

The Index Investor

Why Pay More for Less?

Model Portfolios Performance Update

Through October 31st, our benchmark Vanguard S&P 500 index was down (1.8)% for the year, while the Vanguard total bond market index was up 7.7%. Our risk-based portfolios try to match the volatility of different combinations of these benchmarks while providing superior returns. Thus far, they continue to perform as we had expected.

Our high risk portfolio attempts to match the risk of a benchmark made up of 80% S&P 500 and 20% Total Bond Market Index while generating superior returns. Thus far, it is up 6.3% on the year, versus .1% for its benchmark. It has benefited from the very strong performance delivered by the Oppenheimer Real Asset Fund (up 36.3%), as well as the Vanguard Mid-Cap Index (up 18.5%) and the Vanguard Small-Cap Value Index (up 13.9%).

Our medium risk portfolio attempts to match the risk of a benchmark made up of 60% S&P 500 and 40% Total Bond Market. Year to date, this benchmark is up 2.0% through the end of October. Our medium risk portfolio is up 5.1% year to date, largely on the strength of its holdings of the Real Assets Fund, the Vanguard Long Term Bond Market Index (up 10.8%), and the Vanguard Small Cap Value Index. Our biggest disappointment with this portfolio is the performance of the T. Rowe Price International Bond Fund, which is down by (9.5)% year to date (as are all other international bond funds, largely due to the continued strength of the U.S. dollar). However, we continue to believe strongly in the long-term value of this asset class, because of the protection it provides in case of a substantial drop in the value of the dollar (remember that in the past we said the same thing about the Oppenheimer Real Asset Fund...).

Our low risk portfolio attempts to match the risk of a benchmark made up of 20% S&P 500 and 80% Total Bond Market. It is up 8.3% year to date, versus 5.8% for its benchmark. The overwhelming story here is once again the performance of the Oppenheimer Real Asset Fund. As we have said before, the power of having an asset class in a portfolio whose returns are negatively correlated with all its other holdings is difficult to overstate. A key lesson to learn is that when optimizing the allocation of a portfolio to different asset classes, you should spend more time trying to limit risk than trying to find the last ten basis points of return. Why? Because risk (as measured by standard deviation of returns and their correlation with returns on other asset classes) tends to vary less over time than do returns. For this reason, it makes for a far more stable optimization solution.

Our return based portfolios are structured to maximize the probability of achieving a specific target rate of return while taking on the lowest possible amount of risk. They are designed for investors who have a very clear idea of the minimum average annual rate of return they must earn on their portfolio to fully fund their liabilities over a specified period of time. In the case of these portfolios, our decision to prevent them from investing in the Oppenheimer Real Assets Fund (due to the fact that many readers were, in the past, uncomfortable with this asset class) has, every month, come back to haunt us. While their returns are in line with their relative risk, they are still well below where we would like them to be. For the 12% target return portfolio (that is, the portfolio which, over a twenty year holding period, has the highest probability of achieving compound returns of 12% per year, with the lowest possible risk given the asset classes it can invest in), performance year to date is (5.5)%. For the 10% target return portfolio, the year to date return is (2.7)%. For the 8% target return portfolio, the year to date return is (.9)%, and for the 6% target return portfolio, the year to date return is 1.8%. Clearly, these portfolios will be substantially restructured next year (in fact, that process is currently underway, as we attempt to learn from and quantify the size of our mistake in judgement).

Finally, a number of readers have requested that we update our recommended portfolios before the end of 2000, to help them with their year end tax planning. This makes a lot of sense, so we'll be getting the revised portfolios up on the site with our November letter.

The New Equity Market Dynamics

To help put our product and strategies piece this month in a clearer context, we're switching our normal order around, and covering our investment theme first this month.

Over the past three months, we've taken a look at three factors that have a critical impact on the way markets operate and prices are determined within them. In July (and before that in April), we looked at individual investor psychology. In August, we looked at the way information flows through markets, and how people make (or don't make) use of it. And last month, we looked some of the more important structural aspects of today's market, and how they affect people's behavior. This month, we're going to try to pull this all together for you, and paint a "coarse grained picture" of how today's equity market operates, and what that means to us as investors.

We'll state our conclusions up front: first, the changes that have taken place in the equity market have probably changed its fundamental dynamics, and made it much more prone to inefficient pricing and episodes of sharp downside volatility. Second, given this change, historical data may be even less reliable as a guide to future performance. Third, this means that we should be more skeptical about "fine grained strategies" (see below) than we were in the past.

