Perspectives

LEGG MASON
GLOBAL ASSET MANAGEMENT

March 2012
Legg Mason Capital Management
Shaking the Foundation

LEGG MASON
CAPITAL MANAGEMENT

Mauboussin on STRATEGY

March 23, 2012

Michael J. Mauboussin

mmauboussin@lmcm.com

Shaking the Foundation

Revisiting Basic Assumptions about Risk, Reward, and Optimal Portfolios

An Interview with Ole Peters

Investing is fundamentally a bet on the future. The problem is that we only get to live once: we have to make choices in the face of uncertainty and deal with the consequences for better or worse. The mean-variance approach, the most common guide to build a portfolio, basically says that you should get the most return for a given level of risk based on the expected values of the individual assets and on how they correlate with one another.

One of the great aspects of an affiliation with the Santa Fe Institute (SFI) is the opportunity to exchange ideas with world class scientists who are self-selected to be curious about the world. Ole Peters, trained as a physicist and a visiting researcher at SFI, is a great example. Ole has been discussing his concerns about the standard theories in economics for years, including portfolio theory, and he gave a talk last fall that prodded me to share his thoughts with a wider audience.

At first, I considered writing about this myself. But then I figured it would be better to interview Ole and let you hear the story directly from him. Fair warning: this is not easy material. But I believe working through these ideas and their implications is time well spent.

Here's a brief summary:

- *Distinction between ergodic and non-ergodic systems*. Ole starts us off with this crucial distinction. An ergodic system is one where the ensemble average and the time average are the same. For example, the proportion of heads and tails is the same either if you ask 1,000 people to flip a coin at the same time (ensemble average because an ensemble of people are flipping simultaneously) or if you flip it yourself 1,000 times (time average because it takes time to flip sequentially). His point is that many of the models that economists use were designed to deal with ergodic systems, yet the reality is that we live in a world that is non-ergodic. That mismatch is problematic for portfolio construction.

- *How theory evolves is important*. Over the last 150 years, economists have borrowed ideas from physics—generally with the goal of making economics into a harder science. In the late 1800s, Ludwig Boltzmann described ergodicity, but recognized that it worked under very narrow conditions. Economists embraced the "ergodic hypothesis" and applied it to systems that are not ergodic. The stock market is a good example.

- *Optimal leverage*. In a mean-variance approach, risk and reward are related to one another in a linear fashion: more risk, more potential reward. And the main way to increase risk is to add leverage through debt. The key is that the ensemble average gets "rid of fluctuations before the fluctuations really have any effect." If you assume a time average and a multiplicative process—that is, you parlay your bets—then there is an optimal amount of leverage. Too little leverage leaves potential profit on the table, but too much leverage assures ruin.

MM: Ole, thanks for taking the time to share your thoughts with us. In many fields, there are foundational assumptions that, once established, are often not discussed or critically examined. Because these foundations are laid by the intellectual leaders in the field, they tend to go unquestioned.

You have been doing some fascinating work that challenges some of the foundational assumptions in economics and finance. One such assumption is called the "ergodic hypothesis." Can you explain the difference between an ergodic and a non-ergodic system and tell us why this is so relevant in economics and finance?

OP: Thank you, Michael. It's a pleasure to discuss this with you.

Ergodicity is this complicated-sounding word, and a lot of highly technical work has been done in ergodic theory, but as usual the real meat is conceptual. The conceptual part is deep and subtle and utterly fascinating, but it's not actually very hard to understand. That doesn't mean we can leave out the mathematics. If you want to use the concept, you have to understand the mathematics—there's "no royal road", as Luca Pacioli put it. Without the mathematics, you can't get beyond what I call an "incomplete treacherous intuition" of the meaning. But since this is an interview, not a mathematics lecture, I will try to convey that incomplete treacherous intuition.

Here it is: In an ergodic system time is irrelevant and has no direction. Nothing changes in any significant way; at most you will see some short-lived fluctuations. An ergodic system is indifferent to its initial conditions: if you re-start it, after a little while it always falls into the same equilibrium behavior.

For example, say I gave 1,000 people one die each, had them roll their die once, added all the points rolled, and divided by 1,000. That would be a finite-sample average, approaching the ensemble average as I include more and more people.

