Efficient Frontier

An Online Journal of Practical Asset Allocation

Edited by William J. Bernstein

Winter 2001

Table of Contents

Home E-Mail

Efficient Frontier

William J. Bernstein

Dunn’s Law Review

The Life and Times of "Core and Explore"

When an asset class does well, an index fund in that asset class does even better.—Steven Dunn

One of the silliest bits of conventional financial wisdom is the notion that while indexing works well with the efficient U.S. large-cap market, there is benefit from active management in the "less efficient" small-cap and foreign arenas. In fact, Charles Schwab enshrined this dubious notion with its "Core and Explore" concept—index ("core") the former, and actively manage ("explore") the latter. This comforts greatly the legions of active-management-associated investment advisors and pension consultants, to whom it grants brief respite from the dustbin of financial history.

Obviously, the good folks at Schwab haven't yet heard the Gospel According to Dunn: that the fortunes of indexing a particular asset class depend on its performance relative to other asset classes. I’ve already covered this ground in a previous piece. John Rekenthaler also looked at the data from a somewhat different perspective (and, to my chagrin, came up with bigger t-stats than mine). For example, take a gander at the summary graph from his article:

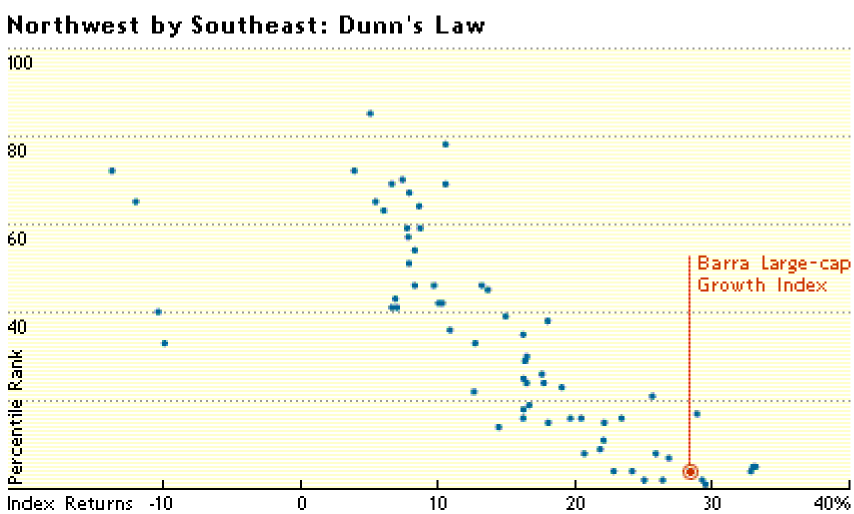

According to "Dunn's Law", when an asset class does well, an index fund in that asset class does even better. Index percentile rankings are calculated versus relevant fund categories.

Mr. Rekenthaler’s graph is a bit confusing since his y-axis is conventionally plotted, meaning that the best index performances are at the bottom (the best performers have the lowest numbers: 1^{st} percentile is the top percentile, 100^{th} the worst). But it is quite clear that there is a strong relationship between how well the asset class does and how well indexing it works. In fact, if you closely examine his plot you'll see that the relationship is curvilinear; there isn't much difference between indexing the best and the middling asset classes. Index performance only begins to suffer with the very worst asset classes.

I've looked at Dunn's Law (DL) both in the U.S. and abroad. Bottom line: DL works very well domestically but not abroad, because of Rekehthaler's "Purity Theory": It's easy for a money manager to stray from his or her style box domestically, but harder for a foreign manager to stray across national boundaries.

It's always nice to test one's hypotheses "out of sample" to guard against data mining: maybe results were just an accident of the time period studied. So I decided to look at DL prospectively, starting with a survivorship-bias-free sample of monthly data beginning January 2000, for the following asset classes and their respective indexes:

Asset Class	Abbreviation	Index Used	Available Fund
Large Cap Growth	lg	Barra LG	Vanguard Growth Index
Large Cap Blend	lb	S&P 500	Vanguard 500 Index
Large Cap Value	lv	Barra LV	Vanguard Value Index
Mid Cap Growth	mg	Barra MG	None
Mid Cap Blend	mb	S&P 400	Vanguard Mid Cap Index
Mid Cap Value	mv	Barra MV	None
Small Cap Growth	sg	Barra SG	Vanguard Small Cap Growth Index
Small Cap Blend	sb	S&P 600	None
Small Cap Value	sv	Barra SV	Vanguard Small Cap Value Index
REIT	re	DFA REIT	DFA Real Estate Securities

As in the previous study, I ranked from 1 (best) to 10 (worst) the performance of each index, and plotted it against the percentile performance for the index in each style box versus the active funds. Purists will chafe at the use of an index instead of an actual fund, but since funds were not available for all of the indexes, I wanted to be internally consistent. Further, of the six funds available from Vanguard, four have managed to equal or surpass their benchmarks by small amounts. This analysis produced 120 data points, plotted below:

