

Reasoning about Criminal Evidence: Revealing Probabilistic Reasoning behind Logical Conclusions

Michelle Cowley (M.Cowley@soton.ac.uk)
University of Southampton, UK

Abstract

There are two competing theoretical frameworks with which cognitive science examines how people reason. These frameworks are broadly categorized into logic and probability. This paper reports two applied experiments to test which framework explains better how people reason about evidence in criminal cases. Logical frameworks predict that people derive conclusions from the presented evidence to endorse an absolute value of certainty such as 'guilty' or 'not guilty' (e.g., Johnson-Laird, 1999). But probabilistic frameworks predict that people derive conclusions from the presented evidence in order that they may use knowledge of prior instances to endorse a conclusion of guilt which varies in certainty (e.g., Tenenbaum, Griffiths, & Kemp, 2006). Experiment 1 showed that reasoning about evidence of prior instances, such as disclosed prior convictions, affected participants' underlying ratings of guilt. Participants' guilt ratings increased in certainty according to the number of disclosed prior convictions. Experiment 2 showed that participants' reasoning about evidence of prior convictions and some forensic evidence tended to lead participants to endorse biased 'guilty' verdicts when rationally the evidence does not prove guilt. Both results are predicted by probabilistic frameworks. The paper considers the implications for logical and probabilistic frameworks of cognitive science for real world reasoning.

Keywords: Logic; Probability; Prior convictions; Criminal evidence; Legal reasoning.

Reasoning about Evidence in the Real World

Much of everyday life presents people with circumstances that require speculation. Speculations in each case may lead to the generation of explanations called hypotheses (e.g., Evans, Over, & Handley, 2002). But it is only by checking if hypotheses reflect the encountered evidence that leads to a true understanding (e.g., Klayman & Ha, 1987). For example, people can reason about evidence in relation to questions which they find curious for their own sake, such as trying to discover why their dog tries to eat bees when he knows they sting, or a young child may try to discover if a kiss will turn a frog into a prince (e.g., Cowley, 2006). People can also reason about evidence in relation to progressive questions, such as how to generate environmentally friendly energy sources (e.g., Kuhn, 1996). And sometimes there are sensitive contexts in the real world in which people reason about evidence out of duty to society, such as ascertaining guilt and responsibility when on a jury (e.g., Loftus; 1996; Pennington & Hastie, 1981). It is this latter example of reasoning about criminal evidence which is of concern in this paper. Consider that jury decision making studies have shown that the verdict thought

about by the majority at the first ballot was the jury's final verdict in at least 90% of trials (Kalven & Zeisel, 1966; cited in Devine et al., 2001). What sorts of reasoning processes could lead to such effects and how can the effects be theoretically explained? One key example of a factor that may affect the consistency of this deliberation process is eye witness evidence (e.g., Loftus, 1975). People tend to believe a defendant is guilty especially if the eye-witness is confident (e.g., Loftus, 1996). Despite a person's confidence, eye-witness memory has shown to be malleable and prone to several sorts of errors, including the misinformation effect in which subsequent false information can be added to the initial memory (Loftus, 1996). This applied research on eye-witness memory error has played a crucial role in the theoretical debate about the reconstructive nature of memory (see Loftus 2003 for a review). The purpose of this paper is similar in that it intends to investigate an applied legal context in which it is important to understand how people may be prone to errors in real world reasoning, and to theoretically explain those errors (e.g., Genn, Partington, & Wheeler, 2006).

One context in which it is becoming increasingly important to theoretically explain people's reasoning is criminal cases of child abuse (Sedlak *et al.*, 2006). This paper intends to examine whether the evidence of a defendant's prior convictions for similar offences should be disclosed in judicial proceedings (e.g., Petrosino & Petrosino, 1999), and whether such disclosure would bias people's reasoning with subsequent evidence. To illustrate the serious manner in which evidence of a prior conviction may affect people's reasoning about a given case, consider the following sensitive real world example:

Megan Kanka was killed by Jesse Timmendequas, a released sex offender, who unknown to her parents, lived nearby. An enormous public outcry led to emergency legislation in New Jersey requiring that every sex offender with a prior conviction be subject to a mandatory community notification system. Within several years the remaining 49 states brought in Megan's Law.

