

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Dual Processes and Training in Statistical Principles

Permalink

<https://escholarship.org/uc/item/07d9w3x4>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 27(27)

ISSN

1069-7977

Authors

Handley, Simon J.
Neilens, Helen L.
Newstead, Stephen

Publication Date

2005

Peer reviewed

Dual Processes and Training in Statistical Principles

Helen L. Neilens (hneilens@plymouth.ac.uk)

Department of Psychology, University of Plymouth, Drake Circus
Plymouth, PL4 8AA UK

Simon J. Handley (shandley@plymouth.ac.uk)

Department of Psychology, University of Plymouth, Drake Circus
Plymouth, PL4 8AA UK

Stephen E. Newstead (snewstead@plymouth.ac.uk)

Department of Psychology, University of Plymouth, Drake Circus
Plymouth, PL4 8AA UK

Abstract

A statistical training study is reported which demonstrates that analytic responding on everyday reasoning problems can be increased after instruction in statistical principles. Participants were given training on the Law of Large Numbers (Fong, Krantz and Nisbett, 1986). Bias was eliminated, but only on written justifications of their responses. Belief-based responding was still utilized when participants were asked for a quick indication of argument strength on a rating scale, thus demonstrating a dissociation between analytic and belief-based responding. Findings are discussed in terms of dual process theories of reasoning.

Keywords: Dual processes, critical thinking, belief-motivated reasoning, Law of Large Numbers, training.

Introduction

In recent years, many researchers have adopted a two-process model of reasoning to explain the heuristic and belief-based influences on reasoning and decision making (Evans & Over, 1996; Stanovich, 1999; Epstein, 1994). The hypothesis is that there are two distinct systems underlying reasoning and decision making. System 1 consists of both innate and domain-specific knowledge acquired through learning and System 2 is related to intelligence and analytic reasoning. It is believed that belief-based and pragmatically cued responses can be interpreted under System 1 processing and logical or analytic responses can be interpreted under System 2. The two systems compete with each other for control dependent on the context of the problem or argument.

It is argued that it is System 2 that is sensitive to instruction and permits abstract hypothetical thinking that cannot be achieved by System 1 (Evans, 2003). The aim of the experiment reported here was to investigate whether training in an intuitive rule system, the Law of Large Numbers, could impact on System 1 belief-based responses.

The Law of Large Numbers is proposed to be an intuitive version of a statistical rule which people use to solve inferential problems in everyday life. The LLN rule states, 'the certainty with which an inference about a population can be drawn increases as the size of a sample drawn from that population increases' (Klaczynski, Gordon & Fauth, 1997). In

order for an individual to apply this intuitive knowledge of the LLN principle it is important that these rules are cued by elements of a problem, such as the sampling process is made clear in order for the information to appear relevant, i.e. it is clear that there is a single trial or many repeated trials as with throwing a dice. Secondly, the role of chance in producing events is clear (Nisbett, Krantz, Jepson & Kunda, 1983).

More recently a series of studies have been reported which demonstrate that LLN reasoning can be evoked on everyday reasoning problems when the conclusion is incongruent with a person's beliefs (Klaczynski & Fauth, 1997; Klaczynski, Gordon & Fauth, 1997). Klaczynski et al. ascertained participants' occupational goals and then presented them with nine reasoning problems which consisted of conclusions that were either enhancing, threatening or neutral to their desired goal. They found that more sophisticated reasoning strategies, such as law of large numbers reasoning, were utilized when the conclusion was threatening to their occupational goal. These analytic strategies were employed to discredit the evidence presented. Klaczynski et al. concluded that strategies are changed to suit the goal of the individual at an intrinsic level and the biases participants are displaying are self-serving.

