Probabilistic reasoning in the two-envelope problem

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

Abstract

In the two-envelope problem a reasoner is offered two envelopes, one containing exactly twice the money in the other. After observing the amount in one envelope it can be traded for the unseen contents of the other. Until recently it was argued that it did not matter whether the envelope was traded or not, but Abbott, Davis & Parrondo (2010) showed that gains could be made if trading was a probabilistic function of amount observed. Three experiments varied where the observed and maximum amounts fell in a possible distribution and tested whether this affected choices. Trading was less likely for lower observed amount than higher, but this effect differed depending on the stated distribution. This suggests that participants' trade decisions were affected by where observations fit in the distribution, and thus their probabilities. The modeling tools used here may be applicable to other reasoning phenomena.

Keywords: Probabilistic reasoning, two-envelope problem, mathematical modelling, decision making.

Introduction

The overwhelming evidence of heuristics and biases affecting people's reasoning has often been seen as evidence of irrationality in human thought (Stanovich, 1999). Stanovich pointed out that this conclusion relies on the apparent gap between normatively correct decisions and actual behaviour. However that such a gap indicates irrationality has been challenged by those suggesting that such norms are inappropriate. For example, Oaksford and Chater (1994) suggested that what is seen as an error in the well-known Wason's 4-card selection task is not an error in terms of information gain if you make appropriate assumptions about the distribution of the relevant variables in the environment. Such probabilistic reasoning approaches have gained increasing acceptance (Oaksford & Chater, 2007). Setting normative standards against which to judge rationality is especially difficult when formal analysis of a problem is difficult, such as has been the case for the two-envelope problem (Zabell, 1988). However a recent analysis (Abbott, Davis & Parrondo, 2010; McDonnell & Abbott, 2009) supported by simulations suggests that distributions are critical to analysing that problem, so it is reasonable to ask whether people act rationally by showing sensitivity to distributions when faced with what has sometimes been considered a paradox.

The two-envelope problem

Versions of the two-envelope problem were proposed by Kaitchik (1953, pp. 133-134) and attributed by the mathematician Littlewood to the physicist Erwin Schödinger (Littlewood, 1953/1986, p. 26). Although he does not claim authorship of it, Zabell (1988) stated a two-envelope version with the following characteristics: the contents of the two envelopes are x and $2x$; no distribution or limit is given for x; the reasoner is handed an envelope (randomly) and opens it; however 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 is a 50% possibility then trading appears to result in an expected outcome equal to $5y/4$. Worse, if the two envelopes were held by two different people (as Zabell proposed), then after opening their own envelopes both 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, has defied a satisfactory mathematical solution (Albers, Kooi, & Schaafsma, 2005). So the paradox was resolved but the problem of whether to trade remained.

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). Maslov and Zhang (1998) modelled how switching between volatile assets and non-performing cash reserves can produce a net gain.

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 were told that Envelope 1 (E1) contained a random amount between \$1 and \$100, and Envelope 2 (E2) contained either twice or half that amount depending on the result of a coin toss. They varied whether the participant was given E1 or E2, whether the participant knew which it was, and whether they opened the envelope. If participants

