FAJ | Financial Analysts Journal
Volume 69 • Number 2
©2013 CFA Institute

The Arithmetic of Investment Expenses

William F. Sharpe

Recent regulatory changes have brought a renewed focus on the impact of investment expenses on investors' financial well-being. The author offers methods for calculating relative terminal wealth levels for those investing in funds with different expense ratios. Under plausible conditions, a person saving for retirement who chooses low-cost investments could have a standard of living throughout retirement more than 20% higher than that of a comparable investor in high-cost investments.

A recent issue of this publication featured an editorial by Charles Ellis (2012), titled "Investment Management Fees Are (Much) Higher Than You Think," in which Ellis argued that as a percentage of *assets*, such fees do look low, but

> calculated correctly, as a percentage of returns, fees no longer look low. . . . Investors should consider fees charged by active managers not as a percentage of total returns but as *incremental* fees versus risk-adjusted *incremental* returns above the market index. (p. 4)

Why? Because

> extensive, undeniable data show that identifying in advance any one particular investment manager who will—after costs, taxes, and fees—achieve the holy grail of beating the market is highly improbable. (p. 6)

The latter point is consistent with decades of research in both academe and the investment industry. As the director of fund research at Morningstar, a leading provider of mutual fund data and analysis, noted,

> If there's anything in the whole world of mutual funds that you can take to the bank, it's that expense ratios help you make a better decision. In every single time period and data point tested, low-cost funds beat high-cost funds. . . . Investors should make expense ratios a primary test in fund selection. They are still the most dependable predictor of performance. (Kinnel 2010, pp. 2–3)

William F. Sharpe is professor emeritus of finance at Stanford University and director emeritus at Financial Engines, Inc., Sunnyvale, California.

Investment expenses have long been reported by mutual funds, although their impact is undoubtedly ignored or misunderstood by many investors. In the United States, those saving for retirement with 401(k) defined contribution plans often have insufficient information about such expenses. However, rules recently issued by the U.S. Department of Labor require that such 401(k) investors receive reports detailing many of these charges. Given this renewed focus, it is more important than ever for investors to understand the possible impact of investment expenses on their future wealth.

In my research, I concentrated on comparisons of two possible investments—one with high expenses and the other with low expenses.[1] I considered the impact of investment expenses in two settings: when a single investment is held for a number of years and when recurring investments are made over many years. For each setting, I looked at cases in which both the low-cost and the high-cost investments provided the same gross returns (before expenses) and cases in which such returns might differ.

Lump-Sum Investments with Equal Gross Returns

The Vanguard Group provides investors with computational tools to assess the possible impact of mutual fund expenses in cases in which a lump sum is invested for a number of years with no withdrawals or contributions. **Figure 1** depicts an example of the Vanguard Total Stock Market Index Fund Admiral Shares (with a minimum investment of $10,000), designed to track the broad U.S. stock market. As Figure 1 shows, the Vanguard fund's expense ratio is 0.06% a year, whereas the average expense ratio of similar funds is assumed to be 1.12% a year. Note that the key outputs are the dollar fees paid over the holding period. These results depend on the main assumptions (the expense ratios and the number of years the investment is held) as well as the

