

Monday, July 03, 2006

The Limits To Learning

Investors Insight Publishing

Introduction

This week we look at mistakes and why we don't learn from them, at least not initially. Good friend James Montier explores the limits to learning we all have and offers some help on how to overcome them. Investors are constantly facing these challenges against their own biases when making sound decisions.

James is the Director of Global Strategy at Dresdner Kleinwort Watterstein, a London and Frankfurt based investment bank. He is also a prolific writer and author of the book "Behavioral Finance - Insights into Irrational Minds and Markets."

Maybe for us to be able to think more "Outside the Box" we must first look within our own "boxes." I hope that you find some insight into your own learning process.

John Mauldin, Editor

The Limits To Learning

By James Montier

Everybody thinks they are experts at learning. After all, most of us have gone through years of university education and emerged on the other side with a piece of paper 'proving' our ability to assimilate information. However, I'm not concerned with book learning; I am far more interested in learning from our own errors and mistakes or, somewhat more accurately, why we often fail to learn from our own past failures.

But first I ought to present just a couple of examples of the evidence we have of people not learning from past mistakes. The first comes from the work of Max Bazerman of Harvard. He regularly asks people the following question:

You will represent Company A (the potential acquirer), which is currently considering acquiring Company T (the target) by means of a tender offer. The main complication is this: the value of Company T depends directly on the outcome of a major oil exploration project that it is currently undertaking. If the project fails, the company under current management will be worth nothing (\$0). But if the project succeeds, the value of the

company under current management could be as high as \$100 per share. All share values between \$0 and \$100 are considered equally likely.

By all estimates, Company T will be worth considerably more in the hands of Company A than under current management. In fact, the company will be worth 50 percent more under the management of A than under the management of Company T. If the project fails, the company will be worth zero under either management. If the exploration generates a \$50 per share value, the value under Company A will be \$75. Similarly, a \$100 per share under Company T implies a \$150 value under Company A, and so on.

It should be noted that the only possible option to be considered is paying in cash for acquiring 100 percent of Company T's shares. This means that if you acquire Company T, its current management will no longer have any shares in the company, and therefore will not benefit in any way from the increase in its value under your management.

The board of directors of Company A has asked you to determine whether or not to submit an offer for acquiring company T's shares, and if so, what price they should offer for these shares.

This offer must be made now, before the outcome of the drilling project is known. Company T will accept any offer from Company A, provided it is at a profitable price for them. It is also clear that Company T will delay its decision to accept or reject your bid until the results of the drilling project are in. Thus you (Company A) will not know the results of the exploration project when submitting your price offer, but Company T will know the results when deciding on your offer. As already explained, Company T is expected to accept any offer by Company A that is greater than the (per share) value of the company under current management, and to reject any offers that are below or equal to this value. Thus, if you offer \$60 per share, for example, Company T will accept if the value under current management is anything less than \$60. You are now requested to give advice to the representative of Company A who is deliberating over whether or not to submit an offer for acquiring Company T's shares, and if so, what price he/she should offer for these shares. If your advice is that he/she should not to acquire Company T's shares, advise him/her to offer \$0 per share. If you think that he/she should try to acquire Company T's shares, advise him/her to offer anything between \$1 to \$150 per share. What is the offer that he/she should make? In other words, what is the optimal offer?

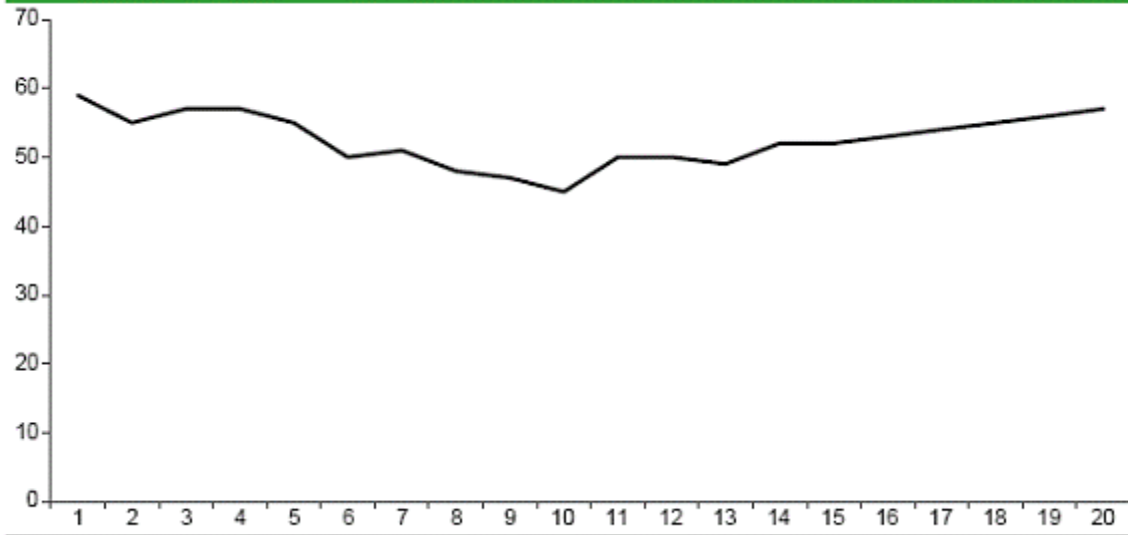
The correct answer is zero. The reasoning is as follows: Suppose the acquirer offers \$60. From the above we know that all points are equally likely, so by offering \$60, Company T is assumed on average to be worth \$30. Given that the company is worth 50% more to the acquirer, the acquirer's expected value is $1.5 * \$30 = \45 . So a bid of \$60 has a negative expected value. Any positive offer has a negative expected value, so the acquirer is better off making no offer.

