THE JOURNAL OF
INVESTMENT CONSULTING

A reprinted article from Volume 20, Number 1, 2020

THE MASTERS SERIES

E. Robert Fernholz, PhD: Stochastic Portfolio Theory

INVESTMENTS & WEALTH INSTITUTE®

E. ROBERT FERNHOLZ, PHD

Stochastic Portfolio Theory

Robert Fernholz is a mathematician and financial researcher who specializes in the mathematics of finance. He founded the institutional equity management firm Intech in 1987 and served as its chief investment officer. Now retired from Intech, he is president of Allocation Strategies, LLC, a company he founded in 2012. Before joining the investment industry, he held academic posts at the University of Washington in Seattle, Hunter College of the City University of New York, and Princeton University. But his interests changed from pure to applied mathematics, and in 1982 he published the article "Stochastic Portfolio Theory and Stochastic Market Equilibrium," which formed the basis for his investment strategies at Intech.

E. Robert Fernholz, PhD

Fernholz is best known for his pioneering research monograph Stochastic Portfolio Theory, *which was published in 2002. Stochastic portfolio theory is a mathematical framework for constructing investment portfolios, analyzing their behavior, and understanding the structure of equity markets. Significant for both academicians and practitioners, the theory is central to the field of modern mathematical finance and is particularly useful to managers of stock portfolios.*

Fernholz earned a PhD in mathematics from Columbia University and a BA, magna cum laude, in mathematics from Princeton University. He is a trustee of the Institute for Advanced Study in Princeton, New Jersey.

In June 2020, Robert Fernholz spoke with members of the Journal of Investment Consulting*'s editorial advisory board about stochastic portfolio theory, its contributions to the investment industry's knowledge of equity markets, and its applicability to creating and monitoring investment portfolios. Taking part in the discussion were Inna Okounkova, Columbia University and editor-in-chief of the* Journal; *Edward Baker, Mesirow Financial; Ludwig Chincarini, University of San Francisco and United States Commodity Funds; Philip Fazio, Merrill Lynch; and Geoffrey Gerber, TWIN Capital Management.*

Inna Okounkova: Let's start with some introductory questions. What were the major factors that helped shape your career? What do you regard as your major achievements? And on the flip side, what was your greatest disappointment?

Robert Fernholz: I was born in 1941 and grew up in Princeton, New Jersey. Princeton is not a standard town—the university and the Institute for Advanced Study give it an interesting atmosphere—and among the kids I grew up with, particularly a good friend, George Akerlof,[1] the idea was for you to think outside the box and try to come up with something new.

Our teachers weren't always happy with this attitude, and that, actually, established an important direction in my career. Eventually, I went to Princeton University, which may seem odd when you've grown up in the town, but it's certainly not a bad university, and they gave me a good scholarship. I ended up studying probability with William Feller, and that was a wonderful opportunity.[2] He was really a very inspiring guy.

I later studied math at Columbia University for a PhD, but really the next major influence on my career occurred when I had the opportunity to work with Harry Markowitz[3] at a hedge fund he had helped to set up at Arbitrage Management Company.[4]

In 1979 I replaced Harry as the research director of the hedge fund, and he had arranged things so that the research director had the option of taking one day a week off for their own research or consulting. This was a spectacular advantage, because normally, if you work for a brokerage, the company owns any idea you have, but Harry had set up this fund in this particularly benign fashion. Markowitz continued as a consultant, but his day job at that time was with IBM developing his simulation language, simscript.

I worked at Arbitrage Management Company for a number of years, and during that period, on my research days, I developed the ideas for stochastic portfolio theory (SPT). So, perhaps the three factors that most influenced my career were my childhood in Princeton, inspiration from Feller, and my work with Markowitz.

As for achievements, I would have to consider the development of SPT as my major achievement. It encompasses a mathematical theory as well as practical applications to investment technology. The Intech strategy was a fairly straightforward application of what Brian Shay and I described in our 1982 paper. Following the publication of my book in 2002, many other applications arose and continue to be discovered, both in theory and in practice.

I can't say I've had many disappointments. Perhaps my transition back to applied math was too slow after my involvement in pure math. But I don't really have that much to complain about.

Inna Okounkova: What triggered your switch from pure to applied mathematics of finance and your switch from academia to industry?

Robert Fernholz: I was interested in probability. When I went to Columbia, I studied pure math. It was beautiful stuff—algebraic geometry and several complex variables. My PhD thesis was about the structure of complex analytical spaces, but I was actually more interested in probability and its applications, so transitioning back to applied math was quite natural. I find that the real power of math shows up in its applications.

