Three Guys Win 2013 Sveriges Riksbank (a.k.a. Nobel) Prize in Economic Sciences: Who is the Best?

I'll tell you who is the best of the three guys in a minute, but first let's be real about the whole business. This prize is not like the prizes given to physical scientists. Those prizes have a long history. The prize for economics was created in 1968 by the Swedish Central Bank. Because all Nobel Prizes must go to living people, they couldn't give them to the likes of Adam Smith or Leon Walras. You likely have heard of the former but not the latter, because the latter employed higher mathematics in his work. The same is true of the living Nobel Laureates; some are known to the public, but the more technocratic are not.

Enter the three guys given prizes yesterday. Eugene Fama is well-known by anyone who has suffered through enough business school finance courses. He wrote papers and texts that have been widely adopted by business school faculty. Perhaps he is best known for his work attempting to rigorously formulate notions of market efficiency. He tried to do this by postulating links between the types of information that analysts, investors, and traders may have available to them, and the resulting stock, bond, and other assets' price movements that could result from that information flow. He then proposed statistical ways to test whether or not these linkages hold. His other best-known work is similar. For example, other professors proposed a linkage between the long-run average rate of return from a particular stock and the so-called :"market portfolio", the linkage mediated by the "beta coefficient". This proposition is taught in all business school finance departments. During the 1990s, Fama and colleagues challenged the relevance of this proposition, by proposing and implementing a statistical test that highlighted two statistics that seem anomalous from this perspective. One is the relatively high long run average of rates of return to hypothetical portfolios composed of stocks with unusually high book (equity) value relative to their market values (which they dubbed "value stocks") -- especially if one simultaneously shorts stocks on the other end of that spectrum. The other seemingly anomalous statistic is the relatively high long run average rates of return to hypothetical portfolios composed of stocks with very low market capitalization ("small stocks") -- especially if one simultaneously shorts stocks on the other end of the size spectrum.

Those with Ph.Ds in finance, economics, or statistics find Fama's work relatively easy to comprehend, and more importantly, relatively easy to extend by use of different statistical tests and/or data sets than he employed. As a result, many now-tenured faculty established their own research credentials by publishing findings that call his work into question. I hope that this will be his long-run impact on the profession, rather than the current emphasis textbook authors place on his findings.

Robert Shiller has had a smaller impact on the finance curriculum, but has become well-known by practitioners through his propensity to give media interviews, appear in advertising for Fidelity, etc. His best-known academic work challenged the Fama-like, business school mantra that stock prices are most heavily influenced by the companies' stream of future payouts, e.g. dividends and stock repurchases. He and colleagues proposed and implemented statistical tests that seemed to show that stock prices were more volatile than they should be if that were the case.[1] This started a cottage industry of alternative ways to both formulate the question and to interpret the findings (e.g. is the phenomena due to time-varying risk aversion, or time-varying "discount rates" induced by time-varying risk aversion, etc.). As often happens when someone receives acclaim for influential findings, others may have been slighted out of that acclaim for similar findings. In Shiller's case, my former Univ. of Minnesota colleague Steve LeRoy is one of them. In addition, he wrote at least one other well-cited article that challenged the legitimacy of some of Fama's influential papers. Steve has a good career going, but deserves more visibility for the originality and depth of his work.

Notwithstanding the reservations noted above, both Fama and Shiller did good technical work earlier in their careers that is less well-known. For example, Fama took nascent steps to theorize about the market pricing of effects of fatter-than-normal tails in asset return distributions. Such theorizing has come back into vogue following our unsettling severe market crashes. And Shiller did earlier work in time-series econometrics on the estimation of autoregressive models. Autoregessive models soon grew to challenge the large-scale -- albeit intellectually bankrupt -- modeling orthodoxy promoted by Otto Eckstein's Data Resources Inc. and similar entrepreneurial faculty.

We now turn to Lars Hansen, the least-publicly known co-winner of this year's Prize. I went to grad school with Lars at the University of Minnesota's Economics Department. It might surprise you to read that during my Minnesota years, I also studied from and/or worked with (either at the Minneapolis Fed and later as a University of Minnesota professor) four others who won that Prize: Leo Hurwicz, Ed Prescott, Tom Sargent, and Chris Sims. Lars and the rest of them are truly great people and scholars who have been kind to me over the years. When I was a student, I wasn't the best-behaved guy (but hopefully not the worst). Hurwicz said I distracted him on occasion. Sargent summoned me to tell me that another Ph.D student wanted me banned from class (it isn't my fault that the guy wound up teaching somewhere in South America), and Sims got annoyed when I stood up at a large conference to challenge something that he asserted about nonlinear dynamics research (my challenge impressed enough of the audience to irk him). But over the many years that followed, all of them have gone out of their way to speak and act favorably about my subsequent work.

