

## Morningstar Conversations What Does Harry Markowitz Think?

Sam Savage and Paul Kaplan built a new approach to the pioneering economist's models. Now, they ask for his blessing. by Paul D.

Kaplan 06-15-10 | Comments 1

In the April/May 2010 issue of *Morningstar Advisor* magazine, Sam Savage of Stanford University and I wrote an article titled "Markowitz 2.0," which outlines our vision for the next generation of portfolio construction. Subsequently, I asked Harry Markowitz, the father of Modern Portfolio Theory, who was awarded the Nobel Prize in Economics for his pathbreaking work on portfolio construction tools, to join Savage and me for a discussion on the origins of MPT and where it is going in light of the sometimes erratic behavior of financial markets. Our conversation took place on April 14 and has been edited for clarity and length.

**Paul Kaplan:** Harry, while much has been written about how you developed the mean-variance efficient frontier, our readers would enjoy hearing the story directly from you. Please briefly recall for us how you first developed the mean-variance model.

**Harry Markowitz:** The magic moment--the moment of epiphany--happened while I was reading John Burr Williams' *Theory of Investment Value*. I was looking into the possibility of doing a Ph.D. dissertation at the University of Chicago, applying mathematical, statistical, or econometric techniques to stock market investment problems.

I was working off of a reading list a professor of finance had supplied. I'd already read Graham and Dodd's *Security Analysis*. I read Wiesenberger's *Investment Companies and Their Portfolios*, and I was reading *Theory of Investment Value*. Williams asserted that the value of a stock should be the present value--the discounted value of future dividends. Because future dividends are not certain, he said you should use the expected value of future dividends.

Now, I knew that if you were maximizing an expected value, the way you would do that was to put all your money into just one stock. That didn't make sense. People do diversify. You could see it in Wiesenberger's *Investment Companies and Their Portfolios*. They diversified because they were worried about risk as well as seeking return.

So, I postulated that investors were interested in expected return and standard deviation. I drew a trade-off curve, like economists always do, and so that afternoon in the business school library at the University of Chicago, I came up with the first efficient frontier.

Williams asserted that with sufficient diversification, risk would disappear; you would receive the expected value. But risk only disappears with diversification if you have uncorrelated risk and, of course, markets are not uncorrelated. So, that was when and how the moment of epiphany happened.

### De Finetti's Scoop

**Kaplan:** In 2006, the *Journal of Investment Management* published the first English translation of a paper by the Italian mathematician Bruno de Finetti on mean-variance optimization. The Italian version of the paper was published in 1940, 12 years before your famous paper. In the same issue of the *Journal of Investment Management*, you published an introduction to the de Finetti paper, which you graciously titled, "De Finetti Scoops Markowitz." In your view, what was the historical significance of de Finetti's work?

**Markowitz:** The historical significance was nil because the importance of de Finetti's work was not understood. If it had any impact at all, it was among Italian actuaries. It was a response to the problem of optimum reinsurance, so it was an actuary thing. I do not believe that it ever got into the American actuary literature, and it certainly didn't get into our financial literature. So, it was a dead end, not because it deserved to be a dead end, but that was, in fact, its historical destiny.

In terms of the merits of de Finetti's contributions, it turns out that he correctly posed the problem of mean variance. He understood that you had to take into account the correlations. Williams didn't understand that something happened differently when you diversify among correlated risk. De Finetti did. He was able to solve the problem of tracing out mean-variance efficient frontiers assuming uncorrelated risk. He would've liked to have solved it using correlated risk, but he couldn't solve that one.

As I explain in my paper, de Finetti had some conjectures about what the solution must be, and at least one of the conjectures was wrong. So he gets a gold star for posing the problem, and I get a gold star for solving it.

**Kaplan:** Sam, I believe you have some personal recollection of de Finetti.

**Sam Savage:** I do. In the 1950-51 academic year, when I was six, my dad (Leonard J. Savage) was on sabbatical in Paris working on his book, *The Foundations of Statistics*. It was at that time that he discovered the random walk thesis of Bachelier and passed this information on to Paul Samuelson, which started the whole "Random Walk Down Wall Street" craze.

I don't know why he knew of de Finetti, but he sure did. And at one point, we visited Italy. I'll never forget Venice in 1951. We met de Finetti. I remember him well from then. I have to say, Paris and Venice are very exciting places for a 6-year-old.