(from Pawson, 2006).

The effects of disclosing information about prior conviction evidence on people's reasoning is evident by how quickly Megan's Law was endorsed by the 50 states (for a review see Pawson, 2006). However, the US sex offender registration and community notification programs remain emotionally charged and contested areas of public policy (e.g., Matson & Lieb, 1996). The debate rests on whether policies such as Megan's Law can serve the balance

of protecting potential victims, while not leading to the conviction of potential innocent defendants who happen to have a prior conviction. Thus this paper has two specific goals: (i) to investigate whether this balance can be preserved in two experiments that examine the effects of reasoning with prior convictions as evidence (in criminal cases of child abuse adapted from a real world case), and (ii) to investigate whether these effects can inform the debate about whether logical or probabilistic frameworks of cognitive science are more or less appropriate for explaining human cognition.

Applying Logical and Probabilistic Frameworks to Real World Reasoning

To predict how prior conviction evidence may affect how people reason about criminal evidence, it is necessary to outline the main theoretical frameworks used to examine how people think. There are presently two competing frameworks with which cognitive science tends to examine human cognition; including how people reason. More generally these frameworks have tended to relate to the dissociation between logicist accounts that value the investigation of symbol manipulation at the computational level (e.g., Fodor & Pylyshyn, 1988); and connectionist accounts that value the investigation of probabilistic implementations of activation at the neural network level (e.g., McClelland & Rumelhart, 1986). In psychological categories of reasoning these frameworks may be *broadly categorized* as logical frameworks of reasoning (e.g., Johnson-Laird, 1999) and probabilistic frameworks of reasoning (e.g., Chater, Tenenbaum, & Yuille, 2006; Griffiths & Tenenbaum, in press; Evans, Over, & Handley, 2002).

Certainty and Uncertainty

Logical frameworks predict that people draw conclusions from the presented evidence to endorse an absolute value of certainty, such as 'true' or 'false' (e.g., Johnson-Laird, 1999), or in the case of reasoning about criminal evidence 'guilty' or 'not guilty'. The prediction that people will reason about presented evidence and draw conclusions of an absolute value of 'guilty' or 'not guilty' may be partially the legacy of the foundations of contemporary theories of reasoning.

For example, the mental models theory in the psychology of reasoning suggests that people's deduction involves the consideration of mental representations (mental models) of possibilities in the world considered to be true (e.g., Johnson-Laird, 1999). For example, the conditional statement 'If Justina is in Dublin, then James is in Belfast' (if p then q) is called Modus Ponens, and is derived from the logic of the propositional calculus (for a review see Johnson-Laird & Byrne, 2002). When people are told 'Justina is in Dublin', people may consider this statement to be true and mentally represent the possibility that 'Justina is in Dublin', and conclude that 'James is in Belfast'. What is

critical is that they are predicted to come to this conclusion with certainty.

What about uncertain contexts in which we cannot assume that a piece of information is absolutely true? There has been concern that reasoning theories built from logical frameworks, such as the propositional calculus, may not be able to fully account for reasoning in uncertain contexts, such as judicial settings, where it is not always possible to reach a conclusion that is certainly true or false (e.g., Oaksford & Chater, 2003). Some researchers believe that there is a probabilistic revolution taking place, especially when we consider the uncertainty with which people have traditionally been considered to reason with in real world decision making (e.g., Kahneman & Tversky, 1982).

For example, there has been the suggestion that people may in fact represent the subjective probability of a conditional statement (e.g., Evans et al., 2002). Consider the statement 'If David is left-handed, then he is guilty of the crime'. People may consider the possibility 'If David is left-handed' (if p), which can be hypothetically added to a person's prior knowledge so that they consider the possibility that 'he is guilty of the crime' (then q). They may actually consider the possibility of guilt given that David is left-handed ($P(q|p)$), and they can conclude with varying degrees of certainty that David is guilty. They may, for example, think about the possibility of other people being left-handed.

