Words From the Wise

Robert Engle

Rob Engle, a 2003 recipient of the Nobel Prize in Economics, recently sat down with Antti Ilmanen, Lasse H. Pedersen and Rodney N. Sullivan of AQR to discuss his research and its many uses in addressing important challenges facing investors. This is part of an ongoing series of "Words From the Wise" interviews published on AQR.com. Following is an executive summary along with the full interview with Dr. Engle.

Robert Engle is the Michael Armellino Professor of Finance at New York University Stern School of Business. He was awarded the 2003 Nobel Prize in Economic Sciences for his research on methods for analyzing economic time series with time-varying volatility. His now famous autoregressive conditional heteroskedasticity (ARCH) model and its many generalizations have become indispensable tools for academic researchers and investment practitioners alike. He is the Director of the NYU Stern Volatility Institute and many of his methods are now featured in the innovative public web site, V-Lab, where daily estimates of volatilities and correlations for more than a thousand assets can be found. Professor Engle is a co-founding president of the Society for Financial Econometrics (SoFiE), a global non-profit organization housed at NYU. Before joining NYU Stern in 2000, he was Chancellor's Associates Professor and Economics Department Chair at the University of California, San Diego and before that he was an Associate Professor of Economics at MIT. He is a member of the National Academy of Science. He received his Bachelor of Science from Williams College and his MS in Physics and PhD in Economics from Cornell University.

AQR Capital Management, LLC
Two Greenwich Plaza
Greenwich, CT 06830

p: +1.203.742.3600
f: +1.203.742.3100
w: aqr.com

Executive Summary

In financial markets and in the economy, estimating risk — how much prices and other economic variables will fluctuate over time — is critically important. Indeed, the tradeoff between risk and return is central to Modern Portfolio Theory, and risk plays a crucial role in valuing share prices and option prices, among other financial instruments. Risk estimates also have broad implications for the macroeconomy, for instance in estimating systemic risk associated with financial crises, monetary policy, inflation, and so on. As we know well, economic and market risk can vary dramatically over time, fluctuating between calm periods and turbulent periods. In this interview, we discuss with Nobel Laureate Rob Engle his breakthrough research on analyzing economic time series with time-varying volatility. He explains his concept of autoregressive conditional heteroskedasticity (ARCH) and its generalizations and how they have become indispensable tools for both researchers and investors in the statistical modeling of time-varying volatility. Here, he discusses using ARCH-type models in estimating volatility and returns in the short-term and long-term as well as their limitations in doing so. Along the way, we hear about his favorite variation of ARCH used in such applications. We then delve into the underlying drivers of risk and the interesting inverse relationship between risk and returns.

The conversation then turns to a discussion on how volatility estimation has important applications to risk management and options pricing and how the Volatility Laboratory, or V-Lab, which was created at NYU Stern School of Business under the direction of Dr. Engle, has grown since its inception from a rather basic laboratory for studying volatility to a highly sophisticated systemic risk measurement and monitoring tool for the global economy.

We then hear from Dr. Engle how receiving the Nobel Prize changed his life and how he maintains a healthy work-life balance, something he believes is very important to ensuring a high quality of life. We conclude by learning about his biggest regret and his mentors and heroes.

AQR

The Father of Risk Modeling: Lessons for Investment Practice

ESTIMATING MARKET VOLATILITY

Sullivan: You've mentioned the scientific method and how this relates to theory and models building on prior results. How does your work in risk and volatility build on the work of Markowitz[1] and others that laid the initial foundations for the theory of finance?

Engle: Well, those models certainly motivated my research. In fact when I looked at these early models, I said, "Volatility's assumed constant, yet when I measure it I get lots of different answers at different times." So, really, what I've done is to take the big themes of finance — the biggest being the tradeoff between risk and return — and ask, "How do you build on these models in order to better incorporate volatility?" Success requires ways of measuring risk that's not just constant. If risk is time-varying, then we need models that are also time-varying, and that was my contribution.

Ilmanen: Moving from the theory to practice, can you discuss why investors find, for instance, your ARCH model so useful?

Engle: Well, if you want to measure risk, then you have to have some tools to say where it is now and what it might be in the future. For instance, if you think there is a risk premium associated with volatility, which we all think must be true, then you have to be able to forecast it to some degree. If you can't forecast it, how can you possibly charge a premium to people to take this risk? So while I don't think that individual investors are ever going to be running ARCH models to look at their portfolios, I do think institutional investors are all managing risk in very sophisticated ways, using models that have been improved greatly since the ARCH model. This is a natural direction, especially given how hard it is to measure expected returns. It means that volatility is particularly useful in investment management.

Ilmanen: Can you discuss for our readers what ARCH, and more broadly, the whole family of ARCH related models, means for estimating volatility?

