

Adam Butler: [00:00:01](#)

Okay, we're now recording. Mark, I want to welcome you to the ReSolve Podcast series. This is specifically the institutional series, so we are intending to go fairly deep on some topics, and really dive into some of the nitty gritty of the underlying concepts. Today, we are here with Mark Kritzman, who is the founding partner and Chief Executive Officer of Windham Capital Management, out of Boston.

Mark, I was going through your bio page, and it is clear that you are truly a prolific author and industry pioneer in the subject of risk management, asset allocation, and, as I scan through the topics, you touch on virtually every corner and niche of the finance domain. It's really remarkable. Over 40 years of industry experience, you've contributed significantly to the world of academic research, your works have been showcased in prominent investment-related publications, and I see you've received significant industry awards, such as seven Bernstein Fabozzi/Jacobs Levy Awards for outstanding articles including one Best Article Award in 2003.

You've co-authored ten books, including "A Practitioner's Guide To Asset Allocation", which we're going to spend a fair amount of time on today. This really is a truly astonishing body of work. So, before we get started on the technical topics, I think it would be helpful to take us through the arc of your career, and maybe talk about what motivated you to launch Windham in the late 1980s.

Mark Kritzman: [00:01:47](#)

Sure, Adam, thank you. I think the arc of my career has been determined in large part by the year that I was born, because that determined the year that I graduated from college and start looking for a job. And so, I graduated at a time when there was a lot of turmoil in the markets - not just in the investment markets, but in the industry. So, I started my career around the time the market was in free fall. In fact, I think it was the largest price adjusted drop in stock prices in history - more so than the Depression, and I'm not sure how it stacks up against the financial crisis, but it was a mess.

And, at the same time, Congress enacted ERISA, the Employee Retirement Income Security Act, and that changed the liability of investment professionals. It made them fiduciaries. And so we had this sort of confluence of really bad market performance with a more stringent regulatory environment, and that prompted people to look for sounder, more defensible ways of managing assets.

And, up until then, I think it's fair to say that asset management was largely a sort of heuristic-driven business. So professionals then started looking for guidance from the academic community. So they looked at the work that Harry Markowitz did on portfolio selection. Of course, he published his paper in 1952, but it really wasn't until the early to mid-70s that the industry started paying attention to it. And they also started paying attention to the work of Bill Sharpe and Jack Treynor, and then of course later, there was Black, Scholes and Merton. And I can go on and on.

But in any event, I was starting my career right at the time that the industry was looking to academia for guidance, and so I had an opportunity to learn finance beyond

what I learned in school, and I think I had a pretty good education, but a lot of what was taught in school when I went there ... It's always tough to acknowledge how old you are. Thank you for reminding everybody earlier. You failed to mention that I started my career when I was in high school.

But in any event, I had the opportunity to actually learn finance directly from these people, from Harry Markowitz, and Bill Sharpe, and Jack Treynor, and Fischer Black, and Merton. And because I was going to seminars and conferences that they were going to, they were trying to get the industry to take them seriously, and I was trying to learn what it is they were doing, and I just happened to be in the right place, right time. So that's sort of ...the arc of my career - at least my intellectual career.

And I've always been intellectually curious. Every day, I feel if I learn something new that day, it's a good day. And I still feel that way. So, I worked for a few large institutions early in my career, and then was persuaded by some people to start a firm called New Amsterdam Partners, and we started that, and that still exists today, and I stayed there for a few years, but then I moved on and co-founded Windham Capital Management.

And I really feel very, very happy in retrospect for having done that. I feel being in a small company, being your own boss, not having to deal with lots of bureaucracy and politics, really enables you to spend more time on research.

Adam Butler: [00:06:40](#)

I can certainly relate. I know that many of our listeners will also relate to that entrepreneurial drive and fulfillment. What does Windham focus on, exactly?

## Windham's Focus

Mark Kritzman: [00:06:52](#)

Well, we have three business lines. Asset management, and then we develop and distribute portfolio construction and risk management software, and we also provide advisory services to a couple of large institutions. On the asset management side, we have two general products. One is global tactical asset allocation, and that is a risk-driven strategy. And then, we also have a risk premia strategy, and that's also largely risk-driven.

So that's our asset management business. Those are the two broad types of strategies that we manage. There are lots of variations within those two flavors. And then, as I mentioned, we have a software business where we distribute mainly a suite of software applications that is based on the research that I've done with some of my colleagues throughout my career.

Adam Butler: [00:08:09](#)

So, based on the narrative arc of your most recent book, "A Practitioner's Guide To Asset Allocation", and just as in our earlier conversations and in perusing your website, it seems like you do have a core focus on the asset allocation decision, and less of an emphasis on security selection. Would you say that's fair?

Mark Kritzman: [00:08:36](#)

Yes.

## Why Asset Allocation?

Adam Butler: [00:08:38](#)

Why is that? Why have you chosen to focus on asset allocation?