I went into industry instead of academia because I found that industry was where the most interesting things were going on. In industry I had more liberty to think outside the box. At the Arbitrage Management Company, I did data-adaptive time-series analysis applied to foreign-exchange trading, and I also did a lot of work on options pricing that was ahead of what was available in the academic literature. Options pricing, with real options in real time, required an understanding of different interest rates for borrowing and lending, non-constant variances, dividends, and other things that hadn't appeared in the academic literature. The Black-Scholes model was the general basis for option pricing,[5] but the application to real markets was more exciting than the research that was happening in the universities at that time.

My time at Arbitrage Management gave me time to pursue, and publish, my own research. The paper on SPT that was eventually published in 1982 was at first summarily rejected by the *Journal of Finance*. Harry Markowitz was unhappy with that because he thought it was an interesting paper. When he became president of the American Financial Association, I think his first job was to set up the organization's annual meeting, and he told me, "Set up a session on portfolio theory, include your paper, and the *Journal of Finance* will publish all the papers from that session." I don't know whether our paper would have ever seen the light of day if Harry hadn't intervened.

In a peculiar sense, SPT evolved from option pricing. In the 1970s and 1980s, holders of institutional stock portfolios would sometimes write call options against the stocks in their portfolios. The strategy was called "option overwriting." It was attractive in that period because the interest rates were very high, 15 to 18 percent, and option overwriting had the effect of creating a portfolio equivalent to a mixture of cash and stock. With interest rates around 15 to 18 percent, stocks and cash had essentially the same compound rate of growth, and simulations that people had run showed that these mixed portfolios would actually grow faster than either of those asset classes individually over the long term.

So, I asked myself, "How is this happening?" I started using stochastic calculus to analyze the performance of these portfolios, and lo and behold, it turned out that a portfolio could indeed have a higher growth rate than the average growth rate of the component stocks, which I thought was sort of astounding. The extra performance was exactly the excess growth rate of the portfolio. SPT began with the discovery of the excess growth rate.

Inna Okounkova: Please tell us more about Intech and its history.

Robert Fernholz: Intech is a company that offers institutional investment management services. We started it in June 1987, which was a particularly propitious time to start—ironically speaking—because the stock market suddenly crashed in October of that year, so it gave us an excellent opportunity for out-of-sample testing. I put SPT into practice around the first of June, and we were running a portfolio on July 1. The crash was an interesting phenomenon. Stock prices were going all over the place, futures were moving opposite to the index, and we couldn't tell what was happening. We didn't—we couldn't—trade. But all in all, our returns were up a couple percentage points one day, then down a couple the next day, and by the end of the week, things had evened out, and we were [randomly] ahead a by few basis points.

Until 2002, Intech was owned by the Prudential Insurance Company. At first, Prudential was selling guarantees that would pay the S&P 500 return plus five basis points, and they backed the guarantees with our strategy. Our strategy was tuned to yield about 2 percent a year above the S&P 500 return, and that's about what it did. Prudential incurred some risks, but they made quite a bit of money with their guarantees over a number of years. However, after a few years of guarantees, Prudential wanted to go public, and technicalities regarding the statutory surplus for insurance companies caused them to move away from equities and into bonds.

After the guarantees, the Intech strategy was sold as a standard institutional investment product, but Prudential really didn't have much of a dedicated institutional sales force. In 1991, Bob Garvy[6] came on as the chief executive officer to build Intech's institutional marketing and distribution capacity, which subsequently grew significantly. However, Intech had been set up as a separate subsidiary, and in 2002 Prudential sold Intech to Stillwell/Janus/Berger Financial. Janus soon took over the whole Stillwell/Janus/Berger conglomerate, and at that point, sales went up sharply because they were much more focused on selling our products than Prudential had been. Now Intech is a subsidiary of Janus and offers a number of different products, all based on variations of the SPT strategy.

Inna Okounkova: Please tell us more about how the strategy works. What is the ideal holding period? How long and deep are drawdowns relative to the market? What is the typical turnover?

Robert Fernholz: It's a long-term strategy. It's for investments that institutional investors will hold onto for a long time. The information ratio for the whole class of Intech strategies probably averages about 0.8. The information ratio in this case is measured by log return of the portfolio above the log return of the benchmark, usually an index, divided by the standard deviation of those log returns [tracking error].

You can't push log return much above about 3.5 percent [per year] over the benchmark. You could push it below 1.5 percent, I suppose, if you just wanted to mix it with the benchmark index. Another strategy we developed had to do with functionally generated portfolios. I discovered the concept of functionally generated portfolios in the mid-1990s.

In the mid-1990s there was a huge concentration of investment into the big stocks. At one point during that period, we were performing badly in some portfolios, and we were trying to determine if our bad performance was caused by that concentration. Moreover, our measurements of the effect of this concentration were coming out differently from the canonical measurements of that period. So, we began studying various types of diversity and concentration measurements, including entropy, which is well known as a natural measurement of diversity.