Later, in 1958, we spent a year in Rome, and my dad visited the Università di Roma and worked even more closely with de Finetti. So I knew him personally. I knew his daughter, Fulvia. I was not following the probabilistic arguments at all in those days.

**Kaplan:** Harry, much of the criticism leveled at the mean-variance approach is the allegation that it assumes that returns follow a normal distribution. In the late 1970s, you did some work with professor Haim Levy at the Hebrew University of Jerusalem that shows that the mean-variance optimization does not rely on this assumption. Please explain your work with Levy and how it answers this criticism.

**Markowitz:** The work that I did with Levy that was published in 1979 in *The American Economic Review*, called "Approximating Expected Utility by a Function of Mean and Variance," is an amplification of ideas I already published in my 1959 book [Portfolio Selection: Efficient Diversification of Investments]. I never ever assumed that probability distributions were normal. I never justified mean-variance analysis in terms of probability distributions being normal.

My basic assumption is that you act under uncertainty to maximize expected utility. In fact, I was a student of Sam's father, and he convinced me, as well as half the rest of the world, that the way to act under uncertainty is to maximize expected utility using probability beliefs where there weren't any objective probabilities.

So, how do I get off peddling mean-variance analysis when I believe in expected utility? An illustration of what I believe is given in chapter six of my 1959 book, where I recommend maximizing expected logarithm of one-plus-return if you're in for the long run. I have a table showing various levels of return, the log utility of that return, and the value of a simple quadratic approximation.

Between a 30% or 40% loss and a 40% or 50% gain on the portfolio as a whole, the quadratic approximation is very close to the log utility. So, the expected value of the one has to be close to the expected value of the other, as long as we're talking about portfolios that rarely lose much more than 30% or 40% or gain much more than 40% or 50%. Of course, the expected value of the quadratic is a function of mean and variance.

So, for the nine or 10 securities that I had available for my illustrative mean-variance analysis, I show that not only does the approximation look good in the table, but, in fact, if you knew the mean and variance, you could guess the expected utility quite well.

What Levy and I did was to do that same experiment using historical returns on investment companies and a variety of utility functions. We confirmed that if your probability distribution is not too spread out, and the annual returns on investment companies turn out to be not too spread out, then if you know mean and variance, you can guess expected utility quite well.

**'It's Jimmie Savage's Son!'**

**Kaplan:** Sam, as Harry just mentioned, he was a student of your father. Is that how you developed an interest in portfolio theory?

**Savage:** Not at all, that's just a fluke. I had heard of Harry and the Markowitz optimization model but had not paid much attention to it. I was not in that line of work.

But in 1989, I got a phone call from a guy with a Texas drawl thick enough to cut with a knife. He asked me an interesting question. He said that when oil companies choose portfolios of exploration sites, they sort these sites from best to worst on something like expected net present value. Then, they start at the top and go down the list until they run out of money for investing. This guy said, "Why aren't they using Markowitz portfolio theory to do this?"

This was my introduction to a fellow named Ben Ball, who was an adjunct professor at MIT at the time, and a former VP of planning at Gulf Oil. My response to Ben was that this was indeed a good question.

But I wondered how the heck you would actually calculate the covariance matrix for a bunch of oil prospects. I knew enough about the Markowitz model to know that it took covariances. It was several years before the answer appeared for how you would, in fact, do portfolio analysis with these things.

Ultimately, I applied this approach with some success at Royal Dutch/Shell. The approach is known, I believe in the profession, as scenario optimization. Certainly, I was not the person who discovered scenario optimization, but it shaped my career. The notion is, you don't describe a probability distribution parametrically, like with a mean and standard deviation. You describe it as a bunch of realizations, perhaps generated by Monte Carlo.

**Kaplan:** How did you meet Harry?

**Savage:** I met him at George Dantzig's 80th birthday party. Dantzig was a famous mathematician at Stanford who invented linear programming. I remember I was introduced to Harry and said, "Oh my God, it's Harry Markowitz!" As I recall, Harry said something like, "Oh my God, it's Jimmie Savage's son!" He told me he had been indoctrinated at point-blank range in expected utility theory by my dad.



