Confounding Changes in Averages with Marginal Effects: Anchoring within Strategic Investment Assessments

Zur Shapira* and J. Myles Shaver**

October 27, 2009

Draft, comments welcome

* Stern School of Business
New York University
40 W 4th St
New York, NY 10012
zshapira@stern.nyu.edu

** Corresponding author:
Carlson School of Management
University of Minnesota
321 19th Avenue South
Minneapolis, MN 55455
mshaver@umn.edu

Author ordering is alphabetical. We appreciate helpful comments from seminar participants at Washington University. We also wish to thank Benjamin Ilhardt for excellent research assistance.

Confounding Changes in Averages with Marginal Effects: Anchoring within Strategic Investment Assessments

Abstract

Profit maximization requires that decision makers assess marginal profits. We demonstrate that decision makers often confound marginal profits with changes in average profits (e.g., changes in return-on-investment). This results in systematic deviations from profit maximization where decision makers forgo profit-enhancing investments that reduce average profits or engage in loss-enhancing investments that decrease average losses. In other words, average profit becomes an anchor by which new investments are assessed. We conduct two decision-making experiments that show this bias and demonstrate it is pronounced when average profit data are accessible or task-relevant. Moreover, we find within-subject effects across experiments, which aids in demonstrating the mechanisms that invoke the bias.

Making strategic investment decisions is not a task that should be taken haphazardly. Managers and MBA students spend time studying appropriate decision criteria such as net present value (NPV), which once calculated properly should guide the decision maker in making the decision. Yet, despite enormous attempts to introduce into corporations' methods of evaluating investment alternatives by discounted cash flows, professionals who deal with capital budgeting still favor inferior techniques such as pay back period calculations (Aggarwal and Gibson, 1989; Klamner, 1972; Mao, 1970; Petty, Scott and Bird, 1975). (Brown, 1970) showed that attempts to convince managers to use formal decision analysis in making decisions have not fared better (see also, Ferraro and Taylor, 2005). One reason that may drive managers to use pay back period calculations based on the internal rate of return (IRR), has to do with the ease of calculation. If one project has a shorter pay back period for the same size of investment than another project, it is ranked higher. The IRR can also be conceived as similar to mortgage rate and hence give a manager a good sense about the potential problems associated with the cash flow of a project.

In this paper we analyze a related but somewhat different aspect of strategic investment decisions. When making a decision whether to invest in a new project, managers often look for two inputs regarding the profitability of such investment. The first is the expected return on the project (i.e., profit) and the second is some scaled version of return, which is actually a measure of average profit (i.e., return on investment). Scaled return data can be beneficial because they facilitate the comparability of investment choices. For example, the preference of two investments that result in $10 million profit likely differs if one requires an investment of $1million and the other an investment of $100 million.

In discussing investment decisions with practicing managers over the years, we

found that managers often equate changes in these scaled profit measures (i.e., changes in ROI) with changes in total profits (*i.e.*, marginal profits) and this causes them to systematically deviate from profit maximization with respect to strategic investment decisions (*e.g.* R&D investments, capital investments, acquisitions). The mechanism that leads to this bias is based on the fact that changes in averages (*e.g.*, return on investment) cause decision makers to anchor on existing performance and reject investments that, although potentially profitable, are relatively less profitable than existing investments.[1] To examine the pervasiveness of the phenomenon we ran a pilot study where we asked managers enrolled in an Executive Master of Business Administration (EMBA) class the following question:

"Assume that you work for a company that currently, after accounting for the cost of capital, has total profits of $1 million on $10 million of invested capital. You are asked to make a recommendation whether the company should invest in Project A, in Project B, in both or in neither.

Project A, which after accounting for the cost of capital, has an expected discounted return of $100,000 on a $1 million investment.
Project B, which after accounting for the cost of capital, has an expected discounted return of $50,000 on a $1 million investment."

Many of the managers who responded to the above choice question recommended investing in Project A but not in B. It is clear that Project A dominates Project B; however, the positive (albeit lower) return on investment (ROI), is a sufficient reason to invest also in B and rational decision makers undoubtedly make the investment. Although managers might shy away from recommending B for many "good" reasons such as

[1] In parallel, this can lead to decisions to make investments that increase total losses yet decrease average losses. Namely, investments that are loss-generating but relatively less loss-generating than current investments.

holding the money for better alternatives or investing the money in riskless alternatives, in the experiments we report in this paper we provided information so as to control for all these reasons, but the phenomenon did not go away. Moreover, such reasons should apply to both projects; however, managers' willingness to invest in A and not B suggests a pronounced anchor on past performance.

The mechanisms that underlie such behavior are likely to have features that are described in the judgment under uncertainty literature. Some of the pioneering work in this area was done by Edwards (1954) who showed that the way people process probabilistic information departs from Bayes rule. Edwards' (1954) major finding was that people do not update their prior probabilities sufficiently when they get additional information in the way prescribed by Bayes rule; he coined the term "conservatism" to describe such behavior.

