

Causality and Reasoning: The Monty Hall Dilemma

Bruce D. Burns (burnsbr@msu.edu)

Mareike Wieth (wiethmar@msu.edu)

Department of Psychology; Michigan State University
East Lansing, MI 48824-1117 USA

Abstract

In the Monty Hall Dilemma (MHD) contestants try to choose which of three doors conceals a prize. After selecting a door, one of the other doors is opened by a host who knows where the prize is, but always reveals a dud. Contestants are then asked if they want to stay with their first choice, or switch to the other unopened door? Switching yields a two-thirds chance of winning, but most people have difficulty accepting this answer. Glymour (2001) points out that central to the MHD is a particular causal structure, the *collider principle*. We hypothesized that an isomorphic version of the MHD that would help participants understand its causal structure would improve their performance. Making the MHD a form of competition should be one way to achieve this. We confirmed this by showing that in a competition version of the MHD participants were much more likely to solve the problem, and more likely to answer a counterfactual question indicating a correct understanding of the problem's causal structure. Furthermore, regardless of MHD version, participants who solved the problem were more likely to also answer the counterfactual correctly. Thus the MHD can be seen as an example of people's difficulties with understanding even relative simple causal structures.

Causality and Reasoning

Wilson and Keil (2000) discuss how explanation pervades all human activity from the simplest and most mundane to the most sophisticated and unusual, such as explanations in science. Yet they go on to point out that explanation remains one of the most under-explored topics in the cognitive sciences. When explanation is explored, virtually all notions of it invoke the concept of cause (Wilson & Keil, p.105) and the centrality of cause in explanations has also been emphasized by developmental psychologists. Baillargeon and Wang (2002) argue that babies as young as three months are reasoning about causality. Thus, causality and how it is used is an important topic in cognitive science as it is critical to understanding how people reason about events.

Despite its importance, people often find causality hard to reason about. In this paper we will demonstrate that the difficulty people have with the well known Monty Hall dilemma is due to it involving a causal structure that people find hard to recognize and reason with. The theoretical basis of our claim that misunderstanding of causality is the key to this problem will be a recent analysis by Glymour (2001).

The Monty Hall dilemma

Like optical illusions, cognitive illusions often persist long after evidence to the contrary has been presented. This makes them interesting because, as with visual illusions, such tenacity should tell us something about the underlying cognitive system. The Monty Hall Dilemma (MHD) is such an illusion. Named after the host of the television show "Let's Make a Deal", it presents a participant with three doors, one of which hides a prize. The participant selects a door hoping it conceals the prize, as behind the other two doors are duds. The host then opens one of the other doors to reveal a dud. The host knows where the prize is, but of course never opens the door hiding it or the door the contestant selected. The participant is then given a choice: stay with the door initially selected, or switch to the other unopened door?

When faced with the choice between the two remaining doors, most people have the very strong intuition that it's a 50% chance either way, and usually they stay with their first choice. However, participants have a two-thirds chance of winning if they switch (Selvin, 1975) therefore one should always switch. The strength of this illusion was perhaps best illustrated by Marilyn vos Savant who in 1990 published the dilemma and this answer in Parade magazine (vos Savant, 1997). In response she received 10,000 letters of which 92% from the general public disagreed with the answer. Empirical reports of this illusion confirm its strength (Friedman, 1998; Granberg, 1999; Granberg & Brown, 1995; Granberg & Dorr, 1998). In these papers, seven studies that used isomorphic versions of the MHD were reported with switch rates ranging from 9% to 17%.

Another noteworthy aspect of vos Savant's letter writers was that education was not necessarily an antidote to the illusion. Of letters that had university addresses 65% disagreed, including ones from statistics and mathematics professors (vos Savant, 1997). Schechter (1998) in his biography of Paul Erdős, one of the greatest mathematicians of the twentieth century and a man who lived to solve mathematical problems, relates what happened when Erdős was told the correct solution to the MHD by a fellow mathematician. Erdős insisted that the answer was incorrect, and could not be convinced even when it was explained to him in the language of mathematics. Eventually a computer simulation convinced him that switching was correct, yet he remained frustrated by his inability to intuitively understand why. This peerless authority on probability was only mollified when several days later another mathematician

friend made him see his error (Schechter, pp. 108-109). These anecdotes illustrate an important point about the MHD, that knowledge about probability is useful but not sufficient for solving the MHD. Thus the tenacity of the MHD seems not so much due to misunderstanding probability, but due to a misrepresentation of the problem.