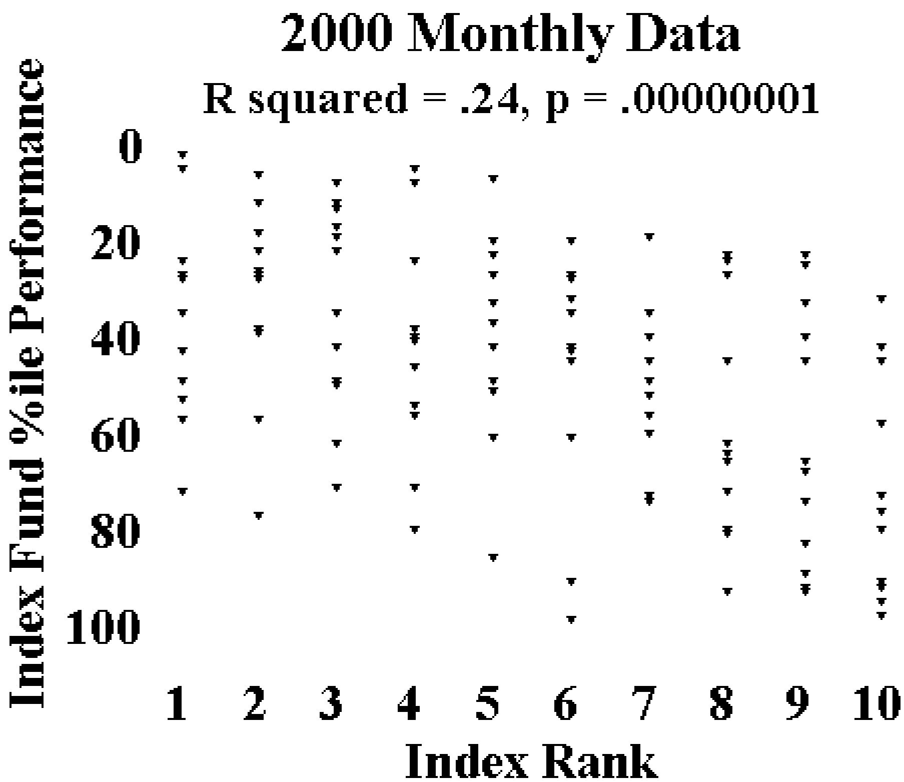

This is a bit of a scattergram, but the "northwest to southeast" bias of the results is still fairly discernible. (Note that I've inverted the percentile scale on the y-axis, placing the "good results" at the top of the graph. This is the opposite of the convention used in Mr. Rekenthaler's graph.) The statistical power of the data, even after only 12 months, is *staggering*: a t-stat of 6.15 and a p value of 10^{-8}. The individual monthly plots are fascinating. In some months, like January, March, August, and December, the points line up like the Rockettes, with highly statistically significant results. (In these four months, the t-stats were 4.00, 3.53, 4.44, and 6.53, respectively, with p values of .004, .008, .002, and .0002.) In other months, the relationship is nonexistent. I've posted the individual plots on a separate page. Those of you who would like the study data may contact me.

But the punchline is the relative ranking of the asset class groups. Below I've plotted the full-year graphs for 1997 and 2000:

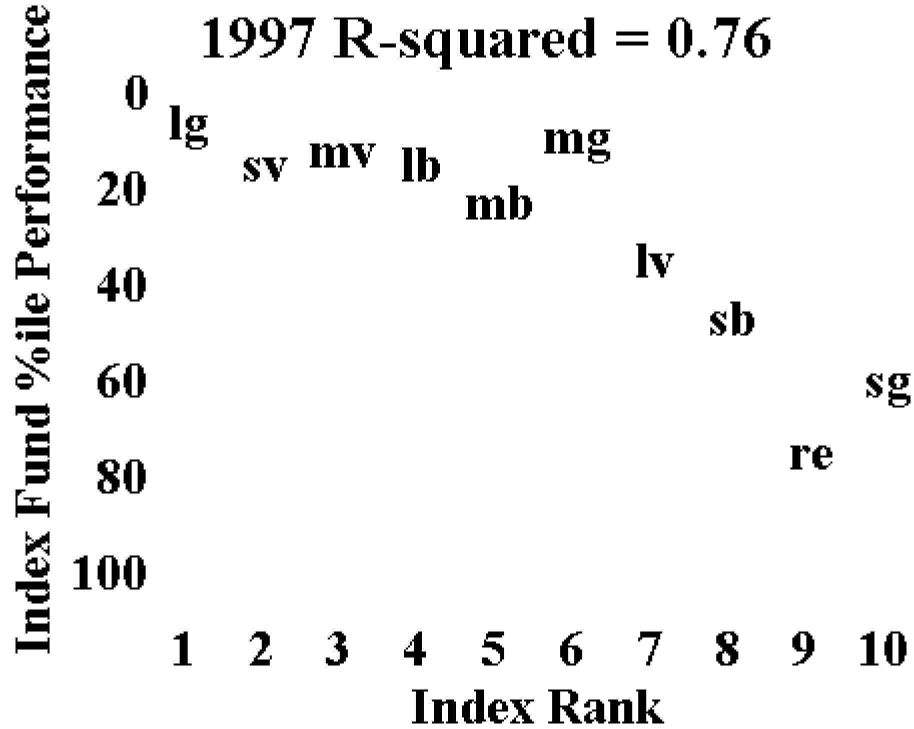

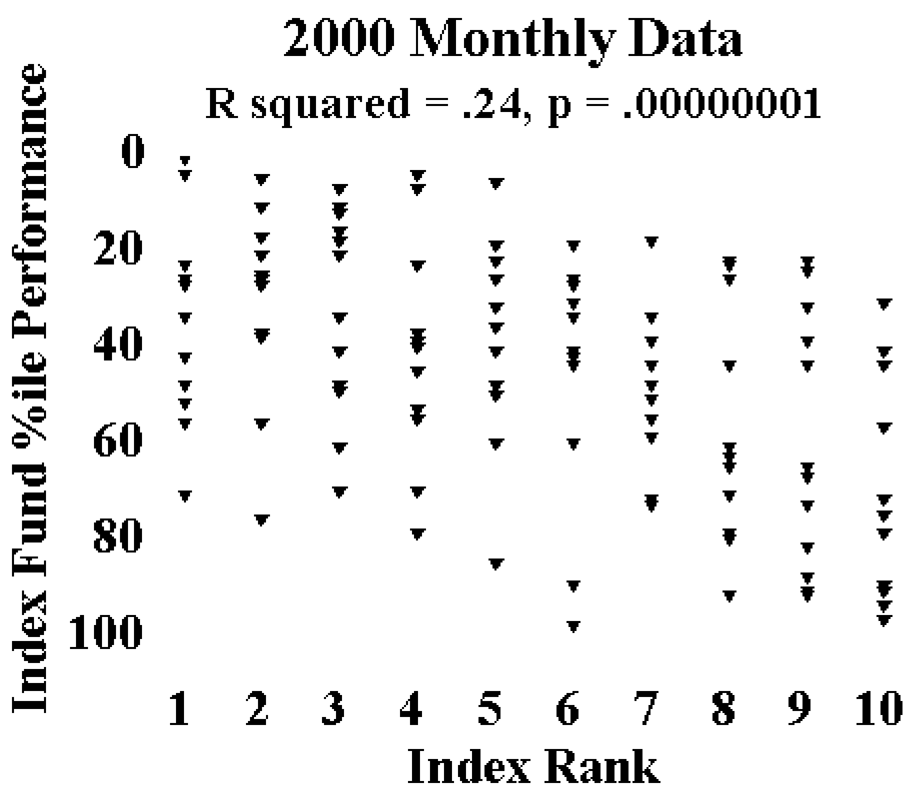

As you can see, in 1997 indexing large caps (the lg, lb, and lv data points) worked better in general than indexing small caps (the sg and sb data points); this certainly did not occur in 2000 when small-cap stocks outperformed large. So good-bye "Core and Explore." In fact, to the extent that there's a small-cap premium, indexing small caps should actually work *better* than large caps. And, to add insult to injury, indexing foreign stocks in 2000 also rebounded: DFA's large-cap foreign and value portfolios ranked in the 37th and 6th percentiles of all foreign funds, respectively, and its long-suffering international small and small value strategies ranked at the 14th and 9th percentiles. To complete the picture, here's the plot for the past five years:

