Harvard Business Review

HBR Blog Network

Why Those Guys Won the Economics Nobels

by Justin Fox | 10:00 AM April 2, 2014

When the Riksbank Prizes in Economic Sciences (a.k.a. the economics Nobels) were announced last fall, the news was greeted with some confusion and amusement. The Swedes had given the award to one guy, Eugene Fama, who is best known for originating something called the efficient market hypothesis, another guy, Robert Shiller, who once called the efficient market hypothesis "one of the most remarkable errors in the history of economic thought," and a third guy, Lars Peter Hansen, whose work is so dense that even academic economists couldn't satisfactorily explain it or its connection to Fama and Shiller. The prizes were awarded "for their empirical analysis of asset prices," but what the three had been doing looked from the outside less like a common endeavor than a not-all-that-coherent argument.

It turns out, though, that there *is* significant common ground between the three winners. His name isJohn Campbell.

Campbell is an economics professor at Harvard and one of the most prominent figures in modern financial economics. He got his PhD at Yale under Shiller's supervision in 1984, but since then he has also done a lot of work expanding on Fama's ideas about risk and return, some of it co-authored with Fama's son-in-law and University of Chicago finance colleague, John Cochrane. Campbell's work has also made liberal use of the analytic tools developed by Hansen. In the long-version explanation of the prizes published by the Nobel committee last fall, Campbell was cited more often than anybody else, apart from the three winners.

Others, most notably money managers and former Fama students Cliff Asness and John Liew in an epic *Institutional Investor* article, have done a lot recently to clarify how Fama's ideas and Shiller's can at least co-exist peacefully. But even Asness and Liew threw up their hands when it came to Hansen. So I wanted to see if Campbell could make sense of the prizes and the current state of academic knowledge about asset prices. It took us a while to line up a time to talk, in part because Campbell was working on an article about the Nobels for the *Scandinavian Journal of Economics*and wanted to finish that first. Now a draft of the article is available on Campbell's website, and what follows is an edited transcript of a conversation I had with him Monday.

It's a quite technical article, and while our conversation was formula-free it got pretty wonky as well. Still, Campbell is a great explainer. So do not be discouraged to learn that we start with something called the stochastic discount factor, which Campbell describes as the central idea of modern asset-pricing theory.

Many lay readers are familiar with John Burr Williams and the dividend discount model, or the discounted value of future cash flows. This stochastic discount factor model is the modern economic update of those, correct?

Yes. There's fairly broad understanding of discounting when you don't have to worry about risk. You know, the future value of money, the present value of money — money today is worth more than in the future because you can invest it and get interest.

If that interest rate is just a number that's constant over time, a super-simple world, then we get the dividend-discount model with a constant discount rate. The tricky question is what to do about risk. We all know that there needs to be some adjustment for risk. If you have a risky proposition you don't want to discount the returns at the same rate that you would for a safe proposition. But how do you do that?

The naïve thing, which in certain circumstances can be right, is you still think of discounting as a single number but you adjust that number for the risk in the payoffs of this deal you're being offered. So if it's a very risky deal you say I'm going to discount using a high discount rate, and if it's a safe thing I'll discount at a lower rate. Back in the '60s, people developed the capital asset pricing model[CAPM] as a way to do that. You'd have this beta with the market, so you have the riskless rate plus beta times the equity premium. That's still widely taught in business-school classes and it's easy for people to understand. *[A mini-glossary: beta is the amount that an individual stock fluctuates relative to the overall stock market, and the equity premium is the difference in expected return between stocks and a "riskless" asset such as Treasury bonds.]*

Now it turns out that if you think a little more deeply about this, it's really not right. Instead, what you need to do is scenario analysis. There's the good scenario where everything works out and your investment makes lots of money. There's a bad scenario where it doesn't work out and you end up losing money. What you should do is take each possible scenario and discount that scenario at a rate appropriate to the scenario. Then at the end, when you've brought everything to the present, scenario by scenario, you average. That's called stochastic discounting, because the discount rate that you use is different for every scenario, and thus in a certain sense it's random, or stochastic.

That's kind of a deep insight. It's not something that necessarily resonates a lot with people in the markets or people in the world. It is used in the more technical end of the financial industry, for example in derivatives valuation. A lot of people in what I'll call the lower-tech part of the financial industry, equity analysts or traditional fundamental investors, they don't think about this.

But it is a very basic rewrite of how to think about finance, and it has swept all before it in the academic world. It goes back to the 1970s. Steve Ross, who was at Yale and now is at MIT, he didthe basic theory. But it's taken some time, decades really, to work out all of the ramifications. And that's what these guys [the 2013 Nobel winners] did. They took this notion of the stochastic discount factor and turned it into something empirically useful.