Probabilities in reasoning research more readily refer to degrees of belief, which are subjective to an individual, rather than the mathematician's frequentist interpretation of probability (e.g., Oaksford & Chater, 2003). For example, if a coin were tossed and landed beneath a table, a person's degree of belief that the coin has landed heads may be different to the mathematician's .5 frequentist interpretation. As the person's visual contact with the coin gets closer their degree of belief may come closer to 1, as the visual cues communicate that it is increasingly likely that the coin has come up heads (Chater et al., 2006). Analogous to this example is how a person's degree of belief may increase along a continuum of guilt in a criminal case as they encounter more evidence indicative of guilt. But whereas probabilistic frameworks predict that the degree of belief in guilt may fluctuate between 0 and 1, logical frameworks predict that belief in guilt can be assigned an absolute value of 'not guilty' or 'guilty' corresponding to 0 and 1 respectively, and not to any value in-between (e.g., Tenenbaum, Griffiths, & Kemp, 2006). Thus probabilistic frameworks predict that people draw conclusions from presented evidence by adding it to their prior knowledge. That is, when people consider evidence such as prior instances in the real world, they may not only endorse a conclusion of guilt which varies in certainty, but they may be biased in their consideration of subsequent evidence relevant to a case.

The paper examines if predictions derived from probabilistic frameworks of reasoning can explain how people reason with evidence of prior convictions in a

criminal case (adapted from a real life example). Experiment 1 investigates whether the consideration of prior instances, such as prior convictions, may lead to conclusions varying in certainty such as ‘more or less guilty’ rather than lead to conclusions of absolute certainty such as ‘guilty’ or ‘not guilty’. Experiment 2 investigates whether the consideration of prior convictions could bias the consideration of subsequent evidence relevant to a case. See Table 1 below for a comparison of predictions between logical and probabilistic frameworks relevant to the experiments that follow:

Table 1: Predictions made by Logical and Probabilistic Frameworks when people reason with prior convictions in real world criminal cases.

Logic	Probability
1. The endorsement of absolute values such as ‘Not guilty’	The endorsement of uncertain values such as ‘Cannot decide’
2. Underlying ratings of guilt should not increase towards ‘guilty’ as prior convictions are disclosed.	Underlying ratings of guilt should increase towards ‘guilty’ as prior convictions are disclosed.
3. The consideration of prior convictions should not bias reasoning about subsequent evidence.	The consideration of prior convictions should bias reasoning about subsequent evidence.

Prescriptive and Descriptive Considerations It is important to note the distinction between prescriptive and descriptive predictions of both frameworks in light of the following experiments. This distinction is relevant to the question of prior convictions, because logically a prior conviction could be viewed as a coincidence that could be ‘an accidental correspondence between independent facts’ (Horwich, 1982; p. 104; cited Griffiths & Tenenbaum, in press). On the one hand logic *prescribes* that evidence must be logically related to the facts of a present case in order for it to be considered rational to consider (e.g., Roberts & Zuckerman, 2004), and prior convictions present an opportunity for people to demonstrate their understanding of this logical distinction. On the other hand probability may be able to *describe* how people might interpret a prior conviction as a meaningful coincidence, and thus be unable to exclude this meaningful coincidence from their reasoning (see Griffiths & Tenenbaum, in press).

Experiment 1

The purpose of this experiment was to examine how people reason with prior convictions as evidence in criminal cases.

A case scenario was created from information of a real life case in which a child died, and the man responsible had similar prior convictions.

Method

Participants Fifty-four people took part. There were eight men and forty-six women. Their mean age was 20.5 years and their ages ranged from nineteen to thirty-three years. They took part on a voluntary basis and were recruited from the University of Southampton campus.

