difference." *Annals of Mathematical Statistics*, v. 12, no. 2 (June 1941), pp. 153-162.

Wald, Abraham. *Sequential Analysis*. New York: Wiley, 1947.

Woodall, William. "The Distribution of the Run Length of One-Sided CUSUM Procedures for Continuous Random Variables." *Technometrics*, v. 25, no. 3 (1983), pp. 295-301.

Yashchin, Emmanuel. "Performance of CUSUM Control Schemes for Serially Correlated Observations." *Technometrics*, v. 35, no. 1 (February 1993), pp. 37-51.

Yashchin, Emmanuel, Thomas K. Philips, and David M. Stein. "Monitoring Active Portfolios Using Statistical Process Control." In Hans Amman, Berc Rustem, and Andrew Whinston, eds., *Computational Approaches to Economic Problems*. Dordrecht, Netherlands: Kluwer Academic Publishers, 1997.