The Real World is Not Normal

Introducing the new frontier: an alternative to the mean-variance optimizer.

by
Cindy Sin-Yi Tsai, CFA, CAIA
Senior Research Analyst

Asset allocation is the process of dividing investments among different kinds of asset categories based on an investor's specific investment objective, risk tolerance, and other constraints. It is one of the most important decisions an investor makes, no matter whether one believes in the conventional wisdom (from Brinson, Hood, and Beebower 1986[1]) or the newer research (such as Xiong et al. 2010[2]). Asset allocation is commonly determined using a software tool that optimizes risk and return trade-offs, and the Markowitz mean-variance optimization has been the standard for creating efficient asset-allocation strategies for more than half a century.

But MVO is not without its shortcomings. The MVO process requires forming asset-class assumptions (namely expected return, standard deviation, and correlation coefficients), which ultimately result in an efficient frontier of the best combinations of those asset classes to achieve the highest portfolio return for each level of risk. Two limitations of the MVO are associated with making asset-class assumptions (normal distribution and linear correlation assumptions) and two more with the optimization methodology itself. Now there are solutions to overcome some of these limitations.

The Limitations of Assuming a Normal Distribution

In an MVO, we use the normal distribution when forming asset-class assumptions. What is nice about the normal distribution is that it is very intuitive: Roughly two thirds of the time, returns are within one standard deviation away from the mean (average) return; more than 95% of the time, returns are within two standard deviations; and returns are within three standard deviations of the mean about 99.7% of the time. This means, according to normal distribution mathematics, there is approximately a 0.13% probability of an extremely large gain or loss (100% less 99.74% divided by 2).

The normal distribution is flawed, however, in that it is a bell-shaped curve that assumes symmetry (a loss is just as probable as a gain) and thin tails (trivial probabilities assigned to three-sigma events, those greater than three standard deviations away from the mean). Because investors are more averse to negative surprises resulting from underestimating extreme losses, as opposed to positive surprises of unexpected large gains, we focus on the normal distribution's ability to model three-sigma losses. When we examine the actual historical monthly data of the S&P 500 Index going back to 1926, we observe that three-sigma losses happened in 10 of the past 1,026 months (over 85 years). This is almost a 1% frequency, which is almost 8 times what a normal distribution predicts. This means that a normal distribution fails to model the "tail risk" in the real world.

As indicated in Xiong and Idzorek (2011)[3], many asset classes empirically exhibit return distributions that are skewed to the left of the mean (negative skewness) and that have fatter tails (excess kurtosis) than a normal distribution. The authors demonstrate that accounting for skewness and excess kurtosis in return modeling and optimization makes a significant impact on the asset-allocation decision, especially in terms of performance during a crisis, such as the one that occurred in 2008.

Another limitation of the traditional MVO is that it assumes correlation coefficients among asset-class returns are linear—in other words, the same correlation coefficient applies in both up and down markets. This is unrealistic, as it it is commonly observed that during crisis, markets tend to go down together.

CONTINUED ON NEXT PAGE

1 Brinson, Gary P., L. Randolph Hood, and Gilbert L. Beebower. 1986. "Determinants of Portfolio Performance." *Financial Analysts Journal*, (July-August):39–44.
2 Xiong, James X., Roger G. Ibbotson, Thomas M. Idzorek, and Peng Chen. 2010. "The Equal Importance of Asset Allocation and Active Management." *Financial Analysts Journal*, (March-April):22–30.
3 Xiong, James X. and Thomas M. Idzorek. 2011. "The Impact of Skewness and Fat Tails on the Asset Allocation Decision." *Financial Analysts Journal*, (March-April):23–35.

Modeling Asset Classes in Practice

To form asset-class assumptions, we selected index proxies[4] to represent 12 asset classes. These include traditional investments such as equities (U.S. large capitalization, U.S. small capitalization, international developed, international emerging), debt (U.S. investment-grade, U.S. high-yield, and international), and cash. We also incorporated alternative investments such as U.S. real estate, international real estate, commodities, and hedge fund arbitrage. We added arbitrage for the potential diversification benefits of its "alternative beta."

