

Individual differences in anchoring: Numerical ability, education and experience.

Matthew B. Welsh¹, Paul Delfabbro², Nicholas R. Burns² and Steve H. Begg¹
([matthew.welsh, paul.delfabbro, nicholas.burns, steve.begg] @adelaide.edu.au)

¹Australian School of Petroleum, ²School of Psychology,

University of Adelaide

Adelaide, SA, 5005 Australia

Abstract

Anchoring is a well-known effect leading to bias in estimation in various decision-making contexts. A question, however, is whether individuals with greater numerical and academic ability would be less prone to this effect than others because of greater ability to discern the value being estimated. In light of growing interest in the role of individual differences in bias susceptibility, anchoring was examined in a simulated poker-like card game, using people with varying levels of academic achievement and psychometric reasoning scores. The results showed that anchoring susceptibility was unrelated to education levels, but negatively associated with numerical reasoning and cognitive reflection scores. This result, however, was mediated by task expertise because participants with higher cognitive abilities were those more likely to display improvements in anchoring performance over the course of the experiment.

Keywords: anchoring, individual differences, numerical ability, experience.

Introduction

Anchoring (Tversky & Kahneman, 1974) describes a robust effect in which the estimates people make are affected by other numbers that they have recently seen. Specifically, people tend to anchor on such numbers and fail to adjust sufficiently away from them when making estimates. This has been shown to affect expert as well as naïve estimators (Northcraft & Neale, 1987); and can be influenced by both relevant and irrelevant anchor values (Thorsteinson, Breier, Atwell, Hamilton, & Privette, 2008) and by obviously wrong anchors (Quattrone, et al., 1984).

Although of theoretical interest in its own right, anchoring also has practical consequences in applied settings. For example, in oil and gas exploration, ‘analogue’ data (i.e. data from a location judged to be analogous to the current location in some way) is regularly used as a starting point for discussions regarding the probability of making a discovery and on assessments of its likely size, value and cost to develop; therefore, anchoring can have a significant impact on decisions (Bratvold, Begg, & Campbell, 2002). As such, considerable time and effort is dedicated to making people aware of anchoring in decision making courses offered at university and in industry settings.

The efficacy of this in reducing susceptibility, however, is doubtful. Research into anchoring has shown the effect to be highly resistant to awareness-based debiasing (Chapman & Johnson, 2002; Welsh, Begg, & Bratvold, 2006). As a result there is increasing interest in whether people’s susceptibility to anchoring might be related to individual differences in cognitive abilities (Bergman, Ellingsen, Johannesson, &

Svensson, 2010; Oechsler, Roider, & Schmitz, 2009; Stanovich & West, 2008) as, if this is the case, decision makers can be *selected* for their resistance to this bias.

Individual Differences

Plausibly, people with greater expertise in a particular area of decision-making should be less prone to biases such as anchoring. However, there is clear evidence to suggest that experts as well as novices are affected by anchors (Northcraft & Neale, 1987).

In their study, Northcraft and Neale divided participants into two groups (expert and non-expert) based upon whether or not they were employed in real-estate. Participants were then asked to value houses after being shown a ‘listing price’ that acted as the anchor. All participants’ estimates were affected by the anchoring values. The researchers also demonstrated that less reasonable anchors had less impact on the responses provided by their non-experts than more reasonable anchors. Somewhat surprisingly, though, they did not examine whether this effect was also observed in experts. In particular, they did not consider the possibility that ‘less reasonable’ anchors may have had an even lesser impact on the responses of experts.

Another limitation of Northcraft and Neale’s (1987) study was that the division of their sample into expert and non-expert was done entirely on the basis of whether or not a participant was employed in real estate and this, potentially, failed to capture more refined differences that could have been considered. For example, it might have been possible to have used years of experience as a (poor) proxy for expertise (for a discussion of the problems in defining expertise, see Malhotra, Lee, & Khurana, 2005).

Alternately, one could follow on from recent work by Frederick (2005), itself building on a tradition of work by Stanovich and West (see, e.g. 1998, 2008), that has shown cognitive abilities to be predictive of susceptibility to a range of decision making biases.