Figure 1. Vanguard Estimates of the Effects of Different Expense Ratios

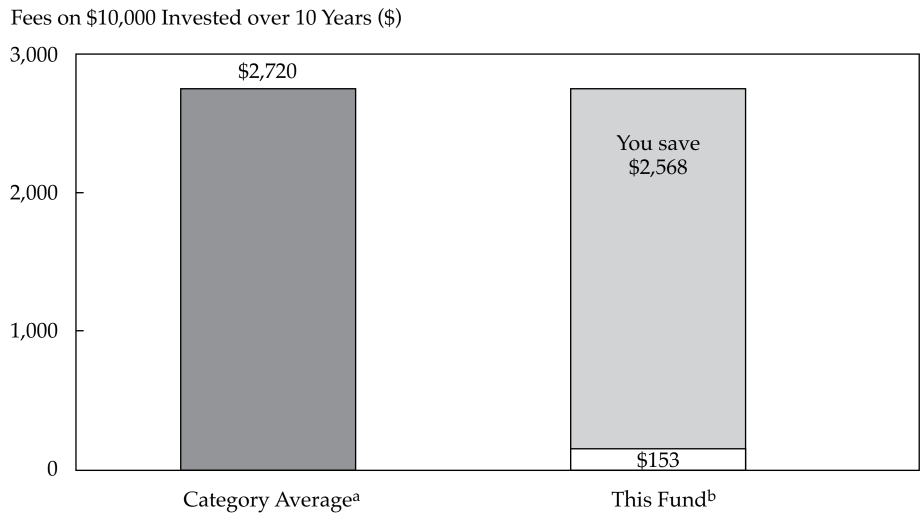

[a]Average expense ratio of similar funds = 1.12%.
[b]Expense ratio = 0.06%.
Source: Vanguard Group (2012).

initial amount to be invested and an assumed constant gross return on the investment to be earned by each fund.

I propose a different measure for the effects of expenses—one that is both simple and likely to be more meaningful for many investors. For reasons that will become clear, I call it the *terminal wealth ratio* (TWR). For cases such as those considered by Vanguard in which the low-cost and high-cost funds are assumed to have equal cumulative gross returns, no assumptions need be made about the levels or patterns of such returns over time.

Let us consider an investment that returns $\tilde{r}_i$ in year i. Returns are measured in proportions; thus, for a return of 8%, $\tilde{r}_i = 0.08$. The tilde over the variable indicates that the value may not be known in advance. A dollar invested at the beginning of year i will grow to $1+\tilde{r}_i$ by the end of the year. Assume that expenses are then paid using an expense ratio of x (e.g., if the expense ratio is 1.12%, $x = 0.0112$). As a result, a dollar invested at the beginning of the year will grow (net of expenses) to

$$(1-x)(1+\tilde{r}_i).$$

Now let us consider an investment held for n years. The terminal value per dollar of initial investment is

$$\left[(1-x)(1+\tilde{r}_1)\right]\left[(1-x)(1+\tilde{r}_2)\right]\ldots\left[(1-x)(1+\tilde{r}_n)\right].$$

After rearranging, the terminal value is

$$(1-x)^n\left[(1+\tilde{r}_1)(1+\tilde{r}_2)\ldots(1+\tilde{r}_n)\right].$$

The expression in brackets is the terminal value per dollar invested that would have been obtained had there been no expenses—that is, the *compounded gross return*, or G_n. The initial term in parentheses $(1-x)$ is the proportion of the value retained each period (if $x = 0.0112$, the proportion retained is 0.9888). Thus, the terminal value per dollar of investment is

$$(1-x)^n G_n.$$

Now let us consider a comparison between two funds with potentially different compounded gross returns and different expense ratios—x_1 and x_2. The final values are

$$(1-x_1)^n G_{n1}$$

and

$$(1-x_2)^n G_{n2}.$$

The ratio of the two final values is the TWR:

$$\text{TWR} = \frac{(1-x_1)^n G_{n1}}{(1-x_2)^n G_{n2}}.$$

Rearranging slightly gives

$$\text{TWR} = \left[\frac{(1-x_1)}{(1-x_2)}\right]^n \frac{G_{n1}}{G_{n2}}. \quad (1)$$

In its calculations, Vanguard assumes that the two investments have equal constant gross returns in each year of the holding period—likely to be a reasonable assumption for expected values, although not necessarily for actual values, a possibility that we will examine later. In any event, when the compounded gross returns are the same, the second ratio equals 1. In such a case, the ratio of the terminal values does not depend in any way on the actual returns during the holding period. Only the number of years the investment is held and the two expense ratios matter. Equivalently, the key ingredients are the *retention ratio* (the bracketed expression in Equation 1) and the number of years the investment is held (the exponent n).