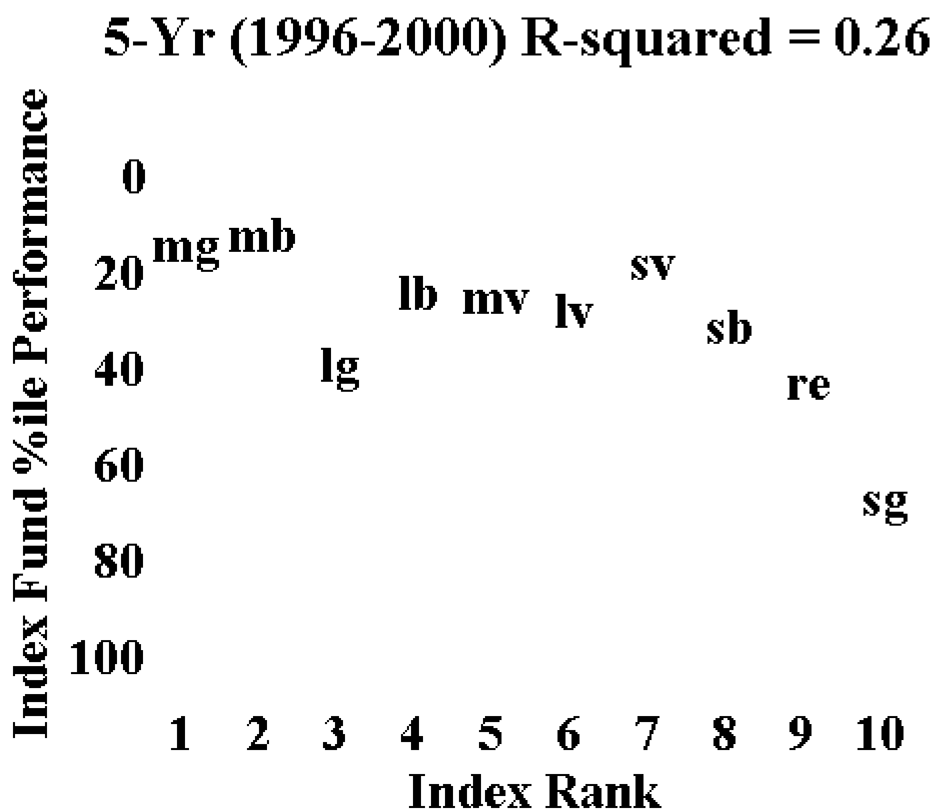

Finally, the phenomenon of "percentile compounding" shows up well in this sample: for the 120 monthly data points in 2000, the average index fund percentile rank was 46th. But over the whole year, the 10 asset classes averaged out to 40th. (So much for the bromide that "indexing didn't work in 2000.") And for the past five years, the average performance of the indexes was 30th percentile.

Out of habit, I'll probably continue this study another year, but given the stunning statistical significance of the above results, I'm probably just amusing myself. Oh yes, and send a memo to the folks at Schwab: Forget about exploring; stick to coring.

Efficient Frontier

William J. Bernstein

The Cowards' Update

In the Cowards' Update for 1999, we looked at the beleaguered condition of the cowards, pummeled again by one more year of S&P 500 outperformance. The asset-class scene reversed course in the year 2000, however, with many "unconventional" assets outperforming the S&P 500. Here's the situation for the three-, five-, and ten-year periods:

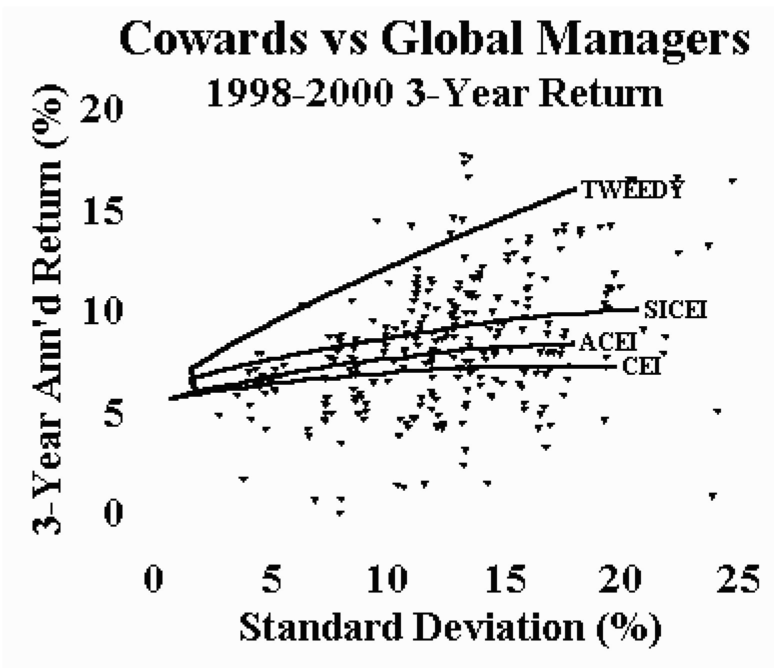

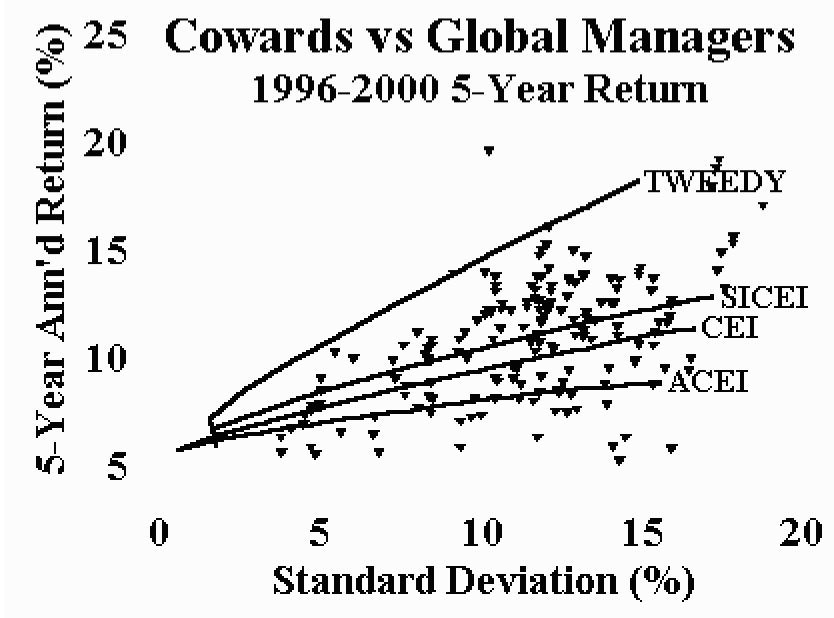

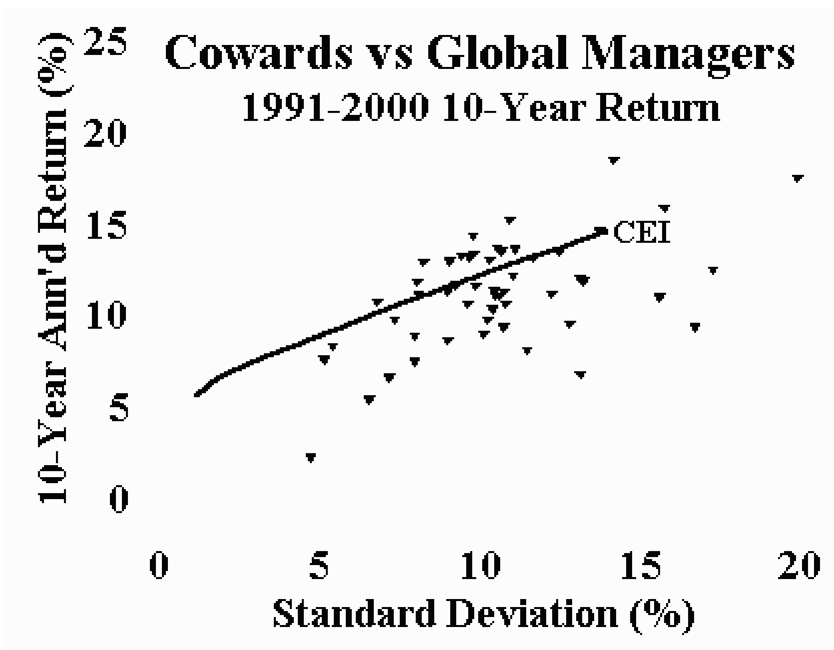

(Those of you unfamiliar with the four coward portfolios, take a look at the July 1997 update for a detailed description of them and the methodology used.)

As before, the robots are doing pretty well for the full ten-year period, but over the three- and five-year periods the picture is a little bleak. With the revival of value investing in 2000, things have begun to look up a bit, in particular, for the long- suffering academic cowards. What's truly amazing, though, is the rightward explosion in the scatter of the active managers (into the high-risk region), particularly noticeable in the three-year plot, as many investors are rudely awakened to the bets being taken in their portfolios.

Even though 2000 was a good year for the well-diversified investor and 1999 excellent for the global investor, these two years clearly do not make up for the relative devastation preceding them.

Consider the three- and five-year returns for the index funds tabulated below. Large-cap growth still leads the pack by a wide margin:

Index	**Index Fund Sampled**	**3-Year Return**	**5-Year Return**
Continental Small Companies	DFA Continental Small Company	6.66	9.15
Emerging Markets (equally weighted)	DFA Emerging Markets	3.28	-0.09
Small Japanese Stocks	DFA Japanese Small Company	6.32	-15.94
EAFE Index	DFA Large-Cap International	9.32	7.96
Pacific Rim Small Companies	DFA Pacific Rim Small Company	3.95	-5.75
U.S. Small and Medium Companies	DFA U.S. 6-10 Small Company	6.68	12.16
U.S. Small Companies	DFA U.S. 9-10 Small Company	5.06	10.87
U.K. Small Companies	DFA United Kingdom Small Company	6.26	10.03
REITs	DFA Real Estate Securities	2.12	11.21
S&P 500	Vanguard 500 Index	12.30	18.31
Emerging Markets (cap weighted)	Vanguard Emerging Markets Stock Index	-1.41	-1.58
EAFE-Europe	Vanguard European Stock Idx	11.33	15.76
Precious Metals Stocks	Vanguard Gold & Precious Metals	4.68	-7.01
U.S. Growth Stocks	Vanguard Growth Index	12.51	19.17
EAFE-Pacific	Vanguard Pacific Stock Index	6.10	-3.93
U.S. Value Stocks	Vanguard Value Index	11.04	16.69

The best automaton conitnues to be the Tweedy Browne coward's portfolio with its worldwide Graham-and-Dodd approach. Once again, the more passively managed a coward's portfolio, the worse it performed.

According to Morningstar, the average "global" and "asset allocation" fund has only about 20% of its allocation abroad and is overly concentrated in large growth stocks, with an average PB of 6.4 and a median market cap of $37B. The analagous index composition over the past five years would be an 80/20 S&P/EAFE mix. By way of comparison, I plotted the five-year performance of the active managers versus a porfolio consisting of the relevant DFA funds, mixed with DFA's one-year bond fund:

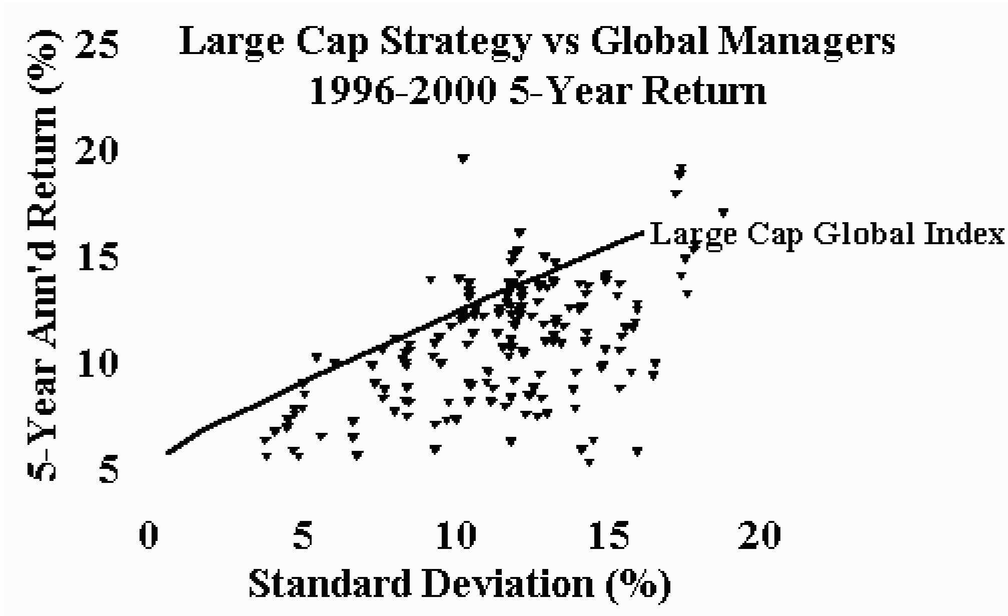

This puts the performance of the active-management crowd into high relief: Clearly, their good fortune is due almost entirely to asset allocation. Further, it would appear that most of them couldn't transact their way out of a paper bag. Of course, it is just possible that the active managers are exhibiting skill in asset-class management by virtue of choosing this winning allocation. So let's bring on a few more years of outperformance by small, value, and foreign stocks, and see how good these folks really are.