Tversky and Kahneman (1974) went further and provided several heuristics to describe the ways people process probabilistic information. The heuristic which they named representativeness has been used to show that when making inferences in situations where prior probabilities are given and then additional information is provided, people do not use Bayes rule and tend to focus on the element that represent the problem in a clearer manner. They proposed another heuristic called *anchoring and adjustment* to describe a situation where an unrelated numerical figure is given to people who are then asked to estimate the frequency or probability of a certain event. They showed that the original numerical stimulus operates as a strong anchor from which people do not adjust enough even though the original stimulus is unrelated to the question at hand.

Anchor accessibility. Numerous studies have replicated the effects of anchoring.

A numerical anchor is a value that has the property that any other estimate of a particular uncertain quantity or probability converges to it or gets swept into it. Recent studies (Mussweiler and Neumann, 2001; Epley and Gilovich, 2006) attempt to understand why and when adjustment is insufficient. These researchers argue that when asked to estimate an uncertain numerical value, the ease with which people can come up with an initial estimate is inversely related to the degree to which they would engage in adjustment, which is usually effortful. For instance, Epley and Giolvich (2006) argue that when asked to estimate the year when Washington became president, people can easily generate an anchor, 1776, which seems plausible, and therefore their inclination to make an adjustment will be smaller compared with another case where generating an anchor is effortful or when an anchor is provided by an experimenter which is not easily accessible to the subject.

One of the surprising aspects of the effect of anchoring on judgment is that the anchor does not have to be related to the question at hand. However, what happens when an anchor is related to the particular judgment or decision task? Extrapolating from Epley and Gilovich (2006) who emphasize the importance of anchor accessibility, one can argue that an anchor that is related to the task and is also accessible, is likely to have an even stronger effect on insufficient adjustment. If, in addition to being accessible, the anchor is relevant to the task at hand, it may add more credibility to the anchor and the decision maker may not even attempt to adjust her judgment or decision.

In the domain that we examine, we expect that the degree to which average profit measures are accessible or task relevant will affect the likelihood that decision makers anchor on the current level of average performance and fail to assess marginal

performance. When faced with a single decision, we predict that decision makers will more likely anchor on current average performance when we present profit data as averages (e.g., ROI) versus totals. Namely, accessibility to average profit data invokes the bias.

We also predict that task benefits of using averages will invoke the bias. In particular, we predict that decision makers will be more likely to anchor on average current performance when they have multiple investments to compare. The reason is that averages facilitate the comparison of investments and this heightens the task relevance of using average performance – and hence the likelihood of observing the bias. Therefore, under either of these conditions we predict that average performance will loom much larger in the decision maker's mind and she will be more likely to assess the new project based on changes in average performance versus marginal performance.

Method

Pilot Study

With the general topic in mind, we went through a process of piloting the experiments we wished to perform. The focus of the pilot study was to see if we could document the bias and ensure that subjects were not mislead in our description of the investment. We provided subjects with an investment decision and a set of additional information to aid their assessment. We asked them to make the decision and explain what they did. We also asked them if they had to make assumptions or required additional data to make their assessment. We wanted them to find the material clear, unambiguous, and complete.

We went through four waves to refine and hone the instrument. Three of these waves involved classes of fulltime MBA students and the fourth was a class of executive MBA students. Through these iterations we altered the instrument so that by the end of the process subjects indicated little evidence of ambiguity or need for additional information.

Experiment Overview

To assess the existence of this bias we designed two decision-making experiments. The goal of the experiments was to assess if conditions where decision makers would more likely anchor on existing average performance would systematically lead to decisions based on changes in average profits, rather than marginal profits. Due to the nature of these experiments, it was possible to have subjects take both experiments at one point in time.

Experiment 1a (what we subsequently refer to as the single-decision experiment) has only one manipulation. We present the profile of an investment and a description of the firm's current performance. In one condition subjects see the performance data presented as averages (i.e., return on investment). In the other condition, subjects see the performance data presented as totals (i.e., profits or loss). Appendix 1 presents the instrument. Our expectation is that subjects who see the data as averages will more likely anchor on changes in average performance than those that see the data presented as totals.[2]

As mentioned previously, averages are an important decision aid when decision makers have to compare investments. As a result, experiment 1b presents subjects with

[2] We recognize that there are plausible alternative interpretations for the results we will present (e.g., presenting averages leads subjects to believe this is how they would be evaluated). The section entitled, alternative explanations, discusses why we do not believe such factors drive the results across studies.

two investments to make (we subsequently refer to this as the multi-decision experiment). We design this experiment such that both investments increase firm profits. With the complexity of two investments from which to decide, we can introduce more manipulations and produce a 2x2x2 design. Appendix 2 presents this instrument.

