Academics vs. Operators: Why the EMT Comes Crashing Down

This article is authored by Massimo Fuggetta, MOI Global instructor and director of Bayes Investments. The article includes two posts from Massimo's [Bayes blog](#).

As he wrote his [*Challenge to Judgment*](#) on the first issue of the Journal of Portfolio Management in 1974, Paul Samuelson expected 'the world of practical operators' and 'the new world of academics' – which at the time looked to him 'still light-years apart' – to show some degree of convergence in the future.

On the face of it, he was right. The JPM recently celebrated its [40th anniversary](#). The Financial Analysts Journal, started with the same bridging intent [30 years earlier](#) under Ben Graham's auspices, is alive and well on its 73rd Volume. Dozens other periodicals have joined in the effort and hundreds of books and manuals have been written, sharing the purpose of promoting and developing a common language connecting the practice and the theory of investing.

Academic vs. Operator

While presuming and pretending to understand each other, the two worlds are still largely immersed in a sea of miscommunication. At the base of the Babel there are two divergent perspectives on the relationship between risk and return. Everybody understands return. You buy a stock at 100 and the price goes up to 110 – that's a 10% return. But this is *ex post*. What was your *expected* return before you bought? And what risk did you assume? The practical operator does not have precise answers to these questions. I looked at the company – he would say – studied its business, read its balance sheet, talked to the managers, did my discounted cash flow valuation and concluded that the company was worth more than 100 per share. So I expected to earn a good return over time, roughly equal to the gap between my intrinsic value estimate and the purchase price. As for risk, I knew my valuation could be wrong – the company might be worth less than I thought. And even if I was right at the time of purchase, the company and my investment might have taken wrong turns in myriads different ways, causing me to lose some or all of my money.

Is that it? – says the academic – is that all you can say? Of course not – replies the operator – I could elaborate. But I couldn't do it any better than Ben Graham: read his books and you'll get all the answers.

But the academic would have none of it. As Eugene Fama recalls: 'Without being there one can't imagine what finance was like before formal asset pricing models. For example, at Chicago and elsewhere, investment courses were about security analysis: how to pick undervalued stocks'. (My Life in Finance, p. 14). Go figure. Typically [confusing science with precision](#), the academic is not satisfied until he can squeeze concepts into formulas and insights into numbers. I don't know what to do with Graham's rhetoric – he says – I need measurement. So let me repeat my questions: what was your expected return *exactly*? How did you *quantify* your risk?

Give me a break – says the defiant operator – risk is much too complex to be reduced to a number. As for my expected return, I told you it is the gap between value and price, but I am under no illusion that I know it exactly. All I know is that the gap is large enough and I am prepared to wait until it closes.

Tut-tut – Fama shakes his head – Listen to me, you waffly retrograde. I will teach you the CAPM. 'The

CAPM provides the first precise definition of risk and how it drives expected return, until then vague and sloppy concepts'. (p. 15).

Risk vs. Return

The operator listens attentively and in the end says: Sorry, I think the CAPM is wrong. First, you measure risk as the standard deviation of past returns. You do it because it gives you a number, but I think it makes little sense. Second, you say the higher the risk the higher is the expected return. That makes even less sense. My idea of risk is that the more there is the more uncertain I am about my expected return. In my view, the relationship between risk and expected return is, if anything, negative. So thank you for the lecture, but I stick with Graham. As Keynes did *not* say ([again!](#)): It is better to be vaguely right than precisely wrong.

Writing ten years after Samuelson's piece, Warren Buffett well expressed the chasm between academics and practical operators: 'Our Graham & Dodd investors, needless to say, do not discuss beta, the capital asset pricing model or covariance in returns among securities. These are not subjects of any interest to them. In fact, most of them would have difficulty defining those terms. The investors simply focus on two variables: price and value'. (Buffett, Superinvestors, p. 7).

But operators are rarely so blunt. Such is the intellectual authority of the Efficient Market Theory that the identification of risk with the standard deviation of returns – a.k.a. volatility – and the implication that more risk means higher returns are taken for granted and unthinkingly applied to all sorts of financial models. Hilariously, these include the same valuations that investment practitioners employ to justify their stock selection – an activity that makes sense only if one rejects the EMT! It is pure schizophrenia: investors [unlearn at work](#) what they learned at school, while at the same time continuing to use many of the constructs of the rejected theory and failing to notice the inconsistency.