More generally, when the two investments have the same compounded gross returns, the TWR equals the compounded retention ratio:

$$\text{TWR} = \left[\frac{(1-x_1)}{(1-x_2)}\right]^n \tag{2}$$

if $G_{n1} = G_{n2}$.

In the Vanguard example shown in Figure 1, $x_1 = 0.0006$, $x_2 = 0.0112$, and $n = 10$. In such a case,

$$\text{TWR} = \left[\frac{(1-0.0006)}{(1-0.0112)}\right]^{10} = 1.1125.$$

Breaking the computation into its components, we see that the retention ratio is 0.9994/0.9888, or approximately 1.0107 (an advantage of more than 1% a year for the lower-cost fund). This ratio, compounded over 10 years, is approximately 1.1125—the TWR for the two investments. Thus, an investor who obtains a given terminal gross return over 10 years with an expense ratio of 0.06% a year ends up with 11.25% more wealth than one who obtains the same terminal gross return but must pay 1.12% a year in expenses. To paraphrase Ellis (2012), the difference between an expense ratio of 0.06% a year and an expense ratio of 1.12% a year may seem small, but the cumulative result of the lower expense ratio after 10 years is an increase in wealth of more than 11%. And this outcome is the case no matter what the actual investment returns are from year to year so long as the investments with the differing expense ratios have the same compounded gross returns.

Figure 2 shows the TWRs of lump-sum investments held for 10, 20, and 30 years with retention ratios of 1.000–1.020. Not surprisingly, the effects of a given retention ratio are greater if assets are invested for longer periods. For two funds with the expenses in our example, the TWR for a holding period of 30 years is slightly less than 1.38. An increase in terminal wealth of almost 38% is huge by any standard.

Real Returns

Most investors are interested in the purchasing power of their terminal wealth, not its nominal value. Thus, returns should be measured in real (inflation-adjusted), not nominal, terms. Fortunately, so long as such returns are computed in the correct manner, Equations 1 and 2 continue to apply. If we let $\tilde{r}_i$ be the nominal return on an investment in period i and $\tilde{c}_i$ be the proportional change in a measure of the cost of living during the period, then the real return is computed as follows:

$$\left(1+\tilde{r}_i'\right) = \frac{\left(1+\tilde{r}_i\right)}{\left(1+\tilde{c}_i\right)}.$$

Dividing each value of $1+\tilde{r}_i$ in each of the equations used to derive Equation 1 by $1+\tilde{c}_i$ and simplifying as before again give Equation 1, with the compounded gross returns measured in real terms. Thus, the previously discussed arguments suggest that Equation 1 holds whether the two investments have the same gross compounded *real* returns or the same gross compounded *nominal* returns—as is also true for Equation 2. More simply, if one alternative is worth more than another by a given percentage in nominal terms, it will exceed the latter by the same percentage in real terms.

Recurring Investments with Equal Constant Annual Returns

Let us now turn to cases in which recurring investments are made over N years. Unfortunately, the simple formula in Equation 1 is not directly applicable, and so calculations must be made on the basis of assumptions about investment returns. We can follow the Vanguard example by assuming the same constant rate of return from year to year for both the low-cost fund and the high-cost fund.

Let us concentrate on cases in which each year for N years, an investor devotes a particular portion of income to investment in a retirement plan, with the goal of using the total value of the fund N years hence to finance spending during the retirement years. To simplify, let us assume that the annual contributions are equal in real terms, although the calculations can be easily adjusted to incorporate changes in the real values saved each year, monthly investments, and other possible aspects of a retirement savings plan.