**Markowitz:** I have in my cabinet several copies of the *Foundations of Statistics*, which I give out like Gideon Bibles to people that I'm trying to convert to Bayesianism.

**Kaplan:** And just to complete the circle, Sam, when you were teaching at the University of Chicago, one of your students was Joe Mansueto, the founder of Morningstar. Harry, in the early 1960s, the French mathematician Benoit B. Mandelbrot and one his doctoral students at the University of Chicago, Eugene Fama, were publishing papers that claimed that changes in security prices were so erratic that they were best modeled using a fat-tail distribution, known technically as stable Paretian distributions. Under this model, variance is not a meaningful concept. So, even your work with Levy cannot justify a mean-variance approach.

Largely ignored for decades, there has been renewed interest in Mandelbrot's line of research, especially after the global market meltdown of 2008. As readers of this magazine know, I have been working with these fat-tail distributions myself for the past two years. What is your view of this research?

**Markowitz:** When I was at Baruch College, they gave me a research assistant, Nilufer Usmen. Because I'm a theoretician, I wondered, what am I going to do with a research assistant? What I decided to do was experiment with Bayesian inference.

What we did was to look at how much Bayesians should shift their belief among various hypotheses concerning the probability distribution of the log of the S&P 500 today divided by the S&P 500 yesterday. What we used was a very broad class of probability distributions called the Pearson family. We found that Bayesians should shift their beliefs massively against normal or Gaussian distributions in favor of a Student's t-distribution with between four and five degrees of freedom. (Harry M. Markowitz and Nilufer Usmen, "The Likelihood of Various Stock Market Return Distributions, Part 2: Empirical Results," *Journal of Risk and Uncertainty*, November 1996.)

A Student's t-distribution has fat tails, but as long as it has more than two degrees of freedom, variances exist. Therefore, the central limit theorem works. If you average together 250 trading days' worth of daily returns, the distribution will get a lot closer to normal. However, the Pearson family does not include stable Paretian distributions.

The interesting thing about stable Paretian distributions is that they are either normal, Gaussian, distributions or distributions with infinite variance. So if you assume stable, then you either have to accept normality, which is no way true, or you have to assume infinite variance, in which case you cannot do mean-variance analysis.

Subsequent to doing this research with me, Nilufer Usmen married Tony Tessitore, who was my Ph.D. student. Tony is now a client of mine. Tony took it about himself to go back and bring the Markowitz-Usmen research, on shifting beliefs among different probability distributions, up to date. While he was at it, he found out how you should shift your belief between the stable Paretian distribution and the Student's t-distribution. We concluded that you should massively shifted against the stable Paretian distribution in favor of the Student's t-distribution.

So, I'm a Bayesian, and I am not worried about Mandelbrot's stable Paretian distributions. For daily returns, Student's t-distribution with four-and-a-half degrees of freedom, plus or minus a half degree, is a good model of stock returns. When you aggregate them up to lower frequencies, things start looking normal. So, I go back to the Levy-Markowitz paper: As long as your probability distributions aren't too spread out, mean variance is a good approximation to the expected utility.

**Kaplan:** Sam, you mentioned that when you started working with the oil industry and doing optimization models to form diversified portfolios, you used the scenario approach. Explain the scenario approach and how that squares with this whole debate about stable distributions versus t-distributions versus normal distributions.

**Savage:** The beauty of the scenario approach is that it's completely agnostic. Suppose I'm trying to analyze the outcomes of the roll of a die. If you study the physics, you'll say that you can have any of the faces show up with equal probability. But instead of that, you roll the die a thousand times and store those thousand die rolls. Those would be scenarios.

Now, let's do an example that you couldn't easily solve with theory. Suppose I had two dice that were statistically related by having magnets imbedded in them. I could roll the dice together and the number that appeared on one die would, in fact, have implications for the number on the other die. I could store these numbers in two columns of thousands of rows in Excel. The scenarios in the two columns capture the joint distribution of the dice.

In his 1959 book, Harry does something similar in a beautiful way. He has spinners with concentric rings of numbers to model joint distributions of statistically dependent things. That is, there is an inner ring of numbers representing one uncertain asset and an outer ring of numbers representing a statistically related asset. That actually is very much in the spirit of the scenarios of dice I just described. I must point out that when I first looked at this book, I expected it to be entirely theoretical. To my delight, I discovered that much of it could be understood by a bright high school student. I highly recommend it.

