
Does Your Benchmark Measure Up?

Carl Bacon, CIPM
Chairman
StatPro Group plc
London

Neil E. Riddles, CFA, CIPM
Founder
Riddles Investment Consulting
Pompano Beach, FL

Selecting an appropriate benchmark is critical to portfolio management and performance measurement. Differences in benchmark methodology and construction rules can lead to large differences in performance, even among benchmarks that appear similar. Understanding these methodologies and construction rules, as well as avoiding behavior associated with common benchmark myths, can aid investors in making wise benchmark choices.

A benchmark is a standard used to evaluate or measure something. The term “benchmark” can be traced to the medieval era; when stonemasons needed a standard measure for building castles, walls, and cathedrals, they would put a mark on the master stonemason’s workbench to serve as the benchmark.

Benchmarks are critical to the investment industry. Asset owners, such as endowments, pension funds, and sovereign wealth funds, use benchmarks to help set asset allocation strategies. Benchmarks are also used to proxy for the past performance of asset classes, to help determine correlations between different asset classes, and to model liabilities.

Benchmarks are essential to portfolio managers as well. Managers’ benchmarks are explicitly stated in investment management agreements. Benchmarks are used to monitor relative returns as well as relative risk. Clients often articulate their risk tolerance relative to a benchmark and benchmark-related metrics, such as tracking error.

Recently, boundaries have blurred between active and passive management relative to benchmarks. In alternative or smart beta strategies, the investment strategy is in the design of the index. There has been an evolution toward using indexes that are designed to be passively followed to deliver active returns to clients rather than using investment strategies that are designed to beat passive indexes.

I will talk about the attributes of a good benchmark, the types of benchmarks that exist, and how to evaluate a benchmark based on differences in

benchmark methodology and construction rules. I will also discuss the Freedom Index, a not-for-profit index company I created to facilitate a more competitive equity marketplace and to encourage advances and customization in index design. Then, Neil Riddles will discuss common benchmark myths and how to avoid making poor benchmark choices based on false beliefs.

Attributes of a Good Benchmark

A good benchmark is investable, accessible, transparent, independent, and relevant. Investable refers to the ability to buy the securities in the benchmark and hold the benchmark portfolio. If a benchmark is not investable, it is questionable whether a portfolio manager can be expected to outperform that benchmark. Information relating to a benchmark’s securities and construction should be accessible and transparent so that portfolio managers can readily understand the benchmark’s components and design. Benchmarks should also be independent. Portfolio managers should not design, monitor, and calculate their own benchmarks; they may well design one that is relatively easy to beat. Finally, a benchmark should be relevant to the investment strategy, even if the strategy requires a high degree of customization.

Some investors question whether it is appropriate to use reference benchmarks, such as target returns or opportunity cost indexes, in a presentation compliant with the Global Investment Performance Standards (GIPS). The GIPS standards state that the benchmark must be appropriate for the investment strategy. If the strategy of the composite is in line with a target return, then that benchmark is appropriate.

This presentation comes from the GIPS Standards Annual Conference held in Boston on 18–19 September 2014.

Best benchmark practices require that a benchmark be unambiguous, clear, and well documented. Portfolio managers should be given one benchmark to outperform; giving a portfolio manager multiple benchmarks simultaneously creates a potential conflict of interest. A classic example is assigning an index and a peer group as benchmarks and asking a portfolio manager to outperform both at the same time. This type of objective is muddled. The performance target should be unique and always set in advance. Benchmarks should never be changed retrospectively.

Types of Benchmarks

At least five different types of benchmarks exist: indexes, peer groups, random portfolios, exchange-traded funds (ETFs), and target returns. Commercially produced indexes include those by Standard & Poor's, MSCI, JPMorgan, Freedom Index, Barclays, and so on. A lot of assets under management are invested in these indexes, and they are generally well designed. Peer groups are different in that a peer group is a collection of competitors—typically mutual funds but also perhaps other institutional managers. Peer groups tend to be country specific; that is, they are direct competitors within a particular market.

