REGRESSION TO THE MEAN AND VALUE INVESTING

Part I

1 June 2002

The child inherits partly from his parents, partly from his ancestry … The further his genealogy goes back, the more numerous and varied will his ancestry become, until they cease to differ from any equally numerous sample taken at haphazard from the race at large … This law tells heavily against the full hereditary transmission of any gift … If it discourages the extravagant expectations of gifted parents that their children will inherit all their powers, it no less discountenances extravagant fears that they will inherit all their weaknesses and diseases.

Sir Francis Galton
Inquiries Into Human Faculty and Its Development (1883)

Graham's Application

Leithner & Co. is a Graham-style "value" investor. Benjamin Graham (1894-1976) was the author of the classic books *Security Analysis* (1934) and *The Intelligent Investor* (1949), a director of Graham-Newman Corp. (1926-1955), an instructor at Columbia University and a mentor and colleague of outstanding investors such as Warren Buffett, Thomas Knapp and Walter Schloss. He is widely regarded as the founder of financial analysis. Graham's key insight is the premise that "investment is most successful when it is most businesslike. An investment operation is one which, upon thorough analysis, promises safety of principal and a satisfactory return. Operations not meeting these requirements are speculative."

Graham sharply distinguished two critical notions: price is what is paid and value is what is received. He observed that at any given point in time the value and the price of a financial asset often diverge (sometimes by a wide margin), but that over time they gravitate towards one another. Value investors, as practitioners of Graham's approach are often called, thus reject today's prevailing view that the price and value of a security (i.e., stock, bond, title to real estate, etc.) necessarily coincide at all times. Graham lamented that most people rarely recognise, and more than a few wilfully ignore, the fundamental difference between – and relationship among – value and price. He regretted, in effect, that a foundation notion ("regression to the mean") is often discounted or ignored entirely by most market participants.

Galton's Insight

Sir Francis Galton (1822-1911), a cousin of Charles Darwin, was keenly interested in heredity and not at all in business and finance. Yet his studies of "the average ancestral type" uncovered a statistical regularity that provides a basis for Graham-and-Buffett style value investing. Heredity is the transmission of key characteristics, such as height, weight, eye colour, susceptibility to and protection from certain diseases, etc., from parents to offspring. The study of heredity takes into account "outliers" – individuals whose characteristics do not conform to the group's norms – but its focus is upon the basic tendency of all members of a species to resemble one another.

Consider peas in a pod. In 1875 Sir Francis conducted an experiment proposed to him by his cousin. After weighing and measuring thousands of sweet peas, ten specimens of each of seven categories of weight were sent to nine associates across the British Isles, including Darwin, who were instructed to plant the peas under carefully specified conditions. After analysing the results, Galton found that the spread of weights and diameters among parent peas within a given category was slightly wider than the dispersion among their offspring. The offspring of peas from each category of weight, in other words, had an overall distribution that was slightly tighter than the distribution of their parents. In an analysis of the heights of human parents and their children, Sir Francis found that tall parents tended to bear tall children and that short parents tended to bear short children. Heredity clearly matters. But it matters in a particularly interesting way: on average, the offspring of tall parents were not as tall as their parents; and the offspring of short parents were not as short as theirs. These and other experiments led Sir Francis to formulate a principle that has become known as reversion (or regression) to the mean.

According to Galton, **"reversion is the tendency of the ideal mean filial type to depart from the parental type, reverting to what may be roughly and perhaps fairly described as the average ancestral type."** If this process of reversion were not at work, i.e., if large peas produced ever-larger peas and small peas produced ever-smaller offspring, the world would eventually comprise nought but midgets and giants. With every passing generation nature would produce ever more freaks and extremes. More generally, "regression to the mean" refers to an inverse correlation among roughly normally distributed observations that are made repeatedly over time. An extreme observation at one point in time ("outlier") tends to be followed by a less extreme observation; **extremes, in other words, revert or regress over time towards mean or average measurements.** Consider as another example a student's test scores. If the student's scholastic aptitude remains constant over some (say, six-month) period of time, then an extreme test result observed at a particular point will probably be followed by ("regress towards") a result that is much closer to the student's average. All else equal, an abnormally high score is likely to be followed at the next test by a lower score; and an abnormally low score is likely to be followed next time by a higher score.

