origin Financial Planning for Every Employee, Everywhere.

Origin Investment Methodology and Strategy

David Barta, CFP®

Lead Origin Planner

Table of contents

Introduction

At Origin, we strive to be the single location where you can manage all of your finances. Critical to this is having access to suitable investment options. We have done considerable research to create multiple portfolios that you can use to invest. Here, we will describe the methodology used to create these portfolios.

First, let's review indexes, active management, and factor investing. Investing in an index, like the S&P 500, is a low-cost way to invest in a set list of companies (the 500 largest companies in the US, in this case), without spending time or energy to determine which companies should perform better than others, (and then choosing to just invest in those). Conversely, an active manager decides which companies to purchase based on extensive research. This process typically involves higher costs. One major critique of active investing is that, after the higher fees are deducted from performance, most active funds fail to beat their index benchmarks (Fama & French, 2010). As such, investors tend to do better by investing in index funds instead (Bogle, 2009). At Origin, we fully believe in this strategy. When it comes to whose particular index fund you purchase, there are many index funds that invest in the same thing and have comparably low fees. For example, for the S&P 500 index, you could purchase SPY from State Street, VOO from Vanguard, or IVV from BlackRock, and they are practically identical. Whose index fund you purchase is of little significance.

So what exactly is factor investing? The Fama and French Three-Factor Model paper was published in the early 1990s and quantified what many investors had intuitively known for a while: 'small' companies and 'cheap' companies tended to outperform 'large' companies and 'expensive' companies. This formed the research basis for factor investing, starting with the 'size' and 'value' factors (the 3rd factor is 'market exposure,' basically that stocks outperform bonds). Most indexes, including the S&P 500, are market-weighted, meaning they invest proportionately to the market-cap of the company (invest more in larger companies). Indexes need not be market-weighted though, for example, they could be *equal-weighted* instead (DeMiguel, Garlappi & Uppal, 2009). Research has shown that equal weighting companies in the S&P 500 can perform better than market-cap weighting them, in significant part because you are investing in greater proportions in the smaller companies of the S&P 500, and are thus getting higher returns because you are exposing yourself more to the *size factor*. Dimensional Fund Advisors began implementing the research of Fama and French. As a result, investors could, for example, buy a fund of the largest companies in the US that was weighted by each company's respective *exposure to factors*, instead of merely market-cap, and as such, investors expect to get a higher return. Currently, many different fund companies offer funds and ETFs with exposure to various factors.

The Carhart 4-factor model paper was published in 1997, and suggested that in addition to the market, size, and value factors, there was a 4th factor that *persists across all asset classes, across all time, and across almost every different market we have data on*, and that factor is momentum. Momentum is the idea that something that is performing well over the recent past will continue to perform well, and something that is performing poorly will continue to perform poorly. Again, this was something intuitively known earlier but had not been called a factor yet (Jegadeesh & Titman, 1993). Momentum is a fascinating factor, because it has less to do with the intrinsic risk properties of the investment, and more to do with the *behavior of investors*. What is the common thread running between all asset classes, across all time, and across every different market? People. People are predictable. For example, people tend to exhibit FOMO (fear of missing out) toward a rising investment, which leads to investors piling into something sexy like Bitcoin. And similarly, people tend to exhibit panic over a falling investment, which leads to investors selling for fear of 'losing everything.'

These papers ignited a financial arms race to discover the thousands of potential secret factors that could exist that might give investors the slightest edge. We chalk the majority of this up to *data-torturing* (data manipulated enough can 'prove' anything) as the original well-researched factors help explain the majority of asset price returns, (Harvey, Liu & Zhu, 2016; Chordia, Goyal & Sarett, 2020).

Here at Origin, we use indexes and the momentum factor. Indexes give us great exposure to the majority of global asset classes. However, we believe we can do slightly better by also using the momentum factor. Why are we not using the size and value factors? Since other companies have started using the size and value factors and publicizing them, a significant number of investors have started to invest in them. As a result, they have started to decay the excess return from the factors i.e. prices are pushed up and forward expected returns drop. Our analysis indicates that over the last 25 years, the size and value factor did not perform well (something Ken French has also admitted). If you are interested in taking a deeper look, you can download Ken French's data. By contrast, the momentum factor continues to look robust. We hypothesize this is because, as more people invest using momentum, they should amplify and strengthen the effect. By contrast, most other factors display a weaker effect as more investors use them.

We do not use active management, as the costs do not repeatedly and consistently outweigh the benefits. Warren Buffet famously made a $1M bet with a hedge fund, the epitome of active management, to see if they could beat the S&P 500 index over a 10 year period. They failed miserably and Warren Buffet's $1M win went to a charity for girls (Girls' Inc of Omaha).

Asset classes and Investment Vehicles

Let's review our asset class universe and how we invest. We could buy the underlying securities ourselves, but with such a plethora of ETFs and mutual funds at our disposal, we do not need to. Research also suggests combining asset classes, instead of individual securities, maximizes return per unit of risk (Brinson, Hood & Beebower, 1986; Brinson, Singer & Beebower, 1991; Ibbotson & Kaplan, 2000). ETFs have the added benefit of ease of access from any brokerage account, unlike mutual funds. ETFs can also be slightly lower cost than mutual funds and slightly more tax efficient. For these reasons, ETFs are our primary investment vehicle.