Materials A short scenario was adapted from a real life case in which an eighteen month old girl was killed by a man who had two similar prior convictions. The names in the scenario were changed to protect the identities of the victim and offender. The consent and debriefing underwent meticulous ethical reviewing with the sensitivity of the study being flagged in the initial person-to-person description and on the consent form on which options to withdraw were clearly stated. The debriefing contained hand-out information with contact details of the principal investigator and relevant organisations.

Participants were asked to consider the following scenario:

On January 2, 2006, David Baxter had been arrested. He had been accused of killing eighteen month old Joanna Connolly. Joanna’s skull had been fractured when she received a physical blow to the head. Joanna was the daughter of Susan Connolly, the woman David Baxter had been seeing.

Design and Procedure A 3 x 1 between-subjects design was used in this experiment. This experiment simply examined the effects of considering prior convictions as evidence. The three conditions included: a control condition in which no evidence of prior convictions was given to participants; an experimental condition in which evidence of one prior conviction was given to participants; and an experimental condition in which evidence of two prior convictions was given to participants.

Participants were randomly assigned to conditions and they first read the above scenario about Joanna Connolly. In the experimental conditions participants were given evidence of prior convictions in the following format (in the control condition they did not receive this sentence and in the experimental condition with two prior convictions they received an identical sentence with an earlier date):

David Baxter had previously served a three year sentence for being physically abusive towards an ex-girlfriend’s three year old in 2002.

Participants then had to answer two sorts of questions that have become a standard format in trying to capture logical or probabilistic reasoning in psychological research on reasoning (e.g., Johnson-Laird & Byrne, 2002). One of the questions asked for a categorical response to capture logical reasoning. Participants were given the options of ‘guilty’ and ‘not guilty’. Participants were also given a categorical

response to capture probabilistic reasoning by giving the option 'cannot decide'.

See this format below:

Please tick whether you think

David Baxter is 'guilty' ___

David Baxter is 'not guilty' ___

'You cannot decide' ___

The other question asked participants to record a rating response on a scale of 1 to 10. This scale was used to capture probabilistic reasoning if it was taking place. Participants were given values on a continuum between the absolute certainty of 'guilty' and the absolute certainty of 'not guilty'. The scale could also capture logical reasoning by making the options 'guilty' and 'not guilty' explicit:

On a scale of 1 to 10, circle the number that best reflects how guilty you think David Baxter is (0 represents Not Guilty; 10 represents Guilty):

Not Guilty
 Guilty 0 1 2 3 4 5 6 7 8 9 10

The experiment lasted six minutes on average.

Results

The results showed that participants chose the response 'cannot decide' (89%) significantly more often than the responses 'guilty' (9%) and 'not guilty' (2%) regardless of how many prior convictions they considered ($\chi^2 = 75.44(2)$, $p < .0005$). There are three possible implications of this result. First, participants demonstrate that they do not tend to endorse the legal logical distinction by concluding 'not guilty' when there is no conclusive evidence to convict. Yet their responses may indicate that they do have some understanding that similar prior convictions, even when two are present, cannot be held as conclusive evidence to convict. Second, the endorsement of an uncertain response such as 'cannot decide' rather than 'not guilty' or 'guilty' may indicate that people are reasoning probabilistically and not logically about criminal evidence.

But there is a third possibility; the result could also lead to the interpretation that people may not understand the materials. This possibility does not hold as the underlying ratings of guilt show that people's reasoning is affected by prior convictions, as Table 2 shows:

Table 2: Underlying ratings of guilt

	No prior conviction	One prior conviction	Two prior convictions
Mean underlying rank of guilt	3.35	5.86	6.33

The underlying mean rating of guilt was higher when a prior conviction was present (Mean rank = 5.86 and Mean rank =

6.33, for one and two prior convictions respectively) than when it was absent (Mean rank = 3.35, Kruskal-Wallis $\chi^2 = 16.162(2)$, $p < .0005$). The result that people's underlying ratings of guilt increase according to the number of disclosed prior convictions tends to suggest that people are reasoning probabilistically.

