Adaptive Uses of Random Criterion: The Largest Number Problem, the Two-Envelope Problem, and the Anchoring and Adjustment Heuristic

Bruce D. Burns (bruce.burns@sydney.edu.au)
School of Psychology, Brennan MacCallum Bldg, A18
University of Sydney, NSW 2006, Australia

Abstract

Many cognitive processes appear to incorporate threshold criteria, but when criteria are know to be random their use may appear irrational. For example, when people's estimates are influenced by random anchors (Tversky & Kahneman, 1974). However Cover (1987) showed that choosing whether a seen or unseen number was greater is improved by using a random number as a criterion. Such *Cover functions* are also the basis for solving the two-envelope problem. This solution suggests that people's responses should be influenced by where a value falls in its distribution, a hypothesis supported empirically. The anchoring and adjustment heuristic can also be seen as application of a Cover function. Simulation can demonstrate that adjustment towards a random anchor from an initial random estimate will on average improve the final estimate. Anchoring and adjustment is an example of how Cover functions can contribute to understanding cognitive phenomena, and may have wide applicability.

Keywords: Bayesian reasoning; anchoring and adjustment; two-envelope problem; decision making.

Introduction

It is axiomatic that a random number cannot provide information about another number. Unless generating unpredictability is crucial, it therefore seems obvious that basing a decision on a number known to be random is at best a mistake, and at worst evidence of irrationality in human thinking. However this is not necessarily true, and this fact may have interesting implications for cognitive phenomena. This can be demonstrated by considering three problems that on the surface are quite different.

Starting with the *pick the largest number* problem (Cover, 1987) the concept of a *Cover function* has been developed: a probabilistic decision function based on a random number. Abbott, Davis, and Parrondo (2010) used this concept to produce a novel solution to the *two-envelope problem*. *Anchoring and adjustment heuristics* are of wider psychological interest, but the same tools can be used to demonstrate that adjusting towards a random anchor can improve an estimate.

Demonstrating that Cover functions are a tool that can lead to new insights into phenomena opens up the potential for wide application. Many models of cognitive processes involve comparisons, and under appropriate conditions some of these may be seen as involving Cover functions.

Chater and Oaksford (2007) describe the Bayesian approach to reasoning as proposing that it is inherently probabilistic, although this does not necessarily imply that the mind is a probabilistic calculating machine (p. 92). Cover functions may in this way contribute to cognition.

Pick the largest number

Cover (1987) presented the *pick the largest number* problem in which two numbers are written on slips of paper, Slip A and Slip B. No information is provided about the distribution of numbers on these slips. The solver then randomly chooses Slip A and reads the number written on it, and then must decide whether that number is higher or lower than the number written on the unseen Slip B. It appears that this task cannot be done with greater than 50% success, however Cover asserts that there is a strategy that raises the expected rate of success above 50%.

Cover (1987) proposes that the solver randomly selects a *splitting number N* according to the density function $f(n)$, $f(n) > 0$, for $n \in (-\infty,\infty)$. If the number on Slip A is less than N then decide it is the lower number, and if greater than N then decide it (A) is the higher number. That this will yield expected success greater than 50% is illustrated in Table 1. Essentially the problem concerns three random numbers A, B, and N. The critical issue is the ordering of these three numbers from smallest to largest, and there are six orders of three numbers.

Table 1: Description of the six different orders of numbers *A*, *B*, and *N* with the decision regarding Slip A for each case and whether the outcome of that decision was correct or incorrect.