But back to the scenarios. Recently, we've developed a technology which takes thousands of scenarios and packs them into a single data element. I call this the Distribution String, or the DIST for short.

As to this question about Mandelbrot versus the Student t-distribution with four degrees of freedom, if probability distributions were beer, then the Distribution String is a beer bottle. You can put anything in it. But let's not forget you can also put poison in beer bottles, and there are some distributions that are so fat-tailed that even if you had a million scenarios, it would not be enough for dependable results. Essentially, the DIST is a pre-computed Monte Carlo simulation, and if your distribution is well behaved enough to be simulated in the first place, then the distribution string is completely agnostic.

**Kaplan:** So, if Harry were using it, he would put in the output of a Monte Carlo simulation of a Student's t-distribution with four degrees of freedom.

**Savage:** Not only that, but I would be happy to promote Markowitz in a bottle in the free market--that is to say, people should be free to choose which distributions to use in a particular setting. If these things were widely available on the Internet, the Markowitz brand would be a hot seller.

#### **Asset-Allocation Infrastructure**

**Kaplan:** Harry, I want to go back to you. Sam and I called our article "Markowitz 2.0." In that article, we said that what you did for the quantitative portfolio construction in 1952 was what the Wright brothers did for aviation. You made the key breakthrough that made it all possible.

But while the aviation industry has made enormous advances since the Wright brothers, the investment profession has largely done nothing beyond your 1952 work. Yet you, yourself, have been advocating major enhancements to your original model, beginning in your 1952 book. But your enhancements don't seem to have been embraced by our industry. Why do you think the industry has been so slow in making changes to your basic model?

**Markowitz:** Let me point out that there are things that the aviation industry does just as the Wright brothers did, and there are other things that they have advanced considerably beyond the Wright brothers.

In our industry--the portfolio theory industry in the broadest sense--there are things that are just exactly like 1959. My thinking through of why to use mean-variance analysis, how to use mean variance, and so on, was not in the 1952 paper. That proposed mean variance; the rationale for it was presented in my 1959 book.

So, let's talk for a moment about what's new and what's old in the aviation industry. We still fly basically like the Wright brothers flew. We have a means to pull the airplane through the air, and then the air flowing over the wings lifts the plane. You've got thrust and you've got the lift. We still don't know how to flap our wings and fly around trees in a dense forest and land lightly on a branch. That's flying! We do not know how to dive screaming down from the heavens, almost vertically straight down, pull out in time to grab a mouse in our talons, and fly off with it. That's flying! But we do know how to put a lot of people in a big jet and serve them coffee, tea, or milk. That's what's old and new in aviation.

What is still true in 2010, as true as it was in 1959, is the table on Page 121 of my book showing that if probability distributions are not too spread out, then mean-variance approximation will work very well.

What is new is data, for example, like Morningstar data, formerly known as Ibbotson data. That has been very valuable to us. Another thing that's new are the models of covariance, like Barr Rosenberg's, and experimenting with different models of expected return, like the Fama-French Three Factor Model.

I think the most important thing that happened between 1959 and the present is the notion of doing your analysis on asset classes in the first instance. This has become part of the infrastructure that we now rely on. In 1959, I had a theory. I had a rationale, and so on. Now, we have an industry.

#### **Applying Scenarios to Portfolios**

**Kaplan:** Sam mentioned that his introduction to the portfolio problem was not in putting together a portfolio of securities, but rather putting together a portfolio of oil exploration projects. In that context, he turned to the scenario approach, which, as we discussed, is flexible--you don't have to make any particular assumption about the shape of the distribution. You can assume normal. You can assume Student's t. You can assume many, many distributions, and you don't have to assume that the covariance matrix is the best way to characterize the relationships between these securities.

Now that we have the computational power, what is your view of applying this scenario approach to the portfolio-construction problem that you first worked on in 1952?

**Markowitz:** I have no objection to what you and Sam do with oil and gas. I might very well do the very same thing. But when it comes to whether David Swensen, in running the Yale endowment, ought to be using scenario analysis instead of doing a mean-variance analysis at the asset-class level and then implementing it himself--I think David is doing just fine.