Beta and CAPM: Unethical?

But here is the biggest irony: after teaching it for forty years – twenty after Buffett's piece – Fama finally got it out of his system: 'The attraction of the CAPM is that it offers powerful and intuitively pleasing predictions about how to measure risk and the relation between expected return and risk. Unfortunately, the empirical record of the model is poor – poor enough to invalidate the way it is used in applications ([Fama and French, JEP 2004](#)). Hallelujah. Never mind that in the meantime the finance world – academics and practitioners – had amassed a colossal quantity of such applications and drawn an immeasurable variety of invalid conclusions. But what is truly mindboggling is that, *in spite of it all*, the CAPM is still regularly taught and widely applied. It is hard to disagree with [Pablo Fernandez](#) – a valiant academic whose work brings much needed clarity amidst the finance Babel – when he calls this state of affairs unethical:

> "If, for any reason, a person teaches that Beta and CAPM explain something and he knows that they do not explain anything, such a person is lying. To lie is not ethical. If the person 'believes' that Beta and CAPM explain something, his 'belief' is due to ignorance (he has not studied enough, he has not done enough calculations, he just repeats what he heard to others…). For a professor, it is not ethical to teach about a subject that he does not know enough about."

Two books that I think are particularly effective in helping operators move from practical unlearning – erratic, undigested and incoherent – to proper *intellectual* unlearning of the concept of risk embedded in the EMT and its derivations are David Dreman's [Contrarian Investment Strategies: the Next Generation](#) (particularly Chapter 14: What is Risk?) and Howard Marks' [The Most Important Thing](#)

(particularly Chapters 5-7 on Understanding, Recognizing and Controlling Risk).

The Example of Amazon: Volatility = Risk?

Besides the EMT's predominance, unlearning is necessary because, at first glance, measuring risk with the standard deviation of returns makes intuitive sense: the more prices 'fluctuate' and 'vibrate', the higher the risk. Take Amazon:

Amazon.com Stock Price (logarithmic scale, 1997-2017)

Source: Massimo Fuggetta, Bayes blog.

If you had invested 30,000 dollars in Amazon's IPO in May 1997 (it came out at 18 dollars, equivalent to 1.50 dollars after three splits), after twenty years – as the stock price reached 1,000 dollars (on 2^{nd} June this year, to be precise) – your investment would have been be worth 20 million dollars. Everybody understands return. But look at the chart – in log scale to give a graphic sense of what was going on: 1.50 went to 16 in a year (+126% in one month – June 1998) to reach 85 in November 1999. Then in less than two years – by September 2001 – it was down to 6, only to climb back to 53 at the end of 2003, down to 27 in July 2006, up to 89 in October 2007, down to 43 in November 2008 and finally up – up up up – to 1000. Who – apart from [Rip van Winkle](#) and Jeff Bezos – would have had the stomach to withstand such infernal rollercoaster?

So yes, in a broad sense, volatility carries risk. The more violent the price fluctuations, the higher is the probability that, for a variety of psychological and financial circumstances – he may get scared and give up on his conviction, or he may need to liquidate at the wrong time – an investor might experience a catastrophic loss. But how can such probability be measured? The routine, automatic answer is: the standard deviation of returns. Here is the picture:

Cumulative Standard Deviation and *Twelve-Month Rolling Standard Deviation*

Source: Massimo Fuggetta, Bayes blog.

The graph on the left is the cumulative standard deviation of monthly returns from May 1997 (allowing for an initial 12-month data accumulation) to May 2017, for Amazon and for the S&P500 index. The graph on the right shows the 12-month rolling standard deviations. The cumulative graph, which uses the maximal amount of data, shows that while the monthly standard deviation of the S&P500 has been stable at around 5%, Amazon's standard deviation has been, after an initial peak, steadily declining ever since, although it still remains about four times that of the index (18.4% vs. 4.4%). The 12-month

rolling version shows a similar gap, with Amazon's standard deviation currently about three times that of the S&P500 (5.1% vs. 1.8%).