Stanovich and West (2008), in fact, concluded that anchoring was one bias *not* lessened by higher cognitive abilities, a result replicated by Oechsler et al (2009); however, a more recent experiment (Bergman, et al., 2010) found that people scoring higher on Frederick’s (2005) CRT measure (a measure of ‘cognitive reflection; i.e., how likely a person is to engage rational rather than intuitive reasoning) and on a general cognitive ability test were less susceptible to anchors. These inconsistent findings, combined with the earlier insights of Northcraft and Neale (1987), suggest a need for further consideration of the

association between expertise, cognitive ability and anchoring effects.

One possibility is that cognitive ability, in and of itself, plays no role in reducing susceptibility to anchoring but instead acts only as a mediating factor in the development of expertise. If this is the case, then this would predict that the relationship between anchoring and cognitive ability be visible only sometimes (where expertise has been developed).

Research Goals

The first aim of this project was to examine whether increased expertise (loosely defined here as greater experience with a specific task) is associated with decreased susceptibility to bias resulting from anchoring. Second, we intended to establish whether cognitive ability was related to anchoring susceptibility, or expertise, or both. Finally, we were interested to see whether educational level predicted bias susceptibility – as university courses seem the most likely place for a person to have previously encountered the concept of anchoring.

Recognizing the difficulties in defining expertise within any given field, we also wanted to create a task on which we could measure participant's actual expertise so that this could be compared with self-rated expertise. For this reason, we chose a card-game with similar rules to poker (see below). This task enabled us to run a large number of trials and calculate the exact probabilities that the participants would be estimating. It also made it possible to observe people's actual expertise (as reflected in their task performance) and whether this was related to how much prior experience they had with games of this nature.

Method

Participants

Participants were 102 university students and members of the general public, recruited via posters and research participation email lists from around the University of Adelaide. The sample had a mean age of 22.5 (SD = 4.89) and consisted of 34 males and 68 females. All participants received \$50 for their completion of a three-hour battery of tasks including those described in this paper.

Materials and Procedure

Participants completed an online questionnaire, which included gathering demographic details, prior to coming in to the laboratory for testing.

Demographics. In addition to age and sex, a number of further measures were gathered as part of the survey of the sample. These included the level of education on a 7-point scale indicating various levels of academic achievement. As a result of the small sizes of several of these categories, responses were recoded on a 3 point scale: 1, have not attended university ($n = 22$); 2, current university student ($n = 58$); and 3, university graduate ($n = 22$).

Additionally, participants were asked to rate their familiarity with poker and other card games (again a 7 point scale ranging from 1, no familiarity, to 7, very familiar).

Cognitive Reflection Task (CRT). As described by Frederick (2005), the CRT is designed to measure a person's level of "cognitive reflection"; that is, how likely they are to engage rational and reflective System II reasoning rather than relying on the fast and intuitive System I reasoning that leads to bias (Stanovich & West, 2000). Participants answer three questions and are scored either right or wrong. CRT score is simply the number of questions that a person gets right – from 0 to 3.

In addition to predicting susceptibility to a number of biases, this measure shows a strong relationship with educational level in Frederick's (2005) data.

Numerical Abilities Test (NAT). A short version of the Numerical Abilities scale from the Differential Aptitude Test (Bennett, Seashore, & Wesman, 1989) was prepared, asking 12 rather than 48 questions and restricting time to 9 minutes rather than 36. A participant's score was simply the number of questions they answered correctly.

This task was computerized as a Matlab GUI and conducted in the laboratory immediately prior to the anchoring task described below.

Anchoring Task. The anchoring task was designed as a card game based on poker but utilizing a deck of 16 cards consisting of the cards 1 through 4 in four suits (blue, red green and yellow). Three cards comprised each hand.

The value of hands was determined from their probability of occurrence and described in poker terminology – thus a "straight flush" (a run of adjacent numbers in a single suit) was the best hand, followed by "3-of-a-kind", a "flush", a "straight" and then a "pair of [4s, 3s, 2s or 1s]".

