

Nichols, William, & Danks, David. (2007). Decision making using learned causal structures. In D. S. McNamara & J. G. Trafton (Eds.), *Proceedings of the 29th annual meeting of the cognitive science society* (pp. 1343-1348). Austin, TX: Cognitive Science Society.

Decision Making Using Learned Causal Structures

William Nichols (wn@andrew.cmu.edu)

Department of Philosophy, Carnegie Mellon University, 135 Baker Hall
Pittsburgh, PA 15213 USA

David Danks (ddanks@cmu.edu)

Department of Philosophy, Carnegie Mellon University, 135 Baker Hall
Pittsburgh, PA 15213 USA; and
Institute for Human & Machine Cognition, 40 S. Alcaniz St.
Pensacola, FL 32502 USA

Abstract

Decision making and causal reasoning are clearly relevant for one another, but have often been studied in relative isolation. In this paper, we report the results of two experiments that investigated whether people can make appropriate decisions using causal beliefs learned from sequences of cases. We found that people behave close-to-optimally for various causal and payoff structures, even though they are relatively poor at providing verbal reports of causal strength.

Keywords: causal learning; decision making; Bayes nets; intervention; causal reasoning.

Introduction and Related Research

Causal knowledge and reasoning are clearly relevant to our decision making, as we take various actions because we think that they will cause desired effects (Sloman, 2005). At the same time, our decision making is relevant for our causal learning and reasoning, both because our decisions shape the data we observe and because we may adjust our causal learning in light of anticipated future decisions. In light of the close connection between causal reasoning and decision making, it seems entirely natural to aim for an integrated theory of the two. Such a theory has emerged in the artificial intelligence and machine learning literatures through the combination of Bayesian networks as a representation of causal knowledge, and causal decision theory. In contrast, examination of an integrated model along these lines has only recently been explored in cognitive psychology (Hagmayer & Sloman, 2005, 2006; Sloman & Hagmayer, in press).

At the coarsest level, decision theory recommends that decision makers choose the option that maximizes the subjective expected utility. That is, given actions A_1, \dots, A_n , do the A_i with largest $\sum P(O_j | A_i) \times U(O_j)$, where O_j ranges over the possible outcomes, and $U()$ is a utility function. For example, suppose I have the choice of Thai food or steak for dinner. I enjoy good steak the most, but 10% of the time it is overcooked and so quite awful. In contrast, the Thai food is always pleasant.

We can think about the problem as involving three possible outcomes corresponding to a very enjoyable (good steak), pleasant (Thai), or unpleasant (bad steak) dinner. The

expected value of eating Thai is $U(\text{pleasant})$, since that outcome is guaranteed; the expected value of steak is $0.9 \times U(\text{very enjoyable}) + 0.1 \times U(\text{unpleasant})$. Standard decision theory prescribes that I do the action with greater expected value, where that clearly depends on the precise utility function $U()$.

Much of the work in, for example, behavioral economics aims to determine features of the utility function. In contrast, the long-running philosophical debate between evidential and causal decision theorists centers on the proper method to calculate the probabilities (e.g., Glymour & Meek, 1994; Hurley, 1994; Joyce 2000; Seidenfeld 1984). Evidential decision theorists argue that the probabilities should be based on straightforward conditionalization; causal decision theory holds that the relevant probabilities are causal ones that depend on the results of the action, understood as an exogenous force on the causal system.

In recent years, Bayes¹ nets have emerged as a relatively standard representation of causal structures. A causal Bayes net is composed of two related components: a directed acyclic graph that encodes the qualitative causal structure, and a joint probability distribution over the variables in the network that encodes the quantitative strengths of the causal relationships. These components are connected by a Markov assumption and there is a rich technical literature on inference and search for Bayes nets; details about these are not required for our purposes. Causal Bayes nets require only minimal metaphysical assumptions about the nature of causation; no strong theory is presupposed.