What does this mean? Why is it relevant? What can such information tell us about the probability that, if we buy Amazon today, we may incur a big loss in the future? A moment's thought gives us the answer: very little. Clearly, today's Amazon is a completely different entity compared to its early days in the '90s. Using any data from back then to guide today's investment decision is nothing short of mindless. Amazon today is *not* four times as risky as the market, as it wasn't five times as risky in November 2008. Nor is it three times as risky, as implied by the 12-month rolling data. The obvious point is that the standard deviation of returns is a backward-looking, time-dependent and virtually meaningless number, which, contrary to the precision it pretends to convey, has only the vaguest relation to anything resembling what it purports to measure.

The same is true for the other CAPM-based, but still commonly used measure of risk: beta. Here is Amazon's beta versus the S&P500 index, again cumulative and on a 12-month rolling basis:

Amazon.com Beta vs. S&P 500

Source: Massimo Fuggetta, Bayes blog.

Again, the cumulative graph shows that Amazon's beta has always been high, though it has halved over time from 4 to 2. So is Amazon a high beta stock? Not according to the 12-month rolling measure, which today is 0.4 – Amazon is less risky than the market! – but has been all over the place in the past, from as high as 6.7 in 2007 to as low as -0.4 in 2009. Longer rolling measures give a similar picture. What does it mean? Again, very little. According to the CAPM, Amazon's beta is supposed to be a constant or at least stable coefficient, measuring the stock's sensitivity to general market movements. But in reality it is nothing of the kind: like the standard deviation of returns, beta is just an erratic, retrospective and ultimately insignificant number.

Higher Risk = Higher Return?

Volatility implies risk. But reducing risk to volatility is wrong, ill-conceived and *in itself* risky, as it inspires the second leg of the CAPM misconception: the positive relationship between risk and expected return. 'Be brave, don't worry about the rollercoaster – you'll be fine in the end and you'll get a premium. The more risk you are willing to bear, the higher the risk premium you will earn.' Another moment's reflection is hardly necessary to reveal the foolishness – and to commiserate the untold damage – of such misguided line of reasoning. The operator's common sense view is correct: once risk is properly defined as the probability of a substantial and permanent loss of capital, the more risk there is the *lower* – not the higher – is the probability-weighted expected return. This also requires unlearning – often, alas, the hard way.

Despite Samuelson's best wishes, then, there is far less authentic common ground between operators and academics than what is pretended – in more or less good faith – in both camps. Operators are right: there is much more to risk than volatility and beta, and actual risk earns no premium.

Investment risk is the probability of a substantial and permanent loss of capital. We buy a stock at 100 expecting to earn a return, consisting of appreciation and possibly a stream of dividends. But our expectation may be disappointed: the price may go down rather than up and we may decide to sell the stock at a loss, either because we need the money or because we come to realise, rightly or wrongly, that we made a mistake and the stock will never reach our expected level.

How does investment risk relate to volatility – the standard deviation of past returns, measuring the extent to which returns have been fluctuating and vibrating around their mean? Clearly, we prefer appreciation to be as quick and smooth as possible. If our expected price level is, say, 150, we would like the stock to reach the target in a straight line rather than through a tortuous rollercoaster. On the other hand, if we are confident that the price will get there eventually, we – unimpressionable grownups – may well endure the volatility. In fact, if on its way to 150 the price dropped to 70 it would create an inviting opportunity to buy more.

Volatility increases investment risk only insofar as it manages to undermine our confidence. We might have rightly believed that Amazon was a great investment at 85 dollars in November 1999, but by the time it reached 6 two years later our conviction would have been brutally battered. Was there any indication at the time that the stock could have had such a precipitous drop? Sure, the price had been gyrating wildly until then, up 21% in November, down 12% in October, up 29% in September and 24% August, down 20% in July, and so on. The standard deviation of monthly returns since the IPO had been 33%, compared to 5% for the S&P500, suggesting that further and possibly more extreme gyrations were to be expected. But to a confident investor that only meant: tighten your seatbelt and enjoy the ride. A 93% nosedive, however, was something else – more than enough to break the steeliest nerves and crush the most assured resolve. 'I must be wrong, I'm out of here' is an all too human reaction in such circumstances.