In contrast to this rational logic, the overwhelming majority of responses fall in the range \$50-\$75. The 'logic' behind this is that on average the company must be worth \$50, thus be worth \$75 to the acquirer, so any price in this range is mutually beneficial. However,

this ignores the rules of the game. Most obviously, the target can await the result of the exploration before accepting or rejecting, and the target will only accept offers that provide a profit.

The first chart below shows 20 rounds of the game. Across twenty rounds, there is no obvious trend indicating that participants learned the correct response. In fact, Ball et al find that only five of the seventy-two participants (MBA students from a top university) learned over the course of the game.

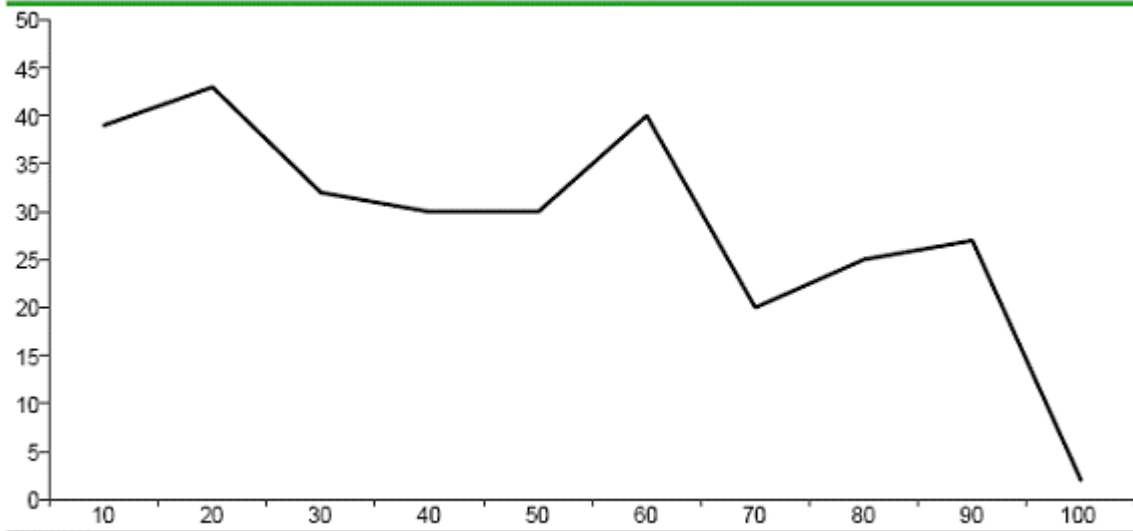
20 rounds of acquiring a company



Source: Ball, Bazerman and Carroll (1991)

The second chart shows the results over a 1000 rounds of the game from a study by Grosskopf and Bereby-Meyer. Players didn't learn from hundreds and hundreds of rounds!

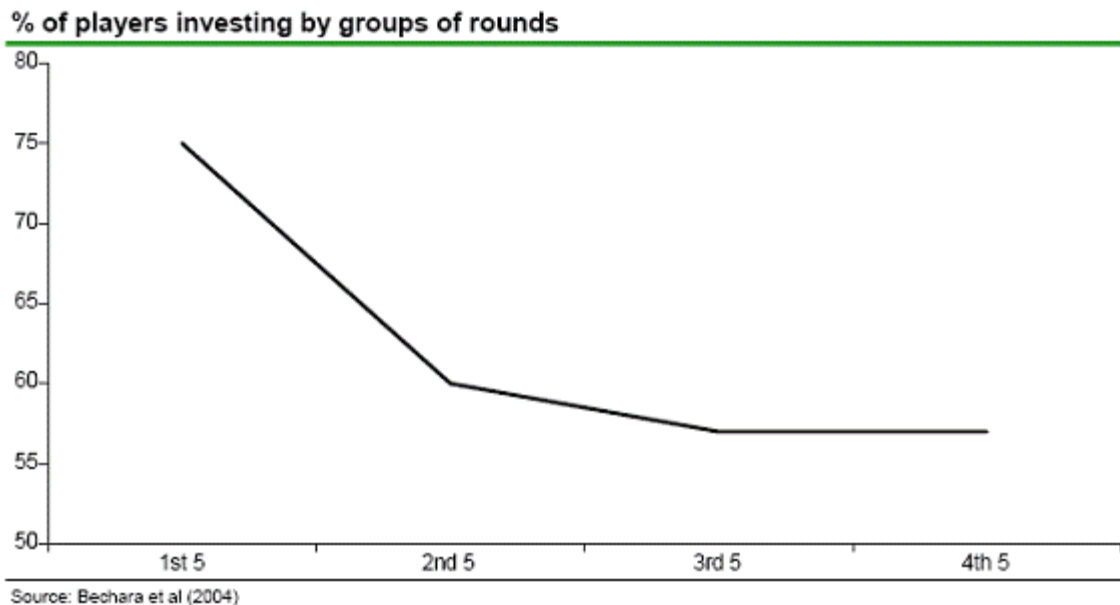
1000 rounds of acquiring a company (blocks of 10 games)



Source: Grosskopf and Bereby-Meyer

The second example of a failure to learn comes from a simple investment game devised by Bechara et al. Each player was given \$20. They had to make a decision on each round of the game: invest \$1 or not invest. If the decision was not to invest, the task advanced to the next round. If the decision was to invest, players would hand over one dollar to the experimenter. The experimenter would then toss a coin in view of the players. If the outcome was heads, the player lost the dollar. If the outcome landed tails up then \$2.50 was added to the player's account. The task would then move to the next round. Overall, 20 rounds were played.

The chart below shows there was no evidence of learning as the game went on. If players learnt over time, they would have worked out that it was optimal to invest in all rounds. However, as the game went on, so, fewer and fewer players continued: they were actually becoming worse as time went on!



I'm sure you can think back and remember many mistakes that you should have learnt from, but didn't (or perhaps I shouldn't judge everybody by my standards). The above is merely setting the scene for our discussion over why people fail to learn. It is the impediments to learning that we now turn our attention towards.