In this case, assuming that equal real investments are made for N years,

$$\text{TWR} = \frac{(1-x_1)^N G_N + (1-x_1)^{N-1} G_{N-1} + \ldots + (1-x_1)^1 G_1}{(1-x_2)^N G_N + (1-x_2)^{N-1} G_{N-1} + \ldots + (1-x_2)^1 G_1}.$$

Figure 2. Terminal Wealth Ratios for Lump-Sum Investments: Alternative Retention Ratios for Investments over 10, 20, and 30 Years

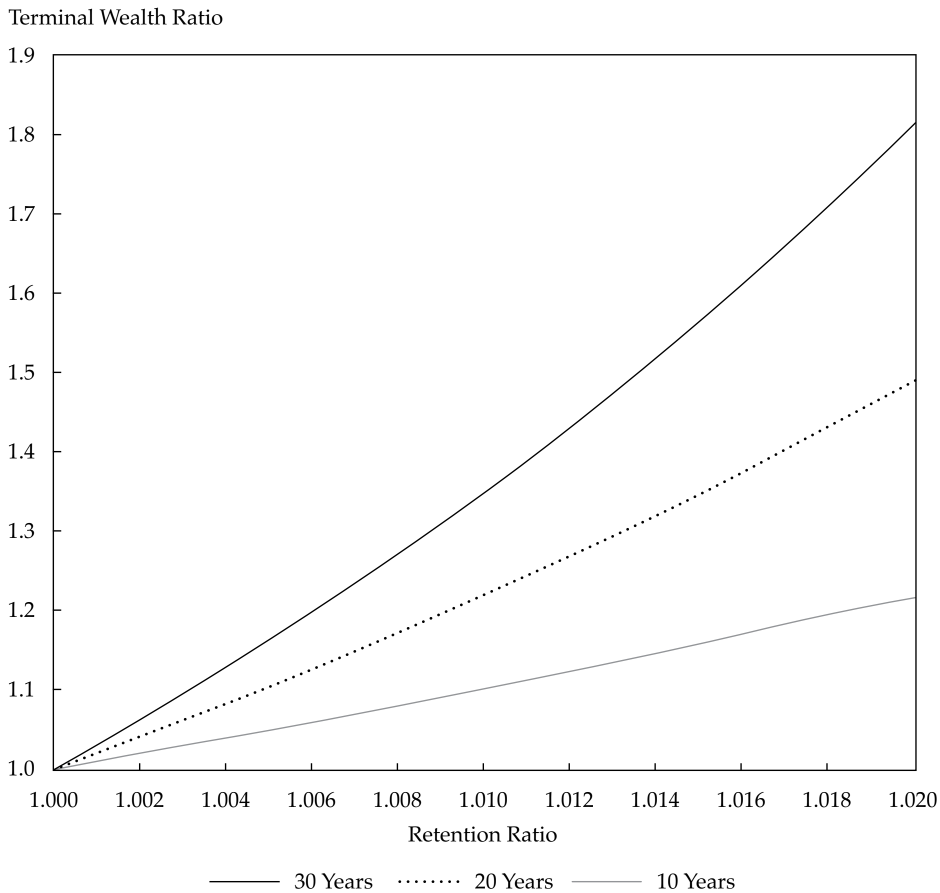

The relative wealth obtained with one expense ratio vis-à-vis that obtained with a different expense ratio depends on the returns provided by the underlying investments. To assess the outcomes, assumptions must be made about likely future investment returns.

Figure 3 shows terminal real wealth ratios for cases in which equal real amounts are invested each year for *N* years while the underlying gross investment real return remains constant. The horizontal axis shows the alternative annual real rates of return; the vertical axis shows the terminal real wealth ratios. Curves are shown for 10-, 20-, and 30-year investment periods. Not surprisingly, the TWRs are higher the longer the period over which investments are made because the effects of higher expense ratios compound from year to year. Moreover, for a given number of years, the higher the real rate of return on the underlying assets, the higher the TWR. Why? Because the higher the return, the greater the terminal values of the earlier contributions relative to those of the later contributions—and the former provide higher TWRs than the latter.