Given some fully specified Bayes net, there is a precise method for predicting post-intervention probability distributions (Pearl, 2000; Spirtes, Glymour, & Scheines, 1993). We can compute the probability of any variable (or set) in the system conditional on an intervention on any other variable (or set). In the Bayes net framework, interventions are most commonly understood as exogenous manipulations of particular variables. These interventions—sometimes called ‘hard’ interventions—change the

¹ There is nothing inherently Bayesian about ‘Bayesian networks.’ The name arises from the original uses of Bayes nets in Bayesian updating, and not because of any necessary connection between the framework and Bayesianism more generally.

underlying causal structure by eliminating the influence of all normal causes upon a variable within the system.

Continuing the food example from before, the underlying causal structure is a simple one: *Food Choice* \rightarrow *Enjoyment*. Suppose, however, that I apply anesthetic to my tongue before dinner so that I cannot taste anything. In that case, food choice is no longer a cause of enjoyment; rather, it is entirely determined by the intervention. Graphically, we remove (“break”) the edge. Hard interventions—those that control the state of a variable—are the most commonly discussed interventions, but the Bayes net theory of interventions is also well-defined for weaker types of interventions, such as those that simply perturb some variable away from its current value (e.g., adjusting my enjoyment by having a particularly sour fruit before dinner).

Bayes nets thus provide a natural complement to causal decision theory, as they provide both a robust framework for modeling causal structures, and the methods required to compute the relevant post-action probabilities for causal decision theory. Various AI and machine learning models use Bayes nets (or related ‘influence diagrams’) and causal decision theory in exactly this way (e.g., Jensen, 1996, and references therein).

A psychological model that integrates causal decision theory and Bayes nets in the natural way has only recently emerged (Hagmayer & Sloman, 2005, 2006; Sloman & Hagmayer, in press). Sloman and Hagmayer’s theory models choices as hard interventions, and expected utilities are all computed conditional on those interventions.

To date, this psychological model has been tested almost entirely by experiments in which participants are explicitly told the causal structure. Sloman & Hagmayer (2005) found that people make different choices about intervening on *A* to obtain *T* if they are explicitly told that an *A*—*T* correlation is due to direct causation ($A \rightarrow T$), versus an unobserved common cause ($A \leftarrow B \rightarrow T$). In other words, people want to intervene on *A* when it is causal, but are comparatively indifferent when *A* is not actually a cause. Importantly, participants are simply told the causal structure; they do not have to do any learning besides text processing.

There are a number of recent studies arguing that people can learn causal structure—represented as a Bayes net—from observational data, and in particular, from sequences of cases (e.g., Gopnik, Glymour, Sobel, Schulz, Kushnir, & Danks, 2004; Griffiths & Tenenbaum, 2005; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003). There is also evidence that people use that causal knowledge to predict the outcome of interventions (e.g., Gopnik, *et al.*, 2004; Kushnir, Gopnik, Schulz, & Danks, 2003; Steyvers, *et al.*, 2003; Waldmann & Hagmayer, 2005), though this research has not been done in a standard decision-theoretic setting.

In this paper, we aim to test the beginnings of an integration of these two literatures by asking: are people capable of using the products of causal *learning* from sequences to make well-informed choices? If so, are people sensitive to perceived causal *strength* (and not just structure)

when making decisions? Experiment 1 aims to begin to answer these two questions.

The previous research has also focused on cases of equal intervention cost/outcome payoffs. It thus does not provide a strong test of causal decision theory, as expected utility maximization is not separated from payoff probability maximization. That is, “correct” choices might simply be due to maximizing the probability of some payoff, rather than taking the utilities of the payoffs into account. In Experiment 2, we use a causal structure and payoff system for which these two methods make differing predictions.

Finally, we wanted all participants to have a strong interest in the outcome of the decision making. We thus provided participants with cash payoffs depending on whether their intervention was successful in bringing about the desired outcome. While the amounts of money involved are small (\$1 to \$3), we believe that the use of outcome-based payoffs provides important incentive for participants.

Experiment 1

Experiment 1 had two conditions with different causal structures; all participants did both conditions in random order. In each condition, participants were shown cases with two potential causes of a specified effect variable. All participants learned causal structures through the sequential presentation of cases, where the sequence ensured that participants saw exactly the desired frequency distribution. In condition A, one potential cause was a distractor variable that was uncorrelated with the effect. In condition B, both potential causes were actual causes, but one was much stronger. Condition B should be more difficult, as both variables are actual causes, and so participants need to track more information to make the final decision.