Cases	Order from smallest to largest	Decision for Slip A	Outcome
1	N A B	Higher	Incorrect
2	N B A	Higher	Correct
3	A B N	Lower	Correct
4	B A N	Lower	Incorrect
5	A N B	Lower	Correct
6	B N A	Higher	Correct

In four of the six cases the solver makes the correct decision. If A, B, and N were all selected from the same uniform distribution then each case would be equally likely, and thus the solver would be expected to be correct 66.67% of the time. However the distributions of A and B do not have to be known, all that is required to produce an expected success rate of above 50% is that Cases 5 and 6 are possible. When N is less than both A and B (Cases 1 and 2) there is a 50% chance of success; when N is greater than both *A* and *B* (Cases 3 and 4) there is also a 50% chance of success; but when *N* splits the two numbers (Cases 5 and 6) then the correct choice is made 100% of time. Therefore if *N* has a distribution such that it is possible for it to split A and B, the overall expected rate of success must be greater than 50%. Of course if nothing is known about the distribution of *A* and *B*, the widest possible distribution is required for *N*, thus Cover proposed that it could be any real number. Alternatively if something is known about the distributions of *A* and *B* then a function for maximizing success could be derived. Functions proposing decision rules based on random criterion have been referred to as *Cover functions* (McDonnell & Abbott, 2009).

Note that whereas the criterion value is random, the decision rule that utilizes this value is not. The rule has to specify the appropriate direction for the decision. Essentially this solution works because it provides the decider with a way to utilize the information content of the seen number. There is a monotonic function between the magnitude of a number and its probability of being the higher number, though the shape of this function depends on the distribution. Using a random number as input to the decision rule provides an appropriate monotonic function regardless of the distribution of the numbers, as long as the distribution of *N* guarantees that it is possible for values of *N* to be between the two numbers (i.e., Cases 5 and 6 from Table 1 are possible). *N* could also be an arbitrary rather than a truly random number, if enough is known about the possible values of *A* and *B* an arbitrary number could be chosen that might split them. As long as Cases 5 and 6 have a nonzero probability, then the expected outcome will be greater than 50%.

Informally, it is clear why the Cover function succeeds for the largest number problem: it provides a way to exploit the fact that the higher a number is the more likely it is to be greater than another number. However it is hard to say how high a given number has to be in order to decide it is higher than an unknown number. In this way the Cover function is a heuristic in the sense that Kahneman and Frederick (2002) define it: it is a way of substituting an answerable question for an unanswerable question. As long as the answerable question is consistent with a regularity of the world, then that heuristic should be adaptive.

The two-envelope problem

The two-envelope problem has a long history as a mathematical puzzle. Versions were proposed by Kaitchik (1953, pp. 133-134), but these were not the earliest. Although he does not claim authorship of it, Zabell (1988) stated a two-envelope version with the following characteristics: 1) the contents of the two envelopes are x and $2x$; 2) no distribution or limit is given for x; 3) the reasoner is randomly handfed one of the envelopes and opens it; 4) then the reasoner is given a choice: keep the amount observed, or trade it for the contents of the other envelope. Before the envelope is opened the expected outcome is:

(1) $E = ½(x + 2x) = 3x/2$

Opening an envelope cannot change the amounts in the envelopes so it should not matter whether you keep or trade envelopes because to trade is equivalent to changing your initial random choice. However, opening an envelope containing y means that trading yields either $2y$ or $½y$. If each possibility has a 50% possibility then trading results in an expected outcome equal to $5y/4$. Thus, if the two envelopes were held by two different people (as proposed by Zabell, 1988), then it might appear that after opening their own envelopes both people would expect to gain from trading. This cannot be true so the problem has sometimes been called a paradox. As Zabell and others have pointed out, the resolution of this paradox is that the envelopes contain two possible pairs of amounts [$2y$, y] or [y, $1/2y$] but they are not equally likely. The $p(y|\text{pair})$ is not equal to $p(\text{pair}|y)$; the first probability is known but it is the second that the reasoner needs. Analyzing what that probability is, and thus what the reasoner should do, was considered to have defied a satisfactory mathematical solution (Albers, Kooi, & Schaafsma, 2005). So the paradox was resolved but the problem of whether to trade remained.