Engle: Well, ARCH is an acronym which stands for autoregressive conditional heteroskedasticity. In brief, it is a method that seeks to forecast what the volatility of some particular time series would be, using only the past of that same time series. What it observes is that high volatility — that is, large moves in either direction — tends to predict that future volatility will continue to be high. So whenever the stock market has a big move — whether it's up or down — there is a better-than-average chance that it's going to have a big move the next day. It might be in the same direction or it might be in the opposite direction. We call this volatility clustering in the data.

Pedersen: What does it say about mean reversion?

Engle: The first thing you see is the clustering. Then over longer periods, you find that when the cluster ends, the data mean tends to revert back to some long-term average level. The cluster might be one of low volatility, and that lasts for a while and then it ends, and the volatility goes back up again and reverts back to the average. That's the prediction from these ARCH-type models. It says if volatility is high today, it's likely to be high tomorrow. But if you go further out in time, it's probably going to go back down toward normal times.

[1] For more, see Words from the Wise: Harry Markowitz on Portfolio Theory and Practice (2016).

I think that this is the most useful aspect of the ARCH model; that it's possible to calibrate it just by using the past series and treat it as an econometric problem, so that if you have a better idea of what the model ought to look like, you can revise it as needed.

Sullivan: On the flip side, what are the limitations of an ARCH framework?

Engle: Well, first of all, it assumes that the dynamic time shape of the volatility clusters and that the mean reversion is relatively stable over time. I think that's a reasonable reflection of what we see in the real world. This class of models typically uses only information on the past returns of the same series. So, although it uses only a limited information set, for instance, it doesn't look at other asset classes, it does amazingly well in predicting volatility. We find that it's hard to make a model much better by bringing in other information, because the other information is hard to use.

Of course, there are other ways to forecast volatility, for instance there are implied volatility models used by options traders. We have compared these implied volatility forecasts with ARCH volatility forecasts, and find that the implied volatility forecasts typically have more information in them, but they also have some inherent biases. So I think the most important limitation of the ARCH models is they lack full information.

Ilmanen: A main empirical result of implied volatility forecasts is that they tend to be biased?

Engle: There is an enormous literature on this, and it's very confusing because it comes up with lots of different answers, and I think that the different answers are more related to the methodology of evaluation than to the actual models and methods used. My reading is that if you could do a good, statistically sound, head-to-head comparison between implied volatility and a GARCH (adding a "G" for generalized) model, maybe over a very short forecast horizon, you'd find the GARCH model was better. But for intermediate terms, the implied volatility model is better. So, even though implied volatility has biases, it just has more information in it.

Pedersen: Yes, the implied volatility models seem remarkably informative.

Ilmanen: Yes, for example, they can embed information about upcoming earnings announcement dates and other eventful information such as highly volatile days.

Engle: Yes, those could be incorporated, but it's complicated to incorporate them into a GARCH model. You have to be able to identify some events in the past that you think are going to be like the ones in the future so that you can model them, whereas an options trader doesn't have to do that.

Ilmanen: The ARCH family of models is a big family, do you have a favorite among them?

Engle: It is indeed a big family, and it's getting bigger all the time. My colleague Tim Bollerslev has a glossary of ARCH-style models which has some 50 entries.

More importantly, I think that's a lot like naming regression models, which we no longer try to do. For instance, each time you have a new dataset, you might have a different set of regressors in a regression model, and we don't really need to name these models. What we're looking for is the best predictor, and besides we use the same type of likelihood function almost all the time. Some of these predictors are high-frequency predictors and some are low-frequency predictors, and we might have univariate predictors and multivariate predictors. It's an open class of models, so there are always going to be more, but you just don't bother naming them anymore.

Ilmanen: Broadly speaking then, do you have a favorite?

Engle: The model that I use most often is what is typically called the GJR model, or the Glosten, Jagannathan, and Runkle (1993) model. It's a model which allows for asymmetry, and so it incorporates the fact that negative returns typically have more predictability for future volatility than positive returns. The EGARCH model is very close to that in terms of its performance and fit. I think there are some cases where one is a little better than the other. It's not a big difference.

Ilmanen: Are you applying this to equities?

Engle: Yes, I'm answering this particularly for equities, but in our V-Lab, we apply this to lots of asset classes. And even in asset classes where you don't expect any asymmetry, you tend to find it in the data. Take interest rates and exchange rates, for example, which often have some asymmetry in the volatility.

Ilmanen: How about for longer horizon forecasts?

Engle: Another point that I think I should mention in terms of the weaknesses of ARCH models is that they're not particularly good for long-term forecasts. They tend to look a lot like longer-term implied volatilities in that they mean revert fairly systematically for longer horizons. Also, the path of mean reversion is not all that informative.

So you might ask, are there other ways to forecast long-run volatilities? Well, I think it's much more likely that you can make progress by building a structural model of what's really going on in, say, this sector or this economy or whatever. But I think a simple univariate volatility model like EGARCH is not likely to have that kind of information in it.

Ilmanen: Can we say anything about the underlying fundamental reasons for the observed large time variation in volatility?