A computer program was written in Matlab to conduct the card game. To ensure comparability between participants, the computer used the same sequence of hands for all participants. Given hands of 3 drawn from a deck of 16, there were 560 possible hands that a player could receive and, with 4 suits, this enabled us to select one quarter of these hands as a representative sample – that is, including exactly the proportion of straight flushes, flushes, and so on that one would expect from 140 hands.

The rules of the game were explained to participants in advance as follows:

1. The rarer the card combination, the greater the likelihood of its winning. The computer will tell you what type of hand you have.
2. You must place a bet at every stage – of either 10c or 30c. Bet small when you expect to lose and high when you expect to win.
3. If you and the house (computer) have exactly the same type of hand, the house wins.

The computer, at each of 140 trials, showed the player their hand (from the representative sample described above), described it using the terminology given above, and asked them to answer two questions. The first was whether they believed that their chance of winning was greater or less than a given number – the anchor. (Anchors were a randomly generated sequence of values ranging from 5 to 95%; however, to ensure comparability across participants, the same sequence of anchors was used for each.)

After selecting the greater than or less than option, the participant was asked to make their estimate of how likely they thought they were to win the current hand given the cards that they had, and to place a bet of either 10c or 30c.

Once the participant's bet had been placed, the computer's hand was revealed and the winner was declared. Participants started the game with a \$10 'stake' and the computer updated this after each hand according to the bets they had made and whether they had won or lost. These winnings did not affect the amount that participants were paid but, rather, were included as another means of tracking the participants' performance.

Results

Measures

As described above, most of the measures were self-explanatory and were based on, either, response scales or from the raw total scores from cognitive tests. The anchoring task measures, however, are described below.

Anchoring. Each participant completed 140 trials of the anchoring task and a person's susceptibility to anchors was measured by calculating the rank order partial correlation between the anchors and their estimates (controlling for the true value) across all trials. This was done for the experiment as a whole but also for each of the four quarters of the task in order to examine learning across the task.

Higher values on this measure indicate greater susceptibility to the anchoring effect – that is, greater agreement between their estimates of the likelihood of their winning a hand and the anchor values they were shown prior to making those estimates.

Accuracy. The accuracy measure was calculated as the rank order partial correlation between the participant's estimates and the true values while controlling for the effect of the anchor. As for the anchoring measure, this was done across all trials and by quarter.

Higher values on this measure indicate greater accuracy – that is, agreement between a person's estimates and the true likelihood of their winning a hand. This measure serves as a check on people's performance and thus expertise.

Winnings. Rather than using the raw amount that a participant won over the course of the task as the measure, we compared their winnings with the optimal winnings. That is, we determined what an optimal decision maker, able to calculate the probability of winning each hand

accurately, would bet at each point and calculated their winnings. For the sample as a whole and for each quarter of the trials, we calculated a participant's winnings as the percentage of the optimal winnings.

Anchoring Task Statistics

Table 1 summarizes the descriptive statistics for our sample across all 140 trials of the anchoring task.

Table 1. Mean and 95% CI for each measure from the anchoring task, calculated across 140 trials. The second column shows the scale of the measure.

	Scale	Mean	CI ₉₅
Anchoring	0 to 1*	0.24	[0.20 0.27]
Accuracy	0 to 1*	0.86	[0.83 0.90]
Winnings	0 to 100%	80.1	[75.2 85.0]

* Note – while the anchoring and accuracy measures can, technically, take negative values, these were not predicted to occur and observation of the sample confirmed this with no accuracy scores and only 7 of the 102 anchoring scores being below zero – and none of these being lower than -0.09.

As indicated in Table 1, there is evidence of the anchoring effect in the data. Participant estimates correlated with the anchors at an average of 0.24 and examination of the confidence interval around this indicates that this effect was reliably non-zero. Notwithstanding, the results suggest that people were quite good at estimating the actual probability, with the correlations between their estimates and the true value averaging above 0.8 and their winnings being around 80% of what an optimal player would earn.