Mark Kritzman: [00:08:43](#)

Well, I think partly, by historical accident, my first job in the industry was in the asset allocation department. But I basically have come to respect the efficiency of the securities markets, and I've come to embrace what people refer to as the Samuelson's Dictum. So Paul Samuelson, who was a professor at MIT, and I knew him, I've taught at MIT, I'm going on my 17th year there, and I overlapped with him for quite a few years and got to know him, but in his life, he's a remarkable human being.

But he put forth the notion that the markets are micro-efficient and macro-inefficient. And by micro-efficient, he meant that it is very difficult to add value by choosing individual securities, and by macro-inefficient, he meant you're more likely to add value by making broad asset allocation decisions.

And his reasoning was that, if a particular company is mispriced, some smart investor will notice this and will trade to take advantage of that mispricing. But by trading, that investor will correct the mispricing. But if an aggregation of securities is mispriced, such as an asset class, and some smart investor recognizes this and trades to exploit that opportunity, most investors just don't have the scale to revalue an entire asset class by trading.

So it typically requires some kind of exogenous shot to jolt many investors to act in concert to revalue an entire asset class. So what that means is that these macro-inefficiencies tend to persist for longer periods of time, and that enables more investors to take advantage of it. That, in a nutshell, is the Samuelson's Dictum.

Adam Butler: [00:11:12](#)

So, is it exclusively a function of the amount of capital that's required in order to drive large asset classes back toward equilibrium, or are there potentially other structural barriers to arbitrage for most investor constituencies that also make that challenging?

Mark Kritzman: [00:11:35](#)

Well, the Samuelson's Dictum, at least his reasoning, what we can read about, has to do with the amount of capital required to correct a mispricing. Now, there may be other macro-inefficiencies that make it easier to add value through asset allocation than security selection, but-

Adam Butler: [00:12:04](#)

I'm thinking more along the lines of, your typical large institution may not have the portfolio agility, the structural portfolio agility, in terms of how decisions are made, or the mandate flexibility in terms of allowable tracking error from a policy portfolio, and probably similar challenges and constraints at the advisor and retail investor level. Have you observed that throughout your career, or is that less of a phenomenon?

Mark Kritzman: [00:12:36](#)

Yeah, that's very perceptive of you. I think that's exactly right. I think, to the extent there's active management going on, it's largely security selection. Asset allocation is done more as a long term strategic policy, and it's not driven by an investor's perception of the relative value of different asset classes, but rather by their expectations about long term returns and risks of different asset classes. So there's just not a lot of capital that's brought to bear on mispricing at the asset class level, or at least historically, that's been the case.

It was discouraged by an influential article by Bill Sharpe, many years ago I believe, that was published in the Financial Analyst Journal, where he basically said, "You shouldn't do market timing." And at the time, I think his reasoning was quite solid, because at the time, it was very costly to do so. So those of us who were doing market timing said, "Okay, we won't do market timing anymore. We'll just do tactical asset allocation."

So then we went through periods where tactical asset allocation went out of favor. It's sort of out of favor now, so we don't do that anymore, we just do risk management. It's all the same activity. We just come up with more euphemistic labels. But in any event, I think you definitely are onto something by saying that the focus of institutional asset management has not been, historically, to exploit inefficiencies across asset classes, but rather within asset classes.

And if I can expand upon that-

Adam Butler: [00:14:43](#)

Please do.

Mark Kritzman: [00:14:46](#)

The institutional industry for many years, probably decades, has embraced this notion of a policy portfolio. So they would do their ... They'd estimate the expected returns and risk of the major asset classes along with their correlations, and they'd come up with an asset mix policy. If it were two asset classes, it would be 60% stocks, 40% bonds, but they obviously diversified across real estate and commodities and other things.

But if you think about it, why do they want an asset mix policy? Well, investors want two things, whether you're an institutional investor or a private investor, you want two things. You want to grow your wealth, and you want to avoid large drawdowns along the way. Well, these are conflicting objectives. The more you design a portfolio to grow wealth, the more you subject it to large drawdowns, and vice versa.

So the asset mix policy was thought of as a way of balancing these two conflicting objectives. But if you think about, investors don't really want the asset mix. They want the return distribution that they think is associated with that asset mix. So that's how they justify an asset mix policy. The problem with that, and this is recognized more and more today, is that the return distribution associated with a static asset mix policy is highly unstable.

So, going into the financial crisis, the trailing annualized, monthly volatility over the prior year for the typical institutional portfolio, say something like a 60-40 portfolio, was only about three or four percent. Coming out of the financial crisis, it was approaching 30%. So that's obviously a unique period. But it still ranges from that three percent to maybe 15 or 20%, depending on what period you look at, and that's pretty unstable.

Now, Peter Bernstein, who is a very dear friend of mine, in fact, there's a picture in my office of him right over there. And he ... put forth this notion, he basically said, "Policy portfolios are obsolete." But I don't think he articulated it very clearly. People didn't really know what he meant, and they would say to him, "Well, what's the alternative?" And I think what he really meant is what I've just described, is that a set of ... A portfolio of fixed weights is very unstable in terms of the risk profile that it delivers.

