

The Index Investor

Why Pay More for Less?

Performance Update

Our first set of model portfolios are designed to deliver returns that are superior to their respective benchmarks, while taking on the same amount of risk (that is, having the same expected standard deviation of returns). Thus far this year, they have underperformed their benchmarks. Our first portfolio is benchmarked against a mix of 80% U.S. equities (as measured by the Dow Jones Total Market iShare, IYY) and 20% U.S. bonds (as measured by the Vanguard Total Bond Market Fund, VBMFX). Year to date, this benchmark is down (6.5%), while our model portfolio is down (11.5%). The major cause of this underperformance is the model portfolio's allocation to European equities, which have underperformed the U.S. market so far this year. A secondary cause is the model portfolio's allocation to real assets, which is also down on the year. On the other hand, this part of the portfolio was a major reason for last year's outperformance versus the benchmark. One is reminded of Crash Davis' comment in the movie Bull Durham: "Some days you win, some days you lose, and some days it rains."

The second portfolio in this group is benchmarked against a mix of 60% U.S. equities and 40% U.S. bonds. Year-to-date, this benchmark portfolio is down (4.3%), while our model portfolio is down (8.7%). Again, European equities and commodities are the culprits. The third benchmark portfolio is a mix of 20% U.S. equities and 80% U.S. bonds. Through the end of July, it is up .1%, while our model portfolio is down between (3.7%) and (2.0%), depending on the international bond fund used to measure its performance. The main story here has been the surprising strength of the U.S. dollar this year, and the consequently weak returns delivered by non-U.S. dollar bonds.

Our second set of model portfolios are designed to match the returns of their respective benchmarks, while taking on less risk. They have also underperformed. While the 80/20 benchmark is down (6.5%) year-to-date, the model portfolio is down (11.8%). The 60/40 benchmark is down (4.3%) year-to-date, while our model portfolio is down (6.9%). Finally, the 20/80 portfolio has a .1% return through the end of July, while the model portfolio is down between (3.6%) and (1.8%), depending on the international bond fund used.

Our last set of model portfolios are designed differently. They assume that an investor wants to maximize the probability of achieving at least a minimum target level of return, while taking on the least amount of risk possible. Year-to-date, our 12% target return portfolio is down (11.5%), our 10% target return portfolio is down (11.8%), our 8% target return portfolio is down (7.8%), and our 6% target return portfolio is down (4.9%).

Product and Strategy Notes

In this month's In-Focus, we will be looking at the possible uses -- and abuses -- of country funds and ETFs. Before we do that, however, some background information is in order.

Let's start by quantifying the big picture. At the end of last year, total world GDP (gross domestic product, which is a measure of an economy's total output) was on the order of 45 trillion U.S. dollars, calculated on a purchasing power parity basis (that is, using exchange rates that equalize across countries the cost of a standard basket of goods). By comparison, at year end 2000, the total market capitalization of the 49 equity markets that are included in the FTSE All-World Index was around 35 trillion U.S. dollars. However, because the use of equity markets varies around the world, the distribution of equity market capitalization between countries was very different than the distribution of GDP.

To get a better handle on how the distribution of GDP and equity market capitalization differ (and also on why it matters), we need to start with the companies that provide the major indexes.

The best known provider of global equity market indexes is Morgan Stanley Capital International, or MSCI. In the past, their country indexes included sixty percent of a given equity market's full-market capitalization (that is, the full equity market value of the companies whose shares traded on the market in question). In recent years, however, other index providers (see below) have moved to what is called "free float" capitalization weighting of their index products. They did this because in many countries, not all of a company's equity shares (or "float") are freely available to foreign investors. For example, some may be locked up in the form of long term strategic investments by other companies, some may be owned by the government itself, and some classes of shares may not be available for purchase by non-citizens of the country in question. The difference between a company's or country's "free float" capitalization (that is, the number of shares foreigners can buy, times their price) and its "full market" capitalization (the total number of shares outstanding, times their market price) often made it difficult or impossible for foreign investment managers to match or equal the performance of the MSCI index for the country in question. In countries such as China and India, where thirty percent or less of market capitalization is actually available to foreigners, this was a significant problem.. As a result, this year MSCI has made two major changes to its indexing approach. First, it will now use free float, rather than full market capitalization to calculate its indexes. And second, it will now attempt to cover 85% of free float capitalization with its index, versus the previous objective of covering 60% of full market capitalization.

MSCI's indexes include the All Country World Index, which tracks equity market returns across 49 countries, the EAFE, which represents 21 developed equity markets outside North America, and the Emerging Markets Index, which covers 26 developing country markets around the world.