Experiment 2

The aim of this experiment is to extend reasoning about prior convictions to contexts in which people must also consider forensic evidence. This experiment examines whether forensic evidence is reasoned about in a biased way subsequent to the disclosure of a prior conviction.

Method

Participants Seventy-two people took part. There were twenty-four men and forty-eight women. Their mean age was 22.4 years and their ages ranged from eighteen to fifty-three. Participants were given the choice of accepting £4 for taking part, and they were recruited from the University of Southampton campus.

Materials The Joanna Connolly scenario and the information about David Baxter's prior conviction from Experiment 1 were used in this study. The consent and debriefing underwent meticulous ethical reviewing, with the sensitivity of the study being flagged in the initial person-to-person description and on the consent form on which options to withdraw were clearly stated. The debriefing contained hand-out information with contact details of the principal investigator and relevant organizations.

Design and Procedure A 2 x 3 between-subjects design was used in this experiment. The first between-subjects factor was presence of prior conviction (evidence of one prior conviction or no evidence of one prior conviction). The second between-subjects factor was forensic evidence related to handedness of the perpetrator (evidence that the perpetrator was right-handed, or evidence that the perpetrator was left-handed, or no evidence of handedness). The six conditions were: control (no evidence of prior convictions or forensic evidence related to handedness); evidence of right-handedness, evidence of left-handedness; one prior conviction and no evidence of handedness; one prior conviction and evidence of right-handedness; one prior conviction and evidence of left-handedness.

Handedness was used because it presents: (i) an objective measure to compare people's subjective underlying ratings of guilt (left-handedness occurs in approximately 10% of the population and right-handedness occurs in approximately 90% of the population), and (ii) a measure to facilitate an examination of whether people's reasoning about rarity is biased, subsequent to the disclosure of evidence of prior convictions. (Note also that twelve people were randomly assigned to the six conditions to aid potential comparisons with a jury).

In the forensic evidence conditions people were given a sentence related to handedness in one of the following formats:

Forensic evidence showed that the blow was delivered by a left-handed person. David Baxter is left-handed.

Or

Forensic evidence showed that the blow was delivered by a right-handed person. David Baxter is right-handed.

Participants answered the same two categorical and scale questions as in Experiment 1 in each of the six conditions.

Results

The results showed that people chose the conclusion ‘cannot decide’ (78%) significantly more often than the ‘guilty’ (16%) or ‘not guilty’ conclusions (6%, $\chi^2 = 65.333(2)$, $p < .0005$), as Table 3 shows below. This result suggests that people can reason towards uncertain conclusions, perhaps indicating a preference for probabilistic reasoning in this case. Table 3 shows the number of people who endorsed each sort of conclusion per condition (where: C = control; RH = right-handed evidence only; LH = left-handed evidence only; PC = prior conviction evidence only; PCRH = prior conviction and right-handed evidence; and PCLH = prior conviction and left-handed evidence). For this preliminary analysis it may be instructive to focus on the number of people out of twelve, perhaps akin to analyzing a jury’s choices, who chose each conclusion (i.e., Cowley, 2007, in preparation).

Table 3: The number of people who endorsed each conclusion per condition

	C	RH	LH	PC	PC RH	PC LH	Total
Cannot decide	10	11	11	8	9	7	56
Guilty	1	1	1	3	1	5	12
Not Guilty	1	0	0	1	2	0	4
<i>Total</i>	12	12	12	12	12	12	72