And what Peter had in mind, and I'm just putting words in his mouth, is that rather than adhere to a rigid asset mix policy that delivers a very unstable risk profile, why not have a flexible asset mix that generates a more stable risk profile? I think investors are coming to embrace that notion more and more and I think they're also realistic in that they do not believe that the risk premium is going to be as high as it has been historically, and that a static exposure to risky assets just isn't going to meet their needs, their spending needs, and that there needs to be some kind of reallocation to try to increase the return you could get.

So I think you're exactly right, that historically, industry did not engage meaningfully in asset allocation. But I think today, industry is much more open to the notion of dynamic asset mix.

Adam Butler: [00:19:36](#)

At the margin.

Mark Kritzman: [00:19:37](#)

Yeah.

Adam Butler: [00:19:38](#)

Right. That seems to be a bit of a common theme throughout your work, and certainly, your most recent book. This idea that investors think about the portfolio problem oftentimes from the perspective of wanting highly stable weights and tolerating an unstable return profile, when really, we should be primarily concerned with a more stable return profile, and less concerned with unstable weights, notwithstanding, obviously, the impact of transaction costs, and tax, that sort of thing.

Mark Kritzman: [00:20:14](#)

Right. And yeah, that's exactly right. And that gets back to my earlier observation that the policy portfolio is just a proxy for a risk profile. The weights aren't, in and of themselves, that percent - it's what return and risk do you expect those weights to produce?

And people have implicitly assumed that they're going to generate the same distribution year in and year out. Now, clearly, even if they did generate the same distribution, you'd still have ... the distribution can be quite wide. But the fact is that risk is time varying, and I think people are much more ... of that today.

Adam Butler: [00:21:06](#)

Yeah, I would agree. And this actually dovetails really nicely into a discussion on one of the most common arguments against the use of portfolio optimization, which is this idea that it is an error-maximizing process, and people typically describe it as an error-maximizing process because of the instability of weights, when, in fact, what we observe of course, is that while the weights are maybe unstable in certain conditions, that the actual character of the portfolio may not change very much. Can you elaborate a little bit on that?

Mark Kritzman: [00:21:53](#)

Yes. I think there are many misunderstandings about optimization. And you've touched upon one of them, which is that optimizers are error-maximizers. This notion is based on the observation that, if you are allocating across some set of assets and you change, you make a mistake in your estimates of say, the expected returns, that you get a vastly different portfolio than if you didn't make those mistakes.

The examples that people who promote this view ... Which I reject, I completely reject it - those who promote this view typically will provide an example in which the investor is presumed to be allocating across a set of assets that are very close substitutes for each other. So, clearly, if they're close substitutes, and you have a small error in one of the inputs, the weights are going to shift quite a bit.

However, it doesn't matter, because they're close substitutes. Although the weights shift a lot, the distribution of returns associated with the correct portfolio and the incorrect portfolio are very similar to each other. The exposure to loss is very similar to each other, the opportunity for gain is very similar to each other.

So that's sort of a disingenuous way of trying to make this point. Now, if instead, you consider allocation across a set of assets that are quite different from each other in terms of their expected return and risk, and you introduce a small error, say, to the expected returns of some of the assets, then the incorrect portfolio will not have very different weights.

And again, you're going to have return distributions that are pretty similar. So I think that argument is over hyped, which is not to say that estimation error isn't important. There are a lot of techniques for dealing with an estimation error when we do optimization. In fact, a couple of colleagues and I introduced one approach a couple years ago. So, I guess, my view about error-maximization is that it's more hyped than reality, that estimation error is a challenge, there are very solid ways of dealing with that challenge. May I comment on a couple of other misunderstandings about optimization?

Adam Butler: [00:25:07](#)

Please.

Mark Kritzman: [00:25:08](#)

So, another assumption that we hear quite often is that mean-variance optimization requires return distributions to be normal, and also, it requires investors to have quadratic utility. Okay, that is not true. Mean-variance optimization requires that returns are elliptical, which is a much more flexible description of returns, and I'll describe what that means in a minute, or that investors have preferences that can be reasonably approximated by mean and variance.

So, normality is much stricter than necessary for mean-variance analysis. So let me see if I can get this right. So, we have what we call symmetric distribution, so if you think of a normal distribution, that's obviously symmetric. And then, a special case of symmetric is an elliptical distribution. And then a special case of an elliptical distribution is a normal distribution. It's the elliptical distribution that mean-variance depends on, because if you have an elliptical distribution, that means that mean and variance can describe the entire distribution.

Now, I'll give you an example of why a symmetric distribution isn't good enough. So, think of two assets that have equal variances, and half the time, they're 80% correlated, and the other half of the time, they're -80% correlated. That will give you a symmetric distribution. If you think of a scatter plot, you have one cluster that's leading to the right, and it's like an X, another cluster is leading to the left. Well, mean-variance can't deal with that.

An elliptical distribution means that if you were to draw concentric circles, or ellipses, around the average of the returns, that the observations would be evenly spread around those circumferences. Okay, that's what's required. Normality is a special case of that, but it's much stricter than is required by mean-variance.