Engle: The gold standard for volatility is to understand the economics of why volatility is whatever it is. My favorite interpretation is that the primary driver of volatility is a response of the markets to new information. How fast the news comes in and how important the news is on any particular asset or asset class. So it's the news process itself that's fundamentally driving volatility. So when we measure volatility, I think we're really measuring the intensity of the news flow.

Sullivan: So then we need to better understand what drives the intensity of the news flow?

Engle: Yes, although I and a few others are looking at this, we haven't done nearly enough work on it. Our work was focused on a cross-country model which looked at what makes the financial markets in one country more volatile than another, and we came up with some interesting macroeconomic findings. We found that when inflation is high, volatility tends to be high, and when the economy is slow, volatility again tends to be higher. When the economy is more volatile — that is, one quarter it goes up, and the next quarter it goes down — financial markets are more volatile. We also found that when interest rates jump around, the equity market is more volatile.

So there are some fairly natural, intuitive connections between the macroeconomy and volatility and it's nice to be able to make a quantitative connection.

Pedersen: Has there been any research that attempts to link news flow, say from newspapers articles, to volatility in prices?

Engle: We (Engle, Hansen, Lunde (2012)) have been working on that for a couple of years now. Our paper counts news stories or words in news stories that include a ticker, and analyzes whether that helps explain the volatility in prices. The answer is yes, but of course, it's hard to know for sure whether the news stories are reporting the volatility or whether they're causing it. But the interesting thing about it is that the news stories affect volatility not only on the day they're reported, but also several days after that. So it seems like the full information in the news story is not immediately incorporated in the volatility. There's some dynamic that is going on in the markets, which maybe is price discovery, although often we don't think price discovery takes quite that long.

So I think there's some interesting things to be said there, and there are a lot of people looking at textual analysis in news stories and so forth. But mostly, they're trying to predict the market return, not the market volatility.

Ilmanen: Some get this idea in the wrong direction, that traders cause volatility.

Engle: That's right. Many believe that the volatility is just something that's manufactured by a lot of beliefs and whatever is currently going on in the marketplace. I think that once you've connected market volatility with news and information, it's much easier for typical investors to understand why markets are so volatile. If Janet Yellen says they're unlikely to raise interest rates and the market goes up, well, there's a good story there that makes sense.

Sullivan: Fischer Black reportedly said that as markets become more efficient, they become more volatile. Do you agree with that? I take it that he meant that as news arrives, it's being incorporated more rapidly into market prices thus making markets more volatile.

Engle: I very much respect Fischer's opinion on anything, including this, but I have a different opinion. I think at least the speed at which information is received in the markets doesn't really affect the daily volatility as measured end-of-day. It may affect which trader gets to capitalize on this volatility and make money on it. But the fact that you see the price move within 10 seconds of a news announcement versus two hours is not going to show up in end-of-the-day volatility. It might show up if you measure second-by-second volatility. But I don't really think it changes where the markets go from day to day.

It might be that price discovery takes three days, in which case day-to-day volatility has changed. But my view is that price discovery probably is on a pretty high-frequency basis, especially now when trading is on such a high-frequency basis. So I don't think that the speed of news shows up very much in the way we measure daily volatility.

And after all, it does not seem that the trend of market volatility is generically rising. Of course it went up during the financial crisis, but then it came back down again. So I don't see any secular drift in volatility, which is what you'd expect if the speed of information drove volatility.

Ilmanen: Black may have meant that in markets with stale prices (say with corporate bonds of the past versus today), volatility tends to be underestimated; then, one can certainly see that you do get greater volatility as markets become more liquid and efficient.

Engle: Black was big on the role of noise traders as slowing down the market. If everybody was given the same information, then prices would adjust immediately to the new values.

RELATIONSHIP BETWEEN MARKET VOLATILITY AND RETURNS

Sullivan: You alluded earlier to the idea that predicting volatility is easier than predicting returns. Why do you think this is the case?

Engle: I think one of the reasons volatility is easier to predict than returns comes from the fact that when you look at a sample path, as time goes on you get more and more information about volatility, but not about returns. In fact, one of the paradoxes of modern finance is that these diffusion processes we use have the uncomfortable implication that you actually don't have to estimate volatility; you just record it by using high-frequency data. So we are getting more and more information about volatility all the time, but it doesn't give us any more information about expected returns, because you need a long span of data, and then there has to be no structural change, everything has to be constant over this long period, and all that's probably not true. Consequently, there is no hope of getting fairly precise estimates of expected returns.

Ilmanen: What about the contemporaneous relation?

Engle: Somewhat surprisingly, volatility and returns are negatively correlated, something that quite a few people have written about. The two general hypotheses seeking to explain it are the leverage effect and the risk aversion effect. It turns out that the leverage effect certainly is not sufficient to explain all the asymmetry in volatility, partly because, for example, many firms don't have any debt. So there's no leverage effect at all in their fundamental asset prices, yet we still see negative correlation between volatility and their stock prices.