**Kaplan:** Well, Harry, let me give some more context to this. In your 1959 book, you mentioned semi-variance as an alternative risk measure. Later, in a paper you cowrote in 1993, you solved the mean-semivariance optimization problem--using the scenario approach.<sup>2</sup> You made a list of scenarios, and under each scenario, there's a return for each security. As I recall, you set it up as a conventional quadratic optimization problem, except it had lots and lots more variables than the typical mean-variance problem because there had to be like at least one additional variable for each scenario. Do you recall the paper I'm referring to?

**Markowitz:** Yes. It was a paper about how to compute mean semi-deviation frontiers using a mean-variance optimizer. It's a much slicker algorithm for solving that problem than what I published in 1959.

I tell my class that one of the problems with mean semi-variance is that you can't just take a matrix full of semi-co-variances and trace out a frontier. There are semi-co-variances, but they're local to the portfolio and its neighboring portfolios that have the same pattern of "profitable years" versus "non-profitable years."

So, you have to use finite samples. You can use historical returns on securities, or if you have some kind of model of covariance, you have to sample synthetic years and use these in the algorithm. You're perfectly right that there you have to Monte Carlo, and then you optimize.

**Savage:** What Harry just described is almost exactly what we did at Shell, except that we store thousands of finite samples in beer bottles.

**Markowitz:** Yes, you have a tool kit full of techniques--optimization, Monte Carlo, present value, linear programming, quadratic programming, and so on. A person who only has one tool in his kit is a danger. You know the old saying, "If all you have is a hammer, everything looks like a nail."

So, I agree. We have many techniques, and we've have to decide where are they applicable.

**Savage:** I have a comment to put this in another perspective.

The thing about the mean-variance model is that it is a magnificent, simple theory, like Isaac Newton's F = ma. Oh, it doesn't really equal ma because of Einstein's theory of relativity. But do you know what? F = ma is a foundation on which we have built modern physics because it is a simple, beautiful model that is good in most real world situations. Another advantage of theory over computation is that it provides its own sorts of insights that foster further scientific thought. The mean-variance model did this in spades. It is a beautiful, simple theory.

**Markowitz:** Let's take the universal theory of gravity of Sir Isaac Newton versus Einstein's theory of gravity. Under almost all circumstances that you and I encounter, like what is going to happen if my left foot trips over my right foot, I probably don't have to use Einstein's theory to figure out the result.

So under certain conditions, and those are very broad, the Newton theory works wonderfully. Similarly, under certain conditions, mean-variance works perfectly well. What are those conditions? Well, they have to do with the probability distribution returns not being spread out too much, like Levy-Markowitz says.

**Kaplan:** This scenario-based approach that Sam and I have been working on we like to call Markowitz 2.0. Are you okay with that?

**Markowitz:** I'm fine with that. How can I refuse anything to a man who has named an important contribution the "Markowitztron."

**Savage:** I'm delighted, Harry. I started looking at this stuff in the oil patch, where the distributions were incredibly far from normal--that is to say, there's an 85% chance of getting zero and some tiny chance of getting big payoffs. And yet, using the scenario approach led to creating these distribution strings. Then, I met Harry Markowitz along the way, and that was a thrill. I've learned a tremendous amount from Harry over the years. But it's a sweet irony that these distribution strings now are beginning to be used in finance, at least by you folks at Morningstar. It closes the loop for me, which is personally gratifying.

**Kaplan:** Harry, you have the last word.

**Markowitz:** I'll try not to make it inflammatory. I do believe in scenarios and Monte Carlo. I like the way Sam packages them and makes them convenient. For situations where mean variance is applicable, I still would trace out mean-variance frontiers, but in deciding which point on the frontier one should recommend, I would use scenario or Monte Carlo for that.

Let me characterize Sam in my mind. There is a conversation we had that I think is prototypical Sam Savage. Sam has these wonderful ways of displaying things, of making things accessible to Excel users. He has this wonderful sense of humor. I think of him as sort of a graphical artist and the salesman of our field.

One time, he told me what his Ph.D. dissertation was. I said, "That's incredible. You could be proving great theorems like your father. Why aren't you?" He replied, "Because I'm making more money and having more fun."

*Paul D. Kaplan. Ph.D., CFA, is director of quantitative research with Morningstar Europe*