The FTSE/Actuaries series of indexes were broadened in recent years when they acquired Barings' international indexes when that company went bankrupt. As a result, the FTSE now also covers the same 49 equity markets as MSCI (though they call their broadest index the "All World" instead of the "All Country World". Go figure...). However, it has already moved to free float weighting, and its indexes aim to cover 90% of free float market capitalization.

The DowJones Global Indexes are similar to the FTSE in that they also attempt to capture 95% of free float market capitalization in the developed country markets they cover, and 80% emerging markets. A significant difference between DowJones and the other providers, however, is that they cover only 34 countries. More specifically, the DJGI leaves out fifteen developing countries that are covered by MSCI and FTSE, including: Argentina, Peru, and Colombia; Israel, Jordan, Egypt, and Morocco; Russia, Turkey, Poland, Hungary, and the Czech Republic; and China, India, and Pakistan.

Does this matter? Maybe yes, and maybe no. Consider these figures. At the end of last year, the 49 countries covered by MSCI and FTSE accounted for about 92% of the world's GDP, while the 34 countries covered by DJGI accounted for just 68% of world GDP. On the other hand, the 34 countries covered by DJGI accounted for 99.6% of the total market capitalization (using FTSE's measure of market cap) included in the other two world indexes. As long as equity markets don't further develop in the countries left out of the DJGI (a good way to think of this is as long as market capitalization as a percentage of their GDP doesn't increase), then it probably doesn't matter which global equity index you use. However, if those markets do develop, then it is highly likely that we will see significant differences between the reported returns on the different indexes, and on the funds that track them.

This issue of index coverage is important, so let's extend it further, using the MSCI indexes as the focus of our discussion.

Traditionally, the most popular measure of non-U.S. equity market performance has been the MSCI EAFE (which stands for Europe, Asia and Far East). The EAFE covers 21 developed equity markets outside of North America. These include France, Germany, the Netherlands, Italy, Spain, Finland, Ireland, Belgium, Greece, Portugal, Austria, Denmark, Norway, Sweden, Switzerland, and the U.K. in Europe, Japan, Hong Kong, and Singapore in Asia, and Australia and New Zealand in the Far East. Collectively, the EAFE countries accounted for about 31% of world PPP GDP in 2000, and about 44% of the total capitalization of the 49 equity markets in the MSCI All Country World Index.

The EAFE is further subdivided into two indexes. The Pacific Index includes Japan, Hong Kong, Singapore, Australia and New Zealand, which accounted for about 9.5% of world GDP and 13.5% of ACWI total capitalization, or, put another way, about 31% of the EAFE's capitalization. The remaining EAFE countries comprise the Europe Index, which account for 21.2% of world GDP and 30.1% of the ACWI's capitalization. Finally, Europe itself can be subdivided, into countries that have joined the European Monetary Union (and agreed to use the Euro as their currency), and those that have not. The countries that have joined the Euro (France, Germany, the Netherlands, Italy, Spain, Finland, Ireland, Belgium, Greece, Portugal, and Austria) are included in the EMU Index. The EMU countries accounted for 16.3% of world GDP and 15.9% of ACWI capitalization at the end of last year.

The Emerging Markets Index covers twenty six countries that collectively account for 36.5% of the world's GDP in 2000, but only 2.3% of the ACWI's total capitalization. These countries include Brazil, Mexico, Chile, Argentina, Peru, Venezuela and Colombia; Russia, Poland, Hungary, and the Czech Republic; Turkey, Israel, Jordan, Egypt, and Morocco; South Africa; and South Korea, Taiwan, Thailand, China, India, Indonesia, the Philippines, Malaysia and Pakistan.

While we are on the subject, we should also note that Vanguard's Emerging Markets Index Fund (VEIEX) does not track the EMF. Rather, it tracks a subset, called the EMF-Select, which includes only thirteen countries: Argentina, Brazil, and Mexico; Hungary,

Poland, and the Czech Republic; Turkey, Israel, and South Africa; and South Korea, Indonesia, the Philippines and Thailand. Together, these countries accounted for 14% of world GDP in 2000, and 1.6% of the ACWI's total capitalization.