Participants

48 Carnegie Mellon University students participated and were compensated \$5 for participation, plus \$0 to \$2, depending on the outcome of their choices.

Design and Materials

The experiment was done on computers in the Laboratory for Empirical Approaches to Philosophy at Carnegie Mellon University. The cover story placed participants in the role of plant biologists attempting to get certain flowers to bloom. Participants were first provided an introduction to the information they would be given, and then instructed that their goal was to learn what causes blooming so that they could subsequently intervene to produce blooming. During the learning phase, participants were (passively) shown a series of cases that they examined in a self-paced manner. In both conditions, the potential causes were (potential) fertilizers. Each had a distinct name; for simplicity, we refer to them below simply as ‘Fertilizer 1’ and ‘Fertilizer 2’.

In condition A, the underlying causal structure was: *Fertilizer 1* \rightarrow *Blooming* *Fertilizer 2* (i.e., Fertilizer 2 was not a cause). Participants saw 48 cases in this condition. The fertilizers were uncorrelated with each other, and the

unconditional frequency of each was 0.5. Table 1 shows the conditional frequency of blooming given the fertilizers.

Table 1: Distribution of blooming for condition A

Fertilizer 1	Fertilizer 2	$P(\text{Bloom})$
Present	Present	0.75
Present	Absent	0.75
Absent	Present	0
Absent	Absent	0

Blooming occurs only if Fertilizer 1 is present. Fertilizer 2 does not affect the probability of the bloom, and is simply a distractor. Participants were told that the Fertilizers were each applied before the bloom (if applied at all).

In condition B, the underlying causal structure was $\text{Fertilizer 1} \rightarrow \text{Blooming} \leftarrow \text{Fertilizer 2}$; both fertilizers are actual causes of blooming. Participants saw 40 cases in this condition. The fertilizers were again uncorrelated with each other and occurred with an unconditional frequency of 0.5; the conditional frequency of blooming is given in Table 2.

Table 2: Frequency distribution for condition B

Fertilizer 1	Fertilizer 2	$P(\text{Bloom})$
Present	Present	0.8
Present	Absent	0.6
Absent	Present	0.2
Absent	Absent	0

During the test phase of each condition, participants were asked “To get the [PLANT NAME] to bloom, what do you want to apply?” where the actual plant name was used and participants were shown both the pictures and names of the two fertilizers. After choosing a fertilizer but before being told the outcome of their choice, participants were asked to rate the causal power of each variable for blooming. Ratings were provided on a slider that ranged from +100 (the cause always produces the bloom) to -100 (the cause always prevents the bloom), with 0 indicating no relationship. The slider moved in increments of 5, and participants were required to move the slider to give a response (i.e., they could not simply “click through”). The outcome of the intervention was determined by a pseudo-random sample from the underlying probability distribution, conditional on the participant’s choice. If the flower bloomed, participants immediately received \$1. After completing one condition, participants moved on to the other condition.

Results and Discussion

There were no significant order effects, either for the ratings or the choices, and so we ignore condition order in these analyses. 45 of the 48 participants (93.75%) chose to intervene on Fertilizer 1 (i.e., the actual cause) in condition A. This pattern is significantly different from random choice ($p < .001$, binomial test). In condition B, 39 of 48 participants (81.25%) chose to intervene on Fertilizer 1 (i.e., the stronger

cause), which is significantly different from random ($p < .001$, binomial test). The difference in choice percentages between conditions is significant ($p = .041$, McNemar’s test). Participants almost universally act to maximize $P(\text{Bloom})$, and therefore expected utility.

Beyond simple choices, we were interested in whether participants were internally coherent: did they choose the fertilizer to which they subjectively assigned greater causal strength? 47 of the 48 participants (97.91%) gave coherent responses in condition A, and 43 (89.58%) were coherent in condition B. Both of these percentages are significantly different from random ($p < .001$ for both, binomial test). No participants were incoherent in both conditions. There is a trend towards greater coherency in A than B, but it is not a significant trend ($p = .22$, McNemar’s test). In any case, participants were clearly highly coherent in their choices.

