# Asset allocation discussion with Mark Kritzman, MIT Sloan School of Management

Below is an edited transcript of a discussion with Mark Kritzman, senior lecturer in finance at the MIT Sloan School of Management, and a leading researcher on asset allocation. He is also an expert on risk management and has developed many approaches for evaluating and addressing risk. Our conversation with Mark was an opportunity to draw from his insights on asset allocation for the short and long term with a focus on the unique challenges faced by sovereign wealth funds. He also addressed questions relevant to sovereign wealth funds about risk management.

**Q:** *The members of the IFSWF working group tasked with researching asset allocation for the short- and long-term has identified five key asset allocation challenges faced by SWFs:*

1. *Balancing the tension between long-term and short-term investment objectives*
2. *Dealing with uncertainty when constructing portfolios*
3. *Developing frameworks to incorporate alternative asset classes into the portfolio*
4. *Communicating with stakeholders*
5. *Optimising the organisational structure*

*Do you think we are missing anything here? Is there anything you would like to add?*

**MK:** Those are all good topics. Another topic that is of interest is the issue of estimation error. Whenever we build portfolios we need to estimate returns and risk. When we do that, we are, unfortunately, exposed to various sources of error.

**Q:** *That is an excellent point and a topic we should cover. If you don't mind, perhaps we can begin with an initial step of the asset allocation process - defining and selecting asset classes. You have done some work in this area. Could you share any insights on what an asset class is, how you define an asset class, and how you might decide whether to incorporate an asset class in a portfolio or not?*

**MK:** That is a great question and I don't think many people have formally addressed it. There are three key criteria for distinguishing an asset class. The first is that is it should be something that raises the efficiency of the portfolio. To be more technical, it would raise the portfolio's expected utility. This means that it either raises the portfolio's expected return or it reduces the portfolio's risk. It should accomplish this without requiring investors to have skill in identifying superior managers. In other words, a passive exposure to the asset class should increase the expected utility of the portfolio.

The second criterion is that the components within an asset class should be homogeneous. They should be similar. The reason for that is that if you combine components that are not very similar within an asset class, then you are imposing an unnecessary constraint on the asset allocation process. You are saying that I must hold these two components in the fixed weights that they appear within the asset class. If the components are very different from each other you should split that asset class into two separate asset classes and that will enable you to achieve a more efficient outcome.

The third criterion for qualifying as an asset class is that it should be sufficiently large to absorb a meaningful fraction of one's portfolio. If you were to invest in an asset class that did not have adequate capacity, you would drive up the cost of investment and reduce the portfolio's liquidity. That wouldn't be a very good outcome.

There are a couple of categories of assets that are a bit tricky. For example, many investors consider hedge funds to be an asset class. I don't believe hedge funds are an asset class. Hedge funds invest in all different kinds of asset classes. So what we are really investing in is perceived manager skill. It is unlikely that you could invest in the average hedge fund without any ability to distinguish between a good fund and a bad fund and raise the portfolio's expected utility.

Another possible asset class is private equity. Our research shows that the average private equity fund, measured on a risk-equivalent basis, has produced a premium relative to public. Therefore, one should expect the average private equity fund to raise a portfolio's expected utility, without the benefit of selection skill. My inclination is to say that private equity is an asset class and hedge funds are not.

**Q:** *One question does come to mind regarding SWFs investing in private equity. Given the size of some SWFs, could you provide any insights regarding how they might approach investing in private equity? It is likely that there isn't sufficient capacity in any particular fund to represent a meaningful allocation within a SWF's portfolio. How might they go about developing a private equity allocation?*

**MK:** Obviously, what you want to do is look at the private equity universe; perhaps you might want to sort it by venture capital, buyout funds, or other sub-categories of private equity. Then you would try to identify those funds that you think are going to generate the best performance and figure out a plan for getting exposure to those funds. To get started or as an alternative to private equity you can invest in liquid private equity. It has been shown that approximately three-quarters of the premium of private equity over public equity, on a risk equivalent basis, can be explained by the sector exposures of private equity funds. What that means is that you can invest in public equity sector ETFs or index funds and expect to receive about 75 per cent of the premium of private equity over public equity; at least that has been the case historically.

