

What Is An Index?¹

Andrew W. Lo²

This Draft: 12 October 2015

Abstract

Technological advances in telecommunications, securities exchanges, and algorithmic trading have facilitated a host of new investment products that resemble theme-based passive indexes but which depart from traditional market-cap-weighted portfolios. I propose broadening the definition of an index using a functional perspective—any portfolio strategy that satisfies three properties should be considered an index: (1) it is completely transparent; (2) it is investable; and (3) it is systematic, i.e., it is entirely rules-based and contains no judgment or unique investment skill. Portfolios satisfying these properties that are not market-cap-weighted are given a new name: “dynamic indexes.” This functional definition widens the universe of possibilities and, most importantly, decouples risk management from alpha generation. Passive strategies can and should be actively risk managed, and I provide a simple example of how this can be achieved. Dynamic indexes also create new challenges of which the most significant is backtest bias, and I conclude with a proposal for managing this risk.

¹I thank Manita Singh for excellent research assistance, David Goodsell, Alex Healy, David Lafferty, Caren Leedom, and Rob Sinnott for helpful comments and discussions, and Jayna Cummings for editorial assistance. Research support from the MIT Laboratory for Financial Engineering is gratefully acknowledged. The views and opinions expressed in this article are those of the author only, and do not necessarily represent the views and opinions of any institution or agency, any of their affiliates or employees, or any of the individuals acknowledged above.

²Charles E. and Susan T. Harris Professor, MIT Sloan School of Management, 100 Main Street, E62-618, Cambridge, MA 02142; Principal Investigator, MIT Computer Science and Artificial Intelligence Laboratory; and Chief Investment Strategist, AlphaSimplex Group, LLC.

Contents

1	Introduction	1
2	A Brief History of Indexes and Index Funds	3
3	What Is An Index?.....	7
4	The Brave New World of Strategy Indexes.....	8
5	Disbanding the Alpha Beta Sigma Fraternity	10
6	Beware of Backtest Bias.....	18
7	Conclusion	22
	References	24

1 Introduction

The standard definition of a financial index such as the Standard & Poor's (S&P) 500 is a market-capitalization-weighted average of a specific and relatively static list of securities. The financial index was first devised in the late nineteenth century as numerical shorthand for market activity. Today, indexes serve many purposes. In addition to their original function of information compression, indexes act as indicators of time-varying risk versus reward, and as a benchmark for performance evaluation, attribution, and enhancements. And since the advent of the Capital Asset Pricing Model (CAPM), indexes have been used to construct passive investment vehicles and as building blocks for portfolio management.

However, recent technological advances in computing, trading, trade processing, telecommunications, and derivative securities have greatly increased the scope of financial products and services that are possible, including new forms of indexes that have little resemblance to a static market-cap-weighted portfolio. In this paper, I revisit the notion of an index in light of these new possibilities and propose a new perspective in which a broader definition offers many advantages but also some potential pitfalls.

Charles H. Dow published the first market index in the era of the telegraph, when markets moved unbelievably slowly by modern standards, and a simple average of stock prices was enough to communicate important information about the health of a market. But thanks to Moore's Law, technology has advanced exponentially since Dow's time, which has important implications for what we mean by an index. For example, consider the MSCI World index, a market-cap-weighted average of over 1,600 stocks from what are sometimes called the "developed markets," founded in 1969. Why does this index take this specific form? Market-cap weighting emerged largely because of older technological limits on trading and portfolio construction. It implies a buy-and-hold portfolio that does not need to be rebalanced except when securities enter or leave the index list. It is a very specific and limiting way of constructing an index, an artifact of its financial and technological era.

Technology and the information revolution have changed so much of our daily lives that it should not be surprising that it has also changed the way we think about saving and investing. Modern trading technology opens up a whole new spectrum of possibilities for defining indexes and creating financial products around them. Some examples are: target-date and life-cycle funds, which change their asset allocation characteristics as they approach their target dates; hedge-fund replication strategies, which replicate the betas of entire classes of hedge funds; trading-strategy indexes, which use transparent, mechanical rules to implement specific trading strategies such as currency carry trades or risk arbitrage; and "fundamental" or so-called "smart beta" indexes, where stocks within a portfolio are weighted according to their fundamentals or other non-market-cap factors.

Mutual funds and active exchange-traded funds (ETFs) associated with these indexes have already captured the imagination of many investors, with over \$544 billion invested in “smart beta” ETFs alone as of February 2015.³ In the aftermath of the Financial Crisis, these new forms of investment resemble a Cambrian explosion of new species, an eruption of financial innovation and diversity after a long fallow period.

But this Cambrian explosion also contains some potential concerns. As many of its critics have commented, “smart beta” need not be smart at all. For the inexperienced investor, “smart beta” is often accompanied by “dumb sigma,” unnecessary and unanticipated kinds of portfolio risk that do not carry a positive risk premium. One obvious example is the idiosyncratic risk of a highly undiversified portfolio, but there are other examples of risks that are not adequately rewarded, especially in the face of market distress.

A new framework is needed for thinking about indexes, indexation, and the distinction between active and passive investing that reflects the new reality of technology-leveraged investing. Although many useful new financial products and services have been made possible by technology, any form of leverage—including technology-leveraged investing—can create new and greater risks. In the financial industry, Moore’s Law must be weighed against the technology-updated version of Murphy’s Law: Anything that can go wrong, will go wrong, and will go wrong faster and bigger when computers are involved.

The starting point for this new framework is to generalize the definition of a financial index by focusing on its basic function. If an index is to be used as a benchmark against which managers are judged, it must have three key characteristics: it is transparent, investable, and systematic. Traditional indexes such as the S&P 500 clearly satisfy this definition, but so do other portfolio strategies that involve more active trading such as target-date funds and publicly disclosed, rules-based 130/30 strategies. Under this new definition, financial indexes can vary greatly in complexity. To distinguish the more complex versions from the traditional market-cap-weighted indexes, I will refer to the traditional indexes as “static” and more complex indexes as “dynamic.”

Dynamic indexes may contain more subtle risks, such as tail, illiquidity, or credit risk. As a result, these types of indexes will require more sophisticated consumers, those with the education and experience to properly assess these risks and use these indexes responsibly. However, one of the most important implications of this new definition of an index is that investing and risk management can be decoupled—passive investing need not, and should not, imply passive risk taking, as it currently does. The two pursuits are distinct, and it is important to separate them, especially during turbulent market conditions.

³ Evans (2015).

Moreover, due to their complexity and method of construction, dynamic indexes will be much more prone to backtest bias, so investors will need to use more sophisticated judgment to evaluate them. But if used properly, dynamic indexes can greatly benefit both investors and portfolio managers by allowing them to construct more highly customized portfolios that can achieve long-run investment objectives by managing short-run risks more effectively.

2 A Brief History of Indexes and Index Funds

Dow published the first U.S. stock market index in 1884, the Railroad Average, which still exists today as the Dow Jones Transportation Average (DJTA). The DJTA was a precursor to the Dow Jones Industrial Average (DJIA), which Dow began publishing in 1896. Today, the DJIA is one of the three most recognized indexes in the world, along with the S&P 500 and the NASDAQ Composite. Dow created these averages to illustrate his theories in what is today called technical analysis.⁴ Even practitioners who disdained Dow theory, however, found the Dow indexes were tremendously useful for following the stock market as a whole.

The Dow indexes were initially calculated as a simple average of stock prices. In fact, they were the earliest version of what is now called an equal-weighted index. In 1928, this method was revised to a price-weighting system in which each stock is weighted by its price relative to the sum of all stock prices within the index. Even before then, however, economists had suggested a market-cap-weighting system to be a more effective method to represent the overall market, since the price of a stock could be moved by splits or merges with little or no financial effect, while market capitalization directly apportioned a stock's importance to its size in the market. Market-cap-weighting (sometimes called value-weighting) solved those problems. In 1923, the Standard Statistics Company used market-cap weighting, apocryphally at the urging of the economist Irving Fisher,⁵ to compute an index to compete with the Dow, the precursor to today's S&P 500. Today, the large majority of financial indexes follow a market-cap-weighting scheme.

