

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

The relationship between blocking and inference in causal learning.

Permalink

<https://escholarship.org/uc/item/4wj7h7zr>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 35(35)

ISSN

1069-7977

Authors

Livesey, Evan
Lee, Jessica
Shone, Lauren

Publication Date

2013

Peer reviewed

The relationship between blocking and inference in causal learning.

Evan J. Livesey (evan.livesey@sydney.edu.au)

Jessica C. Lee (jlee4128@uni.sydney.edu.au)

Lauren T. Shone (lsho0771@uni.sydney.edu.au)

School of Psychology, University of Sydney

Sydney, NSW 2006 Australia

Abstract

The blocking effect in causal learning, once taken as a hallmark of associative learning, has recently been explained in terms of an explicit deductive reasoning process. Yet when the conditions necessary for deduction are removed, a small blocking effect is often still present. We examined the relationship between blocking and participants' performance on analytical thinking and probabilistic reasoning measures. Inferential processes predict blocking or an absence of blocking in this situation, depending on the observer's consideration of conditional probabilities. Although Bayesian inference predicts blocking, most individuals are not inclined to use this form of probabilistic reasoning explicitly, an observation we confirmed using a logical problem with similar properties to the relationships present in the blocking effect. Furthermore, participants who showed the greatest capacity for analytical reflection were *less* likely to show a blocking effect, suggesting that blocking in causal learning is the product of an intuitive and unreflective thought process.

Keywords: Blocking; causal learning; inferential reasoning; associative learning; Bayesian inference.

Introduction

Many theories of causal learning assume that when individuals make judgments about the relationship between putative causes and their effects, some form of inferential reasoning is involved. However, theories differ substantially in how they place inferential reasoning amongst other contributing mechanisms. Some authors have argued that all causal judgments are necessarily the product of explicit inferential processes based on consciously mediated propositions about the relationships between events (Mitchell, De Houwer and Lovibond, 2009). Others assume that in making causal judgments about a cue, relatively automatic memory retrieval mechanisms based on associative learning play a much greater role, bringing to mind the events that were previously paired with that cue. According to this account, inferential thoughts of an analytical nature – for instance based on formal logic and reasoning – play a smaller role, in some cases perhaps only when strongly encouraged.

Blocking in causal learning

The blocking effect has become an important test bed for these arguments. In a typical blocking experiment, one cue (A) is presented and is reliably followed by a particular outcome. In a second stage, A is presented with another cue (B) and this compound of two cues is followed by the

outcome. B is never presented by itself and its relationship with the outcome is thus ambiguous. When asked to give a rating of the extent to which each of a number of cues causes the outcome, participants often given a lower rating for B than for control cues (C and D) that were also presented in compound and followed by the outcome but were never presented on their own.

The cues and outcomes are often presented within a hypothetical scenario. For instance, in the allergist task, the participant assumes the role of a doctor trying to determine the cause of a patient's allergic reactions. The participant might observe that when the patient eats Fish they suffer from a reaction (A+), and later when the patient eats Fish and Rice (AB+), they suffer from the same allergic reaction. The patient might also suffer from an allergic reaction after eating Mushrooms and Pasta (CD+), but does not suffer a reaction after eating various other foods (e.g. E-). After learning to predict what will happen after certain meals, through a process of trial and error, the participant must then make an explicit judgment about the extent to which a food or foods cause the allergic reaction, or the likelihood that a reaction will occur given that certain foods have been consumed.

The blocking effect is well documented in causal learning experiments using the allergist task and other similar scenarios. Its presence was originally taken as evidence that a similar associative learning process was responsible for causal learning and conditioning in humans and other animals because blocking in classical conditioning is widely replicated and well explained by associative learning theories (Dickinson, Shanks, and Evenden, 1984). Several other prominent theoretical approaches to causal reasoning also provide explanations of blocking (e.g. Cheng, 1997; Griffiths, Sobel, Tenenbaum, Gopnik, 2011; Waldmann, 2000). Whether based on associations or statistical computation, many theories of causal learning share an assumption that causal judgments partly reflect an implicit sensitivity to the contingencies between observed events. This sensitivity allows the observer to make judgments about causation with little deliberate mental effort, even when the causal relationships between cues and outcomes are ambiguous and must be inferred indirectly, as in the case of blocking (e.g. see Sternberg and McClelland, 2011).