Table 1 shows key characteristics of historical return distributions in the common time period among these asset classes. Most asset classes have negative skewness and excess kurtosis, but U.S. high-yield bonds, U.S. real estate, and hedge fund arbitrage have much larger figures than others. Xiong and Idzorek (2011) found that variety in skewness and kurtosis among assets makes a significant difference in allocation when an optimizer penalizes downside risk instead of standard deviation. To demonstrate, we generated two sets of asset-class return assumptions, one using normal and one using fat-tailed and skewed distribution models.

We modeled asset-class return assumptions using the log-normal distribution, the natural logarithmic version of the normal distribution that reflects the (unleveraged) real-world experience where investors cannot lose more than 100% of their investment but can make more than 100% on the upside. There are several methods to derive log-normal return assumptions. We selected the building blocks method, outlined in the *Morningstar Ibbotson Stocks, Bonds, Bills, and Inflation*℠ yearbook. In real life, though, the building blocks method serves only as a starting point. Investors should incorporate their own forecasts into return assumptions. To model standard deviations and correlation coefficients, we used historical data covering the common period of the asset-class index proxies (February 1994 to June 2011) for simplicity, even though long-term historical data is preferable.

Histogram graphics allow users to see how their distribution model choice fits historical returns, which can in turn help users further fine-tune assumptions. A histogram is a bar graph in which returns are sorted into bins, and the height of the bin illustrates how often that particular range of returns occurs. Figure 1 shows that the standard log-normal distribution fails to model historical returns of two asset classes: U.S. large-capitalization stocks and U.S. real estate.

In both instances the log-normal distribution curves do not have fat-enough tails or negative-enough tilt to cover the largest losses, represented by the three left-most bars. In other words, this tail risk is completely ignored.

CONTINUED ON NEXT PAGE

Table 1: Historical Return Distribution Characteristics—February 1994 to June 2011

Asset Class	Return %	Std Dev	Skewness	Excess Kurtosis
U.S. Large Cap	7.83	15.56	–0.72	1.03
U.S. Small Cap	7.89	20.00	–0.56	1.08
International Developed	5.19	16.63	–0.68	1.70
International Emerging	6.79	24.31	–0.76	1.98
U.S. Inv Grade	6.18	3.79	–0.26	0.96
U.S. High Yield	7.49	9.32	–1.16	9.33
International Bond	6.49	8.59	0.17	0.58
U.S. Real Estate	10.58	20.54	–0.87	8.67
International Real Estate	6.89	20.12	–0.50	2.61
Commodity	6.99	15.68	–0.53	2.29
HF Arbitrage	8.17	3.54	–3.50	22.59
Cash	3.34	0.58	–0.35	–1.43

Figure 1: Curves of Log-Normal Distributions and Histograms of Historical Returns

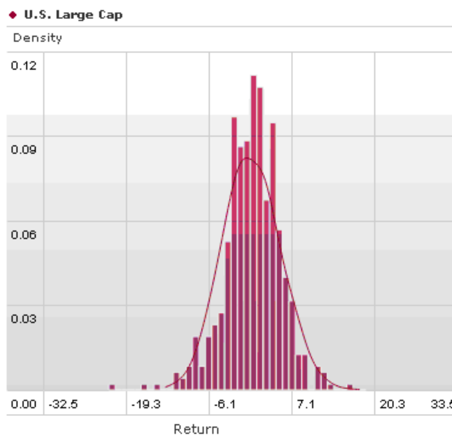

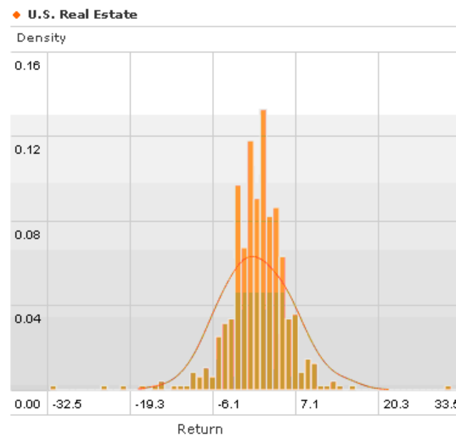

4 IA SBBI S&P 500, Russell 2000, MSCI EAFE, MSCI EM, Barclays US Agg Bond, Barclays US Corp High Yield, Citi WGBI NonUSD, FTSE NAREIT All Equity REITs, FTSE EPRA/NAREIT Dev Ex US, DJ UBS Commodity, Morningstar MSCI Relative Value, Citi Treasury Bill 3 Mon.