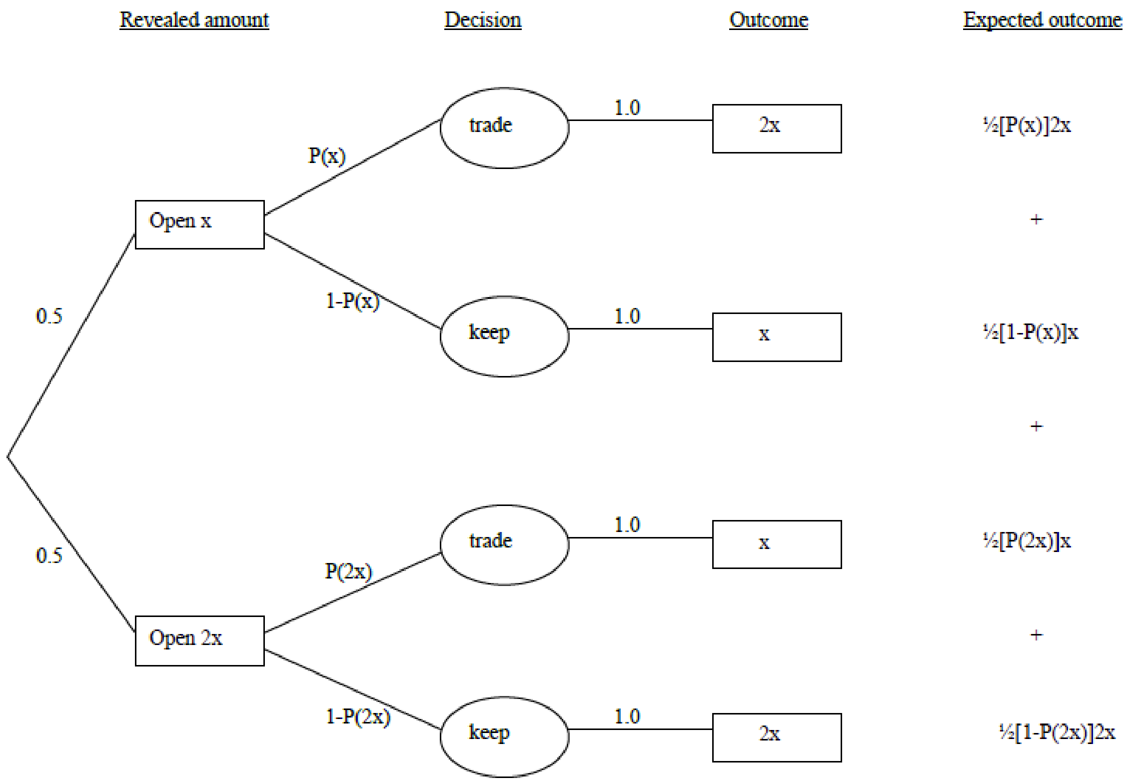

Figure 1: Markov model based on Abbott el al's (2010) analysis. P(x) representing the probability of trading if the value in the opened envelope is x, and P(2x) representing the probability of trading if the observed value is 2x.

observed the amount then Butler and Nickerson asked them what they would do if it had various values ($1, $20, $40, $60, $80, $100). Nickerson and Falk's (2006) analysis of these different versions showed whether it was optimal to always trade, trade depending on the observation, always keep, or to be indifferent. For example, if you know you have been given E1 then it is optimal to trade because E2 was generated from E1 with a 50% chance of two outcomes. They found that if participants observed the amount then there was a tendency towards trading when the amount was less than $50 and keeping when it was above $50, but this was irrespective of whether the conditions should influence their decision. Consistent optimal decision making was rare, so Butler and Nickerson concluded that participants were largely insensitive to the logical structure of the problem; instead they applied simple heuristics or folk wisdom.

A general mathematical solution

Different predictions regarding human performance in the two-envelope problem may arise if there was an accepted mathematical analysis of it. 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 of what it contains breaks the symmetry that leads to Equation 1. Their starting point was Cover's (1987) switching function used to solve 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 converts their analysis to a Markov model. From the model it can be seen that the expected return (E) when x represents the smaller of the two amounts and $2x$ the larger, will be:

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

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 be only be improved if $P(x) > P(2x)$, that is, when the trading function is such that trading is less likely the higher the observed amount is (i.e., the more likely it is to be $2x$ rather than x). 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, but their analysis demonstrates that a simple negative monotonic tendency to trade as a function of observed amount can increase expected outcomes.

Goals

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 in two experiments. A third tested whether the distribution itself was critical. Participants may be acting more rationally than Butler and Nickerson (2008) suggested.

Experiment 1

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 trading rates would be affected by 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 focused on reasoning.

Materials and Procedure. 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. You are told that neither envelope could hold more than $200 [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 $100 [$10]. You are told that you can either keep the $100 [$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 1 presents the proportion of participants choosing to trade in each condition. A logistic regression analysis (using the "Logistic Regression" procedure in SPSS) was performed on choice (0=keep, 1=trade) entering the factors of limit, observed amount, and their interaction. This yielded the following equation for trading:

$$Log(odds) = 1.355 + -0.385*limit + -2.128*observed + 1.39*limit*observed$$

The parameter for limit was not significant, Wald $\chi2(1) = 0.525$, $p = .469$, but that for observed was, Wald $\chi2(1) = 16.224$, $p < .001$, and so was the interaction, Wald $\chi2(1) = 3.885$, $p = .049$.

As predicted, these results showed that participants' choices were affected by the observed contents of the envelope, in that overall there was a strong effect of observed amount. 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.

Table 1: 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)

Experiment 2

An alternative explanation for the interaction between observed amount and limit in Experiment 1 could be that the observed is perceived as worth less in the context of a limit that it is close to. Butler and Nickerson (2008) alluded to such a context effect. So in Experiment 2 participants were directly asked to judge the prior probability of the amount in the envelope. These probabilities should also be lower when the observed amount is half the limit, but such a pattern could not be due to perceptions of monetary value.

