Financial Analysts Journal®
EDITOR'S CORNER

Robert D. Arnott
Editor

What Cost "Noise"?

If we earn 50 percent this year and then lose 50 percent next year, are we back where we started? Of course not. We are down 25 percent. This trivial truism has many implications, some simple and some profound. The "cost" of risk is proportional to the variance of returns: If we double our volatility, and so quadruple our variance, we need four times the skill to merely recover the costs of the risks that we are taking. If we double our risk, we must be more skillful; we must have more wins—and larger wins—to earn profits.[1] This challenge is the central peril of leverage and the central benefit of uncorrelated diversification. The arithmetic is a simple, yet core, element of the CFA Program curriculum.

The same arithmetic holds true in the cost of "noise" in market pricing (Treynor 2005). We know that, after the portion of the return that is attributable to beta on market movements is subtracted, the residual return on the typical stock has an annual standard deviation of 30–40 percent. This "idiosyncratic risk" includes both new developments specific to the company that actually change its investment value—relating to production, marketing, research, key personnel—and temporary departures from the unknowable "true fair value" of the company, or noise. Some stocks have much less idiosyncratic risk than 30–40 percent, some vastly more.

If all true fair values were fixed, so that all of this volatility were mere noise, the consequences for capitalization-weighted portfolios would be horrific. One year in six, they would overweight the most overvalued stocks by an average of 40 percent more in their up years than their down years and, reciprocally, underweight the most undervalued by 40 percent more in their down years than their up years.[2]

Benjamin Graham was fond of saying that in the short run the market is a voting machine but in the long run it is a weighing machine. The volatility tied to the noisy short-term departures from the rational quest for true fair value dissipates in the very long run and thus has little impact on annualized returns relative to the constant search for true fair value. But if merely one-eighth of the idiosyncratic volatility is noise, rather than sensible reaction to the changing fundamentals that set the ultimate, unknowable true fair value, then we will see *14 percent noise in pricing*.[3]

This result would cut the return on any portfolio that tracks this noise—as if it were reality—by 2 percent a year. Cap-weighted "market portfolios" track prices, so they may be pulled down by 2 percent a year as a consequence of simple noise in market pricing. Even if the noise is in the opposite direction, *there is no way for this noise to boost returns* for the cap-weighted market portfolio. Underpricing is just as damaging to the return on a cap-weighted portfolio as overpricing. Mispricing relative to an unknowable true fair value is still mispricing. This noise is pure loss.

What about Market Efficiency?

How can this effect be true? In an efficient market, isn't the market-clearing portfolio assuredly efficient? No! Even the most fervent advocate of efficient markets will readily acknowledge that market prices do not match true fair values, the net present value of all future cash flows produced by an investment. We have no way of knowing the true fair values, but without doubt, market prices differ

Editor's Note: I am indebted to Jack Treynor for his detailed editorial comments and suggestions on this subject and to Fischer Black, whose 1986 presidential address to the American Finance Association, called "Noise," suggested that much of the trading that takes place in the capital markets is merely noise trading—based on presumed information that is already reflected in asset prices. Here, however, we are not exploring the implications of trading on noise but the implications of noise on the market-clearing portfolio.

The Editor's Corner is a regular feature of the Financial Analysts Journal.
It reflects the views of Robert D. Arnott and does not represent the official views of the FAJ *or CFA Institute*.

wildly from these hypothetical values. Even in a market efficient enough to defeat every active manager's attempts at information trading or bargain hunting, the share price consists of the true fair value *plus* a substantial error term.

Statisticians recognize that these error terms magnify the tails of a sample distribution. If a coin lands "heads" three times, a statistician will not say that that the observed mean is 100 percent heads, that the observed standard deviation is 0, and therefore, that the odds of heads must assuredly be 100 percent and the coin must have two heads. They have simple tools to convert distorted sample statistics into maximum-likelihood estimates.

In finance, we see the same sorts of adjustments. Robert Engle's GARCH (generalized autoregressive conditional heteroscedasticity) estimation of risk truncates the most recent sample standard deviation toward a longer-horizon historical mean, sometimes with a further truncation by way of another technique known as a "Bayesian adjustment."[4] Such adjustments can be seen in simple tools to adjust the mean, the variance, the skewness, and the kurtosis of a distribution of returns, to dampen the effects of the errors in the outliers. For example, Harry Markowitz and others have shown that the sample kurtosis in a return distribution is so unreliable that an infinite kurtosis has a higher likelihood of being accurate than the observed kurtosis (Fabozzi, Markowitz, and Kostovetsky 1987). Such adjustments can be seen in the betas from Elroy Dimson (1979), who recognized that the best estimate of the true beta combines the measured beta with leads and lags truncated toward the overall market mean of 1.0.