Therefore, while volatility may well contribute to raise investment risk, it is not *the same as* investment risk. It is only when – rightly or wrongly – conviction is overwhelmed by doubt and poise surrenders to anxiety that investment risk bears its bitter fruit.

Amazon is a dramatic example, but this is true in general. Every investment is made in the expectation of making a return, *together with* a more or less conscious and explicit awareness that it may turn out to be a flop. Every investor knows this, in practice. So why do many of them ignore it in theory and keep using financial models built on the axiom that volatility equals investment risk? As we have seen, the reason is the intellectual dominance of the Efficient Market Theory.

What Prevents Academics from Seeing the Obvious?

Why is it that, according to the EMT, investment risk coincides with volatility? The answer is as simple as it is unappreciated. Let's see.

If the EMT could be summarised in one sentence, it would be: The market price is right. Prices are always where they should be. Amazon at 85, 6 or 1000 dollars. The Nasdaq at 5000, 1400 or 6400. At each point in time, prices incorporate all available information about expected profits, returns and discount rates. Prices are never too high or too low, except with hindsight. Therefore, an investor who buys a stock at 100 because he thinks it is worth 150 is fooling himself. Nobody can beat the market. If the market is pricing the stock at 100, then that's what it's worth. The price will change if and only if *new* information – unknown and unknowable beforehand and therefore not yet incorporated into the current price – prompts the market to revise its valuation. As this was true in the past as it is true in the present and will be true in the future, past price changes must also have been caused by no other reason than the arrival of information that was new at the time and unknown until then. Thus all price changes are unknowable and, by definition, unexpected. And since price changes are the largest

components of returns – the other being dividends, which can typically be anticipated to some extent – we must conclude that past returns are largely unexpected. At this point there is only one last step: to identify risk with the unexpected. If we define investment risk as anything that could happen to the stock price that is not already incorporated into its current level, then the volatility of past returns can be taken as its accurate measure.

Identifying investment risk with volatility *presupposes* market efficiency. This is part of what Eugene Fama calls the *joint hypothesis problem*. To be an active investor, thus rejecting the EMT in practice, while at the same time using financial models based on the identification of investment risk with volatility, thus assuming the EMT in theory, is a glaring but largely unnoticed inconsistency.

The False Assumption That Crashes the EMT Edifice: Common Priors

What is it that practitioners know and makes them behave as active investors, and EMT academics ignore and leads them to declare active investment an [impossible waste of time](#) and to advocate passive investment?

Again, the answer is simple but out of sight. In a nutshell: Practitioners know by ample experience that investors have different priors. EMT academics assume, by theoretical convenience, that investors have common priors.

Different priors is the overarching theme of the entire Bayes blog. People can and do reach different conclusions based on the same evidence because they [interpret evidence](#) based on different prior beliefs. This is blatantly obvious everywhere, including financial markets, where, based on the same information, some investors love Amazon and some other short it. In the hyperuranian realm of the EMT, on the other hand, investors have common priors and therefore, when faced with common knowledge, cannot but reach the same conclusion. As [Robert Aumann](#) famously demonstrated, *they cannot agree to disagree*. This is why, in EMT parlance, prices reflect all available information.

Take the assumption away and the whole EMT edifice comes tumbling down. This is what Paul Samuelson was referring to in the final paragraphs on the Fluctuate and Vibrate papers. More explicitly, here is how Jonathan Ingersoll put it in his magisterial [Theory of Financial Decision Making](#), immediately after 'proving' the EMT:

> In fact, the entire "common knowledge" assumption is "hidden" in the presumption that investors have a common prior. If investors did not have a common prior, then their expectations conditional on the public information would not necessarily be the same. In other words, the public information would properly also be subscripted as φ_k – not because the information differs across investors, but because its interpretation does.
>
> In this case the proof breaks down. (p. 81).

Interestingly, on a personal note, I first made the above quotation in my [D.Phil. thesis](#) (p. 132). A nice circle back to the origin of my intellectual journey.