Random portfolios are rarely used as benchmarks. An example of a random portfolio is a probability opportunity distribution. This type of portfolio generates a hypothetical set of distributions from the allowable investment universe, ranging from being 100% invested in the worst asset to being 100% invested in the best asset. The manager's objective is to generate performance within the top half of the range of potential returns. Although some industry participants argue that probability distribution is a pure form of measurement, clients do not understand these benchmarks and asset managers dislike them.

ETFs may seem like a surprising choice for a benchmark because unlike traditional benchmarks, ETFs are investment products. ETFs are designed to track indexes, and although they have varying degrees of tracking error, they have the advantage of being investable. Investors cannot actually invest in an index and get the index return because indexes do not suffer bid–offer spreads or transaction costs, but they can buy an ETF and achieve returns that include the impact of fees and transaction costs.

A common type of benchmark for certain strategies is a target return. Examples include the risk-free rate, inflation plus a certain percent, or a funding requirement that is appropriate to the strategy. Target returns are not particularly good benchmarks because they are arbitrary numbers and contain very little information. The great value of an index return, by contrast, is that it can be disaggregated; that is, the

total benchmark return is the sum of its parts. The ability to break down an index's return is essential for attribution analysis.

A question often asked is, which type of benchmark is best? Target returns make poor benchmarks, and random portfolios are rarely used. ETFs can function as benchmarks, but they are not true benchmarks. The debate, then, centers on whether indexes or peer groups make the best benchmarks. Most indexes are investable, whereas peer groups are not. It is very difficult for a manager to structure a portfolio against a peer group because the manager does not have timely access to a peer group's underlying exposures. But it is becoming more difficult for investors to obtain this information from indexes as well. Five or ten years ago, such information was readily accessible, particularly from the larger index providers. Today, index information is becoming expensive. Larger asset managers can easily absorb the cost, but smaller asset managers cannot. This situational disparity creates an environment that is anticompetitive.

One advantage of peer groups is that, like indexes, they are independent. Some investors argue that peer groups are more relevant benchmarks than indexes because they consist of active managers that suffer transaction costs, whereas index performance does not reflect transaction costs. But the disadvantages of peer groups outweigh the advantages for several reasons:

- The return of a peer group is not an achievable target return.
- All peer groups suffer from survivorship bias. Unsuccessful funds close, and thus, the performance track record of the peer group consists of successful and moderately successful funds.
- The quality and consistency of peer groups differ among providers.

Evaluating Indexes

Indexes differ according to weighting methods, coverage, concentration, turnover, and construction rules. Several weighting schemes exist, and the differences in weighting methods create large differences in performance. The degree of universe coverage determines whether an index provides broad coverage with less liquidity or fewer constituents and thus more liquidity.

Turnover also differs among indexes. Indexes are not stable, and a high-turnover index will be expensive to follow. Index owners differentiate their indexes through construction rules to create intellectual property so they can generate revenue or sell other services. Variations in construction rules include the methods used to select stocks, such as sampling, filtering for liquidity and ownership, and

buffering techniques used to reduce turnover. Index providers often use these rules to try to make their indexes cheaper to follow. Other rules that differ among indexes include calculation methodologies pertaining to dividends, corporate actions, and constituent suspensions. Sometimes index design can even be political. For instance, a debate currently exists in the United Kingdom about whether to include predominantly foreign-owned stocks in the FTSE 100 Index.

Weighting Schemes. Indexes also differ according to float weights, or the percentage of a company's stock that is actually available for purchase. In a price-weighted index, higher-priced securities have larger weights and vice versa. In an equal-weighted index, all securities have the same weight. Proponents of equal-weighted indexes say investors have equal opportunities to buy securities, so it makes sense to weight securities equally rather than arbitrarily. A drawback to equal weighting is that it requires regular rebalancing, which can be costly. As time passes, some stocks in an equal-weighted index will increase in value and some will decrease. Calculation rules must be in place to rebalance the index.