The other 25 per cent of the premium of private equity over public equity is attributable to illiquidity. The fact that there are lock-ups and fewer disclosure requirements enables private equity managers to do things, or restructure companies, in ways that publicly traded companies cannot. I would think that liquid private equity would be a very good substitute for private equity while you are waiting for your ultimate investment in private equity funds. If you have private equity and have committed capital that has not yet been called, liquid private equity is also a pretty good repository for that committed capital because at least it is delivering a very similar risk profile as you would expect to get from private equity.

**Q:** *Like your comments on private equity and hedge funds we can advance into the next topic on producing estimates of returns, risks, and correlations for asset allocation. Those investments, as well as real estate and infrastructure, do present some issues with producing estimates. Can you touch a bit on each of these assets classes and provide some insights as to how you might deal with some of the issues that each of these investments presents?*

**MK:** If I were to conduct an asset allocation analysis with both publicly traded and less liquid asset classes, I would estimate expected returns, as a starting point, to be equilibrium returns. Those are the returns that you would expect to earn if all asset classes were fairly priced. In the case of publicly traded assets, the equilibrium returns are those returns that are proportional to their betas. That implies that if a particular asset class is mispriced, investors can trade that asset class and correct the mispricing so that the expected returns are proportional to beta. In the case of illiquid asset classes, if you perceive an illiquid asset class to be mispriced you can't simply just trade and expect that mispricing to be corrected. This is because illiquid asset classes are very expensive to trade. In that case, I would think that the equilibrium returns would be more proportional to the variance of the illiquid asset classes. That is just how I would get started. Then you might have views that you want to incorporate. You may think one asset class should have a return higher than its equilibrium for one particular reason or vice versa.

In any event, dealing with illiquid asset classes is tricky for a variety of reasons. One is that, in many cases, the managers are paid performance fees. This has two effects. One is that the measured or the observed volatility of the returns net of fees is lower than the returns gross of fees. So the volatility that you observe actually understates risk. The reason for that is that performance fees cut off the upside. Reducing upside volatility, which is what performance fees do, is not lowering risk. When a manager outperforms and you give some of that outperformance back to managers you are reducing the upside you get…you are not lowering risk. It lowers volatility but it doesn't lower risk. The first thing you need to do is reverse engineer the fee calculation so that you get a proper measure of downside volatility.

The other problem with some of these asset classes, such as private equity and real estate and, in some cases, hedge funds, is that the values are based on fair value pricing. These prices are typically anchored to prior period prices so they are smoothed, there is positive autocorrelation. That also understates the true risk of these investments. So what you ought to do is de-smooth the returns. If you do that you get estimates of risk that make much more sense.

On the return side, performance fees also cause you problems. For example, it turns out that if you have many managers who charge performance fees, the expected returns of those managers as a group will be less than the average of their individual expected returns. The reason for this is that when a manager outperforms, they collect a performance fee. When a manager underperforms, they do not reimburse you for that underperformance. So the actual average return of the managers is lower than the average of the individual expected returns of the managers.

Now, you might argue that there are clawbacks that would prevent that from happening. That is true in principle but, in fact, that is hardly ever the case. It is typically the case that the manager either gets terminated, if the manager underperforms significantly over some period or, if you really like the manager, you are going to reset the high-water mark. I would say a good rule of thumb is that the expected return of a group of managers who charge performance fees is about 80 basis points less than the average of their expected returns.

When you conduct your asset allocation analysis and you have corrected these issues you'll have lower expected returns and higher risks. That is going to cause your optimal allocation to these types of assets to be lower than if you had not taken these issues into account.