These new indexes, in turn, stimulated new thinking about their possible uses. In 1960, Edward Renshaw at the University of California and his graduate student Paul Feldstein first proposed the creation of index funds in their paper, "The Case for an Unmanaged Investment Company."⁶ Their work compared 89 diversified mutual fund returns to those of the DJIA, demonstrating that only 11 funds had higher returns than the DJIA. This was an idea slightly ahead of its time, since there was as yet no compelling theoretical reason for

⁴ Lo and Hasanhodzic (2013, pp. 82–84).

⁵ Fox (2011, p. 27).

⁶ Renshaw and Feldstein (1960).

their result to be more than a numerical coincidence—managers of the time remained confident in their ability to beat the market.

Two revolutionary financial theories catalyzed much broader acceptance of the index fund: the CAPM, introduced independently by Sharpe (1964) and by Lintner (1964) in 1964, and the Efficient Market Hypothesis (EMH), independently proposed by Fama (1965) and by Samuelson (1965) a year later. The CAPM allowed investors to construct a mean-variance-efficient portfolio simply by holding a basket of all stocks in proportion to their market capitalization, i.e., the market portfolio. The EMH, meanwhile, had an obvious corollary that, after accounting for transactions costs and fees, active investing could not outperform passive investing on average. It is no exaggeration that, in combination, these two theories democratized personal investing by taking the reins of portfolio management from the active stock-picking gunslingers of the day and handing them over to a broadly diversified index fund which served as a proxy for the market portfolio. Almost overnight, typical retail investors were given new tools that allowed them to invest on their own, albeit passively.

Although academic research provided the seeds from which the index fund business grew, many credit John Bogle as the pioneering practitioner who planted these seeds and cultivated its first harvest in 1976: the Vanguard Index Trust. However, this was only the first index mutual fund; Bogle generously ascribes the roots of his business to others:⁷

The basic ideas go back a few years earlier. In 1969–1971, Wells Fargo Bank had worked from academic models to develop the principles and techniques leading to index investing. John A. McQuown and William L. Fouse pioneered the effort, which led to the construction of a \$6 million index account for the pension fund of Samsonite Corporation. With a strategy based on an equal-weighted index of New York Stock Exchange equities, its execution was described as “a nightmare.” The strategy was abandoned in 1976, replaced with a market-weighted strategy using the Standard & Poor's 500 Composite Stock Price Index. The first such models were accounts run by Wells Fargo for its own pension fund and for Illinois Bell.

The Wells Fargo group was directly connected to the innovations of academic finance: McQuown was friends with the Fama circle at the University of Chicago, while Fouse knew Sharpe personally and persuaded him to consult for Wells Fargo in the 1970s.⁸

Although academic finance sowed the seeds for the index fund industry, it needed the proper environment to flourish—financial advances do not occur in a technological vacuum. Constructing cash portfolios of broad-based indexes was an extremely difficult

⁷ Bogle (1997).

⁸ Bernstein (1993, pp. 236–248).

and costly task in the 1970s. It is easy to forget the formidable challenges posed by the back-office, accounting, and trade-reconciliation processes for even moderate-sized portfolios in the days before personal computers, FIX engines, and electronic trading platforms, now a distant memory to the managers of today's multitrillion-dollar indexing industry. From the practitioner's perspective, fixing the set of securities and value-weighting them in an index reduced the amount of trading needed to replicate the index in a cash portfolio. Apart from additions and deletions to the index, a portfolio weighted by market capitalization never needs rebalancing since the weights automatically adjusted the proportions as market valuations fluctuated. These "buy-and-hold" portfolios were attractive not only because they kept trading costs to a minimum, but also because they were simpler to implement from an operational perspective. Moreover, weighting stocks in proportion to their relative importance appealed to common sense.

The success of the index mutual fund led to an evolutionary explosion of financial innovation, centered on the concept of the index. Three different stock market index futures debuted in 1982, based on the New York Stock Exchange (NYSE) Composite, the S&P 500, and the Value Line index, respectively. Indexes for each asset class emerged, and additional index funds to track them: the first bond index fund for retail investors in 1986, the first international share index funds in 1990, and the first ETF in 1993. ETFs were similar to index mutual funds in that they closely tracked an index, but could be bought and sold throughout the day on exchanges. These served three broad purposes for the investor: performance indicators; vehicles for direct investment through the use of index funds; and vehicles for hedging, speculation, and investment through the use of derivative instruments based on the indexes.

At the same time, however, the concept of the index itself was evolving. Barr Rosenberg first proposed the notion of a "normal portfolio" in the late 1970s, and implemented it at BARRA in the 1980s.⁹ This was an attempt to construct customized indexes to describe the investment activities of more specialized managers, using a set of securities "weighted as the manager would weight them,"¹⁰ in order to provide insight into their unique risk exposures. This was conceptually expanded in 1990, when Sharpe (1990) defined a new distinction between investment style and investment selection in performance attribution and measurement. Passive fund managers exposed the investor to an investment style, while the active fund manager provided both style and selection. Sharpe reasoned that an actively managed fund's performance should be determined by its selection return, the difference between the fund's return and that of a benchmark with the same style. Sharpe

⁹ Kritzman (1987), Divecha and Grinold (1989), Christopherson (1998).

¹⁰ Christopherson (1998, p. 128).

listed four conditions for a strong style portfolio: it should be “1) a viable alternative, 2) not easily beaten, 3) low in cost, and 4) identifiable before the fact.”

Sharpe emphasized the difference between passive and active management in his exposition. However, with the proliferation of automated trading algorithms, it became clear that the fifth, implicit condition was not necessarily passivity. Benchmark algorithms for high-performance computing blurred the line between passive and active. The key distinction was a lack of human intervention and discretion. If a trading strategy could be codified into a set of transparent rules that gave self-consistent results on similar datasets, as did the benchmark algorithms used to test the machines that implemented them, how did this differ in spirit from an index constructed using a passive portfolio?

Academics and managers both questioned the original assumptions behind the CAPM portfolio. Merton (1973) extended the model intertemporally, while Ross (1976) broke down the CAPM’s beta into a multifactor structure. Some statistical tests suggested that market capitalization was not the most optimal weighting system in CAPM. Goldman Sachs even proposed an earnings-weighted portfolio in the early 1990s. More recently, Arnott, Hsu, and Moore (2005) follow this line of inquiry to create the original set of “fundamental indexes,” portfolios of stocks weighted by book value, cash flow, revenue, sales, dividends, and employment, seeking to capture the value premium.

But perhaps the most obvious illustration of the changing nature of indexes is the proliferation of target-date and life-cycle funds. These funds are designed for specific cohorts of investors sorted by their planned retirement dates, changing their asset mix to become more conservative as they approach their target date.¹¹ Life-cycle funds are not static, but neither are they completely actively managed. Current trading and portfolio-management technology can create passive portfolios capable of capturing complex risk-return profiles that change through time, such as those of an aging population preparing for retirement.

Today, indexes are at the forefront of the building tsunami of financial innovation. Just as the previous financial generation saw “markets in everything,” we currently see “indexes in everything,” as well as funds and derivatives based on those indexes. The technological environment facilitated this efflorescence, but these innovations would never have flourished had the investing public not found indexes useful. What virtues does the modern index have that make it so attractive to investors?

¹¹ Bodie, Merton, and Samuelson (1992), Shiller (2005).

3 What Is An Index?

Ideally, financial form should follow financial function; the proper definition of an index depends on its use. The traditional definition of an index as a market-cap-weighted basket of a fixed set of securities persists today not because of its inherent superiority or economy of implementation, but because its past success has led to inertia in considering other alternatives. To understand how the index evolved, it is fruitful to adopt Merton's functional perspective and ask what functions does an index serve?¹² We can identify at least two distinct functions of a modern index. The first largely informational: indexes provide an aggregate measure of investment performance that abstracts from the vicissitudes of individual components to highlight economy-wide drivers of the market.