In terms of missing GDP, the two biggest omissions on this list are obviously China and India. Given the low levels of free float available in these countries, we can well understand the logic of leaving them out of Vanguard's Emerging Markets Index (given the fund's potential need for liquidity to cover redemptions). Still, we were curious if there was another way we could get exposure to the equity market index in these two countries. The obvious way to do this would be to buy closed end funds that invest in these two countries. Fortunately, a number of these exist. We looked at two: The India Fund (IFN) and the China Fund (CHN). These funds are not cheap; in addition to brokerage commissions on their purchase, they also charge relatively high expense loads of 1.55% and 2.12%, respectively (versus the .59% charged by Vanguard on the VEIEX). The key question we looked at was how closely these funds tracked their respective MSCI Index over the last five years. The results were borderline at best. The India Fund's correlation with the MSCI India Index was .80, while The China Fund's correlation with its MSCI Index was .77. Not bad, but not an index fund. Still, if you want to include some exposure to these two countries in a portfolio, these two funds seem a good way to do it.

Last but not least, we also looked at the wide range of iShares that track individual country indexes (as well as the EMU index). There are currently twenty single country iShares available (not counting the U.S.), including: Australia, Austria, Belgium, Brazil, Canada, France, Germany, Hong Kong, Italy, Japan, Malaysia, Mexico, the Netherlands, Singapore, South Korea, Spain, Sweden, Switzerland, Taiwan and the United Kingdom. Collectively, these iShare countries accounted for 38.5% of world GDP, and 44.8% of the ACWI's total market capitalization. When you consider that the U.S. accounted for 51.8% of the ACWI's market cap, then iShares total coverage is 96.6% of total ACWI capitalization. In short, you can do a lot with iShares. If they have a weakness, it is in the emerging markets area. The five iShares countries in this group account for only

1.3% of the ACWI's capitalization, versus 1.6% for Vanguard's Emerging Markets Index Fund (which doesn't include Malaysia or Taiwan), and 2.3% for the full MSCI Emerging Markets Free Index.

In Focus: Country Tilts

As you may remember, when we rebalanced our recommended portfolios at the end of last year, we used a fairly broad definition of an asset class. Specifically, because the benefit from diversification comes from risk reduction, we required that the “asset classes” we used could have no more than a .60 correlation of returns with each other. That definition eliminated from use a number of groupings of stocks and bonds that other commentators call “asset classes.” Examples of these include small cap stocks or large cap growth stocks, and short-term bonds. In our view, all of these represent various “tilts” that one can make in order to enhance the risk/return trade-off within an asset class. At the time of our rebalancing, we promised that we would be taking a closer look at these “tilts” to see which, if any of them, made sense. Last month we looked at tilts based on economic sectors. This month we look at country fund investing. Next month we’ll combine our insights about sector and country tilts, and reach some conclusions about their implications for asset allocation. In September, we’ll look at investing in different bond maturities, and in October we’ll look at momentum investing. As was the case last month, the fundamental question we’re trying to answer is whether or not you can improve on the risk/return trade-off for the asset class as a whole by making a country tilt in your portfolio.

The data set for our analysis covered the period from January, 1988 to December, 2000. During this period, the EAFE generated average annual returns (in U.S. dollars) of 8.77%, with a standard deviation of 18.32%, or .48% of return per unit of risk taken on. Within the EAFE, the Europe Index had average annual returns of 15.01%, with a standard deviation of 17.06%, or .88% of return per unit of risk, while the Pacific Index had returns of 3.26% with a standard deviation of 23.79, or a measly .14% of return per unit of risk. Finally, the combined EAFE+EMF Index (our proxy for the Vanguard Total

International Market Fund) generated returns of 8.56% with a standard deviation of 18.19%, or .47% of return per unit of risk. Our challenge was to see if it was possible to do better than these results, while taking on no more risk, and investing only in iShares.

First we did an obvious analysis, using relative GDP to weight the available Europe and Pacific iShares (for example, we calculated the percentage of Spain's GDP relative to the total GDP of all the iShares included in the Europe Index, and used that as our allocation to Spain). In the case of Europe, this had a very beneficial impact: average annual returns increased to 17.88%, while standard deviation grew to only 18.67%. As a result, return per unit of risk grew to .96%. The main difference from market cap weighting was that we had relatively more exposure to Germany and Italy, and less to Switzerland and the U.K.

In the Pacific region, this approach didn't improve things much. Average annual returns were only 3.54%, with a standard deviation of 19.21%, or .18% of return per unit of risk. The fundamental problem here is that in both GDP and market cap terms, Japan dominates this index.