The key manipulation we use is how the two investments relate in average performance (i.e., ROI) to the firm's current performance. Keeping in mind that both investments increase profits, in one manipulation (a) one of the investments equals the firm's current ROI and the other exceeds the current ROI; and in the other manipulation (b) one of the investments equals the current ROI but the other is lower than current ROI. Our expectation is that because subjects have two investments to compare, they will tend to scale the investments (namely, use averages). Therefore, we expect subjects will tend to make both investments when the investments equal and exceed current ROI and will be less likely to make both investments when the investments have equal and lower ROI.

The other two manipulations in the 2x2x2 design provide conditions that will magnify the anchor in the event that subjects are not sensitive to the focal manipulation. These additional two manipulations (a) present the data as ROI versus profit levels and (b) presents investments of equal versus different magnitude. The rationale for manipulating the presentation in terms of ROI or profit levels was – as in the single-decision experiment – to induce the consideration of averages. The rationale of varying the magnitudes of the investments was that investments of different magnitude would be more difficult to compare, subjects would be more inclined to calculate averages to compare the investments, and calculating the averages would induce the anchor.

Experiment

Subjects. Two hundred and eighteen MBA students from five MBA classes at a Midwestern university participated in the experiments. Eighty one were fulltime MBA students, 137 were evening MBA students. The average age of students in these programs was approximately 28 years. Participants had an average of 5.27 years of experience in their last two fulltime positions.

Procedure. The experiment was given to students at the beginning of a class session. Prior to the experiment it was verified that no student was in multiple classes where the experiment was being given. Students were introduced to the experiment by a standardized script prepared for IRB approval of the experiment and were assured that this was not a graded element of their class. Students were able to opt out by simply not completing the experiment. The instructor of the class was not in the room at the time the experiment was done. Once all students had finished the experiments – approximately 15 minutes – the packets were collected and the students were debriefed.

We assembled packets that contained the two experiments. The conditions for the single-decision experiment and for the multi-decision experiment were randomly matched. Likewise, the order of presentation was randomly assigned. Due to the nature of data collection, participation in each cell of each design is not equal. However, it is relatively balanced, which reinforces that randomization underlying the distribution of the instrument.

Experiment 1a: Single Investment Experiment. Of the 218 subjects, 216 answered this question. Table 1 summarizes the responses of the subjects who answered the single-decision experiment. Here, subjects evaluated whether or not to make an investment that

would decrease total profit to $-1.08M from -$1.00M. Clearly, the profit maximizing decision is to forgo this investment. However, this investment increases ROI to -9% from -10%.

Our hypothesis is that subjects who were presented with investment performance as averages (i.e., ROI data) would be more likely to recommend the investment than those who were presented with the profit data. The results in Table 1 confirm our expectation. Eighty-two percent of the subjects forgo the investment when presented the profit data; 70% forgo the investment when presented the ROI data. This difference is statistically significant (Fisher exact test, $p < 0.02$).

Two points regarding the magnitudes in each experimental condition are noteworthy. First, although 30% of the subjects in the ROI condition made the incorrect choice, the majority of subjects did make the correct choice. What we want to emphasize is that subjects in this condition were twice as likely to make the incorrect decision compared to those presented the profit data.

Second, 18% of the subjects made the incorrect decision when presented with the profit data. In order to understand why they would have made an investment that clearly decreased profits, we examined the response sheets of the subjects in this condition that chose to make the investment. In 14 of the 21 cases where the decision was made to recommend the investment, we found that the subjects had calculated the average return on the page and subsequently decided to make the investment because it increased average profit! In essence, we observed that calculating averages and using them as the decision criteria was an important source of incorrect answers in the condition where we did not present the ROI data. It appears that the subjects induced the experimental

mechanism and subsequently make the decision consistent with the experimental mechanism that they induced.

In summary, when we presented subjects with data in the form of averages versus totals they were more likely to make an investment that decreased total profits, yet increased average profits. This is consistent with subjects evaluating changes in average profits rather than marginal effects and anchoring on past average performance.

Experiment 1b: Multi-Decision Experiment. Of the 218 subjects, 217 answered this question. Again, our expectation for the multi-decision experiment was that multiple comparisons would invoke the calculation of averages, because it aids comparability of the two investments. This in turn would enhance the bias and strengthen reliance on the anchor of past average performance.

Recall that in the multi-decision experiment subjects were presented with two investments. In all conditions, both investments increased firm profits. Moreover, in all conditions one of the investments had the same average return as the firm's existing investments. However, one of the experimental manipulations was that the other investment had (a) lower average profits (i.e., ROI) than the firm's existing investments or (b) higher average profits (i.e., ROI) than the firm's existing investments.