Other changes were also made to help generalize the results of Experiment 1. Having a definite limit changes some analyses of the two-envelope problem, so instead of "no limit" a large limit ($10,000) was used. It is unlikely this makes much difference to participants but it removes a potential difference between the two limit conditions. Another possibility is that using such a small amount ($10) for the lower observed amount may have led to trading because it was perceived as a trivial amount. So in Experiment 2 the lower observed amount was set to $50.

The 2x2 design of Experiment 2 was the similar to that for Experiment 1, with factors for limit ($200 or $10,000) and observed ($50 or $100). Again I predicted an interaction between trading and observed such the lowest rate should be when the observed amount was close to the limit.

Method

Participants. A total of 235 senior psychology students participated during practical classes focused on reasoning.

Materials and Procedure. Unlike Experiment 1, the task was presented on a computer. Participants read on-screen instructions that were the same as in Experiment 1 (with appropriate variations for the condition) except that now the envelope they opened was referred to as “Envelope A” and the unopened as “Envelope B”.

Participants were asked the following four questions ($50 replaces $100 in the appropriate condition):

QUESTION 1. First, to check if you understand the instructions correctly, can you type what is the MAXIMUM dollars that Envelope B could contain: $_______

QUESTION 2. What would you do? (click one)
Keep the $100 [$50] in Envelope A
Take whatever is in Envelope B

QUESTION 3. Approximately what do you think is the percentage chance that Envelope A (the one you FIRST opened) contains the LARGER amount of money? ________%

QUESTION 4. In this situation, before any envelopes had been opened, what do you think would have been the probability that the first envelope opened contained $100 [$50] or more? _______%

Results & Discussion

Question 1 was designed to check that participants had correctly understood the problem. Most participants (84.3%) gave the correct answer (either $100 or $200, depending on condition), but rates of correctness were not affected by condition. It was decided that participants who did not answer this question correctly either misinterpreted the instructions or were not paying attention. Either way their responses could not be relied on, so only the 198 participants who answered correctly were analysed.

Table 2: Proportion of participants in each condition of Experiment 2 choosing to trade. Samples sizes are in parentheses.

	$50 in opened envelope	$100 in opened envelope
$200 limit	.65 (n=51)	.30 (n=61)
$10,000 limit	.55 (n=38)	.60 (n=48)

Table 2 shows the proportion of participants in each condition choosing to trade envelopes in response to Question 2. (Sample sizes varied because participants were randomly assigned to a condition by their individual computer.) A logistic regression analysis was performed on choice (0=keep, 1=trade) entering the factors limit, observed amount, and their interaction. This yielded the following equation for trading:

Log(odds) = 0.534 + -0.483*limit + -1.374*observed + 1.728*limit*observed

The parameter for limit was not significant, Wald $\chi2(1)$ = 1.190, p = .275, but that for observed was, Wald $\chi2(1)$ = 11.190, p = .001, and so was the interaction, Wald $\chi2(1)$ = 8.420, p = .004. So the Experiment 1 interaction pattern was replicated despite changing the lower observed amount, the specification of the higher limit, and mode of presentation.

In response to Question 3 most participants (92.2%) thought there was exactly a 50% chance that the other envelope would contain more money. The overall mean response was 49.49%, and there were no effects of condition. Thus, despite choosing to keep or trade their envelope, very few participants seemed to think the odds of the other envelope containing more was other than 50%. Even if participants act as though sensitive to a distribution, this does not necessarily mean they are aware of it (e.g., Bargh & Ferguson, 2000).

Table 3: Mean judgments (with standard deviations) of prior probabilities (percentages) for each condition.

	$50 in opened envelope	$100 in opened envelope
$200 limit	57.25 (sd=17.3)	46.70 (sd=14.1)
$10,000 limit	54.41 (sd=29.1)	58.22 (sd=25.5)

In response to Question 4 most participants thought that there was about a 50% probability that their envelope could have contained an equal or higher amount before it was observed, but Table 3 shows that this varied with condition. A 2x2 ANOVA found no main effects of limit, $F(1,194)$ = 1.980, p = .161, no effect of amount observed, $F(1,194)$ = 1.190, p = .277, but a significant interaction, $F(1,194)$ = 5.409, p = .021. Thus consistent with the observed $100 and limit $200 condition being the one least likely to lead participants to favour trading envelopes, participants in this condition were also least likely to think that their envelope could have contained more a priori. Why Question 4 but not 3 showed a difference may be because it does not so starkly ask participants to contradict their intuition that two coin-flip like choices should mean 50% each.