I think a more important observation is that the negative correlation is strongest in broad indices, which is exactly where investors are most worried about volatility because that's systematic volatility, and it's unavoidable. So if there is a belief that future volatility is going to be high, coming from either autocorrelation in the volatility clustering or from other information, you see a very strong effect of higher volatility when the whole market is falling sharply. I take that pattern as evidence of risk aversion, which quite possibly is getting stronger as the data seems to indicate.

Sullivan: Do you believe that the risk aversion effect is more likely to be driving the negative volatility/return correlation?

Engle: Yes, but I don't think that means that there is no leverage effect. I just think that the risk aversion effect is the most powerful contributor, but we also see asymmetry due to leverage.

Emil Siriwardane and I have a paper called "Structural GARCH" (2015) that I presented at AQR in 2015, which builds the leverage effect exactly into stock prices. We think about stock prices as being an option on the value of the firm, then what's important is the volatility of an option price, not the volatility of a fundamental asset. Thus, the moneyness of this option matters and so we estimate the moneyness of the option and the elasticity, or the delta elasticity of it. Therefore what we do is basically de-lever the stock returns, and we find that there's still asymmetry in returns after accounting for all this.

Ilmanen: Has the asymmetry in returns become more pronounced since 2008 or did the rise begin even earlier?

Engle: What I have noticed is that the asymmetric volatility models that we estimate, even using all the historical data, have stronger asymmetry now than they used to. But it's not quite clear when the rise started exactly. Something to look into more closely.

Ilmanen: Sometimes practitioners use simpler modeling approaches such as a rolling average in estimating historical volatilities. How do you compare these simpler methods to ARCH and GARCH model approaches?

Engle: I think there are some important criticisms of using historical volatilities to be aware of. One is that such models contain no information that's older than the window width that you have to work with. The second is that they treat each observation within the window as equally important, whereas if we think that these models are better at short-term forecasting than long-term forecasting, you ought to give more weight to the more recent observations than the older ones. And the third thing is that it has no criterion function for choosing the best estimation window.

Having said that, rolling historical volatilities are a kind of first generation ARCH model. So while the historical volatility approach gets it roughly right, it doesn't quite have the statistical model set out right and for instance can give some erroneous information related to the tail part of the distribution. But once you get the statistical model set out right, with an ARCH-type framework, you can get a good picture of volatility.

BROADER APPLICATIONS OF MARKET VOLATILITY ESTIMATES

Sullivan: Earlier you mentioned how volatility estimation is important to risk management and options pricing. Let's begin with options pricing; can you discuss the role of volatility estimation in options pricing?

Engle: To figure out what an option should be worth, of course, we need to know quite a bit about the underlying price process including its volatility, its distribution and so forth. Over the years, some have sought to use volatility estimates from GARCH models as an input into the Black-Scholes model. However, it didn't prove all that useful. The main reason, I think, is that if the GARCH model is really descriptive of the underlying return process, then Black-Scholes is not the options pricing formula that you want to use, because Black-Scholes assumes volatility is constant over the life of the option.

So this leads us to more sophisticated options pricing models. Suppose your favorite GARCH model really does describe well what's going on in the underlying return process. Then what does your option pricing formula look like? You can now answer that question because if the GARCH model does provide the data generation process for the risk-neutral distribution, then you can then simulate it and price the options.

I, along with two colleagues from Lugano and Lausanne (Barone-Adesi, Engle and Mancini (2008)), took this approach with S&P Index options. We used a GARCH model to price all the different strikes for short-maturity, middle-maturity and long-maturity options, all using the same four-parameter GARCH model — some 250 different options contracts — and it gives quite good prices. Now, this doesn't mean that it is the true underlying process, but it does mean that it's fairly consistent with an underlying process for what the risk-neutral distribution ought to look like. It gives you fat-tails, flatter smiles for the longer maturities, and all the things that we know about from looking at implied volatility pictures. So, yes, you can get all that out of a risk-neutral asymmetric GARCH distribution.

Ilmanen: Can you discuss how you can estimate both asymmetries and fat tails at both short-horizon and longer-horizon?

Engle: I should say that our GARCH model uses bootstrap shocks from the historical distribution, so-called filtered historical simulation. So they're fat-tailed shocks with a little bit of negative skewedness in them already. And then you pick the parameters to do the best job of fitting all these options prices. So it is very computationally intensive; you're fitting a lot of option prices with a small number of parameters.

Ilmanen: Very interesting. Let's now move on to discuss how estimating volatility helps with risk management.

Engle: I think that is actually one of the important applications of GARCH models. And whether it's exactly a GARCH model, or some other approach used by various software providers doing something quite similar, I think it provides a systematic and excellent way to assess risk in a large number of assets, factors, and so forth over the short run.

