Gustav Bamberger

'Tis an Equity Puzzlement

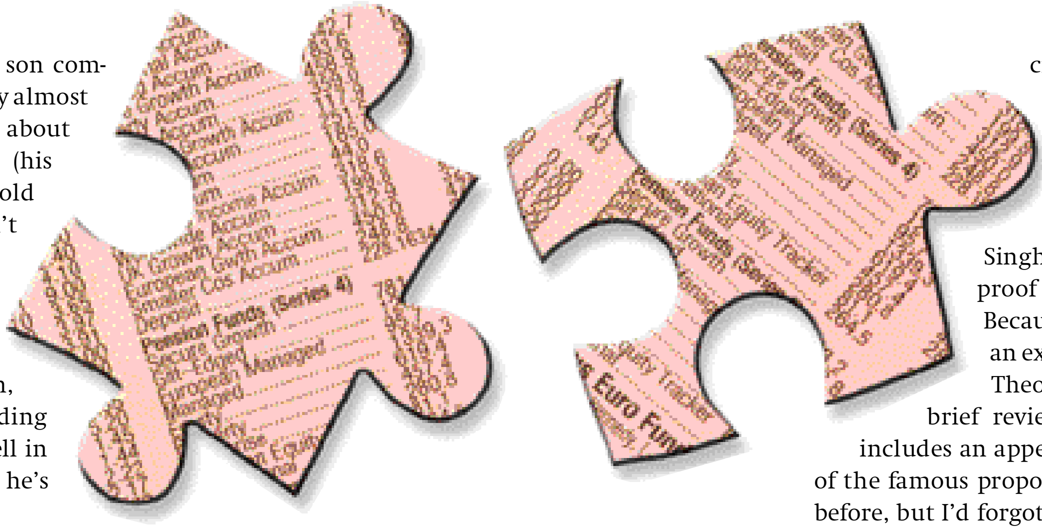

My 11-year old son complains bitterly almost every night about homework (his eight-year old brother isn't old enough for homework yet, which is especially infuriating). Although he doesn't like any homework, he's especially unhappy when it's math, his least favorite subject. According to his teachers, he does very well in the subject, but he's convinced he's not any good at it.

I'm sympathetic, because as a kid I was tagged with the 'good at math' label. I was skeptical. My test scores were always good, but I recognized that my approach was basically 'brute force' math; I memorized whatever I needed to know, but a deep understanding of what I was doing usually escaped me. Still, all those teachers told me I was good at math; they must know something, right? (One of the admittedly few benefits of getting older is that I eventually learned that just because everyone tells me something doesn't mean that everyone knows what they're talking about.) So in my first year of college, I decided to major in mathematics.

That plan lasted about two months. I started with a course in Real Analysis, which was a struggle. I enjoyed it, but I realized that I was in the wrong place when I got an early exam back. On one of my answers the professor wrote a note telling me that my approach was the mathematical equivalent of counting the number of cow legs in a field and then dividing by four if you want to know how many cows are out there. Apparently, my solution was especially inelegant.

The equity premium puzzle has long kept mathematicians and economists enthralled

I changed my major to economics. At some level, that was too bad, because I always enjoyed mathematics. What I found especially intriguing and gratifying were those infrequent flashes of insight when something that seemed completely impenetrable one minute became blindingly obvious the next. But that was the part creativity and insight – that I seemed to be missing. I was reminded of my previous life as a failed mathematician recently while reading Simon Singh's book on Andrew Wile's proof of Fermat's last theorem. Because Fermat's last theorem is an extension of the Pythagorean Theorem, Singh starts with a brief review of Pythagoras, and he includes an appendix that contains a proof of the famous proposition. I must have seen it before, but I'd forgotten it long ago. It's an elegant, simple proof; a diagram and a line of algebra is all you need and it gave me that forgotten thrill of all-of-a-sudden getting it it's so obviously right once you see how it's done.

Something else I liked about mathematics is that people can spend years studying something just because they find it interesting, even if there are no obvious practical applications. For example, I was startled to learn in Singh's book about the fate of Euler's conjecture. Leonhard Euler, a great 18th–century mathematician, besides being familiar to first-year graduate students in economics from the Euler equation (having to do with homogeneous equations), was also responsible for an early contribution to the study of the Fermat problem. Euler showed that Fermat's last theorem was true for the case $n = 3$.