Sir Karl Pearson, Galton's biographer and an outstanding scientist in his own right, observed that Sir Francis had created "a revolution in our scientific ideas [that] has modified our philosophy of science and even of life itself." Galton transformed the notion of probability from a static concept based upon randomness into a dynamic process in which the successors of outliers are predestined to join the crowd. Change and motion from the outer limits towards the centre are incessant and (assuming no change to the system) inevitable. Regression to the mean underlies homilies such as "what goes up must come down" and "from shirtsleeves to shirtsleeves in three generations." **It is a principle that Grahamite investors exploit to great advantage.**

[The mutual-fund manager interviewed on 8 April stated] that there is no reason not to expect 8% to 10% annual returns over the next eight years. I do not mean to pick on him specifically since many of his peers seem to have similar outlooks, but I have not seen anyone say how this is going to happen.

If dividend yield of 1.38% remains constant, to achieve a total return of 10%, the compounded capital return on the S&P 500 index would have to be 8.62% per year. The index would be at 2,223 eight years from now. If earnings on the S&P 500 grow at their 60-year average of 6.2%, total earnings eight years from now would be $62.60, giving the S&P 500 a P/E ratio of 35.5.

What type of returns would be realized assuming a return to the long term average P/E for the index, which is 16? With earnings of $62.60 and a P/E of 16, the S&P would sell at 1,002, which is 12.6% below its current level. Adding in dividends would give a total compounded return of less than 1% per year.

Unless P/E ratios go to and stay at levels never seen before in history, we are looking at much lower returns from large-cap stocks for a long time to come. It is time for all of us in the business to wake up and acknowledge this reality.

Mr Scott Berglund, Roanoke, Virginia
Barron's (29 April 2002)

Sir Francis' Inadvertent Gift to Value Investors

The folklore of investment and financial markets is full of catchcries such as "buy low and sell high." Many are variations on a simple theme: if you allocate your investment capital in a manner consistent with the assumption that the tightly-circumscribed bounds of normality revealed over the twentieth century (i.e., the average annualised return on listed Australian equity over the past 100 years is 9.6% and has a standard deviation of 1.6%) extend into the future, then you will accumulate wealth sooner and face a smaller risk of permanent loss of capital than if you run with today's crowd and this month's fad. Yet many and probably most market participants violate this principle every day.

For whatever reason, and as Ben Graham's parable of Mr Market shows, they are emotionally incapable of insulating themselves from their emotions and embracing their faculties of reason. **Instead, comforted by conformity and bedevilled by greed and fear, they run with the crowd and do not stop, clear their minds of detritus and think for themselves.** Seemingly for psychological reasons, then, it is hard to keep Sir Francis, his pea pods and the principle of regression to the mean in plain view. Even though we never know what is going to happen tomorrow, let alone next year, it seems to be easier to focus upon and plan for the next day than for the more distant future.

From this tendency follows another: the proclivity to assume that tomorrow, the next day and the day thereafter will resemble today, i.e., to focus upon recent events and "case rates" rather than long-term trends and "base rates." In 1930-32 and 1973-74, then, the assumption was that the storm would never end; and in the late 1920s and late 1990s the assumption was that inclement conditions could never appear. **Hence the tendency to lose sight of Sir Francis's priceless gift to investors: the likelihood that investment results obtained tomorrow and into the more indefinite future will more closely resemble the average of results obtained over the past one hundred years than the results obtained today and yesterday.** The negative extremes of the Great Depression and the 1970s were unsustainable and subsequently regressed upwards towards historical averages. The positive extremes of the late 1920s and late 1990s were also unsustainable, and subsequently regressed (or are in the process of regressing) downwards towards those same historical averages.

This latter point helps to clarify much of what is presently discombobulating many market participants. In the words of Barrie Dunstan (*The Weekend Australian Financial Review*, 4-5 May), "what is happening is probably a drawn-out phase during which overpriced stocks of all types are coming back to earth … Forget about projections of economic growth in the United States or elsewhere. Stockmarkets at the moment aren't about business prospects; they're about the excessive valuations investors have been placing on shares. These valuations went far too high until early 2000. Now they are in the process of adjusting in what the experts call 'reversion to the mean.'" According to Robert Fuller (cited by

Dunstan), "you won't hear much about reversion to the mean from investment managers – one never does when they are on the wrong side of it – but the process is one of the most logical and predictable long-term cyclical developments in markets."