Now on to our analysis. One thing that hasn't changed is human psychology, and the way it seems to affect many investors decision making behavior. Of the many studies we have reviewed over the past few years, a number of conclusions stand out, not only for their analytical robustness, but also because of their fit with what we observe around us every day. First, most investors suffer from what is called a "confirmation bias." In a

nutshell, this means two things: (1) it takes much less information for us to form an initial opinion about something than it takes to get us to change that opinion and (2) we have a natural tendency to pay more attention (and give greater weight to) information that supports our beliefs than we do to information that contradicts them. Here's a concrete example. Consider four situations, and how investors typically react to them today. (1) A company with a negative earnings surprise last quarter has a positive one this quarter: "Let's wait and see before we invest." (2) A company with a negative earnings surprise last quarter delivers another one. "Toxic. Bail." (3) A company with a positive earnings surprise last quarter has a negative one this quarter. "Let's wait and see before we sell." (4) A company with a positive earnings surprise last quarter delivers another one this quarter. "Buy more." With respect to equities, the implication of our confirmation bias is that prices will tend to overshoot their "fair" or "fundamental" values (in either direction), because people only slowly incorporate information that is contrary to the view they hold.

Second, the impact on prices of our confirmation bias is accentuated by our tendency to herd, or to make investment decisions not on the basis of the quality of information we have available, but on the basis of what we infer from market price changes about information that other investors have that we don't.

Third, herding operates very differently on the upside than it does on the downside. On the upside, herding takes more time to build up momentum, for two reasons (1) as previously mentioned, it takes a lot of information to get us to change an opinion we have formed, (2) we tend to be loss averse -- that is, we will take gains more quickly than we will losses, and (3) we tend to be myopic, and evaluate our investment performance more often than we should. For example, even though we are saving for a retirement that is twenty years away, we still review our investment performance annually, and make decisions on the basis of whether we're up or down on the year. For these reasons, bubbles form slowly.

Crashes, however, happen much more quickly. First, we start with a natural human tendency to avoid recognizing losses, because they hurt more than gains feel good (2.7

times more, some researchers have estimated!). Given this, you would think that crashes would develop more slowly than bubbles. If investors received information at the same speed in both cases this would probably be true. However, when a stock price starts to fall, a lot of new information quickly comes out, in the form the presence or absence of covering purchases by short sellers (because unlike long purchases, the price of short sales aren't publicly disclosed). This can speed up the process of changing people's opinions, and cause them to sell more quickly than they bought. However, when it comes to the speed with which crashes can develop in today's markets, human nature is only part of the story. We also have to take a lot of new structural factors into account.

Let's look at how these have changed in recent years. In the past, investors had access to less information (not only about companies, but also about what other investors were doing), which came to them more slowly (remember the quaint practice of reading the business section of a newspaper/), and which cost much more to act on (i.e., trading costs were much higher). If you were an institutional investor, in the past you had fewer competitors (there were fewer mutual funds), your performance wasn't compared to theirs (and to the indexes) on a daily basis, it was much more expensive for your clients to switch their funds away from your firm, you didn't have to contend with hoards of individual investors, and, before SEC Regulation FD was issued this year, you still (if you were a big player) had preferential access to information about different companies. Most of your peers followed the same investing approach, using fundamental analysis to estimate the fair value of different stocks, and buying those that seemed undervalued.

But that world, genteel and enjoyable as it was (at least in retrospect), has, as they say, "gone to hell in a handbasket." Today, you're under the gun to deliver performance on a quarter to quarter basis. If you're an active manager, that means at minimum beating your peers, and ideally beating your index. Today, lots of individual investors receive the same flood of information just as fast as the institutions, and can act on it just as quickly. Today, market moves that would have taken a quarter or more to develop happen in a morning or a day. Today, a lot more money is chasing trends than has ever been the case before. How does an investor survive in this world?

A while back (though today it seems much longer), the most popular answer (among professionals, at least) seemed to be "stick to the fundamentals", or, more crudely, "bet against the dumb money." There was a belief that value investing would triumph over less sophisticated trend followers (also known as "momentum investors") by arbitraging them out of existence. Two things were required for this to work: (1) enough patience for the arbitrage strategy to work, and (2) sufficient capital to back up that patience. Unfortunately, neither has been available in the amounts required. Over the past three years, untold billions have been lost by short sellers, put buyers, and others who bet on fundamentals and against the trend. The net result is that we now have a market in which more money than ever before is probably invested on the basis of trend following (or, as one calls it in polite circles, "growth stocks") rather than fundamental valuation. From an individual's perspective, trend following looks like an efficient use of information; from many an institutional manager's perspective, it is the key to survival and next month's mortgage payment.