Recall that for a lump-sum investment held for 30 years, the effect of investing with an expense ratio of 0.06% rather than 1.12% is very large, with a TWR of approximately 1.38%. When money is invested in equal annual real amounts over 30 years, the effects are smaller because not all investments experience the fees for the full period. As Figure 3 shows, for 30-year investment periods, TWRs range from approximately 1.175 to 1.26, depending on the gross return on the underlying assets. For plausible assumptions about the return on a diversified investment portfolio, the TWR is significantly greater than 1.20. Almost certainly, most investors would consider it extremely desirable to be able to look forward to having the funds saved for their retirement provide 20% more purchasing power.

Monte Carlo Analysis

Let us now turn to more realistic cases, in which future gross returns can differ between the two funds, leading to uncertainty about the terminal

Figure 3. Terminal Wealth Ratios with Constant Real Rates of Return: Equal Annual Real Investments over Alternative Investment Periods for Funds with Expense Ratios of 0.06% and 1.12%

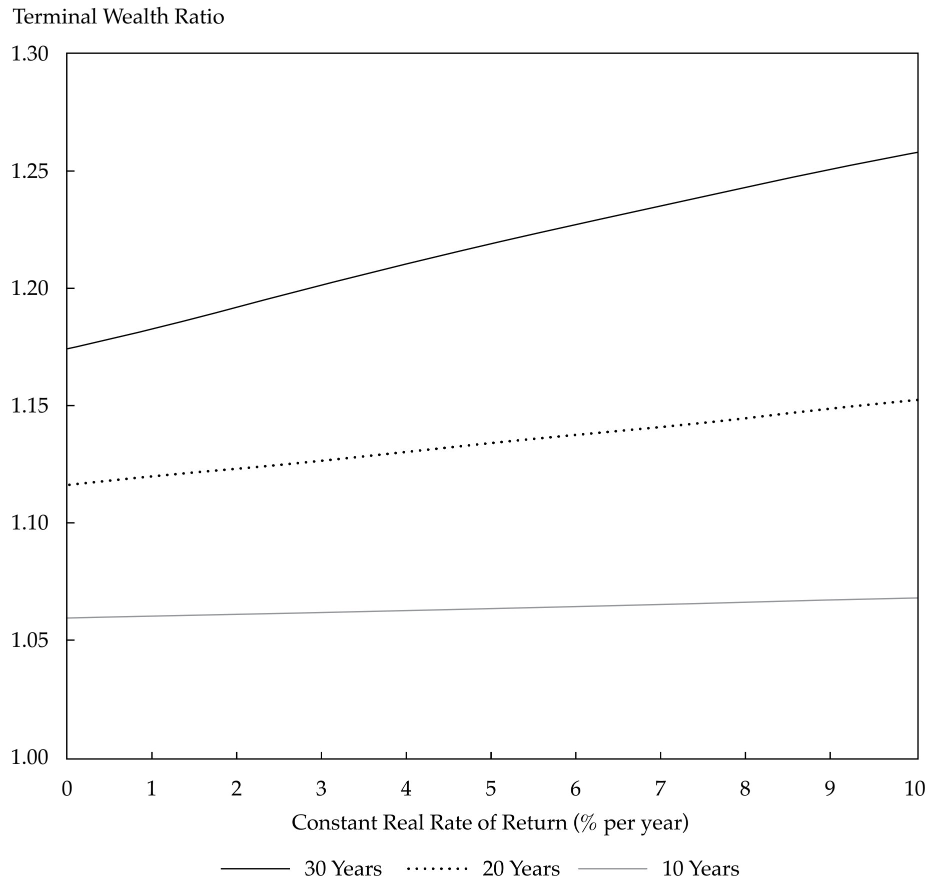

wealth ratio. In this case, I used Monte Carlo analysis to generate a million possible 30-year scenarios, computed the TWR for each one, and then calculated the range of the ratios across scenarios.

Although it is convenient to assume that funds with different expense ratios but comparable investment styles (and/or benchmarks) will have the same compounded gross returns, this scenario is unlikely to be the case in practice. Instead, the fund with higher expenses is likely to engage in a more active management style than the fund with lower expenses. A useful way to represent this difference is to assume that in each year, the higher-cost fund (*h*) will provide a gross return equal to that of the lower-cost fund (*l*) plus a tracking error, which will vary from year to year around a zero value. In such a situation, the TWR of the two funds will be uncertain beforehand, with the actual result dependent on the realized differences between their gross returns.