In this experiment, the profit maximizing decision is to make both investments. Our hypothesis was that subjects with one investment that had lower average profits would be less inclined to make both investments than subjects with one investment that had higher average profits. Panel A of Table 2 presents the results.

As expected, we found that just under half of the subjects in the condition where one of the investments had lower ROI than existing investments made both investments

(49.5%). In contrast, over 70% of the subjects recommended both investments in the condition where one of the investments had higher ROI than existing investments. This difference is statistically significant (Fisher exact test, $p<0.001$).

It is also noteworthy that in the multi-decision experiment – where comparisons likely invoke the use of averages – the degree to which subjects focus on changes on averages appears more pronounced than in the single-decision experiment. In the single-decision experiment, subjects in the condition where we expected the bias made incorrect choices 30% of the time. In contrast, subjects responding to the multi-decision experiment subjects in the condition where we expected the bias made incorrect choices 50% of the time.

Returning to the results in panel A of Table 2, we were curious why almost 30% of the subjects did not make both investments in the condition where one of the investments had higher ROI than existing investments. In analyzing these responses, the vast majority of subjects chose to undertake the investment with ROI greater than current investments. In addition, many stated that they forwent the other investment because it *did not increase* ROI. For subjects in the condition where one of the investments had lower ROI, the subjects wished to make neither investment in some cases. In other cases, the subjects chose to make only the investment with the same ROI as existing operations.

Our motivation for the other manipulations in this experiment was to try and invoke stronger inducements for subjects to assess changes in averages versus marginal effects. Because we find the effects in the multi-decision structure overall, we do not necessarily expect to find further effects based on the other manipulations. Panel B of Table 2 breaks down the investments where we expect the bias (one of the investments

has a lower ROI) by whether or not the data are presented as averages (i.e., ROI) or profits.

The data show a slightly higher incidence of bias among subjects who had the data presented as averages. Here 54% of the subjects did not make both investments. This is in contrast to 47% of the subjects not making both investments when having the profits presented. This slight difference is not statistically significant (Fisher exact test, $p<0.32$).

Panel C of Table 2 breaks down the investments where we expect the bias (one of the investments has a lower ROI) by whether or not the data the investment magnitudes were equal versus different. Here we find no discernable difference in whether the subjects made both investments. 48% of the subjects made both investments when the investment magnitudes were the same and 51% of the subjects made both investments the investment magnitudes differed (Fisher exact test, $p<0.46$).

Within-subject analyses. A possible concern about interpreting our results is that subjects might differ in their ability to assess investments or in their susceptibility to focus on changes in averages versus marginal effects. Although the experimental conditions were randomized among subjects, there exists the possibility that less adept subjects fall into the conditions where we expect the bias to be invoked and more adept subjects fall into the other conditions.

An advantage of our design is that most subjects (215 of 218) completed both the single-decision experiment and the multi-decision experiment. Therefore, it is possible to assess if the results we previously report hold within subjects across the two experiments.

The data presented in Table 3 are not consistent with the interpretation that subject differences drive the results because there appears to be significant within-subject

variation in showing the bias. Columns of the table reflect whether an individual made the profit maximizing decision in the single-decision experiment. The rows of the table reflect whether an individual made the profit maximizing decision in the multi-decision experiment – provided the subject was in a condition that would subject to the bias.[3]

If the results were driven by subject differences, then we would expect subjects to either answer both questions correctly or answer both questions incorrectly. We would not expect a subject to answer one correctly and the other incorrectly. In other words, we would expect no observations in the off-diagonal (top-right and lower-left cells). A test that the count in the off-diagonal cells is zero rejects this null-hypothesis ($p<0.001$).

A further examination of the within subject effect is to assess if subjects are sensitive to the experimental manipulations designed to invoke the bias regardless of how they answered the other question. Table 4 presents data that are consistent with the conclusion that subjects remain sensitive to the manipulation regardless to how they answered the other question. In all four instances (2 questions x correct or incorrect answers), the pattern of results is consistent with what we presented previously – there is a greater likelihood of incorrect answers in the condition subject to the bias. In three of the four cases, the difference is statistically significant at the 90% level.

More specifically, in panel A of Table 4 we see that subjects who answered the single-decision experiment incorrectly were subject to the manipulation of the multi-decision experiment. For these subjects, 39% made both investments in the condition where we expected the anchor to be stronger and 62% made both investments in the condition where we expected the anchor to be weaker ($p<0.10$, Fisher exact test).

[3] Assessing if a subject that was not in the condition subject to the bias did not exhibit the bias does not add to the analysis.