There is a great deal of evidence in previous research to suggest that belief-based influences impact on everyday reasoning (Holland, Holyoak, Nisbett, & Thagard, 1986; Nisbett & Ross, 1980; Lord, Ross, & Lepper, 1979). Nisbett and Ross describe research findings as far back as the 17th century illustrating people's tendency to cling to preconceived beliefs and theories in the face of new evidence that should discredit them. Lord et al. provided evidence to show that supportive evidence serves to strengthen a person's initial belief whereas opposing evidence using the same methods does not much affect belief.

One explanation proposed to account for such findings is the selective scrutiny theory of belief bias in syllogistic reasoning (Evans, Barston, & Pollard, 1983). Individuals focus on the conclusion and if it is consistent with their beliefs they will then accept it but if it is inconsistent with their beliefs then they will examine the logic of the problem to see whether it is valid or not. On critical reasoning

problems they will use whatever cognitive resources are available to discredit the evidence.

Klaczynski et al. proposed a depth of processing explanation for the moment-to-moment shifts in reasoning behavior. Cognitive-Experiential Self Theory (Epstein, 1994) is a dual-process theory which claims that reasoning involves using two parallel, independent systems; the rational and the experiential systems. It is suggested that reasoning is an interaction between the two and in the case of being presented with evidence that is contrary to belief, an individual's analytic system (rational system) is triggered which results in information being processed at a deeper level. This involves more cognitive expenditure but results in the activation of more sophisticated reasoning strategies such as the law of large numbers. When the evidence is consistent with the beliefs of the individual, then the information is processed at a shallow level by the experiential system. The evidence is assimilated to the already pre-existing beliefs and the conclusion accepted with little or no cognitive expenditure.

The aim of the experiment reported here was to investigate whether individuals can be taught or instructed to utilize the same strategies on all the problems, regardless of the direction of belief-laden content.

Fong, Krantz and Nisbett (1986) conducted a series of experiments to investigate law of large numbers reasoning. Participants were either trained in the LLN rule system which consisted of a description of the concept of sampling and the law of large numbers; given examples training which consisted of three problems in a given domain i.e. probabilistic, objective or subjective, followed by an explanation of how to solve the problem in LLN terms; or Full training which consisted of both types together.

They found that both the rule training and examples training improved statistical reasoning and enhanced the quality of the reasoning for problems across all three domains. However the rule training plus examples was found to have an additional effect. The studies did indeed provide evidence for the domain-independence of training. Participants who were given examples only training using problems from the objective domain were able to utilize the statistical principles across all three problem domains on testing. Fong et al. concluded that participants were able to map the LLN rules they had learnt onto a pre-existing set of abstract intuitive rules that they then used on problems in different domains to the one that they had been taught.

Support was provided by Fong and Nisbett (1991) who used more tightly defined domains of sports and ability testing and found use of statistical principles was still improved after a two-week delay. Some domain-specificity of training was observed over the delay however participants still applied the LLN heuristic more than participants who had received no training.

A key question that the experiment addresses is whether training in an inferential rule system transfer to problems involving belief-laden content? Can we get people to utilize their analytic reasoning strategies on the belief-consistent as well as belief-inconsistent evidence? This is the first time that

this training has been tested on belief-laden materials therefore we turn to the literature on instruction effects on belief bias in syllogistic reasoning to derive our predictions.

Belief bias effects are noted by the much higher acceptance rates of believable rather than unbelievable conclusions (Evans, Barston & Pollard, 1983). The belief bias effect is more marked on invalid problems. That is people will readily endorse it as valid due to its believability. This is in line with participants' performance on the everyday reasoning problems. When they are given a problem that is consistent with their beliefs they process the information at a very cursory level and accept the conclusion without searching for the flaws in the argument.

Evans, Newstead, Allen and Pollard (1994) found that bias was reduced by instructional manipulation which has since been interpreted under a dual process account. This account attributes System 2 and 1 processes to the logical and belief-based processes respectively that are influencing the task (Evans, 2003). Instruction was found to reduce bias on these tasks with the assumption that System 2 processing inhibited the automatic System 1 processes.