I started looking at market entropy and trying to get a handle on the stochastic behavior of this measure. I found that if you consider the log entropy of the market, this function actually will generate a portfolio, which I called the entropy-weighted portfolio. Astoundingly, I found that the log return on the entropy-weighted portfolio relative to the market portfolio could be decomposed mathematically into the change in the log-entropy function plus a non-stochastic drift term, a drift term that always increased. This meant that over any period in which the market entropy did not decrease, the entropy-weighted portfolio would, with certainty, outperform the market. I generalized this idea to functionally generated portfolios and got a patent on their application in finance. Intech later offered a diversity-weighted index based on this patent. The diversity-weighted index product was offered for a number of years, and it performed exactly according to theory. The diversity-weighted index could generate about 70 basis points above the S&P 500. But Intech had other portfolios that were performing brilliantly, and the marketers sold those other portfolios because they could make more money on them as opposed to an index product. I suppose you could call the failure to fully develop the diversity-weighted index a disappointment, but from a business perspective, the other products probably generated more revenue for the company.

Prudential invested its own pension fund in the diversity-weighted index for a number of years, and although it worked exactly as it was supposed to work, we never managed more than a few hundred million dollars with it. It did what it was supposed to do, but it wasn't glamorous enough to generate widespread appeal.

Geoffrey Gerber: More-diversified portfolios have higher excess growth rates and logarithmic returns, but from a practical perspective, what is the optimal amount of diversification? Would you want to hold all 500 S&P stocks in different weights? Is that more optimal than holding 400 of the 500?

Robert Fernholz: It depends what you mean by optimal. The active strategies, which we structured to get a couple percentage points ahead of the S&P 500, did not hold all 500 stocks in the index; I think they usually held about 300. But the term optimal is not well defined. The diversity-weighted portfolio, by the way, held everything; it just weighed the smaller stocks a bit more. So, what's optimal in a particular situation depends on which type of strategy is used.

The other consideration is that there is a huge estimation problem in figuring out the covariance matrix for a stock market. Estimation of variances and covariances is quite complicated, so we were actually trying to estimate the tracking error rather than trying to estimate the entire covariance matrix. You can't just carry out MVO [mean-variance optimization] and expect it to work very well if you don't know what the parameters are. The covariance parameters may be difficult to estimate, but the expected return parameters are almost impossible to estimate. This is why MVO has generally failed in applications to large stock portfolios.

It turns out that if you use some type of active portfolio management, especially something like the diversity-weighted index that holds all the market, you get a phenomenon called leakage, or stocks dropping out of the investment universe. Leakage occurs with any large stock portfolio. Real stock markets are open in the sense that stocks can enter and leave the market. The CAPM [capital asset pricing model] and most of the rest of financial theory is based on closed markets, where the universe of stocks stays the same. To deal with the stocks that enter and leave the market, I introduced the use of semimartingale local times, which measure how long the amount of time a random process spends near zero. These local times can be used to measure the effect of stocks entering and dropping out of the investment universe. With the local times came the concept of rank, so I began to study portfolios that depended on the rank, by capitalization, of the stocks they held.

This brings us to another consideration, which is connected to the article by Banner et al. (2019) that was published in the *Journal of Investment Consulting*. The work reported in that paper was based on rank-based analysis, which introduces a new way to look at portfolios. If you try to estimate the expected growth rates of stocks by their rank rather than by their name, you'll have a much greater level of success because the growth rates of stocks are fairly constant by rank but vary widely by name. In the standard Markowitz–Sharpe type of optimization, you have to estimate all the parameters by name, and these parameters are really difficult to calculate—for example, who can accurately estimate the future alphas of stocks?

If you look at growth rates and variances by rank, everything suddenly becomes simple. The growth rates of the top 5,000 stocks are all almost the same, and the stock variances go up almost linearly with rank. This was the basis for the paper by Banner et al. (2019). It's interesting that this paper was rejected by the *Journal of Portfolio Management,* the journal that published the original "surprising alpha" paper that we were critiquing (Arnott et al. 2013). Our paper eventually was published in the *Journal of Investment Consulting,* evidently a more enlightened publication. The difficulty seemed to be that we were in competition with the original journal's referee, who had an explanation for surprising alpha based on a new "factor," as they call epicycles these days. But factors cannot measure excess growth, because factors come from regression analysis, but regression is a quadratic operation, and excess growth has no quadratic variation. Excess growth is invisible to regression-based analysis.

Inna Okounkova: What about the other way around? Can the performance of any other portfolio relative to market performance be explained by measuring or estimating its excess growth rate? You did that for the portfolio of small-cap stocks described in that paper, but what about value or profitability? Can any portfolio's excess growth rate be measured, and can its outperformance relative to the market be explained by the difference in excess growth rate?