Regression to the Mean in Financial Markets

Investors such as David Dreman and behavioural economists such as Richard Thaler and Werner De Bondt have uncovered strong evidence that regression to the mean, or something akin to it, occurs on financial markets. It occurs at both individual and aggregate levels, i.e., with respect to both individual securities and markets as a whole. Using data for the period 1926-1982, Thaler and De Bondt studied the securities of those companies whose prices over a three-year interval had either increased or decreased more than the market average. They found that "extreme returns of stocks listed on the New York Stock Exchange were subsequently followed by significant price movement in the opposite direction."

If investors are either unduly optimistic or pessimistic about a particular company's securities, and if that company's fundamentals remain unchanged, then their stance will likely be reversed over time. Very fashionable stocks and market segments thus become less exalted, and highly unfashionable companies and sectors return to average favour. Dreman, in particular, has demonstrated how the crowd's exaggerated reactions, unwittingly aided and abetted by the mass media, occasionally offer tremendous investment opportunities to those who are prepared to stand apart from the crowd.

If the price of a financially sound company's stock is savaged by pessimistic investors, mass media coverage and commentary to the point where it falls considerably below its intrinsic value, then – as long as the company's operations and prospects remain sound – the price of its securities will eventually recover. Conversely, if a company's shares are inflated above their intrinsic value by euphoric (or simply commission-based) brokers, advisors and media supporters, then – even when its operations and prospects remain unchanged – at some point it will fall from its exalted status.

Although DeBondt and Thaler's test methods have been subjected to some criticism, their findings have been confirmed by others using different methods. When investors overreact to new information and ignore long-term trends, regression to the mean turns the average winner into a loser and the average loser into a winner. This reversal tends to develop with some delay, and this delay creates opportunities for alert entrepreneurs. It appears that at first buyers and sellers overreact to short-term news and then under-react whilst awaiting additional short-term news of a different character. As with peas in a pod, so too with companies: it could not be otherwise. If it were, i.e., if the winners kept on winning and the losers kept on losing, then the economic and financial landscape would consist of a shrinking handful of companies with colossal market capitalisations and virtually no small enterprises.

The Market as a Whole

If the DeBondt-Thaler hypothesis of overreaction to recent news applies to the market as a whole and not just to individual stocks, then regression to the mean should manifest itself as longer-term realities make themselves felt. If, on the other hand, investors are more fearful in some economic environments than in others – say, 1932 or 1974 in contrast to 1968, 1986 or 1999 – stock prices would fall as long as investors are afraid and would rise again when circumstances changed and justified a rosier view of the future. Both possibilities suggest a Graham-like stance vis-à-vis the broader financial market, market news and other market participants: one should ignore short-term volatility, discount the latest news and disregard "case rates" and concentrate upon "base rates" enduring developments and phenomena such as regression to the mean. To do so is to ignore the crowd, embrace Emersonian self-reliance and hold for the long term.

A More Recent Study

The applicability to financial markets of regression to the mean is thus a salutary reminder that at critical junctures the mass media and the crowd are mistaken. At these junctures their views do not reflect reality, but rather exaggerated (i.e., unduly optimistic or pessimistic) perceptions of that reality. This simple but profound notion was recognised at least 2,000 years ago. (Horace, the Latin lyric poet and satirist, wrote in his *Ars Poetica* that "many shall be restored that now are fallen; and many shall fall that now are in honour"). Today, however, it appears either that market participants have forgotten it or that their methods obscure it.

The remarkable levitation of financial markets during the closing years of the twentieth century prompted economists John Campbell of Harvard University and Robert Shiller of Yale University to revisit a topic they discussed in a 1998 paper based upon their joint testimony before the Federal Reserve's Board of Governors in 1996. Their revised paper is entitled "Valuation Ratios and the Long-Run Stock Market Outlook: An Update." Using annual data from 1871 to 2000 and quarterly data for twelve major countries since 1970, Campbell and Shiller (author of *Market Volatility*, MIT Press, reprint edition 1992, ISBN: 0262691515; and *Irrational Exuberance*, Broadway Books, 2001, ISBN: 0767907183) examined price-to-earnings and dividend-to-price ratios as forecasting variables. In their words, "various simple efficient-markets models of financial markets imply that these ratios should be useful in forecasting future dividend growth, future earnings growth or future productivity growth."