Likewise, in panel B of Table 4 we see that subjects who answered the single-decision experiment correctly were also subject to the manipulation of the multi-decision experiment. For these subjects, 54% made both investments in the condition where we expected the anchor to be stronger and 73% made both investments in the condition where we expected the anchor to be weaker ($p<0.01$, Fisher exact test).[4]

In panel C of Table 4, we find that subjects who answered the multi-decision experiment incorrectly were subject to the manipulation in the single-decision experiment. For these subjects, 79% forwent the investment when presented the investment data as profits and 61% forwent the investment when presented the investment data as averages (ROI) ($p<0.06$, Fisher exact test).

Finally, in panel D of Table 4, we find that subjects who answered the multi-decision experiment correctly were subject to the manipulation in the single-decision experiment. For these subjects, 83% forwent the investment when presented the investment data as profits and 76% forwent the investment when presented the investment data as averages (ROI). Although the difference is in the predicted direction and of meaningful magnitude, it does not test significant.

Summary. With two paper and pencil decision-making experiments, we find supportive evidence that many subjects anchor investment decisions on the past average performance of investments. In a single decision experiment, as expected, we find that the bias is pronounced when subjects are presented performance data as averages (ROI) versus totals (profits). We also find that the bias appears when subjects are presented with two investment decisions. This is consistent with our expectation that subjects use

[4] The difference in the proportion of subjects sensitive to the manipulation in the multi-decision experiment across panels A and B (39% versus 54%) does not test different.

averages to make comparisons and this invokes the anchor.

An advantage of our research design is that we can assess within-subject effects. We find that our results are not driven by astute/naïve subjects because there is sensitivity to the experimental manipulations within individuals. Moreover, it appears that individuals were sensitive to the manipulations that invoked the bias in one experiment, regardless of how they answered the other experiment.

Discussion

We set out to analyze the possible reasons for the tendency of managers to make decisions on investment in new projects based on changes in ROI (i.e., change in average performance) rather than on NPV (i.e., marginal profits). Several earlier studies documented the tendency to focus on ROI or IRR without providing a mechanism to explain it. We argue that the anchoring heuristic is at work here. Practicing managers are anchored on their firms' average profits so strongly that they tend to ignore the fact that positive NPV projects can increase total profits.

Although many studies of anchoring show that irrelevant anchors can affect people estimates' of uncertain events, we believe that a clear and relevant anchor such as firm's average profits provides a very strong anchor from which managers find it hard to deviate. Furthermore, as argued by Epley and Giolvich (2006) anchors which come easy to mind tend to develop resistance to adjustment due to new data, as compared with anchors that require some effort to come by. Company's average profits are very accessible to managers. Marginal expected profits, even if their expected NPV is positive, may be perceived as less stable and hence a less reliable indicator of a firm's value. Consequently, we argue that managers equate changes in average performance as

marginal effects and may eventually make sub-optimal decisions.

Some examples illustrate the bias central to our arguments and its effects on observed business practices. Consider, for example, a recent decision by Neville Isdell the CEO of the Coca Cola Company to forgo several investments and focus on Coke's core business despite the company's declining market share. Isdel said that "there just aren't many businesses for sale that produce the lush margins-around 30%, some analysts estimate that Coke makes from selling its proprietary concentrate to bottlers." (Frout and Byrnes, 2004). His statement suggests that Coke's current performance provides a strong anchor that has affected his decisions. Namely, one can imagine that investments below 30% would still be profitable to Coke; however they appear not to want to make such investments.

Likewise, the mechanism we identify also provides a way to reconcile the discussion of why managers state that they are unwilling to make acquisitions that dilute earnings per share (Lynch, 1971 as cited in Matsusaka, 1993), even though there is no evidence that doing so leads to superior market reactions on announcement (Matsusaka, 1993). Earnings per share is an average performance measure and by acquiring a target with lower earnings per share, a firm lowers earnings per share of the combined entity. Nevertheless, total profits can often be increased under this situation. Therefore, the decision bias would be reflected in the managerial statement and not in the market reaction (if markets are efficient).

Another example comes from our conversations with the Chairman and CEO of an S&P midcap company. He mentioned that "no-brainer" decisions were the high return investments. However, they were often small investments that would only marginally

affect his firm's overall profits. The difficult investment decisions were large dollar investments that returned close to the cost of capital. These investments had the ability to have large profit impacts. However, they had to be managed carefully.

Finally, an example of the prevalence of the effect we demonstrate was reflected in the comments of two Executive MBA students who participated in the pilot study. In that study, we also asked subjects, to further describe if there would be other considerations in selling their decision to the board of directors or their senior management. Interestingly, two subjects mentioned that although they would want to make the investment that increased total profits yet decreased average profits, they would not recommend the project to senior management because they didn't think they could convince senior management that this was the correct thing to do. Apparently, they had experiences similar to the Coke example where there was a focus on only considering projects that enhanced average profits.