**Q:** *When producing estimates for both alternative and traditional assets there are a couple of other things we might want to consider. We know that markets exhibit regime type behaviour, so it could be important to incorporate this regime information. Also, institutional investors are often tasked with managing to long-term objectives while also being evaluated over shorter intervals. There is a tension between long-term and short-term objectives that they have to manage. Could you provide some insights as to how we might approach considering risk regimes as well as understanding and addressing risk across different investment horizons when constructing portfolios?*

**MK:** The implicit assumption in the way portfolio theory is usually described in the text books is that returns are generated from as single regime, so there is a single distribution that you have to pay attention to. It turns out that, empirically, that has not been the case. One way of categorizing history is to try to categorise it in terms of fragile or turbulent periods versus resilient or calm periods. The way to distinguish these periods is not best done with volatility and correlation. Those are the traditional ways of measuring instability of returns and risk concentration. What I would do is try to describe two regimes, at a minimum. One would be a fragile regime. That would be characterised by market instability and high risk concentration. The other would be a resilient regime. That would be characterised by calm or very stable returns and low concentration of risk. In recent years, there have been two measures that have evolved in the literature to measure instability and risk concentration that are better or, at least, more informative than volatility and correlation. In terms of market instability, there is a measure called financial turbulence. This is literally a measure of how statistically unusual a set of returns is in a given period given their historical pattern of behaviour. Where standard deviation deals with one asset class at a time, financial turbulence looks at a whole cross section of asset class returns. It takes into account extreme price moves. That, in a sense, is capturing the same information that you get from volatility. It also considers the decoupling of correlated assets and the convergence of uncorrelated assets. So, it is capturing the interaction among the assets as well. You can think of this as capturing two things: One is unusual volatility and the other is correlation surprise. Financial turbulence is a much better measure of market instability than conventional measures such as volatility or credit spreads. It is also the case that it has some very nice empirical features. One is that returns to risk, measured in a variety of different ways, are much lower when markets are turbulent than when they are calm. Furthermore, losses occur when markets are turbulent not when they are calm.

The other component of fragility is what we call risk concentration. Literature has shown that you can compute something called the absorption ratio to measure how concentrated risk is. The way this works is that you conduct a principal components analysis to identify the factors that are driving the variability of returns. You then compute the fraction of total variability that is explained by a few of the most important factors. So, if this ratio is high, in other words, if these few factors explain a high percentage of the variability of returns, that tells us that markets are very tightly coupled; they are unified. When risk is concentrated that way, conditions are very fragile because shocks travel quickly and broadly.

When the same few factors explain a small percentage of the total variation of returns, which means that the absorption ratio is low, that indicates that risk is distributed across many different sources. When that is the state of the world, markets are more resilient. For example, imagine if you had a situation where the absorption ratio was very low, risk was very widely distributed, and you got a shock such as an unexpected jump in oil prices. It might be the case that airlines stocks go down because their operating expenses have gone up unexpectedly. But you wouldn’t necessarily expect that shock to travel to other parts of the market where there is no fundamental connection to the price of oil. But if the market were very tightly coupled where returns are moving in unison and you got a shock like that, it would not be at all unusual for the entire market to sell off or to have a system wide response. This is why the absorption ratio is also used by policymakers to measure systemic risk.

Just to sum up, you can distinguish fragile market conditions from resilient market conditions by monitoring these measures of financial turbulence and risk concentration. This is something that one should take into account not only in modifying your exposure to risk through time but also in figuring out what your policy portfolio is in the first place. For example, if you want to build a portfolio that is diversified against losses, you don’t want to look at the correlations and volatilities that prevailed, on average, across the entire history of returns. It is much more effective to pay attention to the volatilities and correlations that prevailed during these periods of market fragility because that is when losses typically occur.

This also leads to this issue of policy portfolios. You think of a policy portfolio as a set of weights that you are going to hold, on average, through all market environments. It turns out that a set of fixed weights delivers a very unstable risk profile. For example, the typical institutional portfolio going into the financial crisis in 2008 had a trailing annual standard deviation of monthly returns of about three percent. Coming out of the crisis that same portfolio had a trailing annual volatility of about 30 percent. When you think about it, what is the purpose of a policy portfolio? Well, investors want two things, whether they are a SWF or a private investor. They want to grow wealth and to avoid large drawdowns along the way. The purpose of a policy portfolio, or at least one of the purposes, is to balance those two trade-offs which conflict with each other. The more you structure a portfolio to grow wealth the more you expose it to large losses. So the idea of a policy portfolio is to come up with how you want to balance your desire for growth with your aversion for these large drawdowns. Well, it's not really a set of weights that you want. You want the risk profile that you think that set of weights is delivering. What makes more sense in my view, rather than having a policy portfolio of rigid asset class weights, is to have a flexible investment policy. The idea is to target a certain risk profile and then modify your portfolio in some structured and dynamic way to try to maintain that risk profile. What that means is that in periods when markets are very fragile you would try to skew your portfolio towards more defensive assets and in periods when markets are very resilient you would try to orient your portfolio towards growth assets.