Now, then the other assumption that people impose upon mean-variance, or requirement, is that investors have quadratic utility. Nobody has quadratic utility, because what that means is that at a certain level of wealth, you prefer to give money away. I'm not talking about being charitable. You prefer to get less money than more money, because it bends down. Quadratic function bends down.

But mean and variance can ... What optimization requires is that we have what's called power utility, and that never bends down. It's like a ... wealth utility function, something that where you have utility on the vertical axis, you think of utility as just happiness or satisfaction, wealth on a horizontal axis, it increases at a decreasing rate.

Adam Butler: [00:28:56](#)

Diminishing marginal benefit from each incremental dollar.

Mark Kritzman: [00:29:00](#)

... wealth is the prototypical example, and other variations of power utility just have different degrees of curvature. The greater the curvature, the more risk of ... you are. But it never bends down. Well, it turns out that mean and variance can approximate almost ... Just all power utility functions within a range of returns of -40% to +60%. So it's good enough, okay.

Now, where does mean-variance fail? It fails if either returns are not elliptically distributed, or if investors have preferences that cannot be described by mean and variance. So what kind of a preference would that be? It would be a situation where there's a kink in the utility function. Another way of saying that is, investors face a threshold, and if that threshold is breached, their utility changes abruptly and sharply.

You can think of the situation for a private investor, if you lose a certain amount of money, your spouse is going to divorce you. For an institution, like a pension fund, if you lose a certain amount of money, you may go from a surplus to a deficit in terms of your asset liability ratio, and all kinds of trouble can kick in when that happens.

So it's these kinds of ... Or if there's a company that ... money, and they have certain covenants, it's when there's a qualitative change in the satisfaction or dissatisfaction that you experience at a particular threshold. So mean-variance can't handle that.

Adam Butler: [00:30:57](#)

So just sort of thinking through the example you gave of two assets that have a positive correlation, a .8 in one regime, and a correlation of -.8 in a different regime. And it doesn't need to be that extreme, but just sort of directionally that sort of character, is there any static portfolio or policy portfolio that can be constructed that would be optimal, given the extreme shifts in behavior through time?

Mark Kritzman: [00:31:40](#)

There is, and it's pretty complicated to describe, but basically, what you're saying is, there's an estimation error, right? You're estimating the correlation to be positive .8, and that's an extreme example. But let's say you're estimating some correlations to be positive, they turn out to be negative, et cetera. So, what you want to do is ... This can get very long-winded and technical, let me just go off on a tangent just for a minute, and then I'll come back to this other issue that you're raising.

An alternative, of course, is rather than have a static portfolio that's designed to weather all kinds of different regimes, is to build a dynamic strategy that predicts the relative likelihood of being in one regime or the other. So that would be my first choice, and there's really good technology to do that.

If you're not going to do that, and just say, "I want a portfolio," then the way you deal with that is, you estimate the relative stability of the covariances within your portfolio. So, to make this a little bit more intuitive, most people don't have a good intuitive grasp of covariance. They understand what it is, but they don't know, if they give you a certain number, whether that's a high or low covariance.

So a covariance is just a combination of a correlation in the two standard deviations. So let's just think about standard deviation, as a simplified way of trying to get at this issue. So let's suppose that you have two assets that you're allocating to, that you're trying to optimize across several assets, but right now, focused on these two assets.

And let's say one of the assets, and let's say that you're estimating the standard deviations based on what they were historically. And let's say one of the assets has a

higher standard deviation than the other, but the one with the higher standard deviation, that standard deviation has been pretty constant, historically. The one with the lower average standard deviation has been pretty inconstant, historically.

So you have estimation error around those two estimates. It turns out that the asset with the lower standard deviation can be riskier, because its standard deviation is less reliable. The same thing is true of covariances. So one way of dealing with this problem that you've brought up, this extreme example of the plus and -80% correlation, is to determine the relative stability of the covariances and then introduce that into the optimization process as a separate, as a distinct source of risk.

So the risk of an asset is a combination of its standard deviation and correlation, but this other risk is how unstable those estimates are. And it turns out, this almost sounds like a contradiction, but it's not, it turns out that instability, relative instability, is relatively stable. So in other words, if a particular asset pair, historically, had a very stable or reliable covariance, let's say a more stable, more reliable covariance than the covariance of another pair of assets, that's likely to persist in the future.

That can be shown empirically. Let me try to give you an intuition about why you should expect that. It's that if two assets have a very large, in an absolute sense, correlation, a very big positive correlation or a very big negative correlation, those large correlations are relatively stable, because there's probably some structural reason that those correlations are so big, right? But if assets have relatively small correlations that fluctuate between plus and minus, 10 or 20%, there may not be any structural reason for those assets to have any particular correlation, and they tend to be noisier.

So, what we want to do is, when we build a portfolio, is rely more on the assets that have stable covariances, and less on the assets that have unstable covariances. So that's a long-winded way of saying it, and then let me just say, if you want to know precisely how to do that, it's in the book.