Interestingly, participant performance at the rating task was comparatively worse. Mean causal strength ratings for both conditions are shown in Figure 1. In Condition A, the mean strength rating of Fertilizer 1 was 63, which is significantly lower than the true strength of 75 ($p = .01$, t-test). The mean strength rating of Fertilizer 2 was -23, which is significantly lower than the true strength of 0 ($p < .001$, t-test).

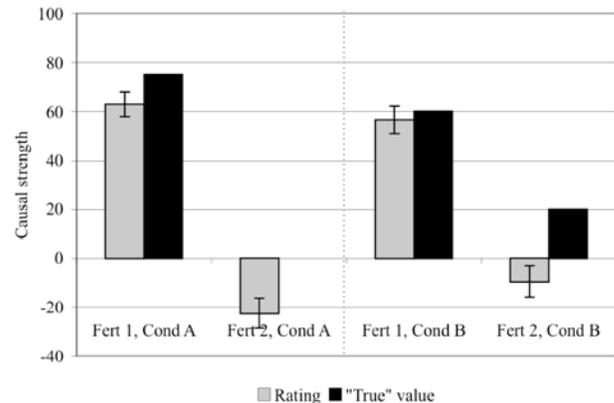


Figure 1: Mean strength ratings for both conditions

In Condition B, the value for the “true” causal strength is not obvious. There is significant debate in the causal learning literature about whether conditional ΔP (Shanks, 1995; Spellman, 1996) or causal power (Cheng, 1997) provides the most appropriate measure of causal strength. In condition B, the two methods yield different values; we focus on ΔP since the causal power value depends on the focal set. The mean reported causal strength of Fertilizer 1 was 57, and is not significantly different from the ΔP value of 60 ($p = 0.542$, t-test). The mean strength rating of Fertilizer 2 was -10, which is significantly lower than the ΔP value of 20 ($p = .002$, t-test).

Despite the largely inaccurate mean ratings, 45 out of 48 participants (93.75%) in condition A gave the correct rank order for the causes (i.e., rating for Fertilizer 1 is greater

than the rating for Fertilizer 2). 42 out of 48 participants (87.5%) gave the correct rank order for the causes in condition B. In both conditions, the numbers of participants with correct rank order are significantly better than random ($p < .001$ for both, binomial tests); participants were accurate about the strength ordering, though they did not get the exact numbers correct. The number with correct rank order in A was not significantly different from in B ($p = .37$, McNemar’s test). The one incoherent participant in condition A also gave an incorrect rank order, but only one individual (of six) who gave an incorrect rank order in condition B acted incoherently. There is no clear evidence of a correlation between incoherent behavior and incorrect rank ordering (in this very small sample).

Experiment 2

Experiment 2 had the same domain as Experiment 1, but used the structure: *Fertilizer* \rightarrow *Nitrogen* \rightarrow *Blooming*, where both *F* and *N* are potential targets of intervention. This causal structure is more difficult to learn than either of those used in Experiment 1 (e.g., Lagnado & Sloman, 2002), and so provides a stronger test. More importantly, however, this experiment aimed to distinguish between two plausible decision strategies: (i) maximize expected utility; and (ii) maximize the probability of payoff.

Participants were paid more if they caused blooming by intervention on *F*; for our probabilities, the larger payoff meant that intervention on *F* maximized expected utility (assuming a natural utility function). At the same time, an intervention on *N* necessarily had a higher probability of success. Thus, participants who seek to maximize expected value should intervene on *F*; those who seek to maximize the probability of a payoff should intervene on *N*. Of course, these different predictions are based on the true probabilities and payoffs; participant behavior will depend on their subjective beliefs.

Since prior research has not examined this type of causal or payoff structure, we used two conditions: a “Stepwise” condition in which participants were shown a sequence of cases (as in Experiment 1); and a “Story” condition in which they were explicitly told the causal story using exact statistics. The Story condition connects this experiment with the research of Sloman and Hagmayer, who have not previously considered a causal structure such as this one.