The idea of the "joint hypothesis" that Eugene Fama framed in the late 1960s and early 1970s was that there had been all this research showing that markets reacted really quickly to information and that professional investors didn't beat the market. Fama said that if you want to say something more about how efficient markets are, how good a job they do of pricing assets, you need some theory of what asset prices should be. And the theory that was available then was CAPM.

One can distinguish between what we call time series models and cross-sectional models. Time series models say how these risk adjustments move over time, and cross-sectional models say how they vary across different assets at a point in time. Now, the world has both, so at the end of the day we need a model that tells us everything. But academics like other people tend to want to work on one thing at a time. So when Fama started, the cross-sectional model was the CAPM, and for the time series people just said well let's just assume that risks don't change over time, so whatever risk adjustment there is it's the same today, tomorrow, and thereafter.

And in those early tests, it seemed like market prices mostly obeyed both CAPM and the efficient market hypothesis.

The early tests that looked at the stock market and looked at short periods of time generally found pretty decent results consistent with the market efficiency insight. There was one exception even back then, post-earnings-announcement drift, which Ball and Brown pointed out in the late 1960s. If a stock has just had surprisingly good earnings, it jumps up but then it tends to keep on going up, and that's not consistent with market efficiency.

A lot of other stuff that looked at stocks and looked at short periods of time didn't find much. Financial economists got very cocky in the '70s. They said the market's efficient, discount rates are constant, we know what's going on, we have this theory, it's great. But then as is so often the case when people become too cocky, cracks appear in the structure.

There were several types of cracks. One was that when you looked outside the stock market, when you looked at interest rates, you found predictabilities. My first published paper with Shiller back in 1983 was on that, and I did my PhD dissertation on that. And then Fama published results on it. So that was one problem — fixed income. Currencies was another. What we now call the carry trade, that currencies with high interest rates give you excess returns. That was another discovery of the 1980s that was inconsistent with the paradigm.

So that's problem No. 1. Problem No. 2 is this distinction between short-term vs. long-term predictability. This is what Shiller made his name on, and I later helped him with it. His research agenda was to say that if stock prices are just dividends discounted at a constant rate, then they can't be more volatile than the stream of dividends that they're supposed to be forecasting. So it's mysterious why market prices are moving around so much. Put another way, it turns out that ratios like the dividend-price ratio or a smoothed-earnings price ratio don't forecast future dividends at all well, but they do forecast future returns. When you look at the market extremes, whether it's 1929 or 1966 or 2000, those very high prices relative to dividends and earnings are not followed by rapid dividend growth. They tend to be followed by declining prices.

Shiller hammered away on this point in the '80s, and in fact Fama also published some of the same observations. He didn't want a behavioral explanation. His view was that, "Oh, we're learning that the risk premium over long periods of time actually does move around."

To me that's one of the most interesting things about this whole story. Fama had a hypothesis and he went out and tested it, and found that it didn't really work. As did other people. But there are these two classes of explanations for why it didn't work. One is that markets are still pretty efficient but risk premia change over time, and the other is that the explanations have to do with behavior. Reading your paper it does feel a little bit like, *you say tomahto, I say tomayto.*

The economics profession is still struggling with the balance between these things. Most people, if you get them in a room and give them a truth serum will say that there's some mix of the two. My view, for what it's worth, is that with some phenomena like the long swings in the market and the premium for value stocks you can get quite a long way with rational explanations. People just get more cautious in recessions than in booms. I have a well-known paper with John Cochrane which argues that this is because people judge their well-being relative to their past experience — the standard of living to which they have become accustomed, as the divorce courts would say.

So in a time like the 1960s, when there's been a lot of growth, people are feeling rich and they're willing to take risks because they've got a cushion of comfort above their baseline expectation. At a time like the present when things have not been so great, people's standard of living is much closer to the baseline minimum that they expect, and they don't feel like they have a big cushion of comfort.

That's a model in which people have reasonable expectations about the future, they just worry about risk a lot more in bad times than in good times. So my view is that we don't necessarily have to have irrational beliefs to explain these long swings in the market, although I think Bob is absolutely right that some people do have irrational extrapolative beliefs, and believe the hype when the market goes up.

Even if it's not irrational to base your judgments on recent experience, it doesn't feel like the rational man of '60s and '70s rational-expectations economics. It feels like it's got a little bit of Kahneman and Tversky in it.