The most disturbing thing about the above data for the efficient marketeer is the dominance of the Tweedy Browne Global strategy. For those of you unfamiliar with Tweedy Browne, this management team employs a highly diversified, disciplined, nation-blind value strategy. As such, the authors of "What Has Worked in Investing" and "Ten Ways to Beat an Index" seem to be fulfilling their promise. Nevertheless, appearances can be deceiving. Their portfolio—hedged and largely European—has been in the "sweet spot" of global equity and currency markets for the past half decade. (Hedging the EAFE has added 3% per annum to returns over the past five years.) It remains to be seen how they will do when those worms turn.

William J. Bernstein

Rolling Your Own: Three-Factor Analysis

If you're responsible for overseeing a gazillion dollars of pension money, it's not enough to open up the quarterly mutual fund supplement of your local paper and compare your manager to her peers or the indexes. The reason is simple: you really don't know whether her good/poor performance was due to a lucky/unlucky style tilt in her security selection. Maybe her large-growth top-heavy bias added several percentage points to her returns, but beyond that she couldn't pick a stock out of a hat if her life depended on it. How do you correct for these biases? The best way to do so is to use a "factor-based analysis" (FBA). FBA is to active money managers, what a light switch is to cockroaches.

There are as many different kinds of FBA as there are portfolio analysts, but the simplest, best known and probably most powerful is the so-called "three-factor model" (3FM) of Fama and French. In April 1999, I touched on this method in an article on small-growth investing, but didn't provide a lot of details. The reason was, the returns series necessary for these calculations were not yet in the public domain.

Now, thanks to Ken French's wonderful online data library, the 3FM is available to the small investor for the analysis of manager returns. The "Fama-French Research Data Factors" file contains the monthly and quarterly returns for the four necessary data series:

- The 30-day T-bill return
- The return of the total market (CRSP 1-10) minus the T-bill return (Mkt)
- The return of small company stocks minus that of big company stocks (Small minus Big, or SmB)
- The return of the cheapest third of stocks sorted by price/book minus the most expensive third (High minus Low, or HmL)

Once you have the returns series for the manager you want to look at, you're on your way. Here's the cookbook:

1. First, you need a spreadsheet package capable of performing multiple regression. Much as it pains me to admit it, Microsoft Excel does this best. (If you need help with either statistical analysis or Excel, I cannot recommend highly enough John Neufeld's ***Learning Business Statistics with Microsoft Excel 2000***.)

2. Next, you need to download from Ken French's site the regression series.

3. Third, you need to extract the monthly returns for your money manager. This is the tricky

part. Ideally, you'll have Morningstar Principia Pro Plus, but even then the extraction process is not completely painless. You have to output the monthly returns from the Advanced Analytics section (Risk and Reward/Rolling Return - Table, and make sure you select the one-month rolling period). Then, once you've opened this in Excel you need to use the copy/edit/paste special/transpose sequence to get it into vertical format. And if you only have the plain-vanilla Principia Pro package, you'll have to extract the monthly returns point-by-point with the graph function.

Lay this series next to the T-bill returns, and subtract the latter from the former. You now have a column that represents the difference between the two, which is the "risk-adjusted return" (RAR). As an example, I've chosen the Dodge and Cox Stock Fund. Column B contains the raw returns; column C, the T-bill returns; column D, the RAR (column B minus column C); and columns E, F, and G, the three regression series. The six rows of the spreadsheet should look something like this:

	A	B	C	D	E	F	G
1		Dodge & Cox St	T-bill	Fund Minus T-bill	Mkt	SmB	HmL
2	Jul-90	-1.231	0.68	-1.911	-1.62	-2.9	-0.36
3	Aug-90	-9.873	0.66	-10.533	-9.83	-3.34	1.63
4	Sep-90	-5.683	0.6	-6.283	-5.99	-3.62	0.56
5	Oct-90	-2.009	0.68	-2.689	-1.92	-5.76	0.27
6	Nov-90	8.516	0.57	7.946	6.03	0.27	-2.87
7	Dec-90	3.519	0.6	2.919	2.35	0.88	-1.63
8	Jan-91	5.646	0.52	5.126	4.38	3.58	-1.65

4. Last, multiply regress the RAR versus the Mkt, SmB, and HmL series. I recommend including the labels in the regression menu.

Here's what the regression dialog box looks like:

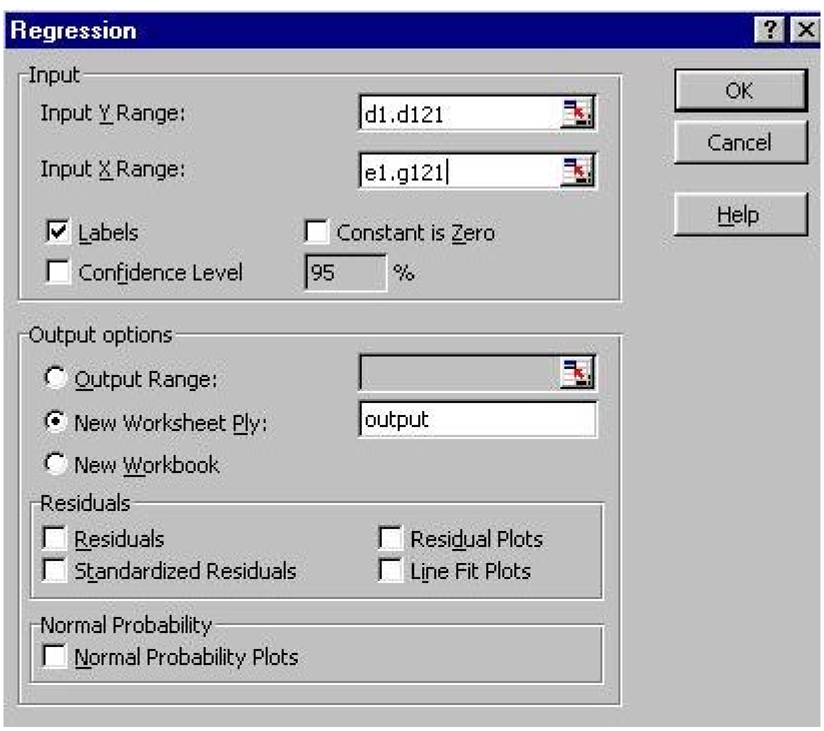

Note that I've started the regression in row 1, where the labels are, and checked the "label" box. Finally, the abbreviated output is shown below:

Regression Statistics	
Multiple R	0.9118679
R-Squared	0.8315031
Adjusted R-Squared	0.8271454

	Coefficients	*Standard Error*	*t-Stat*	*p-Value*
Intercept	-0.2058513	0.1590456	-1.2942908	0.1981357
Mkt	1.0575511	0.0450486	23.475762	6.991E-46
SmB	-0.0538972	0.0453912	-1.1873937	0.2374981
HmL	0.4876222	0.0638117	7.6415827	6.747E-12

The first group of statistics shows us that we have a fairly decent fit with the data, with a raw R-squared of 0.91. This tells us that 91% of the fund's monthly returns are explained by the three factors.