Next, let's review the asset classes we analyzed. We cast as wide a net as possible, and then trimmed down the number of assets in the final portfolios to reflect those that were the most beneficial. We started with 19 asset classes: the Total US stock market (VTI), Developed market stocks (VEA - international stocks from developed countries ex-US, think Japan, Europe, etc.), Emerging market stocks (VWO, think China, Brazil, etc.), Real Estate (VNQ), US stock momentum (MTUM) and Developed stock momentum (IMTM), Gold (GLD, because it has some really interesting counter-cyclical properties), Developed bonds (BNDX), Emerging market bonds (VWOB), TIPS (SCHP, treasury inflation protected securities), long-term treasuries (VGLT), intermediate-term treasuries (VGIT), short-term treasuries (VGSH), municipal bonds (VTEB), US investment-grade corporate bonds (LQD), US below-investment-grade corporate bonds (HYG), mortgage backed securities (VMBS), Cash/ultra-short-term treasuries (SHV), and the total US investment-grade bond market (BND). For a continued discussion on these asset classes, see the appendix below.

You will notice natural resources, energy, and commodities are absent from this list.

We do not think there is a good reason to call out these or any subsector versus invest in the total market. The noticeable exceptions are real estate because the market is so large and unique (offering diversification benefits), and gold, which has very specific risk-off properties (offering volatility dampening benefits).

These are the asset classes we analyzed, not the final list used in our investment portfolios.

Brief overview of risk and return

Of these 19 asset classes, how do you determine which to invest in? And in what proportions?

Intuitively, we want our portfolio to only consistently go up, and never go down. If we could, we would create a portfolio with 20% average annual returns and no volatility! But alas, we live in a world where higher risk usually accompanies higher returns. But, intuitively, we want to build portfolios with as much return potential for the least amount of risk potential.

What is the return potential of an asset class? For example, what is the forward looking return potential of VTI, the total US stock market? This is an extremely difficult question to answer. One option is to look at past returns to give an indication, but that doesn't take the current state of the economy, politics, unemployment, inflation, etc. into consideration. We used research from BlackRock and JP Morgan to inform our forward looking expectations (and standard deviations) of each asset class.

Next we determined the risk of investing in each asset class. This is extremely difficult - how do you define risk? One option is to try to minimize the historical standard deviation of returns, but it is impossible to determine whether that will be accurate in the future. Standard deviation measures the dispersion of data around its mean, so a higher standard deviation of returns indicates returns vary more wildly about its average return (booms and busts). While that wild volatility could be a good proxy for risk, does one worry about the booms? If your asset class abruptly shoots up in value and does not go down, does that cause concern? No, you are overjoyed and you do a dance. So in fact, if we could split standard deviation between upside volatility (volatility of positive returns) and downside volatility (volatility of negative returns), we could more accurately optimize for the thing that causes the most anxiety: downside volatility.

Another risk measure is CVaR (conditional value-at-risk). Professor Philippe Jorion, is a major proponent of VaR analysis, which is a way of summarizing an expected loss over a certain time horizon at a given probability (for example, we have a 1 in 20 chance of losing 3% over the course of a day in this certain investment). CVaR takes this a step further to try to answer this: given that this bad left-tail outcome has occurred, how bad does it get on average? Conditional on breaching that 1 in 20 chance, what is our expected loss? For example, using the same investment as before, assuming we have breached our 1 in 20 chance, the CVaR of this investment is 5.4%. The difference between VaR and CVaR is that VaR is the expected loss at exactly the 1 in 20 case (the 5% case), CVaR is the expected loss given you are in the bad timeline (past the 5% case). So another option to minimize risk could be to minimize CVaR.

Everything we previously discussed is at the 'per asset class' level, what happens when you start to cocktail investments together into a portfolio? Does the math get much harder? Expected returns of a portfolio is easy, it is just the weighted average of the expected returns. The interaction between the assets makes the risk a bit harder to quantify as it will depend on how *correlated* the assets are to each other (and how that correlation could change).

For example, say you invest in Gold and in Stocks. You notice that they have a negative correlation (so when one goes up the other goes down). If you put them in a portfolio, the standard deviation of the portfolio will go down significantly thanks to the negative correlation. Similarly, if you invest in Real Estate and Stocks, you notice that they are positively correlated, but the correlation is less than 1 (so not perfectly correlated). Combining these investments in a portfolio will yield a lower standard deviation than the greatest of the two's individual standard deviations because their correlation is less than 1.

We can generalize and use historical correlations with standard deviation (historical or projected) to get a covariance matrix to see how each asset class might 'co-vary' with each other.

Brief overview of optimizations

Optimizers take risk, return, and relationships between asset classes as inputs and return the 'best' portfolio they can find. We'll first review the underlying theory of optimizers and then discuss some different types.

Modern Portfolio Theory (MPT) as pioneered by Markowitz (Markowitz 1952), suggests that the characteristics of an investment (risk and return) should not be viewed in isolation, but in how the addition of that investment *affects your portfolio*. And if we assume investors don't like risk, then for a certain level of acceptable risk, we can mix investments together in a way that ***maximizes** the expected return for that level of risk*. We can also do it in reverse, for a certain desired expected return, we can find a combination of assets that minimizes risk. It becomes less important how an individual asset performs, and more important how (given returns, risk, and covariances) the portfolio performs.