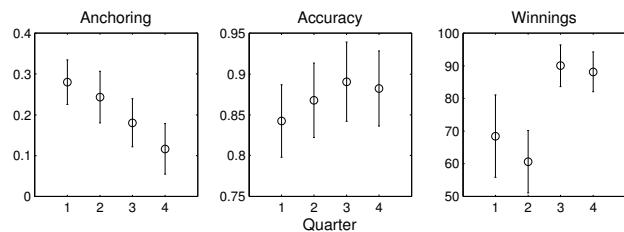


Figure 1. Mean scores on the three anchoring task measures with 98.7% CIs, divided by task quarter (i.e., 35 trials).

Figure 1 summarizes the three anchoring task measures across the four quarters of the task. The anchoring measure, in particular, shows a clear improvement quarter-by-quarter. Given that there is only a one-in-24 chance of seeing the pattern of results observed here – constant reduction across four stages – the data therefore provides strong evidence for a general trend of improvement. The evidence for improvements in accuracy and winnings, by comparison, is less clear. In particular, while the accuracy measure seems to show a trend toward improvement across the task, the corresponding confidence intervals largely overlap indicating that not too much should be drawn from this. The winnings data, by comparison, seem only to show a distinct

difference between the first and second *halves* of the task.

Analyses compared performance between the first and last quarters, looking at the mean differences on each of the three measures and finding evidence for improvement in all three cases as none of the confidence intervals around these mean differences contain zero: $MD_{anch} = -0.10$ $CI_{95} = [-.14 -.06]$; $MD_{acc} = .04$, $CI_{95} = [.01 .06]$; and $MD_{win} = 19.7$, $CI_{95} = 10.5$ 28.9].

We also calculated the size of the changes in the three measures using the non-parametric common language effect size measure A (Ruscio, 2008), which indicates the probability of a randomly chosen datum from one condition being better than a randomly chosen datum from the alternate condition. This indicated a large effect of learning on anchoring, a medium to large effect on accuracy and a small to medium effect on winnings, $A = 0.71$, 0.67 and 0.61 , respectively, when comparing 1st and 4th quarters.

That is, there is a clear trend of improvement on performance on the anchoring task as the task progresses through its 140 trials; specifically, as expertise (measured in terms of accuracy and winnings) increases, susceptibility to anchoring decreases.

Predicting Anchoring

To test for relationships between our independent measures and those from the anchoring task described above, we conducted a series of rank order correlations. These are given in Table 2.

Table 2. Rank order correlations between cognitive tasks, experience measures and performance on the anchoring task.

	1	2	3	4	5	6
1. Anchoring	-	<.01	.10	.13	.20	.98
2. Accuracy	-.41	-	<.01	.55	.04	.09
3. Winnings	-.17	.49	-	.03	.03	.04
4. NAT	-.15	.06	.22	-	<.01	.01
5. CRT	-.13	.20	.22	.44	-	.24
6. Poker Exp.	.00	.17	.20	.25	.12	-

Note: the bottom left triangle shows the correlation strengths while the top right shows the corresponding p-values (two-tailed). $N = 102$ in all cases. **Bold** correlations are significant at the .05 level. *Italic* correlations are significant at .05 if considered as directional hypotheses.

When one considers Table 2, it can be seen that the majority of effects (14 of 15) are in the directions one would expect. Anchoring is negatively related to people's accuracy and winnings and also to both of the cognitive measures. Interestingly, however, poker experience has no noticeable effect on susceptibility to anchors in this task. Accuracy, similarly, shows the expected relationships, being the best predictor of people's winnings and positively related to the CRT measure, people's self-rated poker experience and even, weakly, to the NAT. Similarly, the correlations between winnings and all the other measures were in the expected direction. That is, in addition to the correlations

noted above, winnings correlates positively with CRT, NAT and poker experience.

Focusing on the individual difference variables, it is evident that they show logically coherent relationships. CRT and NAT scores were most strongly related to one another and all of their correlations with the anchoring measures, while weak, were in the expected direction.

Anchoring and Expertise

The above results show that our independent measures do, weakly, predict performance on the anchoring task. However, given that learning is occurring during the task, we need to establish how this increasing 'expertise' and our measures relate. Importantly, we need to ascertain whether differential ability on these measures predicts *changes* in susceptibility to anchoring as the task progresses.