Pedersen: Maybe you have heard this before, but some people will say, "I buy something cheap and yes it may be volatile. The price may even drop below where I bought it. But as long as I don't sell it, and as long as I remain confident that ultimately it will return me more money than I paid for it, then volatility is not really risk." Such investors believe their capital is not permanently impaired. Have you thoughts on such comments that volatility is not a relevant measure of risk?

Engle: I can understand why one might say that, but it's unclear whether one can really wait out the downturns. Even if so, it certainly doesn't mean that volatility is irrelevant. It just means that your decision process — not to sell an asset that falls in price, or at least hang on to it — ignores short-term volatility. So I think volatility is still useful whether it's synonymous with risk or not. In this case, one might believe that a little bit of volatility isn't bad, but a lot of negative return might be enough to have a margin call, and then it's not so benign anymore.

With that said, I am sympathetic with the idea that short-run volatility is not the only measure of risk. A key feature of the financial crisis, I think, was that short-run volatility was very low while longer-run volatility was high. A lot of investments were made as if the short-run volatility was going to stay low forever. But if you're going to invest in an asset that you're going to hold for a long period of time, you shouldn't use short-term volatility as a measure of risk. Instead, you should use a longer-run measure of volatility, because that's what this asset is going to face eventually when you liquidate it.

I should mention that we're not so good at knowing what long-run volatility really is. A big weakness in risk management these days is that we don't have good tools for looking at longer-run volatilities or longer-run risks. The main tool that we have is using scenario analysis, but we have a hard time associating them with probabilities. So we don't exactly know whether this is a risk that you want to accept or avoid because you don't know what its probability is.

One of the research projects we've taken on with the Volatility Institute relates to the regulatory environment for insurance companies. Instead of stress testing insurance companies, what is being recommended under the EU's Solvency II is to not do just a single stress test, but do thousands of stress tests. So you have many scenarios which are equally likely, and you simulate what you think the whole financial landscape is going to look like. So insurance companies can give thought to what their strategy — their portfolio or business model — might look when we simulate these paths out for 50 years. You can look at the ones with a low P&L, and then those are the ones to worry about in considering whether the company is going to be able to weather the storm. This type of work is a frontier that's really important in determining how best to build models to assess long-run risks.

Pedersen: Are companies now being pushed towards more conservative investments by regulators with a primary focus on short-term risks?

Engle: Yes. I think that's why this type of analysis has been adopted for insurance companies. A lot of the risk management for insurance companies has been looking at liability risk because unlike banks, their liabilities are more stochastic. Yet the real risks of the insurance business have to do with assets and liabilities. If you have short-run risks on one side and long-run risks on the other side, well, there are a lot of bad things that can happen.

Sullivan: Can you discuss V-Lab and the Volatility Institute and what you hope to accomplish through them?

Engle: The Volatility Institute was founded in 2009 here at the Stern School with a goal to leverage academics' and practitioners' knowledge so that we can better understand risks and volatility in financial markets and to help the broader public better understand these risks.

One very successful part of this effort has been the Volatility Laboratory, or V-Lab, which has grown from a primitive laboratory for studying volatility to a highly sophisticated risk measurement and monitoring tool for the whole global economy. We have a powerful website to help disseminate this information.

The original idea was to take all these different volatility models that we were discussing just now and apply them to thousands of assets and update them without a human hand quite regularly (daily or weekly) and try to figure out which models were the most successful.

We started with volatility and then built correlation models. We have models that estimate correlations between thousands of assets. This is quite an undertaking. Then, when the financial crisis came, we developed a systemic risk model. It's like a stress test based on market prices for more than 1,000 financial institutions around the world, updated once a week. All this information is published on the Internet. It's very cool.

More recently, we've developed sections for market liquidity and fixed income based on this risk model scenario generator that we talked about earlier. Finally, we have a long-run risk measure, which essentially asks, "What's the value at risk, not today, but a year from today?" This is our first entry into answering the question we just discussed, "How do you estimate long-run risk?"

Ilmanen: And what approach did you take in measuring market liquidity?

Engle: We used the so-called Amihud (2002) ratio, which is the absolute return on a day of some particular asset divided by the dollar volume that's traded on that day. The ratio can be interpreted as a price impact, it answers the question, "How much volume was traded and how much did the price move?" The bigger the ratio, the more illiquid is the stock. It has the great advantage that it doesn't need bid-ask quotes, it doesn't need price impact, you don't need to know trade directions and a lot of things which are quite complicated.

So, it’s a nice measure. Our innovation is to use it to build a forecasting model for the time series of illiquidity. We now have a series of liquidity forecasting models for thousands of assets and people are finding it interesting, I think.

Ilmanen: Back to your earlier point, what about the long-run risks?

Engle: In thinking about this, we see that there's another risk factor that you have to take into account — the risk that the risk will change. So if you think the risk will change, you have to ask yourself, how much can it change? And how do you estimate that? Well, that's what a volatility model does; it measures how much the risk changes. So if you simulate your favorite GARCH model, say 1,000 times, you get a histogram of what the future outcomes might be, and you can calculate the value at risk from this histogram.