The second and more practical function is as a standard against which active managers can be compared, i.e., "This is how you would have performed, if you had invested in this particular passive manner." We see this explicitly in Rosenberg's normal portfolio and Sharpe's style return, as well as in how active managers construct their portfolios subject to "tracking error" constraints relative to some index. For such comparisons to be economically meaningful, the index must be able to serve as the basis for an investment vehicle, a transparent portfolio with a plausible risk/reward profile, as McQuown, Fouse, Bogle, and other indexing pioneers envisioned.

To achieve this second function, we can reverse engineer three fundamental properties this form of index must have. First, it must be transparent, meaning that every aspect of the index must be public information and verifiable by any interested third party. Second, it must be investable, meaning that an investor should be able invest a large amount of capital into the portfolio over a short period of time and realize the return reported by the index. Finally, it must be systematic, meaning that the construction of the index must be rules-based and not dependent on any discretion or human judgment. No alpha should be needed to implement the index in a live portfolio; any investor should be able to do it (subject only to technological constraints).

This more general definition of an index may seem innocuous enough, but it does have several significant implications for how we think about indexes and indexation. For example, this definition excludes certain well-known indexes such as the Case-Shiller Home Price Indexes as well as most hedge-fund indexes (they are not based on liquid instruments and are, therefore, not investable in large size). However, these quantities still play important roles with respect to the first function, even if they are not investable. Moreover, they can often serve as the basis for other financial securities that *are* investable. For example, futures contracts on the Case-Shiller Indexes do trade on the Chicago Mercantile

¹² Merton (1989, 1995a,b) and Merton and Bodie (2005).

Exchange, and real-estate investment trusts and hedge-fund beta replication funds are also offered through investable liquid securities.

Our new definition also includes all the traditional market-cap-weighted indexes where the constituents are liquid securities, which we shall call “static indexes” for the obvious reason. However, the main motivation for redefining an index is to cover the case of “dynamic indexes,” which refers to all portfolios satisfying the three conditions of our new definition but which are not market-cap-weighted and, therefore, require more frequent rebalancing.

This is more than a subtle distinction in semantics. Just as a new formalism in mathematics can sometimes provide a mathematician more insight into proving a theorem, we can gain greater financial insight through grouping functionally similar financial ideas and concepts together, and explore how they may interact in the financial system. Passive market-cap-weighted portfolios are the simplest form of index in common use, but it is clear how variations on this theme readily emerge from their conceptual potential. New trading technology has now given us the ability to create indexes that are not necessarily market-cap-weighted, that are not necessarily even “passive” in the traditional sense. In the next section, we discuss one of these new variations: the strategy index.

4 The Brave New World of Strategy Indexes

The strategy index is a dynamic index that embodies a particular investment strategy. Many of these dynamic indexes are not typically considered strategies, such as life-cycle or target-date funds. However, consider the simplest version of such a fund, one that implements the “100-minus age” rule of thumb. It is entirely transparent—maintain a percentage equivalent to 100 minus the investor’s age in a broad-based index fund, and the remainder in bonds, and update this portfolio yearly. Its components are fully investable, and it is systematic in the sense that it requires no discretionary intervention to update it. Much more complicated strategies, like those underlying 130/30 funds, have been codified into dynamic indexes.¹³ Meanwhile, dynamic indexes for previously esoteric hedge-fund strategies like merger arbitrage are now available to the average investor through rules-based portfolios that invest in publicly announced takeovers of certain pre-defined characteristics.

The theoretical underpinnings for these dynamic indexes are straightforward enough, and flow from variations on the original CAPM formulation. For example, CAPM can be generalized to multiple factors, such as Merton’s Intertemporal CAPM or Ross’s Arbitrage Pricing Theory. Meanwhile, the efficient frontier can be estimated using weightings other

¹³ Lo and Patel (2008).

than market capitalization, as with Arnott, Hsu, and Moore's fundamental indexes, equal-weight indexes seeking to capture small-cap premia, or low-volatility indexes to capture low-volatility premia. In this brave new multifactor, algorithmic world, nearly any plausible strategy can be broken down into components of investment style, weighting, and other conditions. In fact, the burgeoning literature and industry applications involving hedge-fund beta replication takes this observation to its logical conclusion: if a hedge-fund strategy's returns contains common factors that can be "cloned," i.e., identified, quantified, and replicated using liquid futures contracts without the need for active management, why not use them as benchmarks for comparison?¹⁴

However, the key question for investors is whether a strategy index carries a sustainable risk premium, and if so, under what conditions. It is here that the other financial theory behind the modern index fund, the EMH, is relevant. The EMH implies that no investor should be able to generate a consistent return in the market above the risk-return relationship defined by the CAPM or a similar equilibrium asset-pricing model. Any sustainable risk premium above this benchmark should be arbitrated away by investors in pursuit of profit. However, there is compelling evidence that the EMH is only the limiting case of a more complex reality. The Adaptive Markets Hypothesis (AMH) suggests that a sustainable risk premium may be available to investors for a period of time, given the financial environment and the population history of the market.¹⁵

As an illustration, the AMH explains that behavioral biases are likely responsible for many market anomalies, and therefore are a possible source of risk premia for dynamic indexes meant to exploit such anomalies. A naïve critique of this possibility is that behavioral biases are often correctable: if one points out a behavioral bias to an individual, he or she will often stop exhibiting such a bias. But the AMH is a hypothesis about the dynamics of the marketplace, not the statics. In fact, the "discipline of the market" should punish investors with this bias until they exit the market, or adaptively change their strategy. If a behavior is innate, however, and there is a flow of new investors coming into the market, then a behavioral bias premium may be sustainable. As P.T. Barnum almost said, "There's a new investor born every minute."

While the question of sustainability is *prima facie* important to investors, the ultimate sources of expected return—risk premia or alpha—lie at the heart of this issue. Competition suggests that alpha should be capacity-constrained, hard to come by, and expensive. In theory, alpha is either competed away to nothing, or it becomes commoditized to some level where the returns are just enough to compensate investors for the risks associated with the activity, i.e., beta, which should be less constrained, easy to

¹⁴ Hasanhodzic and Lo (2007).

¹⁵ Lo (2004).

come by, and cheap. The dynamic properties of these risk factors and their expected returns require a framework other than the static assumptions of the EMH to interpret.

5 Disbanding the Alpha Beta Sigma Fraternity

The CAPM broke important new ground in teaching investors how to distinguish between unique investment acumen that justifies active management fees versus commoditized risk premia that can be captured much more cheaply. But this dichotomy treats risk in a very rigid manner: active strategies manage risk actively, and passive strategies do not manage risk at all. The reason for this distinction is largely historical—as described in Section 2, passive investing became synonymous with market-cap-weighted indexes to minimize the amount of trading needed to manage an index portfolio. Over time, the faithful reproduction of an index's returns by a portfolio of securities has become the over-riding concern of an index manager, irrespective of the rollercoaster ride that the index imposes on its investors. A more cynical perspective is that misery loves company; an index manager will not be punished for suffering losses if all index funds experience similar losses.

While it may be the case that the manager may not be punished, the investor is not so fortunate. On October 24, 2008 when the S&P 500 volatility reached a record level of 89.53 as measured by the VIX index, passive investors in this benchmark were exposed to extraordinary amounts of risk that they surely did not intend to take. At an annualized volatility of 89%, the probability of loss is 58.7% assuming lognormally distributed index returns with an annualized mean of 10%, and the probability of losing 25% or more is 42.7% (see Figure 1).

One measure of extreme loss commonly used in the hedge-fund industry is the maximum drawdown (MDD), defined to be the largest percentage decline in the net asset value of a fund over any investment period over the fund's entire history. The MDD of the S&P 500 index from 2007 to the present occurred between October 9, 2007 and March 9, 2009, when the index declined by 56.8%, a loss not easily absorbed by any investor, especially investors who aren't monitoring their portfolios on a regular basis.