Now say I rolled a die 1,000 times in a row, added all the points rolled and divided by 1,000. That would be a finite-time average, approaching the time average as I keep rolling that die.

One implication of ergodicity is that ensemble averages will be the same as time averages. In the first case, it is the size of the sample that eventually removes the randomness from the system. In the second case, it is the time that I'm devoting to rolling that removes randomness. But both methods give the same answer, within errors. In this sense, rolling dice is an ergodic system.

I say "in this sense" because if we bet on the *results* of rolling a die, wealth does not follow an ergodic process under typical betting rules. If I go bankrupt, I'll stay bankrupt. So the time average of my wealth will approach zero as time passes, even though the ensemble average of my wealth may increase.

A precondition for ergodicity is stationarity, so there can be no growth in an ergodic system. Ergodic systems are zero-sum games: things slosh around from here to there and back, but nothing is ever added, invented, created or lost. No branching occurs in an ergodic system, no decision has any consequences because sooner or later we'll end up in the same situation again and can reconsider. The key is that most systems of interest to us, including finance, are non-ergodic.

It's probably best to step back and start at the beginning. The word "ergodicity" was coined during the development of statistical mechanics. Ludwig Boltzmann, an Austrian physicist, invented it. He called ergodic systems "monodic" at first. Here's the etymology: a "monodic" system is one that possesses only one (mon-) path (-odos) through the space of its possible states. The word became "ergodic" because Boltzmann considered systems of fixed energy (or work = ergon), and the idea was that the one path covers all the states allowed by that energy (the energy shell).

If this is the case, Boltzmann argued that the system will over time—in his case a second may be long enough—visit all the relevant states that it can access. The system will visit each state with some relative frequency. We can mathematically treat those frequencies as probabilities, which has the incredibly nice consequence that the long-time averages of the quantities we're interested in are formally the same as expectation values arising from the relative frequencies that we interpreted as probabilities. Just like in the dice example.

This makes the mathematics extremely simple. But it is a trick (Boltzmann literally called it a trick): we calculate expectation values that are *a priori* [1] irrelevant. In this very special case of an ergodic system, expectation values are the same as time averages, and that's why we're interested in these special cases.

Time averages are interesting because they are what we observe in physics. We usually measure some sort of macroscopic property, like the pressure in a balloon. That pressure, in terms of the microscopic model, is the rate of momentum transfer per unit area to the balloon membrane resulting from a gazillion collisions of molecules with the membrane. Any device that we use to measure this is so sluggish that it will only give us a long-time average value of that momentum transfer. The very clever insight of Boltzmann was that under very special conditions, we just have to calculate an expectation value of the rate of momentum transfer per area, and that will coincide with the time average pressure that we actually observe.

So, practically, ergodic means that time averages are the same as ensemble averages, or expectation values. Non-ergodicity means that they are different. Since there are many more ways of being different than there are of being identical, it comes as no surprise that most systems are non-ergodic.

Why would this be important in economics and finance? Quite simply because Boltzmann's trick doesn't work. We cannot equate the behavior over time (i.e., what really happens) to this elusive mathematical object, the expectation value.

I'm not arguing for new models, I'm just arguing that we should have a look at the very sensible models that economists have devised and to see whether they are ergodic. The ones I'm interested in are non-ergodic, and so my work has focused on pointing that out and asking about the consequences. Where did we miss the lessons of these models because we were wearing the wrong glasses?

MM. Let me jump in here and note that Paul Samuelson, a Nobel-prize winner, claimed that the “ergodic hypothesis” is essential for advancing economics from the realm of history to the realm of science. So he thought that assuming ergodicity is essential to economics and finance.[2]

OP: Your Samuelson quote really gets to the heart of the epistemological issue. Samuelson said that we should accept the ergodic hypothesis because if a system is not ergodic you cannot treat it scientifically. First of all, that's incorrect, although I think I understand how he ended up with this impression: ergodicity means that a system is very insensitive to initial conditions or perturbations and details of the dynamics, and that makes it easy to make universal statements about such systems. In physics we know this all too well—ergodic equilibrium systems naturally fall into strongly attracting universality classes.