Most indexes are market-capitalization weighted—that is, size weighted. A stock's weighting is determined by the number of shares issued multiplied by the stock's market price. Another problem with market-cap-weighted indexes is that they overweight overpriced stocks and underweight underpriced stocks. Fundamentally weighted indexes were designed to overcome this drawback. These indexes use more stable factors to weight securities, such as book value, free cash flow, sales, dividends, and number of employees.

Fixed Income. Fixed-income indexes have a unique set of problems. They contain a significantly greater number of securities than equity indexes, many of which are difficult to price and illiquid. The duration of a fixed-income index, or its sensitivity to interest rates, is the primary driver of index performance. Index duration is determined by issuer, rather than investor, preference. Companies and countries that issue the most debt have the largest

weightings in the index. Debt weighting is the equivalent of market-cap weighting. An index might have more exposure to risky countries, such as Italy and Spain, simply because these countries have more debt.

Customization. Customized indexes are also widely used. As the use of more specialized investment strategies has increased throughout the years, so has the need for customized indexes. The use of these indexes requires independent calculations, consistent application, and thorough documentation. Customized indexes can be constructed so that the asset class weights are either fixed or floating. The issues pertaining to fixed versus floating weights are complex and can lead to large performance differences.

Consider a hypothetical index composed of 50% equities and 50% bonds. After the first month, equities increase to 52% of the index and bonds fall to 48%. If the index weights are fixed, regular rebalancing is required. Fixed-weight indexes are badly designed because they not only introduce transaction costs but also force investors to sell indexes that are rising in price and buy indexes that are falling in price. Such transactions lead to poor performance over time. Investors should accept floating weights, even if an index's rules specify fixed weighting. At the very least, clients and consultants should set a predetermined rebalancing schedule—for example, monthly, weekly, or quarterly.

Hedging. For international indexes, a market standard foreign exchange rate has been developed by WM/Reuters. It allows investors to derive spot currency returns from base and local returns in international indexes. Hedged currency returns can be derived in a similar way. **Table 1** illustrates a sample international index with performance shown in local currency returns, US dollar returns, and US dollar-hedged returns. Performance in an international index might be shown in a base currency as well as the local currency. In this example, the base currency is US dollars. The index might show the return hedged back to that base currency, which is not equivalent to the local return because of the costs or benefits associated with hedging.

Table 1. Sample Index Data

Country	Weight (%)	Local Return (%)	Dollar Return (%)	Dollar-Hedged Return (%)
Australia	8.46	0.81	6.28	0.56
France	23.31	-3.03	0.53	-3.17
Japan	48.46	2.68	6.23	2.84
Singapore	2.16	-3.22	-1.88	-3.19
Switzerland	<u>17.61</u>	<u>0.47</u>	<u>3.89</u>	<u>0.57</u>
Total	100.0	0.67	4.32	0.72

In Japan, the hedged return is higher than the local return, indicating that Japan's interest rates were lower than rates in the United States. Hedging was beneficial in Japan, but in Australia, the hedged return is lower, indicating that Australian interest rates were higher than US rates. There is a cost of hedging as investors borrow money in Australian dollars to sell them and buy US dollars for the hedge. The data in this table are monthly, so the annual difference is 25 bps multiplied by 12. This table demonstrates that interest rates in Australia were 3% higher than they were in the United States. This information is implicit in the data embedded in the index.

One final note is that a partially hedged index, such as an index 50% hedged to the US dollar, may provide more than 50% dollar exposure if there is any dollar exposure from US equities in the underlying index. Clients who want 50% exposure to the US dollar may need to have a separate index to obtain that exposure.

The Freedom Index

The Freedom Index is a not-for-profit index limited by guarantee.¹ I formed the Freedom Index out of frustration. There is no share capital and no distribution of profits. I am not paid. I believe that index information should be free, or nearly free, because it is basically market barometer information. Having access to that information creates a competitive market composed of small players as well as big players.