Finding a solution to this problem may well be more than just puzzle solving. McDonnell and Abbott (2009) point out that the envelope problem has attracted wide interest in game theory and probability theory, and that it is paradigmatic of recent problems in physics, engineering and economics which involve probabilistic switching between two states. For example, it has been shown in stochastic control theory that random switching between two unstable states can result in a stable state (Allison & Abbott, 2001).

There is only one published paper on how people respond to the envelope problem. Butler and Nickerson (2008) presented participants with six different versions of the problem. They found that participants were largely insensitive to the logical structure of the problem. They concluded that instead participants applied simple heuristics or forms of folk wisdom.

A solution

Recently McDonnell and Abbott (2009) and Abbott, et al (2010) propose a strategy that can increase the expected outcome above that in Equation 1. The key to their approach is to recognize that once an envelope is opened the information regarding what it contains breaks the symmetry that leads to Equation 1. Their starting points were Cover's (1987) solution to the *pick the largest number* problem, and the analysis of Parrondo's games in which two losing strategies can be combined to produce a winning strategy if the current state of the problem is used as a criterion (Harmer & Abbott, 1999). Solving these types of problems requires probabilistic switching between states.

Abbott et al (2010) supposed that opening the envelope reveals y dollars, and the player then trades envelopes with a probability $P(y) \in [0,1]$. Figure 1 illustrates their analysis with a Markov model. From the model it can be seen that their analysis shows that the expected return (E) when x represents the smaller of the two amounts and $2x$ the larger, will be:

$$\text{(2)} \quad E = \tfrac{1}{2}[2x\,P(x) + x\,[1\text{-}P(x)] + xP(2x) + 2x[1\text{-}P(2x)]]$$
$$= \tfrac{1}{2}[3x + xP(x) - xP(2x)]$$
$$= 3x/2 + x/2[P(x) - P(2x)]$$

Figure 1: A Markov model based on Abbott el al's (2010) analysis. P(*x*) representing the probability of trading if the value in the opened envelope is *x*, and P(2*x*) representing the probability of trading if the observed value is 2*x*.