This leads to the first optimizer: Mean-Variance optimization. At a basic level, we assign random weights to each asset in a collection of assets to make a random portfolio. We then estimate the return and risk of that portfolio by using historical or projected data. We then plot that (risk, return) data point on a chart. We then repeat the process thousands of times to get thousands of points plotted on a chart, as the example shows below:

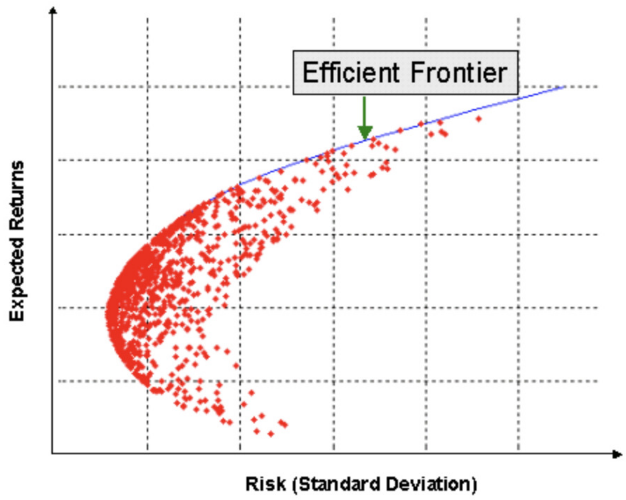

Something interesting starts to occur - there appears to be a 'limit of returns,' called the Efficient Frontier (the blue curve). This curve is upward sloping, as we would expect: for more risk (moving to the right) we should get a higher expected return (moving up). This illustrates our earlier argument that higher returns are typically accompanied with higher risk. Our portfolios lie on the efficient frontier - they are portfolios that have the maximum expected return for a given level of risk. Standard performance measures that we would try to maximize are the Sharpe ratio - expected returns above the risk free rate divided by standard deviation (essentially return per unit of standard deviation risk), and the Information Ratio - returns over a benchmark divided by the standard deviation of that difference (essentially 'smart returns' per unit of 'smart risk').

So what is wrong with this optimizer? Primarily, it is incredibly sensitive to inputs (return, risk, covariance). Change one input by 1% and your optimized portfolios could be entirely different, not merely different by 1%. This is not ideal because most of these inputs are based on estimates, and they could easily be off by 1%. To increase confidence in this optimizer, there are a few different approaches - the Black-Litterman (1992) approach to improve estimates of returns, Shrinkage to improve estimates of risk/covariance (e.g. Ledoit and Wolf, 2004), or Resampling/Monte-Carlo simulation to try to improve all estimates (Michaud, 1998).

Another issue is that this optimizer uses standard deviation as a proxy for risk, which again penalizes upward volatility and downward volatility, when in fact, investors are primarily concerned by downward volatility. Post-modern portfolio theory (PMPT) (Rom & Ferguson, 1993) attempts to improve on MPT by minimizing downward volatility. Additionally, the relationship between assets changes depending on market conditions i.e. we cannot use a static covariance matrix. Specifically, when the market is in a selloff panic, most correlations increase and go to 1. Again, this is not ideal because these are the exact events we are trying to minimize and we cannot allow the optimizer to be blind to them.

Although these problems exist, the mean-variance optimization remains the industry standard. Therefore, we conducted this optimization. However, we wanted to improve this optimization by utilizing PMPT. If investors primarily care about minimizing downside volatility, and are particularly worried about the worst-case selloffs, why not optimize based on CVaR? CVaR focuses on 'the average bad scenario,' which is intuitively something beneficial to minimize. Also, depending on implementation, there is no need to make any assumption about the distribution of returns. We also don't have to worry about an unstable covariance matrix, as historical data (modified slightly as we discuss below) contains information on how asset classes move together in times of market selloffs, as well as how they move 'regularly' in other times. Given these were our primary concerns with mean-variance optimization and mean-CVaR has a response to each concern, we also conducted a mean-CVaR optimization. The performance measure we optimized here is expected return per unit of CVaR risk.

Determining inputs to optimizers

Both optimizers require expected returns and expected volatility per asset class as well as a covariance matrix for the mean-variance optimizer. As noted above, for forward looking returns (ignoring taxes and fees) and expected volatility per asset class, we leveraged BlackRock (20 year) and JP Morgan (long-term) estimates.

For historical data, we used 25 years of monthly returns (October 1995-September 2020). We used the first 20 years for analysis and the last 5 years for out-of-sample testing. For the covariance matrix we used the first 20 years of historical returns.

Beginning with BlackRock and JP Morgan estimates of asset class returns, we subtract expense ratios of our respective ETFs (not the stated expense ratio, but the drag in performance in practice that each ETF has from its respective index, which is different from its stated expense ratio, thanks to securities lending and other factors), to get an after-expense return. We ignore trading costs as they are de minimis. Next we subtract taxes due to distributions and taxes due to selling the positions (helpfully provided by each ETF provider, assuming the highest federal tax bracket and no state costs). This gives us an after-expense, after-tax return per asset class (this is our optimization for taxable accounts, not for tax shelters). As is standard, when dealing with factors, it is best to assume the premium harvested declines in the future, so we cut the US Momentum premium by ⅓, and International Momentum by ⅔, which led to a similar forward-looking factor premium across both asset classes.

All of these inputs are then used in the mean-variance optimizer. We optimized twice, once with Shrinkage and once with Resampling.

For the mean-CVaR optimization, we based our risk factor (CVaR) on historical returns. However, historical return/risk data is different from forward looking return/risk projections. We modified our historical returns so the mean/standard deviation matched forward looking expectations (levering to match standard deviations, and then shifting to match means). If we assume relationships between asset classes remain similar to how they have been historically, we have an estimate of future CVaR based on historical CVaR. We have not needed to force normalcy on our data and we have made no covariance stability assumptions. Like with the mean-variance optimizer, we use projected returns for our positive variable (not historical). We use all these inputs to perform the mean-CVaR optimization. Although it was not the primary risk measure we minimized, we closely watched downside volatility.