Explanations of the illusion

There have been various attempts to explain this illusion. One claim has been that the MHD contains ambiguities and unstated assumptions. Nickerson (1996) pointed out that the mathematical analysis supporting switching relies on two assumptions: the host always (or randomly) makes the switch offer; and, the host always opens a door hiding a dud prize. Without the first assumption the probability is indeterminate; without the second it is 50%. Granberg and Brown (1995) found that most participants gave the probability of winning if they switched as 50%, rather than indeterminate or any other number. Therefore participants seemed to be making the first assumption as were Vos Savant's (1997) letter writers who strongly insisted that the probability was 50% because there are two doors left. There is also no reason to think people fail to make the second assumption. Whereas thinking that the host could open the door concealing the prize would make switching a 50% chance; such behavior would not fit with the concept of a game show. Empirical studies of the MHD have carefully made clear that a door is always opened, that the door to be opened cannot hide the prize, and that a knowledgeable host always offers the switch choice; yet the illusion persists.

Reliable increases in switching have been observed when participants experience multiple trials. Granberg and Brown (1995, Study 2) and Granberg and Dorr (1998, Exp. 1) both found that after 50 trials participants' mean switch responses increased to 55%. Friedman (1998) gave participants money for winning and was able to increase switching to 40% but only after participants had already been exposed to seven trials. It is known that people are good at frequency detection (Hasher & Zacks, 1984) so it is unsurprising that people learn to increase their rate of making a response which is successful on two-thirds of the trials. Yet turning the MHD into a learning problem does not explain why so few people make switch decisions to start with.

Vos Savant (1997) suggested that if the problem is modified so that the contestant is faced with many doors, of which all but one is opened after an initial selection, then people will better understand that they should switch. Granberg and Dorr (1998, Exp. 2) found that increasing the number of doors from three to seven (opening five of seven means switching yields a 83% chance of winning) increased switching rates from 11% to 25%, although this difference was not statistically significant. This modification changes the problem substantially, but it suggests something interesting: participants may need to recognize that the process resulting in the unopened door is crucial.

Why stay?

When asked, most experiment participants say that there is a 50% chance of winning if they switch; yet they almost always stay. If they truly thought of it as a 50% chance, then

one would expect half the participants to switch, but this is not the case. Just how strong the stay response is was illustrated by Granberg and Brown (1995) in a modified card game version of the MHD. Even if the payoff from winning after a switch was made substantially greater than that after a stay choice, most participants chose to stay.

This tendency to stay with one's first choice is more general than the MHD; for example, Mathews (1929) found that although students believe that it is best to stick with their first choice on multiple choice tests, 53% of changed answers gained points and only 21% lost points. This tendency to stay may be interesting in its own right, but probably explains little about why the MHD is so hard. The stay bias does not explain why people say that the probability of winning is 50% either way. Rather it appears that two phenomena may combine in creating so many incorrect answers to the MHD: 1) people assess the probabilities as 50%, 2) they give staying a higher value. We will focus on the first phenomenon: why does it seem so compelling that the probabilities should be 50% once there are two options? However, the bias to stay will be useful methodologically as it will allow us to use switch responses as our dependent variable. As the empirical studies have shown, most participants will stay in the absence of a strong and compelling reason to switch.

Why 50%?

Johnson-Laird Legrenzi, Girotto, Legrenzi, and Caverni (1999) discuss the propensity of people to give an equal chance to all available possibilities they are presented with, unless individuals have beliefs or knowledge to the contrary. They pointed out that an analogous principle of indifference or insufficient reason has a long history in probability theory. Johnson-Laird, et al. (1999) also presented some empirical evidence that people tend to follow this principle, as did Falk (1992) who calls it the "uniformity assumption." Of course, people do not necessarily see all possibilities as equally probable, as the odds-making in sports-betting abundantly demonstrates. However when people confidently assign each of two options a 50% chance they are indicating that they can see no reason to differentiate the options. In the MHD the mathematical analysis shows that the choices are different, so we sought a way to make this fact clear.