Some of my published work (*J. Econometrics* 1995, *Econometrica* 1997, *J. Econometrics* 2002) work, the latter two done in conjunction with (now) Yale Economics Professor Yuichi Kitamura, is intended to augment and extend the framework that Lars devised. So it is especially gratifying that he has been supportive. I'll do my best to explain what Lars achieved, without using statistical theory (OY!).

Let's start by considering a seemingly unrelated issue: how to measure the Earth's mass. In science class, all of us have placed something on a balance to measure its mass. But

---

[1] Here, as in my assertions about Fama's work, I am oversimplifying what he did, in order to make this missive more accessible to those of you who don't have Ph.Ds in finance, economics, or statistics.

even the kids who skipped that day know that there isn't a balance big enough to hold the earth. So how is it done? The answer is to use a theory that relates the earth's mass to other quantities that can be measured, and then work backward to figure out the value of the earth's mass that the theory implies. I used a websearch to immediately find that one such way is to drop a small object of some mass (dubbed "object mass"), and measure its acceleration while falling. Newton's 2nd Law of Motion relates these quantities to the gravitational force acting on it:

F = object mass x its acceleration.

Newton's Universal Law of Gravitation relates the same force to the Earth's Mass, the distance "r" from the object to the earth's center of gravity, the gravitational constant "G":

F = (G/ r^2) x Earth Mass x object mass

Now subtract the two equations and divide through by the object mass to obtain the equation:

(1) acceleration of a small object - (G / r^2) x Earth Mass = 0

A physicist might just plug-in an estimated value for the gravitational constant "G", an estimate for r^2, and his/her measurement for the dropped object's acceleration. The only number left is Earth Mass, so the physicist could solve the above equation to estimate: Earth Mass =acceleration of a small object x r^2 / G. Repeated dropping of the same small object, or perhaps other small objects, would result in (hopefully only slightly) different estimates of the Earth's Mass. The slight differences could arise from the error in measuring r and/or G and/or the acceleration of what was dropped. The average or perhaps the median of those numbers could be adopted as the estimate of the Earth's Mass.

But this procedure might only be as good as the estimate used for G / r^2. For most purposes that might not be a problem, but for others it might be. Estimates of the earth's radius r have been made since Erastothenes in the 3rd Century B.C., and of course have been improved subsequently. The gravitational constant "G" has a very long history of importance, yet the author of its Wikipedia entry notes that

"*G* is quite difficult to measure, as gravity is much weaker than other fundamental forces, and an experimental apparatus cannot be separated from the gravitational influence of other bodies."

The author mentions the original measurement technique Cavendish in 1798, and that attempts to improve on it persist to this day.

So let's try to avoid the whole issue by defining a *parameter* θ_1 (the Greek letter theta) representing G / r^2, another parameter θ_2 representing the Earth's Mass, and posit that over repeated measurements of a dropped small object or objects, the long-run average or "mathematical expectation" of (1) would equal zero:

(2) E[acceleration - θ_1 x θ_2] = 0.

The observed data are the set of acceleration measurements on the same or different small objects. Because the law of large numbers implies that a hypothetical infinite number of measurements would average to the above expected value E, the statistician with a finite amount of data might just look for a value of θ_1 x θ_2 that makes the measured numbers' finite average equal zero, i.e. set θ_1 x θ_2 equal to the average acceleration. But that will only provide an estimate for θ_1 x θ_2, which cannot as yet be decomposed to find the desired estimate for the Earth's Mass θ_2. In econometrics, this is called a problem of "*underidentification*".

If a physicist could use theory to posit another expected value relationship, that like (2), depends at most on θ_1 and θ_2, this would provide another "moment condition":

(3) E[data; θ_1, θ_2] = 0

Then, with an infinite amount of observation of both the acceleration of small objects from (2), and whatever observed data is needed for (3), the law of large numbers would imply that the "long-run" average of the expressions inside the brackets of (2) and (3) will both be zero when evaluated at the correct values for θ_1 and θ_2. With a little luck, the two equations in those two unknowns will uniquely determine values for those two parameters, a property known as "*exact identification*". But again, the real world only yields a finite amount of data, so the econometrician could specify a *criterion function* used to quantify the meaning of two finite averages both being close to zero. For example, one might just square the value of each average, add the squares together, and search of values of θ_1 and θ_2 that make this sum of squares as small as possible. These values are the econometrician's estimates of the parameters θ_1 and θ_2, with the latter (i.e. the Earth's Mass) the parameter of most interest -- recall that his was what we wanted to know in the first place!.