When we move away from equilibrium, or ergodicity, everything becomes more complicated. But that does not mean that we can't make meaningful statements about non-ergodic systems. Actually, that's what the contributors to statistical mechanics have been doing for the last 100 years or so—Boltzmann only started the field, he didn't finish it. So, there's a grain of truth in what Samuelson said, but he severely overestimates and overstates the problem, and I'm struggling to understand why he did that.

Another problem with Samuelson's statement is the logic: we should accept this hypothesis because then we can make universal statements. But before we make any hypothesis—even one that makes our lives easier—we should check whether we know it to be wrong. In this case, there's nothing to hypothesize. Financial and economic systems are non-ergodic. And if that means we can't say anything meaningful, then perhaps we shouldn't try to make meaningful claims. Well, perhaps we can speak for entertainment, but we cannot claim that it's meaningful.

In what sense would saying something that's patently false be "meaningful," or "scientific" rather than "historical"? You can see where I'm going with this. Important models that economists use are not ergodic, so what's this debate about? In physics, Boltzman hypothesized ergodicity because it's not possible to compute time averages for a system as complicated as 10^{24} molecules bouncing around. He had to simplify the mathematics, even if that meant resorting to fiendish tricks. Many would argue that we cannot justify these tricks in physics, but scientists confirmed their validity indirectly many times by experimentally testing predictions based on them.

In finance or economics the situation is different. Take the most basic model of a stock market, Louis Bachelier's random walk. Is that model ergodic? No. A little later, in the 1950s, maybe starting with M. F. M. Osborne, the popular model in finance became geometric Brownian motion—basically a random walk in log-space. That's a very sensible model, and except for some details about fat tails and correlations it fits stock market data pretty well. Let's say that none of its deficiencies make it any more or less ergodic. It's also not a bad description of the world economy or of national economies.

Since geometric Brownian motion is a mathematical model, you can answer the question of whether that's ergodic by scribbling a few lines of equations. Of course it is not. It's a model of growth, after all, so it can't be ergodic, but you can actually make this completely formal and do the math, and not even the expectation value of the growth rate is equal to the time average of the growth rate. At the end of the day, what's more important in finance than growth rates?

So Samuelson's comment makes little sense. A hypothesis is about something we don't know, but in the case of finance models this is something we do know. There's no reason to hypothesize—the system is not ergodic. It's like hypothesizing that 3 times 4 is 0 because it makes the mathematics simpler. But I can calculate that the product is 12. Of course, a formalism that's based on the 3-times-4 hypothesis will run into trouble sooner or later. In economics, that happens with the ergodic hypothesis when we think about risk, or financial stability. Or inequality, as we're just working out at the moment.

The reason this is so important is quite simple, and stems from a basic question: what does risk mean if the notion of time is not irreversible? The only reason risk exists is that we cannot go back and make decisions over again. Economics got very confused about the point of dealing with risk, and had to resort to introducing psychology and human behavior and all sorts of things. I don't mean to say that we don't need behavioral economics. What I mean is that there are lots of questions in economics that we can only answer behaviorally at the moment, but at the same time we have a perfectly formal natural physical analytic answer that's very intuitive and sensible and that comes straight out of recognizing the non-ergodicity of the situation.

To be blunter, I'm pointing out that economics is internally inconsistent. I accept all the models that economists have developed. I could critique them, but I'm not worried about that. I didn't make them up, the economists did. But when the economists treat the models as if they were ergodic, that's when someone has to say "stop, that's enough."

MM: You've mentioned Boltzmann and Bachelier and their roles in the development of thought about random systems. Can you take a step back and trace the history of randomness and thinking about risk? In particular, I am thinking of the classic St. Petersburg Paradox.[3] In this game, the house (with infinite wealth) flips a fair coin. If it lands on heads, you receive $2 and the game is over. If it lands on tails, the house flips again. If the second flip lands on heads you get $4, if it lands on tails, the game continues. For each successive round, the payoff for landing on heads doubles (i.e., $2, $4, $8, $16, etc.) and you progress to the next round until you land on heads. This is a game with an infinite expected value (expected value = 1 + 1 + 1 +… = ∞) but few people are willing to pay more than a few dollars to play.