Robert Fernholz: My book includes one section on value portfolios, and they seem to be highly correlated with small-stock portfolios. The decomposition described in the *JIC* paper probably should be applied to all portfolios. Investment managers might benefit by considering the average-growth/excess-growth decomposition we propose in the paper. If managers can get higher average growth rates with their stocks, that should show up. They should be able to demonstrate that they're doing better than average market performance. This would mean that they're good stock pickers.

On the other hand, if they're manipulating the variances, that also would show up. This type of analysis could be applied to any portfolio. A manager who is really picking stocks should be able to get a better average growth rate than the market itself, and a manager who takes advantage of the variances also should outperform. The portfolios used as examples in that paper were all naïve portfolios, so they didn't have higher average growth rates, and all of the outperformance came from the variances.

A manager who is really picking stocks should be able to get a better average growth rate than the market itself, and a manager who takes advantage of the variances also should outperform.

Philip Fazio: What were the major developments in SPT, and how did the theory translate into investment practice?

Robert Fernholz: Everything of course started with Harry Markowitz, with mean-variance optimization. The next major advance was with Bill Sharpe and his single-period equilibrium model, CAPM. Then Robert Merton generalized CAPM to continuous-time models using stochastic calculus. However, Merton's model had constant mean and variance parameters, and Barr Rosenberg and James Ohlson (1976) showed that led to an internal inconsistency, a fundamental flaw in the continuous-time CAPM (the Rosenberg–Ohlson paradox).

Here's what the inconsistency was. If you have constant parameters and your portfolio weights are a function of these constant parameters because you're doing MVO, then you get constant portfolio weights. Mathematical functions of constants are constant; this is precalculus. Since the only thing that changes is a mixture of cash and stocks, as in Sharpe's CAPM, the relative weights of the stocks in the market must remain constant.

In part, I developed stochastic portfolio theory to address this internal inconsistency at the center of continuous-time finance—it's interesting, very few people seem to be aware of this inconsistency. But after the 1982 paper, this aspect of SPT was not studied further, as far as I know. Far more important seemed to be the development of a descriptive theory of stock markets and portfolios based on continuous-time stochastic processes.

The theory really started with publication of my monograph *Stochastic Portfolio Theory* in 2002. I didn't publish anything for a while because Intech was managing money and we didn't really want to attract attention, but things changed when I got the patent on diversity-weighted indexing. When you have a patent, you want to go out and tell the world, so that's when we started publishing academic papers.

SPT is composed of two branches. One has to do with functionally generated portfolios and arbitrage. These portfolios demonstrate that in a market that looks like a real market, you actually can have arbitrage. You can beat the market over a certain period of time with a probability of one—with zero risk. The condition you need, aside from the standard conditions assumed in academic finance, is that the market does not spend all its time concentrated in a single stock.

If the average weight of the biggest stock over the whole lifetime of the market is less than, say, 99.999 percent, then that market has arbitrage. When I first tried to get that idea published, it was rejected. The 99.999-percent assumption seemed so innocuous that the referee didn't notice it and announced that I must have found a "paradox." When I explained that the 99.999-percent assumption was actually an additional condition on the market, the referee rejected the paper claiming that this was "a very strong assumption." So, the first publication of arbitrage in a stock market was in my book, where the publisher, Springer, apparently had a more perspicacious worldview.

I believe we have finally captured exactly what's going on with stock-market arbitrage in the latest paper with Yannis Karatzas and Johannes Ruf (Fernholz et al. 2018). After this paper, it appears that in mathematical finance, the condition of no-arbitrage is being replaced by no-instantaneous-arbitrage. This paper uses a portfolio with three dimensions as an example; it's a toy portfolio. If you can construct this type of portfolio, and if you hold it for any period of time, you'll be ahead of the market.

The nonexistence of instantaneous arbitrage is equivalent to the possibility of MVO. I believe most everybody thinks that MVO is reasonable and therefore instantaneous arbitrage can be outlawed, but the absence of longer-term arbitrage may be an unreasonable assumption. Yannis and his students are developing the theory.

A number of researchers in mathematical finance are looking at situations where long-term arbitrage exists. Even with the weaker assumption of no-instantaneous-arbitrage, you can still get option pricing. You can get everything you need, but you can no longer use the no-arbitrage hypothesis.

Arbitrage, which depends on functionally generated portfolios, is one branch of the SPT. The other branch splits off from leakage and rank-based analysis. Rank-based analysis has a number of applications, for example, in economics. My son Ricardo is an economist at Claremont McKenna College, and we co-authored a paper on Zipf's Law.[7] In 1935, George Kingsley Zipf, a Harvard linguist, noticed that if you look at the English language, the frequency with which a word occurs is inversely related to the rank of the word. Hence, the most common word shows up a certain number of times, the second one shows up half has many times, the third, a third as many, and so on. This phenomenon occurs in a lot of different systems, including the distribution of wealth, income, and the capitalization of companies. I'd never heard of Zipf's Law, but economists knew about it, and Ricardo told me about it.