Figure 2: Curves of Johnson Distributions and Histograms of Historical Returns

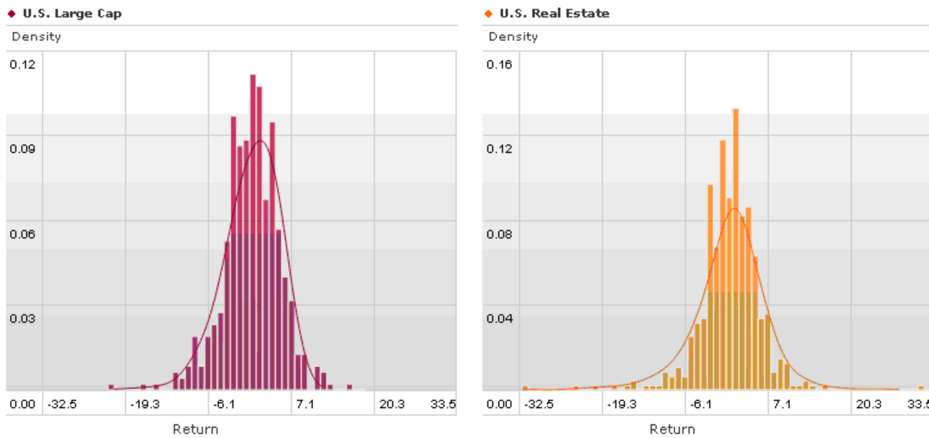

This is not surprising for U.S real estate, given its historical skewness and excess kurtosis. (See Table 1). But the log-normal model is just as poor in representing a traditional asset class such as U.S. large- capitalization stocks, which anchor the portfolios of most U.S. investors.

To model the second set of assumptions for the same asset classes using a fat-tailed distribution, we chose the Johnson distribution. The reason is twofold: first, to offer a different viewpoint than Xiong and Idzorek (2011), who use the Truncated Lévy-Flight model, and second, the Johnson distribution is more intuitive than the TLF model. (The Johnson model's primary limitation, however, is that it is less useful for modeling daily or weekly returns). This is because, in order to model tail risk, the Johnson method requires only two additional inputs—skewness and kurtosis—beyond the traditional expected return, standard deviation, and correlation coefficient inputs required for MVO. Skewness and kurtosis are easily understood when illustrated visually. When making skewness and kurtosis assumptions, one can start with historical skewness and excess kurtosis (the kurtosis above and beyond a normal distribution's kurtosis, which is 3) as a baseline for further refinement. Modeling each asset class' tail risk individually is preferable, as equities and alternative assets have more tail risk than plain-vanilla fixed income (at least historically).

Figure 2 shows how much better the Johnson distribution models historical U.S. large-capitalization equity and U.S. real estate data relative to the log-normal distribution in Figure 1, when using historical skewness and excess kurtosis from February 1994 to June 2011 as parameters. The bars on the left side of histograms, those that represent the largest losses, are better covered with the Johnson distribution curve. Moreover, both the placement of and the height of the curve's peak fall better in line with the tallest bar. This little extra effort makes a compelling argument dramatic modeling improvement with very to incorporate tail risk into the asset-allocation process. Therefore, we ran two optimizations, one with assumptions generated with the log-normal distribution and another with the assumptions based on the Johnson distribution.