Daniel Bernoulli, who presented this game in 1738, came up with a clever way to resolve the paradox by adding something called "utility"—a measure of satisfaction. The key is that as your wealth increases, the amount of satisfaction you gain from each incremental dollar declines. So an incremental dollar of wealth will have great utility for a poor person but little utility for a billionaire. So he used the theory of utility to explain why people would not sacrifice their net worth to play the game even though the expected value is vastly in excess of their net worth.

Is it possible to resolve the St. Petersburg Paradox without resorting to utility theory? What does the St. Petersburg game teach us about how we approach randomness and risk?

OP: You're absolutely right—the history of formal thought on randomness is older than Bachelier (1900) and older than Boltzmann (1870s). The late 19th century is the time when randomness entered physics, and that's when people started asking questions about ergodicity. But formal mathematical treatments of random systems started with an exchange of letters between Pierre Fermat and Blaise Pascal in the summer of 1654.

In this exchange of letters Fermat developed the concept of expectation values. It's incredibly important to remember that Fermat developed the idea of expectation value as a moral concept. It's a formal notion of fairness, and initially had nothing to do with predictions. The question Fermat and Pascal were working on was how to split the pot in a fair way if two players are in the middle of a game of craps, each player has posted a wager on the outcome of the game, and the police burst in and halt the game unfinished. Nothing about predictions here.

Fermat invented the notion of the ensemble. He said that we should imagine that everything that could have happened really did happen in a collection of parallel worlds. He said it would be fair to take an average over those parallel worlds. He also insisted that the worlds would have to be chosen so that each one has the same likelihood, meaning if event A is twice as likely as event B then we create twice as many hypothetical worlds where event A happens than worlds where B happens.

It didn't take long for people to realize the relevance of expectation values in predictive contexts. This works when there really are lots of systems that run in parallel and end up sharing their resources in using some sort of average. An example is life annuities, an early financial product that Edmond Halley (of comet fame) first priced in 1693. The idea is that if a king needs money, whether for war or a new palace, he can sell this financial product. You give the king some money today, and he will pay you a pension for the rest of your life. Halley reckoned correctly that if the king sells this product to lots of people at something close to the expected payout, some people will die early, others later, and the total payout will be such that the king is likely to break even.

But by that time, people had already forgotten where the expectation values come from. What remained was a consensus that the price of any product with a random payout should equal its expectation value.

Then the St. Petersburg paradox came up. This was introduced by Nicolaus Bernoulli in 1713 in a letter to the French mathematician, Pierre Rémond de Montmort. Nicolaus Bernoulli just said look at this game, and you'll find something curious. What he meant was that the expected payout is infinite. The problem becomes a paradox only if we forget that expectation values are only a correct price under special circumstances.

So how can St. Petersburg be solved without utility? The paradox is that people don't behave the way the mathematics at the time suggested they should. But that's not because people are strange, but rather because mathematics was not very advanced in 1713. If you think about what matters to people, and what evolution has taught us, it's pretty clear that it's our performance *over time* that matters, not some average over parallel copies of ourselves. Nicolaus Bernoulli was making fun of the idea that the expectation value is the only thing that matters. Today it should be obvious that the idea is wrong. Let's say we roll a die with the agreement that if it comes up a one, I shoot you, and if any other number appears, you'll live. The mathematical expectation value of the game is 3.5. So according to the expectation value, you'll live. But I suspect that I'd be correct in saying that you wouldn't play the game.

Back to St. Petersburg. Instead of computing the expectation value of your increase in wealth, you can compute the time average of your increase in wealth.[4] Remember, that's what really matters. The only reason we use expectation values (under special conditions) in physics is that they can be the same as time averages, but time averages is what we're really after. There's one more little thing. In order to compute a time average, we need a dynamic. A dynamic links probabilities to what happens in time. This is what makes non-ergodic systems less universal—now that time has a meaning, we need to specify a dynamic, and the answers we find will depend on that dynamic. But that's just how it is, we can't just compute meaningless quantities only because it's easier.