They find, however, that "overall, the ratios do poorly in forecasting any of these. Rather, the ratios appear to be useful primarily in forecasting future stock-price changes, contrary to the simple efficient-markets models … When stock-market valuation ratios are at extreme levels by historical standards, as dividend/price and price/earnings ratios have been for some years in the U.S., one naturally wonders what this means for the stock-market outlook. It seems reasonable to suspect that prices are not likely ever to drift too far from their normal levels relative to indicators of fundamental value, such as dividends or earnings." **Hence the simple reversion-to-the-mean theory: "when stock prices are very high, relative to these indicators, as they have been recently, then prices will eventually fall in the future to bring the ratios back to more normal historical levels."** Indeed, and as Campbell and Shiller testified in 1996, "despite all the evidence that stock returns are hard to forecast in the short run, this simple theory of mean reversion is basically right and does indeed imply a poor long-run stock-market outlook."

A Conclusion and Three Cautions

The simple and powerful notion of regression to the mean provides the underpinning of many decision-making systems. And for good reason: there are few situations when large things continue without interruption to become infinitely large and small things become infinitesimally small. Trees grow upwards but never reach the heavens. Accordingly, when we are tempted – as we so often are – to extrapolate past trends into the future, we should remember the genetics of Sir Francis's humble sweet peas.

Yet if regression to the mean follows such a constant pattern, why is forecasting such a frustrating (and ultimately useless) activity? The simplest answer is that the forces at work in nature are not the same as the forces at work between people's ears. The accuracy of most forecasts depends upon decisions made by people rather than Mother Nature; and nature, with all its vagaries, is much more dependable than an individual or committee of individuals trying to make a decision.

Accordingly, there are three reasons (Peter Bernstein, *Against the Gods: The Remarkable Story of Risk*, John Wiley & Sons, 1998, ISBN: 0471295639) why regression to the mean can be a fallible and frustrating guide to decision-making. First, sometimes regression proceeds at such a slow pace that an exogenous "shock" will disrupt or reverse it. Second, the regression may be so strong that matters do not come to rest once they reach the mean; rather, they "overshoot" before fluctuating irregularly around the mean. Finally, and perhaps most importantly, the mean itself may be unstable, so that yesterday's normality may be supplanted today by a new normality that we cannot anticipate and know little or nothing about.