## Stability Adjustment

Adam Butler: [00:37:48](#)

So it's your stability adjusted, it's the optimization, right? Which I found really fascinating. And I was wondering whether or not you applied the same principles of the stability adjusted optimization to construct dynamic portfolios as you do to construct strategic portfolios. It seems to me that the estimation error is similar and present, but just at a different time frame, and so you can get maybe in less trouble by explicitly acknowledging the error term from a tactical sense, as well as from a strategic sense.

Mark Kritzman: [00:38:31](#)

I haven't thought too much about that, but I think that when you do something like using a hidden Markov switching model, that process implicitly takes into account the

relative stability of the covariances. So it's not done explicitly like we described instability adjusted optimization, but it's implicit in the way Markov process works.

Adam Butler: [00:39:05](#) Right. And you said a specific type of Markov process in the book that seems to involve some sort of cross-validation process, the Baum-Welch algorithm, it looks like.

Mark Kritzman: [00:39:23](#) Yeah, you really want to get technical today, don't you?

Adam Butler: [00:39:26](#) No, we don't need to go into it, but I just think some readers, or some listeners, rather, may be interested in digging a little bit more deeply into this specific algorithm. I was fascinated by it, because I haven't had much luck with Markov switching, but I really like the way that the Baum-Welch algorithm approaches that. I think it's uniquely suited to time series analysis, which could be highly non-stationary. The cross-validation steps, I think, make a lot of sense.

Mark Kritzman: [00:39:55](#) Yeah, I think it's explained reasonably clearly in the book, and there is actually a chapter in the book. You know, we should've put this up front. But it's chapter ... Section 4, chapter 17. And what we do in somewhere between six and eight bullet points, we just cover every chapter in the book. It's the Cliff's Notes. It's pretty efficient, so you might want to start, as your sort of point of entry to the book, with chapter 17, and then to the extent you're more or less interested in a particular topic, you can then go to the chapter that covers it in detail.

Adam Butler: [00:40:42](#) Yeah, just go a little deeper. Continuing on the theme of optimization, I know there was a paper, I want to say 2009, maybe, DeMiguel, Garlappi, Uppal, that made the case that optimization doesn't work in practice, and that the one over N portfolio is a highly legitimate and attractive contender. We've written some on that, about some of the flaws, but I'm interested in your perspective on that.

Mark Kritzman: [00:41:12](#) Yeah, I think it's flawed. There are three fundamental reasons it's flawed. First reason is it's not true. And by that, I mean the way they approached it, as I recall, is that they extrapolated historical means and covariances from relatively short samples, and then they built a portfolio, and then they looked at how it performed out of sample, and then they continued that process, and they found that on average, the one over N portfolio was more reliable out of sample than the optimized portfolios.

The flaw is in their methodology and, in particular, that they extrapolated means to serve as expected returns. Nobody in industry uses short term historical means as expected returns, and in fact, in many of the samples that they had in their experiment, the historical means of less risky assets were greater than the historical means of riskier assets, which means that, if investors then assume that that was going to carry forward to the future, it would imply that they were risk seeking, rather than risk averse.

So I attributed it to the fact that they didn't have a lot of experience, or knowledge with how practice actually conducts portfolio construction. So that, I think, is the

reason they came up with the result that they came up with. And we wrote a paper called "The Fallacy of One Over N," and I think that we then made it into a chapter in our book, and what we showed is that if you just have reasonable assumptions for the expected returns, and it could be equilibrium returns, it could be just returns that are proportional to risk, it could be anything that's reasonable, or just use a longer history to estimate the expected returns, then the optimized portfolios, on average, were much better out of sample than the average weight portfolios.

So that's, I think, the flaw in their methodology. But let's say it was true that equally weighted portfolios were shown to be more reliable out of sample than optimized portfolios. You still don't want to use equally weighted portfolios, because it assumes that all investors have the same risk aversion. There's only one portfolio, so if you're very conservative, you're forced to hold the same portfolio as someone who's very aggressive. So it's not such a good idea.

And then the other way to think about it is, there's a certain arbitrariness to it, so that if you have five asset classes, and then you decide that it makes sense to divide one asset class into two asset classes, then by equally weighting, you double the weight in the asset class that you began with. That strikes me as pretty arbitrary. So there's lots of good reasons not to rely on equally weighted portfolios.

Adam Butler: [00:45:10](#)

Yeah, there's a clear fallacy of composition which negates it as a macro-consistent option, anyway. And I was especially fascinated at the fact that the authors, they used sector and industry portfolios primarily, if I recall, and so of course, you've got highly co-integrated, highly correlated time series that you're trying to optimize anyway, so it's going to be unstable. And then, with each of the investment universes, they also added back the market cap-weighted index, just to confound the optimization even further. They really went far out of their way to make it very difficult for the optimizer to be able to add value and produce stable solutions.

Mark Kritzman: [00:46:07](#)

Right, and then on top of it, they didn't end at all toward the notion of addressing estimation error.

Adam Butler: [00:46:16](#)

Right. Of course.

Mark Kritzman: [00:46:18](#)

So you can say if you do this really, really naively, you're better off not doing it at all and just equally weighting your portfolio, okay. But you don't have to do it really, really naively.