The collective result of all this is that we now have a market in which stock prices are much more prone to overshooting their "fundamental value", and to staying that way for long periods of time. But not forever, for bubbles always burst. And when they do, today's market makes sure they do so much more violently than in the past. Why? Four reasons. First because more money is herding (oops, I meant trend following, or "growth stock" investing...whatever), while information is moving more rapidly. Second, because markets are more interconnected than they ever have been in the past, so there are more potential trouble spots, and trouble spreads more quickly. Third, because markets today are much more leveraged than they ever have been in the past (think margin accounts, hedge funds, and derivatives). And fourth, because that leverage rests on a capital base that is relatively smaller, and, perhaps more important, relatively more concentrated (in a limited number of commercial and investment banks) than it ever has been in the past. As we wrote last month, the dirty little secret in the market today is the disappearance of liquidity when things hit the fan (contrary to what the theory that underlies all those "risk management" models assumes). When that happens, prices can and do drop like a rock.

In short, it's not your parents' (or even your older siblings') market any more. Is it an irrational market? It depends on what you mean by rational. One meaning of rational is "efficient", in the sense that prices generally reflect fundamental values, and there are no patterns in prices that can be exploited to make a profit. If this is your test, then no, the market today isn't rational. However, if by rational you mean a market composed of rational people, rationally trying to figure out the best approach to use to make the most money they can, then the market is, indeed rational. All that has changed is the dominant model that people use to identify situations where prices appear higher or lower than they should be. In the past, the most popular model was rooted in economics, and company fundamentals. Differences of opinion about the "right price" for a stock were rooted in different views of a company's future prospects, and the outlook for interest rates. Today, those differences are rooted in psychology, and different expectations about the future actions of other investors. You may not like it. It may not be efficient (in the economic sense). But it is, in its own way, completely rational. If we're playing hockey instead of football, it doesn't pay to try kick a field goal.

Which brings us to our last point in this section. If the market has fundamentally changed, then you have to be very, very careful about using investment approaches that are based on data about past market relationships. If those relationships don't exist anymore, the approaches that are based on them probably won't work (or, if they do, will do so only by luck). With that in mind, let's move on to our next section.

Sizes, Styles and Sectors, Oh My!

Over the past three months, we've compared the different indexes that subdivide the equity market based on size (market capitalization), style (growth versus value), and economic sector. The question we're going to address this month is whether or not it makes sense to use any of these, instead of simply investing in an index that covers the

market as a whole (e.g., the Russell 2000, the Dow Jones Total Market Index, or some combination of the S&P 500, 400, and 600 indexes).

Let's start with a basic axiom. At the end of the day, there are only two reasons people choose one investment over another. Either they think its price is less (more) than its fundamental value, or they see other people buying (selling) it. Given the relatively uncertain pricing dynamics in the current market that we have noted above, we think you have to make a pretty good case in one of these areas to justify departing from the "no brainer" option of simply buying the market index. With that in mind, let's look at the arguments for making size, style and sector bets.

Let's start with buying the market index as a whole, which is the alternative against which the size, style, and sector tilts will be compared. First (as is our prejudice), some data. Our benchmarks for the following analysis will be the Dow Jones Total Market Index, the Russell 3000 Index, and the S&P 1500. The first two can be purchased as iShares on the Amex. The S&P can be constructed by purchasing the right amounts of the S&P 500, 400, and 600 indexes. The DJTMI maintains a constant size based market capitalization based weighting. Companies in the top 70% of market cap make up 74% (by market cap) of the DJTMI, companies in the next 20% of market cap make up 21%, and companies in the bottom 10% of market cap make up 5% of the index. By comparison, in the Russell 3000, the top 200 companies (by market cap) make up 70% of the index, the next 800 companies make up 22% of the index, and the bottom 2000 companies (aka, the Russell 2000) make up 8% of the index by market cap.