By replicating the interaction found in Experiment 1, Experiment 2 further supported the hypothesis that participants are less likely to trade when the higher amount would be at the end of the distribution. Adding support to the claim that this was because of where they felt the observed amount fell in the distribution the manipulations had a similar effect on a direct measure of how likely they thought that the observed amount could have been higher.

Experiment 3

Abbott et al’s (2010) solution suggests that people may be less likely to trade when the observed amount is higher in the distribution, but working out the optimal trading strategy would depend on knowing the details of the distribution of amounts. If people act consistent with this analysis, then people’s tendency to trade should be affected by what they believe about the distribution. So far the results suggest that that people’s responses reflect a sensitivity to the distribution of amounts, so explicitly stating a different distribution could affect their choices.

In Experiment 3 participants were told that the envelope amounts had either a flat or a bimodal distribution. It is likely that many participants assumed an essentially flat distribution in the previous experiments, in which case explicitly stating that the distribution is flat should produce similar results to Experiments 1 and 2. However explicitly stating that there was a bimodal distribution could lead to a different pattern of results. By increasing the chances of high amounts in envelopes this distribution should increase trading when the other envelope potentially contains an amount at the top of the distribution. A 2x2 design was used with factors for distribution (flat or bimodal) and observed ($50 or $100). The limit was always $200.

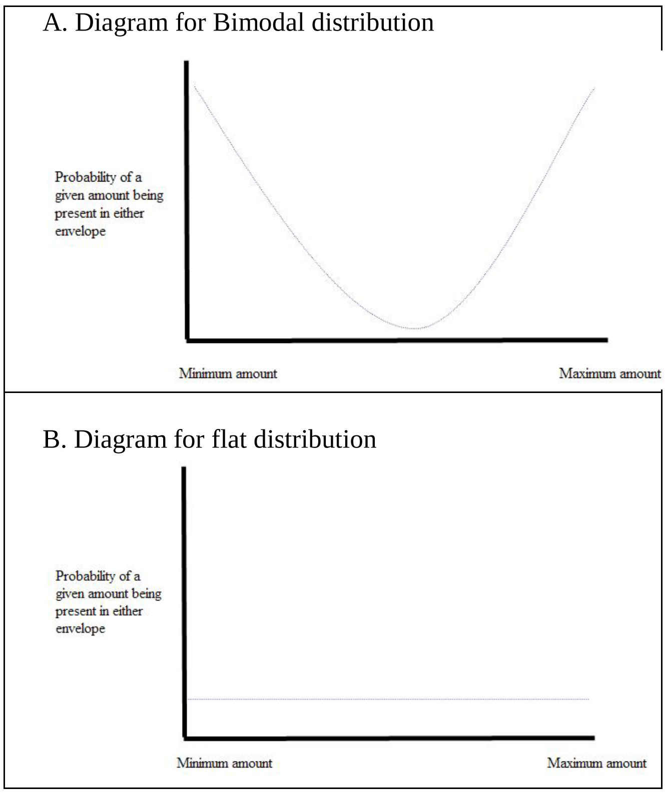

Figure 2: Diagrams accompanying the instructions for the bimodal (Panel A) and flat (Panel B) distributions.

Method

Participants. One hundred and three first-year psychology students completed the experiment for partial course credit.

Materials and procedure. Materials and procedure were identical to Experiment 2 except for the addition of the distribution manipulation. In the flat condition participants read that "the probability of any amount is equal to any other" and saw the graph in Panel B of Figure 2. In the bimodal condition they read "the probability of any amount is not equal, in that amounts closer to the minimum or maximum amounts are more likely" and saw the graph in Panel A of Figure 2. These graphs were intentionally vague in order to give a general shape to the distribution rather than provide a precise way to calculate the probabilities. The observed amount in the opened envelope was either $50 or $100, but the maximum possible was always $200.

Results & Discussion

Most (83.5%) participants correctly identified the maximum amount the unopened envelope could contain, but as in Experiment 2 only these 86 were analysed. The proportion in each condition choosing to trade is presented in Table 4.

Table 4: Proportion of participants in Experiment 3 choosing to trade. Maximum amount was always $200.

	$50 in opened envelope	$100 in opened envelope
Bimodal distribution	.25 (n=20)	.42 (n=24)
Flat distribution	.58 (n=19)	.26 (n=23)

For the flat distribution the trading proportions were similar to the same conditions in Experiment 2 in which no distribution was specified, with more trading when the observed amount was $50 than when $100, $\chi2$ (1) = 4.37, p = .037. In the bimodal condition, the direction of the effect of revealed amount was the opposite, but this effect was not significant, $\chi2(1)$ = 1.35, p = .246. A logistic regression analysis was performed on choice entering the factors distribution (0=bimodal, 1=flat), revealed amount, and their interaction. This yielded the following equation for trading:

Log(odds)= -1.099+1.417*distribution + 0.762*observed + -2.122*distribution*observed

The parameter for distribution was significant, Wald $\chi2(1)$ = 4.161, p = .041, but not that for observed, Wald $\chi2(1)$ = 1.326, p = .250. The interaction parameter was significant, Wald $\chi2(1)$ = 5.120, p = .024.

