Heisenberg* uncertainty principle…or

Why You Shouldn't Trust Most Financial Research

The standard admonition must be repeated, however: it is not known what the nature of that relationship is. For example, a numerical relationship such as that depicted in Figure 1 does not establish a causal relationship. It is, nevertheless, common practice in finance research to ignore this admonition by saying that the factor "explains" the rates of return.

Figure 1 shows only a simple linear regression – one set of rates of return against one factor. It is also possible to run a multiple regression – one set of rates of return against more than one factor. Imagine three or more columns of numbers lined up next to each other. It is harder to visualize than the simple graph in Figure 1 because the least-squares fit becomes a plane instead of a line.

If you line up those columns in Microsoft Excel and highlight them, then select from the Data menu the Analysis option (if you have activated it) and select "Regression" – Presto! You'll have run the regression in about two seconds.

The point is that if you had the numbers for many factors lined up in your spreadsheets, you could run hundreds of regressions in a single day. And that's if you just used Excel. If you wrote a program to do it, you could run many more. More people in finance than you may realize do that.

Hypothesis testing

The idea of lining up the factors with the one-month rates of return is to formulate a hypothesis that the rates of return depend in some way on the factors – and maybe can even be predicted by the factors, the Holy Grail of the investment world.

As many or most readers will know from statistics classes – but may have forgotten – there is a well-defined procedure for testing such hypotheses.

The procedure is as follows. You start by assuming that the rates of return *don't* depend on the factors; that is, that the betas are zero. This is known as the null hypothesis.

But since – as can be seen from Figure 1 – the data are spread out in a random pattern because of unknown factors, they might accidentally fall in a recognizable pattern. They might accidentally make it look like a coefficient of the regression – a beta, a slope of the regression line – isn't zero, even if it really is.

There are ways to estimate the probability that this could happen. The steeper the slope of the fitted line, the less likely it could have happened; the less spread out in a random cloud around the fitted line the points are, again the less likely it could happen.

The probability it could have happened at random – even if there is no relationship, that is, if the null hypothesis is true – is known as the *p-value*.

Conventionally, if this p-value is less than 0.05 (5 percent), it is common to say that *the test rejected the null hypothesis at the 0.05 significance level*. That allows the researcher to say that the null hypothesis was rejected in favor of the alternative hypothesis – that the rates of return did indeed depend on the factor – and that this result was significant.

Multiple testing

Now, remember that the p-value measures the probability that the slope of the regression line – the beta – could have come out looking steep because of randomness. If the p-value were 0.05, or one-twentieth, that means that if you conducted 20 tests, in all of which the beta was really zero, then on average, in one of the tests, you would get a p-value of 0.05 or less. In that case you would have proclaimed the null hypothesis rejected at the 0.05 significance level.

So what happens if researchers are conducting hundreds or thousands of tests? For each hundred tests in which there is in reality no relationship between the factors and the rates of return, they will declare that in five tests they found a significant relationship.

Harvey et al.'s paper addresses this problem. They go through the procedures that have been developed in the field of statistics to deal with what is called the multiple comparisons or multiple testing problem.

In each procedure, the way to deal with the problem is simply to raise the bar, making it more difficult to declare a null hypothesis rejected. The solution is not to allow a researcher to declare a result significant if the p-value was only 0.05 or less. The p-value might have to be 0.005 or 0.0005 or even 0.0000005 before it can be declared significant. It depends on how many tests are being done by researchers everywhere.

The problem is that we don't know how many tests are being done because researchers rarely if ever disclose this in their papers. Harvey et al. own that they can only review what has been published in the literature, and there are undoubtedly many tests that have been carried out that have not been published. In fact, published literature is biased in favor of publishing research in which the p-value is low. Tests in which the p-value does not meet the test of significance don't see the light of day.

And as we have seen, such tests could number in the thousands, even in the tens of thousands. For example, in an interview with Stanford economist Russ Roberts, Harvey pointed to an experience he had "at a high-level meeting, at one of the top 3 investment banks in the world." A person at the company was presenting his research, "and basically he had found a variable that looked highly significant in beating the market, with a regression analysis… And it turned out that this variable was the 17th monthly lag in U.S. Industrial Production." And Harvey said, "the 17th--

where's that coming from?" And then the researcher basically said, "Well, that's the only one that worked."

A large number of factor regressions like this are surely being done – an enormous number -- and we don't know how many. Some of them are sure to "work." But most of those that work are accidents of randomness. How shall we know which ones? Harvey et al.'s approach – the multiple-testing approach – of simply raising the bar of significance won't suffice because there might be thousands upon thousands of tests being done, and most of the results have been placed off the record. Harvey et al. mention that in the field of medicine attempts to replicate previously published research are frequently performed and often fail to confirm the previous results. But in the field of finance, attempts at replication are uncommon. Research in the field is a free-for-all without adequate rules or discipline.