Adam Butler: [00:46:31](#)

We always say the out of the box solution for mean-variance optimization is rarely useful in practice, but just with a few fairly minor and intuitive adjustments, it can be extremely useful.

Mark Kritzman: [00:46:43](#)

Ah, yes.

## Optimal Portfolios

Adam Butler: [00:46:47](#)

Just on the topic of optimal portfolios, we in finance get caught up a lot on how to construct optimal portfolios in terms of absolute returns and absolute risk. And we know from our experience with investors, and from a lot of the research in the behavioral finance literature for example, that investors don't just experience risk in the dimension of absolute means and absolute volatilities, or even absolute downside risk.

But investors tend to also be highly conscious of how their portfolio is tracking relative to some largely arbitrary perceived benchmark. And so, in practice, the optimal portfolio is some mix of the objective of maximizing return, minimizing risk, subject to the constraint that there's a high probability that the investor will stick with it when things get challenging, or when the character of returns diverges materially from the benchmark.

And there's a chapter you have in the book, and you've written, I know, several papers on this same topic, but how do you think about this dual problem, and this sort of dual nature to optimization?

Mark Kritzman: [00:48:22](#)

Yeah, that's a good question, Adam. I wrote a paper years ago called "Wrong and Alone," to get at this issue. And what I was getting at is, why is it that people ... So you brought up the fact that if you just run optimization without any tweaks at all, sometimes you get a corner solution, and it's not anything that anybody would actually invest in, so you tweak it and you get a better portfolio. ...

Adam Butler: [00:49:00](#)

Which happens all the time in practice.

Mark Kritzman: [00:49:02](#)

Yeah, yeah, and the way that typically happens is, they impose constraints. So then, when you ask people, "Why are you imposing constraints?" The answer that you will get is, "Because I don't have sufficient confidence in my inputs." And that's not the real answer. They may think that's the real answer, but if they reflect carefully, they'll come to the understanding that it's not the real answer.

The way to show it is to simply say, let's presume that these inputs are going to be delivered to you by some divine source, so that they are incontrovertibly correct inputs. The inputs aren't point estimate returns. They're distributions, right? And distributions, especially for stocks, are quite wide.

So even though you have the absolutely correct inputs, you're going to go through periods where your portfolio is performing very badly, and also very differently, from what you might consider to be the norm, or your peer group, or whatever it is you pay attention to.

Adam Butler: [00:50:25](#)

The wrong and alone situation.

Mark Kritzman: [00:50:27](#)

So that's why people use constraints. They don't want to be wrong. Wrong in the sense that they're losing money, and alone in the sense that they happen to be the only one that's losing money. It's okay to lose money as long as you're losing money alongside everybody else. And it's okay to be different than everybody else, as long as you're not losing money, but if you're losing money and you're different from everybody else, that's what gets you in trouble. That's what I call wrong and alone.

But there's a better way of dealing with that than constraints. Constrained mean-variance is inefficient. I like to just make the general assertion that constraints are bad in all things in life. But I just say that to be provocative. But in talking about optimization, I can say constraints are bad in a mathematically objective way.

I can come up with a portfolio ... Let's say we care about absolute performance and relative performance. So we care about standard deviation and we care about tracking error. Tracking error relative to a benchmark, could be tracking error relative to the average portfolio of a group of peer investors, could be tracking error relative to a particular factor profile.

And what I can state as just mathematically true is you can come up with a better portfolio by doing mean-variance tracking error optimization than you can by doing constrained mean-variance optimization. So, let me just tell you what I mean when I say mean-variance tracking error optimization. I should point out that this concept was introduced by a gentleman named George ... I think in 1995. He used to work with us, he's a good guy. Smart guy.

And so, the objective function for mean-variance optimization is to maximize expected return, minus risk aversion, times portfolio risk. Mean-variance tracking error optimization would simply append to that minus aversion to tracking error, times tracking error squared.

And so whereas mean-variance optimization gives you an efficient frontier, mean-variance tracking error optimization gives you an efficient surface. An efficient surface, the three dimensions are expected return, standard deviation, and tracking error. You don't need a separate dimension for relative return, because total return and relative return are linearly related to each other, but tracking error and standard deviation are not.

So on this efficient surface, the upper boundary of the surface would be the traditional mean-variance efficient frontier. The right boundary would be the mean tracking error efficient frontier. And the lower boundary would simply be combinations of the minimum risk asset and the benchmark portfolio.

So all portfolios that lie on that surface would be efficient in those three dimensions. Doesn't make them efficient if you go back to two dimensions, though. Now, what I'm saying is that the portfolios that lie on that surface are better than if you did mean-

variance optimization and then said, "Subject to the constraint, ... no less than 40% in stocks, and no more than 10% in real estate, et cetera."

Better in the sense that, for the same standard deviation and tracking error, you get a higher expected return, or for the same expected return ... you'll get a lower tracking error. And for the same expected return in tracking error, you'll get a lower standard deviation.

So, that's what I would recommend to take into account the fact that investors care not just about absolute outcomes, but about how they perform relative to others. And it's not just a matter of competitiveness, or pride, that there are often times real economic consequences associated with relative risk.