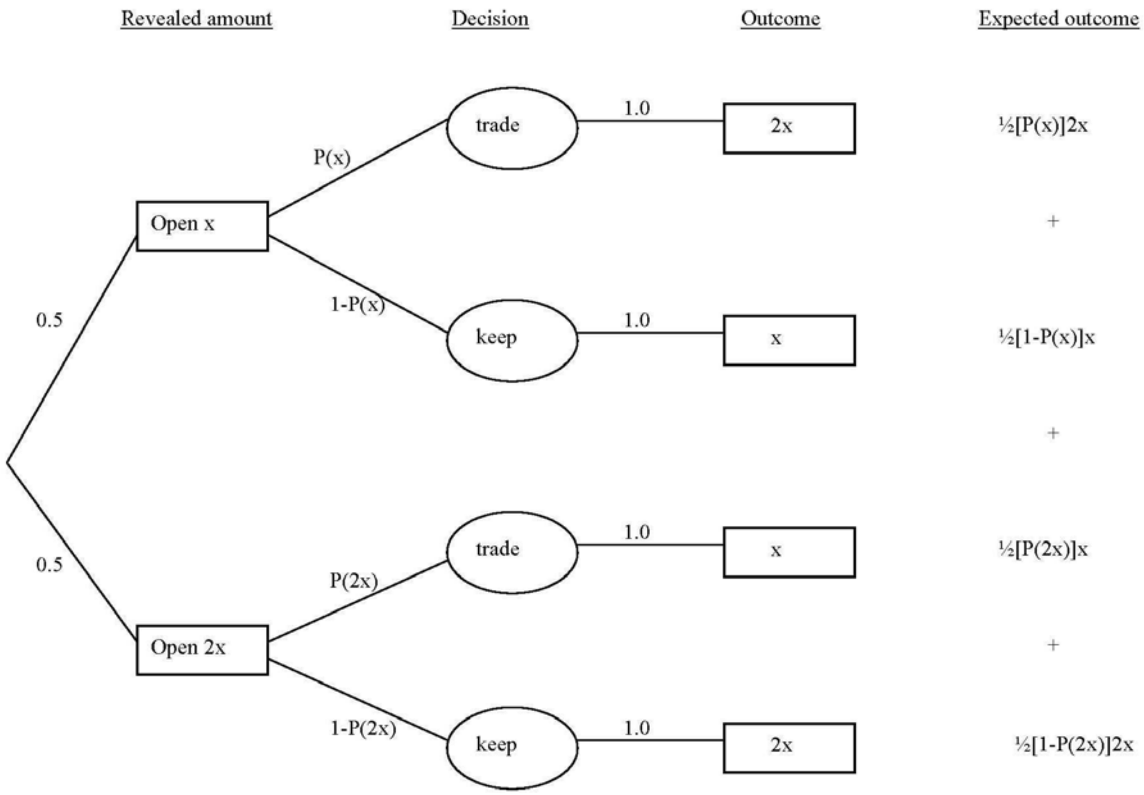

Equation 2 shows that probabilistic trading as a function of *x* can raise the expected value above that expected from either trading or keeping regardless of the observed amount (i.e., Equation 1). Returns can only be improved if $P(x) > P(2x)$, that is, when the trading function is such that trading is less likely the higher *x* is. Abbott, et al. (2010) show that a monotonically decreasing function will increase the expected outcome, and that this does not presuppose any particular probability density function for *x*. Calculating the optimal trading function requires knowing the probability density function for *x*, but their analysis demonstrates that a tendency to trade that is a negative monotonic of observed amount can increase expected outcomes.

Thus the two-envelope problem can be seen as a variation of the *largest number problem*, except now instead of the two numbers having no specified relationship to each other one number is twice the other. A third random criterion number *N* could be proposed resulting in three numbers with the same six possible orders as in Table 1. So a Cover function in which the envelope is only swapped if the observed contents are less than the random number will be a monotonic function that yields better expected outcomes than Equation 1.

An empirical prediction

Abbott, et al's (2010) model shows that the higher an observed amount sits within the distribution of amounts, the less likely trading should be. Thus adaptive behavior for people faced with the two-envelope problem would be to be less likely to trade the higher the observed is within the distribution of possibilities. This prediction was tested by specifying a range so that it could be said with some confidence that participants saw an amount as high in the distribution.

An Experiment

Where the observed contents of an envelope sit in a distribution of possible amounts depends both on what the amount is and what are the upper and lower limits of possible amounts. So in Experiment 1 both the observed amount ($10 or $100) and the limit ($200 or no limit) were manipulated. It was predicted that the extent to which people traded would depend on the interaction of the observed and limit factors, such that they would be least likely to trade when the observed was $100 and the upper limit was $200.

Method

Participants. A total of 160 senior psychology students at the University of Sydney participated during a practical class focussed on reasoning.

Procedure & Materials. Participants read and responded to the following scenario on paper (the italicized text in the squared brackets replaced the underlined text in the relevant condition):

> Imagine that you given a choice between two envelopes each containing a sum of money. <u>You are told that neither envelope could hold more than $200</u> *[You are told that the envelopes could contain any amount of money]*, but one envelope contains exactly twice as much money as the other. You randomly choose one of the envelopes and open it, revealing that it contains <u>$100</u> [*$10*]. You are told that you can either keep the <u>$100</u> [*$10*] or take whatever is in the other envelope. What would you do?

Participants circled whether they would keep the $100 [$10] or trade it for whatever was in the other envelope.

Results & Discussion

Table 2 presents the proportion of participants choosing to trade in each condition. Given that choices were dichotomous, a logistic regression analysis was performed on choice (0=keep, 1=trade) entering the factors of limit, observed amount, and their interaction. The parameter for limit was not statistically significant, Wald $\chi^2(1) = 0.525$, $p = .469$, but that for observed was, Wald $\chi^2(1) = 16.224$, $p < .001$, as was the interaction, Wald $\chi^2(1) = 3.885$, $p = .049$.