The second group is far more interesting. First, the "intercept" is the fund's alpha, negative 0.206 per month, or about 2.5% per year. In other words, the managers, after expenses, underperformed the regression-based benchmark by that amount. However, the t-stat and p-value tell us that this is not statistically significant. Next, we have the "loadings" for the three factors. The Mkt loading is 1.06. This is the traditional beta of the fund. Most equity-only funds have values very close to 1.0. The SmB loading is -.05. This means that the fund is primarily large cap. (A zero value signifies large cap, and a value of greater than 0.5, small cap.)

Finally, the HmL loading is 0.49, which tells us that we're looking at a value fund. (A zero value defines a growth portfolio, a value of more than 0.3, a value fund.)

You cannot overestimate the power of this model. Wirehouse reps and pension consultants are fond of "attribution analysis," which usually consists of qualitative excuses for a manager's lackluster performance in a given period. The 3FM throws investment results into much sharper relief. Let me give you a real-life example:

I sit on a committee overseeing a local institution's pension portfolio and was treated to our consultant's sheepish exposition of our small-cap fund's obviously dismal performance. But when its returns were tossed into the 3FM grinder, a stunning picture emerged: market and size loadings of 1.0, a value loading of zero, an R-squared of 0.99, and an alpha of *minus 10% per year*. In other words, we had been burned by a manager doing his best to mimic the CRSP 9-10 index (which has almost identical loadings), but who couldn't transact his way out of a paper bag.

To be fair, very few managers have highly significant positive alphas; among them is none other than John Bogle Jr, who accomplishes this by riding momentum (probably the fourth factor!) for all it is worth. Unfortunately, positive prior alpha does not predict future positive alpha. Robert Sanborn of Oakmark Fund manufactured spectacularly positve results in his first few years, only to achieve similarly spectacular negative results just before he "retired."

As should be evident from the above methodology, this technique is not for the faint-of-heart. Even with my fingers flying, the export-transform-import-regress cycle required usually takes about ten minutes for each fund. But master it and you've gone most of the way toward minimizing your chances of being sandbagged by a consultant or manager.

Efficient Frontier

William J. Bernstein

The Retirement Calculator From Hell - Part II

One of the joys of this Web site is the appreciation I get for dispensing investment commentary. And by far, the piece that has elicited the greatest response is "The Retirement Calculator From Hell."

Long story short: Almost all retirement calculations are done with a straightforward amortization algorithm (via a handheld financial calculator or a software product) that does not take into account the fluctuation in returns. The risk being, if you get a string of bad years at the front end of your retirement, you could run out of money long before the calculated number of years based on an overall rate of return.

For example, the real return on an 80/20 mix of large- and small-cap stocks for the 30-year period 1966-1995 was 5.9%. Plugging this return into a financial calculator tells us that an annual real withdrawal (of the beginning principal, inflation-adjusted) of 7.2% over this 30-year period is possible. Yet the following two plots tell a more worrisome story: Even a real 5% withdrawal rate leads to an Alpo diet long before 30 years is up. Paradoxically, in this scenario, because bonds offset poor stock returns in the first years, adding intermediate treasuries, with only a 2.9% real return for the period, helps stretch the money. These portfolios were rebalanced annually. (Readers of the original piece may notice that this is even grimmer than the plots first published. Thanks to reader Alan Beeman for catching a computational error in the first article. The corrected graphs now appear in both articles.) I apologize for the messy plots: all five of the portfolios, from all-stock to all-bond, decay at nearly the same rate, and it is difficult to visually separate them out.