But there is nothing sacred about that particular criterion function. For example, one might seek parameter values that make the sum of the absolute values of the two bracketed expressions' averages as close to zero as possible.

Of course other physical theory could be brought to bear, resulting in additional moment conditions with the form (3). Then we would have more moment conditions than there are parameters, a situation dubbed "*overidentification"*. Again, a hypothetically infinite amount of data would then result in more equations than unknown parameters. One might discover that there are no values of the two parameters that satisfy all the moment conditions, in which case we reject the underlying theory or theories that led to those moment conditions. Presumably with an infinite amount data of the right kind, Einsteinian gravitational effects would cause this to be the case.

The types of asset pricing theories that have captivated economists' attention are also overidentified. While it is not taught this way outside of Ph.D programs, most "risk vs. return" type of theories predict some relationship between an individual asset's expected gross return (denoted R_i for the ith asset whose gross return can be measured accurately) and an obscure quantity alternatively called a "*stochastic discount factor*" or "*pricing kernel*" m. Different asset pricing theories lead to different pricing kernels, each of which is a function of some data (e.g. aggregate consumption in the economy, or the "market portfolio's" return) and some parameters (e.g. the "coefficient of risk aversion", the "rate of subjective time discount", the "elasticity of intertemporal substitution" and others).
Hence each theory yields a different specification for the function m, which we write as

(4) $m(\text{data};\ \theta_1, \theta_2, \theta_3, \ldots)$

Lars Hansen and others noticed that in such asset pricing theories, the moment conditions can be placed in the form:

(5) $\mathbf{E[R_i\, m(data;\ \theta_1, \theta_2, \theta_3, \ldots)-1] = 0\ ,\ i = 1, 2, 3, \ldots}$

The immense database for the hordes of different assets' returns implies that there are a very large number of equations in (5) relative to the number of parameters in asset pricing theories. As such, the situation is inherently overidentified.

Lars Hansen won the Nobel Prize in large part for his econometrically innovative approach to using finite data to estimate parameters subject to the moment conditions (5). He posited a general class of criterion functions used to measure closeness of the resulting group of averages to zero. He then derived a way to specify a particular criterion function in that class, dependent on the data and the posited moment conditions, that would lead to good parameter estimates when the moment conditions are true (i.e. the theory implying them is true). He also provided a statistical test for the latter, used in an attempt to settle the question of whether a posited theory holds water.

All this was Lars Hansen's breakthrough, which he called the "Generalized Method of Moments".His GMM method was fairly rapidly adopted, in no small part due to computer software developed and made available by others, that permitted use of the method without deep understanding of it. I doubt that the Great Gauss' own Method of Least Squares (i.e. "regression") estimator would be widely adopted if not for the same access to "user-friendly" software that does not require deep knowledge of the underlying statistical theory. For evidence, just look at the bulk of misleading statistical work done in departments ranging from sociology to epidemiology.

Considering the bulk of published finance studies using the GMM method, I am sad to report that most of this work has rejected the asset pricing theories still taught in Ph.D programs. From this we learn that:

Economics $\neq$ Physics. Ain't No Einstein Here.

My work done in conjunction with Prof. Yuichi Kitamura (cited earlier here) devised and promoted a different way of addressing the same class of problems that he did -- i.e. problems requiring estimation of fewer parameters than the number of equations that the theory implied should equal zero. When a subsequent researcher develops and promotes an approach that differs from an innovator, the innovator often engages in turf defense by denying the legitimacy of the proposed alternative. But an intellectually honest and humble innovator will judge the proposed alternative solely on its merits. That is what Lars Hansen did with our work. He doesn't view it as superior, but rather as an interesting alternative approach that has intrinsic merit, particularly via other, yet to be found applications of the same mathematics. He has used some of that mathematics (i.e. the statistical theory of large deviations, and its entropic interpretation) in seemingly unrelated theories he has devised with Tom Sargent and with Jose Scheinkman.

When I return home from work today, I plan to raise a chilled glass of akvavit to Lars Hansen, and to the Swedish Central Bank for funding the prize given to Lars. Another student once told me that very long ago Lars said something to the effect that I would either win the Nobel Prize or never be heard from again. He was correct there, too!