So obviously, the value at risk is greater when the volatility is greater, but because volatility is mean-reverting, it may not stay higher. And similarly, if volatility is low today, because it's mean-reverting, the forecast over the next year will show it rising.

Sullivan: What sort of interest have organizations such as the Systemic Risk Council and the Treasury shown in this work?

Engle: We have done a lot of work with systemic risk. Our measure, we call it SRISK, is an estimate of the amount of capital a firm would have to raise in order to continue to function normally given another financial crisis. So if you're a regulator, you would hope the answer to this question is zero, that this firm would not need to raise capital if we have another financial crisis. So it is a useful framework from the point of view of the regulatory process.

However, the approach that we use employs stock market valuations of these firms, and asking how much the firm’s stock market valuation will likely deteriorate in a financial crisis. Unfortunately, regulators seem to have a mandate to use only detailed supervisory data on each of these financial institutions to estimate what their value would be under the counterfactual assumptions of a proposed negative event scenario. This makes for a very time-consuming and complicated job, because there are so many different kinds of assets on the books. For instance, there are both simple and complex loans and there are bank subsidiaries and brokerages, and so on.

So what is interesting is that the questions are really quite similar, and at least in the US, our rankings are pretty similar to what the Fed comes up with. For the global economy, we have a list of the 30 most systemic banks in the world, which has about a 75% overlap with the Bank of International Settlement’s (BIS) list of the systemically most important banks. So again, it's pretty similar, but not the same. On the other hand, the ECB did a stress test of capital adequacy of banks, which completely disagrees with our results. The major difference being that we think that the major banks in Europe need to strengthen their balance sheets and the ECB is more focused on smaller banks across southern Europe and Cyprus. Overall, by having our information publicly available on the Internet we are a gadfly in pushing regulators to do a better job.

By the way, it’s important to mention that the theory behind all this was developed here at Stern, and Lasse is one of the coauthors on the research. The researchers are Acharya, Pedersen, Richardson, and Philippon (2016) who did the primary work behind this, and then with V-Lab we applied the econometrics work.

Ilmanen: So after the global financial crisis you and other academics decided to focus attention on important practical financial issues?

Engle: Yes, I would say it was a brand-new opportunity to start research in a new area where there was almost nothing and it was viewed as tremendously important. That is a gold mine for an academic, it’s exactly what we want.

Ilmanen: Are the V-Lab systemic risk models useful for the early detection of financial crises?

Engle: Our risk measure for banks depends on three things: the size, leverage and risk of the bank. Early detection basically requires finding some combination of these variables that provides meaningful risk estimation.

Prior to the financial crisis, it was hard for us to detect it because the stock market loved the banks and so bank capitalizations were fairly high, which means that leverage did not look so high for them as shown in the debt-to-equity ratio. So, because we use the market value of equity, we were not able to see the bust before the regulators saw it. So the model didn't give us a two-year advanced warning that this was a problem, but we have the data to show that it measured it in an important way; we can see the ranking of systemic risk of financial institutions before the crisis which is very close to the ranking of institutions that got into trouble.

Ilmanen: How important are correlations when it comes to forecasting systemic risks?

Engle: Our bank risk measure is based on both correlations and volatilities; more specifically their beta. We use beta, because say the global market collapses; the question is, what's that do to a bank? The answer to that question lies in the company's beta. We estimate the beta using dynamic conditional correlations, volatilities of the bank stock and volatilities of the market as a whole. We call this dynamic conditional beta, measured with a GARCH-style model. So correlations are a key part of this, but they interact with the leverage and the size in an important way.

Ilmanen: Can you discuss your financial consulting work during your career?

Engle: In the beginning of my time exploring finance I did some consulting for Salomon Brothers. It was very exciting for me. I loved my time there. I worked with Joe Mezrich and Eric Sorensen and we built volatility and correlation models and used them in developing derivative strategies. It was very exciting for me to see how all these statistical techniques could be implemented in practice.

Later, I moved to New York and began to consult for Morgan Stanley. At that time, I turned from doing volatility research to work on market microstructure. I developed some new statistical approaches that we used to measure market liquidity. I also worked on how best to incorporate transaction costs into trading platforms, both pre-trade and post-trade. All this work was for Morgan's equity trading lab, a group that was doing algorithmic trading and really enjoyed the challenge.

Ilmanen: How does time-varying market liquidity affect execution strategies?

Engle: One of the key things about understanding execution costs is that there is a risk-return tradeoff between time and execution. You're going to have to pay somebody to take the other side of the trade, and you want to get the best price you possibly can, so you're going to have to let the trade run for some period of time,which means taking some risk. So, there's always a choice to be made.

THE EARLY YEARS

Sullivan: What drew you to finance and what were the challenges at the time that inspired you?