Alternative explanations

Before concluding, we wish to highlight two alternative explanations for why managers would avoid making investments that increased profits, yet decreased average profits. Although these reasons are valid and consistent with profit maximization, we also discuss why we do not expect them to hold in our experimental setting.

First, managers might face investment constraints, such as limited capital or managerial resources to oversee projects. Here, optimizing managers would rank investments on their contribution to profitability and forgo profitable investments once constraints became binding. We do not believe this effect drives our results for the following reasons. (i) We are explicit in the instructions that such a constraint is not

binding. (ii) To the extent that subjects did not read or believe that constraints were not binding, we find that subjects equate the constraint with current average profits. This reflects the anchoring mechanism central to our arguments; however, there is no reason to believe that this would necessarily be the point where the constraint binds.[5] (iii) Experiment 1 reflects a loss condition where the investment constraint would not be an alternative mechanism. With the loss-making investment in Experiment 1, it would not make sense to make this investment whether the firm was constrained or not.

Second, firms might have better non-investment alternatives than low profit investments. Here, firms would forgo profitable investments to find other uses for the capital. We do not believe this effect drives our results because we were explicit in the non-investment alternatives. Moreover, to the extent this trade-off occurred, again would appear to occur at the level of current average profits. There is no reason to believe that this would necessarily be the point where the alternative becomes more attractive and is consistent with the anchor of using changes in average profits.

Another alternative explanation not consistent with firm profit maximization but manager self-interest is that when we present average profits in experiment 1a this primes subjects into thinking that average profits are what the company assesses or rewards. Therefore, making investments based on average profits reflects that decision makers believe they would be assessed on this metric. There are two reasons why we do not believe this is the primary determinant for the results that we find. First, in the multi-decision experiment, subjects were not more likely to make the incorrect judgment when

[5] Average profits would reflect the payoff cutoff for current investments only if all investments have the same average return. Consider a situation where the average payoff was 11%. Unless all investments pay 11% return then at least some investment have a return less than 11%. For example, if the firm has 3 investments with an average 11% return and one returns 12%; then at least one of the other investments must return less than 11%.

presented with average versus total profit figures (Table 2C). Here it appeared the task advantages to using averages drive the result. If presenting averages lead subjects to believe that average profits are the performance metric they should assess, then we would expect to see this across experiments – and we do not. Second, in the single-decision experiment we observed that some subjects calculated averages when presented total profit figures and subsequently make the incorrect decision. Apparently it was the use of averages – but not being primed to use averages – that drove the outcome.

Conclusion

"Average" profitability measures such as ROI play an important and legitimate role in managerial decision making because they facilitate the comparison of investment alternatives. However, changes in average profit measures are not marginal effects and attention to changes in averages can lead to systematic deviations from profit maximization. We demonstrate how focusing on changes in averages causes managers to anchor investment decisions on current performance.

We provide experimental evidence that MBA students – many of whom were practicing managers – were susceptible to this anchoring bias when asked to make a strategic investment decision. In the case of making an isolated decision, we find the bias more pronounced when decision makers have average performance metrics presented to them. We also find that the bias is pronounced when decision makers have to compare investment alternatives. Moreover, we find evidence that our experimental manipulations showed within-subject effects. This aids in ruling out alternative explanations based on subject skill as the driver of the results that we find.

Our findings suggest that the ways in which managers process investment

decisions can lead to systematic biases in the outcomes. Anchoring has been demonstrated in many studies of managerial decision making. For example, Lant and Shapira (2008) show that the term expectation is a major organizing concept for professional economists; whereas, managers tend to focus on targets and aspirations. Targets are often set in a relative manner where a company's performance is evaluated in relation to historical profitability (Baum et. al., 2005; Greve, 2000). Recognizing that random fluctuations in the environment can lead to variations in performance over time, average performance is perceived as a much more stable and valid indicator of performance. Indeed, many companies advertize their average performance over different periods of time to help investors evaluate their value. This information is very accessible to managers and potentially enhances its role as an anchor when and evaluating future projects.

As we show, anchoring on average performance is very powerful due to its accessibility and relevance. Anchoring has been recently shown to affect the way consumers pay their credit card debt (Stewart, 2009) and the way consumers and professional insurance agents evaluate the value of alternative insurance policies (Shapira and Venezia, 2008). Our study suggests that anchoring affects strategic decision investments as well.

We believe that better understanding the psychological processes underlying the ways managers make decisions within their job context can help explain strategy choices. This has benefits beyond enhancing our understanding of the strategic decision making process. Understanding systematic determinants of strategy choice is also important to accurately discern how firm strategies affect performance (e.g., Shaver, 1998).

References

Aggarwal, R. and C. Gibson. (1989). *Discounting in Financial Accounting and Reporting: Issues and Literature.* Morristown, NJ: Financial Executives Research Foundation.