In any case, one branch of SPT exists in rank-based analysis, and the other in functionally generated portfolios and arbitrage. Both branches were introduced in the monograph, but most of the development has taken place since then. Many of the processes are more complicated than I would like—I like simplicity—but mathematicians love to have complicated problems to solve.

The intersection of the two branches occurs in money management. If you're managing a portfolio that looks like it might be functionally generated, you have to be able to measure the leakage, and leakage is a ranked-based phenomenon.

Inna Okounkova: To continue with practical applications, a number of companies have applied your ideas, not just Intech. Jason Hsu listed those companies and even their results at the Society of Quantitative Analysts seminar where you presented your article by Banner et al. (2019). So, how scalable is this investment approach? You said it cannot be done for everyone. Someone needs to be on the other side to keep outperforming the market. So, how much money can be deployed in this strategy before it stops working?

Robert Fernholz: It's getting tougher for the people who are using the theory. The biggest problem, essentially for all mathematical and quantitative equity managers, is front-running by the hedge funds. Hedge funds that invest in equities frequently achieve their main profits through front-running active managers, and mathematical managers are more predictable and hence easier to front-run.

Brokerage houses sometimes do it too. When they say they want to get "market color," they're actually trying to figure out when you're going to do your next trade. This practice started about 2005 or 2006, and it's become very complicated. So now you need strategies that allow you to hide what you're doing and to move from one strategy to another at random times.

Other phenomena—factors, in particular—are complicating things now. I don't think factors were terribly important until they were made important; they became important because people think they're important.

ETFs [exchange-traded funds] are now frequently based on a particular factor, and I think these ETFs are actually moving the market. The cart is now pulling the horse. These things make life more complicated for stock managers than it used to be.

Edward Baker: What role does judgment play in quantitative processes—either in terms of actual implementation or simply in the way things must be adjusted over time to reflect market changes and instabilities?

Robert Fernholz: I'm retired from Intech, though I've been a consultant for a number of years, but here's the way we did things. We did our research, we set up our strategy, and we then did cross-validation and whatever other testing was necessary to determine whether the strategy was going to work.

Then we implemented the plan. We didn't override it unless we discovered that it was significantly underperforming our expectations, in which case we had to update it or abandon it. With individual trades, though, we had a computer that generated all our trades, because generating them by hand is too complicated. At least when I was there—and I presume they've continued this—the trades were scanned by the traders for possible errors before they were sent to the brokers.

First, the computer checks to see whether anything looks anomalous, and then a trader checks to make sure any trade the computer flagged makes sense. Somebody needs to look at the trades carefully, so there's judgment involved, but only to check for anomalies. If there's some huge crash, as seems to occur from day to day now, you might want to postpone trading for a few days because you can't tell whether the prices are valid.

Aside from that, you do not override what the machine says because otherwise you're just flying blind. If you've done your work and you have the theory straight, you follow what the machine says. If the computer says something you think is wrong, maybe you stop the process for a little while, but you don't start making decisions on the basis of suppositions such as you think the market's going to go up tomorrow so you buy more stocks.

Edward Baker: What happens when market conditions change—for example, when there's a structural change such as the introduction of those factor portfolios you referred to—and the models need to be adjusted?

Robert Fernholz: When I was at Intech, all the portfolios were reoptimized on a quarterly or monthly basis, depending on the portfolio. When you get new data, you re-optimize everything. And of course, you track an index so your positions are going to change according to how the index has changed. You try to capture the changes through that monthly or quarterly re-optimization, which incorporates all the new data.

Edward Baker: And that takes care of it?

Robert Fernholz: To some extent. We don't have zero tracking error, but you do the best you can. You get new information and incorporate it. Some of it is incorporated daily because it tracks the trading system that follows the index and keeps an eye on where you're supposed to be. And then you change the parameters on a monthly or quarterly basis.

You could change the parameters more frequently, but you don't get that much new information from more frequent data, at least for the type of strategies we were using. Plus, the parameters are the result of data analysis, not a question of somebody's judgment. Of course, judgment goes into the data analysis techniques you choose, and those also will change from time to time.

Geoffrey Gerber: In developing SPT as a descriptive rather than a normative mathematical theory for analyzing stock market and portfolio returns, you use volatilities and correlations. As we know, volatilities and correlations can be nonstationary—low beta stocks like those of financial firms or utilities can become high beta stocks and vice versa. So how does SPT deal with changing volatilities and correlations over time? And why is it more effective than modern portfolio theory?

Robert Fernholz: Basically, with its normative framework, modern portfolio theory explains why one cannot construct a portfolio that will outperform the market portfolio without incurring increased risk, while stochastic portfolio theory, with its descriptive framework, shows how one can construct a portfolio that will outperform the market portfolio without incurring increased risk.