As an initial test of the relationship between expertise (on the anchoring task) and our other measures, we calculated rank order correlations between the anchoring measure and each of these for just the first quarter of trials – when expertise was lowest. This confirmed that, across the first 35 trials of the task, the correlations between anchoring and all three of NAT, CRT and poker experience were zero, $\rho = -.02$, $.00$ and $.04$, respectively.

To further examine this question, the sample was divided into people who improved their performance on the anchoring measure ($n = 69$) and those who did not ($n = 33$) and these groups were compared (Figure 2).

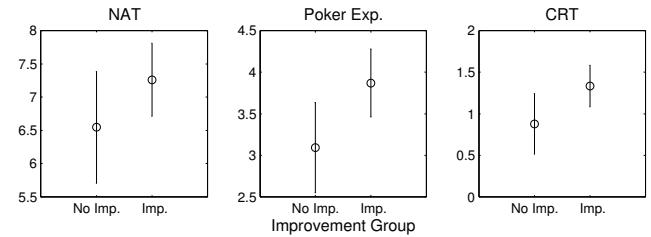


Figure 2. Mean NAT score, Poker Experience and CRT score, with 95% CIs, by improvement group (i.e. whether participants reduced their susceptibility to anchoring).

Inspection of the confidence intervals in Figure 2 suggests that all of these measures show some difference between the participants who did and did not improve their performance (that is, reduced their susceptibility to anchoring). Specifically, the people who became less susceptible to anchoring were those who scored higher on the NAT and CRT and rated their own poker experience more highly.

These effects are small to medium for the cognitive measures $A = 0.59$ and 0.63 for the NAT and CRT, respectively, and medium for poker experience, $A = 0.64$.

We also examined the number of participants from each education group who improved or did not improve their anchoring score across the course of the task. Table 3 shows the observed and expected sizes of each group for the six combinations of education level and improvement.

Table 3 shows marked deviations from statistical

expectations. Specifically, group 1 (non-university educated) have fewer improvers than would be expected while group 3 (university graduates) almost all showed improvement on the anchoring measure. A χ^2 -test confirmed that the probability of seeing such divergence in the absence of an effect was very low, $p = .014$. The effect, comparing educational levels of the group showing improvement and those not, was of medium strength, $A = 0.66$.

Table 3. Observed and expected sub-group sizes, comparing education level and improvement on anchoring measure.

		Education Group		
		1	2	3
Improve	No	12 / 7.1	18 / 18.8	3 / 7.1
	Yes	10 / 14.9	40 / 39.2	18 / 14.9

Overall, these results suggest that the relationships observed above between anchoring and the individual differences are largely mediated by expertise. That is, people who score better on the cognitive measures, while no less susceptible at the beginning of the task, *become* increasingly less susceptible to anchoring as the task progresses, with the result that, overall, a relationship is seen between cognitive ability and anchoring. Similarly, while prior experience with poker offered no immediate advantage, it did predict improvement on this, related, task.

Education and Anchoring

A potential confound in tasks such as these is participants' level of education. As noted above, this is known to be strongly related to Frederick's (2005) CRT measure and is likely to relate to numerical ability as well. A university education also seems the most likely way that a person might have become aware of the anchoring effect and, thus, potentially, be less susceptible. Therefore, participants' scores on the anchoring measures, cognitive measures and poker expertise were examined in relation to their educational attainment (Figure 3).

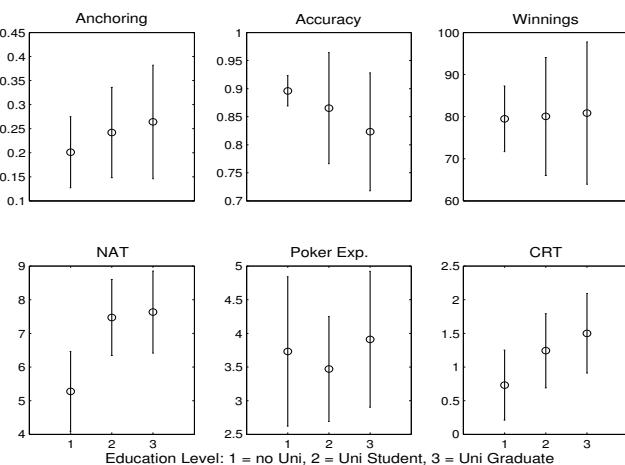


Figure 3. Means for anchoring task and independent measures with 98.3% confidence intervals, by education group.