The objective of the Freedom Index is to provide free, open, and independent benchmark indexes that are designed by the asset management community for the asset management community. The indexes are funded by voluntary subscription. We provide free access to underlying weights and returns at the security level. Daily data are available on the website for each security—daily weights and returns for the complete equity market, including the United Kingdom, the United States, South Africa, and so on. There are basic starter indexes in local currencies, sterling, and local currencies hedged to sterling. Investors can choose indexes with 50, 100, 250, 500, or 1,000 stocks or even indexes that include all stocks in a universe, but the main objective of the Freedom Index is to encourage asset managers and asset owners to design their own indexes. The Freedom Index Company will confirm for any index that uses our data that it has used our data consistently and has appropriate methodology that is consistently applied and well documented. We will publish that information on our website. Our

¹See www.thefreedomindex.com.

goal is to facilitate the design of customized indexes that are made available to the general asset management community. We hope the website will generate community discussions about data quality, index design, sector definitions, and so on.

Benchmark Myths

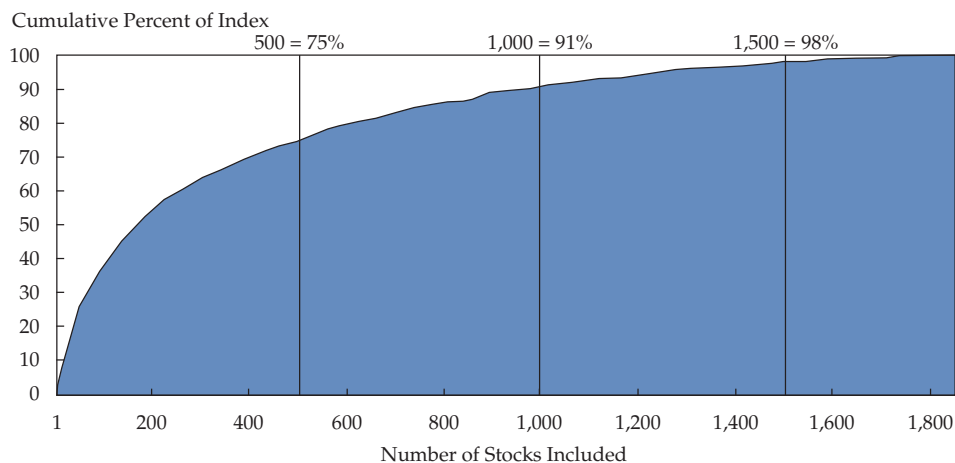
“Benchmark myths” are assumptions about benchmarks that arise from false beliefs. There are some common areas of confusion about benchmarks and their structure. We will outline some points that are frequently misunderstood.

Myth: More Coverage Is Always Better. In an ideal world, a benchmark would contain all securities in the universe it is measuring. This type of coverage is called “full replication.” Some indexes contain a representative subset of the universe they are measuring. Those indexes are called “sampled” indexes. The accuracy of sampled indexes depends on the effectiveness of the sampling methodology. Nonsampled indexes are not necessarily full replication. In reality, very few true full-replication benchmarks exist. Most benchmarks have a minimum market-cap requirement.

Figure 1 shows the diminishing value of adding additional stocks to a benchmark. The index shown is the MSCI ACWI ex USA Index as of 31 July 2014. The total number of stocks in the index is 1,829. The figure shows that 500 stocks cover 75% of the benchmark universe. At 1,000 stocks, coverage is 91%; at 1,500, 98%. Index coverage does not increase by the same rate as the number of stocks. Coverage is a function of the stocks' weights in the index multiplied by their returns. At 98% coverage, the remaining 2% of stocks would have to perform in a significantly different manner from the rest of the market to make any difference at all. These stocks are also the most expensive and difficult stocks to include because they are the smallest and most obscure stocks. Owing to the cost of including the very smallest stocks combined with their minimal influence on the index return, most index providers limit inclusion to some minimum size.