By bringing together the two major bodies of research, Fong et al.'s training studies and Klaczynski et al.'s individual differences research, it will be possible to investigate whether the same pattern of findings may be obtained on a different type of reasoning task. According to the belief bias literature, bias may be reduced after training. Training will impact on System 2 and a function of this will be to override or inhibit the belief-based responses cued by System 1. However, according to Klaczynski et al. the effects of belief and the level of LLN reasoning are independent and associated with different systems. Training will impact on the amount of statistical reasoning utilized but it will have no impact on the bias. Hence System 2 instruction may impact only on the level of LLN reasoning, but not belief.

Prior to the training experiment, a pilot study was conducted in order to develop the belief-laden conclusions required to strongly engage the belief-based and analytic reasoning strategies.

Pilot Study

Eighteen participants were administered a questionnaire involving thirty-four different professions and occupations. The aim was to determine typical, untypical and neutral character traits for the different occupations. The inventory consisted of 34 familiar occupations and professions. Following each occupation was a list of six traits or behaviors that may or may not be typical of people in that profession. For example, "Nurses are Caring, aggressive, thoughtful, intelligent, lazy, healthy". Participants were required to indicate on a scale of 1 to 5 (1 being very untypical, 5 being very typical) how typical they rated each trait/behavior for each particular occupation.

The six highest, six lowest and six neutral character traits (most typical, least typical and neutral related to professions) were identified by the mean scores. A repeated measures ANOVA (typical x untypical x neutral) was performed which illustrated that the mean ratings for each set of

professions/traits were significantly different from each other ($F(2, 34) = 327.53$, $MSE = 42.51$, $p < .001$). These items were then used to design the law of large numbers problems employed in the main experiment.

Experiment

Method

Design. A between subjects design was utilized involving two conditions, Training and Control. Participants under the training condition received Fong et al.'s full training.

Participants. 60 undergraduates from the University of Plymouth, 51 female and 9 males (mean age 22.1, st.dev. 5.84) took part in the experiment. Groups were randomly allocated to each condition resulting in 30 participants in each.

Materials. The instructions for both conditions were taken from Fong et al. (1986). The instructions from the control group read:

We are very interested in studying how people go about explaining and predicting events under conditions of very limited information about the events. It seems to us to be important to study how people explain and predict under these conditions because they occur very frequently in the real world. Indeed, we often have to make important decisions based on such explanations and predictions, either because there is too little time to get additional information or because it is simply unavailable.

On the pages that follow, there are a number of problems that we would like you to consider. As you will see, they represent a wide range of real-life situations. We would like you to think carefully about each problem, and then write down answers that are sensible to you.

Participants in the Training condition were presented with the first paragraph of the above prior to the training and the second paragraph was presented prior to the test materials, and ended in the sentence, "In many of the problems, you may find that the Law of Large Numbers is helpful".

Full Training. After the instructions described above, a paragraph introducing the law of large numbers was given:

"Experts who study human inference have found that principles of probability are helpful in explaining and predicting a great many events, especially under conditions of limited information. One such principle of probability that is particularly helpful is called the Law of Large Numbers".

Following this participants read a two-page description of the concept of sampling and the law of large numbers using examples of red and white beads in a jar, 30% red and 70% white. The beads in the jar represented the population, the proportion of red and white beads the population distribution and a selection of beads from the jar a sample. After the

concept of sampling was explained, the law of large numbers was presented:

"As the size of a random sample increases, the sample distribution is more likely to get closer and closer to the population distribution. In other words, the larger the sample, the better it is as an estimate of the population".

Participants were then given a demonstration of the law of large numbers, using a jar containing red and white beads with the same population distribution as that of the written description – 70% white, 30% red. The experimenter stated the main concepts again and then proceeded to draw samples from the jar, four of size 1, four of size 4 and four of size 25, to demonstrate that the average deviation of a sample from the population would decrease as the sample size increases as the law of large numbers predicts. The experimenter and the participants summarized each sample on a table, keeping track of the deviation between each sample and the population.