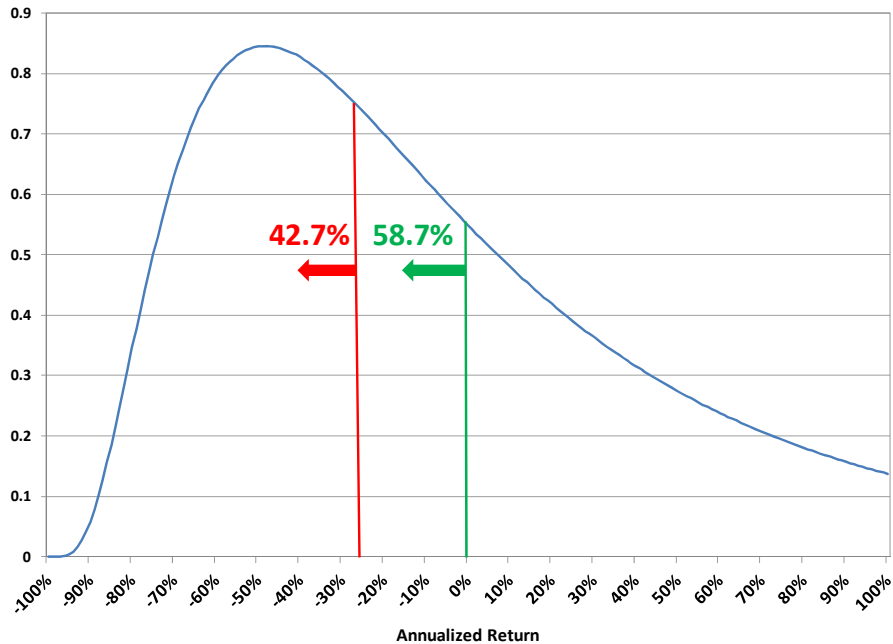


Figure 1. Lognormal distribution for asset returns, assuming that simple (not continuously compounded) returns have an annualized mean and standard deviation of 10% and 89%, respectively. The probability of a negative realization is 58.7% and the probability of a loss greater than -25% is 42.7%.

Traditional passive investments are especially problematic when volatility can change dramatically, as it has over the last decade. Figure 2 illustrates the dynamic nature of volatility as measured by the daily closing values of the VIX index from January 2, 1990 to March 20, 2015. Also plotted are the probabilities of a greater-than-25% loss assuming an annualized expected return of 10% and an annualized volatility equal to the VIX. Portfolio managers and financial advisers often chide retail investors for not focusing on the long run and being too sensitive to short-term fluctuations—is it reasonable to expect any rational investor to stay the course in an investment that can lose more than half its value over a 16-month period and experience an increase in volatility by a factor of 4 over 63 trading days?¹⁶

This lack of risk management—the fact that smart beta can be accompanied by dumb sigma—is perhaps the greatest weakness of traditional passive investing. When market volatility is relatively stable, risk management may not be as important for static index products. Therefore, during the “Great Modulation” (Lo, 2012)—the period from the 1930s to the early 2000s when U.S. stock market volatility was both low and relatively stable—it

¹⁶ In 2008, the VIX increased from 18.81 on August 22 to 80.86 on November 20.

is not surprising that static index funds did an admirable job of allowing investors to capture the equity risk premium and manage their overall risk exposures by adjusting their asset allocation between the stock index, bonds, and cash. But when the volatility of volatility becomes significant, as it has over the past decade, forgoing risk management can be devastating to investors.

However, this is no longer a necessary evil of indexation thanks to the many technological advances in algorithmic trading, securities exchanges, derivatives, telecommunications, and back-office and accounting systems infrastructure. What took McQuown and Fouse at Wells Fargo a month to implement in the 1970s can now be done almost instantaneously, and at much lower cost. Moreover, one of the most mind-numbing aspects of professional portfolio management—monitoring the portfolio in real time and deciding when to act in response to rapidly deteriorating market conditions—can be automated to a significant degree, especially for passive strategies that are dedicated to achieving the returns of an index. The traditional pairing of active risk management with active investing and passive risk management with passive investing can be severed.

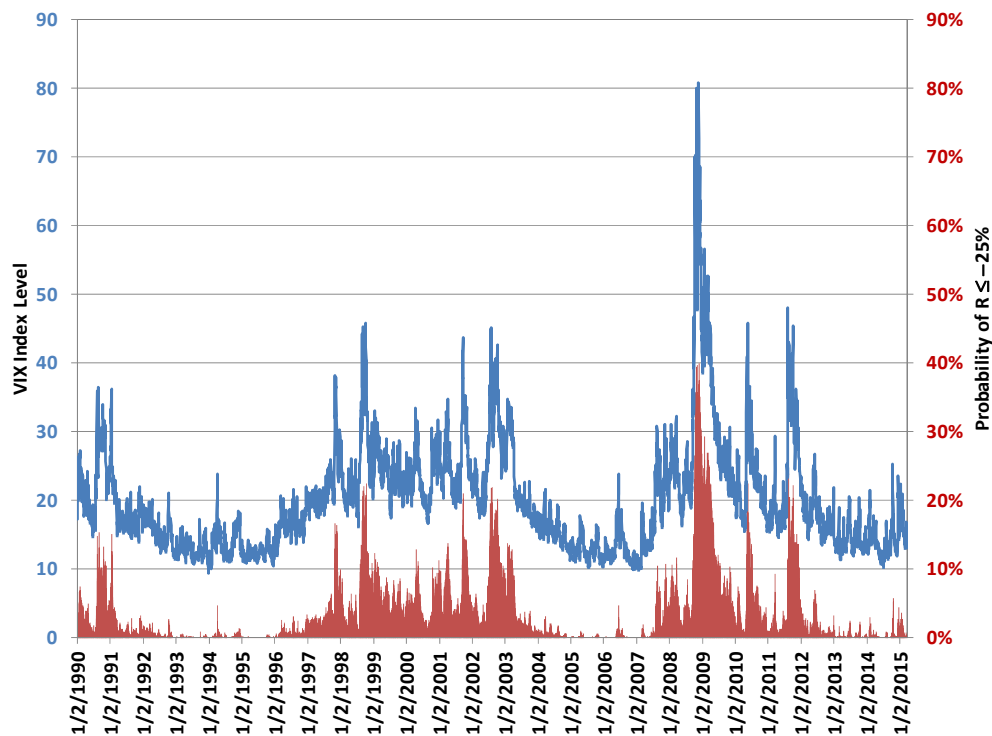


Figure 2. Daily closing values of the VIX index from January 2, 1990 to March 20, 2015 and loss probabilities of greater than or equal to -25% assuming lognormally distributed returns with annualized expected simple return of 10% and volatility equal to the VIX. Source of VIX Index: <http://finance.yahoo.com> (accessed March 22, 2015).

One simple example of how to sever this link is to create a dynamic index fund that contains no alpha but is actively risk-managed to a target level of volatility σ_o . This can be accomplished by investing a portion of the fund in cash if the estimated volatility $\hat{\sigma}_{t-q}$ of the index at date t exceeds σ_o and investing more than 100% of the fund in the index, i.e., leveraging the fund, if $\hat{\sigma}_{t-q}$ falls below σ_o :

$$\tilde{R}_t = \kappa_t R_t \quad , \quad \kappa_t = \text{Min} \left[\frac{\sigma_o}{\hat{\sigma}_{t-q}} , \bar{l} \right] \quad , \quad \hat{\sigma}_{t-q}^2 = \frac{1}{n-1} \sum_{j=q}^{q+k-1} (R_{t-j} - \hat{\mu}_{t-q})^2 \quad , \quad q , \bar{l} \geq 1 \quad (1)$$

where $\hat{\mu}_{t-q}$ is the rolling-window mean return between $t - q - k + 1$ and $t - q$, and \bar{l} is some fixed upper bound on the amount of leverage that can be used. By setting the leveraging/deleveraging factor to be the ratio of σ_o to $\hat{\sigma}_{t-q}$, and if $\hat{\sigma}_{t-q}$ is a reasonably accurate measure of short-term volatility, this algorithm will yield returns with volatility closer to the target σ_o than the static index. And if σ_o is chosen with an investor's risk tolerance in mind, the actively risk-managed index \tilde{R}_t will be more palatable as a long-term investment than the static index. Such a portfolio is "passive" in that there is no alpha, but it is actively risk-managed and not value-weighted.