The same problem afflicts capitalization—but no one has seen fit to measure its effect or to correct it. Capitalization has never been examined from the perspective of the statistical properties of its "error term." Consider the company that deserves to be considered the largest on the basis of the unknowable future cash flows that it will produce. Chances are that it is a well-recognized company and that it has a large market capitalization, whether the pricing error on this company is positive or negative. But if it does *not* have a positive pricing error, isn't it highly likely that *some* stock with (1) lower true fair value and (2) positive error will have a larger market capitalization? So, even in a truly efficient market for the pricing of individual stocks, the largest-cap stock is likely to have arrived at the top of the heap at least partly because it has a positive error term large enough to put the stock at the very top of the heap. In short, assuming the errors are themselves random, the top-ranked stock is considerably more likely to have a positive error than a negative error.

How Do the Biggest Stocks Perform?

Table 1 allows examination of this hypothesis from an empirical perspective. Panel A indicates that, based on history, the largest-cap stock on 1 January of each year has a 38 percent chance of outpacing

Table 1. How Have the Largest-Cap Stocks Fared against the S&P 500

	How Often Did #1 Stock Beat Average?				What Percent of Top 10 Beat Average?			
Statistic	1 Year	3 Years	5 Years	10 Years	1 Year	3 Years	5 Years	10 Years
A. What percent of the largest-cap stocks added value vs. the average stock?								
1926–2004								
Average	38%	30%	25%	24%	45%	41%	38%	32%
Standard deviation	49%	46%	44%	43%	27%	26%	25%	25%
Adjusted *t*-statistic	–2.2	–3.1	–3.0	–2.4	–2.9	–2.8	–3.0	–3.2
1964–2004								
Average	34%	26%	19%	16%	40%	37%	32%	29%
Standard deviation	48%	44%	40%	37%	26%	25%	22%	22%
Adjusted *t*-statistic	–2.1	–2.8	–2.9	–2.6	–5.4	–4.2	–4.8	–3.9
B. What magnitude of relative performance did the largest-cap stocks deliver vs. the average stock?								
1926–2004								
Average relative return	–7.1%	–5.4%	–5.3%	–5.0%	–2.9%	–3.3%	–3.2%	–3.0%
Standard deviation					2.7%	1.6%	1.4%	1.4%
Adjusted *t*-statistic					–3.5	–3.9	–3.2	–2.2
1964–2004								
Average relative return	–9.3%	–6.7%	–6.7%	–6.3%	–3.6%	–4.6%	–4.3%	–3.6%
Standard deviation					4.4%	2.0%	1.7%	1.9%
Adjusted *t*-statistic					–2.6	–4.2	–3.7	–1.9

the average stock in the S&P 500 Index during the year and only a 24 percent chance over 10 years. The largest stock delivers a diminishing "win" over the average stock as time spans increase. The *t*-statistics, adjusted for overlapping samples, range from –2.2 to –3.1. Because of the growing reliance on hugging our benchmarks and the flow toward indexing, the statistics for the last 40 years, from 1964–2004, are even worse.

Moreover, this result is not unique to the single largest-cap company. The top-10 companies exhibit the same pattern, albeit in a less pronounced way. Although 45 percent of the top-10 samples managed to outperform over a single year, only 32 percent outperformed over the next 10. The 79 single-year results for the top stock might be dismissed as data mining, but doing so is tough with 790 (largely independent) results for the top-10 stocks. The likelihood of a –2.9 *t*-statistic with 790 samples is barely more than one in a thousand.

Furthermore, the *magnitude* of underperformance by the largest-cap stock is huge; the average shortfall over the subsequent year is 7.1 percent, expanding to a startling 5.0 percent *a year* at 10 years, which compounds to a 40 percent performance shortfall relative to the average stock in the S&P 500 in 10 years. Results for the 10 largest stocks are, again, milder—but more statistically significant—than the results for the top-ranked stock. The compounded 10-year shortfall is 26 percent. Even with the problem of overlapping samples, we cannot easily dismiss the statistical significance of this result. And, again, because of the growing popularity of indexes and indexing, the results are worse for the past 40 years.[5]

Does Capitalization Error Explain Other "Anomalies"?

If a company has a high or low true fair value, its error term is approximately symmetrical, within a rather wide range. What can we say about the implications of this circumstance for the capital markets?

First, there will be true fair values that are high multiples of earnings, book values, dividends, and so forth, and true fair values that are low multiples of fundamentals. If we assume that the market is efficient with regard to the pricing of growth and value stocks, then the errors in pricing these stocks are also symmetrical. These error terms will magnify the dispersion, creating more high-multiple growth stocks and more low-multiple value stocks than the true fair values would justify. The "growth" stocks, which result from an error, will provide a subpar internal rate of return, and the "value" stocks, which result from an error, will have a superior IRR. In this world, *value beats growth,* on average, over time.

Second, assume that the market is efficient with regard to large and small companies—that is, regardless of company size (*not* capitalization), companies have a symmetrical error term. In this case, the error terms will magnify the dispersion, creating more large-cap stocks and more small-cap stocks than the true fair value would justify. The "large-cap stocks," which are the result of an error, will have a subpar IRR, and the "small-cap stocks," which are the result of an error, will have a superior IRR. In this world, *small beats large*, on average, over time.