Following the demonstration the participants were given a set of three example problems with an answer following each one that provided an analysis of it in terms of the law of large numbers taken from Fong et al. Participants were asked to read each one and then consider it for a few minutes before turning the page to read the law of large numbers answer.

Law of Large Numbers Problems. Nine Law of Large Numbers were adapted from Klaczynski, Gordon and Fauth (1997). Hypothetical individuals presented arguments and evidence that were either neutral, consistent or inconsistent with participants' beliefs using typical or untypical personality traits. Belief-consistent problems involved arguments for a positive correlation between an occupation and a typical personality trait (e.g. firemen are brave), whereas belief-inconsistent problems involved arguments for a correlation between an occupation and an untypical personality trait (e.g. firemen are cowards). Of the nine problems employed, three were belief-consistent, three were belief-inconsistent and three were belief-neutral. See Table 2 for an argument resulting in a belief-inconsistent conclusion.

Table 2. Example of an argument involving a belief-inconsistent conclusion.

An editorial in a local newspaper recently criticized the occupation of being an aerobics instructor. The journalist's argument was:

I've got a friend who's an aerobic instructor and I wouldn't want anyone I know to copy her lifestyle. She is so unhealthy. She drinks and smokes and actually never takes any real exercise herself, she just tells others how to do so! I know her flat mates and they say she never stops eating as well, not healthy food either, fry ups and chocolate are normal. My conclusion? I don't think there's an healthier group of people than aerobics instructors!

Following each problem, participants indicated on two 9-point scales how convinced they were by the argument (1=not at all convinced; 9=very convinced) and how strong they thought the conclusion was based on the evidence used (1=very weak; 9=very strong). These were then referred to as the ‘persuasiveness’ and ‘evidence evaluation’ ratings respectively. Total scores on each rating scale were calculated separately for belief-consistent, belief-inconsistent and belief-neutral problems thus ranging from 3 to 27 for each problem type. Participants were then required to write explanations of why the conclusions were convincing/not convincing and strong/not strong. All test problems were presented in random order for each participant.

Coding of Explanations. A 3-point system developed by Fong et al. (1986) was employed. A score of ‘0’ was given if the response contained no indication of statistical reasoning. For example, ‘My friend is an aerobics instructor and she’s healthy’. A score of ‘1’ indicated that the participant referred to the law of large numbers, but vaguely. The participant implied that he or she was using statistical reasoning, but was not explicit about the statistical basis of his or her reasoning. For instance, ‘but that’s just one aerobic instructor’. If a participant scored ‘2’ for a response, it meant that the LLN principle was clearly applied in the explanation, for example, ‘the conclusion is a bad one as not all aerobics instructors are unhealthy, the journalist is only talking about one instructor. If he was to look at a larger sample of aerobics instructors, he may be able to report a more convincing argument’. For each problem type, scores were collapsed and added together to produce three total reasoning scores that ranged from 0-6.

The principal researcher and a second coder who was blind to condition and problem types independently coded all items for the 60 participants. Total agreement was achieved on 98% of the items. Agreement on the remaining 2% was achieved after discussion.

Procedure. Participants took part in the experiment in groups of 2 to 6. The control group was asked to read the instructions and proceed through the test material booklet. Participants in the Training group were required to attend two sessions. In session one they received full training and in session two, approximately one week later, they completed the test materials booklet. The training session took 40 minutes and the participants were allowed up to 1 hour 30 minutes to complete the test materials booklet.

Results

As may be seen in Figure 1 the level of statistical responding was higher for the participants that received training on the LLN principle on all three types of problems. It is also evident that law of large numbers reasoning is higher on responses involving belief-inconsistent conclusions than on belief-neutral or belief-consistent problems.