The results show also that people chose ‘cannot decide’ significantly more often than ‘guilty’ or ‘not guilty’ in every case except when they consider evidence of a prior conviction and left-handedness, as Table 3 shows. For example, when people consider evidence of right-handedness or left-handedness, in the absence of a prior conviction, they endorse ‘cannot decide’ more often than ‘guilty’ and ‘not guilty’ to the same degree (92% ‘cannot decide’ and 8% ‘guilty’ respectively), ($\chi^2 = 8.33(1)$, $p < .005$). When people consider evidence of one prior

conviction they endorse ‘cannot decide’ (67%) significantly more often than ‘guilty’ (25%) and ‘not guilty’ (8%, $\chi^2 = 6.5(2)$, $p < .05$). Likewise when people consider evidence of one prior conviction and right-handedness they endorse ‘cannot decide’ (75%) significantly more often than ‘guilty’ (8%) and not guilty (17%, $\chi^2 = 9.5(2)$, $p < .01$). These results suggest that people can reason towards uncertain conclusions, and that this pattern of results may be indicative of probabilistic reasoning. The conclusion ‘cannot decide’ can be considered rational in light of the fact that the evidence presented cannot prove guilt beyond doubt.

For example, when there is evidence of left-handedness there is a chance that David Baxter belongs to 10% of the population who are not guilty, and his prior conviction may be a mere coincidence. But when people consider evidence of one prior conviction and left-handedness, they begin to endorse ‘guilty’ (42%) almost as often as ‘cannot decide’ (58%), and they tend to ignore the ‘not guilty’ conclusion (0%, $\chi^2 = .333(1)$, $p > .05$), as the emboldened column in Table 3 shows. Why do people now begin to endorse ‘guilty’ verdicts indicative of certainty? Perhaps because they consider evidence showing a connection between the victim and defendant which is stronger than chance. Thus the suggestion that their reasoning can be biased by the consideration of a prior conviction is supported. Further support for this suggestion comes from an examination of people’s mean underlying ratings of guilt, which were higher when a prior conviction was present (Mean rank = 6.22) than not (Mean rank = 4.8, $\chi^2 = 12.28(5)$, $p < .05$).

The results show that people’s reasoning about forensic evidence can be biased, subsequent to the disclosure of prior conviction evidence.

Conclusions and Current Directions

This paper intended to investigate whether theoretical frameworks of logic and probability, broadly construed, could explain how people reason in the real world. The context of reasoning about prior convictions as evidence in criminal cases made it possible to examine if such real world applications could inform the debate about whether logical or probabilistic frameworks are better for understanding human cognition. Experiment 1 showed that people’s underlying ratings of guilt increased with the number of prior convictions disclosed, while the surface conclusions represented the uncertain ‘cannot decide’. Experiment 2 showed that people’s reasoning was biased when they reasoned about forensic evidence subsequent to the disclosure of a prior conviction. Probabilistic frameworks tend to predict that the consideration of prior instances influence the interpretation of subsequent evidence (e.g., Oaksford & Chater, 2003). The results tend to support the suggestion that people may reason probabilistically towards uncertain conclusions, and that they may also reason probabilistically towards logical conclusions of which they are certain.

The results have broader implications for probabilistic frameworks of reasoning by suggesting that (i) we need to

apply theories of reasoning to real world contexts of uncertainty more often (e.g., Oaksford & Chater, 2003), and (ii) we need to understand the circumstances under which mere coincidences turn into meaningful probabilistic evidence (see Griffiths & Tenenbaum, in press).

Presently content analyses are being carried out to understand the reasons why people endorse the conclusions that they do in these experiments. For example, do people mention anything about 'rarity' when they reason about left-handedness? Further experiments are being carried out to examine how people reason about similar and dissimilar prior convictions in group contexts.

The tentative suggestion that people reason probabilistically towards logical conclusions may have broader implications for cognitive science. Future work on how probabilistic reasoning with uncertainty leads to reasoning towards certainty may inform key debates of cognitive science. For example, the question of whether computational levels of explanation, concerned with how people reason with logical representations (e.g., Johnson-Laird, 1999) and symbols (e.g., Fodor & Pylyshyn, 1988), are more or less appropriate than connectionist levels of explanation, concerned with how people acquire knowledge at the probabilistic neural network level (e.g., McClelland & Rumelhart, 1986).

Acknowledgments

This research was funded by the School of Social Sciences' Small Grant Scheme for new researchers at the University of Southampton. Thanks to Laura Pennicott for testing the participants in the experiments.