Another of Euler's contributions to mathematics has come to be known as Euler's conjecture. Euler claimed that there were no whole number solutions to an equation that looks like Fermat's: $x^4 + y^4 + z^4 = w^4$. Just like Fermat's last theorem, no one had been able to prove Euler's conjecture for hundreds of years. But in 1988, Harvard mathematician Naom Elkies showed why. It turns out that Euler's conjecture is wrong. For example,

$$2,682,440^4 + 15,365,639^4 + 18,796,760^4 = 20,615,673^4 .$$

I have no idea why Elkies was investigating Euler's conjecture (he also showed that there were an infinite number of other solutions), but it seems clear that it wasn't so he could take his finding and use it to launch a multi-billion dollar IPO (not that there's anything wrong with that).

In my working life as an economist, that type of elegance and certainty (Euler's conjecture is wrong, period) is rare. I know there are many economists writing theory papers which contain numerous 'proofs,' but that's not what I do. I'm typically analyzing data of one sort or another, and data are invariably messy and less well behaved than you might hope.

Take the 'equity premium puzzle.' For the last 15 years or so, there's been a stream of articles analyzing why stockholders have historically earned so much more than owners of bonds. At first blush, this doesn't seem like much of a puzzle. One of the basic tenets of modern finance theory is the 'risk-return tradeoff' in equilibrium, relatively more risky assets must have a higher expected return than relatively less risky assets, and stocks are riskier than bonds. But in 1985, economists Rajnish Mehra and Edward C. Prescott showed that, based on standard theory, the equity premium should be about 0.35 percent per year far lower than the observed risk premium. For example, in the US over the period 1889-1978, the equity premium as usually calculated was over 6 per cent.

The risk-return tradeoff is a qualitative proposition there's nothing in the theory that tells you how large the risk premium should be. To estimate what the equity premium should be, Mehra and Prescott start with a standard theory and apply real-world data to it. First things first: what precisely do Mehra and Prescott mean by 'risk'?

For the last 15 years or so, there's been a stream of articles analyzing why stockholders have historically earned so much more than owners of bonds

In the capital asset pricing model, the (undiversifiable) risk of a stock depends on the correlation of its return with the return of 'the market.' The CAPM can be written as:

$$E[R_S] - R_f = \beta(E[R_m] - R_f) .$$

where $E(R_S)$ is the expected return on a particular stock; R_f is the 'risk-free' rate; β measures the covariance of the stock's returns with the market's; and $E(R_m)$ is the equity risk premium. So the CAPM (and other asset pricing models) use the equity risk premium as an input, the equity risk premium is 'outside the model.'

The equity risk premium has to be derived from a more basic model. So: why do people save and invest? Fundamentally, you have to move consumption from the present to the future (whether you or your heirs get to consume in the future, of course, is a different issue). Economists typically assume that individuals are risk averse and maximize expected utility. Mehra and Prescott start with a model where utility depends on consumption, so what matters to a consumer is how the return to an asset is correlated with consumption over time. The basic idea of risk is, the same as in the CAPM the only risk that matters is risk that can't be diversified away, but now a low-risk asset is one that isn't very correlated with your consumption. Because risk aversion means you'd like to 'smooth' your consumption, an asset that pays well when your consumption is low and pays badly when your consumption is high is low risk. (Confusingly, for technical reasons these types of models assume, in effect, that consumption is identical to income in every period. Which means there's no saving, so why worry about investing in stocks and bonds? One way to think about the setup is that everyone has labor income plus a wealth endowment, and the question is how to invest the endowment between different assets, like stocks and bonds; risk comes from the extent of correlation between asset returns and labor income.)

Here's where the problem starts, it turns out that even though the return on stocks is more correlated with *per capita* consumption growth than the return on bonds and so more risky it isn't by much. So why does a slightly higher correlation lead to a big equity risk premium? Let's take a step back to theory.

Recall that individuals are assumed to maximize expected utility, which means that you need to know what expected utility depends on. Mehra and Prescott use a standard utility function that they write as:

$$U(c, \alpha) = \frac{c^{1-\alpha} - 1}{1 - \alpha} \quad \text{for} \quad \alpha > 0 .$$

where $U(c, \alpha)$ is the utility derived by an individual in a time period in which consumption equals c, and α measures the individual's 'risk aversion.' The larger the value of α, the more risk averse the consumer.

So a key question is: how big is α, how risk averse are people? If α is big enough, even a small difference in the riskiness of stocks and bonds can explain a large equity risk premium. Mehra and Prescott review several studies that estimate α and report that it appears to be in the range of 0 to 2. (By the way, you may have noticed that the utility function isn't defined for $\alpha = 1$; so if $\alpha = 1$, the utility function is defined as the logarithmic function, which is the limit approached by the utility function as α approaches 1.) But to be sporting about it, they allow the value of α to be as high as 10. Even then, though, the largest equity risk premium they can get is 0.35 per cent a long way from six per cent. Mehra and Prescott conclude that standard 'models that abstract from transactions costs, liquidity constraints and other frictions' cannot explain the equity premium.