To see how this volatility-control algorithm might work in practice, we apply it to daily CRSP value-weighted index returns from 1925 to 2014 using a 21-day rolling-window volatility estimator with one lag, i.e., $\hat{\sigma}_{t-1}$, and a value of 16.9% for the annualized volatility target level σ_o , which is the unconditional volatility of the CRSP value-weighted index returns over the entire sample period. Figure 3 contains a comparison of the volatility of the raw index (in blue) and the volatility-controlled index (in red), where 125-day rolling windows are used to estimate these volatilities (we use a longer window for this comparison to show that the volatility control does have an effect even outside of the 21-day window used to scale the portfolio). A comparison of the two plots confirms the intuition that dynamically adjusting the portfolio as a function of short-term volatility does create a substantially less volatile time series of returns.

However, this stability comes at a cost. Scaling the portfolio on a daily basis requires monitoring short-term volatility and active risk management, i.e., adjusting the portfolio's exposure either by trading the index constituents or, more likely, by implementing a futures or forward contract overlay strategy to dynamically scale the index exposure up or down. The green bar graph in Figure 3 displays the amount of scaling κ_t involved; the straight horizontal red line is set at a value of 1 for the scale factor. For much of the time, κ_t is at its upper bound of 1.3, implying that most of the time short-term annualized U.S. equity market volatility is less than $16.9/1.3 = 13\%$ and the fund is 130% invested in the market. However, occasionally κ_t falls below the red line, indicating that short-term volatility has exceeded the target level of 16.9% and a portion of the portfolio is switched into cash. We assume that the cash earns the yield on one-month U.S. Treasury bills, and

that all changes in portfolio weights incur transactions costs of 0.05% or 5 basis points (bps) of the trade size. For the S&P 500 index, implementing the dynamic index (1) using the Chicago Mercantile Exchange's E-Mini S&P 500 Futures contract would yield considerably lower transactions costs than 5 bps.¹⁷

By actively managing the risk of the fund, this algorithm is reflecting the behavioral predilections of the typical investor—reducing market exposure when risk becomes too high and restoring it when risk returns to normal—but doing so more systematically and at a higher frequency than all but the most active traders can manage. As a result, investors are more likely to stay invested in this strategy rather than exiting after a large loss and waiting too long before reinvesting. Table 1 shows that staying invested in this fund is rewarded: \$1 invested in 1926 becomes \$11,141 in 2014, which compares favorably with the \$4,162 from the unmanaged index. More importantly, the MDD of the risk-managed strategy over this 89-year period is –72% which is severe, but less so than the –84% of the raw index. The difference in excess kurtosis—a measure of the likelihood of “tail” events—also points to a substantial reduction in risk: 4.85 for the risk-managed fund vs. 16.87 for the raw index according to Table 1 (for comparison, the excess kurtosis for a normal distribution is 0).

¹⁷ The contract value for the E-Mini S&P 500 is 50 times the index value, hence an index level of 2,000 yields a contract value of \$100,000. The bid/offer spread for this contract is typically 1 tick which is \$12.50 per contract, hence the one-way cost can be approximated by half this amount or 0.625 basis points (bps). Additional fees for the E-Mini S&P 500 (commission, NFA, exchange, etc.) range from \$1.87 to \$2.46 per contract, depending on the method of execution, which amounts to 0.221 bps on average, hence the total cost of executing a single contract is slightly less than 1 bp as of March 30, 2015.

	Raw Market Returns		Returns w/ VolControl		Stats for VolControl κ			Raw Market Returns		Returns w/ VolControl		Stats for VolControl κ	
1926 to 2014 (Entire Sample)							2010 to 2014 (Most Recent 5 Years)						
Mean	9.26%	10.41%	Mean	1.14	Mean	14.1%	14.5%	Mean	1.13				
SD	16.86%	14.94%	SD	0.26	SD	16.2%	15.7%	SD	0.25				
Sharpe	0.36	0.48	Min	0.17	Sharpe	0.87	0.92	Min	0.34				
Skew	-0.12	-0.54	Median	1.30	Skew	-0.40	-0.51	Median	1.30				
Kurt	16.87	4.85	Max	1.30	Kurt	4.53	1.80	Max	1.30				
MDD	-84%	-72%			MDD	-21%	-21%						
CumRet	\$4,162	\$11,141			CumRet	\$1.94	\$1.98						
1926 to 1935 (First 10 Years)							2012 to 2014 (Most Recent 3 Years)						
Mean	3.5%	5.0%	Mean	0.92	Mean	18.5%	20.3%	Mean	1.24				
SD	26.8%	17.2%	SD	0.36	SD	11.8%	14.5%	SD	0.11				
Sharpe	0.06	0.18	Min	0.17	Sharpe	1.57	1.39	Min	0.81				
Skew	0.44	-0.15	Median	0.96	Skew	-0.26	-0.32	Median	1.30				
Kurt	9.64	3.14	Max	1.30	Kurt	1.09	1.17	Max	1.30				
MDD	-84%	-72%			MDD	-10%	-13%						
CumRet	\$1.51	\$1.77			CumRet	\$1.67	\$1.74						
2005 to 2014 (Most Recent 10 Years)							2014 (Most Recent 1 Year)						
Mean	7.7%	9.2%	Mean	1.07	Mean	10.5%	9.4%	Mean	1.25				
SD	20.4%	15.7%	SD	0.30	SD	11.3%	14.1%	SD	0.10				
Sharpe	0.31	0.50	Min	0.20	Sharpe	0.93	0.66	Min	0.91				
Skew	-0.18	-0.44	Median	1.26	Skew	-0.45	-0.57	Median	1.30				
Kurt	10.09	1.37	Max	1.30	Kurt	1.19	1.39	Max	1.30				
MDD	-55%	-38%			MDD	-8%	-10%						
CumRet	\$2.11	\$2.42			CumRet	\$1.11	\$1.09						

Table 1. Summary statistics for volatility control mechanism applied to daily CRSP value-weighted index returns from January 25, 1926 to December 31, 2014, and for selected subperiods. The volatility control mechanism multiplies raw returns by a scale factor $\kappa_t = \text{Min} \left[\frac{16.9\%}{\hat{\sigma}_{t-1}}, 1.3 \right]$, where $\hat{\sigma}_{t-1}$ is a 21-day rolling-window annualized volatility estimator, and 16.9% is the annualized volatility of the index returns over the entire sample.

This volatility-controlled dynamic strategy is reminiscent of portfolio insurance, except that the objective here is simpler: to maintain a more consistent level of volatility so as to avoid triggering panic selling by investors. Typical portfolio insurance strategies such as Black and Pérold (1992) involve the dynamic replication of a put option on the portfolio's value, which involves reducing equity exposure as the value of equity declines. In the volatility control mechanism (1), equity exposure is reduced in response to increasing short-term volatility, not because of market direction. However, because stock prices and volatility are negatively correlated,¹⁸ a strategy that reduces market exposure in response

¹⁸ Black (1976).

to increasing volatility will, on average, be reducing equity exposure during declining markets and increasing equity exposure during rising markets.