Blocking and inferential reasoning

Recently, several authors have argued for an explanation of blocking that relies only on inferential reasoning based

upon a relatively simple set of propositions (De Houwer, Beckers, & Glautier, 2002; Lovibond, Been, Mitchell, Bouton, & Frohardt, 2003). Proponents of this account point out that there are circumstances under which the observer can logically deduce that the blocked cue (B) is not a cause of the outcome. For instance, this position is reached if one assumes that the effects of the patient's allergies are additive and that a more severe reaction could be observed if it were present. Holding these assumptions, if one does not observe an increase in the severity of the outcome when B is eaten at the same time as the allergenic food A, then one can deduce that B does not contribute to the allergic reaction. For example, if eating Fish causes an allergic reaction of severity 5 (on a fictitious allergy scale with a maximum of 10) and eating Fish and Rice also causes an allergic reaction of severity 5, then Rice has not made the reaction worse and thus probably isn't a cause of the reaction itself. Consistent with this inferential reasoning hypothesis, Lovibond et al. (2003; see also De Houwer et al., 2002; Livesey & Boakes, 2004) observed that pretraining and explicit instructions that encourage this outcome additivity assumption enhance the blocking effect.

Lovibond et al. (2003) also argued that if the observer assumes that the effects of the causal cues do not add to create a larger effect then this deduction is no longer valid and therefore there should be no blocking observed. This "nonadditive" assumption is encouraged by explicitly showing that the addition of two causes does not result in a stronger outcome than one cause on its own. According to this argument, participants with an assumption that the outcome is nonadditive should identify that they cannot be certain of the causal status of B, any more than the control cues C and D, and thus give each of these cues an equivalent causal rating that reflects that uncertainty.

In practice, a statistically robust blocking effect is often observed even after explicit nonadditive pretraining, albeit one that is numerically smaller than after additive pretraining (e.g. Lovibond et al., 2003; Mitchell, Lovibond, Minard, & Lavis, 2006). The presence of this persistent blocking effect has been viewed by some as a problem for the inferential reasoning account of causal learning because blocking after nonadditive pretraining is not the result that a participant would generate when applying inferential reasoning in a rational way (Lovibond et al., 2003).

Yet it is worth noting that, at least from the perspective of classical probability theory, this blocking effect is entirely rational. For both the blocking and control cases, the problem involves determining the probability of the hypothesis that a certain cue, X, is a reliable cause of the outcome, $p(X+)$. Relevant information is gained from observing that X in compound with another cue *does* cause the outcome (XY+). Thus the problem becomes one of calculating the conditional probability that X is a cause of the outcome given the observation that the compound XY causes the outcome, $p(X+ | XY+)$. We can use Bayes' theorem to calculate this conditional probability as follows:

$$p(X+ | XY+) = \frac{p(XY+ | X+) \times p(X+)}{p(XY+)}$$

where

$$p(XY+) = p(X+) + p(Y+) - p(X+)p(Y+)$$

In the case of the blocked cue, B, we can assume that participants are already certain that A causes the outcome the first time they experience AB+ trials, i.e. $p(A+) = 1$. In the case of the control cue, D, there is equal uncertainty about it and cue C, and thus $p(C+) = p(D+)$. In the absence of any further information, these unconditional probabilities, as well as $p(B+)$, are assumed to be equal to the base rate (the probability that the outcome will occur on any given trial or for any given cue). If we assume that, when a non-additive outcome follows a compound of two cues, the outcome is independently caused by at least one of the cues, then $p(XY+ | X+) = 1$. This is a reasonable assumption unless it is explicitly shown to be false, as in the case of patterning discriminations (Harris & Livesey, 2008; Livesey, Thorwart, & Harris, 2011). The predicted blocking effect derived from these assumptions is a function of the base rate probability, as shown in Figure 1. As the base rate approaches zero, $p(B+)$ approaches zero and $p(D+)$ approaches 0.5. As the base rate approaches 1, $p(B+)$ and $p(D+)$ both approach 1. Importantly, for every base rate between 0 and 1, $p(B+)$ is less than $p(D+)$. Most causal learning experiments (including this study) present equal numbers of outcome and no outcome trial types, meaning that the base rate is around 0.5. This means that a modest blocking effect is predicted, as can be seen in Figure 1.