References

Chater, N., Tenenbaum, J. B., & Yuille, A. (2006). Probabilistic models of cognition: where next? *Trends in Cognitive Sciences*, 10, 7, 292-293.

Cowley, M. (2007). Reasoning about criminal evidence: Revealing probabilistic reasoning behind logical conclusions. *In preparation*.

Cowley, M. (2006). *The role of falsification in hypothesis testing*. Unpublished PhD Thesis, Trinity College, Dublin.

Devine, D. J., Clayton, L. D., Dunford, B. B., Seying, R., & Pryce, J. (2001). Jury Decision Making: 45 years of empirical research on deliberating groups. *Psychology, Public Policy, and Law*, 7, 3, 622-727.

Evans, J. St. B. T., Over, D. E., & Handley, S. J. (2002). Suppositions, extensionality, and conditionals: A critique of Johnson-Laird and Byrne. *Psychological Review*, 112, 1040-1052.

Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28, 3-71.

Genn, H., Dame, Partington, M., & Wheeler, S. (2006). Law in the real world: Improving our understanding of how law works. *The Nuffield Inquiry on Empirical Legal Research*. London: Nuffield Foundation.

Griffiths, T. L., & Tenenbaum, J. B. From mere coincidences to meaningful discoveries. *Cognition*, in press.

Johnson-Laird, P. N. (1999). Deductive Reasoning. *Annual Review of Psychology*, 50 (1), 109-135.

Johnson-Laird, P. N., & Byrne, R. M. J. (2002). Conditionals: A theory of meaning, pragmatics and inference. *Psychological Review*, 109, 646-678.

Klayman, J., & Ha, Y-W. (1987). Confirmation, Disconfirmation, and Information in Hypothesis Testing. *Psychological Review*, 94, 2, 211-228.

Kuhn, T. S. (1993). *The structure of scientific revolutions*. (3rd Ed.). Chicago: Chicago University Press.

Loftus, E. F. (1975) 'Leading Questions and the Eyewitness Report'. *Cognitive Psychology*, 7, 560-572.

Loftus, E. F. (1996). *Eye-witness Testimony*. Cambridge, MA: Harvard University Press.

Loftus, E. F. (2003). Our changeable memories: Legal and practical implications. *Nature Reviews, Neuroscience*, 231-234.

Matson, S., & Lieb, R. (1996). Community notification in Washington State: 1996 survey of Law Enforcement. Olympia, WA: Washington State Institute for Public Policy.

McClelland, J. L., Rumelhart, D. E., & the PDP research group. (1986). *Parallel Distributed Processing: Explorations in the microstructure of cognition*. Volume I. Cambridge, MA: MIT Press.

Oaksford, M., & Chater, N. (2003). Conditional probability and the Cognitive Science of Conditional Reasoning. *Mind and Language*, 18, 4, 359-379.

Pawson, R. (2006). Evidence-based Policy: A Realist Perspective. London: Sage.

Pennington, N., & Hastie, R. (1981). Juror decision making models: The generalization gap. *Psychological Bulletin*, 89, 462-492.

Petrosino, A.J., & Petrosino, C. (1999). The public safety potential of Megan's Law in Massachusetts: an assessment from a sample of criminal sexual psychopaths. *Crime and Delinquency*, 45(1), 140-158.

Roberts, P., & Zuckerman, A. (2004). *Criminal Evidence*. Oxford: Oxford University Press.

Sedlak, A. J., Schultz, D., Wells, S. J., Lyons, P., Doueck, H. J., & Gragg, F. (2006). Child protection and justice systems processing of serious child abuse and neglect cases. *Child Abuse and Neglect*, 30, 657-677.

Tversky, A., & Kahneman, D. (1982). Judgment under uncertainty: Heuristics and biases. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases*, 201-208. New York: Cambridge University Press.

Tenenbaum, J. B., Griffiths, T. L., & Kemp, C. (2006). Theory based models of inductive learning and reasoning. *Trends in Cognitive Sciences*, 10, 7, 309-318.