The St. Petersburg case uses the multiplicative dynamic, which is very natural. This encodes the idea that if I'm broke I can't buy any more St. Petersburg lottery tickets. And if I get very rich, I can re-invest my winnings. It's a sensible dynamic, the same as that behind the Black-Scholes equation, or geometric Brownian motion. There's actually a paper about Babylonian commodity prices that were found on clay tablets, and they seem to follow a multiplicative dynamic. [5]

So, all these details aside—we just calculate the time average instead of the expectation value. While the expectation value doesn't exist (that's very sensible mathematics jargon for "is infinite," or "diverges"), the time average is perfectly finite and pretty much in line with how people behave. If you insist on translating it into a utility function, you'd end up with logarithmic utility, but there's no need, really.

As an added bonus to resolving the apparent paradox, the logarithm that Daniel Bernoulli used as his utility suddenly has a meaning. We know where it comes from: it's the assumption of multiplicative dynamics. But now we know that it really is a logarithm and why, and Daniel Bernoulli's expression in his 1738 paper turns out to be only almost correct. From this new point of view he actually made a little error. Daniel Bernoulli did not calculate the expected net change in utility, but no one thought it was a problem because utility was quite an arbitrary concept anyway.

These things are important. Pierre-Simon Laplace, for example, corrected Bernoulli when he recounted Daniel's argument in 1812 but he never mentioned that what he wrote was actually a correction. He probably thought it didn't matter much anyway and was only a matter of taste. Once we have a physical solution to the problem, it's not about taste any more. It's not so much that we know now whether things are correct or incorrect, but we know exactly what we mean and we can check our arguments for consistency.

MM. Ole, how does your work fit into the process of scientific inquiry?

OP: What I'm doing follows a totally normal process of scientific discovery: the expectation value is a good approximation for time averages and a basis for sensible behavior *under certain conditions*. We can characterize these conditions as "small leverage." As long as the decision is about what I'm doing with a negligible fraction of my wealth, the expectation value is sensible. But as leverage increases, i.e., is not close to zero or surpasses one, as was the case in the financial crisis, it is a horrible approximation.

More formally, the ensemble-average growth rate approaches the time-average growth rate in the limit as leverage approaches zero. This is the situation Fermat and Pascal originally considered. They were talking about gambling for fun with some friends—we bet a dollar to make it a bit more interesting. Even earlier work is concerned with that specific case. For example, Gerolamo

Cardano in the 16th century emphasized that "there must be moderation in the amount involved"—that's the limit of leverage being close to zero.

Unfortunately, all of this was forgotten and academics and practitioners applied the rule of thumb for small-leverage bets to large-leverage bets, with predictably disastrous results.

This is actually a totally normal trajectory of scientific understanding because it mirrors the development in many other fields. It's not even uncommon for such development in thought to take centuries:

a) Newton's gravitational law (published in 1686) was considered correct, until it was found to be a rule of thumb valid in a certain regime, and Einstein (more than 200 years later) introduced a more accurate theory that contains the old one as a limit (just like the time perspective contains expectation-value perspective in the limit of leverage approaching zero);

b) developed in the 18th and 19th century, all of mechanics was found to be a rule of thumb valid in a certain regime, namely large masses or something like that. About 100 years later a more accurate theory that worked outside that regime was put forward: quantum mechanics. The Correspondence Principle demands that quantum mechanics contain every-day mechanics as a limit (just like Einstein contains Newton and time contains expectation);

c) in optics, we've worked with the diffraction limit as if it were a hard limit for a long time, until recently—by revisiting Maxwell's equations—we learned (or reminded ourselves) that it is a rule of thumb, valid in some limit.

My work is most similar to optics, c), because a) and b) are actual discoveries of new laws, whereas what happened in optics is "just" a re-visiting of laws that were already known. I didn't discover time. We all know that time goes in one direction. I'm just reminding us of this fact and I'm re-visiting, or exploring further, the consequences of the knowledge that was always there.

MM: Let me shift directions a bit and ask about the practical implications of what you are saying. For example, if I build a portfolio guided by the mean-variance framework, risk and reward are related in a linear fashion: more risk, more reward. But that doesn't appear to be true for a portfolio that is guided by the principle of geometric mean maximization. Can you discuss what your work tells us about building portfolios, and, perhaps, how some large failures (e.g., Long-Term Capital Management) did not incorporate this thinking?