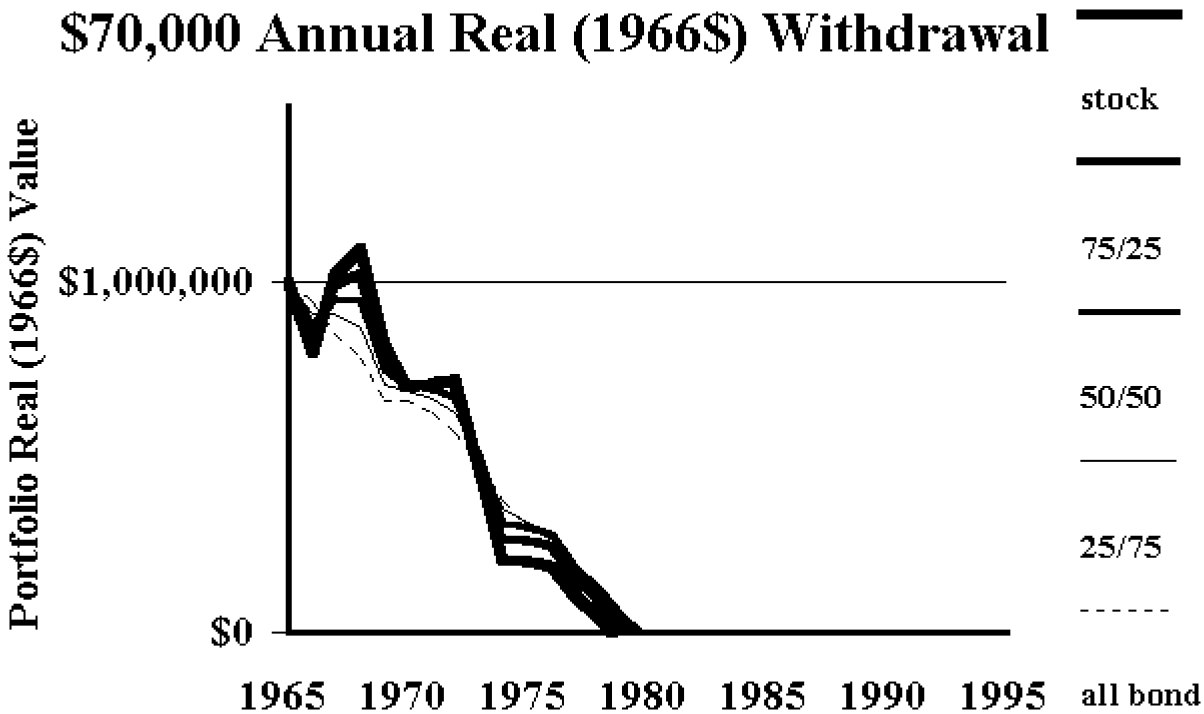

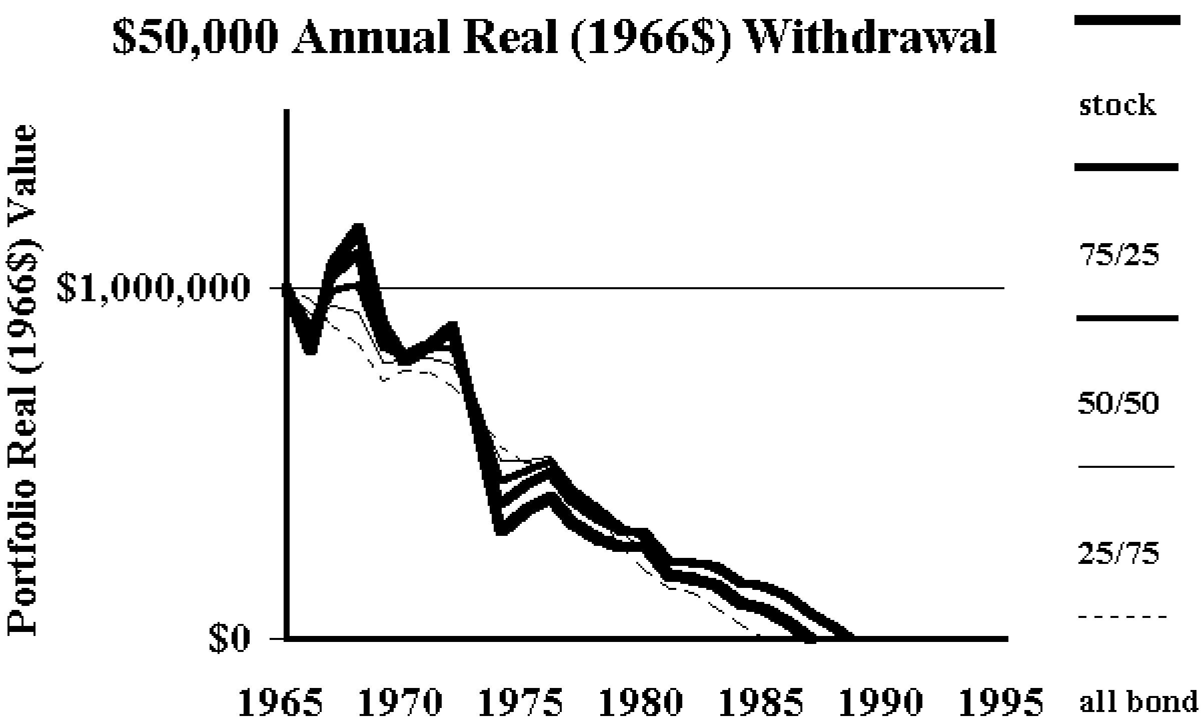

Even this may be an overly rosy scenario. With the S&P yielding just 1.2% in dividends, a 5.9% real return requires either an inflation-adjusted earnings/dividend growth rate of 4.7% or a further expansion of multiples. Neither of these are events I'd want to bet *my* golden years on.

So we have a real problem. First, standard amortization doesn't factor in return variability. Second, historical data do not take into account the reduced equity risk premium we currently expect.

Is there a better way? Yes. I have received quite a few e-mails from engineering types about Monte Carlo analysis, and I've always been dubious about using it for portfolio optimization. As a simple retirement calculator, however, Monte Carlo simulation offers clear-cut advantages over running raw historical data or simple amortization algorithms.

What exactly is Monte Carlo, and how do we deploy it for retirement planning? First, a few assumptions... Let's say that we're examining a 100% stock portfolio for which we've assumed a 7% real return, with a returns standard deviation of 12%. Because of "variance drag" (more on this four paragraphs below), in order to obtain an *annualized* return of 7.0%, we need an *average* annual return of 7.7%. (If you are confused about the difference between average and annualized returns, take a look at the discussion of this about a third of the way into Chapter 1 of ***The Intelligent Asset Allocator***.) Let's further assume that we're withdrawing 5% of our initial capital in real dollars each year.

Monte Carlo simulates returns variability by adding in a random component to each period's return. Say we've decided to look at a 50-year period and are examining annual returns. That's 50 time periods. In each, we generate a random number with a normal distribution. This means that two-thirds of these random numbers will be between –1.0 and +1.0, 95% will lie between –2.0 and +2.0, and nearly 99.9% will be between –3.0 and +3.0. Let's say that the random number that pops up for the first year is –1.15. The return for the first period would then be:

7.7% - (1.15 x 12%) = -6.1%

Remember too, in each period we have to withdraw 5% of the initial capital. We generate 50 different random normal numbers to calculate a run of 50 years. The run is categorized as a "success" if we have any money left after 50 years, and a "failure" if we don't. Repeating the process thousands of times gives a fairly solid statistical estimation of the true success rate.

Monte Carlo simulation requires a large amount of computation. For example, 100,000 runs of 50-year simulations using monthly returns involves 60,000,000 separate monthly calculations. As a practical matter, this is beyond the memory capacity of most desktop spreadsheet systems. (The outer limits of microprocessor/spreadsheet analysis is about 30 annual periods x 10,000 runs.)

There are also some computational nuances to consider. The most important one is this: varying SD also varies return, via so-called "variance drag." Say you begin with an annual return of 10% each and every year (zero SD). You will of course wind up with a 10% annualized return over the long haul. But crank in a normally-distributed random term with 20% SD, and you find that your range of annualized returns falls to a median of about 8%. This needs to be adjusted for, as it was in the above example when the assumed 7% annualized return was replaced by an equivalent average annual return of 7.7%. Changing annual returns and SDs to monthly ones raises other computational issues.

So there's no getting around a dedicated application for this purpose. My colleague David Wilkinson graciously cooperated with a nifty bit of code that clips along at about 4,000 runs per second of 600 months each on my first-generation Pentium system.

(For those who want to explore the spreadsheet option, I have posted a 100-run by 30-year sample here.)

The first table below shows the results for monthly returns and withdrawals for an "old-paradigm" returns scenario, with high (7%) real stock returns, low (2.5%) real bond returns. I assumed a relatively low stock SD of 12%. This is worth a comment here. When one runs Monte Carlo simulations of stock returns using the historical 15%-20% SD, one comes up with higher long-period (>20 years) variability than is actually observed. The reason is, over the long haul stock returns have a tendency to mean-revert. To correct for this, I've lowered the SD a bit.

Next, I added bonds to the mix, assuming a 2.5% real return and an SD of 5% for intermediate-term instruments. A zero stock-bond correlation was assumed. The rebalanced overall portfolio real return/SD data were computed using another of David Wilkinson's programs, MvoPlus. The return and SD data are displayed in the first two columns. These were then fed into David's new MCRetire program, generating the first table. The 30-year success rates for withdrawals of 4%, 5%, 6%, and 7% of the the initial principal amount in real terms are listed in the last four columns.