Participants

The same 48 Carnegie Mellon students participated and were compensated an additional \$0, \$1, or \$3, depending on the outcome of their choice. Sixteen participants were in the Story condition; 32 were in the Stepwise condition.

Design and Materials

The experiment was conducted in the same location, and the cover story was nearly identical. Participants were asked to learn what causes blooming so that they can intervene either on the fertilizer, or on the soil nitrogen, to produce blooming. In the Story condition, participants were told:

If you use nitrogen and the rose blooms, you will receive \$1. If you use fertilizer and the rose blooms, you will receive \$3. Fertilizer makes roses bloom by adding nitrogen to the soil. If you add fertilizer, you have 3 chances in 4 of triggering the nitrogen. If the soil has nitrogen in it, there are 4 chances out of 5 of making the rose bloom. The soil will only have in it what you put it in. There will be no naturally occurring nitrogen or fertilizer. Which would you rather use, nitrogen or fertilizer?

In the Stepwise condition, participants were shown a sequence of cases that captured the relevant frequency distribution. The fertilizer occurred with an unconditional probability of 0.5. The nitrogen never occurred without fertilizer; when fertilizer was present, the nitrogen occurred with probability 0.833. This conditional probability is slightly different from that in the Story condition, and was due to a programming error. Since we do not compare across conditions, we do not believe that the slight change makes a significant difference. Blooming never occurred without nitrogen; when nitrogen was present, blooming occurred 80% of the time. The resulting distribution of cases is shown in Table 3; for reasons of space, we omit cases that never occur. As in Experiment 1, participants passively observed the 48 cases.

Table 3: Case distribution for Stepwise condition

Fertilizer	Nitrogen	Blooming	# of cases
Yes	Yes	Yes	16
Yes	Yes	No	4
Yes	No	No	4
No	No	No	24

In both conditions, participants first gave a response. Before being told the outcome, they were asked to rate the causal strength of each variable on blooming, with the same prompt and rating slider as in Experiment 1. Participants were then told the result of the intervention, which was determined by a pseudo-random draw from the appropriate conditional distribution. If the flower bloomed and the participant used the fertilizer, the reward was \$3; if she used nitrogen, then the reward was only \$1. The objective expected value from using the nitrogen was \$0.80 in both conditions. The objective expected value for an intervention on the fertilizer was \$1.80 in the Story condition, and \$2.00 in the Stepwise condition. At the same time, $P(\text{Bloom} \mid \text{Intervene on } N) = 0.8 > 0.66 = P(\text{Bloom} \mid \text{Intervene on } F \text{ in Stepwise}) > 0.6 = P(\text{Bloom} \mid \text{Intervene on } F \text{ in Story})$.

Expected utility maximization and payoff probability maximization thus make different predictions in both conditions. Note that there is no correct answer for this experiment, as the “right” answer depends on what the participant wishes to maximize. Although we report participant responses below, we are more concerned with their strength ratings, and whether they acted to maximize subjective expected utility or payoff probability (or neither).

Results and Discussion

Story Condition. Five of the sixteen participants chose to intervene on the nitrogen; the other eleven chose the fertilizer. The mean causal strength rating for the fertilizer (60) is identical to the true value of 60 (see Figure 2 below). However, the distribution of responses was not unimodal: six participants (37.5%) gave a response within five units of 60, while seven participants (43.75%) gave a response within five units of 75, which is the strength of the fertilizer *on the nitrogen*. Of the three participants who gave other types of responses, two gave responses very near the middle of the slider, which suggests that they did the minimum required to move to the next question. The last gave a response of 50.