Determining the actual assets to use

To reduce our 19 asset classes to the final ones in our portfolios, we started with a few essential assets and added the next-most-useful asset one at a time. This placed the burden of 'proving usefulness' on each new asset.

We started with a core of assets (Cash, US Stocks, Developed Stocks, EM Stocks, and both Momentums) and proceeded to see which single asset would add the most value if added to these assets i.e. push our mean-CVaR efficient frontier the most. As an example, the below plots efficient frontiers of the core assets with the addition of each individual new asset:

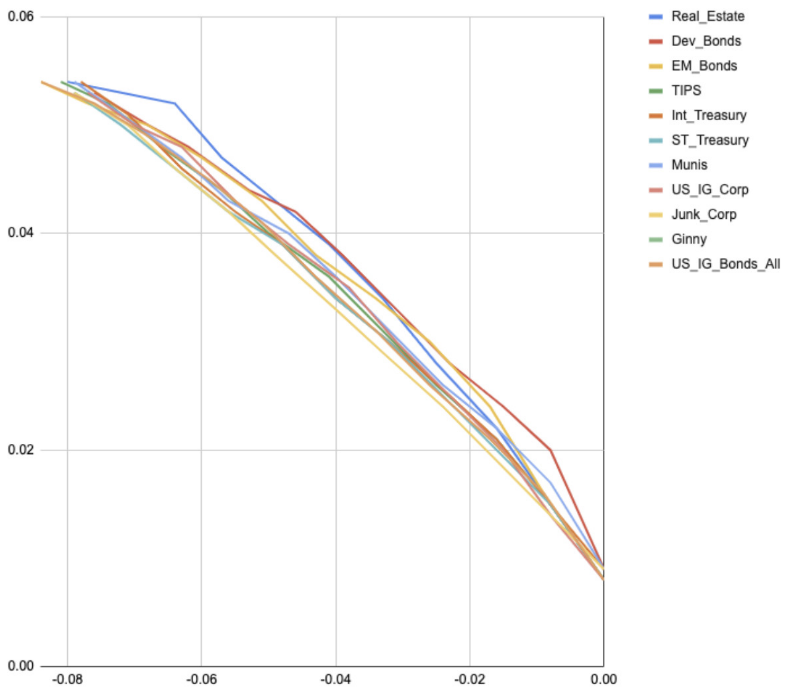

The x-axis is CVaR, and the y-axis is average returns. The higher the line, the more return per unit of CVaR. We see above that adding Real Estate to the core assets increases our expected return in risky portfolios (blue line), and the addition of Developed Bonds adds to our returns in more conservative portfolios (reddish line). We added Developed Bonds to our core assets, and reran the process.

It is worth noting that the caps used for each asset class are incredibly important. If the optimizer runs uncapped, it may create portfolios that are 100% of a single asset, and not use the other assets at all. At Origin, we want to build diversified portfolios, so we added caps to each of our assets. This included both minimums and maximums, sometimes in relation to each other.

We continued to repeat until adding new assets did not materially improve our efficient frontiers. This led to the final list of 11 assets: Cash, US Stocks, Developed Stocks, EM Stocks, US Equity Momentum, Developed Equity Momentum, Real Estate, all US IG bonds, Munis, Developed Bonds, and EM Bonds.

Optimizing

Let's review our final portfolios. We ran 160k different mean-CVaR scenarios using our final 11 assets. This gave us an efficient frontier (in return-per-unit-of-CVaR space) that we used to find portfolios that were evenly risk-spaced apart. We compared these results against their expected down-side volatility projections to ensure we were not blind to what was happening with volatility.

We vetted these results against results from mean-variance optimization. These results included Shrinkage for one optimization, Resampling for another, and a simple uniform cap on most assets in order to determine what the optimizer prefers.

Without using Resampling or Shrinkage, we get portfolios that abruptly add or remove asset classes (shown below). If you draw a vertical line at the far left, you will have a portfolio that is 100% cash (dark green, y-axis is the percent weight), that should have the lowest risk (left on the x-axis) and lowest expected return. By comparison, if you drew a vertical line on the far right, we would have a portfolio of 6 assets that should have the highest expected return, and highest risk.

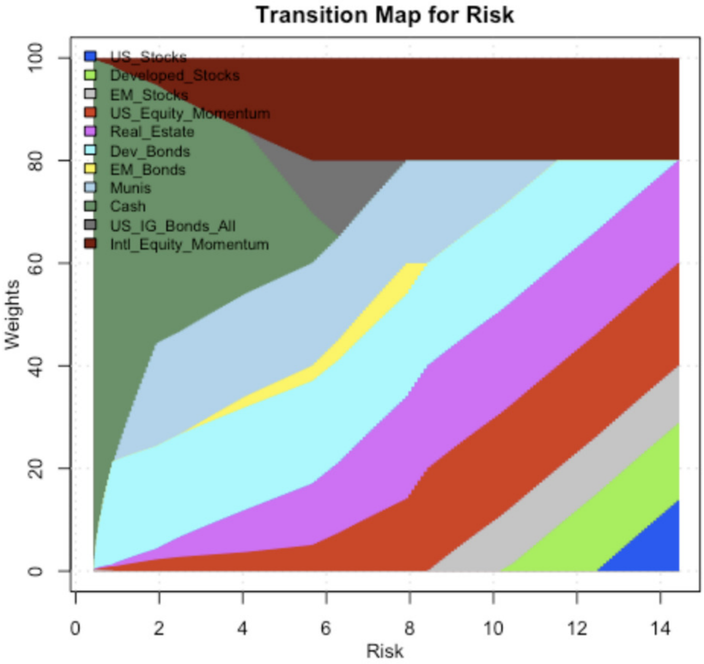

Resampling smooths the asset class weights:

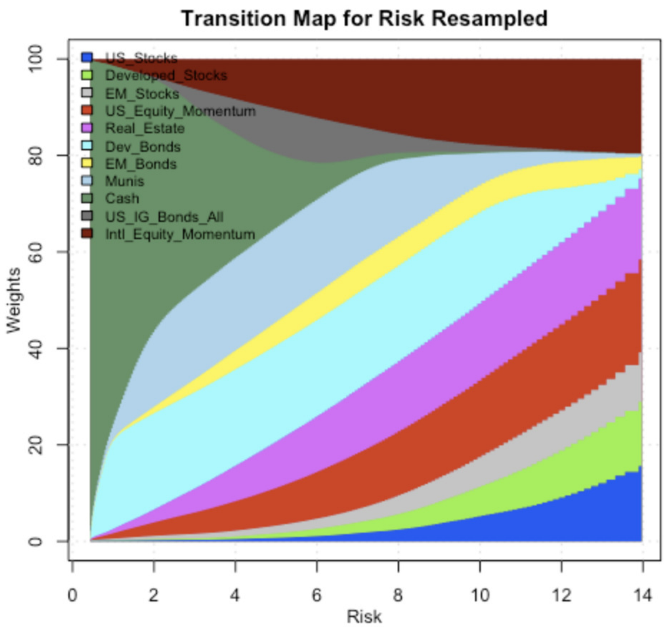

We confirmed what we knew from the mean-CVaR optimization: the optimizer preferred Momentum, didn't like US Stocks, liked Real Estate, liked Developed Bonds (way more than US bonds), found some use for EM bonds, found a significant use for Munis, and of course found that cash is very useful in lower risk portfolios.

Looking at Shrinkage:

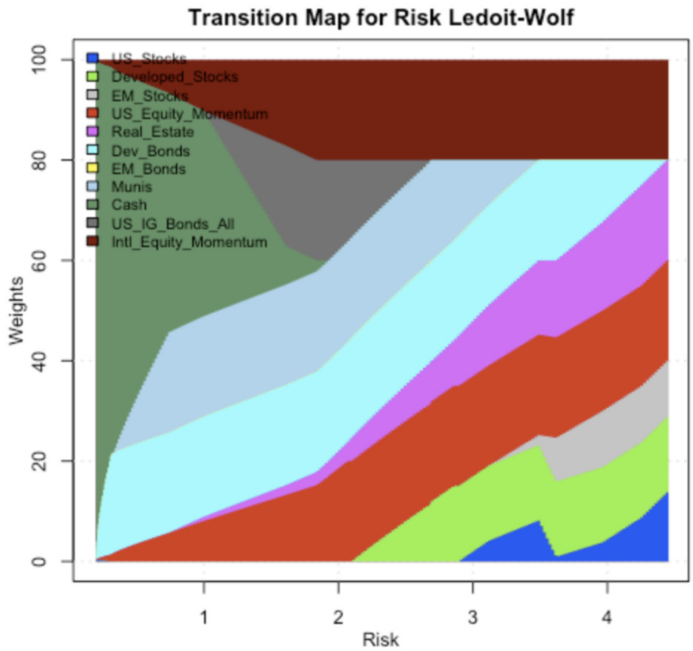

We again confirmed the same trends, but found a slight less preference for Real Estate and no use for EM bonds. We note that Shrinkage produced more abrupt changes in allocations than Resampling, which is to be expected.

We combined all this information, relying more heavily on the mean-CVaR analysis with more robust and relative capping, to arrive at our final 12 portfolios.

Out of sample testing

Now that we had our 12 portfolios based on 20 years of data (first 20 years of our full dataset: October 1995-September 2020), we performed out-of-sample testing to understand how our portfolios would have performed in the last 5 years of this dataset (Kan & Smith 2008).

Looking at returns (y-axis) vs CVaR (x-axis), our 12 portfolios are the black boxes, random portfolios are the blue stars:

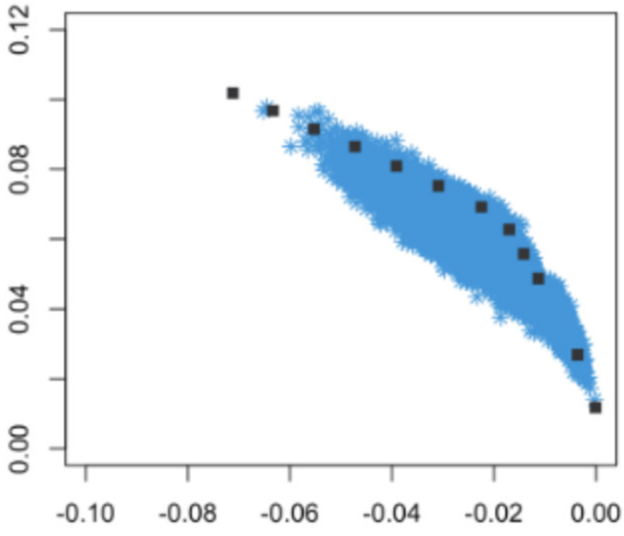

We see that the portfolios perform well. Our portfolios are close to the efficient frontier (top part of the blue). Next we analyze returns vs downside volatility:

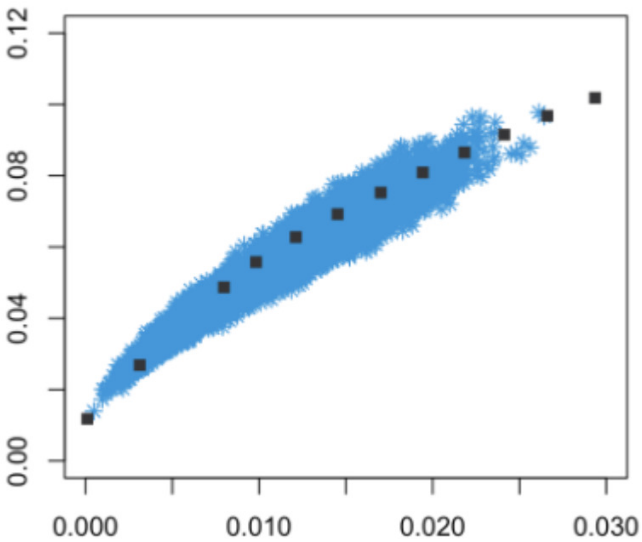

Here we see that the portfolios have also performed well, though not quite as close to the efficient frontier. This is unsurprising as the variable we were primarily optimizing was CVaR, not downside volatility.

Rather than select these 12 portfolios based on the first 20 years of data, we incorporated the last 5 years of information. To solve the out-of-sample issue, we opted to rerun all the optimizations on the full timeframe (all 25 years), and verify that the weights don't change significantly from the original weights. We found only minor changes in the data, so we selected the results from the 25 years. We are aware that there is risk of overfitting, but acknowledge that changes were fairly minor.

This results in our final portfolios:

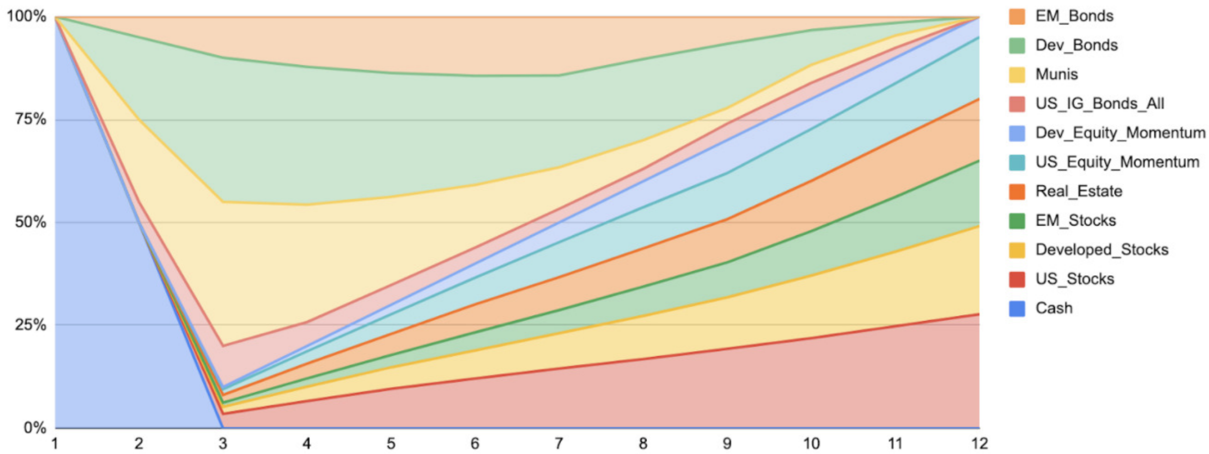

Who uses what portfolio

To determine which of the 12 portfolios a client receives, we use a risk assessment. The risk assessment allows us to determine if a client is aggressive, moderately aggressive, moderate, moderately conservative, or conservative. We combine this with our client's need to spend their money using 5 time horizons (Within 1 year, 1-3 years, 3-5 years, 5+ years, or not sure). Logically, for a more aggressive investor with a longer time horizon, we will recommend a more aggressive portfolio.

We quantify this recommendation. To do this, we looked at basic cash/bond/stock portfolios from December 1986 - November 2020, and analyzed every portfolio split over every time horizon (shifted by one month to include all the data and have the greatest number of data points at each time horizon).This analysis resulted in the following rules of thumb:

1. **An aggressive investor is one that is reaching for the highest return, but could suffer a 5-10% loss 5% of the time.**
2. **A moderately aggressive investor is one that is also reaching for high returns, but is willing to suffer less losses, specifically a 1-5% loss 5% of the time**
3. **A moderately aggressive investor is one that is also reaching for high returns, but is willing to suffer less losses, specifically a 1-5% loss 5% of the time**
4. **A moderate investor is one who is breaking even (0% return) at the 5th percentile**
5. **A conservative investor is one who is breaking even at the 1/10th percentile of returns.**

These rules of thumb led us to these recommendations:

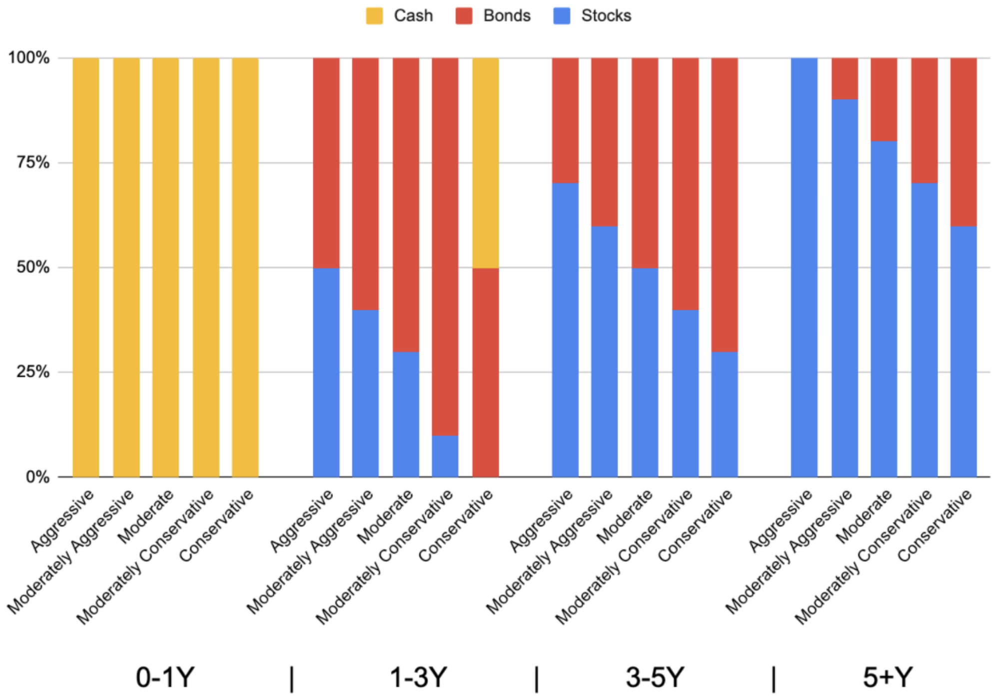

Trading

We use a rebalancing algorithm that tests to see how far off track client portfolios are from their target model. We have a buffer of 5% on each asset class, so if any are 5% away from their target, a trade is triggered. We trade partial ETFs, so even on small client accounts we can match the target model reasonably well. We trade daily near the market open, and we target a small 1% cash buffer in accounts in order to reduce cash drag and ensure orders are fulfilled. The rebalancer takes tax estimates as inputs when deciding what to buy or sell, and attempts to make transactions as efficiently as possible when a client adds or removes cash. We allow wash sales in order to ensure all client transactions occur when requested.

Conclusion

At Origin, we are striving to best serve our clients by providing optimized portfolios for investors of every risk level and every time horizon. We began with a base of MPT and added PMPT to give our clients the most relevant analysis. We hope this article helps our clients understand why we've made the decisions we did. Too often, investment firms use the phrase 'Nobel-prize winning research,' and we want to ensure that our investors are armed with the current state of investment research and do not fall prey to subpar investment strategies. We believe that using the process we outlined above, we have been transparent and aim to serve our clients well.

Appendix - List of Asset Classes

Total US Stock Market - Broad US equity market exposure, including securities traded on the NYSE, NYSE American, NYSE ARCA, NASDAQ, Bats Global Markets, and the Investors Exchange. Nearly 4,000 US-based constituents across mega, large, small and micro capitalizations, representing nearly 100 percent of the U.S. investable equity market. This is a staple of most portfolios, so we included it for analysis.

Developed Stock Market - Approximately 3,700 common stocks of large-, mid-, and small-cap companies located in Canada and the major markets of Europe and the Pacific region (Japan, UK, Canada, France, Germany, Switzerland, Australia, etc.). This is another staple of many portfolios.

E*merging Market Stocks* - Approximately 5,000 stocks of companies located in emerging markets around the world, such as China, Taiwan, India, Brazil, and South Africa. Though this represents a higher-risk and more inefficient asset class, higher expected returns should accompany it. EM stocks are another standard part of many portfolios (appropriately sized), so we included it for analysis.

Real Estate - Approximately 175 stocks of large, mid-size, and small U.S. companies within the real estate sector, typically real estate investment trusts (known as REITs), which includes specialized REITs, and real estate management and development companies. These REITs purchase office buildings, hotels, and other real property. Real estate is an absolutely massive asset class and very unique so we included it in our analysis.

US Equity Momentum - Seeks to track the performance of U.S. large- and mid-capitalization stocks (approximately 125) exhibiting relatively higher momentum characteristics. Screens for 6 and 12 month momentum. As stated above, the momentum factor has been one of the most significant sources of excess returns, so we included it.

Developed Equity Momentum - Exposure to large- and mid-cap developed international stocks (approximately 300) exhibiting relatively higher price momentum. Also screens for 6 and 12 month momentum. Similar to the above, momentum exists across nearly every asset class, and this asset class (developed equity momentum) exists in a reasonably sized ETF (IMTM), so we included it.

Gold - Physically backed investment in Gold, not futures of the commodity. Technically fractional shares in a trust whose sole assets are gold bullion. Gold has some really interesting risk-off properties (flee to 'safety'), so we wanted to see if it would add anything to our portfolios.

US Investment-Grade Bonds - BND provides broad exposure to the taxable investment-grade U.S. dollar-denominated bond market, excluding inflation-protected and tax-exempt bonds. This includes a wide spectrum of public, investment-grade, taxable, fixed income securities in the United States—including government, corporate, and international dollar-denominated bonds, as well as mortgage-backed and asset-backed securities, all with maturities of more than 1 year. This is a standard bond staple in many portfolios.

Developed Market Bonds - A broad-based measure of the global, investment-grade, fixed-rate debt markets. The index includes government, government agency, corporate, and securitized non-U.S. investment-grade fixed income investments, all issued in currencies other than the U.S. dollar and with maturities of more than one year. To minimize the currency risk associated with investment in bonds denominated in currencies other than the U.S. dollar, the fund (BNDX) will attempt to hedge its currency exposures. Top country exposures are Japan, France, Germany, Italy, the UK, etc. This is another standard bond staple (compared to the US version BND), so we included it for analysis.