These results indicated that people were sensitive to the distributions of amounts when deciding whether to trade. For the same amount with the same limit their propensity to trade was influenced by what they were told about the distribution of amounts. When the distribution was flat they responded similarly to how they did in Experiment 2, suggesting that participants had previously assumed a flat distribution. However a bimodal distribution changed the pattern of their responses implying that they took into account the prior probabilities of different amounts.

It should be noted that the Figure 2 distributions are only possible for either the higher or the lower amounts, not the sum of their distributions. Given that participants do not know if they observe the higher or the lower amount they may have been confused as to what exactly was the distribution represented by their diagram. However the main point of the experiment was to test whether the distribution plays a role in participants' choices, and confusion about the distribution should not affect their choices unless they see the distribution as important.

General Discussion

Abbott et al's (2010) analysis suggests that a probabilistic strategy for trading can lead to gains in the two-envelope problem unobtainable by a pure strategy. In general, such a probabilistic strategy can increase expected outcome over an absolute strategy if the probability of trading is a monotically decreasing function of the observed amount. This suggests that people given the two envelope problem may have a tendency to trade that is sensitive to the distribution of amounts. The results of Experiments 1-3 support the claim that people do this when faced with the two-envelope problem. Participants were consistently least likely to trade when the higher alternative would be at the top of the distribution, except in Experiment 3 when the bimodal distribution increased the likelihood of such an amount. Furthermore, in Experiment 2 it was found that participants' assessments of the prior probabilities of amounts had the same pattern. Thus the results suggest that people are consistent with what Abbott et al's model suggests optimizes responses to the two-envelope problem: trading as a function of the observed amount and being sensitive to the distributions. Experiment 3 is critical in showing that not just the size of the observed amounts but their perceived distribution affected choices. However this conclusion is weakened by possible limitations of its methodology, therefore more research is required.

These experiments did not systematically manipulate the amount in the revealed envelope to see what shape there might be to any monotonic function to trade. Inspection of Butler and Nickerson's (2008) data suggests that there is a trend within the large effect of greater/lesser than $50 towards less trading as observed amounts increase. However their sample size is not large enough to expect a post-hoc analysis to show a significant effect. Overall, the results do not dispute Butler and Nickerson's finding that participants often make fundamental errors in analysing the two-envelope problem. The errors they revealed were in understanding the logical implications of the details of different versions, and in this way they are analogous to Wason's (1968) finding that people were poor at understanding the logical implications of his selection task. However Oaksford and Chater's (1994) analysis showed that people's responses may make sense if seen in terms of how information is distributed in the world. Thus my results fit with a more general trend of finding that people are poor at applying formal logic but can be sensitive to the implications of probability distributions. Applying probabilistic inference may be seen as the computational goal of cognition.

McDonnell and Abbott (2009) saw the two-envelope problem as interesting because it embodies a phenomenon that comes up in many domains, that of probabilistic switching between two states. Their analysis demonstrates that an appropriate probabilistic function may improve outcomes even when important characteristics of the distributions are unknown. A number of decision making tasks require a choice between functions whose properties are uncertain, for example, choices between different market options. The demonstration here that the mathematical analysis of such choices can lead to supportable behavioural predictions suggests that these mathematical tools may have value for analysing other types of decisions.

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.

Bargh, J. A., & Ferguson, M. J. (2000). Beyond behaviorism on the automaticity of higher mental processes. *Psychological Bulletin, 126*, 925-945

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.

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.

Littlewood, J. E. (1953/1986). *Littlewood's Miscellany* (B. Bollobas, Ed.). Cambridge, UK: Cambridge University Press.

Maslov, S., & Zhang, Y. (1998). Optimal investment strategy for risky assets. *International Journal of Theory and Applications of Finance, 1*, 377–387.

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

Nickerson, R. S., & Falk, R. (2006). The exchange paradox: Probabilistic and cognitive analysis of a psychological conundrum. *Thinking and Reasoning, 12*, 181–213.

Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review, 101*, 608-631.

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.

Wason, P. C. (1968). Reasoning about a rule. *Quarterly Journal of Experimental Psychology, 20*, 273-281.

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.