Engle: I came to finance late in my career and was actually not so much consciously drawn into it as pulled into it. I had this success with developing the ARCH model, which I thought of as a model for macroeconomics. I kept getting invited to finance meetings to talk about volatility because even though it was originally designed for a macro problem, it actually didn't have any great applications in macro, whereas in finance, it was very popular and useful. So I gradually started going to more and more finance meetings and was fascinated by all the applications of volatility analysis and multivariate volatility analysis to finance. So it was the solution to practical problems that interested me more than actually being interested in finance itself.

Ilmanen: If I recall correctly, the initial application for ARCH was for estimating inflation?

Engle: Yes, that's right.

Pedersen: Going back even further, can you discuss what drew you into economics even before finance?

Engle: Yes, I went to graduate school at Cornell to study physics, and I spent one year in the basement of Rockefeller Hall of Cornell University studying low-temperature physics — superconductivity and so forth — which was a pretty exciting field at that time. I took classes with some famous physicists, including Hans Bethe. Yet it seemed like if I were successful, that it would only be a handful of people in the world who would be aware of my work. I wasn't likely going to have the kind of impact on society that I was hoping for. So I looked around for the most quantitative of the social sciences, and it was economics.

Pedersen: At the time, did you expect that some of the math that you learned in physics would be directly applicable to economics and finance?

Engle: I had thought that would be the case. I had various friends who would say, "I think quantum mechanics might help with estimating economic cycles." For instance, is it a wave or a particle? But I never found, really, that any of those specific methods were useful.

What is useful, I think, is the approach. Physics, like all sciences, is about the scientific method. You have hypotheses, you have experiments, you get data, you analyze the data, and you infer something about the world we live in. Well, that's what econometrics does. Econometrics takes the data and tries to make inferences about the world. That approach came from my physics background, even though the specific methodologies didn't translate very well.

Pedersen: So some of us financial economists, me being one, have physics envy in the sense that in social sciences the theories are not as stable and clear as they are in physics. How do you feel about that difference?

Engle: Well, I do think that's true. And physicists, who tend to have a clear view of the superiority of physics, will be quick to reinforce your physics envy. It would be nice if there were laws of economics that were as strong as they are in natural sciences. But here we're talking about people, and while that makes for interesting problems, I think it's a bit much to expect that we can be as precise as physicists.

THE NOBEL AND OTHER ACHIEVEMENTS (AND THE SECRET TO HAPPINESS)

Sullivan: How did receiving the Nobel Prize impact your life?

Engle: When they called me, the Nobel Committee said this now famous phrase, "Once we hang up from this telephone call, your life will never be the same again." At least in my case, this was true. It completely changed my life. I decided to use the opportunity as a way of furthering my research. I took the time to review the volatility research that had been done since I stopped working in volatility research, especially in the prior decade and decided that it wasn't done yet and that there were some important things that no one had done. One of them, as I mentioned earlier, being the economic underpinnings of volatility.

So I dove back into volatility research with a renewed sense of interest in it, and that's why I established the Volatility Institute. That is not to say that I'm not interested in market microstructure or other things, but volatility is a fundamental part of asset pricing, it's important and useful. By launching the Volatility Institute and in particular, the V-Lab, I can discuss issues through the lens of this research, and it changes every day. It is great tool for public speaking, and it's great for passing information on to people. And I find that it's very exciting. It's one of the things I love to do. And every now and then I'll say, "I have a great idea for something new we could put in here," or a new paper. Now I have a vehicle for delivering it that is different from writing a new paper, and that's wonderful.

So it has really changed my life. I travel all over the world. I get invited to meetings. I do interesting consulting work that I might not have done before. I recommend it highly.

Pedersen: You had a good life before, too.

Engle: Yes, I had a good life before the prize. I was not at all unhappy. I didn't feel any need for this, but I'm just delighted to have it.

Pedersen: Did it ever cross your mind that this could happen?

Engle: Well, I knew that Clive Granger, the co-recipient of the 2003 Nobel Prize, had a good chance for it, and I had nominated him several years before. So, I thought he had a good chance because co-integration was such a big topic. Although I co-authored the famous co-integration paper (1987) with him, it was really his original idea, so I didn't think I would share in the prize with him as we did in 2003, as you know.

Pedersen: But you also had a number of papers with citations that were off the charts.

Engle: Yes. Well, that's true. The ARCH and co-integration papers were very influential, and I think, actually, that is the reason why the Nobel Committee was willing to go to econometrics, which is not a field that they typically consider supporting. But they did so because both of these methodologies were widely used by practitioners. The co-integration is used throughout macroeconomic analysis. All the people that were building vector autoregressive models switched them to vector error correction models so that they could keep the co-integration effects working properly. Meanwhile, everyone who was doing risk management wanted to know about ARCH models. So there was this big community that considered these substantive contributions. At least that's my interpretation of it.

Ilmanen: What other achievements are you most proud of?

Engle: There may be two things in economics that I should mention. We talked about the Volatility Institute and V-Lab, of course I'm very proud of them and I feel they have an enduring presence.