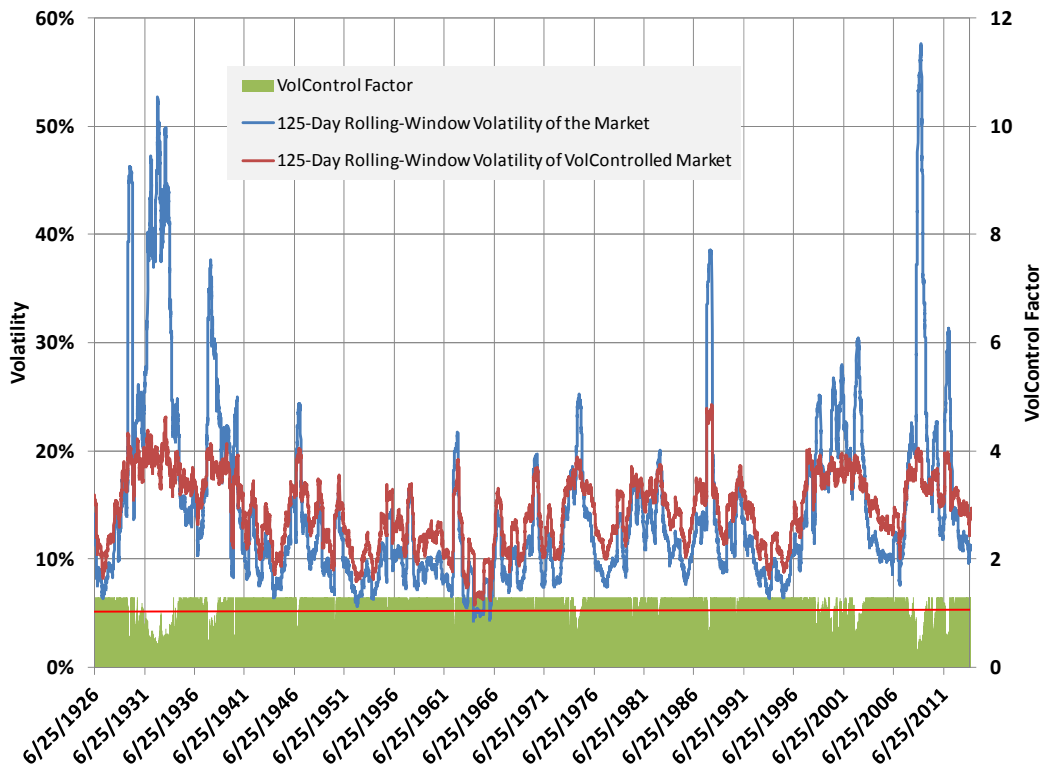


Figure 3. 125-day rolling-window annualized volatility estimates for CRSP daily value-weighted index returns from June 25, 1926 to December 31, 2014 with and without volatility control using a scale factor $\kappa_t = \text{Min} \left[\frac{16.9\%}{\hat{\sigma}_{t-1}}, 1.3 \right]$, where $\hat{\sigma}_{t-1}$ is a 21-day rolling-window annualized volatility estimator.

If the stock-price/volatility relation is persistent, such a volatility control mechanism may actually add to overall performance rather than subtracting from it due to being partially invested in cash during periods when volatility exceeds the target. This intuition is confirmed in Figure 4, which contains a comparison of the cumulative return of a \$1 investment in the CRSP value-weighted index and the volatility-controlled index. Over the 89-year period, the volatility-controlled index is the winner by a factor of 4. By reducing equity exposure when volatility is high, the risk-managed benchmark holds more cash when the equity risk premium is lower than average, and it holds more equity when the equity risk premium is higher than average, thereby exploiting the inverse relationship between stock prices and volatility documented by Black (1976) over four decades ago.

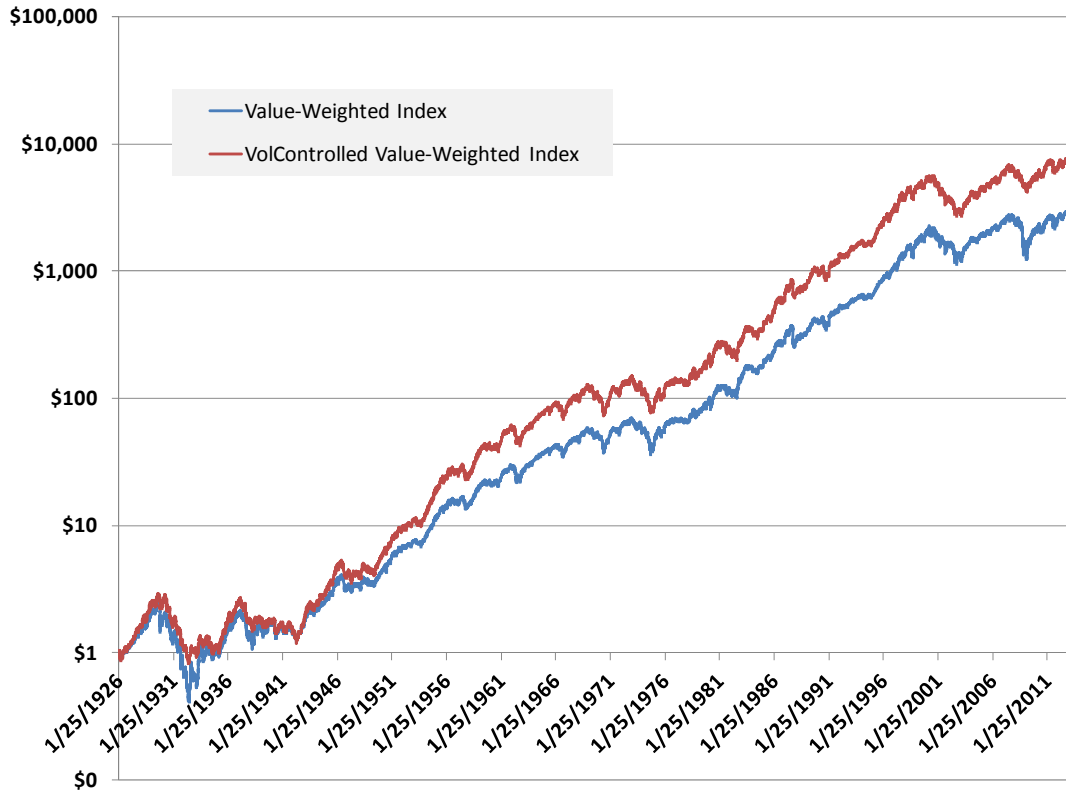


Figure 4. Cumulative return of the CRSP value-weighted index with and without volatility control from January 25, 1926 to December 31, 2014 (logarithmic scale). The volatility control consists of multiplying daily returns by a scale factor $\kappa_t = \text{Min} \left[\frac{16.9\%}{\hat{\sigma}_{t-1}}, 1.3 \right]$, where $\hat{\sigma}_{t-1}$ is a 21-day rolling-window annualized volatility estimator of the value-weighted index return.

This simple example illustrates the potential benefits of separating active risk management from active investment management—one need not be tied to the other, given current trading technology, algorithmic overlay strategies, and a wide spectrum of liquid index futures contracts. Moreover, there are many ways to improve upon the volatility control mechanism (1), leading to a host of new financial products and services that can be tailored to the unique circumstances of each individual investor. This customization process is limited only by the imagination of the portfolio manager and financial adviser, given the powerful trading and portfolio-optimization tools now at our disposal. By applying active risk management overlays to static indexes, we can begin to harvest the benefits of smart beta without also having to suffer the consequences of dumb sigma.

6 Beware of Backtest Bias

The broader definition of an index that technology now supports has spawned a host of financial innovations that would have been impossible a decade ago. At the same time, it has also created some daunting new challenges for investors. Choosing among the dizzying array of financial products available today requires greater education and training—even for professional financial advisers—and the typical retail investor may not be fully equipped to evaluate the potential risks and rewards of the various options with which they are bombarded (consider, for example, the sometimes counterintuitive behavior of double-leveraged inverse S&P 500 ETFs). Moreover, the same technological advances that have brought us smarter betas also allow us to destroy wealth at the speed of light, as the shareholders of Knight Capital Group discovered on August 1, 2012.¹⁹

However, the single biggest challenge created by the smart beta revolution is the potential for misleading investors and portfolio managers through backtest bias. The problem is simple to state but devilishly difficult to address; in fact, it can be argued that backtest bias is an unavoidable aspect of any investment process. Suppose we wish to select the best of n investment opportunities, and each opportunity i can be evaluated by a single summary statistic θ_i such as its Sharpe ratio. Because θ_i is not directly observable, we must estimate it using whatever methods of evaluation we have at our disposal (including both qualitative and quantitative information), ultimately yielding estimators $\{\hat{\theta}_i = \theta_i + \epsilon_i, i = 1, \dots, n\}$, where ϵ_i is the estimation error associated with our evaluation of investment i .