This effect also drives the results in Table 1. Whereas the 500th largest stock should have a reasonably symmetrical pricing error, with a true fair value equally higher or lower than its market price, the very top stocks are likely—by virtue of being the largest stocks—to be the beneficiaries of more positive error than negative error. So, the (sometimes erroneously) largest-cap stocks will tend to have a lower IRR than the average stock, just as Table 1 suggests. Therefore, *capitalization weighting should underperform equal weighting*, on average, over time.

A Steady State of Errors

Now, let's assume that the average scale of the error term is reasonably stable over time but that the market is constantly hunting for the true fair value; that is, errors that become known diminish but the market replaces them with new unknown errors—a process that some might term "noise." This circumstance is a steady-state world, in which companies with large errors, whether positive or negative, are more likely to move toward true fair value than to undergo expansion of the existing error but in which average error across all stocks remains steady over time.

What is the consequence of the steady-state world? One is long-horizon mean reversion in returns. Another is that the relative returns of large-versus-small stocks and value-versus-growth stocks are magnified. The quest for true fair value concentrates the long-term IRR differences into the near future rather than spreading them out uniformly over time. The much-studied factor returns are magnified, as a joint function of the IRR differences between the overvalued and undervalued assets and the pace at which these errors are corrected.

Suppose these assumptions are wrong and this noise introduces neither a structural inefficiency in the cap-weighted indexes nor some of the market anomalies. For these circumstances to be true, large companies and companies with above-average future growth prospects would need a downward bias in their pricing errors and small companies and slow-growth companies would need an upward bias in theirs. If these asymmetries are just sufficient to counter the IRR implications of large-versus-small and of value-versus-growth, why, then, we have an inefficient market in the pricing of individual stocks.

These same results should affect other markets, such as bonds, country allocations in international portfolios, and niche categories, such as convertible bonds. *An efficient market in the pricing of individual assets, with pricing errors relative to true fair value, requires an inefficient market in the cap-weighted indexes—and vice versa.*

Conclusion

The cost of noise trading and the structural bias in our cap-weighted benchmarks may be a major, even primary, driver of the historical return advantage associated with equal weighting, with value stocks, with small-cap stocks, with GDP weighting in international portfolios, and with rebalancing. Does an equally weighted portfolio earn a positive CAPM alpha by capturing the Fama–French value/distress and size/liquidity factors? Or are the Fama–French value/distress and size/liquidity factors merely proxying for the structural incompatibility between the market-clearing cap-weighted indexes and mean–variance efficiency and for the impact of noise trading?[6]

Should we, then, not hold General Electric, Citicorp, ExxonMobil, or Microsoft—simply because they are the four largest-cap stocks in the market? Of course not! But we should not assume that, even in an efficient market in the pricing of individual stocks, we have an efficient market for the cap-weighted indexes. We should scrutinize the large-cap stocks carefully to make sure that we are not buying simply because the herd is buying.

Questions about the effects of noise trading and measuring size by market capitalization provide fertile grounds for investigation, and the answers may have important implications for the $20 trillion in assets that are managed to or are benchmarked against the cap-weighted benchmarks. The answers will be far from suggesting a new equilibrium theory to improve on the CAPM. But they will suggest ways to profit from the structural inefficiencies of a cap-weighted market-clearing portfolio.

Notes

1. Specifically, the optimal holding weight for the *i*th security, h_i, is determined by its idiosyncratic risk–adjusted alpha, u; that is, $h_i = u_i/\sigma i^2$. To justify the same portfolio allocation, a stock with twice the idiosyncratic "volatility" must offer four times the alpha.
2. This conclusion assumes that the market's pricing errors are normally distributed.
3. Forty percent volatility is sixteen percent variance, one-eighth of which is two percent variance, corresponding to fourteen percent noise.
4. Engle won the 2003 Nobel Prize for the GARCH concept.
5. **Table 2**, showing the experience of the largest and the 10 largest stocks over the subsequent 1- to 10-year spans, is available online with this article at www.cfapubs.org/faj/issues/v61n2/toc.html.
6. Jack Treynor (correspondence 2004) supports the latter view, with a further suggestion that "Fama–French value/distress proxies for investment in brand franchise, which accountants don't capitalize and which has a duration so much shorter than that of plant, which accountants do capitalize, that it is more sensitive to central bank shifts between ease and restraint. Size/liquidity is sensitive to the market price of liquidity (the *nominal* overnight rate, which tracks the velocity of money)."

References

Dimson, Elroy. 1979. "Risk Measurement When Shares Are Subject to Infrequent Trading." *Journal of Financial Economics*, vol. 7, no. 2 (June):197–226.

Fabozzi, Frank, J., Harry M. Markowitz, and Leonard Kostovetsky. 1987. *Mean–Variance Analysis in Portfolio Choice and Capital Markets*. London, U.K.: Basil Blackwell.

Treynor, Jack L. 2005. "Why Fundamental Indexing Beats Cap-Weighted Portfolios." Working paper.