Optimization

Besides its faulty assumption process, traditional MVO's optimization process also poses problems. One is that it uses arithmetic mean for expected return. An alternative is geometric mean, which is the time-weighted rate of return over multiple periods. Optimizing on arithmetic mean assumes a single-period investment horizon and maximizes a portfolio's return over this period, based on the premise that one revisits asset allocation at every period. Multiperiod optimization, which has the objective of maximizing long-term wealth, requires the use of geometric mean.

A second limitation of the MVO process is that it uses standard deviation as the measure of risk. Standard deviation measures total risk on both the upside and downside, while many investors are more concerned with downside risk. There are several measures of downside risk, but one that is particularly good at capturing tail risk is the conditional-value-at-risk, or CVaR for short. The easiest way to understand CVaR is to understand its cousin VaR and with an example. When an asset's fifth percentile VaR is 30%, there is a 5% chance of losing at least 30% of its value. The CVaR, on the other hand, is the probability-weighted average loss of all possible losses equal to or exceeding 30%. The CVaR, essentially, captures a distribution's entire left tail after the 30% loss.

Xiong and Idzorek (2011) demonstrate that there is no need to optimize using CVaR if one models asset-class assumptions using a normal distribution, because the allocations will be the same as that of a conventional MVO. Doing so just adds extra complexity. If one believes that certain asset classes exhibit negative skewness and fat tails, however, and if one incorporates these beliefs into asset-class assumptions (using the Johnson distribution, for example), optimizing with a downside risk measure such as the CVaR makes an impact on asset allocation. Therefore, in order to demonstrate the impact of tail-risk modeling, we paired up log-normal assumptions with the conventional MVO and, separately, the Johnson assumptions with a M-CVaR (mean-CVaR) optimizer.

Because optimization is part art and part science, constraints help to ensure the optimizer produces intuitive results. We set three types of constraints for the purposes of this study. The first is a maximum allocation to each individual asset class. For example, for this study, we did not want allocations to

CONTINUED ON NEXT PAGE

international emerging stocks and U.S. high-yield bonds to exceed 30% individually, as these asset classes are particularly risky. We also didn't want international bond, U.S. real estate, international real estate, commodities, and hedge fund arbitrage to exceed 20% each. Next, we wanted to limit the combinations of allocations to alternative investments to 25%. Finally, we didn't want the riskier asset classes' allocations to exceed those of the less-risky assets, so we constrained the weighting of U.S. small-cap stocks and international developed bonds to be less than that of U.S. large-capitalization stocks. Similarly, we limited U.S. high-yield bonds or international bonds to the weightings of U.S. investment- grade bonds, and the weighting of international emerging-markets stocks to 40% of the amount allocated to international developed stocks. These constraints apply to both MVO and M-CVaR optimizer.

When running an optimization, an investor can identify optimal portfolios based on the investor's expected return objective or risk tolerance. For example, Xiong and Idzorek (2011) took this approach, comparing a mean-variance optimized portfolio to a mean-CVaR optimized portfolio of the same mean, or expected return. For this article, we followed a similar process, but we also specified a particular broad asset class mix, of 45% equity, 30% fixed income, and 25% alternative investments. This approach allows us to more easily identify which subasset classes are favored in M-CVaR optimization within each broad asset class.

Figure 3: Comparison of Allocation Spectrums Between MVO and Fat-Tail (Johnson) Optimization