Table 2: Proportion of participants in each condition of Experiment 1 choosing to trade (with sample sizes).

	$10 in opened envelope	$100 in opened envelope
$200 limit	.80 (n=39)	.32 (n=38)
unlimited	.73 (n=40)	.56 (n=43)

As predicted, these results showed that participants' choices were affected by what they observed in the envelope, in that overall there was a strong effect of amount observed. However there was also a significant interaction in that trading was least likely if the highest and observed amounts were such that the largest amount possible was at the limit. This suggests that people's responses were affected by where they saw the possible amounts as falling in the distribution of amounts.

The results of this single experiment are not definitive, but they are consistent with the prediction derived from Abbott, et al's (2010) model. Thus the application of Cover functions to problems can lead to testable predictions about human behavior.

Anchoring and Adjustment

A phenomenon of wider interest to researchers into cognition than the two-envelope problem is anchoring and adjustment. This heuristic was one of the three that Tversky and Kahneman (1974) presented in their canonical paper on biases in human judgement. They proposed that people often make estimates by starting from an initial value and adjusting it to yield a final answer. Such initial values may be suggested by the formulation of the problem or preliminary calculations, but different starting points would yield different estimates biased by the initial values. They called this called *anchoring*.

Anchoring might occur because the anchor is considered a reasonable estimate, so the strongest demonstrations of anchoring as a heuristic are those in which the anchor is explicitly random. For example, Tversky and Kahneman (1974) describe a study in which participants were asked to estimate various quantities such as the percentage of African countries in the United Nations. For each quantity a number between 0 and 100 was determined by spinning a wheel in the presence of the participants. Participants first indicated whether the random number was higher or lower than the correct answer and then they estimated the quantity. When the wheel indicated 10 the median estimates for percentage of African countries was 25%, but when the wheel indicated 65 the median was 45%. Thus participants' estimates were heavily influenced by a number they knew to be random.

There is a substantial body of research on anchoring and it is of wide interest because as Wilson, Houston, Brekke, and Etling (1996) pointed out "it is rare to find a single, relatively simple process that explains such diverse phenomena." (p. 387) By showing that people are influenced by information that they know to be random and arbitrary, such anchoring has been seen as an example of human irrationality (e.g., Ariely, 2008). However there appears to have been no analysis of whether ignoring randomly generated anchors is in fact optimal, it has simply been assumed.

Adjustment as a Cover function

The most impressive demonstrations of anchoring are when a randomly selected anchor influences people's final estimates. These demonstrations usually involve questions for which the participant has little idea regarding the right answers, thus effectively any initial guess is a random number (specific factors may determine guesses for a particular question, but across questions guesses are effectively random). In such cases the true answer does not really matter and thus across questions the true answer can also be thought of as a random number. Therefore demonstrations of the effect of random anchors can be thought of as analogous to the pick the greatest number problem, with two random numbers that may be split by a third random number. However instead of using the splitting number to decide which of the other two numbers are higher, the question is whether shifting the initial guess towards the splitting number yields an estimate closer to the true number than was the initial guess. Thus adjustment can be seen as a type of Cover function.

I will refer to the three numbers as *T* (true answer), *G* (initial guess), and *A* (anchor). As in Table 1, these three numbers can have six different orderings. Table 3 shows which orders would be expected to yield better estimates. As in Table 1, four of the six cases usually yield a better estimate when the rule "shift the guess towards the anchor" is followed. However whereas Cases 2 and 3 are likely to yield better estimates they are not guaranteed to because an guess adjusted towards the anchor may leap-frog the true value and end up further away from it than the initial guess.

Table 3: Description of the six different orders for numbers A (anchor), G (guess), and T (true value) with for each case the direction of adjustment from the G, and whether that adjustment would lead to a better of worse outcome relative to the true value.