Percent of coverage should be just one of the factors that investors consider when choosing a benchmark, in addition to the other criteria mentioned earlier, such as methodology, weighting schemes, and availability of data.

Myth: The Manager Selects the Benchmark. Typically, plan sponsors have specified benchmarks. For example, an RFP (request for proposal) might state that a fund is searching for “an active US equity manager to run a portfolio versus the S&P 500 Index.” It may be difficult for plans to change benchmarks. Also, plan sponsors are somewhat

Figure 1. Index Coverage by Number of Stocks Included

Notes: The index shown is the MSCI ACWI ex USA Index as of 31 July 2014. The total number of stocks is 1,829.

Sources: Based on data from MSCI and FactSet.

skeptical about managers suggesting a change in benchmarks. Broad benchmarks generally have similar returns, so even if a manager believes that the Russell 1000 Index is a better benchmark than the S&P 500, that manager will generally agree to be measured against the S&P 500. In the long term, it is unlikely to make much difference in how the manager is evaluated. Performance differences become more apparent at the sub-index level and for such analytical techniques as attribution.

Myth: Float Is Float. Size-weighted indexes can be market-cap weighted or float weighted. Float weighting adjusts market-cap weights to reflect shares that are actually available to investors. Indexes that are float weighted typically exclude shares that reflect government ownership, controlling shareholders' and their families' ownership, large stakes held by management and other companies, and foreign investor limitations.

Because it better reflects an investor's opportunity set, float weighting is preferable to full market-cap weighting. However, float weighting is not unambiguous. Some of the adjustments are subjective, and it can be difficult to determine where to draw the line when excluding blocks of stock. Different index providers may come to different conclusions. For international investors, float weighting is also relative. For example, the United States restricts foreign ownership of some companies. Non-US investors cannot buy 100% of a US airline stock; they can only buy 25% of the voting stock. A number of these types of exceptions exist. So, an index for international investors may have a different estimation of float for a particular company from an index for domestic investors. Foreign ownership limits can be difficult to

model accurately; the information used to determine float weighting is not always perfect.

Myth: Current Index Definition Explains History. Indexes evolve over time. The index methodology available today does not apply to the entire index history. A good example is the use of float weighting. Float weighting was not common in the past but is now used by most broad equity indexes. Some index histories include a non-float period, then a mix, and then pure float weighting. It can be very difficult to determine an index's historical methodology because although there is comprehensive information available about what an index provider is doing today, it can be difficult to get information about what has been done historically.

As an example, the MSCI World Index included Mexico prior to 1987, South African gold mines until the 1990s, and Malaysia for a period of time as well. Today, none of these constituents are in the index. In another example, the S&P 500 Index had 223 stocks when it was introduced in 1923. Presumably, the index was not called the "S&P 500" at the time. The index expanded to 500 stocks in 1957. Foreign stocks were removed from the index in July 2002. Both the MSCI and the S&P indexes had six-month periods during which they were half float weighted and half cap weighted as they were transitioning to being float weighted. The Dow Jones Industrial Average was composed of 12 stocks in 1896, then 20 stocks, and now 30 stocks. All of these index changes remain in the performance track record.

Myth: Growth and Value Should Sum to Total. The myth that growth and value stocks should sum to the total universe of stocks also applies to small cap, large cap, and a number of other strategies. If

the benchmark objective is style analysis or asset allocation, the benchmark should be exclusive—that is, discrete. If the benchmark objective is manager evaluation or risk control, the benchmark should be nonexclusive—that is, overlapping. For nonexclusive indexes, some stocks will be in both the growth and value indexes and thus be double counted, and some stocks will not be in either index. An overlapping index provides a clearer picture of a manager's value added and risk taken. It allows an investor to assess misfit risk, or areas of the benchmark not covered by the managers, by including the most relevant benchmarks.