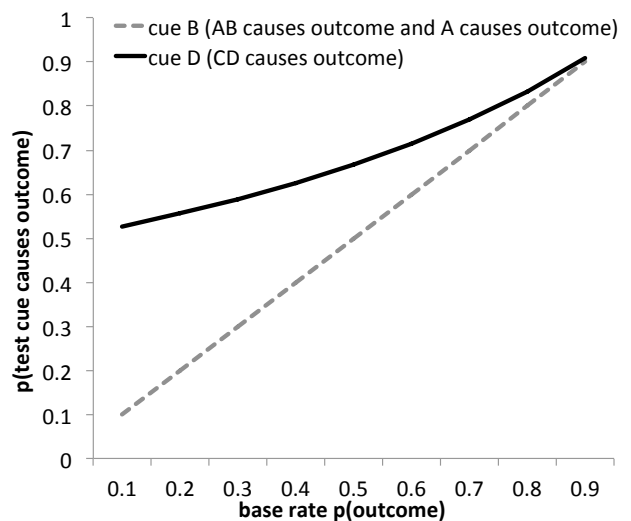


Figure 1. Probability that a test cue (B or D) causes the outcome as a function of the baseline probability that any given cue causes the outcome. Values were calculated by applying Bayes' theorem to the propositions that can be derived from a typical blocking design involving the "blocked" cue B and the control cue D (see parentheses).

The solution can also be derived without the previous equations using a series of simple inferential steps. For ease

of illustration, let us also assume that the probability of any cue shown in the experiment causing the outcome is 0.5. Given these assumptions, for any given compound of two cues A and B, there are four equally likely possibilities; i) A and B are both causal, ii) A only is causal, iii) B only is causal, or iv) neither A nor B is causal. In the case of the blocking cue, we know that A leads to the outcome, which allows us to rule out two of these possibilities (iii and iv), leaving possibility (i) in which B is causal, and possibility (ii) in which B is not causal. Thus the probability the B causes the outcome is 0.5. In the case of the control cues, we observe only that the compound causes the outcome, which allows us to rule out only possibilities (iv) that neither cue causes the outcome. The remaining three possibilities are still equally likely, and D causes the outcome in two of these three possibilities. Thus the probability that D causes the outcome is 0.67 (likewise for C).

Inferences and probabilistic reasoning

Although it may seem surprising to some that blocking under these circumstances is completely logical, the temptation to conclude that blocking is the result of an explicit rational inference based on classical probability theory needs to be tempered by an equally striking observation. In a host of similar situations, most participants are very unlikely to apply this form of reasoning. The rationale applied above to blocking shares formal qualities with other problems involving conditional probabilities, which most normal adults find extremely difficult (e.g. Bar-Hillel & Falk, 1982). A prominent example is the Monty-Hall dilemma (see Burns & Wieth, 2004), in which participants are so resistant to the solution derived from conditional probabilities that the problem is often referred to as a cognitive illusion. Thus, even though the blocking effect under nonadditive assumptions could be described as being rational, one should question whether participants are capable and inclined to explicitly use the inferential process that is necessary to arrive at the judgment in a rational and logical fashion.

If participants do use explicit reasoning processes akin to Bayesian inference, and the nonadditive blocking effect is a consequence of this reasoning, then the participants who show the greatest inclination to engage critically in inferential reasoning will be the most likely to give ratings in line with the blocking effect. Alternatively, Lovibond et al. (2003) assume that the most prevalent rational inference will be one in which the blocked and control cues are treated as being equally ambiguous, and thus no difference in their causal ratings should be observed. If this assumption is correct then those participants who are most likely to engage in *that* rational inference will be the *least* likely to produce a blocking effect in their judgments of causality. This hypothesis also implies that the blocking effect that has previously been observed after nonadditive pretraining is the result of a non-rational process such as a failure to retrieve the outcome associated with the blocked cue (Mitchell et al., 2006).