A 2 (Condition) x 3 (Problem type) mixed ANOVA was performed with condition as a between-subjects factor and problem type as a within-subject factor. Main effects of condition ($F(1, 58) = 39.54$, $MSE = 102.76$, $p < .001$) and

problem type ($F(2, 116) = 6.94$, $MSE = 6.62$, $p < .01$) illustrate that participants utilize more sophisticated reasoning techniques on all problems after training and LLN is utilized on problems that involve conclusions which are inconsistent with prior belief. An LSD follow-up test revealed reasoning scores on problems involving inconsistent information to be higher than scores on consistent or neutral problems (both $p < .01$).

The ANOVA also yielded a significant interaction between condition and problem type ($F(2, 116) = 3.96$, $MSE = 3.77$, $p < .05$). See Figure 1 for the graph of the interaction.

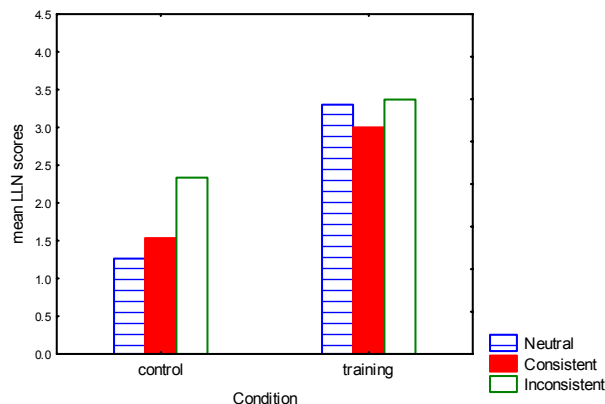


Figure 1. Interaction between Condition and Problem Type.

As illustrated in Figure 1, LLN reasoning is utilized more for inconsistent problems than either consistent or neutral problems in the control condition ($F(1, 58) = 8.9$, $MSE = 9.6$, $p < .01$ and ($F(1, 58) = 17.72$, $MSE = 17.07$, $p < .0001$), whereas with training use of LLN reasoning is improved greatly on all three problem types. There are no differences in sophistication of responses between any of the problem types (all $p > .1$). In other words, when presented with problems or arguments that are inconsistent with one’s prior beliefs, one will utilize a more sophisticated reasoning style to argue with, whereas if the information is consistent or neutral to one’s beliefs then it is much less likely to be evoked. However, with training in statistical principles, statistical reasoning is more likely to be used whatever the problems type.

Rating Scales

Table 4 displays the means for both the ‘Evidence Evaluation’ and ‘Persuasiveness’ rating scales for each problem type under both conditions. For the evidence evaluation scale, the conclusions were rated as stronger for problems that were neutral and consistent with a person’s beliefs. A 2 (Condition) x 3 (Problem type) mixed ANOVA with condition as a between subjects factor and problem type as within subjects found no effect of condition ($F(1, 58) = 1.02$, $MSE = 20.67$, $p > .05$). A main effect of problem type may be accounted for by the low strength of conclusion ratings by participants on problems that were belief-inconsistent ($p < .001$ when compared to consistent and neutral problems). In other words, the problems which contain

conclusions which are inconsistent with a person's beliefs and elicit more LLN reasoning are perceived as weaker arguments than the ones that contain conclusions which are consistent or neutral to a person's belief system.

Table 4. Means on the rating scales under both conditions.

Scale	Problem type	Control	Training
Evidence Evaluation	Neutral	12.53	14.53
	Inconsistent	8.77	9.40
Persuasiveness	Neutral	11.97	13.37
	Consistent	13.57	12.80
	Inconsistent	8.07	8.53

A 2 (Condition) x 3 (Problem type) mixed ANOVA was performed on the ratings on the persuasiveness scale. No effect of condition was found ($F(1, 58) = 0.45$, $MSE = 6.05$, $p > .1$). Participants rated the belief-inconsistent problems as less convincing than consistent or neutral problems ($F(2, 116) = 53.69$, $MSE = 431.82$, $p < .001$; a follow-up analysis found differences between inconsistent and consistent and neutral to be significant at $p < .001$ for both).