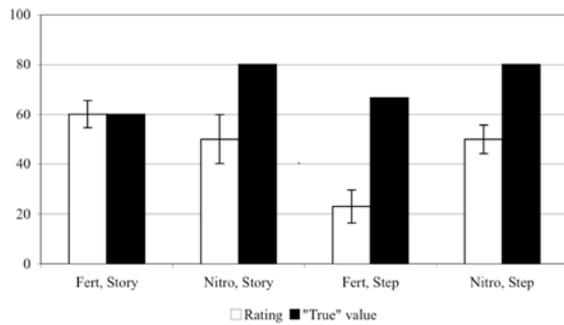


Figure 2: Mean strength ratings for both conditions

The causal strength ratings for the nitrogen were more surprising. The mean rating was 50, which is significantly less than the true strength of 80 ($p=.004$, t-test). Only eight participants estimated the strength of the nitrogen within 5 units of the true strength. Four participants gave causal values very near the middle of the slider. One gave a rating of -20, suggesting that the nitrogen *inhibited* blooming.

Only two participants gave correct answers for the strengths of both causal variables: one intervened on the nitrogen, the other on the fertilizer. The general failure of participants to give the correct—or even plausible—strength ratings is quite surprising, particularly since participants were paid based on the success of their intervention, and they had just finished reading a story that explicitly provided the causal strengths.

Having noted that participants may not have provided accurate ratings, we analyzed participant behavior to determine whether subjective expected utility maximization or payoff probability maximization better explains their behavior. The classification of participant behavior is shown in Table 4.

Eleven participants (68.75%) gave responses that maximize subjective expected utility, and eight (50%) sought to maximize subjective payoff probability. Neither of these is significantly different from chance (respectively $p=.21$, $p=.50$, binomial tests) Notice that some participants were able to maximize both expected utility and payoff probability given their subjective beliefs.

Table 4: Classification of Story condition behavior

	Expected utility maximizer	Not an expected utility maximizer
Payoff prob. maximizer	5	3
Not a payoff prob. maximizer	6	2

Stepwise Condition. Fifteen of the 32 participants in this condition chose to intervene on the nitrogen. The mean causal strength rating for the nitrogen was 50, which is significantly lower than the true value of 80 ($p<.001$, t-test). The mean causal strength rating for the fertilizer was 23, which is significantly lower than the correct value of 66 ($p<.001$, t-test), but the plurality of participants gave a response of 0. This value is the causal strength of the fertilizer *conditional* on the presence of nitrogen, implying that in this condition, many participants reported conditional strengths. Even if ratings near the middle of the slider are removed, both mean strength ratings were significantly less than the actual causal strengths (both $p<.001$, t-tests). Twenty participants (62.5%) gave the correct rank order for the causal strengths, which is not significantly different from chance ($p=.108$, binomial test).

The classification of participant behavior is shown in Table 5. 25 participants (78.13%) acted as if they were maximizing subjective expected utility, which is significantly more than chance ($p=.001$, binomial test). 21 participants (65.6%) acted as if they were maximizing payoff probability, which is not significantly different from chance ($p=.11$, binomial test). Notably, 17 of the 20 participants who gave the correct rank order for the causal strengths acted as expected utility maximizers, which is significantly different from chance ($p=.003$, binomial test).

Table 5: Classification of Stepwise condition behavior

	Expected utility maximizer	Not an expected utility maximizer
Payoff prob. maximizer	19	2
Not a payoff prob. maximizer	6	5

Many participants gave causal strengths that allowed them to both maximize utility as well as the probability of a payoff. While they chose the wrong causal strengths, these participants were at least coherent when they intervened.

Conclusion

These experiments are part of a larger project to try to tie together causal learning and reasoning, and causal decision theory. They provide further support that the causal learning and decision making elements of our cognitive systems are closely connected. In particular, people seem to be quite capable of learning simple causal structures from

experience, and then using those beliefs in sensible ways for decision making in novel situations.

Experiment 1 showed that people can use the results of causal learning from sequences to generate sensible decisions. Although people did not necessarily learn the true causal strengths, they largely used their (incorrect) subjective beliefs in a coherent manner. Not surprisingly, condition A was easier for participants than condition B. This finding suggests that people are not overly distracted by a potential cause that is uncorrelated with the effect, but they are affected by the presence of other actual causes and will rank secondary causes as a prohibitive cause. This result is, in some ways, not particularly surprising in light of empirical evidence that people sometimes focus more on causal *structure* than causal strength (e.g., Griffiths & Tenenbaum, 2005; Steyvers, *et al.*, 2003).