OP: A very relevant question, Michael. Let me first repeat that what I'm doing is geometric mean maximization only in the case where the dynamics are purely multiplicative. It's important to keep this in mind because otherwise one can be side-tracked into a debate about whether geometric means are good or not.

My conceptual statement is deeper than "use geometric means." I'm saying let's optimize portfolios, or really any sort of object, for their behavior over time, not across an ensemble of imagined parallel universes. For this we need a dynamic, and if that dynamic is multiplicative, the technique amounts to geometric-mean optimization. I should also repeat that I haven't made any big discovery—this was done by Boltzmann in the 19th century. I'm only exporting some nuggets of knowledge from physics to economics, where the importance of these 19th- and 20th-century insights apparently has not been fully appreciated. Actually, John Kelly implicitly used the same time arguments in the 1950s. [6] He didn't make them explicit, probably because they seemed too obvious to him, but that made his work very inaccessible to people with a different background.

I'll focus on two problems with the mean-variance framework. First, (the most commonly raised criticism), it only considers the mean and variance of returns. Second, (my main criticism) the framework relies on expectation values.

The mean-variance framework summarizes return distributions by their means and variances. The reason for doing this is the Gaussian central limit theorem: if I generate a lot of random numbers from some distribution, and it almost doesn't matter what distribution, then the sum of those random numbers will be Gaussian-distributed if I rescale it properly. So a lot of observations end up approximately Gaussian-distributed. "Gaussian-distributed" means that you know exactly what the distribution looks like if you have the mean and variance.

But of course not everything is Gaussian-distributed—for example it is impossible for something to be Gaussian-distributed if it is non-linearly related to something else that is Gaussian-distributed, and for those things the mean and variance don't contain all the information about the distribution.

There's also the problem that the Gaussian central limit theorem doesn't work in quite the same way for multiplicative processes (like the simple models of asset price dynamics).[7] That's the theoretical side. These concerns are practically relevant because whenever someone goes out and measures a real return distribution, he comes back with something non-Gaussian, with fatter tails. To summarize, the mean-variance framework, by construction, misses important information about the nature of the return distributions.

In contrast, time-average maximization (geometric mean for multiplicative dynamics) doesn't assume anything about the distributions. You stick in whatever distribution you like, crank the handle, and out comes your optimal investment strategy. It is much more general.

Second, my personal feeling is that everything I just said about non-Gaussian distributions is important but not as important as the conceptual failure of believing that expectation values have any sort of meaning in themselves. I prefer the term "ensemble average" to "expectation value" because it conveys more of the conceptual background. Expectation values are always an average over an ensemble of (usually imagined) systems.

I think the question about Gaussians can be a distraction, but much of the debate focuses on it because the far more basic tenets—relevance of ensemble averages—are not questioned any more. They are so engrained in the way we think that we've forgotten that they rely on the enormous assumptions of equilibrium and ergodicity. So what I'm about to say is about the mean-variance framework, but the real problem is the naive use of ensemble averages in this framework. Maybe we should call it the ensemble-average framework.

Ensemble averages do this funny thing: they get rid of fluctuations before the fluctuations really have any effect. That makes it very difficult to deal with risk because risk is often just another word for fluctuations. Let's say there's an investment whose expected rate of return is more than what I have to pay to borrow money from a bank. If I leverage the investment by borrowing money to invest, then the expected rate of return on my equity just grows and grows—the more I borrow, the better. That's the linear "more risk, more return" you mentioned. Something isn't right with that because we all know that if we borrow too much, we will be wiped out by fluctuations. So here's an instance where the mean-variance framework (or ensemble-average framework) misses an important message about risk. It doesn't naturally account for the effects of leverage.

The time-framework gives a completely different answer: if I leverage the investment I just mentioned, the time average growth rate will increase with leverage for a little while, then it will reach a maximum, and after that it will decrease: eventually, the more I borrow, the worse off I'll be. So time averages just spit out the concept of optimal leverage. Before, it was just "more is better," now there's an optimum.[8]

In the mean-variance framework there's a famous attempt to deal with the fact that we prefer smaller fluctuations: just divide any expected excess return by the volatility—that's the Sharpe ratio. It gets smaller if the fluctuations are bigger, so that's good. Still, when I first saw this ratio I thought it was a typo. Why? Because of the dimensions of this object (that's a fancy work for "units").