But maybe the other studies are wrong and α is larger than 10. Unfortunately, to get an equity premium of around 6 per cent, α needs to be around 100. And that leads to bizarre predictions about individual behavior. Risk aversion means you'd be willing to pay to avoid a 'fair bet' a bet with an expected return of zero. (An example of paying to avoid fair bets is insurance. Not only are you willing to buy insurance, but many of us are even willing to spend time with people who are trying to sell us insurance, can there be more eloquent proof of risk aversion?) But an α of 100 means you're really risk averse. To illustrate the point, Bradford Cornell provides this example: "Consider a family that annually consumes $50,000 and that faces a fair gamble. A coin will be flipped. If it lands heads, the family wins $10,000; if it lands tails, the family pays $10,000. The amount that the family will pay to avoid this bet is a measure of their risk aversion." With an α value of 100, the family is willing to pay $9,700 to avoid the bet. Which means the family is willing to lose 97 percent of $10,000 with certainty to avoid a 50 percent chance of losing $10,000. (If α equals 2, by the way, the family is only willing to pay about $2,000 to avoid the bet.) So extremely high values of α aren't an attractive option for explaining the equity premium.

Mehra and Prescott's findings generated numerous studies that relaxed one or more of the basic model's assumptions in an effort to explain the puzzle. In a review of these studies, Narayana Kocherlakota shows that the equity premium puzzle depends on three basic assumptions: (1) individuals' preferences are explained by the 'standard' utility function; (2) asset markets are 'complete,' which means that individuals can insure against any bad outcome; and (3) trading stocks and bonds is costless. Unfortunately, relaxing these assumptions doesn't seem to help all that much Kocherlakota concludes that "the literature provides only two rationalizations for the large equity premium: either investors are highly averse to consumption risk or they find trading

If you can't explain the equity premium puzzle, maybe the next best approach is to hope real hard that it goes away.

stocks to be much more costly than trading bonds. Little auxiliary evidence exists to support either of these explanations, and so (I would say) the equity premium puzzle is still unresolved."

If you can't explain the equity premium puzzle, maybe the next best approach is to hope real hard that it goes away. Sure enough, that seems to be happening. A recent study, for example, finds that the equity premium in the US was about 7 per cent during the period 1926-70, but only about 0.7 per cent since 1970. Based on returns earned by stocks in the S&P 500 and long-term bonds (different portfolios will generate somewhat different answers), the equity premium was 1.76 per cent during the 1970s; -0.59 per cent during the 1980s; and 0.98 per cent during the 1990s. The authors speculate that the decline in the equity premium has to do with the reduction in market imperfections during the last few decades, lower trading costs, more easily available information, but they concede that they don't have a definitive explanation for what's going on.

But before we get too carried away, keep in mind that we don't really have a good idea of how big the equity premium was even back when we all thought that it was big. The problem is that the equity premium is unobservable; all we can see are realized stock returns, which aren't the same thing as expected stock returns. And realized stock returns can vary dramatically year by year (some of you may have noticed this recently). The difference between stock and bond returns during the period August 1929 to June 1932, for example, was -59.30 per cent, and no one thinks that period's negative equity premium was the result of stocks being especially unrisky. So the equity premium is estimated by measuring the difference between stock and bond returns over long periods Mehra and Prescott use 90 years. The hope is that over long periods, the difference between realized and expected stock returns will average to zero. But that means that even 30 years may not be enough to be sure that the equity premium has changed. We may have to wait another 60 years or so to be sure.

Cornell summarizes the current state of knowledge about the equity premium (which he calculates over the 72-year period between 1926 and 1997 as 7.4 per cent): "72 years' worth of data is not enough to measure the risk premium with sufficient precision to satisfy most investors. Although the historical risk premium over treasury bonds is 7.4 per cent, the data are so imprecise that the hypothesis that the true forward-looking risk premium is 3 per cent or 12 per cent cannot be rejected at standard levels of statistical significance."

But suppose that the equity premium really has fallen to close to zero like the theory predicts. Don't start celebrating just yet. A drop in the equity premium means that stock prices will go up (because future dividends are discounted at a lower rate), but once the new equity premium is fully reflected in stock prices, stocks should be a much worse long-term investment less risk, less return, remember? So if your retirement plans are based on the assumption that your stock portfolio will increase in value at a real rate of return of 8 per cent or so, you may be in for a nasty surprise.

It's safe to say that the equity premium puzzle and its apparent disappearance remain mysterious. Fortunately, I've been thinking about this issue a lot, and I'm happy to report that I have a truly marvelous solution to this problem but, unfortunately, the space I have left is too narrow to contain it.

W