The current study sought to assess exactly what types of reasoned inference participants were inclined to use in this situation and how the inferential skills of individual participants were related to the blocking effect.

Blocking and critical thinking

To test the relationship between inferential thinking and blocking, we coupled a typical blocking task with a test of cognitive reflection developed by Frederick (2005). The test presents three mathematical problems, each of which can be solved with minimal calculation. The problems were specifically designed to provoke an intuitive answer that is incorrect. Deriving the correct answer requires a modest amount of self-reflection and analytical thought in order to reject the first number that comes to mind and to then apply the inferences that are appropriate for the logic of the question at hand. Frederick's (2005) analysis of this cognitive reflection test (CRT) over multiple samples of young American adults revealed that a substantial proportion scored 0 out of 3 on the test, revealing a strong tendency to accept and report the intuitive foil answer for each question. CRT performance is associated with general cognitive ability (Frederick, 2005). However, some studies have shown that performance on the test is influenced by the conditions under which the information is presented; for instance when the questions are more difficult to read they are more likely to be answered correctly (Alter, Oppenheimer, Epley, & Eyre, 2007). This suggests that participants' propensity to engage in critical reflection of the questions fluctuates and can be manipulated. For this study, the CRT was administered immediately after participants had finished making the causal judgments and thus, we assumed would assess their engagement in critical reflection around the time when the key measures of blocking were taken.

Participants were also given an additional problem designed to have similar logical properties to the contingencies in the blocking effect, in particular the presence of relevant conditional probabilities. Participants were instructed to "Imagine you are playing a game where, on every turn, a player tosses two normal everyday coins – a 50-cent coin and a \$1 coin – in the air. The coins are not biased: they are equally likely to show heads or tails. If either of the coins lands heads up, the player wins the round." They were then given two scenarios and asked to provide a probability for each:

1) "It is your turn next and you toss the coins. The \$1 coin shows heads but the 50-cent coin falls out of sight. What is the probability that the 50-cent coin is showing heads?"

2) "Your turn to toss the coins comes around again. This time, when you toss the coins, both coins fall out of sight. The other players in the game say (honestly) that you have won but you cannot see the coins. What is the probability that the 50-cent coin is showing heads?"

The answer to the first of these questions is relatively straightforward. Because the \$1 coin lands heads, the fact that the participant has won has no bearing on the

probability that the 50-cent coin is showing heads. Thus the correct answer is 0.5. The answer to the second question is more difficult because the information indicating that the participant has won is important for the probability that either one of the coins has landed heads. The correct answer is 0.67 because two of the three equally probable circumstances that could lead to the participant winning involve the 50-cent coin landing heads. We anticipated that most participants would say that the probability in this instance was also 0.5. This result would be consistent with the logical inference that Lovibond et al. (2003) assume is most likely to occur in a blocking experiment with non-additive outcomes.

Of most importance in this experiment was the relationship between CRT performance and blocking, and specifically whether blocking was found to be larger or smaller in those individuals that showed greater capacity for cognitive reflection. The coin-toss problem was added to further assess how participants engaged in inferences about similar uncertain events. If, as expected, many participants conclude that the uncertain events in each part of the coin-toss problem are equally likely, then it shows a tendency to use the inferential reasoning described by Lovibond et al. (2003). On the other hand, if participants tend to give the correct answer then it suggests they are very capable of using conditional probabilities in this context and may do so to make explicit inferences in causal learning that would produce a blocking effect.

Table 1: Design of the current Experiment.

Pretrain	Train 1	Train 2	Test
	A+	AB+	B
W-		CD+	C, D
X+	E+		
Y+	G-	F+	E, F, EF
Z-	GH-	GH-	EM, FM
WZ-	IJ+	KL-	H
XY+	L-	L-	L

Note: Letters A-M and W-Z denote randomly allocated foods used as predictive cues. These cues were followed by either no allergic reaction (-) or an allergic reaction (+). Trials above the dotted line in Train 1, Train 2 and Test comprise the blocking contingencies.