Within the S&P 1500, the S&P 500 accounts for 90% of market cap, the 400 for 7% of market cap, and the 600 for 3% of market cap. Despite their different composition, all three indexes delivered similar performance over the period for which comparable data is available. Between January, 1994 and September, 2000, the DJTMI had an average annual return of 21.01% and a standard deviation of 17.12% (or 1.23% of return per 1% of risk); the Russell 3000 had an average annual return of 20.83% and a standard deviation of 16.78% (or 1.24% of return per unit of risk), and the S&P 1500 had an

average annual return of 19.85% and a standard deviation of 16.63% (or 1.19% units of return per unit of risk).

Could you have done better over the 1/94 to 9/00 period by buying a different mix of market capitalization weightings? Let's look at the large caps first. If you bought the S&P 500 in January, 1994, you would have realized an average annual return of 21.62% through September, 2000, with a standard deviation of 17.06%, or 1.27% of return per 1% of risk. Buying the Russell 1000 would have delivered pretty much the same performance: average annual return of 21.62%, and standard deviation of 17.01%, for the same 1.27% of return per 1% of risk. In short, by buying large caps instead of the market as a whole, you could have realized 3 more basis points in return per unit of risk. If there is an explanation for this, it is probably that institutions using a trend following strategy preferred the large caps because of their liquidity, and were willing to pay up a bit for it.

By way of comparison, buying the S&P midcap 400 would have generated average annual returns of 21.09% and with a standard deviation of 19.76% or 1.07% of return per 1.00% of risk, while buying the S&P smallcap 600 would have generated average annual returns of 14.93% and with a standard deviation of 19.96% or .75% of return per 1.00% of risk. You would have done worse had you chosen the Russell 2000 as your small cap index over this period, earning average annual returns of 14.37% with a standard deviation of 20.89%, or .69% of return per unit of risk. So far, it seems that you would have done a little bit better by buying the large caps instead of the market as a whole, but not by much, as the former make up the large majority of the total market indexes. This is proven out by our optimization software. We used the S&P 500, the Midcap 400 and the Russell 2000 as our asset classes. We first attempted to find an allocation that bettered the 20.83% average return delivered by the Russell 3000, without exceeding its standard deviation of 16.78%. The best we could do was 21.03%, using a mix of 92% large cap and 8% small cap. The results were actually worse when we reversed the test, and tried to take on less risk while matching the R3000's 20.78% return. The best we could do was to reduce standard deviation to 16.71%, by taking on 89% large cap exposure and 11% small cap.

What about a style tilt instead of a size tilt? If you bought the Russell 3000 Growth Index, you would have realized average annual returns of 24.04% over the 1/94 to 9/00 period, with a standard deviation of 20.59%, or 1.17% of return per unit of risk. If you had purchased the Russell 3000 Value Index, you would have realized average annual returns of 17.38%, with a standard deviation of 15.89%, or 1.09% of return per unit of risk. In short, by pursuing a pure growth or value strategy, you would have done worse on a return per unit of risk basis than if you had purchased the market as a whole in the form of the Russell 3000 index. Why? Because by taking a pure style approach you not only give up the returns on the other style, you also give up the diversification benefit that comes from investing in both styles. As is so often the case in life, this is another example of the whole amounting to more than the sum of its parts.

Beyond this, there is another issue with respect to style indexes. In a nutshell, they don't correspond well to the two major logics that drive investment decisions. To begin with, price/book (or any of the other proxies used to separate companies into the growth and value categories) tells you nothing about either (a) whether or not the company is undervalued relative to its fundamentals, or (b) whether or not it has substantial price momentum. For example, a firm can have a low price to book because investors don't fully understand its growth prospects, or because its products are obsolete. Similarly, a company with a high price to book could be wildly overvalued by investors in the middle of a bubble, or it could be spectacularly undervalued because investors don't fully appreciate the value of its growth potential. Or it could have just taken a big write off on its way to bankruptcy, and be the farthest thing from a momentum stock.

Well then, what about combining a size and style tilt. If large cap growth was your taste, and you had invested in the S&P 500 growth index, you would have earned average annual returns of 24.82% over the 1/94 to 9/00 period, with a standard deviation of 19.54%, or 1.27% of return per unit of risk. If you had invested via the Russell 1000 growth index, you would have earned average annual returns of 25.02%, with a standard deviation of 20.56%, or 1.22% or return per unit of risk. On the other hand, if your taste

ran to large cap value, and you had invested in the S&P 500 value index, you would have enjoyed average annual returns of 18.35% per year, with a standard deviation of 16.81% or 1.09% of return per unit of risk. If you had invested in the Russell 1000 value index, the comparable figures would have been 17.98%, 16.42%, and 1.10%.