SPT was developed to address the problem of the constant coefficients in the Merton model[8] that result in the Rosenberg-Ohlson paradox. That was what got me interested in this line of thinking. On the theoretical side, all the parameters were assumed to be variable from the beginning. On the practical side, I've talked about the need for re-optimization as a result of changing parameters.

When you look at a market, everything is variable, and stochastic calculus will take care of that. We use continuous semimartingales, fairly standard now in mathematical finance. Sure, there are jumps in prices, but you can approximate those as closely as you want with continuous functions.

As for the question of descriptive versus normative, CAPM is based on the normative assumption that everybody is optimizing their utility function, and that they have complete knowledge of all the market parameters. Normative theories have been prevalent in economics and finance, although they're being replaced somewhat now with behavioral economics, which I think is an important development. SPT gets completely away from these normative assumptions. The natural sciences, at least since Galileo, have been descriptive, and descriptive epistemology is the basis of SPT. In other words, you observe a phenomenon, and you try to explain what's going on. The models in SPT were built to capture as closely as possible what goes on in the market, and that's variable. There are other lesser effects that occur, like splits and dividends, but they are not difficult deal with.

These two branches, finance and mathematical finance, don't interact much. People in industry mostly use finance, except the hedge funds, which, I think, use mathematical finance.

Inna Okounkova: But don't those kinds of natural market changes generate only about 5-percent turnover a year?

Robert Fernholz: Yes, the natural market changes are only about 5 percent, but these portfolios have a much higher turnover than that. I think optimized portfolios would have more like 100-percent turnover a year. I forget the turnover percentage of the diversity-weighted portfolios, but I think their turnover was about 50 percent.

Edward Baker: Why isn't private sector support to the field of mathematics more widespread? I'm the chairman of the Mathematical Sciences Research Institute's board of trustees, and we struggle trying to find sources of private support. Fortunately, the National Science Foundation continues to provide generous support, but it's hard to find individuals who believe in giving money to support mathematics. Why do you think that is?

Robert Fernholz: Math is not flashy or glamorous. On the other hand, math isn't very expensive. Mathematicians traditionally have needed only paper and pencil or a blackboard and chalk. Now, however, they're getting involved in computer science, including the use of algorithms, which can be expensive. More money is going into that because people think they can make money if they have better algorithms.

Math problems are hard to explain to the public, and there aren't any pictures of black holes that can appear in the newspaper, so I think people sort of forget about it. Almost everybody in science understands that mathematics is the basis of everything, but there's not much support for mathematics among the broader population. What you really need is a responsible government that recognizes the importance of mathematics and supports it accordingly.

Edward Baker: I remember reading Feller's book when I was a math student. It was more about applied mathematics than pure math. You mentioned that your interest in applied mathematics is more acute, but a lot of pure mathematics becomes important in applications even though it might have been viewed as having no possible practical applications when initially developed.

Robert Fernholz: It's absolutely critical. Without pure math, you don't have the machinery to do applied math. I'm sure that the National Science Foundation understands that if they don't support math, all of science in the country will suffer.

Edward Baker: Another thing people don't appreciate is just how collaborative mathematics is and how crucial facilitating that collaboration is.

Robert Fernholz: Collaboration is important because mathematicians have different skills and different thought processes. Bringing all these skills together allows you to solve a problem.

Philip Fazio: Why do you believe that excess growth rate as a measure of diversification and manager or fund performance has not become a widely accepted investment analytical tool?

Robert Fernholz: This is an interesting question. I think there are two cultures—mathematics and finance. There's mathematical finance, which is a branch of mathematics, and in this discipline rigorous proofs are required for everything. And then there's finance, which is a branch of economics, and economics is a more normative field.

These two branches, finance and mathematical finance, don't interact much. People in industry mostly use finance, except the hedge funds, which, I think, use mathematical finance. The hedge funds don't reveal exactly what they do, but it seems pretty straightforward that something like that is going on. The papers that are written and the courses that are taught in the majority of the business schools are classic finance, and classic finance is like a priesthood. These people have invested a lot of time in learning the standard rules, and they want to continue in that tradition.

Ludwig Chincarini: What is the most important aspect of successfully managing money, and how do quantitative models perform at a time of unexpected crisis such as what we're experiencing right now?

Robert Fernholz: The most important aspect of successfully managing money is that you tell the truth. Don't make claims you can't substantiate. Tell people what you can do, and don't claim you can do things you can't do. Don't misrepresent what you're doing. Of course, this means you have to have some understanding of what you're doing.

How do these models work in extraordinary circumstances? I don't know exactly. I mentioned the 1987 crash. I think we moved a few basis points away from the benchmarks during the recent crash in the spring of 2020.