Method

Participants. Forty-four introductory psychology students at the University of Sydney participated in the experiment in partial fulfillment of course requirements (32 female, mean age = 18.9 years).

Apparatus and Stimuli. Participants were tested in individual cubicles in a quiet laboratory. The causal learning experiment was programmed using the Psychophysics toolbox for Matlab and was presented using Apple Mac Mini computers attached to 17 inch displays. Experimental stimuli included images of a banana, apple, fish, lemon, cheese, milk, coffee, eggs, garlic, bread, pasta, peanuts,

avocado, meat, mushrooms, olive oil, strawberries, peas, and rice accompanied by written labels. The allocation of foods to cue (A, B, etc.) was randomized for each participant. The CRT and coin-toss problems were administered in paper and pencil format, with each test presented on a single side of A4 paper, printed clearly in 14 point Times New Roman font.

Procedure. Participants were asked to assume the role of a doctor whose task was to ascertain which foods were causing the allergic reactions of a fictitious patient, Mr X. Participants were given general instructions about the scenario and the procedure, as well as explicit instructions about the nonadditive nature of the outcome. The latter was reinforced by presenting a pretraining phase in which two cues (X and Y) had demonstrably nonadditive effects. Here trials with X, Y and the compound XY were presented, each with followed by an identical allergic reaction. The presentation of the reaction outcome was the same throughout the experiment and was always accompanied by a fictitious severity index showing the same level of severity for all allergic reactions.

For each of the Pretrain, Train 1 and Train 2 phases shown in Table 1, each of the trial types was presented 8 times in a randomized order. On each trial, either one or two foods were presented and participants predicted what outcome (“no allergic reaction” or “ALLERGIC REACTION”) occurred by clicking either option. When an outcome was selected the options disappeared and were replaced with feedback about the actual outcome.

In the Test phase, participants were presented with a cue (or cues) and asked to make several judgments. First they were asked to judge “What is the probability that this food (these foods) will cause Mr X to have an allergic reaction?” and were required to make a rating on a linear analogue scale ranging from 0 to 1 with 0.1 increments marked along the scale. They were also asked to rate “How confident are you that your first rating is correct?” and “How severe will the reaction most likely be?” on additional linear analogue scales. The order of presentation of trials within the test phase was randomized, with each trial type presented only once. The critical cues in this phase for assessing blocking were cues B, C and D.

On completion of the allergist task, participants were given the CRT and conditional probability coin-toss problem in paper and pencil form. Participants were told to take as much time as they needed to finish these questions. Two versions of the coin-toss problem were used (counterbalanced between participants), one with the “neither coin visible” question first, the other with the “\$1 coin visible” question first. Above the response line for each question, participants were reminded that “A. If EITHER of the coins shows heads, you win the round” and “B. You know that you have won this round.”

Results

Learning during the pretraining and training phases of the causal judgment task was generally very rapid. In the final

block of pretraining, phase 1 and phase 2 training, mean accuracy exceeded 0.95 for every cue-outcome contingency. All participants performed well above chance. Statistical analyses focused on the critical test data only. All analyses were performed with an alpha level of 0.05.

Test Ratings. Of greatest importance was the probability rating for cue B ($M = 0.52$) compared to the mean probability rating for C and D ($M = 0.65$). The difference between these ratings was statistically significant, $t(43) = 2.83$, $p = .007$, indicating a reliable blocking effect overall.

CRT scores. Performance on the CRT was generally poor. Participants made on average just 0.70 correct responses out of a maximum of 3. The vast majority (84.1%) of errors resulted from reporting of the intuitive foil answers to each item (for further details, see Frederick, 2005). Table 2 shows the number of participants who scored 0-3 on the CRT test, and the mean blocking score for participants with each score.

Table 2: Frequency of CRT scores and blocking score.