Cases	Order from smallest to largest	Direction of movement of initial guess	Outcome relative to true value
1	T G A	Higher	Worse
2	G T A	Higher	Better (probably)
3	A T G	Lower	Better (probably)
4	A G T	Lower	Worse
5	T A G	Lower	Better
6	G A T	Higher	Better

Simulations

Simulation can be used to explore when adjustment towards an anchor would be expected to improve estimates. These simulations are not intended as being for a single question, for which *T* would be fixed, but a population of questions. In each simulation three random numbers *T*, *G* and *A* were selected. To simplify these simulations all three numbers were selected from the same uniform distribution: integers between 1 and 1 million. A new estimate is then calculated using the following formula:

(3) $\text{estimate} = \text{guess} + S(\text{anchor} - \text{guess})$

How much to adjust the guess is not clear, so this was determined by a parameter *S* (with range [0.0, 1.0]) that determines the proportion of the distance between the anchor and the guess that the guess is adjusted by. If $S = 0$ there is no adjustment, if $S = 1.0$ then the estimate becomes the anchor (note that the simulations are neutral with regard to whether the adjustment is from the guess to the anchor, or the anchor to the estimate). Table 4 shows the results of 1 million simulations for each of a range of values of *S*. There are two outcome measures: percentage of the trials in which the adjustment moves guesses in the right direction; and the mean size of all the adjustments (positive adjustments are towards the true answer, negative are from the true answer).

Table 4: Results of simulations with different values of the *S* parameter with percent of trials in which adjustment was in the right direction and the mean size of all adjustments.

Value of *S*	Adjust in right direction	Mean adjustment
.01	66.49%	1650.1
.10	65.02%	14988.4
.20	63.37%	26740.1
.30	61.67%	34916.6
.40	60.12%	40217.2
.50	58.28%	41413.6
.60	56.68%	40012.5
.70	55.04%	35079.0
.80	53.41%	26889.8
.90	51.64%	14664.4
1.0	49.97%	381.5

Table 4 shows that for all levels of *S* up to 1.0, adjustment is expected to improve the final estimate. The highest percentages of adjustments in the correct direction are for the smallest values of *S*. Small adjustments minimize the chance that the adjustment will overshoot the true value when Cases 2 or 3 in Table 3 represent the order. The smaller *S* was, the closer the percentage should approach to two-thirds (of course if $S = 0$ then no adjustment is made and it cannot be in the right direction). However in terms of the size of the improvement, the largest mean adjustment was when $S = .50$. An adjustment of 41413.6 represents 4.1% of the maximum value (1 million) but given that the average guess will be 500,000 this indicates an 8.2% improvement in the average estimate. Thus a substantial expected improvement can result from adjusting towards a random anchor. Informally, it is clear that adjustment works for the same reason the solutions to the greater number and two-envelope problems work: they provide people with a way to exploit the fact that high numbers tend to be high. Thus improvement in each can result from applying a Cover function as an adaptive heuristic

This result is not dependant on *G*, *T* and *A* having the same distributions, however once this assumption is relaxed it is possible to give them distributions in which adjustment will not be expected to yield a better estimate (e.g., when *A* has a higher mean than *G*, and *T* a lower mean that *G*, and these discrepancies are large). Also adjustment will be less effective when G and T are positively correlated. The aim of these simulations is not to show that adjustment is an effective heuristic under all conditions, but instead to show that it is false to assume that adjusting towards a random anchor is inherently irrational.

General Discussion

This analysis has implications both for the particular phenomena considered and more generally for reasoning and decision making.

This analysis is not intended to address most of the issues examined by the large literature on anchoring and adjustment. However an assumption of this literature is that any influence of a random anchor on people's decisions must be an error, and it is this assumption that the analysis demonstrates is incorrect. Seeing such anchoring and adjustment as beneficial may suggest new directions for research. For example, the analysis suggests there are parameters for which adjustment towards an anchor would not be expected to improve estimates, examination of such parameters could be the basis of predictions regarding when people will be most influenced by anchors. The results of the experiment on the two-envelope problem illustrate how Cover function solutions can lead to empirically testable predictions. In my own work on how belief in the hot hand is adaptive (Burns, 2004) I showed how such an approach can change the sorts of questions asked about it.