The major reason we don't learn from our mistakes (or the mistakes of others) is that we simply don't recognise them as such. We have a gamut of mental devices all set up to protect us from the terrible truth that we regularly make mistakes.

Self attribution bias: heads is skill, tails is bad luck

We have a relatively fragile sense of self-esteem; one of the key mechanisms for protecting this self image is self-attribution bias. This is the tendency for good outcomes to be attributed to skill and bad outcomes to be attributed to sheer bad luck. This is one of

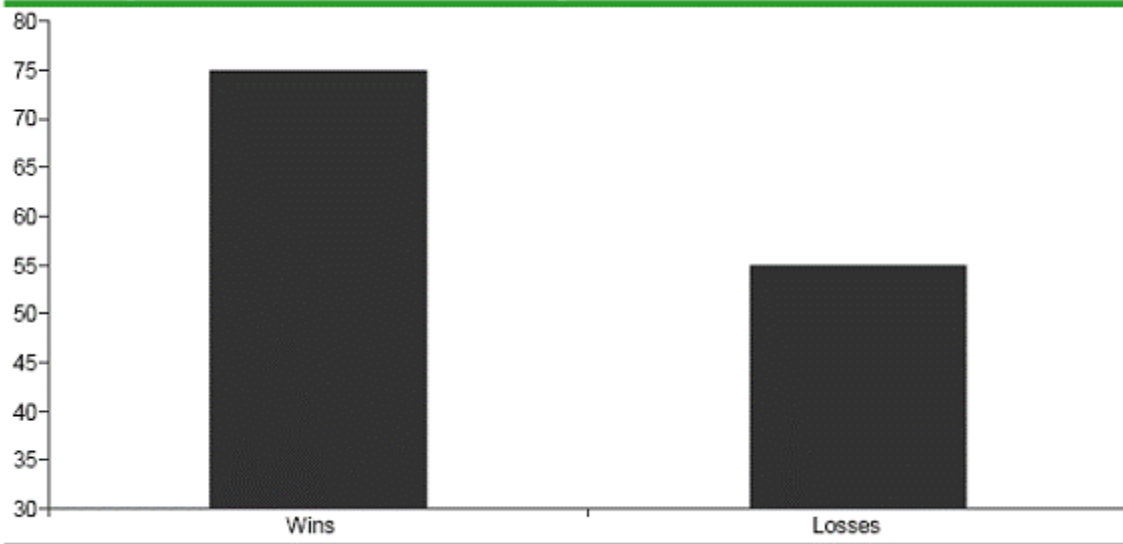
the key limits to learning that investors are likely to encounter. This mechanism prevents us from recognizing mistakes as mistakes, and hence often prevents us from learning from those past errors.

You can't have helped but notice that the football world cup is under way at the moment. Personally I can't stand the sport, but it might just be worth listening to the post match analysis to see how many examples of self attribution one can find.

Lau and Russell examined some 33 major sporting events during the autumn of 1977. Explanations of performance were gathered from eight daily newspapers, giving a total of 594 explanations. Each explanation was measured in terms whether it referred to an internal (something related to the team's abilities) or external factor (such as a bad referee).

Unsurprisingly, self attribution was prevalent. 75% of the time following a win, an internal attribution was made (i.e. the result of skill); whereas only 55% of the time following a loss was an internal attribution made.

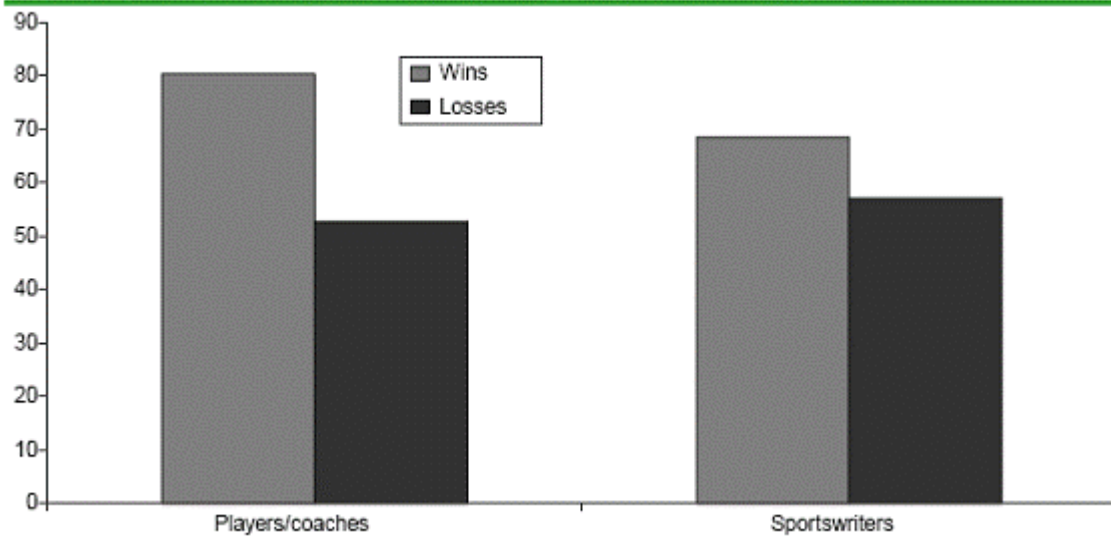
Percentage of internal attributions following wins and losses



Source: Lau and Russell (1980)

The bias was even more evident when the explanations were further categorized as coming from either a player/coach or a sportswriter. Players and coaches attributed their success to an internal factor over 80% of the time. However, internal factors were blamed only 53% of the time following losses. Sportswriters attributed wins to internal factors 70% of the time when it was their home team, and 57% of the time when their home team lost.