CRT score	N	Blocking	
/3	Participants	Mean	SEM
0	26	.207	.059
1	7	.019	.115
2	9	.005	.086
3	2	-.001	0
total	44	.13	.047

Note: “Blocking” refers to the difference in probability rating given for the control cues C/D and the target cue B.

Of greatest interest was whether the number correct was related to blocking (as indicated by the difference in probability ratings for B and C/D). The correlation between CRT score and blocking was negative and significant, $r = -0.304$, $p = .045$. As can be seen in Table 2, this was mainly due to a large blocking effect in those that scored 0 on the CRT, with little variance in blocking scores amongst participants with CRT scores of 1 to 3. Participants who scored 0 on the CRT showed significantly more blocking than those who scored more than 0, $t(42) = 2.27$, $p = .028$. This is illustrated in Figure 2.

Coin-toss problem. Of the 44 participants, 29 responded 0.5 for the answer to both questions. Just two participants gave the correct responses, answering 0.67 for the scenario where neither coin is visible and 0.5 for the scenario where one coin is visibly showing heads (both scored 2 out of 3 on the CRT task and both exhibited a blocking effect). The remaining 13 participants did not systematically assign a higher probability to the “neither coin visible” scenario ($M = .52$) than to the “one coin visible” scenario ($M = .54$). As can be seen in Figure 3, participants who gave the same response to both questions (i.e. 0.5/0.5) produced equivalent blocking scores to those that produced different answers to the coin-toss problem, $t(42) = 0.15$, $p = .88$.




Figure 2. Mean causal probability judgments for the blocked cue B and control cues (mean of C and D), as a function of CRT performance. Left: Mean ratings for participants who failed to correctly answer a single question on the Cognitive Reflection Test. Right: Mean ratings for participants who scored at least 1 on the CRT. Error bars show SEM of the difference between B and C/D ratings.




Figure 3. Mean causal probability judgments for the blocked cue B and control cues C and D, as a function of answers to the coin-toss problem. Left: Mean ratings for participants who answered 0.5 for both items. Right: Mean ratings for participants who gave other answers (including two who gave the correct answers). Error bars show SEM of the difference between B and C/D ratings.

Discussion

Overall, participants showed a modest but statistically reliable blocking effect. This observation is typical of many studies in causal learning, including several that involve non-additive pretraining to discourage participants from deducing that cue B is not causal (e.g. Lovibond et al., 2003; Mitchell et al., 2006).

More importantly, the size of the blocking effect was significantly related to participants’ CRT performance. In particular, participants who scored zero on this test showed a substantial blocking effect whereas those that answered at least one of the three questions correctly gave comparable judgments for cue B and the control cues. The participants that scored zero on the CRT demonstrated the weakest ability to reflect critically on the questions in order to reject the most obvious answer and derive the correct one. These results are consistent with Lovibond et al.’s (2003) assertion that participants who reason carefully about the cues in a

blocking task involving nonadditive outcomes will judge the blocked and control cues to be equally likely to cause the outcome rather than adopting a Bayesian inference that appropriately accounts for conditional probabilities and is actually best aligned with the blocking effect itself.

Thinking about cause and effect under uncertainty is a difficult task and people do not readily adopt the approach typified by classical probability theory. The final probability question we used in this study is an example – with formal qualities similar to the blocking contingencies – where only two participants out of 44 gave the correct answer. Most (29 out of 44) gave the same answer, $p = 0.5$, to both problems, suggesting that they assumed the status of the unseen coin was unaffected by information about the outcome (i.e. winning the round) in both of the examples. This logic is very similar to Lovibond et al.'s (2003) argument about reduced blocking with a nonadditive outcome. They argued that participants will conclude that no information is known about cue B and, likewise, no information is known about either of the cues C or D and, therefore, all three should be given the same rating. However, unlike the coin toss scenario, the conservative logic expressed in this inference was not as prevalent in the causal ratings for the cues B, C and D. Furthermore, participants who gave the 0.5/0.5 response to the coin-toss problem, and thus should not show blocking based on Lovibond et al.'s inference, were just as likely to show a blocking effect in their causal ratings as those who gave different answers to the coin-toss problem. These results suggest that, although the logic described by Lovibond et al. is prevalent in decisions involving uncertain causal relationships, the application of the inference is not necessarily consistent across different scenarios.