We would like to choose the investment i^* with the highest Sharpe ratio $\theta_{i^*} = \max_i \theta_i$, but because we only observe the estimated Sharpe ratios, it is tempting to select the investment i' with the highest estimated Sharpe ratio $\hat{\theta}_{i'} = \max_i \hat{\theta}_i$, and herein lies the problem. By selecting on the basis of imperfect estimates, we may be confounding genuine investment performance with random estimation error. In other words, by selecting the investment with the biggest $\hat{\theta}_i$, we hope to be getting the biggest θ_i but we may, in fact, be getting the biggest ϵ_i instead. And given the inherent noisiness of even the best investment performance evaluation methods, it is impossible to completely eliminate ϵ_i .

However, there are several reasons why backtest bias is especially problematic for the burgeoning smart beta industry. These reasons are related to the key drivers of backtest bias—other things equal, backtest bias becomes more severe as: (1) the number n of managers/models/track records grows; (2) the signal-to-noise ratio $\text{Var}[\theta]/\text{Var}[\epsilon]$ declines; and (3) decisions become more dependent on simulated performance statistics

¹⁹ A programming error caused Knight Capital Group's automated trading system to enter into 4 million trades over a period of 45 minutes, resulting in a \$440 million loss for the firm that ultimately led to its demise and acquisition by Getco, Inc.

rather than live track records. While investors routinely face one or two of these issues in evaluating any investment opportunity, all three issues arise in the case of dynamic indexes. The number of new products is growing rapidly, and because, by definition, new products do not have live track records, estimates of their performance can only be based on simulated returns and are, therefore, noisier than for more-established products. And because simulations are, for many of these new products, the only way investors can develop intuition for the products' risk/return profiles, decisions tend to rely much more heavily on biased performance statistics.

To illustrate the subtleties involved in evaluating simulated returns, consider the volatility control strategy proposed in Section 5 and suppose that we make one small change in the algorithm (1):

$$\tilde{R}_t = \kappa_t R_t \quad , \quad \kappa_t = \text{Min} \left[\frac{\sigma_0}{\hat{\sigma}_{t+2}} , \bar{l} \right] \quad , \quad \hat{\sigma}_{t+2}^2 = \frac{1}{n-1} \sum_{j=0}^{k-1} (R_{t+2-j} - \hat{\mu}_{t+2})^2 \quad , \quad \bar{l} \geq 1 \quad (2)$$

How does (2) differ from (1)? The only change is that the short-term volatility estimator $\hat{\sigma}_{t+2}^2$ in (2) now uses returns from $t - k + 3$ to $t + 2$ instead of from $t - k$ to $t - 1$. The simulated performance of this version, summarized in Table 2, is considerably better than that of strategy (1): over the entire 89-year sample, the compound annual return is higher, the volatility is lower, hence the Sharpe ratio is higher, and the cumulative return of a \$1 investment is \$73,626, nearly seven times the amount generated by (1), and over 17 times the cumulative return of the raw index.

	Raw Market Returns	Lagged Volatility Estimator		Biased Volatility Estimator	
		Returns w/ VolControl	Stats for VolControl κ	Returns w/ VolControl	Stats for VolControl κ
Mean	9.26%	10.41%	Mean 1.14	12.65%	Mean 1.15
SD	16.86%	14.94%	SD 0.26	13.78%	SD 0.26
Sharpe	0.36	0.48	Min 0.17	0.68	Min 0.14
Skew	-0.12	-0.54	Median 1.30	-0.21	Median 1.30
Kurt	16.87	4.85	Max 1.30	0.56	Max 1.30
MDD	-0.84	-0.72		-0.72	
CumRet	\$4,162	\$11,141		\$73,626	

Table 2. Summary statistics for look-ahead-biased volatility control mechanism applied to daily CRSP value-weighted index returns from January 25, 1926 to December 31, 2014. The volatility control mechanism multiplies raw returns by a scale factor $\kappa_t = \text{Min} \left[\frac{15\%}{\hat{\sigma}_{t+2}} , 1.3 \right]$, where $\hat{\sigma}_{t+2}$ is a 21-day rolling-window annualized

volatility estimator that uses contemporaneous and two consecutive future daily returns to estimate short-term volatility.

How is this possible? The seemingly minor change in (2) introduced “look-ahead” bias into the simulation by using contemporaneous returns to scale the portfolio’s leverage. In the event of a large one-day decline in the market on dates t to $t + 2$, the short-term volatility indicator will now increase, causing a decline in the scaling factor κ_t that multiplies the date- t return. This, of course, leads to improved performance, but it is completely spurious—in practice, we cannot scale down a losing position prior to incurring the loss, which is what (2) implicitly assumes. To underscore the impact that this kind of bias can have, if we had used 5 days instead of 21 to compute volatility—which increases the relative importance of the contemporaneous and future returns in the short-term volatility estimator—a \$1 investment in 1926 becomes \$173,444 by 2014, more than 40 times the cumulative return of the value-weighted index.

Although this particular kind of bias is easily prevented by an experienced quantitative portfolio manager, it may not be so easily detected by the investor. And even if this bias is avoided, many others can arise through the process of selecting “optimal” parameters for the strategy. For example, even the simple volatility control mechanism (1) has several parameters: the rolling-window length k , the leverage upper bound \bar{l} , the number of lags q for the short-term volatility estimator $\hat{\sigma}_{t-q}$, and the transactions cost τ . If we consider optimizing this strategy over the parameter space of plausible values for each of these parameters—say $k = 2, \dots, 125$, $\bar{l} = 1.00, 1.05, 1.10, \dots, 2.00$, $q = 1, \dots, 10$, $\tau = 1, \dots, 10\text{bps}$ —we will be selecting the “best” among $124 \times 21 \times 10 \times 10 = 260,400$ models. Including other desirable features such as a lower limit on leverage, a stop-loss policy, and a turnover constraint to control trading costs can yield a combinatorial explosion in the number of models. Searching over these many versions is likely to impart significant backtest bias, especially if the underlying “signal” of genuine performance is small relative to the statistical errors in estimating performance.

Because it is so easy to construct an attractive investment product on paper with the benefit of hindsight, and because such products rarely perform as well when they are implemented, regulators have developed strict rules to address this issue. Rule 206(4)-1 of the 1940 Investment Adviser Act states that:

It shall constitute a fraudulent, deceptive, or manipulative act, practice, or course of business...for any investment adviser registered or required to be registered under [the Investment Adviser Act], directly or indirectly, to publish, circulate, or distribute any advertisement which refers, directly or indirectly, to any testimonial of any kind concerning the investment adviser

or concerning any advice, analysis, report or other service rendered by such investment adviser.

Although the term “testimonial” is never defined, the SEC clearly had in mind the pitfalls of backtest bias when it adopted this rule in 1961 because the commissioners wrote “...such advertisements are misleading; by their very nature they emphasize the comments and activities favorable to the investment adviser and ignore those which are unfavorable.”²⁰ The SEC routinely enforces this rule through cease-and-desist orders, criminal prosecutions, and severe financial penalties against unscrupulous managers using misleading simulations in their marketing efforts. For example, in 2012 the SEC issued a cease-and-desist order against a nationally syndicated radio personality and financial advice author who had been touting his “Buckets of Money” wealth management strategy in seminars he hosted for potential clients. According to Michele Wein Layne, Regional Director of the SEC’s Los Angeles field office, “[the manager and his advisory firm] left their seminar attendees with a false sense of comfort about the Buckets of Money strategy...The so-called backtests weren’t really backtests, and the strategy wasn’t proven as they claimed.”

However, even the SEC has acknowledged the necessity and value of simulating the performance of various portfolio strategies, and has provided guidance in the form of “no-action letters” that describe cases in which the SEC will not pursue enforcement action. The 1986 no-action letter in response to a request by Clover Capital Management is the most relevant for the use of simulated returns in marketing materials.²¹ This letter states:²²

The staff no longer takes the position, as it did a number of years ago, that the use of model or actual results in an advertisement is per se fraudulent under Section 206(4) and the rules thereunder, particularly Rule 206(4)-1(a)(5). Rather, this determination is one of fact, and we believe the use of model or actual results in an advertisement would be false or misleading under Rule 206(4)-1(a)(5) if it implies, or a reader would infer from it, something about the adviser's competence or about future investment results that would not be true had the advertisement included all material facts. Any adviser using such an advertisement must ensure that the advertisement discloses all material facts concerning the model or actual results so as to avoid these unwarranted implications or inferences.