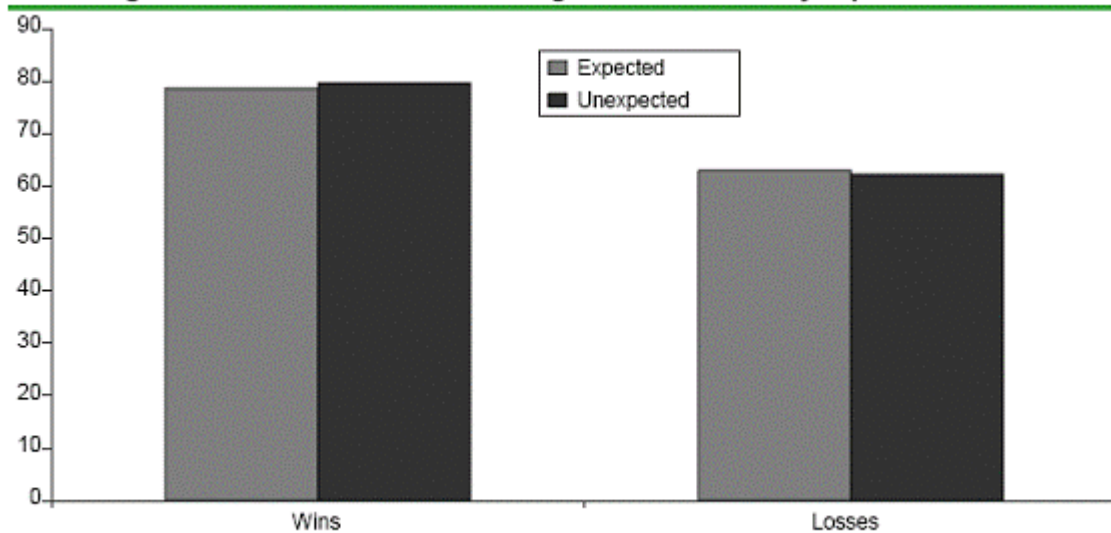
Percentage of internal attributions following wins and losses by function



Source: Lau and Russell (1980)

The expected outcome of the game had no impact on the post match explanations that were offered. Even when one team was widely expected to thrash the other, the attributions of the winners referred to internal factors around 80% of the time, and the attributions of the losers referred to an internal factor 63% of the time.

Percentage of internal attributions following wins and losses by expected outcome



Source: Lau and Russell (1980)

To combat the pervasive problem of self attribution we really need to keep a written record of the decisions we take and the reasons behind those decisions. We then need to map those into a quadrant diagram like the one shown below. That is, was I right for the right reason? (I can claim some skill, it could still be luck, but at least I can claim skill), or was I right for some spurious reason? (In which case I will keep the result because it makes the portfolios look good, but I shouldn't fool myself into thinking that I really knew what I was doing). Was I wrong for the wrong reason? (I made a mistake and I

need to learn from it), or was I wrong for the right reason? (After all, bad luck does occur). Only by cross-referencing our decisions and the reasons for those decisions with the outcomes, can we hope to understand when we are lucky and when we have used genuine skill.

Beating self attribution

	Good outcome	Bad outcome
Right reason	Skill	Bad luck
Wrong reason	Good luck	Mistake

Source: DriKW Macro research

Hindsight bias: I knew it all along

One of the reasons I suggest that people keep a written record of their decisions and the reasons behind their decisions, is that if they don't, they run the risk of suffering from the insidious hindsight bias. This simply refers to the idea that once we know the outcome we tend to think we knew it was so all along.

The best example of this from the investment world is probably the bubble in TMT in the late 1990s. Those who were going around telling people it was a bubble were treated as cretins. However, today there seems to have been an Orwellian re-writing of history, so everyone now thinks they knew it was a bubble (even though they were fully invested at the time).

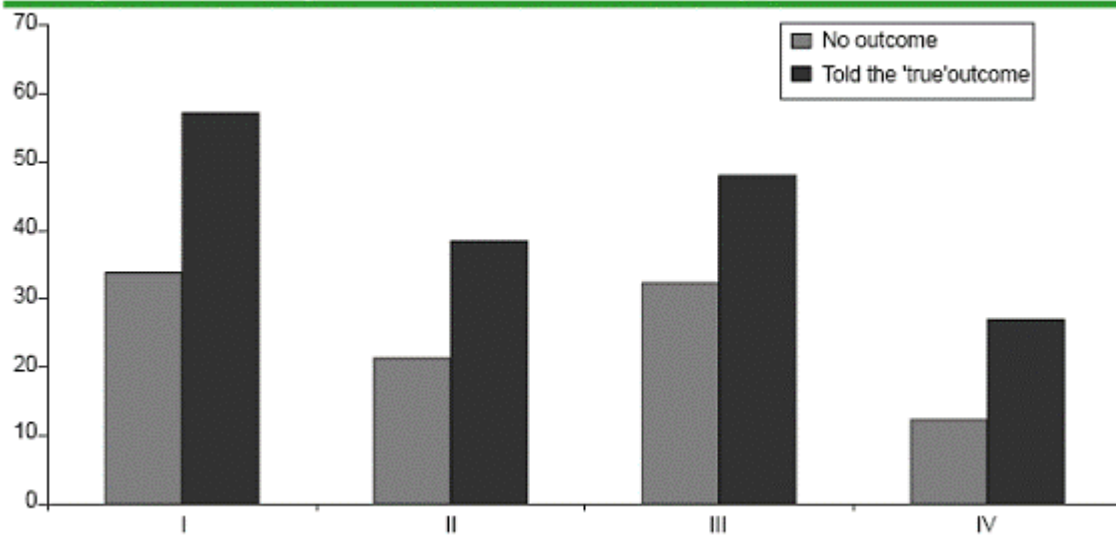
Baruch Fischhoff first noted this strong tendency in 1975. He gave students descriptions of the British occupation of India and problems of the Gurkas of Nepal. In 1814, Hastings (the governor-general) decided that he had to deal with the Gurkas once and for all. The campaign was far from glorious. The troops suffered in the extreme conditions, and the Gurkas were skilled at guerrilla style warfare and few in number, offering little chance for full-scale engagements. The British learned caution only after several defeats.