I think it's fair to say that even those economists who play the rational expectations game have in more recent years written down models of people who may have rational beliefs, but are emotionally volatile. That story I told in the model with Cochrane is of an emotionally volatile rational guy who gets into a funk and is very cautious all of sudden and then a few years later is very aggressive. That's very different in spirit from the 1960s and 1970s constant-discounting guy. The strong distinction between rational economics and behavioral economics, I think it can be overblown. There's a border zone where a lot of the literature is.

Now if I can just rewind a little bit, I was saying that in the sort-of heyday in the '70s there were these cracks in the structure, and the first one was fixed income and currencies, the second one was the long-run behavior of stock prices. The third was anomalies in the cross-section of stock returns. Fama and Ken French said the rewards in the market don't just come from beta as the CAPM would have it. There's also a size factor and a value factor. What Mark Carhart later did was add a momentum factor. *[Mini-glossary: The size factor means that small-cap stocks outperform bigger companies. The value factor means that cheap stocks, as measured by price-to-book, price-to-earnings, or some other such ratio, outperform expensive ones. Momentum means that once a stock heads up or down, it tends to keep going in that direction.]*

In my view, rational models have plenty of ways to explain the value premium, but the phenomenon of momentum is really very hard for a purely rational model to explain. If there's one thing that should make Gene Fama uncomfortable, it's probably momentum. The behavioral story about momentum is that a lot of people aren't paying enough attention to fundamental news, so there's money to be made by, whenever you see prices go up, jumping on it and driving them up more.

We should get back to Hansen. You've got this wonderful line in your paper about meeting him and "sensing that his penetrating insight would require effort to fully understand but would amply reward the undertaking."

Lots of distinguished economists have had this experience of either reading a Lars Hansen paper or listening to a Lars Hansen presentation and feeling that it's a sort of message from the future — like an alien artifact dropped from a flying saucer. That is, potentially amazing technology if you can only figure out how it works. Lars is famous for that.

But at this point a lot of his work can be translated, and is widely used and understood. Gene Fama said we can't test market efficiency unless we have this auxiliary model. Hansen said, well, hang on a minute. Suppose we believe that the market is efficient if we have the right model, but suppose that this right model that we have in mind has some unknown parameters that come from the impatience of investors or the risk aversion of investors or other features of the world — and we want to know what these parameters are. Well, presumably the right parameters are the ones that, when plugged into the model, make returns as unpredictable as possible. Because we know that if the market's efficient and we have the right model, then returns are unpredictable.

You've explained that in a way that I almost understand it. Is that the GMM [generalized method of moments]?

That's the GMM.

It sounds useful.

It's extremely useful. It's sort of the standard method that any of us use when we come up with some model. We say oh, that's a nice model, how are we going to see if it works and what these unknown parameters are? It's a universal tool, and it's very, very important.

So that's the first thing about Lars. The other place where I bring him in is next to the discussion of behavioral finance, because Lars has also worked on models of in-a-way-irrational beliefs. Although what Lars does sort of blurs the distinction between rational and irrational beliefs.

He draws on an engineering literature called robust optimal control. Suppose you're an engineer and you're building a bridge. You have some physical understanding of the way in which the traffic is going to vibrate the bridge and cause stresses on the supports, but there's a range of possible models of how much vibration the trucks are going to make. So what you might do as an engineer is try to build your bridge to be safe in a worst case. You can't literally take the ultimate worst case, because your bridge would be infinitely expensive. You're going to have to take a worst reasonable case. The discipline of engineering has evolved to do precisely that. They have rule-of-thumb methods, but they also in recent years have developed a more mathematical, abstract approach called robust optimal control. And Lars has taken some of these ideas and applied them in finance.

He's saying you don't know how good the return on the stock market is going to be. There's a range of pundits, they all seem to say different things. Maybe you want to invest in a way that will give you good results even in a worst reasonable case. So if you were inclined to put a lot of money in the stock market you say no, I won't do that because the equity premium might be very low and then I'm taking a lot of risk and not getting much reward. So perhaps I'll invest as if the equity premium is lower. He takes this perspective of almost deliberately choosing to have beliefs that are pessimistic to protect yourself in case things go wrong. These beliefs aren't rational in the literal sense, but nor are they crazy. He uses this word "robust": they're defensible beliefs that have this degree of pessimistic conservatism built into them.

They sound pretty, well, rational.

It's another way in which this stuff blurs the distinction between rational and irrational. In a way it's surprisingly close to Shiller, but Shiller's view is that people can't know the true model and then become excessively influenced by social forces, and can get into a herd mentality. In Lars's world people don't have the true model, but they know that they don't have the true model and they react by being very conservative. Lars's guys are irrational in a level-headed way, and Bob Shiller's guys are irrational in a way that is subject to these social fads and fashions.