However this paper is intended to be only a demonstration of how the concept of Cover functions could apply to psychologically interesting phenomena. That decisions can be improved by utilizing a random criterion, is a result that could have implications for understanding a number of phenomena. In both perception and decision making, criteria with little validity might be used with adaptive success, especially initially. Cover functions suggest that criteria do not necessarily need to be based on any knowledge in order to be useful, even if only for bootstrapping a system.

To the extent that phenomena such as anchoring may show that people utilize Cover functions, they support a Bayesian approach to reasoning, even if they are not Bayesian models themselves. As Oaksford and Chater (2007) show, when reasoning is seen as probabilistic and based on taking advantage of the distribution of information in the environment then behavior previously regarded as in error can be shown to be rational. The analyses presented here are consistent with the claim that people may be making decisions consistent with sensitivity to the distributions of the numbers representing anchors and the numbers they are asked to estimate. Stanovich (1999) pointed out that the rationality of human reasoning has often been judged by how close that reasoning has been to what was considered normatively correct. However often little analysis has informed declarations of what is normative, and Bayesian approaches have the potential to correct this. Cover functions support such analyses and provide a potentially useful tool for modeling a range of cognitive functons.

References

Abbott, D., Davis, B. R., & Parrondo, J. M. R. (2010). The two-envelope problem revisited. *Fluctuation and noise letters, 9*, 1-8.

Albers, C. J., Kooi, B. P., & Schaafsma, W. (2005). Trying to solve the two-envelope problem. *Synthese, 145*, 89–109.

Allison, A., & Abbott, D. (2001). Control systems with stochastic feedback. *Chaos, 11*, 715–724.

Ariely, D. (2008). *Predictably irrational: The hidden forces that shape our decisions*. London: HarperCollins Publishers.

Burns, B. D. (2004). The hot hand in basketball as fallacy and adaptive thinking. *Cognitive Psychology, 48*, 295-331.

Butler, S. F., & Nickerson, R. S. (2008). 'Keep or trade? An experimental study of the exchange paradox. *Thinking & Reasoning, 14*, 365 – 394.

Cover, T. (1987). Pick the largest number. In T. Cover & B. Gopinath (Eds.), *Open Problems in Communication and Computation* (p. 152). New York: Springer-Verlag.

Kahneman, D. & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgement. In T. Gilovich, D. Griffin, & D. Kahenman (Eds.), *Heuristics and biases: The psychology of intuitive judgement* (pp. 1-19). Cambridge, UK: Cambridge University Press.

Harmer, G. P., & Abbott, D. (1999). Losing strategies can win by Parrondo's paradox. *Nature, 402*, 864.

Kraitchik, M. (1953). *Mathematical Recreations* (2nd edition). New York: Dover.

McDonnell, M. D., & Abbott, D. (2009). Randomized switching in the two-envelope problem. *Proceedings of the Royal Society A*, *465*, 3309-3322.

Oaksford, M., & Chater, N. (2007). *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford, UK: Oxford University Press.

Stanovich, K. E. (1999). *Who is rational?: Studies in individual differences in reasoning.* Mahwah, NJ: Lawrence Erlbaum Associates.

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

Wilson, T. D., Houston, C. E., Brekke, N., & Etling, K. M. (1996). A new look at anchoring effects: Basic anchoring and its antecedents. *Journal of Experimental Psychology: General, 125*, 387–402.

Zabell, S. L. (1988). Symmetry and its discontents. In B. Skyrms & W. L. Harper (Eds.), *Causation, chance, and credence* (Vol. I, pp. 155 – 190). Dordrecht, the Netherlands: Kluwer Academic Publishers.