Discussion

The aim of this experiment was to investigate whether training on the LLN principle would transfer to problems involving belief-laden content. In summary, use of statistical principles was increased on all the problems designed to elicit belief-based responses after training, even after a one-week delay. Klaczynski et al.'s findings were also replicated in this experiment. Arguments involving belief-inconsistent conclusions elicited more sophisticated reasoning strategies than either belief-neutral or belief-consistent conclusions.

The effects of training on the everyday reasoning problems designed to elicit belief-based responses were surprising. It was predicted that the effects of bias may be reduced as in the syllogistic reasoning literature (Evans et al., 1994). In accord with the belief bias literature, Stanovich (personal communication, October 2004) proposed that training in rule-based strategies would attenuate but not eliminate biases. In dual process terms, instruction would increase System 2 function which would inhibit System 1 responses. In contrast Klaczynski suggested (personal communication, November 2004) that the training may not even transfer to problems involving belief-laden content at all. However the results reported in this experiment illustrate the elimination of bias after training on the concept of the laws of large numbers.

According to Klaczynski's account of belief effects on these problems, heuristic and analytic responding are independent of each other. Hence training on an explicit rule would not impact on System 1's intuitive system. Klaczynski and Gordon (1986) proposed training should increase the sophistication of reasoning responses on all problems; however the difference between responses on belief-consistent and belief-inconsistent arguments should remain

the same. Klaczynski, Gordon and Fauth (1997) argued that analytic responding was related to measures of intelligence and biases were related to thinking styles. In Klaczynski et al.'s view, higher ability participants would acquire the law of large numbers rule more rapidly but they would not be able to utilize it on problems designed to elicit belief-based responses. However in contrast to Klaczynski et al., training increased analytic responding overall and it also reduced the impact of beliefs when participants were asked to generate verbal evaluations of the strength and persuasiveness of the arguments.

Interestingly, the effects of beliefs are still present in the rating scales. The first rating scale asked for an evaluation of the strength of the conclusion based on the evidence presented, an evaluation that can be objectively made based upon the characteristics of the samples being discussed. After training the influence of beliefs on this scale are as strong as the influence of beliefs in the control group. This is startling, given that there is no influence of belief on the written justification for these responses. It is as if asking for a simple evaluation of an argument (such as the rating scale) does not engage explicit and effortful processing and is consequently subject to the influence of beliefs. Whereas asking people to generate a written evaluation activates the analytic System 2 processes that make available the LLN principles that have been taught.

The second rating scale asks about persuasiveness of the argument and it could be argued that it is quite rational to be less persuaded by a conclusion that is inconsistent with beliefs. One piece of evidence that is incongruent with beliefs that may often be based on many pieces of evidence, should not in a Bayesian sense impact drastically in changing or persuading us to change our view.

This is the first time that law of large numbers training has been tested on everyday reasoning problems involving belief manipulations. What is it about the training that facilitates this domain general reasoning? According to Fong et al. people are able to map the LLN rules they have learnt onto pre-existing abstract intuitive rules that they can then use on problems in different domains to the one they have been taught. The results here are consistent with that explanation and indeed add more leverage to it as the effects of training were still very strong after a one week delay between training and testing.

However, the above findings do not explain why LLN training eliminates belief bias. It is possible that people are utilising the rule on the belief motivated arguments as they would on any everyday reasoning problem. The process of using the rule may elicit cognitive decontextualisation on these tasks. Participants have been cued to use the rule which triggers System 2's analytic reasoning strategies. They then read through the problem and identify the small sample size as being a problem, regardless of the conclusion and the direction of belief. Hence when asked for a written evaluation of the evidence participants utilise the rule. However, when they are asked to rate the argument's strength and persuasiveness, their System 1 processes automatically cue

the belief-influenced response. A simple instruction to rate an argument's strength or persuasiveness does not engage System 2's analytic thinking processes.