Despite our predictions of a possible association between anchoring performance and education, the results in the top row of Figure 3 show little difference in scores between the different groups excepting that university students and graduates are more variable in their performance. Similarly, there is no apparent difference between any of the groups in relation to their level of poker experience.

As was expected, however, there were clear differences between the education groups on the cognitive measures. Both of the university groups scored higher than the non-university group on the NAT, whereas scores on the CRT seemingly increase across each educational level. The sizes of these effects were calculated, indicating that the difference between CRT scores in the different groups varied from weak to medium/strong, $A = 0.64$, 0.57 and 0.70 when comparing groups 1 and 2, 2 and 3, and 1 and 3, respectively. By comparison, the difference between education groups on the NAT varied from very small to very strong, $A = 0.77$, 0.52 and 0.78 , comparing groups 1 and 2, 2 and 3, and 1 and 3, respectively.

Discussion

Taken as a whole, the results were generally consistent with our predictions. Participants showed clear anchoring effects on our task but the degree of bias decreased as people played more and more hands of the card game and, presumably, learnt the underlying probabilities of winning with the different types of hand. The average correlations between participants' estimates and the anchors they had seen reduced from 0.28 in the first quarter of the trials to 0.12 in the last.

Of equal interest is the fact that both of the cognitive measures we considered, the NAT and CRT, predicted susceptibility to the anchoring effect. The overall effect of this was weak, however, with correlations of -0.15 and -0.12, and we note that these measures do not predict anchoring susceptibility at the *beginning* of the task. Rather, people who scored higher on numerical ability and cognitive reflection were more likely to show reduced susceptibility to anchoring as the task progressed. Similarly, self-rated poker expertise had no, overall, relationship with anchoring but did predict improvement on the anchoring measure as the task progressed (see Figure 3).

Finally, participant's level of education was also related to anchoring susceptibility. In particular, it predicted whether their susceptibility to anchoring would decrease as they continued in the task – an unsurprising finding given the strong relationships observed between NAT, CRT and education level.

In general terms, then, it seems that the experiment provides evidence for the value of considering individual differences when looking at biases such as anchoring. It also confirms previous findings (Frederick, 2005; Stanovich & West, 1998, 2008) linking bias susceptibility to cognitive ability.

An important caveat, however, is that the larger differences observed here are those linked to experience and

expertise. That is, differences in cognitive ability are, in this study at least, more useful for predicting the rate of increase in expertise rather than direct susceptibility to biases. Consequently, it is the increase in expertise that accounts for the larger differences in susceptibility to biases. For this reason, we are left in the position of having to define expertise within a field of interest before being able to state definitively that 'experts' will be less affected by anchors.

A second concern is whether what we have called 'expertise' herein is representative of expertise in real domains. Clearly we have a practice effect but expertise generally refers to more durable, transferable abilities. Given the difficulties in defining expertise, however, we feel justified in calling this a necessary first step and certainly regard it as an improvement over previous studies where expertise is self-rated or simply assumed from vocation.

Future Research

This experiment was, primarily, concerned with establishing that anchoring susceptibility varies with expertise. As such, and given our choice of a poker-based game, we have concentrated on the small number of cognitive tasks we thought most relevant to such a task – given previous findings. The nature of anchoring, however, where a value acts to bias recall of knowledge about another value being estimated, suggests that a consideration of other cognitive abilities would be valuable. In particular, aspects of memory would seem to be involved and a consideration of individual differences in memory would, therefore, be informative.

More work is also required to establish the relationships between the various types of cognitive bias commonly discussed. While there have been several attempts at creating a taxonomy of biases, these have largely relied on assumptions regarding the underlying heuristic processes (see, e.g., Tversky & Kahneman, 1974). A large-scale factor analysis of bias questions and a range of cognitive measures would enable the relationships between various biases to be placed on a firmer footing.