In reality, most growth managers will buy some stocks that value managers like and most value managers will buy some stocks that growth managers like. Splitting the universe in two makes little sense. Growth managers will not avoid a stock that has good growth prospects but is trading cheap; they would actually prefer that stock. Value managers will not avoid a stock that is trading cheap but also has great growth prospects; they would prefer that stock. Overlap always exists between growth and value, and overlapping stocks are more likely to be given a heavier weighting in a portfolio. Exclusive indexes make sense if the goal is to determine a custom blend of indexes. But if the objective is to simply find a benchmark for an active manager, then an overlapping universe is more appropriate.

Myth: Turnover Is Not a Big Problem. Some types of indexes experience a lot of turnover. Excessive turnover is a problem because indexes trade without transaction costs, whereas real portfolios pay commissions and market impact. Also, indexes hypothetically trade at the closing price. For cap-weighted or float-weighted indexes, turnover tends to be small because they are self-rebalancing; prices change, but the number of shares does not.

A comparison of annual turnover between two MSCI indexes illustrates how turnover increases with active management. The MSCI ACWI ex USA Index is a large, broad international index. It has turnover of 4%. The MSCI ACWI ex USA Value Index has turnover of about 25%, which approximates the degree of trading that some active managers have in their portfolios. The high level of turnover is the result of the classification of growth and value stocks being somewhat arbitrary.

Every six months, a group of stocks moves back and forth between the growth index and the value index, creating turnover. This “passive” index has quite a bit of trading activity, executed at the market close with no transaction costs. Real portfolios have to pay to trade. Similarly, the same stocks have different transaction cost effects depending on which

index they are in. When tested at the end of July, if the largest stocks in a small-cap index had moved into a large-cap index, it would have caused 22 bps of turnover in the small-cap index, but that same transaction would have caused 1 bp of turnover in the large-cap index. The difference in turnover occurs because the largest stocks with the largest weightings in the small-cap index are moving into a large-cap index, where they become the smallest stocks with the smallest weightings. The amount of turnover depends on the index definitions. Small-cap indexes and GDP-weighted indexes, for example, can have quite a bit of turnover.

Myth: Rebalancing Periodicity Does Not Matter. This myth pertains to balanced benchmarks. How often an asset mix benchmark is rebalanced can make a big difference in returns. The asset classes do not have the same returns, nor are they perfectly correlated, so rebalancing will change the overall return. The benchmark should reflect the portfolio's policy asset mix. For example, if the portfolio has a 60/40 policy asset mix, the benchmark should be similarly weighted. Because transaction costs are associated with rebalancing, rebalancing should generally not be done monthly. Also, rebalancing should not be dependent on the portfolio's active weights. A valid benchmark is rebalanced without the need to reference the active portfolio. Turnover owing to rebalancing does not create transaction costs for an index but does create transaction costs for the active portfolio.

Myth: Net-of-Tax Indexes Are More Relevant. Another myth is that “net-of-tax indexes” are more relevant for investors. Many governments withhold taxes on dividends received by foreign investors. Net-of-tax indexes reduce the dividend received by the amount of taxes withheld. Most countries, however, have dual taxation treaties, so investors receive a partial tax refund on dividend taxes. Net-of-tax indexes assume the investor receives none of the tax back. Returns will always be lower for a net-of-tax index than a gross-of-tax index, so net-of-tax indexes are easier to beat. Because most investors receive a partial tax refund, a net-of-tax index can be too aggressive and a gross-of-tax index can be too conservative.

Myth: The Industry Will Move to the Best Index. Indexes are like software in that the newest index tends to be the best because it includes updates and improvements to existing benchmarks' methodologies. But the industry does not automatically move to the newest index. One reason that newer benchmarks are not adopted is that older benchmarks tend to have more history. Other reasons include plans' reluctance to change

benchmarks and managers' reluctance to pay for and manage to multiple benchmarks. At the overall market level, broad benchmarks provide similar returns and information.

This final reason new benchmarks are not adopted is inertia. Often, there is an assumption that whatever index is being used is "good enough." So, investors do not automatically make the effort to change to improved indexes as they become available.