Emerging Market Bonds - Includes U.S. dollar-denominated bonds that have maturities longer than one year and that were issued by emerging market governments and government agencies, as well as government-owned corporations. Primary country exposures include Mexico, Saudi Arabia, Indonesia, Turkey, Qatar, etc. This again offers a diversified bond exposure, so we included it for analysis.

Treasury Inflation Protected Securities (TIPS) - Includes all publicly-issued U.S. Treasury Inflation-Protected Securities (TIPS) that have at least one year remaining to maturity, are rated investment grade (BBB- or higher) and have $500 million or more of outstanding face value. This is an interesting asset class, the principal adjusts with inflation (and then the interest is based on that adjusted principal) to help prevent investors' purchasing power from decaying. Because of its uniqueness, we included this asset class.

Long-Term Treasuries - Invests in US Treasury Bonds with a dollar-weighted average maturity of 10-25 years. Treasuries are effectively loans to the US Government, and are backed by the full faith and credit of the US Government. Treasuries form an important foundation to the bond market (and global economy), so we included it for analysis (not just long-term, see below, intermediate, and short term too).

Intermediate-Term Treasuries - Invests in US Treasury Bonds with a dollar-weighted average maturity of 5-10 years. This again forms an important foundation for the bond market (many rates are linked to the US 10Y rate), so we included it in our analysis.

Short-Term Treasuries - Same as the above, but a 1-3 year average maturity. We included this for completeness so we could analyze the full spectrum of Treasury securities.

Municipal Bonds - Investment grade segment of the US Municipal Bond market. Includes municipal bonds from issuers that are primarily state or local governments or agencies whose interest is exempt from U.S. federal income taxes and the federal alternative minimum tax (AMT). A big reason to include this is potentially higher yield than Treasuries (due to more risk of a municipality vs the US Government) and no income taxes at the federal level. Sample holdings include: 1) San Diego CA Unified School District GO, 2) Grand Parkway Transportation Corp. Texas System Toll Revenue, and 3) Metropolitan Transportation Authority NY Revenue.

US Investment-Grade Corporate Bonds - These represent a further increase in potential risk and return from munis and treasuries because corporations cannot force customers to buy their products (cannot raise taxes like governments). This is a broad representation of the U.S. dollar-denominated liquid investment grade corporate bond market. Top issuers include JP Morgan, Comcast, Verizon, and Apple.

US Below Investment-Grade Corporate Bonds - Lovingly known as 'Junk Bonds,' (rated below BBB-). HYG invests in a broad representation of the U.S. dollar denominated liquid high yield corporate bond market. This represents a further increase in potential risk and return from investment-grade bonds because these have lower ratings. This includes debt from Ford, Sprint, Occidental Petroleum, Tenet Healthcare, etc. This is the only below investment-grade asset class, so we wanted to include it for analysis.

Mortgage-Backed Securities - VMBS invests primarily in U.S. agency mortgage-backed pass-through securities issued by Ginnie Mae (GNMA), Fannie Mae (FNMA), and Freddie Mac (FHLMC) with a dollar-weighted average maturity of 3 to 10 years. Asset-backed securities are pretty interesting, so we wanted to include them in our analysis.

Cash - We wanted to include an ultra short-term treasury option (SHV) for conservative portfolios as a proxy for cash (money markets, high yield savings, etc.). The weighted average maturity is about 0.38 years.

Bibliography

Black, F. & Litterman, R. (1992). Global Portfolio Optimization. Financial Analysts Journal.

Bogle, J. (2009). Common Sense on Mutual Funds. Wiley.

Brinson, G. P., Hood, L. R., & Beebower, G. L. (1986). Determinants of Portfolio Performance. Financial Analyst Journal, 39-44.

Brinson, G. P., Singer, B. D., & Beebower, G. L. (1991). Determinants of Portfolio Performance II: An Update. Financial Analyst Journal, 40-48.

Carhart, M. M., (1997). On Persistence in Mutual Fund Performance. Journal of Finance.

Chordia, T., Goyal, A., & Sarett, A., (2020). Anomalies and False Rejections. The Review of Financial Studies.

DeMiguel, V., Garlappi L., & Uppal, R. (2009). Optimal Versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy? The Review of Financial Studies.

Fama, E. F. & French, K. R. (1992). The Cross-Section of Expected Stock Returns. Journal of Finance.

Fama, E. F. & French, K. R. (2010). Luck versus Skill in the Cross-Section of Mutual Fund Returns. Journal of Finance.

Harvey, C. R., Liu, Y., & Zhu, H., (2016). ...and the Cross-Section of Expected Returns. The Review of Financial Studies.

Ibbotson, R.G. & Kaplan, P.D. (2000). Does Asset Allocation Policy Explain 40, 90, or 100 Percent of Performance? Financial Analyst Journal, 26-33.

Jegadeesh, N. & Titman, S., (1993). Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency. Journal of Finance.

Kan, R. & Smith, D., (2008). The Distribution of the Sample Minimum-Variance Frontier. Management Science.

Ledoit, O. & Wolf, M., (2004). Honey, I Shrunk the Sample Covariance Matrix, Journal of Portfolio Management.

Markowitz, H. (1952). Portfolio Selection. Journal of Finance.

Michaud, R. (1998). Efficient Asset Management. Oxford University Press.

Rom, B. M. & Ferguson, K., (1993). Post-Modern Portfolio Theory Comes of Age. Journal of Investing.

Sharpe, W. (1964). Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risks. Journal of Finance.

origin

Financial Planning for Every Employee, Everywhere.