**Q:** *We haven't yet addressed the issue of long-term versus short-term risk. Could you share your insights on that topic?*

**MK:** It is an important topic and I want to make sure we get this issue out on the table. The way people measure risk relies, typically, on two assumptions. One is that correlations do not change depending on the return interval used to estimate them. The academic literature, as well as the software that practitioners use, assume that, over the same sample, the correlation will be the same regardless of whether you are estimate correlations from monthly or annual or daily returns. That turns out not to be true.

The other assumption that people make is that volatility, particularly standard deviation, scales with the square root of time. So, for example, if you were to estimate the standard deviation of an asset class based on monthly returns you would multiply that standard deviation by the square root of 12 to get an estimate of the volatility of annual returns. That also is not borne out by the data. It turns out that to the extent autocorrelations are not zero then that square root of time rule does not work. If you have positive autocorrelation that means that the risk of annual returns is going to be greater than the square root of twelve times the volatility of monthly returns.

In the case of correlations, not only do you have to pay attention to the autocorrelations of both return series you also have to pay attention to the lagged cross correlations. To the extent any of those are not zero then correlations will not be constant across different return intervals. A good example is the correlation between U.S. stocks and emerging market stocks. During the period starting in 1990 through 2013, both emerging market stocks and U.S. stocks had about the same cumulative annualised returns. One had a return of 9.3 per cent and the other had a return of 9.5 per cent. Moreover, their monthly returns were 69 per cent correlated. Yet, there was one three-year period when emerging market stocks outperformed U.S. stocks cumulatively by 120 per cent and there was another three-year period where they underperformed cumulatively by 60 percent. That is somewhat of a puzzle. How can you have you two asset classes that have the same cumulative returns and monthly returns that are highly correlated and experience such divergent performance in these sub-periods? Well, it turns out that the correlation of monthly returns was 69 percent, the correlation over the same sample period of the annual returns was only 40 percent, and the correlation of three year returns was zero. They were uncorrelated at the three-year return interval. This is a big deal. When you are building your portfolio, typically you are estimating your risk parameters based on monthly returns and then you are converting them to annual inputs. My presumption is that when you build a portfolio, when you do asset allocation, you are designing a portfolio to be optimal over some multi-year horizon. The risk profile over that multi-year horizon is going to be vastly different than what you are going to infer from volatilities and correlations estimated from monthly returns. This is something that one needs to address.

I would also argue that institutions like to say they are long-term investors and they can withstand large drawdowns along the way. I've been in this business for over forty years and that is not the case. People may like to think that they are long-term investors, but people do care about what might happen along the way. You can say that I'll structure my portfolio based on estimates of long-term risk, but then, if you do that you are going to make your portfolio vulnerable to large interim drawdowns. If you focus on just short-term risk, you are going to expose your portfolio to sub-optimal growth over the long-term. This is something that must be balanced.

**Q:** *So you now have two covariance matrices for different time horizons. How do you go about balancing those?*

**MK:** What you would like to be able to do is to estimate correlations and volatilities based on monthly returns and then estimate correlations and volatilities based on three-year returns then come up with two different covariance matrices and introduce both of those into the optimisation process. The problem with that approach is that the lagged correlations are not necessarily constant through time. So you may have some periods where there are positive autocorrelations. In which case, longer horizon risk is going to be greater than you would expect. Then there are going to be cases when there are negative lagged correlations, where longer horizon risk is going to be lower than you would expect.

There is a new approach that just recently appeared in the latest issue of the Journal of Portfolio Management which deals with this kind of estimation challenge by measuring the relative stability of covariances and uses that information in the portfolio construction process as a separate component of risk.