This result is correlational and should be interpreted cautiously. Blocking may arise from other forms of explicit inference, such as deductive reasoning, which is encouraged by additive outcome assumptions (Lovibond et al., 2003). Thus the key relationship observed in this study should only arise if the assumptions that participants bring into the experiment are tightly constrained to prevent deduction.

Conclusion. Although the blocking effect is arguably rational, even when assuming that the outcome is non-additive, it is nonetheless associated with an uncritical mode of causal judgment. Only the minority of participants displaying some critical analytical ability on the CRT gave equivalent ratings to the blocked and control cues, consistent with the type of inferential reasoning outlined by Lovibond et al. (2003). The results are consistent with an account of causal learning that assumes that judgments are based on both explicit inferences and some form of associative learning or other automatic psychological operation that approximates Bayesian inference.

Acknowledgments

This research was supported by a University of Sydney Bridging Support Grant to E.J.L.

References

- Alter, A. L., Oppenheimer, D. M., Epley, N., & Eyre, R. N. (2007). Overcoming Intuition: Metacognitive Difficulty Activates Analytic Reasoning. *Journal of Experimental Psychology: General*, 136, 569-576.
- Bar-Hillel, M., & Falk, R. (1982). Some teasers concerning conditional probabilities. *Cognition*, 11, 109-122.
- Burns, B.D. & Wieth, M. (2004). The Collider Principle in Causal Reasoning: Why the Monty Hall Dilemma is so Hard. *Journal of Experimental Psychology: General*, 133, 434-449.
- Cheng, P. W. (1997). From Covariation to Causation: A Causal Power Theory. *Psychological Review*, 104, 367-405.
- Dickinson, A., Shanks, D., & Evenden, J. (1984). Judgement of act-outcome contingency: The role of selective attribution. *Quarterly Journal of Experimental Psychology*, 36A, 29-50.
- De Houwer, J., Beckers, T., & Glautier, S. (2002). Outcome and cue properties modulate blocking. *Quarterly Journal of Experimental Psychology*, 55A, 965-985.
- Frederick, S. (2005). Cognitive Reflection and Decision Making. *Journal of Economic Perspectives*, 19, 25-42.
- Griffiths, T. L., Sobel, D. M., Tenenbaum, J. B., Gopnik, A. (2011). Bayes and Blickets: Effects of Knowledge on Causal Induction in Children and Adults. *Cognitive Science*, 35, 1407-1455.
- Harris, J.A. & Livesey, E.J. (2008). Comparing Patterning and Biconditional Discriminations in Humans. *Journal of Experimental Psychology: Animal Behavior Processes*, 34, 144-154.
- Livesey, E. J. & Boakes, R. A. (2004). Outcome additivity, elemental processing and blocking in human causality judgements. *Quarterly Journal of Experimental Psychology*, 57B, 361-379.
- Livesey, E. J., Thorwart, A., & Harris, J. A. (2011). Comparing positive and negative patterning in human learning. *Quarterly Journal of Experimental Psychology*, 64, 2316-2333.
- Lovibond P. F., Been S.-L., Mitchell C. J., Bouton M. E., Frohardt R. (2003). Forward and backward blocking of causal judgment is enhanced by additivity of effect magnitude. *Memory & Cognition* 31, 133-42.
- Mitchell C. J., De Houwer J., Lovibond P. F. (2009). The propositional nature of human associative learning. *Behavioral and Brain Science*, 32, 183-246.
- Mitchell C.J., Lovibond P.F., Minard E, Lavis Y. (2006). Forward blocking in human learning sometimes reflects the failure to encode a cue-outcome relationship. *Quarterly Journal of Experimental Psychology* 59, 830-844.
- Sternberg, D. A. (2012). Two Mechanisms of Human Contingency Learning. *Psychological Science*, 23, 59-68.
- Waldmann, M. R. (2000). Competition Among Causes But Not Effects in Predictive and Diagnostic Learning. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 26, 53-76.