Having read a much longer version of the above, Fischhoff asked one group to assign probabilities to each of the four outcomes: (i) British victory, (ii) Gurka victory, (iii) military stalemate without a peace settlement, (iv) military stalemate with a peace settlement.

With the other four groups, Fischhoff provided the 'true' outcome, except that three of the four groups received a false 'true' outcome. Again these groups were asked to assign probabilities to each of the outcomes.

The results are shown in the chart below. The hindsight bias is clear from even a cursory glance at the chart. All the groups who were told their outcome was true assigned it a much higher probability than the group without the outcome information. In fact, there was a 17 percentage point increase in the probability assigned once the outcome was known! That is to say, none of the groups were capable of ignoring the ex post outcome in their decision making.

Hindsight bias: knowing the outcome influences choices (%)



Source: Fischhoff (1975)

Hindsight is yet another bias that prevents us from recognising our mistakes. It has been repeatedly found that simply telling people about hindsight and extolling them to avoid it has very little impact on our susceptibility. Rather Slovic and Fischhoff found that the best mechanism for fighting hindsight bias was to get people to explicitly think about the counterfactuals: what didn't occur and what could have lead to an alternative outcome? In experiments, Slovic and Fischhoff found that hindsight was still present when this was done, but it was much reduced.

Skinner's pigeons

An additional problem stems from the fact that our world is probabilistic. That is to say, we live in an uncertain world where cause and effect are not always transparent. However, we often fail to accept this fundamental aspect of our existence. Way back in 1947, B.F. Skinner was exploring the behaviour of pigeons. Skinner was the leader of a school of psychology known as behaviouralism, which held that psychologists should study only observable behaviour, not concern themselves with the imponderables of the mind.

Skinner's theory was based around operant conditioning. As Skinner wrote, "The behavior is followed by a consequence, and the nature of the consequence modifies the organism's tendency to repeat the behavior in the future." A more concrete example may be useful here.

One of Skinner's favourite subjects was pigeons. Skinner placed a series of hungry pigeons in a cage attached to an automatic mechanism that delivered food to the pigeon "at regular intervals with no reference whatsoever to the bird's behaviour". He discovered that the pigeons associated the delivery of the food with whatever chance actions they had been performing as it was delivered, and that they continued to perform the same actions:

One bird was conditioned to turn counter-clockwise about the cage, making two or three turns between reinforcements. Another repeatedly thrust its head into one of the upper corners of the cage. A third developed a 'tossing' response, as if placing its head beneath an invisible bar and lifting it repeatedly. Two birds developed a pendulum motion of the head and body, in which the head was extended forward and swung from right to left with a sharp movement followed by a somewhat slower return. (Superstition in the Pigeon, B.F. Skinner, Journal of Experimental Psychology 38, 1947)

Skinner suggested that the pigeons believed that they were influencing the automatic mechanism with their "rituals" and that the experiment also shed light on human behaviour:

The experiment might be said to demonstrate a sort of superstition. The bird behaves as if there were a causal relation between its behavior and the presentation of food, although such a relation is lacking. There are many analogies in human behavior. Rituals for changing one's fortune at cards are good examples. A few accidental connections between a ritual and favorable consequences suffice to set up and maintain the behavior in spite of many unreinforced instances. The bowler who has released a ball down the alley but continues to behave as if she were controlling it by twisting and turning her arm and shoulder is another case in point. These behaviors have, of course, no real effect upon one's luck or upon a ball halfway down an alley, just as in the present case the food would appear as often if the pigeon did nothing - or, more strictly speaking, did something else. (Ibid.)

Indeed, some experiments by Ono in 1987 showed that Skinner's findings were applicable to humans. He placed humans into the equivalent of Skinner boxes: rooms with a counting machine to score points, a signal light and three boxes with levers. The instructions were simple:

You may not leave the experimental booth... during the experiment. The experimenter doesn't require you to do anything specific. But if you do something, you may get points on the counter. Try to get as many points as possible.

In fact, participants would receive points on either a fixed time interval or a variable time interval. Nothing they did could have influenced the outcome in terms of points awarded. However, Ono recorded some pretty odd behaviour. Several subjects developed "persistent idiosyncratic and stereotyped superstitious behaviour". Effectively they began to try and find patterns to behaviour, such as pulling the left lever four times, and then the right lever twice, and the middle lever once.

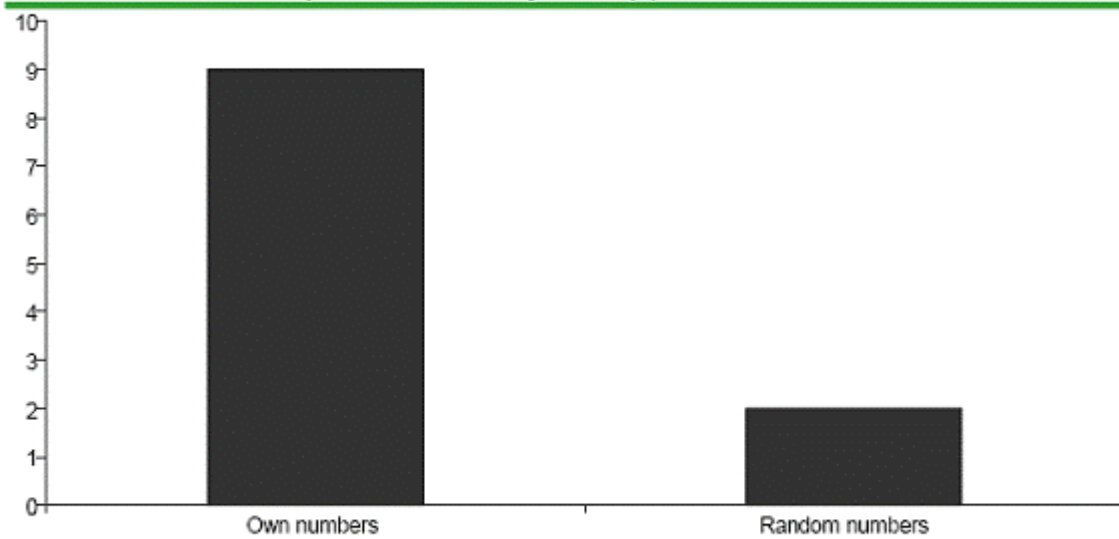
My favourite behaviour was displayed by one young lady in Ono's study. He records, "A point was delivered just as she jumped to the floor (from the table)... after about five jumps, a point was delivered when she jumped and touched the ceiling with her slipper in her hand. Jumping to touch the ceiling continued repeatedly and was followed by more points until she stopped about 25 minutes into the session, perhaps because of fatigue."