END

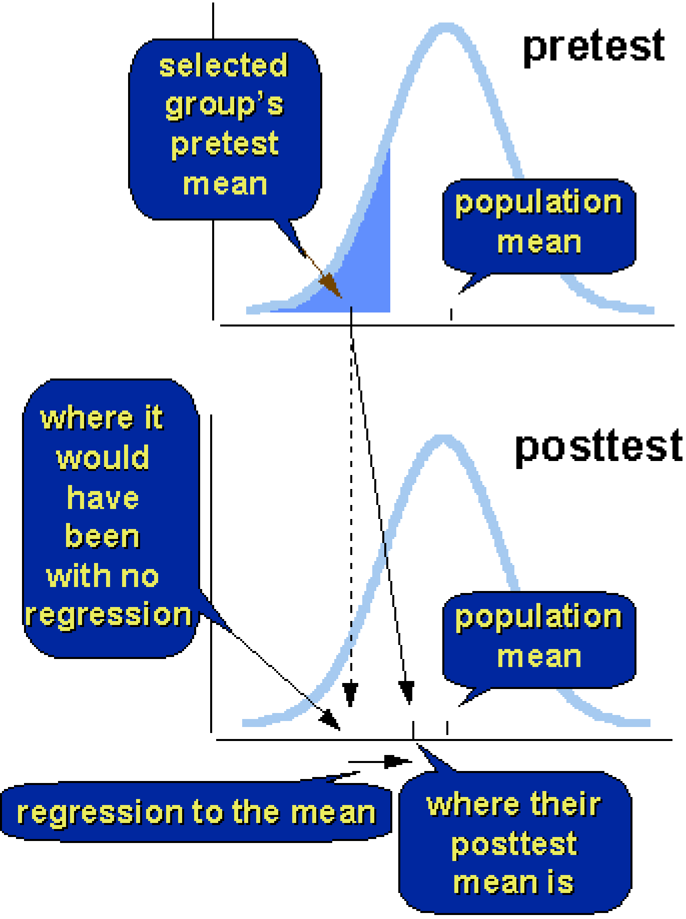

Regression to the Mean

A regression threat, also known as a "regression artifact" or "regression to the mean" is a statistical phenomenon that occurs whenever you have a nonrandom sample from a population and two measures that are imperfectly correlated. The figure shows the regression to the mean phenomenon. The top part of the figure shows the pretest distribution for a population. Pretest scores are "normally" distributed, the frequency distribution looks like a "bell-shaped" curve. Assume that the sample for your study was selected exclusively from the low pretest scorers. You can see on the top part of the figure where their pretest mean is -- clearly, it is considerably below the population average. What would we predict the posttest to look like? First, let's assume that your program or treatment doesn't work at all (the "null" case). Our naive assumption would be that our sample would score just as badly on the posttest as they did on the pretest. But they don't! The bottom of the figure shows where the sample's posttest mean would have been without regression and where it actually is. In actuality, the sample's posttest mean wound up closer to the posttest population mean than their pretest mean was to the pretest population mean. In other words, the sample's mean appears to *regress toward the mean* of the population from pretest to posttest.

Why Does It Happen?

Let's start with a simple explanation and work from there. To see why regression to the mean happens, consider a concrete case. In your study you select the lowest 10% of the population based on their pretest score. What are the chances that on the posttest that exact group will once again constitute the lowest ten percent? Not likely. Most of them will probably be in the lowest ten percent on the posttest, but if even just a few are not, then their group's mean will have to be closer to the population's posttest than it was to the pretest. The same thing is true on the other end. If you select as your sample the highest ten percent pretest scorers, they aren't likely to be the highest ten percent on the posttest (even though most of them may be in the top ten percent). If even just a few score below the top ten percent on the posttest their group's posttest mean will have to be closer to the population posttest mean than to their pretest mean.

Here are a few things you need to know about the regression to the mean phenomenon:

- **It is a *statistical* phenomenon.**

Regression toward the mean occurs for two reasons. First, it results because you asymmetrically sampled from the population. If you randomly sample from the population, you would observe (subject to random error) that the population and your sample have the same pretest average. Because the sample is already at the population mean on the pretest, it is impossible for them to regress towards the mean of the population any more!

- **It is a *group* phenomenon.**

You cannot tell which way an individual's score will move based on the regression to the mean phenomenon. Even though the group's average will move toward the population's, some individuals in the group are likely to move in the other direction.

- **It happens between *any two variables.***

Here's a common research mistake. You run a program and don't find any overall group effect. So, you decide to look at those who did best on the posttest (your "success" stories!?) and see how much they gained over the pretest. You are selecting a group that is extremely high on the posttest. They won't likely all be the best on the pretest as well (although many of them will be). So, their pretest mean has to be closer to the population mean than their posttest one. You describe this nice "gain" and are almost ready to write up your results when someone suggests you look at your "failure" cases, the people who score worst on your posttest. When you check on how they were doing on the pretest you find that they weren't the worst scorers there. If they had been the worst scorers both times, you would have simply said that your program didn't have any effect on them. But now it looks worse than that -- it looks like your program actually made them worse relative to the population! What will you do? How will you ever get your grant renewed? Or your paper published? Or, heaven help you, how will you ever get tenured?

What you have to realize, is that the pattern of results I just described will happen anytime you measure two measures! It will happen forwards in time (i.e., from pretest to posttest). It will happen backwards in time (i.e., from posttest to pretest)! It will happen across measures collected at the same time (e.g., height and weight)! It will happen even if you don't give your program or treatment.

- **It is a *relative* phenomenon.**

It has nothing to do with overall maturational trends. Notice in the figure above that I didn't bother labeling the x-axis in either the pretest or posttest distribution. It could be that everyone in the population gains 20 points (on average) between the pretest and the posttest. But regression to the mean would still be operating, even in that case. That is, the low scorers would, on average, be gaining more than the population gain of 20 points (and thus their mean would be closer to the population's).