Johnson-Laird et al. (1999) argued that people's poor performance on the MHD is due to them not forming the right representation in terms of mental models. Krauss and Wang (2003) presented some interesting evidence that giving people an appropriate representation improved performance. We do not doubt that if people are provided with the right set of mental models, then they will probably be more likely to solve the problem. However, this does not necessarily answer a more fundamental question, why do people have such difficulty constructing for themselves the right representation? Even when told that they are wrong, people will often be unable to recognize that there is an alternative to their current representation. We propose that what blocks generating the right representation is people's failure to understand the MHD's causal structure.

Causality in the Monty Hall Dilemma

Glymour (2001) suggests that the causal structure of the MHD is an example of the *collider principle*, a structure that Glymour speculates people may often have problems with. This principle is illustrated by a causal graph in which two independent variables have an edge ("collide") into a third variable (Pearl, 2000, refers to these as inverted fork causal graphs). For example, for a car to start the tank must not be empty and the battery must have sufficient charge to turn the starting motor. The two variables of the state of the fuel tank and the battery are independent, however both have a causal influence on the car starting. Thus if the car does not start then the fact that the tank is not empty provides information about the battery (i.e., it is probably dead). Thus once knowledge about the car starting is provided, the states of the tank and the battery are dependent conditional on that knowledge. A similar causal structure underlies the Berkson (1946) paradox in epidemiology. Berkson pointed out that if two diseases are unrelated but both are causal with regard to putting someone into the hospital, then the two diseases can be correlated in the hospital population. Pearl (2000, p.17) refers to such selection biases as the *explaining away effect* as, for example, it can lead to music ability and SAT scores being negatively correlated at a college that admits students on the basis of either of these factors. Many people are familiar with the idea of two variables being correlated due to both being caused by a third variable. Such a structure can be represented by a causal graph in which two outcomes have causal links from a single factor. In a sense, the collider principle has the opposite causal structure, one outcome with causal links from two factors. Fewer people seem to understand that this structure can also create correlations between otherwise independent variables.

Glymour (2001) points out that the MHD has a collider causal structure, as the placement of prize and the contestant's initial choice both have a causal influence on which door the host will open. This results in these two, otherwise independent, causal factors being dependent conditional on which door the host opens. Therefore when the host opens a door information is provided regarding the placement of the prize, which leads to it being advantageous to switch.

This analysis leads to testable hypotheses. When presented with the MHD, if problem solvers do not grasp that the placement of the prize has a causal influence on what door the host opens, or they fail to understand the implications of this causal structure, then they will be likely to give the wrong answer. If there truly was no causal link between the placement of the prize and the host's choice, then the contestant would be correct to think that it did not matter which door was the final choice (Nickerson, 1996).

It is unlikely that a single explanation alone accounts for the difficulty of the MHD, and we are not making that claim here. Instead our claim is that a major factor behind people's poor performance in the task is that the standard versions of the MHD obscure the causal structure of the problem. Therefore our overall prediction was that presenting participants with scenarios analogous to the MHD that increase the salience of the causal structure by having the

options compete, should lead to better reasoning than standard versions of the MHD.

Manipulating Understanding of Causal Structure

Glymour (2001) did not attempt to provide any evidence that his analysis explains why the MHD is so hard, nor for his more general claim that collider structures are difficult to understand. Burns and Wieth (2000) set out to test the idea that failure to correctly understand the causal structure of the MHD was a factor behind the failure of people to solve it. They did so by creating isomorphic versions of the problem that varied in terms of likelihood that people would see the problem in terms of causality. Versions with isomorphic structures placed in a context that should be more likely to invoke the right causal structures should lead to more people solving the problem correctly.