Volatility has the dimensions one-over-the-square-root-of-time, whereas excess expected return has the dimensions of one-over-time. So the ratio of the two is not dimensionless. This is probably a physics thing. Physicists are trained to look at the dimensions of any object, just as a sanity check. If something has dimensions, then it's not really fundamental, and the numerical value of that object cannot carry any relevant information because I can change the units, which changes the numerical value but of course doesn't change anything about the physics of the system.

The Sharpe ratio,[9] developed by the economist William Sharpe in the 1960s, also a recipient of the Nobel Prize in Economics, is usually given in units of "inverse square-roots of one year." So if I were to express time in minutes instead of years, the numerical value of the Sharpe ratio would be something completely different. Anyway, that was the first thing that struck me as suspicious. Quantities that have units are just not fundamental, but the Sharpe ratio was presented to me as a fundamental property of an investment. So I tried to understand how this ratio behaves.

For example, what happens to the Sharpe ratio when I leverage an investment? The answer is: nothing. The ratio stays the same. That makes it dangerous if we use it as the sole criterion for judging the quality of an investment. The investment can have a great Sharpe ratio, let's say 50 inverse square-roots of one year (the usual units), but be completely over-leveraged or under-leveraged, and I will either go bust quickly or not gain much from the investment.

From the time point of view, the Sharpe ratio seems very arbitrary. It's just something that becomes smaller if the fluctuations are larger or the excess return is smaller. But the time perspective would just look at optimal leverage to judge the quality of an investment. This is a dimensionless number, and its numerical value is very meaningful because it's actually an indicator of market stability, although that's another story. It's perfectly intuitive too—if an investment opportunity is great, I should leverage it, if it's really bad I should short it, which is just negative leverage. This extra bit of information also sheds more light on the Sharpe ratio. If you compare two investments, the one with the higher Sharpe ratio will have a better time-average return *if both investments are optimally leveraged.*

LTCM is now a classic example of over-leveraging gone wrong. I wouldn't be surprised if the culture of expectation values inherent in the mean-variance framework had contributed to missing the detrimental effects of excessive leverage.

MM: Ole, you've been very generous with your time, and we've covered a lot of ground. Are there any final thoughts you'd like to share that we didn't cover?

OP: I hope they won't be my final thoughts, but there are a few more things I'd like to comment on. We've focused, in this interview, on messages for individual investment decisions. But we're messing with the very basis, the conceptual foundation, of economics, so we can pick almost any area in economics and ask: what does this different way of thinking imply here?

Maybe not surprisingly—because risk is so naturally treated in this framework—we can learn a lot about market stability. I've worked on this aspect with Alex Adamou, and we now have some very interesting results, both theoretical and empirical. Markets don't just adjust prices but also price fluctuations, and this leads to a dynamic coupling between volatility, interest rates, leverage and optimal leverage. The basic statement is that investments that are so good that it's optimal to leverage them must be unstable long-term. This has very serious implications, for example, for the way we think about the housing market.

Perhaps more surprisingly, we've also learned a lot about economic inequality. This is work with Alex Adamou and with Bill Klein from Boston University. It turns out that these two different averages are very helpful in de-politicizing the very emotionally charged debate about how to share and distribute wealth in an economy. Essentially, the time-average rate of economic growth turns out to be the typical individual experience of economic circumstances, whereas the ensemble-average rate of economic growth is relevant to a central government that needs to estimate tax income. So this helps to get to the bottom of the conflict between individual and collective perceptions of economic circumstances. Without having to make any statements about whether inequality is desirable or not or what the right level of inequality is, we can study it and understand how it behaves.

Lastly, I want to say thank you to everyone who has contributed to this work, though it's impossible to name everyone. Many are in some way connected to the Santa Fe Institute.[10] This is not mainstream physics, it's not mainstream mathematics, and unfortunately it's not yet mainstream economics. So it's been absolutely essential to have a lot of supportive people around me. Brian Hoskins, the director of my institute at Imperial College has been immensely helpful, not just in discussions, but also by covering my back and letting me stretch my research remit beyond recognition while I was doing all this. Reuben Hersh, in the course of this project, has really helped me open my eyes to how creative and conceptually rich and vague mathematics is, and how important it is to keep this in mind in applications.