30-Year Success Rates (Monthly Withdrawals)
Stocks 7%, Bonds 2.5%

	Return	**SD**	**4.0%**	**5.0%**	**6.0%**	**7.0%**
Stocks	7.0%	12.0%	98.7%	93.4%	81.0%	63.3%
75/25	6.0%	9.09%	99.3%	93.4%	76.0%	50.3%
50/50	4.9%	6.5%	99.6%	91.4%	61.2%	24.9%
25/75	3.8%	4.8%	99.3%	77.4%	24.8%	2.6%
Bonds	2.5%	5.0%	87.2%	33.4%	3.7%	1.3%

I also examined a "new-paradigm" returns scenario with 4.5%/3.5% stock/bond returns in the second table.

30-Year Success Rates (Monthly Withdrawals)
Stocks 4.5%, Bonds 3.5%

	Return	**SD**	**4.0%**	**5.0%**	**6.0%**	**7.0%**
Stocks	4.5%	12.0%	88.8%	70.4%	48.0%	28.3%
75/25	4.4%	9.09%	94.3%	75.2%	46.6%	22.2%
50/50	4.2%	6.5%	98.2%	80.1%	41.3%	12.0%
25/75	3.9%	4.8%	99.5%	81.1%	28.9%	3.4%
Bonds	3.5%	5.0%	98.2%	68.3%	18.5%	0.2%

Obviously, the new paradigm with its low equity risk premium is less forgiving than the old paradigm. You can safely withdraw about 5% of initial principal under the old paradigm, but only about 4% under the new. No surprise here. And it's also no surprise that under the new paradigm, very high levels (75%) of bonds seem to be optimal. But even under the old paradigm at the safe 4% and 5% withdrawal rates, a 25% bond contribution is well-tolerated. Finally, if your withdrawal rate is "too high" for safety, 100% stocks give you your best (albeit poor) chance of success under both scenarios.

A few philosophical observations are in order. One of finance's ironies is that concentrating your portfolio is the best way to become very rich; it is also the best way to become poor. So an investor's first duty is to decide which of these two goals is paramount. This article is aimed at those who seek safety in retirement, not great riches. (If you haven't saved enough to sustain your chosen retirement lifestyle, then in fact you do maxmize your chances of success by investing aggressively, as can be seen in the last column of the second table. But this is damage control—you are simply optimizing grim odds.)

Obviously, Monte Carlo has its flaws. But playing with your portfolio's return/risk assumptions gives you a good idea of the tradeoffs between these inputs and retirement success. And it’s lot more realistic than using a simple amortization calculation... by a country mile.

For those of you who are interested, David’s MCRetire is currently in beta testing and should be available soon. The estimated price is in the $20-$30 range. If you want details about David's software products, you can get them at Efficient Solutions.

Efficient Frontier

William J. Bernstein

Link of the Month: Ken French's Database

When I first got into the portfolio theory business in the early 1990s, the barriers to entry for the small investor were almost insurmountable: there were no guidelines to practical allocation and, worse, no software. But by far the biggest problem was the dearth of useful returns data. One high-end analyst shook his head and admitted to me, "Even I have problems getting this stuff."

No more. Now, all manner of domestic data is available online from Barra, Wilshire, and even MSCI. And it keeps getting better. Professor Ken French of MIT has recently posted an extensive dataset of returns of U.S. stocks sorted by size, price/book, price/earnings, price/cash flow, dividend yield, and industry.

Below is a simple example of what you can do with his files: a plot of the growth of $1 invested in each of his six style categories (small value, blend, and growth, and large value, blend, and growth).

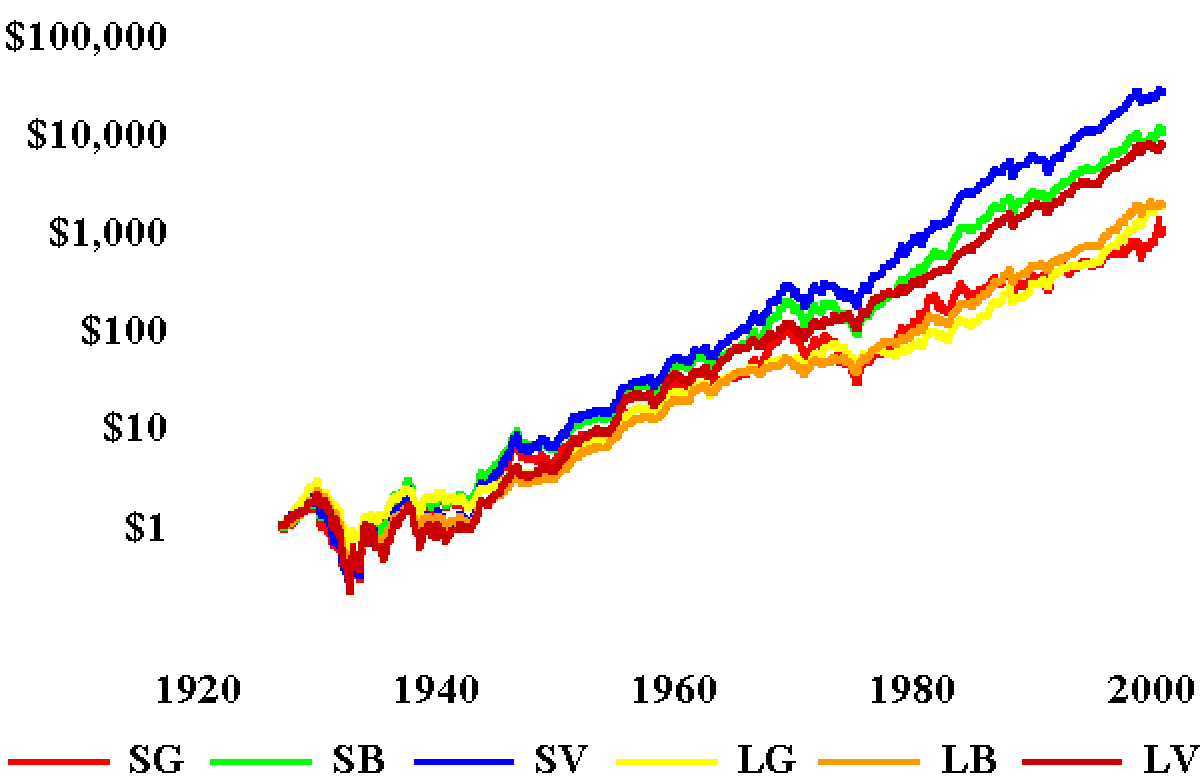

The files are not small; he had to limit the series to text files to conserve size. You'll have to be proficient at spreadsheet importation. (I recommend using space-delimited importation. Column-registration errors will occur, so you'll need to make sure that the columns stack up properly.)

This is a real treasure trove. Explore and enjoy.

Finally, an amusing tidbit from Larry Swedroe: the Internet Wasteland page from the *VTO Report*. Not for those with weak stomachs.

Home E-Mail