In competitions it seems easier to understand the implications of the collider principle, perhaps because people often seem to see competitions in terms of causality (Lau & Russell, 1980; McAuley & Gross, 1983; White, 1993). For example, in a game between a competitor that we know has just defeated an opponent versus one who is yet to play a game, it may seem natural to favor the victor of the completed game. People may vary in how they explain such a preference, but to the extent that this preference is correct (disregarding factors such as practice) it is due to the collider principle. Who won the previous game is caused by two factors: which of the three players was not in the game, and which player is the best (by some relevant criterion). Thus who won the first game creates a degree of dependency between who won the previous game and which of the two remaining options is better, just as the collider principle does in general. Therefore isomorphic versions of the MHD that place it into a competitive context should lead to more solutions.

Burns and Wieth (2000) presented a set of scenarios that varied in terms of how genuine was the competition involved. In the strongest version we replaced the three doors with three boxers who will fight a pair of bouts. One of the three boxers was so good that he was guaranteed to win any bout. After the contestant selects one of the three boxers as his or her pick to be the best, the remaining two boxers fight. The winner of the first bout will then fight the boxer initially selected, and contestants win if they chose the winner of this second bout. However, the contestant is offered the choice after the first bout: stay with the initial selection, or switch to the winner of the first bout. This Boxers version is isomorphic to the MHD: the three boxers represent doors and the best boxer (i.e., guaranteed winner of any bout) represents the door concealing the grand prize. All that varies is how one of the unchosen options is eliminated: in the MHD the host applies a rule (i.e., open an unselected door without the prize, or if the initial selection hid the prize then open a random door) whereas in the Boxers version the option is eliminated by a competition (i.e., the single best boxer eliminates the opponent in the first bout, or if the best boxer is the one initially selected, then it is random which of the other two is eliminated).

In order to manipulate the degree of genuineness of

competition presented by the scenarios, we also created a Wrestlers version. The Wrestlers version was identical to the Boxers version, except that professional wrestlers replaced boxers, and it was pointed out that the results of professional wrestling matches are determined beforehand. So although which of the unselected options is eliminated is decided by competition, the results of all possible competitions have been predetermined. To make the competition even less genuine, a Wrestlers-D version was created in which the wrestlers defended doors, and the door concealing the grand prize (placed there before the matches) was defended by the "best" wrestler. If a wrestler lost a match, then that wrestler's door was opened. A Doors version was created identical to the Wrestlers-D version, except that now the wrestlers were not directly involved in which door was opened. They just stood in front of doors and yelled at each other until the predetermined door was opened by the knowledgeable host. The Doors version was designed to be closest to the standard versions of the MHD.

Burns and Wieth's (2000) preliminary study found that 51% of participants switched in the Boxers condition, 30% in the Wrestlers condition, 37% in the Wrestlers-D condition, and only 15% in the Doors condition. Thus, as predicted, participants were more likely to solve the problem correctly when presented with versions that made the causal structure of the problem easier to understand. However, we did not directly show that our manipulation led people to a better understanding of the causal structure of the MHD.

An Experiment

Burns and Wieth (2000) showed an effect of manipulations that should increase the likelihood of participants understanding the causal structure of the MHD. This new experiment was designed to demonstrate that participants in a competition condition really did have a better understanding of the causal structure, and that within each condition those with a better understanding of the causal structure would be more likely to solve the problem. In addition, the experiment was intended to replicate the difference between the Boxers (competition) condition and Doors (noncompetition) found by Burns and Wieth.

Glymour's (2001) analysis suggests that the critical element to understanding the causal structure of the MHD is the causal link from the location of the best option to which option is eliminated. If participants fail to understand the causal implications of this link, then they will not understand the problem correctly. Thus to probe participants understanding of this link, we presented them with a counterfactual: what if this link was random? Thus in the Doors condition participants were asked whether it would affect their answer if they learnt that the host indicating which door to open did not actually know where the prize was, although fortunately a door without the prize had been opened. In the Boxers condition participants were told that all the boxers had the flu, thus who won each bout was random.