So, for example, if you are an endowment fund, and you're going to be judged based on not only how you do in an absolute sense, but if you're MIT, you're going to be compared to Harvard, and Yale, and Princeton, and Stanford, and Chicago. Well, if you do badly, relative to your peer institutions, then they're going to have a competitive advantage in terms of attracting students and faculty. So there's a real economic consequence to relative performance.

Adam Butler: [00:55:54](#)

Yeah, right. Where you've got a situation where the investment committee, the decision makers at the endowment, have career risk, as well. They may have the right portfolio in terms of maximizing the probability of sustainable distributions at whatever rate is mandated by the policies, but if they underperform their peer group for three or five years by some threshold amount that nobody knows in advance, then they would be required to move on, right?

Or they'll be more heavily constrained. So you've got to balance off this risk of being coached out of the investment team against the needs of the endowment to be able to support its distributions and perpetuity, and it's not an easy balance.

Mark Kritzman: [00:56:52](#)

It's not, and the other thing that's a bit tricky is balancing short term goals with long term goals. I remember once advising the foundation about their asset mix policy, and they posed the question as to find an asset mix that has a very low chance of losing more than 20% over some 20-year horizon, or something like that.

And so, I thought about it, and it occurred to me that as long as the expected return is greater than the spending rate ... Sorry, it's not lose more than 20%, but they didn't want the foundation assets to depreciate by more than 20%. It just occurred to me that if the expected return fund is greater than the spending requirement, then you're not going to lose ... The fund's not going to go down by more than 20%.

But I said, "So, what you're telling me, though, I just want to be clear, you're not going to meet with your investment committee until the last hour of the last day of the 20th year, right? Because there's an awful lot that can happen along the way." And I

remember the woman saying, "No, no, of course, of course not." Do you care if you're down 20% a year from now, or five years from now?

Adam Butler: [00:58:24](#)

Which is the intra-horizon risk that you spend a fair amount of time discussing in the book, as well.

Mark Kritzman: [00:58:28](#)

It was that conversation with her which led me to come up with this notion of within horizon risk, using the first passage probability. To me, that's something that's really overlooked in the industry. People just don't ... They look at what's going to happen after five years. They look at the distribution of five year outcomes, and it looks pretty benign. They don't realize that there might be only a five percent chance of losing a certain amount at the end of five years, but within that five year period, there's probably greater than a 50% chance that you're going to be down that amount at some point, right?

Adam Butler: [00:59:10](#)

Right. So it's sort of a joint objective optimization, where you want to find a portfolio that minimizes the probability of being down more than X percent, and X percent, expect ... )over five years, but also simultaneously, no more than an X percent ... fall over 20 years, right?

Mark Kritzman: [00:59:34](#)

Which becomes very complicated, yes. It does. But it is reality.

## New Research Focus

Adam Butler: [00:59:42](#)

Absolutely. Okay, just in the spirit of wrapping up and sort of looking forward, I'm curious about what you're most excited about from a research standpoint, if you look out over the next, say, one to three years. What are you looking at right now that's really piqued your curiosity and you want to dig into?

Mark Kritzman: [01:00:03](#)

Well, let me tell you a little bit about a paper that we've just submitted, and it's called, I think it's called "Crowded Trades," but it does ... Implications ... I forget what it's called. Here's what we did. We came up with a way of identifying bubbles and then determining whether we're in the run up or sell off phase of a bubble.

So this is very, very controversial, because there are some very famous people who say it's impossible to detect bubbles ... notably ... at the University of Chicago. And so, we came up with a proxy for crowded trading. So we don't actually observe flows, but what we do is we observe ... We applied this to sectors and factors.

We observed the centrality. We observed the centrality, and centrality is ... It's hard to describe. But let me just describe it sort of mathematically. So, first of all, there's this statistic called the absorption ratio. The way to think about the absorption ratio is if you do a principle components analysis on, say, the covariance matrix of sectors in, say, the US stock market, then what you do is you compute the fraction of total variability that's explained by the first couple of factors, the most important factors. That's the absorption ratio.

So if that ratio is high, if these few factors explain a large percentage of variability, what that means is that the market is very tightly coupled, and when it's in that state, it's very fragile, because shocks travel more quickly and more broadly. If the same few factors explain only a small percentage of total variability, what that means is that risk is distributed broadly across many disparate sources, and when that's the state of the world, the markets are much more resilient to shocks.

So that's the absorption ratio, and that's just step one in calculating centrality. So the next step is to say, "Well, let's look at the factors in the numerator of the absorption ratio," and when I use the term factors, I could just as well say principal components, or ... but they're literally just linear combinations of the sectors.

Adam Butler: [01:02:49](#)

Directions of risk.

Mark Kritzman: [01:02:51](#)

I will say, for example, if we want to know how central ... So centrality is really what's driving the variability of returns. So how central is, say, the technology sector, but we'll say, "What's the weight of the technology sector in the first factor?"