Conclusions

In conclusion, we have demonstrated that susceptibility to anchoring is related to the cognitive abilities we have examined and other measures such as educational level and self-rated expertise on related tasks. More importantly, however, it seems that these abilities mediate peoples' development of expertise and it is this expertise within a specific estimation context that actually reduces susceptibility to anchoring.

Therefore, if one is confident that the person making an estimate is an expert, then one can also expect that their estimates will be less affected by anchoring values than less-expert individuals. The use of simple cognitive measures such as the NAT and CRT, however, while somewhat predictive of differences in expertise, will not be sufficient to enable this conclusion to be reached. Instead, a measure of expertise on the task in question will be required.

Acknowledgments

MBW & SHB acknowledge ExxonMobil and Santos' funding of the IBP Group at the ASP. The authors thank Karina Burns for her assistance with data collection.

References

Bennett, G. K., Seashore, H. G., & Wesman, A. G. (1989). *Differential Aptitude Test*. Marrickville, NSW: Psychological Corporation.

Bergman, O., Ellingsen, T., Johannesson, M., & Svensson, C. (2010). Anchoring and cognitive ability. *Economic Letters*, 107, 66-68.

Bratvold, R. B., Begg, S. H., & Campbell, J. M. (2002). SPE 77509 - Would you know a good decision if you saw one? *Proceedings of the Society of Petroleum Engineers 78th Annual Technical Conference and Exhibition*.

Chapman, G. B., & Johnson, E. J. (2002). Incorporating the irrelevant: anchors in judgments of belief and value. In T. Gilovich, D. Griffin & D. Kahneman (Eds.), *Heuristics and Biases: the Psychology of Intuitive Judgment* (pp. 857). Cambridge: Cambridge University Press.

Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives*, 19(4), 25-42.

Malhotra, V., Lee, M. D., & Khurana, A. K. (2005). Domain experts influence decision quality: towards a robust method for their identification. *Journal of Petroleum Science and Engineering, Special Issue*(Oil and Gas Exploration and Production Research in Australia 2005).

Northcraft, G. B., & Neale, M. A. (1987). Experts, amateurs and real estate: an anchoring-and-adjustment perspective on property pricing decisions. *Organizational Behavior and Human Decision Processes*, 39, 84-97.

Oechsler, J., Roider, A., & Schmitz, P. W. (2009). Cognitive abilities and behavioral biases. *Journal of Economic Behavior and Organization*, 72, 147-152.

Quattrone, G. A., Lawrence, C. P., Warren, D. L., Souza-Silva, K., Finkel, S. E., & Andrus, D. E. (1984). *Explorations in anchoring: the effects of prior range, anchor extremity and suggestive hints*. Unpublished manuscript, Stanford University: Stanford.

Ruscio, J. (2008). A probability-based measure of effect size: robustness to base rates and other factors. *Psychological Methods*, 13(1), 19-30.

Stanovich, K. E., & West, R. F. (1998). Individual differences in rational thought. *Journal of Experimental Psychology: General*, 127(2), 161-188.

Stanovich, K. E., & West, R. F. (2000). Individual differences in reasoning: implications for the rationality debate. *Behavioral and Brain Sciences*, 23(645-726).

Stanovich, K. E., & West, R. F. (2008). On the relative independence of thinking biases and cognitive ability. *Journal of Personality and Social Psychology*, 94(4), 672-695.

Thorsteinson, T. J., Breier, J., Atwell, A., Hamilton, C., & Privette, M. (2008). Anchoring effects on performance judgments. *Organizational Behavior and Human Decision Processes*, 107, 29-40.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185, 1124-1131.

Welsh, M. B., Begg, S. H., & Bratvold, R. B. (2006). SPE 102188: *Correcting common errors in probabilistic evaluations: efficacy of debiasing*. Paper presented at the Society of Petroleum Engineers 82nd Annual Technical Conference and Exhibition., Dallas, Texas, USA.