If the competition manipulation was effective in Burns and Wieth (2000) because it led more participants to understand the causal structure of the MHD, then more

participants in the Boxers than the Doors condition should indicate that the counterfactual would, or at least might, affect the answer to the problem that they had already given. Furthermore, if this understanding is the key to solving the MHD, then participants in *either* condition who gave the right answer should be more likely to indicate that the counterfactual might change their answer.

One possibility though is that participants in the Boxers condition simply understood the instructions better, thus they may simply better understand that there is a nonrandom link from the option that represents the prize to the option that is eliminated. To check this possibility, we asked participants to indicate how random was this link.

Method

Participants. A total of 124 members of the Michigan State University subject pool participated in the experiment.

Procedure. The study utilized two versions of the MHD used by Burns and Wieth (2000). In each a situation was first described in which a person is randomly selected at an event (either a boxing or wrestling night). Each time, this person is presented with three options, one of which represents a substantial prize. After an option is selected, one of the remaining two options is shown not to be the critical one (prize or best competitor). The person then always has to make a decision: stay with the first selection, or switch to the other remaining option? Therefore it was made clear that the switch offer was always made, as the competition was conducted in the same way every week. Participants were then told that he or she was the randomly selected person, an option was eliminated, and the choice to "stay" or "switch" then had to be made by the participant. They were then asked to indicate how many out of nine people they expected to win if all nine were to stay, and if all nine were to decide to switch. This gave an indication of whether participants knew the probability of winning if they switched. The two versions of the MHD were as follows:

Competition (Boxers): Three boxers will fight. One of the boxers is so good that he is guaranteed to beat either of the others, no matter what. If you can select this boxer you win the prize, but you have no basis at all on which to evaluate who is the best. First, you select one boxer randomly (participants are told who they selected), then the other two fight. The winner of the first contest will fight the one you initially selected, and the winner of that second bout represents the best boxer. After the first bout you have to decide whether to stay with your first selection, or switch to the winner of the first bout.

Noncompetition (Doors): This condition was closest to the standard version of the MHD. As in Burns and Wieth (2000), the wrestlers did not fight but instead just yelled at each other while standing in front of doors. The host, who knew where the prize was, indicated which door to open.

Further questions were presented on a second page. Participants first answered on a six point scale the question "The process by which Boxer/Door A was eliminated is best described as." On the scale "1" indicated "completely

random" and "6" indicated "completely nonrandom." This scale was used by Burns & Corpus (in press) to show differentiation of scenarios in terms of randomness.

Participants then answered the counterfactual question, which asked whether it would have affected their answer if the way an option was eliminated was completely random. They could answer YES, NO, or MAYBE. In the Noncompetition condition, the counterfactual was explained as the host not knowing where the prize was, and in the Competition condition as the boxers being sick such that who won any bout was random. In both cases it was emphasized that this would make it completely random which option was eliminated, but that luckily the eliminated option was not the one representing the prize.

Results and Discussion

From Table 1 it can be seen that again participants in the competition condition performed better in that they were more likely to decide to switch (42% versus 14%, $\chi^2[1] = 12.20, p < .001$). In addition, participants in the Competition condition were more likely to give the correct probability of winning (16% versus 5%, $\chi^2[1] = 4.46, p = .035$), in that they indicated that 6 out of 9 people would be expected to win if all nine decided to switch. Although relatively few of those participants who decided to switch could calculate the correct probability of success, this could be because people are poor at calculating conditional probabilities in general (see Kahneman, Slovic, & Tversky, 1982). There was a greater impact of the competition manipulation on simply recognizing that the chance of winning when switching would be greater than 50% (45% versus 19%, $\chi^2[1] = 10.14, p = .001$). Thus we replicated the previously found effects of the competition manipulation.