Conclusion

The benchmark decision is the foundation of performance analysis. If the benchmark is incorrect, the analysis and portfolio decisions that flow from it will be incorrect. Evaluating many benchmarks to determine the proper one is time well spent. Information about individual benchmarks, such as construction methodology, can be found online.

Investors should evaluate the most important attributes of a benchmark rather than focusing on subtle changes. An international benchmark should cover the countries the portfolio will be invested in. For a small-cap benchmark, the definition of small cap should be in line with the investor's definition. These types of considerations are more important than subtle ones, such as the way a benchmark handles corporate actions. Investors should also be aware that benchmarks that seem similar may be very different when examined closely. For example, growth and value benchmarks can differ markedly from one provider to another.

In conclusion, faulty benchmarks lead to bad investment decisions, poor client relations, and costly manager turnover. Having the correct benchmark is good for both the manager and the client. Investing the time to pick the right benchmark is critical.

This article qualifies for 0.5 CE credit.

Question and Answer Session

Carl Bacon, CIPM, and Neil E. Riddles, CFA, CIPM

Question: Why is a net-of-tax benchmark too aggressive and a gross-of-tax benchmark too conservative?

Riddles: For investors based in the majority of countries—with the exception of tax havens, such as Luxembourg or the Bahamas—dual taxation treaties exist. When a dual taxation treaty is in place, investors will receive some, but not all, of the withholding tax back at the end of the year. For these investors, a net-of-tax construction methodology lowers the benchmark return too much because the tax drag on the benchmark is overstated. Yet, a gross-of-tax return is too conservative—the returns are higher and thus harder for a manager to beat—because this methodology assumes that an investor receives all withheld taxes back.

Question: How frequently do you recommend rebalancing an asset mix?

Bacon: I would avoid rebalancing because it generates transactions and turnover, which are expensive. I recommend discussing with clients why they want to fix the weights and advising them to allow the weights to float for at least a year. The more frequently rebalancing is done, the more expensive it is. When there is momentum in the market, rebalancing forces investors to sell a rising index and buy a falling index. This strategy does not yield good performance results.

Riddles: A policy portfolio benchmark is supposed to be passive, so rebalancing every 6–12 months is probably fine.

Question: Is it ever appropriate to use T-bills or other risk-free rates as benchmarks?

Bacon: From a target-return perspective, risk-free rates might be used because investors expect higher returns than the risk-free rate from investing in risky securities. But from an attribution standpoint, there is very little information contained in those sorts of benchmarks. This lack of information prohibits any kind of detailed performance analysis. If an asset manager invests in securities other than the risk-free rate, an implicit strategy is being implemented that should be replicable in some type of customized benchmark. From my perspective, risk-free rates are not good benchmarks.

Riddles: For a long–short portfolio, I can understand why an investor would choose the T-bill rate as a benchmark. The strategy seeks to generate alpha but also expects to earn at least the T-bill

rate. One problem with using T-bills as a benchmark for a long–short strategy is that the portfolio may seek to outperform the T-bill rate over time, but the benchmark’s level of risk is nowhere near the risk level in the portfolio—a real mismatch.

Question: What types of benchmarks do you recommend for alternative strategies, given that peer groups are subject to survivorship and self-selection biases?

Riddles: The alternative category covers a huge group of managers. Peer groups may be the best comparison for long–short or go-anywhere managers that can invest in any type of security or strategy. Investors who use a peer group as a benchmark should choose a group of peers that would have met their initial investment criteria rather than a broad group of managers that they happen to be tracking today. For instance, if an investor would not place money with a manager who does not have at least \$1 billion or more in assets under management, then these managers should not be included in the peer universe.

Bacon: Peer groups are particularly susceptible to survivorship bias, but liquid indexes do not work well as benchmarks for alternative asset portfolios either because liquid indexes are more volatile. Investors also have to be very careful when using risk-adjusted statistics with these types of asset categories.

Question: What are the problems associated with leasing index data?