We can represent the relationship as follows:

$$\tilde{r}_h = \tilde{r}_l + \tilde{\varepsilon}. \qquad (3)$$

Equation 3 differs from the common formulation, in which an "alpha" term on the right-hand side represents the expected additional return from active management. In effect, we are assuming that there is no such added return and thus omit the alpha term.

In the Monte Carlo simulations, for each year and scenario, a return on the low-cost fund $(\tilde{r}_l)$ is drawn randomly from a normal distribution. This distribution is the same for every year and scenario; thus, *ex ante*, these returns are uncorrelated from year to year. I based this return distribution on the historical performance of an index of global real stock returns (Dimson, Marsh, Staunton, McGinnie, and Wilmot 2012, p. 14, Table 3). From 1900 through 2011, these stocks provided an annual average real return of 6.9%, with a standard deviation of 17.7% a year. After obtaining each return for the low-cost fund, I then drew a value for the tracking error (ε) from a normal distribution with an expected value of zero and a specified standard deviation. These tracking errors are also uncorrelated *ex ante* from year to year. The return on the high-cost fund is the sum of the return on the low-cost fund and the tracking error, as in Equation 3.

Lump-Sum Investments with Uncertain Returns

The first set of simulations assumes that a lump sum is invested in each of the two alternative funds for 30 years. As before, let us assume that the low-cost fund has an expense ratio of 0.06% a year and the high-cost fund has an expense ratio of 1.12% a year.

Figure 4 shows the results for three possible levels of active risk, with tracking error standard deviations of 0%, 2.5%, and 5.0% a year. The results for zero tracking error are all the same because the conditions required for Equation 2 are met, giving a TWR of approximately 1.38% in every scenario. However, each of the two cases with tracking error risk has a range of possible outcomes. The probability that the TWR will exceed 1.0 is more than 99% for the case with moderate active risk (0.025) andw more than 90% for the case with the greatest active risk (0.050). In each case, there is a 50% chance that the TWR will be roughly 1.38% or greater.[2]

Note that with moderate amounts of active risk (tracking error standard deviations of less than 2.5% a year), an investor in the low-cost fund will almost certainly be richer than an investor in the high-cost fund. Moreover, the disparity is likely to be very large. But there is a small chance that after 30 years, the two will have relatively similar amounts of wealth. And when the active risk is relatively large (tracking error standard deviations of at least 5.0%), there is a chance (albeit a small one) that even after 30 years, an investor in the low-cost fund will be poorer than an investor who chose the high-cost alternative.

Although betting on a relatively active manager with no ability to add value, on average, is a poor choice, the simulations show why a Darwinian process does not weed out such managers with great rapidity. In this case, the odds are even that an investor in the low-cost fund will be

Figure 4. Probabilities of Alternative Terminal Wealth Ratios for Tracking Errors with Annual Standard Deviations of 0, 0.025, and 0.050: Lump-Sum Investments for 30 Years with Expense Ratios of 0.06% and 1.12%