Could it be that investors are like Skinner's pigeons, drawing lessons by observing the world's response to their actions? It is certainly possible. The basic failure with the pigeons and Ono's human experiments is that they only look at the positive concurrences, rather than looking at the percentage of the times the strategy paid off, relative to all the times they tried.

Illusion of control

We love to be in control. We generally hate the feeling of not being able to influence the outcome of an event. It is probably this control freak aspect of our nature that leads to us to behave like Skinner's pigeons. My favourite example of the illusion of control concerns lottery tickets from the classic paper by Langer⁸. She asked some people to choose their own lottery numbers, whilst others were just given a random assignment of numbers. Those who chose their own numbers wanted an average \$9 to give the ticket up. Those who received a random assignment/lucky dip lottery ticket wanted only \$2!

The illusion of control: price to sell lottery ticket (\$)



Source: Langer (1975)

Another great example comes from Langer and Roth. Subjects were asked to predict the outcome of 30 coin tosses. In reality, the accuracy of the participants was rigged so that everyone guessed correctly in 15 of the trials, but roughly one-third of the subjects began by doing very well (guessing correctly on the first four tosses), one-third began very badly, and one-third met with random success. After the 30 tosses, people were asked to rate their performance. Those who started well, rated themselves as considerably better at guessing the outcomes than those who started badly.

In their analysis of a wide range of illusion of control studies, Presson and Benassi summarize that the illusion is more likely when lots of choices are available, you have early success at the task (as per above), the task you are undertaking is familiar to you, the amount of information available is high, and you have personal involvement. Large

portfolios, high turnover and short time horizons all seem to be the financial equivalents of conditions that Presson and Benassi outline. Little wonder that the illusion of control bedevils our industry.

Feedback distortion

Not only are we prone to behave like Skinner's pigeons but we also know how to reach the conclusions we want to find (known as 'motivated reasoning' amongst psychologists). For instance, if we jump on the bathroom scales in the morning, and they give us a reading that we don't like, we tend to get off and have another go (just to make sure we weren't standing in an odd fashion¹¹). However, if the scales have delivered a number under our expectations, we would have hopped off the scales into the shower, feeling very good about life.

Strangely enough, we see exactly the same sort of behaviour in other areas of life. Ditto and Lopez¹² set up a clever experiment to examine just such behaviour. Participants were told that they would be tested for the presence of TAA enzyme. Some were told that the TAA enzyme was beneficial (i.e. "people who are TAA positive are 10 times less likely to experience pancreatic disease than are people whose secretory fluids don't contain TAA"), others were told TAA was harmful ("10 times more likely to suffer pancreatic disease").

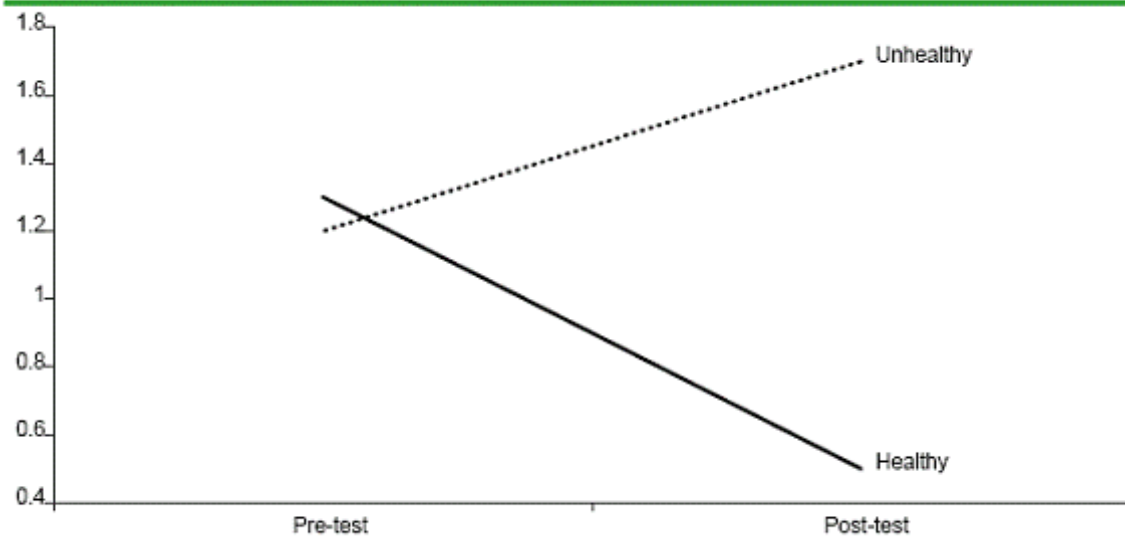
Half of the subjects in the experiment were asked to fill out a set of questions before they took the test, the other half were asked to fill out the questions after the test. In particular two questions were important. The first stated that several factors (such a lack of sleep) may impact the test, and participants were asked to list any such factors that they had experienced in the week before the test. The other question asked participants to rate the accuracy of the TAA enzyme test on a scale of 0 to 10 (with 10 being a perfect test).