We then got a bit more adventuresome, and decided to take advantage of the fact that Canada is not included in the EAFE, and therefore typically gets left out of many portfolios. We decided to include it in our GDP weighted Europe portfolio to see what would happen. At the same time, we substituted the EMU iShare for the iShares of countries included in the EMU, not only to hold down transaction costs, but also to gain exposure to Ireland, Greece, Portugal and Finland, which lack iShares of their own. The results were encouraging. Average annual returns were 15.69% with a standard deviation of 16.69%, or .94% of return per unit of risk.

We then took the same approach to the EAFE, using GDP weights, and including Canada and the EMU iShare. In this case, average annual returns were 11.96%, with a standard deviation of 15.45%, or .77% of return per unit of risk. This was a considerable improvement over the market cap EAFE (without Canada) index.

We then applied the same approach to the full EAFE+EMF, using the emerging market iShares, the EMU iShare, the Canada iShare, and the other EAFE iShares all at their respective relative GDP weights. The results were impressive: average annual returns were 16.31%, with a standard deviation of 17.65%, or .92% of return per unit of risk (almost double the result for the market cap weighted EAFE+EMF index).

Given the consistent improvement in results we were able to achieve, we need to ask the same question we've asked in previous chapters in this series. Are these results just an historical accident, or is there something more fundamental, and possibly sustainable, going on here?

First, the good news. Diversification across countries can improve risk adjusted performance because the correlations of returns between them are often low. Moreover, the cause of these low correlations goes beyond simple exchange rate relationships, as evidenced by the low correlation in many countries between exchange rate changes and equity market performance. In short, many of the differences between countries that may give rise to low equity market correlations -- such as differences in culture, language, ways of doing business, accounting rules, and the like -- are not going to go away any time soon. Equally as important, research has demonstrated that investors tend to have a "home country bias" -- that is, regardless of the potential for improving risk adjusted returns through international diversification, they still prefer to invest the bulk of their portfolios in the country in which they live. Taken together, these factors form a basis for concluding that the potential benefits from country tilts will continue to exist in the future.

Now for the bad news. The period covered by the data was one during which a lot of fundamental changes took place, including the formation of the Euro bloc, a substantial increase in the integration of the Mexican and United States economies, the introduction of significant structural reforms in Brazil, and a sharp increase in the global reach of the Nordic economies. This raises the possibility that some of the processes that gave rise to

the improvement in results from simple GDP weighing may no longer be operating (e.g., economic policies are more similar across countries, and exchange rates are relatively more stable), or may be operating more weakly than in the past. Finally, because the relative importance of different industry sectors differs across countries, some of the difference in country performance may really have been driven by underlying differences in the structure of economic activity. For example, at the end of 2000, the energy sector accounted for about 26% of the Netherland's market capitalization (based on MSCI data), and 22% of Argentina's. In contrast, it accounted for only .6% of Japan's market cap, and 11% of Canada's and the U.K.'s. More complicated statistical analysis would undoubtedly show that industry effects were interacting with true country to produce the differences in country index returns that we observe (we'll look at the implications of this analysis next month).

Still, on balance, it would seem that you could do worse than take a simple GDP weighing approach when it comes to dividing your assets if you want to make country tilts. However, caution is in order: carried to its logical conclusion, this approach would generate relatively high exposures to countries like China where respect for foreign investors' and minority shareholders' rights is still spotty at best. Anybody adopting GDP weighting is therefore well advised to set as a pre-condition that the countries receiving a GDP based allocation have well functioning equity markets.

GDP, of course, is but one approach to weighting countries in ways that differ from their market cap weights in an index. Another approach is to use an optimization model to help you determine the country weights. When we did this, we came up with some interesting results.

Our first analysis focused on the MSCI Europe Index, which had achieved average annual returns of 15.01% (in U.S. dollars) between 1/88 and 12/00, with a standard deviation of 17.06%, or .88% of return per unit of risk. As possible investments, we used the country iShares contained in the Europe Index, as well as the Canada iShare. We set two constraints: (1) no more than 15% of the portfolio could be invested in any single

country, and (2) the expected standard deviation of our optimized portfolio had to equal 17.06%. The portfolio that resulted had an expected annual return of 19.47%, or 1.14% per unit of risk taken on -- quite an improvement over the index, assuming, of course, that history is a good guide to the future (which sometimes it isn't!). The weightings behind this result were as follows: Belgium, France, the Netherlands, Sweden, Switzerland, and the U.K., all 15% each; and 10% to Canada.