Table 1: Number of participants choosing each answer to the counterfactual question depending on their decision to stay with their first option or switch, for both the competition and noncompetition conditions.

		stay	switch
Competition condition			
Counterfactual Response	YES	4	11
	MAYBE	17	8
	NO	14	6
Total		35	25
Noncompetition condition			
Counterfactual Response	YES	2	2
	MAYBE	14	4
	NO	39	3
Total		55	9

Table 1 shows that participants in the Competition condition were more likely to answer yes or maybe to the counterfactual question, $\chi^2(2) = 15.20, p = .001$. In the Competition condition, 67% recognized that making the elimination process random at least might make them reconsider their answer, whereas only 34% did so in the

Noncompetition condition. Thus the competition manipulation did not just produce more correct answers to the problem, but it also affected participants in the way we expected it would: it led participants to be more likely to recognize the implications of the elimination process. Furthermore, participants who solved the problem were also more likely to recognize the implications of making the elimination process random whether they were in the Competition, $\chi^2(2) = 8.27, p = .016$, or Noncompetition conditions $\chi^2(2) = 6.93, p = .031$.

There was no difference in randomness ratings by participants who decided to stay ($M = 4.1, SD = 1.6$) or switch ($M = 4.2, SD = 1.6$), $t(122) = 0.24, p = .81$. Participants in the Noncompetition condition actually saw the process as more nonrandom ($M = 4.5, SD = 1.4$) than did those in the competition condition ($M = 3.7, SD = 1.6$), $t(122) = 3.0, p = .003$. Thus the competition manipulation did not appear to affect participants' understanding that there was a causal link between the prizes' placement and which option was eliminated. Instead participants in the Noncompetition condition were less likely to recognize the significance of this causal relationship.

Discussion

By presenting participants with a scenario in which they were better able to understand the causal structure of the process by which an option was eliminated (as indicated by responses to the counterfactual question), we were able to greatly increase the solution rate for the MHD and the number of participants correctly calculating the probability of winning if they switched. Furthermore, regardless of their condition, participants who gave the correct answer to the MHD were more likely to indicate an understanding of the causal structure of the MHD. Thus we have evidence for the hypothesis that a major reason why the MHD is such a difficult problem is that its causal structure is one that people have difficulty understanding.

The lack of a difference in the randomness ratings suggests that it was not a difference in the extent to which participants saw that process as random that led to any difference in performance. Instead, it is lack of understanding of the significance of the causal link between the elimination of an option and which option was the best. This is consistent with Glymour's (2001) claim which was that people find it hard to understand the implications of the collider principle, not that they fail to recognize the causal links that exist in a situation. We do not claim that people get the standard MHD wrong because they think that Monty acted randomly, but because they fail to understand what significance the causal structure gives to his actions.

However, the illusion does not "disappear" when participants are presented with a competition scenario. This is not surprising given that the counterfactual responses indicate that our manipulation was not successful in leading all participants in the competition condition to correctly represent the causal structure of the MHD. Furthermore, a correct representation is just one step on the path to a correct solution, as has been found with other problems for which an incorrect representation contributes to the difficulties people experience. For example, Weisberg and

Alba (1980) found no participants who solved the nine-dot problem without a hint to go outside the square defined by the nine dots, but even with the hint only 25% of participants solved the problem. Therefore we did not expect to completely eliminate a strong illusion with a simple change to its context.

Giving more direct aids to representation should also help people solve the MHD, as Krauss and Wang (2003) found. However, our focus has been on why do people have such difficulty in setting up the right representation themselves? We suggest that without an understanding of the causal structure, participants will be forever plague by the thought "But why isn't it 50% if there are only two left?"

This first evidence for Glymour's (2001) analysis of the difficulty of the MHD as due to it involving the collider principle, suggests that misunderstanding of this form of causal structure may underlie other reasoning errors. The hundreds of citation in the medical and epidemiology literature of Berkson (1946) suggest there is at least one other form of reasoning error that appears to arise from a difficulty in understanding the collider principle. Yet the collider principle seems a relatively simple causal structure, as it consists of just two causal variables and one outcome variable. As Pearl (2000) points out, making proper inferences about causality is the central aim of the physical, behavioral, social, and biological sciences. However, if the collider principle can cause huge problem for even experts in statistics, it may not be so surprising that Pearl also observes that we have difficulty understanding causality. In further research we will try to investigate the conditions under which the collider principle is hard to recognize and use. In this way we may gain insight into why even simple causal structure can be hard to understand.