We have reasonable statistical techniques to avoid being misled by these kinds of events, and if these strategies are decently robust, they will survive a crash pretty well. Their returns might be a little ahead or behind the market. Then afterward, you have an estimation problem in trying to clean up the data so that you end up with data you think will represent the future rather than just the crash, because these variances are quadratic functions. When you're trying to estimate a variance, it blows up if you have one huge move.

For the type of stock portfolios we managed, you normally measure performance relative to the market or benchmark. It's difficult to measure performance absolutely when you're following a volatile market. It's more reasonable to assess things relative to the market. But there are no miracles. If the whole market goes down and you hold part of the market, you're going to go down with it. You won't be able to get through a crash unscathed, but if you have a fairly robust process, you shouldn't do too badly compared with your benchmark.

Edward Baker: The new principle of no-instantaneous-arbitrage still seems like an important concept in financial markets, but you don't appear to assume it in some cases. Do you find other principles more important?

Robert Fernholz: I think the principle that's most important is the use of descriptive methodology. The absence of arbitrage is a normative principle. No-arbitrage means that even if you have a long period of time, you can't construct a portfolio that will outperform the market with probability one.

There is one normative assumption that seems to be universal, and this is the assumption of strong nondegeneracy for the covariance matrix. I won't go into detail, but this is actually a weak assumption that seems to be reasonable from a descriptive point of view, too. If you make that normative assumption and you add the assumption that the market will not be completely concentrated in a single stock, then you actually get arbitrage over a long enough, but fixed, period of time. Therefore, that model, which seems to be close to reality, would violate the no-arbitrage principle. Some people may claim that we need the no-arbitrage principle, but with no-arbitrage goes the possibility of complete concentration of the market.

I think what we need is the principle of no-instantaneous-arbitrage, which means you can't have a portfolio that will instantly outperform the market. The nonexistence of instantaneous arbitrage is equivalent to the possibility of MVO. MVO is a reasonable, and very weak, normative assumption, but a blanket no-arbitrage condition is not consistent with observed reality.

Also, most theoretical markets are closed. You assume they have a fixed number of stocks and they just stay that way. But real markets are open markets. Stocks drop out and are replaced and leakage occurs. In this case the arbitrage question becomes more complicated. Researchers are just beginning to study open markets, but whether the markets are open or closed, I believe that the no-arbitrage assumption should be replaced by no-instantaneous-arbitrage.

Inna Okounkova: We also would like to hear your views about the future of the investment industry and what advice you would give to someone entering this industry today.

Robert Fernholz: I think active equity management is extremely difficult these days. Active managers can perform exactly as they have promised, and they will still hear things like, "Unfortunately, we have to terminate you because the board has decided to move everything into indexes." So, even if you're an active manager and you achieve exactly what you predict, which is to outperform the benchmark, you might get terminated anyway. Active stock management doesn't seem like a good way to go.

There are too many ETFs and there are too many factors. One approach you could take is to do big-data high-speed trading, but that would require you to put a lot of money into the machinery. Your competitors probably would have already invested billions of dollars in computing equipment at that point, so it's probably pretty tough to get into that, but you might be able to do it.

Another approach you could consider would be asset allocation. Asset allocation can be useful. But with fixed income investments at 0-percent interest rates? I don't know. What else can you do? I read that doctors' practices are being privatized now. Really? I don't know if you'd want to get into that type of private equity.

Venture capital could be exciting, but that requires a completely different skill. Still, I think that's spectacular. Do that if you can.

As for management of large investments like endowments, I believe that David Swensen, Yale's long-time chief investment officer, has said we can expect a yield of about 5 percent in the future. And endowment managers are probably the best asset allocators in the world. I'm not sure what Warren Buffett is saying these days, but it's difficult to do well even in asset allocation these days.

Inna Okounkova: How would you apply your theory of excess growth rate and your diversity-weighted approach to asset allocation?

Robert Fernholz: These approaches are not directly applicable because you have to take different asset classes into account. In that case, you could use statistical techniques and possibly big data applications like neural networks to try to figure out what's going to happen with the different asset classes. The excess growth rate won't be a major factor in asset allocation because you only have a half-dozen different asset classes and they won't have very high variances, so you don't have much to work with from that perspective. I've worked in statistics and adaptive time-series analysis, and I'm actually getting back into neural networks, but that requires a different methodology. It wouldn't actually use SPT in asset allocation as far as I can tell.

Inna Okounkova: What would happen if everyone started investing in diversity-weighted indexing?

Robert Fernholz: Well, if you start with a lot of diversity-weighted indexing and just move money from the bigger stocks to the smaller stocks, this would change the shape of the distribution curve. If you make a log-log chart of the capital distribution, there's a point where the slope of the tangent is minus one. I think it comes after about a thousand stocks, and it's curved. You would flatten that curve a little more if you moved money into the smaller stocks. In other words, if you bring the curve's slope up at the bottom, it's already flattened out a certain amount at the top, so you drive that data point down a bit further. The market probably would be slightly more efficient if more people were doing that type of diversity weighting.