Bacon: Leasing index data creates huge problems. In many cases, investors do not realize that they are leasing, rather than buying, the data. When index information is leased, a firm may have to remove all historical data once payment is discontinued. If the firm is discovered to have some residual data, it may be fined. It is very difficult to purge a firm’s systems of index information because the data are vital to the firm. It is essential for organizations to understand whether they are leasing or buying index data.

Riddles: Many lease contracts specify that if an index provider believes that data still reside on a firm’s servers, the firm will be required to pay for an audit by an auditor of the provider’s choosing. The auditing process is very expensive. An example of data usage limitations that exist under a leasing contract is the typical MSCI client

contract, which permits only a seven-year window of trailing data. MSCI clients are supposed to delete any data beyond that period. Typically, a firm's legal counsel specifies in the contract that the data can be retained if required by law, but any other exception is very costly.

Question: If a firm wants to change index providers, must it remove the 10-year index performance from its GIPS presentation, or does the firm need to stay with the index provider forever to retain the benchmark data used in the calculations?

Riddles: I have encountered this situation before, and no, the firm is not stuck with the index provider forever. Hopefully, the firm's legal counsel has added a clause in the contract that says the firm will abide by the contract to the extent it is not required to keep the data by law. In the case of a GIPS presentation, the firm is required by law to keep the relevant data behind whatever performance information has been shown. The actual index level tends not to be protected information; most index providers put this information on their websites. It is the other, more detailed data, such as industry weights, country weights, and constituents, that are treated as subscription-based intellectual property.

Question: Because asset managers do not select their benchmarks and it is very difficult to compete with established commercial indexes, what is the expected adoption rate of the Freedom Indexes?

Bacon: There is clearly a branding issue with the Freedom Index. When clients pay for a FTSE, MSCI, or Russell index, they are really paying for the brand. The Freedom Index is not sold as a brand; in fact, it is not sold at all. I am giving it away. It is true that there is no real advantage for large asset managers to use the Freedom Index because they have multiple clients using the established brand-name indexes. It would take a very long time for these managers to switch indexes.

But the Freedom Index is gaining traction among smaller managers who do not have an alternative and who find it expensive to use brand-name indexes. Smaller managers cannot compete without this information. Demand is also coming from investors who design customized indexes and want them independently calculated. This type of use is the main objective of the Freedom Index.

I want to encourage people to introduce their own intellectual capital, rather like Wikipedia represents a pooling of intellectual capital.

Question: Is it ever acceptable to compare a gross return with a net benchmark, such as a peer group?

Bacon: It is more appropriate to show net performance against a peer group. A peer group is a net-of-fees benchmark, in essence. But if gross returns must be shown for some reason, then a firm must certainly disclose that the returns are not net of fees.

Riddles: Some small-cap growth managers say a price-only index is the most appropriate benchmark because small-cap stocks generally pay very few dividends. But even in this circumstance, a price-only benchmark is not appropriate.

Question: Are there benchmarks for private equity strategies?

Riddles: A peer universe is really the only option for a private equity benchmark.

Question: How do you address changing a benchmark over time to reflect changes in a portfolio's strategy?

Riddles: One way is to keep the original benchmark for the initial time period and link a second benchmark to it at an appropriate time. If a strategy changes substantially but does not constitute a new product, this approach can be used.

Bacon: A benchmark can be spliced; the history of a series of returns can always be linked together. So, the benchmark can change completely. Whether such a change is appropriate is another issue, but a change in benchmark can be implemented in a way that provides continuous measurement if that change aligns with the strategy.

Question: Can more than one benchmark be listed on a GIPS-compliant presentation?

Bacon: Yes, depending on the context. It is certainly not best practice to give a portfolio manager multiple benchmarks to outperform, but one could argue that more than one benchmark might be relevant in certain circumstances.

Riddles: For example, the MSCI ACWI ex USA Index might be the most appropriate benchmark, but a manager may also include the MSCI EAFE Index because it is so widely known.