Acknowledgements

We would like to thank Fernanda Ferreira and Regina Vollmeyer for helpful comments on an earlier draft.

References

Baillargeon, R., & Wang, Su-hua (2002). Event categorization in infancy. *Trends in Cognitive Sciences*, 6, 85-93

Berkson, J. (1946). Limitations of the application of fourfold table analysis to hospital data. *Biometrics Bulletin*, 2, 47-53.

Burns, B. D., & Corpus, B. (in press). Randomness and inductions from streaks: "Gambler's fallacy" versus "Hot hand." *Psychonomic Bulletin & Review*.

Burns, B. D., & Wieth, M. (2000). *The Monty Hall Dilemma: A causal explanation of a cognitive illusion*. Presented at the Forty-First Annual Meeting of the Psychonomic Society, New Orleans, LA.

Falk, R. (1992). A closer look at the probabilities of the notorious three prisoners. *Cognition*, 43, 197-223.

Friedman, D. (1998). Monty Hall's three doors: Construction and deconstruction of a choice anomaly. *The American Economic Review*, 88, 933-946.

Glymour, C. N. (2001). *The mind's arrow: Bayes nets and graphical causal models in psychology*. Cambridge, MA: MIT Press.

Granberg, D. (1999). Cross-cultural comparison of responses to the Monty Hall Dilemma. *Social Behavior and Personality*, 27, 431-438.

Granberg, D., & Brown, T. A. (1995). The Monty Hall dilemma. *Personality and Social Psychology Bulletin*, 21, 711-723.

Granberg, D., & Dorr, N. (1998). Further exploration of two-stage decision making in the Monty Hall dilemma. *American Journal of Psychology*, 111, 561-579.

Hasher, L., & Zacks, R. T. (1984). Automatic processing of fundamental information: The case of frequency of occurrence. *American Psychologist*, 39, 1372-1388.

Johnson-Laird, P. N., Legrenzi, P., Girotto, V., Legrenzi, M. S., & Caverni, J.P. (1999). Naive probability: A mental model theory of extensional reasoning. *Psychological Review*, 106, 62-88.

Kahneman, D., Slovic, P., & Tversky, A. (Eds.) (1982). *Judgement under uncertainty: Heuristics and biases*. Cambridge, UK: Cambridge University Press.

Krauss, S., & Wang, X. T. (2003). The psychology of the Monty Hall problem: Discovering psychological mechanisms for solving a tenacious brain teaser. *Journal of Experimental Psychology: General*, 132, 3-22.

Lau, R. R., & Russell, D. (1980). Attributions in the sports pages. *Journal of Personality & Social Psychology*, 39, 29-38.

Mathews, C. O. (1929). Erroneous first impressions on objective tests. *Journal of Educational Psychology*, 20, 280-286.

McAuley, E., & Gross, J. B. (1983). Perceptions of causality in sport: An application of the Causal Dimension Scale. *Journal of Sport Psychology*, 5, 72-76.

Nickerson, R. S. (1996). Ambiguities and unstated assumptions in probabilistic reasoning. *Psychological Bulletin*, 120, 410-433.

Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge, UK: Cambridge University Press.

Schechter, B. (1998). *My Brain is open: The mathematical journeys of Paul Erdős*. New York: Simon & Schuster.

Selvin, S. (1975). A problem in probability [Letter to the editor]. *The American Statistician*, 29, 67.

vos Savant, M. (1997). *The power of logical thinking*. New York: St Martin's Press.

Weisberg, R. W., & Alba, J. W. (1980). An examination of the alleged role of "fixation" in the solution of several "insight" problems. *Journal of Experimental Psychology: General*, 110, 169-192.

White, S. A. (1993). The effect of gender and age on causal attribution in softball players. *International Journal of Sport Psychology*, 24, 49-58.

Wilson, R. A., & Keil, F. C. (2000). The shadows and shallows of explanation. In F. C. Keil, & R. A. Wilson (Eds.), *Explanation and cognition* (pp. 87-114). Cambridge, MA: MIT Press.