When we build portfolios, we need to estimate returns and risk. We know that those estimates are made with error. I'm not going to focus on return right now because most people don't extrapolate historical means to estimate expected returns. They typically use equilibrium returns or they have some fundamental approach to doing that. However, most investors do extrapolate historical covariances. To be clear, when I use the term covariance I am using it interchangeably with volatilities and correlations. When they extrapolate historical covariances it exposes them to several different types of errors. For example, typically what we have is some long history of returns for the asset classes that we care about. It could be decades long. What we are trying to do is to build a portfolio that is optimal for some shorter future period, such as one to three to five years. That means that we are exposed to small sample error because the realised covariances in the small sample that reside within this larger sample are going to be much different than the covariances of the large sample. So we have small sample error. We also have independent sample error because the future period that we are designing the portfolio for is distinct from the history we have used to characterise that future period. Finally, we have what we call interval error. This is what I have just been talking about; that the covariances that you estimate from monthly returns are not easily mapped on to covariances of longer interval or longer horizon returns. So we have these three components of error.

What we have developed is a way of measuring the relative stability of the asset class covariances. We are then able to build portfolios that use this information to quantify risk in a more holistic way. One way to think about this is, in terms of standard deviation since getting your head around covariances can be hard, that you can have two assets; one with a higher standard deviation than the other. It could be the case that the asset with the higher standard deviation is more stable. In other words, there is less estimation error around it than the one with the lower standard deviation. In which case, it is possible that the asset with the higher standard deviation is less risky than the asset with the lower standard deviation. This is because out of sample the asset with the lower standard deviation can have a much higher standard deviation. This applies to correlations as well. So, what we are arguing is that the relative stability of the covariances is something that one should account for when building portfolios. The experiments we have done show that this generates much more stable portfolios than ignoring errors. It also generates much more stable portfolios than the conventional approach to dealing with estimation error, which is Bayesian shrinkage.

**Q:** *So you are basically incorporating information about the volatility of volatility in the portfolio construction process?*

**MK:** Yes.

**Q:** *In the time we have left, we did want to address some questions directly from members of the IFSWF. The first question is as follows: It seems that various parts of the world are going to go through a long period of very low interest rates. How do you think this will change the way we look at these types of investments?*

**MK:** You can address that in several ways. One is to define interest rate regimes. You can then characterise your estimates of future return and risk of portfolio components contingent on what regime you expect to be in.

When you conduct an optimisation, what you are doing is maximizing expected return minus some coefficient of risk aversion times portfolio risk. That portfolio risk is characterised as a covariance matrix. So, what you can do is to collect a long history of returns. You have information about when interest rates were low in history and when interest rates were high in history. Instead of basing the risk of the asset classes on the full sample of historical returns, divide the historical returns into two samples. One sample would be returns when interest rates were below some level and the other would be returns when interest rates were above some level. You would then calculate separate covariance matrices and condition expected returns based on what prevailed in the low interest rate regime versus the high interest regime. Then, when you optimise your portfolio, instead of maximizing expected return minus risk aversion times one covariance matrix, you would maximise expected return minus one risk aversion coefficient times covariances estimated from the low interest rate regime minus another risk aversion coefficient times covariances estimated from the high interest rate regime. So you have two interest rate regimes. Earlier I spoke about a fragile regime and a resilient regime. You can take the same approach but have it be conditioned on these different rate environments. Then the risk aversion coefficient that you assign to these two covariance matrices can either reflect the relative aversion you have toward risk during periods of high or low interest rates or, instead, it can reflect your expectation for what the future will hold. You may argue, and I would argue, that interest rates are more likely to be higher in the future than they have been in the recent past and you might put a higher probability on that when you do your optimisation. That is one approach.

The other thing to keep in mind is that interest rates historically, at least in the United States, have gone through very long cycles. We had a declining interest rate environment from 1979 through just about the present. Short rates went from about 20 per cent down to zero. It is unlikely that that trend can continue. It can’t continue without going significantly negative. I think it’s more likely that we’ll have a long and gradual increasing interest rate environment.