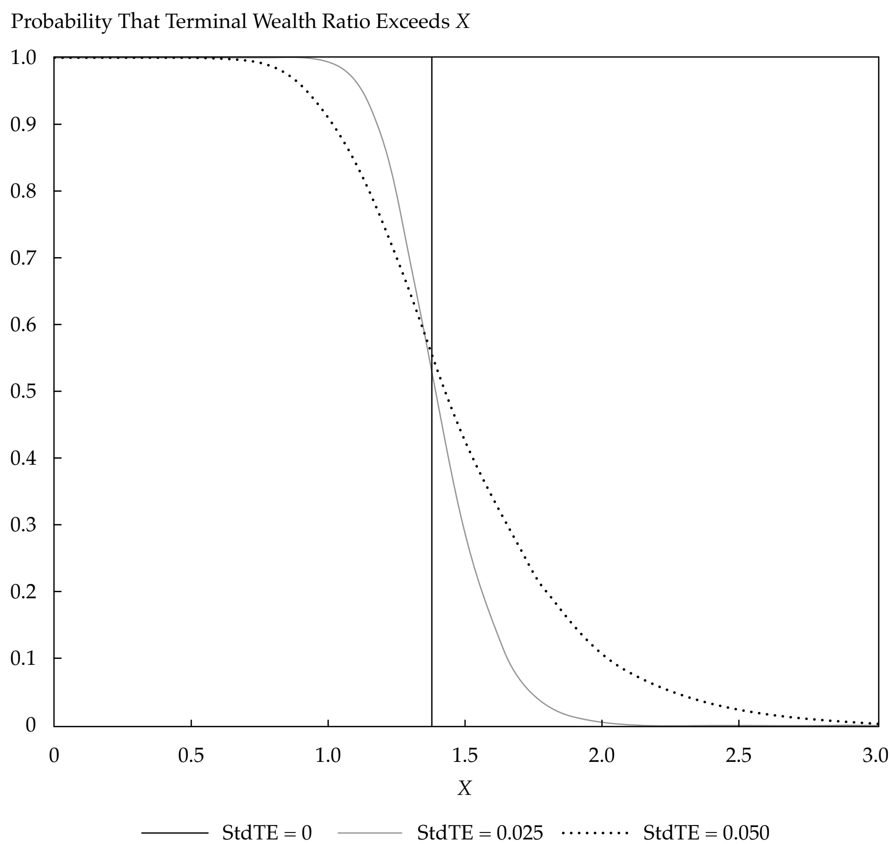

well over a third richer than an investor in the high-cost fund after 30 years. But there is a small chance that an investor in the low-cost fund will regret not having selected the high-cost fund. For those who choose funds with high expense ratios, hope may spring eternal.

Recurring Investments with Uncertain Returns

Finally, we can analyze the ranges of TWRs obtained with equal real investments in each year for a given number of years. **Figure 5** shows the results for investments made over 30 years. It differs from Figure 4 in two respects. First, there is uncertainty about the TWR even when the high-cost fund has zero tracking error. This result is due to the variation in the relative weights of the TWRs of the annual investments, which arises from the variation in the returns of the two funds from year to year. Second, the ranges of TWRs are narrower than in the earlier case because the holding periods for all but one of the amounts invested are less than 30 years.

Despite these differences, there is again a small but significant chance (10%) that after 30 years, a low-cost fund will provide an investor with less wealth than a high-cost fund with a similar investment style but substantial tracking error (5%), even though both funds have the same *ex ante* expected gross returns. That said, the odds are even that the frugal investor will have over 20% more money to spend during retirement.

Conclusion

Many of the numeric results in this article depend on the particular expense ratios used (0.06% a year for the low-cost fund and 1.12% a year for the high-cost fund). But the formulas and Monte Carlo procedures can be applied in other cases. My main goals here were to advocate the use of the terminal wealth ratio—a simple yet meaningful measure of the relative outcomes provided by funds with

Figure 5. Probabilities of Alternative Terminal Wealth Ratios for Tracking Errors with Annual Standard Deviations of 0, 0.025, and 0.050: Equal Annual Real Investments for 30 Years with Expense Ratios of 0.06% and 1.12%