MVO

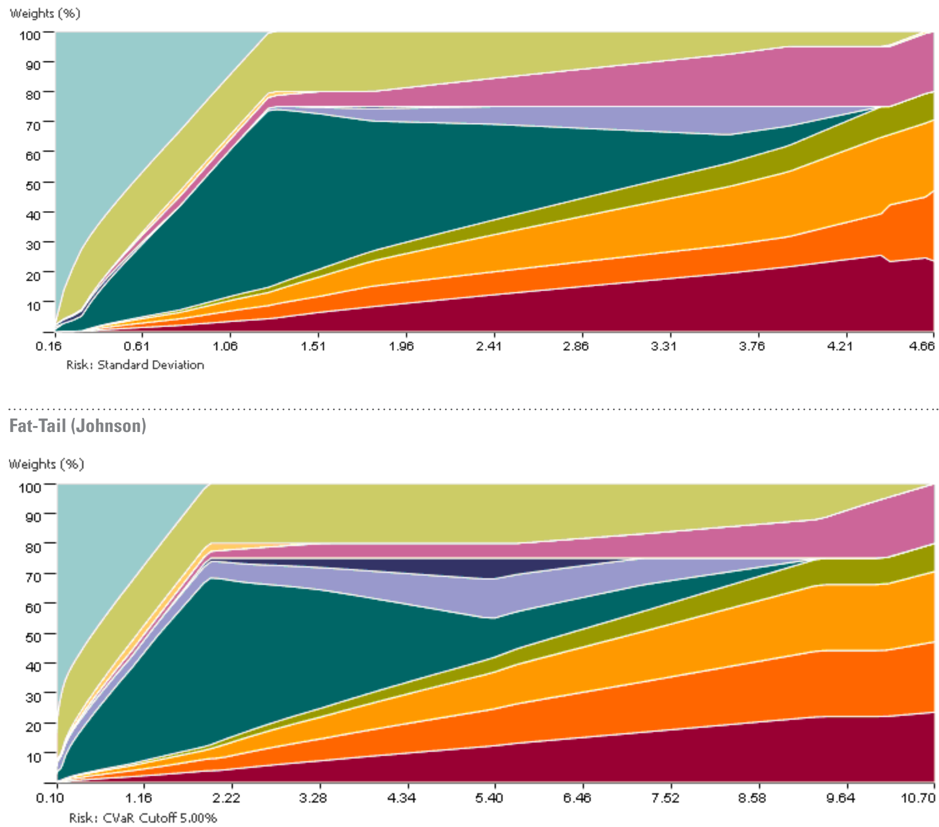

The Results

The allocation area graphs in Figure 3 display the allocation results of our MVO and Johnson M-CVaR optimizations across the entire risk/return spectrum, from lowest risk on the left to highest risk on the right. The MVO allocations were generated with the normal distribution assumptions, and the Johnson M-CVaR optimization used the Johnson (fat-tailed) distribution assumptions. We found the 45/30/25 mix in each graph, the details of which are displayed in Table 2 (next page).

All else being equal, investors ought to favor asset classes with positive skewness and small (or even negative) excess kurtosis, and this bias should manifest itself in the difference between the Johnson M-CVaR optimizer results and the MVO results in Table 2. The Johnson M-CVaR optimizer ought to recommend less allocation to those asset classes with large negative skewness and excess kurtosis, characteristics that are ignored in the MVO. Per the historical skewness and kurtosis statistics in Table 1, we would expect international bonds to be favored by the Johnson M-CVaR optimizer, while U.S. high-yield bonds, U.S. real estate, and hedge fund arbitrage should be relatively unattractive.

The results are generally consistent with what we intuitively expect. Looking at the last column of Table 2, we see that, within the four equity subasset classes (the first four rows), the difference in allocations between the traditional MVO and Johnson M-CVaR optimization is generally unremarkable, although small-capitalization stocks are slightly favored by the Johnson M-CVaR optimizer for having a smaller negative skewness. In fixed income, however, we see a significant difference in allocation. Intuitively, international bonds were

CONTINUED ON NEXT PAGE

Table 2: Allocations Comparison Between Two 45% Equity/30% Fixed Income/25% Alternatives Portfolios Generated With MVO and Johnson (Fat-Tail) Optimization

Asset Class	MVO %	Johnson (Fat Tail) %	Difference %
U.S. Large Cap	15.3	14.1	–1.2
U.S. Small Cap	8.3	11.5	+3.2
International Developed	15.3	14.1	–1.2
International Emerging	6.1	5.6	–0.5
U.S. Inv Grade	22.7	9.9	–12.8
U.S. High Yield	7.3	9.9	+2.6
International Bond	0.0	9.9	+9.9
U.S. Real Estate	12.7	16.0	+3.3
International Real Estate	0.0	0.0	0.0
Commodity	0.0	1.7	+1.7
HF Arbitrage	12.3	7.3	–5.0
Cash	0.0	0.0	0.0
Total	**100.0**	**100.0**	**0.0**