The second one is the Society for Financial Econometrics, SoFiE, which Eric Ghysels and I started in 2008. Next year we're going to have a big meeting here in New York to celebrate the 10 years of SoFiE. It's had conferences that have gone around the globe twice, and now it's coming back for a third year. So those are both academic accomplishments that I'm very proud of.

I also have an outside-of-academic accomplishment that I'm very pleased with, my figure skating. I'm a competitive ice dancer. My dance partner, Amy, and I have competed in the US championships for three or four years in an event that I never used to do, which involves things like twizzles and lifts, among other things that I've only learned recently, although I've been ice skating since my 20s. Last year we competed in the world championships in Oberstdorf, and we placed second in two of the events. That was really fabulous.

Pedersen: I've known you for years and you seem very happy; what is your secret to eternal youth and happiness?

Engle: Well, I believe in a balanced life. I think it's really important to match your research and your job with your family life and with your outside activities so that they all have an important, even if not equal, role. I think that does keep you from getting too stressed out in any one field.

But I do get stressed out. And I find over and over again, I will agree to do something, and it might be to give a talk or turn in a paper, or it might be to go to a skating competition, or it might be to do something with my family that stresses me out. But then when it's over, it usually turns out well, and it just gives you this really nice high from doing what you want to do and getting it done and having it come out well. I do think that does keep you young.

Pedersen: What role does your wife, who is an expert in psychology, play?

Engle: Yes, she's a child psychologist, and she's also a sports psychologist. She does two things that are quite helpful. One is she calms me down when I get upset and really is very supportive, and often she'll come along to skating competitions and helps me relax and not be too nervous. But she also pushes. She says, "You really need to go out and get this done." So she's very committed to my career, but also to her career and wants to see everything be very successful. She doesn't feel like we're competing for time. She wants it both.

Sullivan: How about regrets, things you missed?

Engle: Well, my first job was teaching at MIT in 1969 in the economics department. It turns out that probably 100 yards away, Fischer Black and Myron Scholes were teaching in the Sloan School and they were talking about options pricing. Franco Modigliani was also there, and he was talking about the capital asset pricing model. Also, Bob Merton had just finished his Ph.D. from Paul Samuelson and so he was around a lot. So the greats of the early days of finance, particularly quantitative finance, were all around me, but I didn't pay any attention to them. Unfortunately, it took me 20 years before I came to finance.

Sullivan: We've all learned so much from you; who are your mentors and heroes?

Engle: Well, the first thing I want to say is that some of my mentors are my students. I've had lots of fabulous students, too many to name. Although I start off by teaching them, they end up teaching me which is great.

Other people I'd like to mention would be Jim Hamilton and Hal White from San Diego, who were terrific. And, of course, Clive Granger, my long-time colleague and friend who's passed away, and Hal White has also passed away. Some others I'd include are Joel Hasbrouck, Alain Monfort and Viral Acharya. I get a tremendous amount of stimulation from talking with them. And Lasse, when you were here at Stern, it was great.

Sullivan: You should come around more often, Lasse!

Engle: Yes, absolutely!

Sullivan: Thank you, Rob, for sharing of your time and wisdom with us. We look forward to seeing more of your research, and are grateful for your many and lasting contributions to our profession.

References

Amihud, Yakov, 2002, "Illiquidity and Stock Returns: Cross-Section and Time-Series Effects." *Journal of Financial Markets*. Vol 5, pp.31-56.

Acharya, Viral, Lasse Heje Pedersen, Thomas Philippon, and Matt Richardson, 2016. "Measuring Systemic Risk." *The Review of Financial Studies,* forthcoming.

Barone-Adesi, Giovanni, Robert Engle, and Loriano Mancini, 2008, "A GARCH Option Pricing Model with Filtered Historical Simulation." *The Review of Financial Studies*. Vol 21, Issue 3 (May): pp.1223-1258.

Engle, Robert, Martin K. Hansen and Asger Lunde, 2012, "And Now, The Rest of the News: Volatility and Firm Specific News Arrival." CREATES Research Papers (March).

Engle, Robert and Clive Granger, 1987, "Co-integration and Error Correction: Representation, Estimation and Testing." Econometrica. Vol 55. pp.251-276.

Markowitz, Harry with Antti Ilmanen and Rodney N. Sullivan, 2016, "Words from the Wise: Harry Markowitz on Portfolio Theory and Practice." AQR Capital Management, LLC.

Glosten, Lawrence R., Ravi Jagannathan, and David E. Runkle, 1993, "On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks." *The Journal of Finance*. Vol 48, No 5 (December): pp.1779-1801.

Engle, Robert and Emil Siriwardane, 2015, "Structural GARCH: The Volatility Leverage Connection." Working paper (October).

Disclosures

AQR Capital Management, LLC
Two Greenwich Plaza, Greenwich, CT 06830
p: +1.203.742.3600 I f: +1.203.742.3100 I w: aqr.com