The other thing that this implies is that the risk from fixed income assets is much greater than you think it is. For example, there is a strategy called risk parity. What that means is that you structure a portfolio such that each of the major components of the portfolio contributes the same amount to total portfolio risk. So you should lever up your exposure to bonds and cut your exposure to equities. People have written articles showing that this risk parity strategy approach has outperformed a 60/40 stock/bond portfolio going back to the 1920s. It turns out, that is not true. They based that performance on the Sharpe ratio which has as its denominator standard deviation. They converted the standard deviation of monthly returns to the standard deviation of longer horizon returns using that heuristic I described earlier. If you take into account the lagged correlations of the asset’s returns then the 60/40 portfolio outperformed the risk parity portfolio by as much as they argued it underperformed.

Anyway, the short answer is…and I have trouble giving short answers…I would condition my expected returns and risk estimates on the sub-samples of high and low interest rates and use that information to build my portfolio.

**Q:** *I like the discussion and identification of this richness of risks and I understand the statistical qualities of these other risk measures. What it presents is added complications regarding optimisation and determining what is a best portfolio…especially if you have multiple objectives. Generally, my board is happy if we do well versus public plans, if we don’t have a high risk of losing money, if we show actuarial progress, or the equivalent of actuarial progress, towards some long-term goal. It sounds like what you are describing is that, in general, the profession has made more advancements in risk measurement than on the optimisation side. What is your view?*

**MK:** Well, I think there is a lot of misunderstanding of optimisation. Let's talk about mean variance for a minute. It turns out that mean variance is much more robust than people give it credit for. I am a big fan of Harry Markowitz, and he and I have had many discussions about this. Mean variance optimisation requires one of two things. Either that returns are approximately normally distributed or that investors have preferences that can be reasonably described by just mean and variance. You do not need both of those to be true. You just need one or the other to be true. So, mean variance does a pretty good job.

Now, you can amplify mean variance to take into account multiple objectives like you've just described. For example, you may care about performance relative to your peers and you also may care about your absolute performance. So, just as I described about how you can come up with covariance matrices based on different regimes, you can come up with covariance matrices based absolute returns and covariance matrices based on relative returns. You can specify the objective function of mean variance optimisation to be expected return minus absolute risk aversion times the covariances of absolute returns minus a measure of aversion to relative risk times the covariance matrix of relative returns. So, basically, you are jointly optimizing for both absolute volatility and tracking error relative to some portfolio of peer investors or some benchmark. That is one thing you can do that is trivial to implement.

In terms of pension liabilities or actuarial progress, as you describe it, that is a really interesting question. This is research that we are actually doing right now and I'll be giving a talk at Oxford University in a couple of months on the topic. If you want to hedge the monthly volatility of your liabilities, for example, the best hedge would probably be some kind of fixed income asset. This is because high frequency volatility of liabilities is typically a function of changes in discount rates. To be clear, when I say high frequency I mean monthly versus say yearly rather than milliseconds. Bonds would be the best hedge for that. But over the long term, the low frequency volatility of liabilities is a function of wage inflation and productivity growth. Equities are a better hedge against that. Again you can construct an optimisation process that balances your aversion to large drawdowns along the way versus your aversion to the gradual erosion of your pension assets relative to your liabilities. That is another thing you can do in the optimisation process.

To the extent that you or your committee or stakeholders have preferences that can't be well described by mean and variance, there is another thing you can do. A typical example of this is thresholds. If there is some threshold where if you breached that threshold conditions would be qualitatively worse than if you suffer a loss above that threshold, this is what we call a "kinked" utility function. If you have a situation where your returns are not normally distributed and you have preferences that are affected by these thresholds, then you can't use mean variance optimisation. What you would use is what is called full-scale optimisation. Full-scale optimisation is just plain direct utility maximisation using sophisticated search algorithms.

So you write down your utility function. You have some sample of returns. You plug those returns into the formula for your kinked utility function and then you plug in a portfolio with one set of asset weights and calculate the utility. Then you plug in another portfolio with another set of asset weights and calculate the utility. You do this repeatedly until you find the portfolio that has the highest utility. Now, that is computationally very challenging, especially if you have portfolios that have more than just a few assets in them. However, it turns out that there are optimisers that run full-scale optimisation that can sample as many as half a million portfolios in about 30 seconds. So, this is what I would use instead of mean variance optimisation in the case where you believe returns not to be approximately normally distributed and you have thresholds.