Figure 4: Comparison of Efficient Frontiers With Log-Normal to Fat-Tailed (Johnson) Distribution

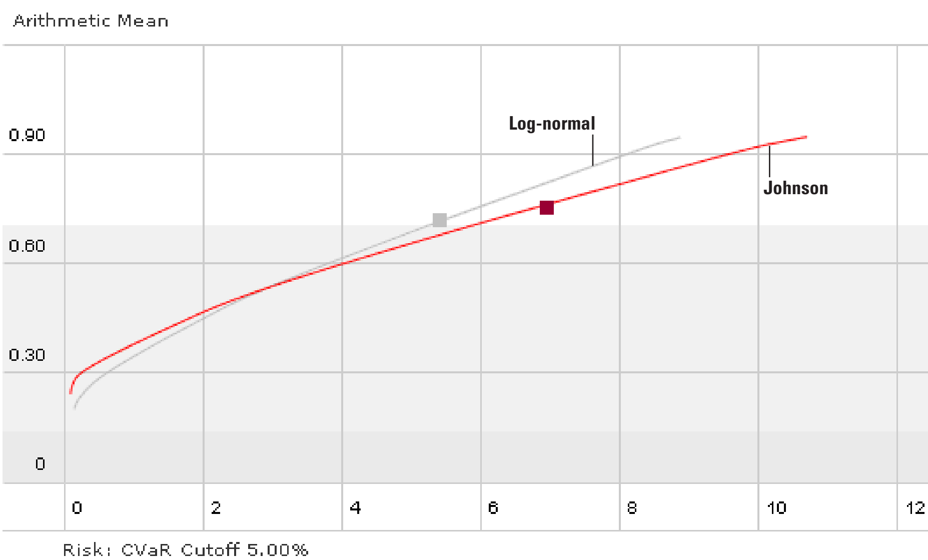

ignored in the MVO portfolio, but the M-CVaR optimizer calls for the maximum 9.9% allocation because of the asset class' positive skewness and low excess kurtosis. (The Johnson M-CVaR optimizer also increased the allocation to U.S. high-yield bonds, which may appear counterintuitive because of their skewness and kurtosis characteristics. The lower correlation benefit of U.S. high-yield bonds to international bonds trumps these characteristics, however). Another area of significant impact is the alternative investments bucket, where the allocation to hedge fund arbitrage is significantly reduced, as we would expect because of this asset's large negative skewness and outsized excess kurtosis. The commodities bucket, however, gets a greater, albeit small, allocation in the Johnson M-CVaR process and is ignored in MVO. U.S. real estate received a larger allocation as well. Overall, we find the results to be consistent with the conclusion in Xiong and Idzorek (2011)—that taking skewness and kurtosis into consideration makes a significant impact in asset allocation.

Figure 4 shows the two efficient frontiers related to the MVO and M-CVaR allocation area graphs in Figure 3. The two dots on Figure 4 represent the two 45/30/25 portfolios discussed in the previous paragraph. We see that, when incorporating non-normal assumptions (of skewness and kurtosis) into the Johnson M-CVaR optimization, our efficient frontier falls to the southeast of the MVO efficient frontier for most of the risk spectrum. This means that our MVO optimization underestimates risk and that the Johnson M-CVaR efficient frontier is more likely to model reality.

To Optimize, or Not to Optimize

Because our allocation experiment produced relatively intuitive results, one might think that it is unnecessary to run an optimization. One might simply obtain the historical skewness and excess kurtosis figures for each asset class and manually reduce the allocations to the unattractive asset classes. Whether or not one chooses to employ an optimizer, the argument for incorporating tail risk into the asset-allocation decision process is clear. Optimization is but one tool to aid in that process. M

This is an abridged and slightly modified version of the article published in October 2011. Morningstar's Alternative Investments Observer is quarterly electronic newsletter for institutional investors and advisors. For a free copy, please email nadia.papagiannis@morningstar.com.