Is it impossible to raise the bar high enough?

Harvey et al.'s paper does not ask a crucial question: Is it even possible to raise the bar high enough? Is there perhaps a **Heisenberg* uncertainty principle** in statistical research in finance, in which, when the signal-to-noise ratio – that is, the ratio of structure to randomness in the data – is too low, it will be impossible to distinguish spurious results from real ones?

The answer is yes, there is. Imagine a sequence of monthly rates of return. Now imagine a virtually endless array of series of factors that could be matched to that sequence – that is, that could be placed in a column next to it. Such a nearly endless array is easy to conjure by combining factors to create new factors for example, or by lagging them by one month, two months, …, 17 months, …, and so on.

Suppose now that there is a real relationship – a regression beta – of the sequence of rates of return against *one* of the factors. But in each case there is a high level of noise.

Here is what will happen. Because of the noise – and no matter how high the significance bar is raised – a regression of the rates of return against some one of the factors will produce a significant result. But the chances, if the signal-to-noise ratio is low enough, that the factor so identified will be the one that *actually does* bear a relationship to the sequence of rates of return will be arbitrarily low.

The identification of spurious factors will be inevitable even if one (or perhaps more) of them really does affect rates of return.

But that's only the beginning of the problems when it comes to investments

All of this rests on **the assumption** that relationships between factors and rates of return – if they exist – are ***persistent*****.** That is, that they persist from one time period to another. This assumption is warranted in almost all scientific fields. If a result of a combination of physical forces is convincingly discovered in experiments, it will work in the future too. If a vaccine is developed that prevents a disease for test subjects, it will work even if everybody uses it.

These assumptions are completely unwarranted in the investment field. If an investment strategy that beats the market is discovered and verified – even if it is not a spurious discovery – it will not work for everybody. **It cannot; it is tautologically obvious that not everybody can beat the market.** Hence, a strategy that is identified as effective, correctly or not, must, at some point, if it becomes popular and widely adopted, stop working and may even reverse to become a bad strategy.

What, then, are all those people running all those regressions doing?

They are creating an impression; an impression of diligent, technologically sophisticated research.

This is called busywork. And the most highly compensated busywork in the world it is.

Michael Edesess, a mathematician and economist, is a visiting fellow with the Centre for Systems Informatics Engineering at City University of Hong Kong, a principal and chief strategist of Compendium Finance and a research associate at EDHEC-Risk Institute. In 2007, he authored a book about the investment services industry titled The Big Investment Lie*, published by Berrett-Koehler. His new book,* The Three Simple Rules of Investing*, co-authored with Kwok L. Tsui, Carol Fabbri and George Peacock, has just been published by Berrett-Koehler.*

Kwok L. Tsui is a distinguished statistician and Head of the Systems Engineering and Engineering Management department and Chair Professor of Industrial Engineering at City University of Hong Kong.

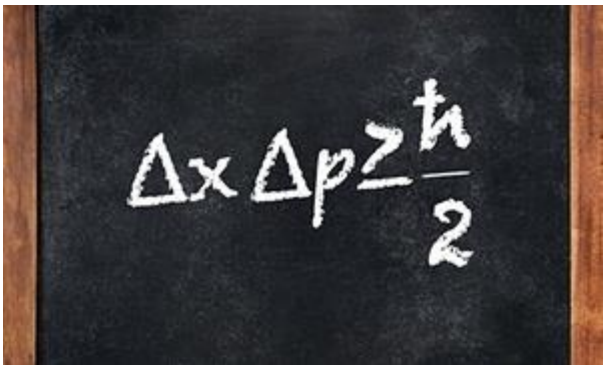

***Heisenberg's Uncertainty Principle.**

Alok Jha Sunday 10 November 2013

The **uncertainty principle** is one of the most famous (and probably misunderstood) ideas in physics. It tells us that there is a **fuzziness in nature**, a fundamental limit to what we can know about the behaviour of quantum particles and, therefore, the **smallest scales of nature**. Of these scales, the most we can hope for is to calculate probabilities for where things are and how they will behave. Unlike Isaac Newton's clockwork universe, where everything follows clear-cut laws on how to move and prediction is easy if you know the starting conditions, the uncertainty principle enshrines a level of fuzziness into quantum theory.

Werner Heisenberg's simple idea tells us why atoms don't implode, how the sun manages to shine and, strangely, that the vacuum of space is not actually empty.