- **You can have regression up or down.**

If your sample consists of below-population-mean scorers, the regression to the mean will make it appear that they move ***up*** on the other measure. But if your sample consists of high scorers, their mean will appear to move ***down*** relative to the population. (Note that even if their mean increases, they could be losing ground to the population. So, if a high-pretest-scoring sample gains five points on the posttest while the overall sample gains 15, we would suspect regression to the mean as an alternative explanation [to our program] for that relatively low change).

- **The more extreme the sample group, the greater the regression to the mean.**

If your sample differs from the population by only a little bit on the first measure, there won't be much regression to the mean because there isn't much room for them to regress -- they're already near the population mean. So, if you have a sample, even a nonrandom one, that is a pretty good subsample of the population, regression to the mean will be inconsequential (although it will be present). But if your sample is very extreme relative to the population (e.g., the lowest or highest x%), their mean is further from the population's and has more room to regress.

- **The less correlated the two variables, the greater the regression to the mean.**

The other major factor that affects the amount of regression to the mean is the correlation between the two variables. If the two variables are *perfectly* correlated -- the highest scorer on one is the highest on the other, next highest on one is next highest on the other, and so on -- there will no be regression to the mean. But this is unlikely to ever occur in practice. We know from measurement theory that there is no such thing as "perfect" measurement -- all measurement is assumed (under the true score model) to have some random error in measurement. It is only when the measure has no random error -- is perfectly reliable -- that we can expect it will be able to correlate perfectly. Since that just doesn't happen in the real world, we have to assume that measures have some degree of unreliability, and that relationships between measures will not be perfect, and that there will appear to be regression to the mean between these two measures, given asymmetrically sampled subgroups.

The Formula for the Percent of Regression to the Mean

You can estimate exactly the percent of regression to the mean in any given situation. The formula is:

$$\mathbf{P_{rm} = 100(1 - r)}$$

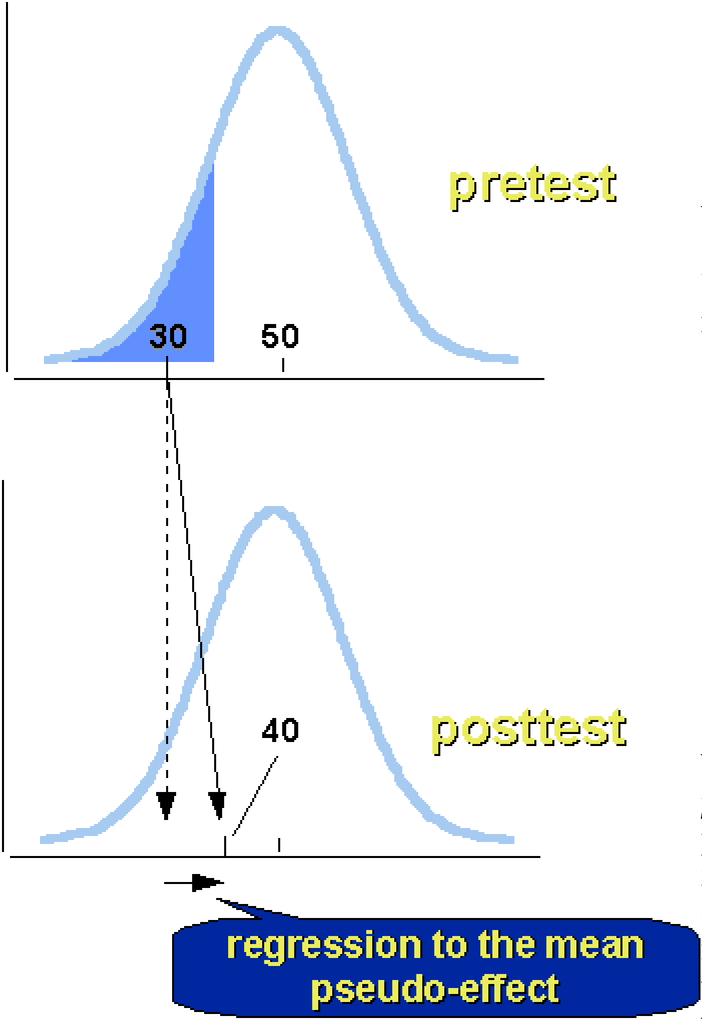

where:

P_{rm} = the percent of regression to the mean
r = the correlation between the two measures

Consider the following four cases:

- if r = 1, there is no (i.e., 0%) regression to the mean
- if r = .5, there is 50% regression to the mean
- if r = .2, there is 80% regression to the mean
- if r = 0, there is 100% regression to the mean

In the first case, the two variables are perfectly correlated and there is no regression to the mean. With a correlation of .5, the sampled group moves ***fifty percent*** of the distance from the no-regression point to the mean of the population. If the correlation is a small .20, the sample will regress 80% of the distance. And, if there is no correlation between the measures, the sample will "regress" all the way back to the population mean! It's worth thinking about what this last case means. With zero correlation, knowing a score on one measure gives you absolutely no information about the likely score for that person on the other measure. In that case, your best guess for how any person would perform on the second measure will be the mean of that second measure.

Estimating and Correcting Regression to the Mean

Given our percentage formula, for any given situation we can estimate the regression to the mean. All we need to know is the mean of the sample on the first measure the population mean on both measures, and the correlation between measures. Consider a simple example. Here, we'll assume that the pretest population mean is 50 and that we select a low-pretest scoring sample that has a mean of 30. To begin with, let's assume that we do not give any program or treatment (i.e., the null case) and that the population is not changing over time on the characteristic being measured (i.e., steady-state). Given this, we would predict that the population mean would be 50 and that the sample would get a posttest score of 30 *if there was no regression to the mean*. Now, assume that the correlation is .50 between the pretest and posttest for the population. Given our formula, we would expect that the sampled group would regress 50% of the distance from the no-regression point to the population mean, or 50% of the way from 30 to 50. In this case, we would observe a score of 40 for the sampled group, which would constitute a 10-point pseudo-effect or regression artifact.

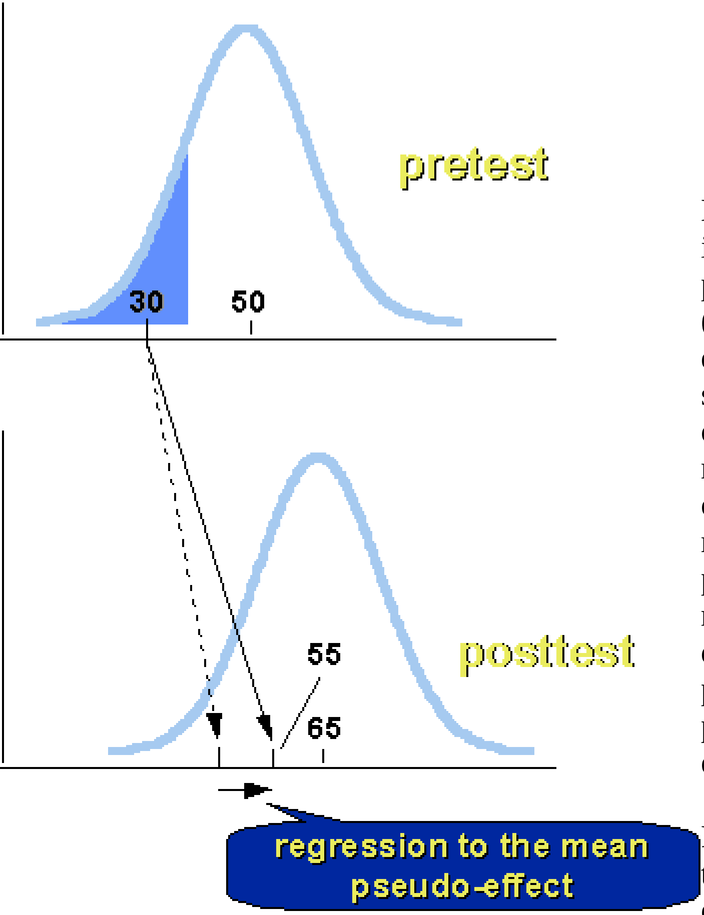

Now, let's relax some of the initial assumptions. For instance, let's assume that between the pretest and posttest the population gained 15 points on average (and that this gain was uniform across the entire distribution, that is, the variance of the population stays the same across the two measurement occasions). In this case, a sample that had a pretest mean of 30 would be expected to get a posttest mean of 45 (i.e., 30+15) if there is no regression to the mean (i.e., r=1). But here, the correlation between pretest and posttest is .5 so we expect to see regression to the mean that covers 50% of the distance from the mean of 45 to the population posttest mean of 65. That is, we would observe a posttest average of 55 for our sample, again a pseudo-effect of 10 points.