It seems like the clearest practical lessons from this academic work have been in asset management. There are people very much steeped in this work, including you, who are out there managing stocks and other assets based on it. The sense I've gotten, and I've talked to Cliff Asness a lot, is that markets are pretty efficient but there are all these little things going on that if you're careful about it you can take advantage of.

The late, great Paul Samuelson once talked about micro efficiency and macro inefficiency. What he meant was that if the inefficiencies you're looking for are relative mispricings of different types of stocks, that will be corrected fairly soon, because if they got big there's a lot of easy money to be made. Cliff Asness would be all over it, and the mispricing would disappear in no time. Whereas if you're talking about macro things, the big, long swings in the market, the fact that stocks were so expensive in the 2000s and so cheap in the fall of 2008 is very difficult to arbitrage away. You have to be a macro hedge fund, and not only do you the hedge fund manager have to have nerves of steel, your clients have to have nerves of steel.

A theme of Bob's work is that these long swings over time in the market can have big effects on prices and be large inefficiencies. Whereas when it's arbitraging away small price discrepancies across similar assets at a point in time, that can be done easily so you're likely to find that the deviations are small.

But because that's less dangerous to do, that's the direction that most quantitative money managers that come out of the academy actually take, right?

Absolutely. A lot of my work is on big asset allocation themes and things that play out over many years. But when I've spent time in the industry it's with a quant equity firm that's trying to pick stocks and beat an index and show results in a reasonable period of time to clients. What we often find is that asset allocation is something the clients themselves want to manage. The results take years to play out and it's hard to set up a contract with an asset manager to hire them to do this because it takes so long to see if they're right or not.

Another area where these debates have some resonance is in policy —monetary policy, financial regulation and the like. There it seems like this macro inefficiency is actually pretty important.

Absolutely. As we think about the stability of the financial system, large swings in asset prices and big changes in risk premia can be very important, and relatively hard to arbitrage away.

It seems like the lesson that came out of the '60s and '70s is that more finance is better. More people out there trying to arbitrage away all these inefficiencies will make markets more efficient.

At a micro level sure, that's probably right. But at a macro level maybe they make the swings bigger.

There's a very meaningful debate about that now. If we come back to the Nobel guys, one of the interesting things about Shiller is that despite being famous for his work on bubbles, he doesn't say let's shut it down. Instead what he says is let's have financial innovation that is actually helpful. His vision of financial innovation is that by designing the right instruments, you could help people be more rational, because you could focus their attention on the things that matter.

One simple example would be if you take a stock and break it into claims to dividends at different parts of the future — the near term and then the next year and then the year after that and so on — and you trade them all separately. It forces people to recognize that the value of the stock should be the value of all these claims to dividends in particular years. If you want to pay a fortune for this stock, you have to recognize that you're paying a fortune for a claim to cash that will be paid in some particular future year. It kind of focuses your belief about when the ship is going to come in. The pot of gold can no longer be just at the end of the rainbow. You have to say where it is exactly. And that might help people be more rational.

The capital asset pricing model was supposed to allow companies to calculate their cost of capital in a consistent way. Now that nobody seems to think that risk premia stay the same over time or that beta really does reflect everything that matters about risk in the stock market, it seems like there is no one way to calculate the cost of the capital anymore. But everybody still uses the method that came out of CAPM.

You're right. It's a weakness of modern finance that we haven't been able to deliver something that is as useful as the CAPM. I remember there was a Fama and French paper with the wonderful title, "The CAPM Is Wanted, Dead or Alive." Which basically says, well, we need a model, and even if this model is in some sense dead, it's still wanted and still used.

I think we're groping our way towards better procedures. It's like the CAPM is an aging champion and there are all these wannabes that would like to replace it, and none of them have quite come to the fore yet. I have an entry in the competition, an intertemporal model in which I break the market movements into permanent shocks driven by cash flows and temporary shocks driven by discount rates, and I show that one of them should have a higher price of risk than the other. It's like in the old days cholesterol was the measure of risk, and now we know there's good cholesterol and bad cholesterol. So in that framework what do is you calculate the beta of your firm or your project with two components of the market return, and one of them is the one that you really worry about. It's a relatively small change in the procedure, but I have some hopes that in the long run it might prevail. But these changes in thinking take a very long time.

JUSTIN FOX
Justin Fox is Executive Editor, New York, of the Harvard Business Review Group and author of *The Myth of the Rational Market*. Follow him on Twitter @foxjust.