What about midcaps? If you had purchased the S&P 400 growth index, you would have earned average annual returns of 25.70% over the 1/94 to 9/00 period, with a standard deviation of 26.18% or .98% of return per unit of risk. The S&P 400 value index would have generated exactly the same amount of return per unit of risk, with an average annual return of 17.02% and a standard deviation of 17.40%.

If your taste had run heavily to small caps, you wouldn't be happy today. By buying the S&P 600 growth index, you would have received average annual returns of 14.55%, with a standard deviation of 24.43%, or only .60% of return per unit of risk. Buying the Russell 2000 growth index instead wouldn't have helped. In this case, you would have earned average annual returns of 16.23%, with a standard deviation of 28.56%, or just .57% of return per unit of risk. The story on the small cap value side isn't much prettier. Buying the S&P 600 value index would have generated average annual returns of 14.93% with a standard deviation of 17.17%, or .87% of return per unit of risk, while buying the Russell 2000 value index would have generated average annual returns of 12.36% with a standard deviation of 14.59%, or .86% of return per unit of risk. When we combined all six inputs in our optimization software, we once again found that we could achieve very little improvement on the performance of the Russell 3000: a .92% gain in return with the same level of risk, or a .67% reduction in risk for the same level of return. In both cases, the weighting of large caps was about the same as in the overall market index.

What have we learned up to now? Buying the market index as a whole, in the form of either the Dow Jones Total Market Index or the Russell 3000 would have generated, respectively, either 1.23% or 1.24% of return per unit of risk. Of all the eight possible size/style combinations, the only one that exceeds this (and only in the case of one of the

two possible indexes in this category) is the S&P 500 Large Cap growth index, with 1.27% of return per unit of risk over the 1/94 to 9/00 period.

Thus far, not exactly an overwhelming endorsement of size and style tilts away from their underlying weights in the market index portfolio.

But couldn't you come up with a weighting of size/style indexes that is different than, and outperforms the market index? Practically speaking, the answer to that question is no. The bottom line is that over the past six years, large cap stocks have been the primary drivers of market index returns (perhaps, as we said, because of their attractiveness to institutional trend followers), and are already quite heavily weighted in the major market indexes. In a nutshell, there is simply very little room to improve risk adjusted performance by overweighting large cap growth relative to its current percentage of the overall market index.

Ah, you say, but what about sector indexes? It is true that had you invested in the S&P Technology Index, you would have realized average annual returns of 40.84% between 1/94 and 9/00. But you would also have experienced a standard deviation of 36.46%, or only 1.12% of return per unit of risk you took -- inferior to the market as a whole. You would have done better, however, with the Dow Jones Health Care Index, which generated average annual returns of 27.57%, with a standard deviation of 20.75%, or 1.33% of return per 1.00% of risk. Then again, had you invested in the S&P Basic Materials Index, you would have realized average annual returns of only 7.23%, with a standard deviation of 21.25% -- that's .34% of return for every 1.00% of risk you took on.

Let's face it: the trend buying was in technology and health care, not basic materials. On the other hand, maybe the latter is more overvalued. Who knows? And that's the point. There is no doubt that we can use historical data and out optimization software to come up with a sector by sector allocation that would have delivered more return per unit of risk than the market as a whole over the 1/94 to 9/00 period.

But hindsight, as they say, is 20/20. Because we already knew what had happened, our backwards looking allocation would capitalize on the bubbles and undervaluations, while limiting the damage from the "low but accurate" valuations and crashes.

The point is that because of the amount of money backing trend following strategies these days, it is very hard to be confident that whatever sector allocation delivered the best performance over the last 81 months will do so over the next 81 months. For example, it is not out of the realm of possibility to guess that technology and health care stocks are in the later stages of a bubble, and that buying these indexes now could very well lead to tears later on, as the momentum traders move into energy and telecommunications. As we said, the market these days seems much more likely to be in a disequilibrium state than it has in the past. And when that is the case, the past is usually a poor guide to the future.

The bottom line is that picking sizes, styles, and sectors within the U.S. equity market are all forms of active management. And winning at active management is very, very hard to do today -- possibly harder than it has ever been, because the past is no longer as good a guide to the future as it once was. Our conclusion is that under these circumstances, discretion is probably the better part of valor. In short, given the nature of market pricing dynamics these days, we are probably better off buying the market index as a whole, and adopting a fully passive indexing approach without any active tilts.