Baum., J.A.C., Rowley, T.J., Shipilov, A.V., & Chuang, Y.T. (2005). Dancing with Strangers: Aspiration Performance and the Search for Underwriting Syndicate Partners. *Administrative Science Quarterly*, 50(4): 536-575.

Brown, R. (1970). "Do Managers find Decision Theory Useful?" *Harvard Business Review*, May-June, 78-79.

Edwards, W. (1954). "The Psychology of Decision Making." *Psychological Bulletin*, 51, 380-417.

Epley, N. and Giolvich, T. (2006). "The Anchoring and Adjustment Heuristic." *Psychological Science*, 17, 311-318.

Ferraro, P. and O. Taylor (2005). "Do economists recognize an opportunity cost when they see one? A dismal performance form dismal science." *Contributions to Economic Analysis & Policy*, 4 (1) no. 7.

Foust, D. and N. Byrnes (2004). "Gone flat: The good old days weren't as good as you thought. That's one more reason Coke's latest CEO needs a bold new formula." *Business Week*, Dec. 20.

Greve H.R. 2002. Sticky aspirations: Organizational time perspective and competitiveness. *Organization Science*, 13(1): 1-17.

Klamner, T. (1972). "Empirical evidence of the adoption of sophisticated capital budgeting techniques," *Journal of Business, 35,* 391-397.

Lant, T. and Shapira, Z. (2008). "Managerial reasoning about aspirations and expectations." *Journal of Economic Behavior and Organizations*, 66, 60-73.

Lynch, H.H. 1971. *Financial performance of conglomerates*. Division of Research, Harvard Business School.

Mao, J. (1970). "Survey of capital budgeting: Theory and practice," *Journal of Finance, 25,* 349-360.

Matsusaka, J. G. 1993. Takeover motives during the conglomerate merger wave. *RAND Journal of Economics*, 24(3): 357-379.

Mussweiler, T. and Strack, F. (2001). "The use of Category and Exemplar Knowledge in the Solution of Implausible Anchors." *Social Cognition*, 19, 145-160.

Petty, J., D. Scott, and M. Bird. (1975). "The capital expenditure decision making process of large corporations," *The Engineering Economist, 20,* 159-172.

Shaver, J. M. (1998). "Accounting for endogeneity when assessing strategy performance: Does entry mode choice affect FDI survival?" *Management Science*, 44, 571-585.

Shapira, Z. and Venezia, I. (2008). "On the preference for full-coverage policies: Why do people buy too much insurance?" *Journal of Economic Psychology*, 29, 747-761.

Stewart, N. (2009). "The cost of anchoring on credit-card minimum payments." *Psychological Science*, in Press.

Tversky, A. and Kahneman, D. (1974). "Judgment under Uncertainty: Heuristics and Biases." Science, 185, 1124-1131.

Table 1 – Single-decision experiment

The correct answer for this investment decision is not to make the investment. The investment increases total losses, yet decreases average losses.

	Presented as total profit. Total profit becomes $-1.08M from $-1.00M.	**Presented as average profit (*i.e.* ROI). ROI becomes -9% from -10%.**
Recommend not to make the investment (percentage)	93 (81.6)	71 (69.6)
Recommend making the investment (percentage)	21[a] (18.4)	31 (30.4)

n=216

p=0.029, one-tailed Fisher exact test

[a] Of the 21 subjects who chose to make the investment and were presented the data as total profits, 14 explicitly showed that they calculated average profit (*i.e.*, ROI) before making the incorrect decision.

Table 2 – Multi-decision experiment

The correct answer under all conditions for this investment decision is to make both investments because every investment increases total profits.

A. Comparison of the condition where the investments equal and increase average profits versus equal and decrease average profits.

	One investment decreased average profits, the other equaled average profits	**One investment increased average profits, the other equaled average profits**
Recommend making both investments (percentage)	52 (49.5)	79 (70.5)
Recommend not to make both investments (percentage)	53 (50.5)	33 (29.5)

n=217
p<0.001, one-tailed Fisher exact test

B. Of the investments where one investment decreased average profits and the other equaled average profits – analysis based on whether the data were presented as totals versus averages.

	Presented as total profits	**Presented as average profit (*i.e.* ROI)**
Recommend making both investments (percentage)	28 (52.8)	24 (46.2)
Recommend not to make both investments (percentage)	25 (47.2)	28 (53.9)

n=105
p<0.32, one-tailed Fisher exact test

C. Of the investments where one investment decreased average profits and the other equaled average profits – analysis based on whether the data were presented as totals versus averages.