An early incarnation of the uncertainty principle appeared in a 1927 paper by Heisenberg, a German physicist who was working at Niels Bohr's institute in Copenhagen at the time, titled "On the Perceptual Content of Quantum Theoretical Kinematics and Mechanics". The more familiar form of the equation came a few years later when he had further refined his thoughts in subsequent lectures and papers.

Heisenberg was working through the implications of quantum theory, a strange new way of explaining how atoms behaved that had been developed by physicists, including Niels Bohr, Paul Dirac and Erwin Schrödinger, over the previous decade. Among its many counter-intuitive ideas, quantum theory proposed that energy was not continuous but instead came in discrete packets (quanta) and that light could be described as both a wave and a stream of these quanta.

In fleshing out this radical worldview, Heisenberg discovered a problem in the way that the basic physical properties of a particle in a quantum system could be measured. In one of his regular letters to a colleague, Wolfgang Pauli, he presented the inklings of an idea that has since became a fundamental part of the quantum description of the world.

The uncertainty principle says that we cannot measure the position (x) and the momentum (p) of a particle with absolute precision. The more accurately we know one of these values, the less accurately we know the other. Multiplying together the errors in the measurements of these values (the errors are represented by the triangle symbol in front of each property, the Greek letter "delta") has to give a number greater than or equal to half of a constant called "h-bar". This is equal to Planck's constant (usually written as h) divided by 2π. Planck's constant is an important number in quantum theory, a way to measure the granularity of the world at its smallest scales and it has the value 6.626 x 10^{-34} joule seconds.

One way to think about the uncertainty principle is as an extension of how we see and measure things in the everyday world. You can read these words because particles of light, photons, have bounced off the screen or paper and reached your eyes. Each photon on that path carries with it some information about the surface it has bounced from, at the speed of light. Seeing a subatomic particle, such as an electron, is not so simple. You might similarly bounce a photon off it and then hope to detect that photon with an instrument. But chances are that the photon will impart some momentum to the electron as it hits it and change the path of the particle you are trying to measure. Or else, given that quantum particles often move so fast, the electron may no longer be in the place it was when the photon originally bounced off it. Either way, your observation of either position or momentum will be inaccurate and, more important, the act of observation affects the particle being observed.

The uncertainty principle is at the heart of many things that we observe but cannot explain using classical (non-quantum) physics. Take atoms, for example, where negatively-charged electrons orbit a positively-charged nucleus. By classical logic, we might expect the two opposite charges to attract each other, leading everything to collapse into a ball of particles. The uncertainty principle explains why this doesn't happen: if an electron got too close to the nucleus, then its position in space would be precisely known and, therefore, the error in measuring its position would be minuscule. This means that the error in measuring its momentum (and, by inference, its velocity) would be enormous. In that case, the electron could be moving fast enough to fly out of the atom altogether.

Heisenberg's idea can also explain a type of nuclear radiation called alpha decay. Alpha particles are two protons and two neutrons emitted by some heavy nuclei, such as uranium-238. Usually these are bound inside the heavy nucleus and would need lots of energy to break the bonds keeping them in place. But, because an alpha particle inside a nucleus has a very well-defined velocity, its position is not so well-defined. That means there is a small, but non-zero, chance that the particle could, at some point, find itself outside the nucleus, even though it technically does not have enough energy to escape. When this happens – a process metaphorically known as "quantum tunneling" because the escaping particle has to somehow dig its way through an energy barrier that it cannot leap over – the alpha particle escapes and we see radioactivity.

A similar quantum tunneling process happens, in reverse, at the centre of our sun, where protons fuse together and release the energy that allows our star to shine. The temperatures at the core of the sun are not high enough for the protons to have enough energy to overcome their mutual electric repulsion. But, thanks to the uncertainty principle, they can tunnel their way through the energy barrier.

Perhaps the strangest result of the uncertainty principle is what it says about vacuums. Vacuums are often defined as the absence of everything. But not so in quantum theory. There is an inherent uncertainty in the amount of energy involved in quantum processes and in the time it takes for those processes to happen.

Instead of position and momentum, Heisenberg's equation can also be expressed in terms of energy and time. Again, the more constrained one variable is, the less constrained the other is. It is therefore possible that, for very, very short periods of time, a quantum system's energy can be highly uncertain, so much that particles can appear out of the vacuum. These "virtual particles" appear in pairs – an electron and its antimatter pair, the positron, say – for a short while and then annihilate each other.

This is well within the laws of quantum physics, as long as the particles only exist fleetingly and disappear when their time is up. Uncertainty, then, is nothing to worry about in quantum physics and, in fact, we wouldn't be here if this principle didn't exist.

Our Research. Your Success.

Best regards,

Lawrence

your

CIO & Investment Strategist

Zürich, October 2015