In the Stepwise condition of Experiment 2, a majority of participants acted to maximize both expected utility and payoff probability. That is, their subjective beliefs led to a choice problem that does not distinguish between these two principles. We are currently developing an experiment that more directly tests these two principles. The Story condition is more troubling, as participants did not seem to read the story carefully (as evidenced by their failure to rate the causes appropriately). Thus, our next experiment will use various incentives to improve participant comprehension, as well as explicit measures of story comprehension. We also intend to provide causal diagrams rather than text, thereby perhaps avoiding typical problems of story comprehension.

Ordering effects over the two experiments may also have played a role, as all participants were first exposed to experiment 1. Finally, we aim to understand better the individual differences that lead to variations in choice behavior. Our participants are relatively high-functioning, and so fine-grained measures of particular cognitive abilities may be required to separate out individual variation.

Acknowledgments

D. Danks was partially supported by grants from the Office of Naval Research and the National Institutes of Health, as well as intellectual support from the James S. McDonnell Foundation Causal Learning Collaborative. We thank four anonymous reviewers for helpful comments.

References

Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, *104*, 367-405.

Glymour, C., & Meek, C. (1994). Conditioning and intervening. *British Journal for the Philosophy of Science*, *45*, 1001-1021.

Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review*, *111*, 3-32.

Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive Psychology*, *51*, 334-384.

Hagmayer, Y., & Sloman, S. A. (2005). A causal model theory of choice. In B. G. Bara, L. Barsalou & M. Bucciarelli (Eds.), *Proceedings of the 27th annual conference of the cognitive science society* (pp. 881-886). Mahwah, NJ: Erlbaum.

Hagmayer, Y., & Sloman, S. A. (2006). Causal vs. evidential decision making in Newcomb's Paradox. In R. Sun & N. Miyake (Eds.), *Proceedings of the 28th annual conference of the cognitive science society*. Mahwah, NJ: Erlbaum.

Hurley, S. L. (1994). A new take from Nozick on Newcomb's Problem and Prisoner's Dilemma. *Analysis*, *54*.

Jensen, F. V. (1996). *An introduction to Bayesian networks*. London: UCL Press.

Joyce, J. (2000). Why we still need the logic of choice. *Philosophy of science*, *67*, S1-S13.

Kushnir, T., Gopnik, A., Schulz, L. E., & Danks, D. (2003). Inferring hidden causes. In R. Alterman & D. Kirsh (Eds.), *Proceedings of the 25th annual meeting of the cognitive science society* (pp. 699-703). Boston: Cognitive Science Society.

Lagnado, D. A., & Sloman, S. A. (2002). Learning causal structure. In W. D. Gray & C. D. Schunn (Eds.), *Proceedings of the 24th annual conference of the cognitive science society* (pp. 560-565). Mahwah, NJ: Erlbaum.

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

Seidenfeld, T. (1984). Comments of decision theory. In *Proceedings of the biennial meeting of the philosophy of science association*, Vol. 2.

Shanks, D. R. (1995). Is human learning rational? *The Quarterly Journal of Experimental Psychology*, *48A*, 257-279.

Sloman, S. (2005). *Causal models: How people think about the world and its alternatives*. Oxford: Oxford University Press.

Sloman, S. A., & Hagmayer, Y. (in press). The causal logic of choice. *Trends in Cognitive Sciences*.

Spellman, B. A. (1996). Conditionalizing causality. In D. R. Shanks, K. J. Holyoak & D. L. Medin (Eds.), *Causal learning: The psychology of learning and motivation*, vol. 34 (pp. 167-206). San Diego, CA: Academic Press.

Spirtes, P., Glymour, C., & Scheines, R. (1993). *Causation, prediction, and search*. Berlin: Springer-Verlag.

Steyvers, M., Tenenbaum, J. B., Wagenmakers, E.-J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, *27*, 453-489.

Waldmann, M. R., & Hagmayer, Y. (2005). Seeing versus doing: Two modes of accessing causal knowledge. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *31*, 216-227.