	Investment sizes the same across investments	**Investment sizes differed across investments**
Recommend making both investments (percentage)	26 (48.2)	26 (51.0)
Recommend not to make both investments (percentage)	28 (51.8)	25 (49.0)

n=105
p<0.46, one-tailed Fisher exact test

Table 3 – Within subject accuracy across experiments:

Frequencies of correct answers across both experiments
(conditioned on the multi-decision experiment being in the condition subject to bias)

	Answered single-decision experiment correctly	**Answered single-decision experiment incorrectly**	**Total**
Answered multi-decision experiment (with one investment decreasing ROI) correctly	42 (40.4)[a]	10 (9.6)	52
Answered multi-decision experiment (with one investment decreasing ROI) incorrectly	36 (34.6)	16 (15.4)	52
Total	78	26	104

[a] Number in brackets is the percentage of all 104 observations.

$p<0.001$, of the test that subjects either made: (a) both decisions correctly or (b) both decisions incorrectly (*i.e.*, that the count in the off-diagonal cells equals zero)

Table 4 – Within subject accuracy across experiments

A. Responses to the multi-decision experiment of the subjects who answered the single-decision incorrectly

	One investment decreased average profits, the other equaled average profits	**One investment increased average profits, the other equaled average profits**
Answered multi-decision experiment correctly (percentage)	10 (38.5)	16 (61.5)
Answered multi-decision experiment incorrectly (percentage)	16 (61.5)	10 (38.5)

n=52 p=0.082, one-tailed Fisher exact test

B. Responses to the multi-decision experiment of the subjects who answered the single-decision correctly

	One investment decreased average profits, the other equaled average profits	**One investment increased average profits, the other equaled average profits**
Answered multi-decision experiment correctly (percentage)	42 (53.9)	62 (72.9)
Answered multi-decision experiment incorrectly (percentage)	36 (46.2)	23 (27.1)

n=163 p=0.009, one-tailed Fisher exact test

C. Responses to the single-decision experiment of the subjects who answered the multi-decision incorrectly

	Presented as total profits	**Presented as average profit (*i.e.* ROI)**
Answered single-decision experiment correctly (percentage)	33 (78.6)	26 (60.5)
Answered single-decision experiment incorrectly (percentage)	9 (21.4)	17 (39.5)

n=85 p=0.057, one-tailed Fisher exact test.

D. Responses to the single-decision experiment of the subjects who answered the multi-decision correctly

	Presented as total profits	**Presented as average profit (*i.e.* ROI)**
Answered single-decision experiment correctly (percentage)	59 (83.1)	45 (76.3)
Answered single-decision experiment incorrectly (percentage)	12 (16.9)	14 (23.7)

n=130 p=0.227, one-tailed Fisher exact test.

Appendix 1: Single-decision

Two conditions: (A) Data presented as totals: (B) Data presented as averages

The following is a decision that we would like you to assess:

You work for a company that currently, after accounting for the cost of capital, has [(A) total losses of $1 million/ (B) a return of -10%] on $10 million of invested capital.

You are asked for your opinion about the following project that is being considered. The project will require $2 million of capital investment. After making this investment and accounting for the cost of capital, the expected performance of the firm will be [(A) losses of $1.08 million/ (B) a return of -9%] on $12 million of invested capital.

Please recommend if the company should invest in this project.

The following information might aid your assessment.

- The risk profile of this project does not materially differ from the risk profile of your existing business.
- This is the only project that you have to consider at this time.
- Any decision that you make is not expected to aid or hinder your ability to make further investments should the opportunity arise.
- Any cash that the company does not use to make investments will be invested in liquid assets that have returns of about 1% (and experts do not think that the return for such instruments will materially change in the future).

Explain your decision (feel free to add any qualifications to your answer).

Appendix 2: Multi-decision experiment

For presentational clarity, we present only one condition of the 2x2x2 design. The condition is: (a) one investment equals current ROI, one investment is greater than current ROI, (b) data presented as totals, (c) investments are of different size. The complete instrument is available upon request.

The following is a strategic investment decision that we would like you to assess:

You work for a company that currently, after accounting for the cost of capital, has total profits of $1 million on $10 million of invested capital.

You are asked for your opinion about the following two projects that are being considered.
Project A, which after accounting for the cost of capital, has an expected discounted return of $150,000 on a $1 million investment.
Project B, which after accounting for the cost of capital, has an expected discounted return of $100,000 on a $1 million investment.

Please recommend if the company should do Project A, Project B, both, or neither.

The following information might aid your assessment.

- Making either of these investments does not affect your ability to make the other.
- The risk profiles of these projects do not materially differ from each other or from the risk profile of your existing business.
- These are the only projects that you have to consider at this time.
- Any decision that you make is not expected to aid or hinder your ability to make further investments should the opportunity arise.
- Making either of these investments is not expected to affect the profitability of your existing operations.
- Any cash that the company does not use to make investments will be invested in liquid assets that have returns of about 1% (and experts do not think that the return for such instruments will materially change in the future).

Explain your decision (feel free to add any qualifications to your answer).