However, when we looked at this result, we were concerned that some of it may have been driven by historical exchange rate factors that have moderated or disappeared since Belgium, France and the Netherlands signed up to use the Euro. To test this, we ran another optimization, this time using the MSCI EMU Index in place of separate indexes for each EMU member country. However, we again included Canada, and the non-EMU European countries for which iShares are available -- Sweden, Switzerland, and the U.K. As before, we set some constraints on our optimization: No more than 20% weighting to any country, except the EMU group, which could be weighted up to 50%, and expected standard deviation equal to 17.06%. This time our optimized portfolio had an expected annual return of 18.53%, or 1.09% per unit of risk taken on -- still an impressive result.

We next attempted to apply our methodology to the Pacific Index, and came up against one of its major shortcomings. In short, our model said, in effect, 50% Hong Kong, 40% Australia, 10% Singapore, and nothing in Japan. Of course, if history repeats itself in the future, this is undoubtedly the right way to maximize return per unit of risk. But what are the chances that the Japanese equity market will perform as poorly over the next ten years as it has over the past ten years? This is a perfect example of the dangers of blindly following the dictates of an optimization model, without first applying the common sense test to its recommendations. In this case, history is probably at best an imperfect guide to the future, not only because of the prospects for change in Japan, but also because of the tendency for investors and markets to overreact in both directions. In this case, the best route to take would probably be to set an additional constraint on the optimization that specified the minimum percentage of the portfolio that had to be allocated to Japan. As long as that percentage is less than its roughly 80% weighting in the market cap based

Pacific Index, the resulting expected returns will be higher than what has been achieved historically (again, assuming the past is a good guide to the future).

This is the approach we took in our next analysis, which was focused on the EAFE+EMF combined index, which had an average annual return of 8.56% between January, 1988 and December, 2000, with a standard deviation of 18.19%. In this case, we set a minimum investment in Japan at 12% of the portfolio, which is equal to the country's share of the total GDP of the countries included in the index. Again, we used the EMU instead of the individual country indexes, and set a constraint that it could account for no more than 50% of the portfolio. Again, we included Canada, Sweden, Switzerland, the U.K., Australia, Hong Kong, Japan, Singapore, Brazil, South Korea, Malaysia, Mexico, and Taiwan as potential investments, with a constraint that no more than 15% of the portfolio could be invested in any one of these countries. Finally, we set the required standard deviation to 18.19%. Once again, our results were impressive: we achieved an expected average return of 24.55%, or 1.35% of return per unit of risk, based on these portfolio weights: Australia, Canada, Switzerland and the U.K., 15% each; 13% to the EMU, 12% to Japan; 10% to Mexico, 3% to Brazil, and 2% to Sweden. Of even more interest was the fact that the underlying industry weightings of our optimized portfolio were very similar to those for the EAFE+EMF Index as a whole.

Industry	EAFE+EMF Weight	Weight in Our Portfolio
Energy	5.99%	5.38%
Materials	5.06%	7.81%
Industrials	9.44%	8.08%
Consumer Cyclicals	13.41%	12.00%
Consumer Staples	6.87%	8.26%
Healthcare	8.85%	9.76%
Financials	26.21%	26.09%
Information Technology	9.77%	8.52%
Telecommunications	10.24%	11.41%
Utilities	4.16%	2.68%

The similarity in industry sector weightings between the two portfolios suggests that the expected performance we achieved through optimization was largely due to country effects, rather than an underlying industry tilt.

The next question we asked was whether these apparent country effects were mostly due to the diversification of exchange rate risks. To test this, we looked at the standard deviation of returns (in U.S. dollars) over the January, 1988 to December, 2000 period for both exchange rate changes and for the equity market as a whole in Australia, Canada, the EMU Group, Japan, Mexico, Sweden, Switzerland, and the U.K. In every case, the standard deviation of returns for the equity market as a whole was significantly larger than for the exchange rate alone. The multiple ranged from a low of 1.67 for the EMU and 1.66 for the U.K. to a high of 4.15 for Canada. This seems to indicate that the benefits of diversifying across countries go beyond simple reduction of exchange rate risks.

Overall, when they are combined with the improved performance we were previously able to achieve through sector tilts alone (as detailed in last month's Index Investor), our country tilt results suggest that still further improvements in expected risk adjusted portfolio returns might be achieved through a combination of sector and country tilts, rather than the simple use of either of these in isolation. This is the topic we will explore next month.

In the meantime, it bears repeating that the factors which gave rise to our country tilt results may not operate in the future as they have in the past. To cite but one example, governments and economic policies can and do change over time, and not always in ways you like. Still, having said that, the results remain intriguing, and well worth the consideration of an investor searching for ways to boost risk adjusted returns on the international portion