Murray Gell-Mann has been phenomenally helpful in many ways—walking through some of the arguments with me and, with his unparalleled experience in science, putting this in context. I've already mentioned Bill Klein and Alex Adamou, who are most actively involved at the moment, and their creativity and attention to detail is absolutely invaluable. It's something money can't buy—the breadth and depth of knowledge that comes from 40 years (sorry Bill) in mathematical physics. Brian Arthur made some very encouraging comments to me to the effect that it's completely normal for economics journals to be sluggish in recognizing a new successful direction of thought. Sam Bowles was very helpful in discussions about economic inequality. I could go on and on, but I'll finish by saying that it's not just academics who are helping to get this off the ground.

George Soros has shared his thoughts with me and put me in touch with his Institute for New Economic Thinking. The interactions I've had with you and others at Legg Mason Capital Management have been very informative and fun of course, and that applies to other businesses too, back in London, in New York and Scotland (you know who you are). A feeling for what's practically relevant is very valuable information in getting the big picture right. This is clearly interdisciplinary work, but I think it's more; it also transcends the divides between academia, policy-making, and private business, and I'm very happy about that. In the end it's not about succeeding in any specific realm but about figuring out how we can make things work for 7 billion people on the planet.

So thank you, Michael, for helping to get the word out. This interview has been a real pleasure.

MM: Thank you, Ole.

Endnotes

[1] "*a priori*" - *a priori* is a Latin phrase meaning "from the former".
[2] Paul Davidson, "Can future systemic financial risks be quantified? Ergodic vs nonergodic stochastic processes," *Brazilian Journal of Political Economy*, Vol. 29, No. 4, October-December 2009, 324-340.
[3] Daniel Bernoulli, "Exposition of a New Theory on the Measurement of Risk," *Econometrica*, Vol. 22, No. 1, January 1954, 23-36.
[4] Ole Peters, "The time resolution of the St. Petersburg paradox," *Philosophical Transactions of The Royal Society A*, Vol. 369, No. 1956, December 2011, 4913-4931.
[5] N. E. Romero, Q. D. Y. Ma, L. S. Liebovitch, C. T. Brown, and P. Ch. Ivanov, "Correlated walks down the Babylonian markets," *A Letters Journal Exploring the Frontiers of Physics,* April 2010.
[6] J. L. Kelly, Jr., "A New Interpretation of the Information Rate," *Bell System Technical Journal*, 1956, 917-926.
[7] Sidney Redner, "Random Multiplicative Processes: An Elementary Tutorial," *American Journal of Physics*, Vol. 58, No. 3, March 1990, 267-273.
[8] Ole Peters, "Optimal leverage from non-ergodicity," *Quantitative Finance*, Vol. 11, No. 11, November 2011, 1593-1602.
[9] Sharpe ratio - Sharpe ratio is a risk-adjusted measure of investment return. The higher the Sharpe ratio, the better the fund's historical risk-adjusted performance.
[10] Ole Peters, "On Time and Risk," *Santa Fe Institute Bulletin*, 2009.

Readers may also find Aaron Brown's new book, *Red-Blooded Risk: The Secret History of Wall Street* (John Wiley & Sons, 2011) useful, especially chapter 5, "When Harry Met Kelly."

Biography

Ole Peters received a Ph.D. in theoretical physics from Imperial College London in 2004. He then moved to the United States, where he held a joint fellowship at the Santa Fe Institute and the Center for Nonlinear Studies at Los Alamos National Laboratory. His work there focused on problems in statistical mechanics with applications to atmospheric physics, which later led him to join the Climate System Interactions group at UCLA. Following academic visits in Budapest, Beijing, and Hamburg, he returned to Imperial College in 2009 and is currently a member of the Mathematics Department and the Grantham Institute for Climate Change. His most recent work is concerned with the conceptualization of randomness in probability theory. In particular, he is interested in non-ergodic random systems whose behavior in time cannot be summarized by a probability distribution.
For more, see: http://tuvalu.santafe.edu/~ole/.

Important risks

All investments are subject to risks, including loss of principal.