The charts below show the results that Ditto and Lopez uncovered. In both questions there was little difference in the answers offered by those who were told having the TAA enzyme was healthy and those who were told it was unhealthy provided they were asked before they were given the result. However, massive differences were observed once the results were given.

Those who were told the enzyme was healthy and answered the questions after they had received the test results, gave less life irregularities and thought the test was better than those who answered the questions before they knew the test result.

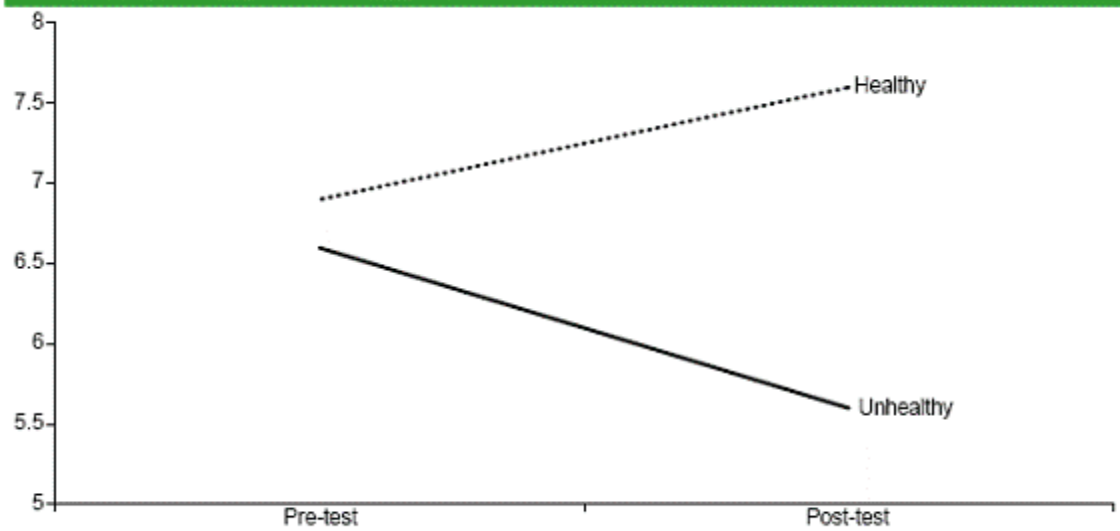
Similarly, those who were told the enzyme was unhealthy and answered the questions after the test results, provided considerably more life irregularities and thought the test was less reliable than those who answered before knowing the test result. Both groups behaved exactly as we do on the scales in the bathroom. Thus, we seem to be very good at accepting feedback that we want to hear while not only ignoring, but actively arguing against, feedback that we don't want to hear.

Number of test-affecting life irregularities cited



Source: Ditto and Lopez (1992)

Perceived accuracy of TAA enzyme test



Source: Ditto and Lopez (1992)

Interestingly, Westen et al found that such motivated reasoning is associated with parts of the brain that control emotion, rather than logic (the x-system, rather than the csystem, for those who have attended one of my behavioural teach-ins). Committed Democrats and Republicans were shown statements from both Bush and Kerry and a neutral person. Then a contradictory piece of behaviour was shown, illustrating a gap between the rhetoric of the candidates and their actions. Participants were asked to rate how contradictory the words and deeds were (on a scale of 1 to 4). An exculpatory statement was then provided, giving some explanation as to why the mismatch between words and deeds occurred, and finally participants were asked to rate whether the mismatch now seemed so bad in the light of the exculpatory statement.

Strangely enough, the Republicans thought that the Bush contradiction was far milder than the Democrats, and vice versa when considering the Kerry contradiction. Similar findings were reported for the question on whether the exculpatory statement mitigated the mismatched words and deeds.

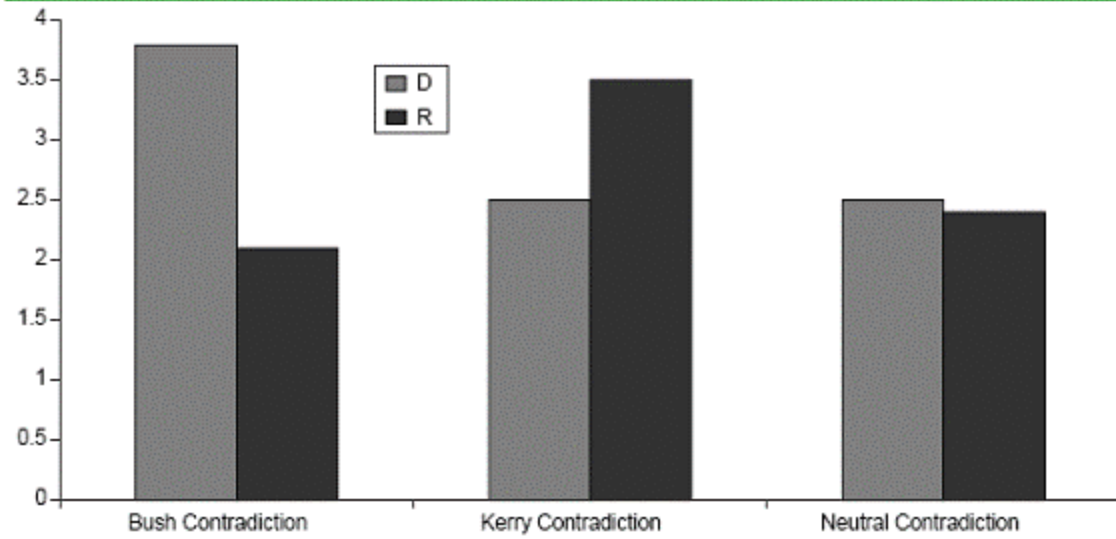
Example of statements

Initial	"First of all, Ken Lay is a supporter of mine. I love the man. I got to know Ken Lay years ago, and he has given generously to my campaign. When I'm President, I plan to run the government like a CEO runs a country, Ken Lay and Enron are a model of how I'll do that" – George Bush, 2000
Contradictory	Mr Bush now avoids any mention of Ken Lay and is critical of Enron when asked
Exculpatory	People who know the President report that he feels betrayed by Ken Lay, and was genuinely shocked to find that Enron's leadership had been corrupt