The letter then provides a number of specific examples of inappropriate practices, which includes the advertising of simulated results that: do not deduct trading costs and fees,

²⁰ SEC Investment Advisers Act Rel. No. 121 (Nov. 2, 1961) (adopting rule 206(4)-1).

²¹ SEC (1986).

²² Original footnotes have been deleted.

highlight the potential for gains without also mentioning the risk of loss, fail to disclose material assumptions underlying the results and the inherent limitations of simulated returns, etc. In short, the use of backtest results must satisfy the same standards of accuracy and disclosure as the use of live track records does, and managers have an affirmative obligation to avoid any false or misleading statements about their simulations.

Seasoned investment professionals have long been aware of backtest bias and have learned to deal with it in several ways. The first and most obvious way is to treat all investment performance records with a healthy dose of skepticism, and acknowledge that even a stellar track record contains some element of sheer dumb luck. How much of historical success is luck versus skill is another way of asking how much of $\hat{\theta}_i$ is θ_i and how much is ϵ_i ?

The second way is to use additional information to distinguish θ_i from ϵ_i . For example, if a manager claims to be a talented stock picker, we can check whether this manager's stock-picking success materially changed during bear markets; if so, then perhaps the manager's "skill" is more beta than alpha. If, on the other hand, the manager's stated strength is asset allocation, a standard performance attribution analysis will allow us to verify this claim by separating the manager's cumulative returns into market and asset-class timing, security selection, and other sources of value-added.

The third, and most direct, way to distinguish between θ_i and ϵ_i is to conduct live out-of-sample experiments. Decide today to follow the manager's performance over the course of the next year or two and evaluate the manager at the end of that period. By collecting new data on the manager's performance that is statistically independent of the past, backtest bias is minimized. And if the manager's out-of-sample record is comparable to the backtest, then $\hat{\theta}_i$ may be more θ_i than ϵ_i .

In each of these three approaches, we are seeking additional information that can confirm the link between $\hat{\theta}_i$ and θ_i , and if we can't find such information, then the more likely explanation for attractive historical performance $\hat{\theta}_i$ is good luck, ϵ_i .

Simulations do play an important role in developing a deeper understanding of new investment products and should not be shunned. Although backtest bias is an unavoidable aspect of financial innovation, its impact can be reduced by practicing good "statistical hygiene" in generating and interpreting backtests.

7 Conclusion

A confluence of technological advances has caused tectonic shifts in the financial landscape, creating winners and losers overnight. The winners are technology-savvy investors who understand their own risk preferences and financial objectives, and can appreciate the full

spectrum of risks and rewards offered by today's dizzying array of smart-beta and index products. The losers are the technophobes and Luddites who don't know and don't care about investing—the investment ecosystem has become much more dangerous for them.

The traditional advice of “equities in the long run” and “buy and hold” worked well during the relative calm of the Great Modulation, and the equity risk premium was remarkably consistent and positive over this period. However, the same advice may not be as effective in the current environment of seesawing volatility and intense financial innovation. There are a great many benefits of dynamic indexes—more sources of diversification and risk-sharing, cheaper ways to meet individual needs, and greater flexibility in reflecting investment views—but they come at a cost of potential pitfalls if abused. Hand saws are not nearly as useful as chain saws in clearing downed trees after a hurricane, but hand-saw accidents are not nearly as dangerous as chain-saw accidents. If investors do not adapt, but persist in treating modern financial tools as if they behaved like their less powerful predecessors, a series of chain-saw accidents will be the outcome. To Keynes' adage that “in the long run, we are all dead,” we should append the further imperative to “make sure the short run doesn't kill us first.”

References

- Bernstein, Peter. 1993. *Capital Ideas*. New York: Free Press.
- Black, Fischer. 1976. "Studies of Stock Price Volatility Changes," *Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economical Statistics Section*, 177–181.
- Black, Fischer and André F. Perold. 1992. "Theory of Constant Proportion Portfolio Insurance." *Journal of Economic Dynamics and Control* 16: 403–426.
- Bodie, Zvi, Robert C. Merton, and William F. Samuelson. 1992. "Labor Supply Flexibility and Portfolio Choice in a Life-Cycle Model." *Journal of Economic Dynamics and Control* 16, 427–449.
- Bogle, John C. 1997. "The First Index Mutual Fund: A History of Vanguard Index Trust and the Vanguard Index Strategy." electronic copy available at http://www.vanguard.com/bogle_site/bogle_lib.html#1997
- Christopherson, Jon A. 1998. "Normal Portfolios: Construction of Customized Benchmarks." In *Selected Topics in Equity Portfolio Management*, edited by F.J. Fabozzi. New Hope, PA: Frank J. Fabozzi Associates.
- Divecha, Arjun and Richard C. Grinold. 1989. "Normal Portfolios: Issues for Sponsors, Managers, and Consultants." *Financial Analysts Journal* 45, 7–13.
- Evans, Judith. 2015. "Smart Beta is No Guarantee You Will Beat the Market." *Financial Times*, February 1. <http://www.ft.com/intl/cms/s/0/2d00969c-a32f-11e4-9c06-00144feab7de.html>
- Fama, Eugene. 1965. "The Behavior of Stock Market Prices." *Journal of Business* 38: 34–105.
- Fox, Justin. 2011. *The Myth of the Rational Market: A History of Risk, Reward, and Delusion on Wall Street*. New York: Harper Business.
- Hasanhodzic, Jasmina and Andrew W. Lo. 2007. "Can Hedge-Fund Returns Be Replicated? The Linear Case." *Journal of Investment Management* 5, 5–45.
- Kritzman, Mark. 1987. "How To Build A Normal Portfolio In Three Easy Steps." *Journal of Portfolio Management* 14, 21–23.
- Lo, Andrew W. 2004. "The Adaptive Markets Hypothesis: Market Efficiency from an Evolutionary Perspective." *Journal of Portfolio Management* 30, 15–29.
- Lo, Andrew W., 2012, "Adaptive Markets and the New World Order." *Financial Analysts Journal* 68, 18–29.
- Lo, Andrew W. and Pankaj Patel. 2008. "130/30: The New Long-Only." *Journal of Portfolio Management* 34, 12–38.
- Merton, Robert C. 1973. "An Intertemporal Capital Asset Pricing Model." *Econometrica* 41, 867–887.

- Merton, Robert C. 1989. "On the Application of the Continuous-Time Theory of Finance to Financial Intermediation and Insurance." *Geneva Papers on Risk and Insurance* 14, 225–262.
- Merton, Robert C. 1995a. "Financial Innovation and the Management and Regulation of Financial Institutions." *Journal of Banking and Finance* 19, 461–481.
- Merton, Robert C. 1995b. "A Functional Perspective of Financial Intermediation." *Financial Management* 24, 23–41.
- Merton, Robert C. and Zvi Bodie. 2005. "Design of Financial Systems: Towards A Synthesis of Function and Structure." *Journal of Investment Management* 3, 1–23.
- Renshaw, Edward F. and Paul J. Feldstein. 1960. "The Case for an Unmanaged Investment Company." *Financial Analysts Journal* 16, 43-46.
- Ross, Stephen. 1976. "The arbitrage theory of capital asset pricing." *Journal of Economic Theory* 13, 341–360.
- Samuelson, Paul A. 1965. "Proof That Properly Anticipated Prices Fluctuate Randomly.:" *Industrial Management Review* 6, 41–49.
- Securities and Exchange Commission, 1986, *SEC Staff No-Action Letter, Clover Capital Management, Inc.*, October 28, No. 86-264-CC. Washington, DC: Securities and Exchange Commission.
- Sharpe, William F. 1964. "Capital asset prices: A theory of market equilibrium under conditions of risk." *Journal of Finance* 19, 425–442.
- Sharpe, William F. 1990. "Asset Allocation: Management Style and Performance Measurement." *The Journal of Portfolio Management* 18, 7-19.
- Shiller, Robert J. 2005. "The Life-Cycle Personal Accounts Proposal for Social Security: A Review." NBER Working Paper No. 11300.