Regression to the mean is one of the trickiest threats to validity. It is subtle in its effects, and even excellent researchers sometimes fail to catch a potential regression artifact. You might want to learn more about the regression to the mean phenomenon. One good way to do that would be to simulate the phenomenon. If you're not familiar with simulation, you can get a good introduction in the The Simulation Book. If you already understand the basic idea of simulation, you can do a manual (dice rolling) simulation of regression artifacts or a computerized simulation of regression artifacts.

The big mo

Nov 20th 2008
From The Economist print edition

Why do share prices move relentlessly in one direction?

Illustration by S. Kambayashi

THESE days almost all stockmarkets seem to be falling inexorably. But in more normal times individual stocks are affected by momentum, which is the tendency for popular stocks to keep rising (and for unpopular ones to keep falling). When it comes to shares, what goes up does not always come down—at least in the short term.

The phenomenon has been noted in a wide range of studies and has often been exploited by fund managers, but it has puzzled academics for decades. It is hard to square with the idea that investors are rational. If it were easy to identify which shares were due to go up and which to go down by looking at their previous price movements, why would a rational investor be willing to sell the former group or buy the latter?

Explanations for momentum have thus tended to focus on the idea that investors are irrational. For example, they may be slow to recognise that the fundamentals of a business have changed for the better (or worse). A company may need to beat profits forecasts for two or three quarters before the market is willing to give the stock a premium valuation.

But a new working paper* by researchers at the London School of Economics (LSE) suggests that the momentum effect is still consistent with the idea that investors are rational. The paper's main insight is that most investors do not buy stocks directly, but give their money to fund managers. This creates an agency problem: how do the clients know that the managers are earning their fees?

In the short term, it is difficult to distinguish management skill from luck. Because the index represents the average return of all investors before costs, some managers will beat the index while others will underperform. There is a natural tendency to assume the outperformers are skilful. So the underperformers will lose clients and the outperformers will gain.

The dotcom bubble was a case in point. “Value” investors (who look for stocks that appear cheap by usual measures) ignored the technology industry. They were dumped by clients who gave money to “growth” investors (who look for companies with a promising future) instead. By itself, that pushed up the value of dotcom stocks and made the relative performance of value investors even worse.

In the academics’ view, nobody was being irrational. The clients thought they were picking the best fund managers; the value investors were avoiding overpriced stocks; the growth managers were doing what they were paid to do. After the dotcom bubble popped in March 2000, the same thing happened. Value managers started to outperform, so clients switched their money away from growth stocks. This continued for several years.

By extension, the theory also explains why momentum effects can occur at the industry level. If there is one industry (oil is a case in point) with a low correlation to the market, fund managers will watch their exposure to it very carefully, to avoid the risk of underperforming the index. So if oil shares are doing well, managers will be forced to buy them, pushing up their prices even further.

What is trickier to explain is why the momentum effect ever stops. Academics have found a tendency for a reversion to the mean (outperformers start to falter, underperformers to recover) over longer periods such as three to five years. The LSE authors suggest that momentum effects eventually take prices to such extreme levels that the gains from betting the other way are irresistible. The tricky question is who has the cash to take advantage.

Take the bursting of the dotcom bubble. Value investors were losing clients and so were selling not buying. Growth investors had a mandate from their clients to buy tech stocks and thus had no incentive to switch. And the index-trackers just bought the stocks in the index.

Reversion thus requires a *deus ex machina* in the form of some superrational investor (Warren Buffett, maybe?) or, the authors suggest, fund managers using their own money, who can take advantage of the opportunity provided.

The theory does provide some insights into how momentum might work. But relying on the notion of rational investors seems to complicate matters. If investors are rational, and cannot be sure whether active managers have skill, why do they not just put their money in index-trackers? The idea that investors can occasionally become irrational seems both simpler and intuitively more appealing, especially in the light of recent events.