And then, of all those factors in the numerator, what fraction of the variability is being accounted for by that first factor, and we'll scale the weight of the technology sector by the relative importance of that factor, we'll go to the second factor, look for technology there, scale it, sum it up, that's centrality.

To make it intuitive, it's how Google ranks web pages, okay? So it's a measure of the extent to which a particular sector is driving the variability of returns, and what we observe about it is that a sector is more central to the extent that it's more volatile, and to the extent it's more connected to other sectors.

Now, why should centrality be associated with crowded trading? So, if investors are crowding into a particular sector, then there are going to be large order imbalances, and that's going to lead to large price adjustments, so that's going to create more volatility.

And also, if investors are crowding into a particular sector, they're not thinking independently about the components within that sector, they're thinking of the sector as a unit unto itself, and that means that the components within it are going to start moving more together, and that's going to make the sector less diversified, and hence, more volatile.

And then, to the extent investors are crowding into a particular sector, that sector is going to become more of a bellwether, and other sectors are going to follow it more and it's going to become more connected to the other sectors. So the features that we observe about centrality are consistent with what we would expect to observe from crowded trading.

And anyway, we then show that, centrality does correlate very nicely with the formation of bubbles. The problem-

Adam Butler: [01:05:14](#)

And you went back and examined bubbles across different asset classes at different times, Japanese ... bubble, the Nasdaq bubble, the housing bubble, that sort of thing?

Mark Kritzman: [01:05:21](#)

The way we approached this is to say, let's look at well-known bubbles, and let's look at how the prices of those sectors, or whatever groups, evolved, and let's look at how centrality evolved, and see. And what we found ... Let me just, before we go there, Adam, what we've discovered is that centrality does a really good job at identifying bubbles. Not just big bubbles, but little bubbles.

When I say bubble, I'm not necessarily referring to the dot-com bubble, whatever. What I'm referring to is any kind of price activity where there's a run up and a sell off that can't be explained by changes in fundamentals, okay?

Adam Butler: [01:06:13](#)

And it could apply to industries, or sectors, or factor portfolios, presumably?

Mark Kritzman: [01:06:19](#)

Yes. And we do it in the paper, we apply it to sectors and factors out of sample. In sample, we just looked at how the famous bubbles looked across these measures. So what we discovered is that centrality is very effective at locating bubbles, but it has a big, big problem. It can't tell the difference between the run up and the sell off. So what happens is, you see an increase in centrality as the bubble begins to inflate, and then it rises even more steeply when the bubble deflates, which to me, just suggests that the crowding on the way out is more intense than the crowding on the way in.

Adam Butler: [01:07:05](#)

Right.

Mark Kritzman: [01:07:05](#)

So what we did is, we combined it with relative value to distinguish the run up phase from the sell off phase. Now, what's interesting is, relative value by itself is not effective, because it cannot separate price activity that's legitimately driven by changes in fundamentals from price activity that's driven by fads or human behavior.

So it by itself is useless. Centrality by itself is useless. When we put the two together, they do a really, really good job of locating bubbles and segmenting the run up from the sell off.

Adam Butler: [01:07:55](#)

I can also see an application alongside momentum, where the measure of centrality might attenuate the signal from the momentum signal.

Mark Kritzman: [01:08:06](#)

Well, we get that question a lot, because, as you know, value hasn't worked for a long, long time, since the crisis. And what people are trying to do is combine it with other factors like momentum to see if the combination works better. And what we find is that, the measure of centrality and relative value is much more effective than value and momentum, and pretty independent to value and momentum.

So we think there's something different here than what you see with the traditional factor. So we did an experiment on the US sector, we got really strong results. And then, we took the same approach, precisely the same calibration, and applied it in five or six stock markets around the world, and very similar results. Except Australia, because there's really no dispersion across sectors. There's just a two or three-sector economy, and there's no opportunities set there.

But in Germany, the UK, Japan, Canada, results very, very similar to the US. And then we applied it to value, size, quality, and low vol.

Adam Butler: [01:09:28](#)

Using the decile portfolios?

Mark Kritzman: [01:09:29](#)

Yeah, and we got very, very similar results there, as well. So we're pretty excited about it.

Adam Butler: [01:09:35](#)

That sounds very exciting. It's one of the unsolved and most impactful problems in finance. So, definitely would look forward to investigating that.

Mark Kritzman: [01:09:48](#)

Yeah, I'll send you the paper if you're interested.

Adam Butler: [01:09:50](#)

Yeah, yeah, I'd really appreciate that. That sounds great. Okay, well, we covered a lot of ground, Mark. I want to really express my appreciation. I have been reading your stuff for well over a decade. We use your absorption ratio directly in the management of our portfolios, and think it's got lots of merit.

Really looking forward to investigating the turbulence ratio, as well, from the book, and the robust optimization methods that you propose. Lots of really great grist for the mill, and again, just thank you very much for your time today, and for sharing so much with us.

Mark Kritzman: [01:10:35](#)

Well, Adam, it's my pleasure. I really appreciate the opportunity, thank you.

Adam Butler: [01:10:40](#)

All right, have a great afternoon.

Mark Kritzman: [01:10:41](#)

You too.