Source: Westen et al

Westen et al found that the neural correlates of motivated reasoning were associated with parts of the brain known to be used in the processing of emotional activity rather than logical analysis. They note "Neural information processing related to motivated reasoning appears to be qualitatively different from reasoning in the absence of a strong emotional stake in the conclusions reached."

Behavioural ratings of the extent of the contradiction (scale 1 to 4, higher equals worse)



Source: Ditto and Lopez (1997)

Furthermore, Westen et al found that after the emotional conflict of the contradiction has been resolved a burst of activity in one of the brain's pleasure centres can be observed (the ventral striatum). That is to say, the brain rewards itself once an emotionally consistent outcome has been reached. Westen et al conclude "The combination of reduced negative affect... and increased positive affect or reward... once subjects had ample time to reach biased conclusions, suggests why motivated judgments may be so difficult to change (i.e. they are doubly reinforcing)."

Conclusions

Experience is a dear teacher - Benjamin Franklin

Experience is a good teacher, but she sends in terrific bills - Minna Antrim

Experience is the name that everyone gives their mistakes - Oscar Wilde

We have outlined four major hurdles when it comes to learning from our own mistakes. Firstly, we often fail to recognize our mistakes because we attribute them to bad luck rather than poor decision making. Secondly, when we are looking back, we often can't separate what we believed beforehand from what we now know. Thirdly, thanks to the illusion of control, we often end up assuming outcomes are the result of our actions. Finally, we are adept at distorting the feedback we do receive, so that it fits into our own view of our abilities.

Some of these behavioural problems can be countered by keeping written records of decisions and the 'logic' behind those decisions. But this requires discipline and a willingness to re-examine our past decisions. Psychologists have found that it takes far more information about mistakes than it should do, to get us to change our minds.


As Ward Edwards notes:

An abundance of research has shown that human beings are conservative processors of fallible information. Such experiments compare human behaviour with the outputs of Bayes's theorem, the formal optimal rule about how opinions... should be revised on the basis of new information. It turns out that opinion change is very orderly, and usually proportional to numbers calculated from Bayes's theorem - but it is insufficient in amount. A convenient first approximation to the data would say that it takes anywhere from two to five observations to do one observations' worth of work in inducing a subject to change his opinion.

So little wonder that learning from past mistakes is a difficult process. However, as always, being aware of the potential problems is a first step to guarding against them.

Conclusion

Your looking forward to fireworks and kids analyst,



John F. Mauldin

johnmauldin@investorsinsight.com

Disclaimer

John Mauldin is president of Millennium Wave Advisors, LLC, a registered investment advisor. All material presented herein is believed to be reliable but we cannot attest to its accuracy. Investment recommendations may change and readers are urged to check with their investment counselors before making any investment decisions.

Opinions expressed in these reports may change without prior notice. John Mauldin and/or the staffs at Millennium Wave Advisors, LLC and InvestorsInsight Publishing, Inc. ("InvestorsInsight") may or may not have investments in any funds cited above.

PAST RESULTS ARE NOT INDICATIVE OF FUTURE RESULTS. THERE IS RISK OF LOSS AS WELL AS THE OPPORTUNITY FOR GAIN WHEN INVESTING IN MANAGED FUNDS. WHEN CONSIDERING ALTERNATIVE INVESTMENTS, INCLUDING HEDGE FUNDS, YOU SHOULD CONSIDER VARIOUS RISKS INCLUDING THE FACT THAT SOME PRODUCTS: OFTEN ENGAGE IN LEVERAGING AND OTHER SPECULATIVE INVESTMENT PRACTICES THAT MAY INCREASE THE RISK OF INVESTMENT LOSS, CAN BE ILLIQUID, ARE NOT REQUIRED TO PROVIDE PERIODIC PRICING OR VALUATION INFORMATION TO INVESTORS, MAY INVOLVE COMPLEX TAX STRUCTURES AND DELAYS IN DISTRIBUTING IMPORTANT TAX INFORMATION, ARE NOT SUBJECT TO THE SAME REGULATORY REQUIREMENTS AS MUTUAL FUNDS, OFTEN CHARGE HIGH FEES, AND IN MANY CASES THE UNDERLYING INVESTMENTS ARE NOT TRANSPARENT AND ARE KNOWN ONLY TO THE INVESTMENT MANAGER.

Communications from InvestorsInsight are intended solely for informational purposes. Statements made by various authors, advertisers, sponsors and other contributors do not necessarily reflect the opinions of InvestorsInsight, and should not be construed as an endorsement by InvestorsInsight, either expressed or implied. InvestorsInsight is not responsible for typographic errors or other inaccuracies in the content. We believe the information contained herein to be accurate and reliable. However, errors may occasionally occur. Therefore, all information and materials are provided "AS IS" without any warranty of any kind. Past results are not indicative of future results.

We encourage readers to review our complete [legal](#) and [privacy](#) statements.

InvestorsInsight Publishing, Inc. -- 14900 Landmark Blvd #350, Dallas, Texas 75254

© InvestorsInsight Publishing, Inc. 2005 ALL RIGHTS RESERVED