But you can't have everybody doing it. Among other things, the advantage, or the drift rate, you get is from absorbing trades because you always sell on upticks and buy on downticks. By the way, you actually could move into high-speed trading, which I wrote about in Fernholz and Maguire (2007). If you measure volatilities over a very short period of time, such as five minutes, you get a very different result than if you measure them over a day or a week or a year. The more frequently you trade, the more volatility you absorb. The more money that goes into diversity weighting, the flatter the capital distribution becomes. Whether this would be good or bad for the market, I don't know.

Nobody knows what good or bad means. Perhaps having more money in the smaller stocks would be good. Perhaps less volatility would be good.

Inna Okounkova: Thank you for sharing your ideas with us and for your contributions to our journal.

ENDNOTES

1. George Arthur Akerlof (1940–) is an American economist who is a professor at the McCourt School of Public Policy at Georgetown University and the Koshland Professor of Economics Emeritus at the University of California, Berkeley. He won the 2001 Nobel Memorial Prize in Economic Sciences (shared with Michael Spence and Joseph E. Stiglitz).
2. William Feller (1906–1970) was a Croatian–American mathematician who specialized in probability theory. He is remembered for championing probability theory as a branch of mathematical analysis in Sweden and the United States.
3. Harry Markowitz (1927–) is an American economist and a recipient of the 1989 John von Neumann Theory Prize and the 1990 Nobel Memorial Prize in Economic Sciences; he is an adjunct professor of finance at the Rady School of Management at the University of California, San Diego. He is best known for his pioneering work in modern portfolio theory, in which he studied the effects of asset risk, return, correlation, and diversification on probable investment portfolio returns.
4. In 1968, Markowitz joined Arbitrage Management Company, founded by Michael Goodkin. Working with Paul Samuelson and Robert Merton, he created a hedge fund that represents the first known attempt at computerized arbitrage trading. He took over as chief executive in 1970. After a successful run as a private hedge fund, AMC was sold to Stuart & Co. in 1971. A year later, Markowitz left the company.
5. Black–Scholes is a mathematical model for pricing an options contract. The model estimates the variation of financial instruments over time. It assumes the prices of these instruments (such as stocks or futures) will have a log–normal distribution. Using this assumption and factoring in other important variables, the equation derives the price of a call option (www.investopedia.com).
6. Robert A. Garvy (1942–), founder and chairman emeritus of Intech, joined the company in January 1991. In partnership with Robert Fernholz, he helped build the firm from a single-product provider into a solutions-based global investment manager with more than $40 billion in assets under management in 2019.
7. In probability, Zipf's Law is the assertion that the frequencies of certain events are inversely proportional to their rank (https://www.britannica.com/topic/Zipfs-law).
8. The Merton model is an analysis model used to assess the credit risk of a company's debt. Analysts and investors use the Merton model to understand how capable a company is of meeting financial obligations and servicing its debt, and to weigh the general possibility that it will go into credit default (www.investopedia.com).

REFERENCES

Arnott, R. D., J. Hsu, V. Kalesnik, and P. Tindall. 2013. The Surprising Alpha from Malkiel's Monkey and Upside-Down Strategies. *Journal of Portfolio Management* 39, no. 4 (summer): 91–105.

Banner, A., R. Fernholz, V. Papathanakos, J. Ruf, and D. Schofield. 2019. Diversification, Volatility, and Surprising Alpha. *Journal of Investment Consulting* 19, no. 1: 23–30.

Fernholz, R. 2002. *Stochastic Portfolio Theory*. New York: Springer-Verlag.

Fernholz, R. T., and R. Fernholz. 2020. Zipf's Law for Atlas Models. *Journal of Applied Probability* 57, no. 4 (December). http://arxiv.org/pdf/1707.04285. [forthcoming in December 2020 issue]

Fernholz, R., and C. Maguire, Jr. 2007. The Statistics of Statistical Arbitrage. *Financial Analysts Journal* 63, no. 2 (September): 46–52.

Fernholz, R., and B. Shay. 1982. Stochastic Portfolio Theory and Stock Market Equilibrium. *Journal of Finance* 37, no. 2: 615–624.

Fernholz, R., I. Karatzas, and J. Ruf, 2018. Volatility and Arbitrage. *Annals of Applied Probability* 28, no. 1: 378–417.

Rosenberg, B., and J. A. Ohlson. 1976. The Stationary Distribution of Returns and Portfolio Separation in Capital Markets: A Fundamental Contradiction. *Journal of Financial and Quantitative Analysis* 11, no. 3 (September): 393–402.

5619 DTC Parkway, Suite 500
Greenwood Village, CO 80111
Phone: +1 303-770-3377
www.investmentsandwealth.org