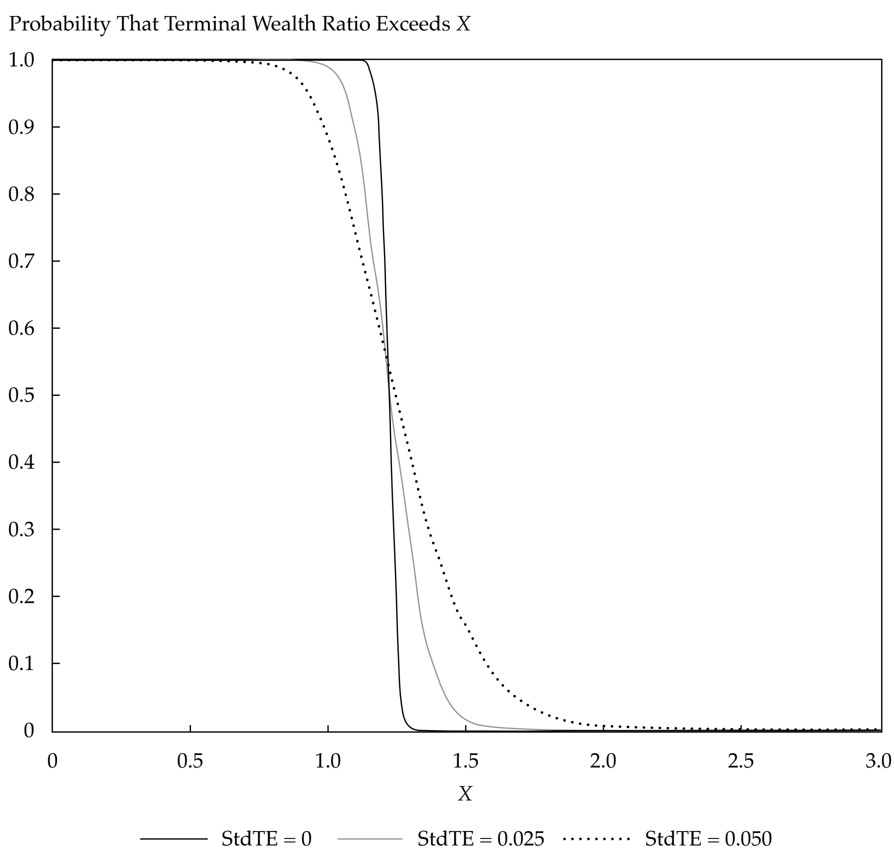

different expense ratios—and to show how to calculate possible ranges of this measure for various types of investments over time.

That said, the results I obtained for the expense ratios considered are dramatic. Whether one is investing a lump-sum amount or a series of periodic amounts, the arithmetic of investment expenses is compelling. Although a long-term investor may be able to find one or more high-cost managers who can beat an appropriate benchmark by an amount sufficient to more than offset the added costs, the reality is that "compared with the readily available passive alternative, fees for active management are astonishingly high" (Ellis 2012, p. 4). Managers with extraordinary skills may exist, but as I argued in this publication many years ago (Sharpe 1991), another exercise in arithmetic indicates that such managers are in the minority. And as Ellis has reminded us, they are very hard indeed to identify in advance. *Caveat emptor.*

I thank Robert Young, John Watson, and Jason Scott of Financial Engines and Steven Grenadier of Stanford University for helpful comments.

This article qualifies for 1 CE credit.

Notes

1. The formulas and simulation procedures also work if one of the investments has zero expenses.
2. Given our (traditional) assumptions about tracking error, the median compounded gross return for the high-cost active strategy is smaller than that for the low-cost passive strategy even though the expected returns for both strategies are the same.

References

Dimson, E., P. Marsh, M. Staunton, P. McGinnie, and J. Wilmot. 2012. *Credit Suisse Global Investment Returns Yearbook 2012*. Zurich: Credit Suisse.

Ellis, Charles D. 2012. "Investment Management Fees Are (Much) Higher Than You Think." *Financial Analysts Journal*, vol. 68, no. 3 (May/June):4–6.

Kinnel, Russel. 2010. *Morningstar FundInvestor*, vol. 18, no. 12 (August):1–3.

Sharpe, William F. 1991. "The Arithmetic of Active Management." *Financial Analysts Journal*, vol. 47, no. 1 (January/February):7–9.

Vanguard Group. 2012. "Vanguard Total Stock Market Index Fund Admiral Shares" (https://personal.vanguard.com/us/funds/snapshot?FundId=0585&FundIntExt=INT#hist=tab%3A3; accessed 5 July 2012).