An alternative explanation is that the reasoning utilised after training is superficial and participants in these experiments are simply transferring by analogy. One claim against Fong et al.'s original (1986) study was that the domains used were too narrow therefore the training was not transferring to very different types of problems. The examples used in the training in the experiment reported here were taken from Fong et al.'s training study and consisted of a completely different structure and content to the test materials.

What does seem apparent from the previous research and from these two experiments is that we do have some degree of conscious control over our reasoning processes i.e. there are effects of training. In the absence of training, the typical effects of belief were also found in this experiment. Arguments involving belief-inconsistent conclusions elicited more sophisticated reasoning strategies than either belief-neutral or belief-consistent conclusions.

By integrating the two bodies of research, the law of large numbers training and the individual differences in everyday reasoning studies, it has been possible to observe the interaction of analytic and belief-based processes. Manipulations of belief within everyday reasoning problems illustrate how individuals' strategies change dependent on whether System 1 or System 2 processes are engaged. The training effects reflect the interactive relationship between the two systems. Under this account, explicit instruction serves to trigger the rational System 2 processes which override the implicit System 1 processes leading to the elimination of biased responding and an increase in analytical reasoning.

The findings of this study are not fully consistent with the specific predictions derived from dual process accounts. However, the data do show a clear dissociation between written justifications and participant ratings. At the very least this demonstrates that people show the moment to moment switches between analytic and belief based judgements that are consistent with a dual systems account.

Acknowledgments

We would like to thank the Economic and Social Research Council for funding this research (Award No. R00429934481).

References

- Epstein, S. (1994). An integration of the cognitive and psychodynamic unconscious. *American Psychologist*, 49, 709-724.
- Evans, J. St. B. T. (2003). In Two Minds: Dual Process Accounts of Reasoning. *Trends in Cognitive Sciences*, 7, 454-459.
- Evans, J. St. B. T. & Over, D. E. (1996). *Rationality and Reasoning*. Hove, England: Psychology Press.
- Evans, J. St. B. T., Newstead, S. E., Allen, J. L., & Pollard, P. (1994). Debiasing by Instruction: The Case of Belief Bias. *European Journal of Cognitive Psychology*, 6, 263-285.
- Evans, J.St.B.T., Barston, J.L., & Pollard, P. (1983). On the conflict between logic and belief in syllogistic reasoning. *Memory and Cognition*, 11, 295-306.
- Fong, G. T. & Nisbett, R. E. (1991). Immediate and Delayed Transfer of Training Effects in Statistical Reasoning. *Journal of Experimental Psychology: General*, 120, 34-45.
- Fong, G. T., Krantz, D. H., & Nisbett, R. E. (1986). The Effects of Statistical Training on Thinking about Everyday Problems. *Cognitive Psychology*, 18, 253-292.
- Holland, J.H., Holyoak, K.J., Nisbett, R.E. & Thagard, P.R. (1986). *Induction: Processes on inference, learning and discovery*. Cambridge, Mass.: M.I.T. Press.
- Klaczynski, P. A. & Fauth, J. (1997). Developmental Differences in Memory-Based Intrusions and Self-Serving Statistical Reasoning Biases. *Merrill-Palmer Quarterly*, 43, 539-566.
- Klaczynski, P. A. & Gordon, D. H. (1996). Self-Serving Influences on Adolescents' Evaluations of Belief-Relevant Evidence. *Journal of Experimental Child Psychology*, 62, 317-339.
- Klaczynski, P. A., Gordon, D. H., & Fauth, J. (1997). Goal-Oriented Critical Reasoning and Individual Differences in Critical Reasoning Biases. *Journal of Educational Psychology*, 89, 470-485.
- Lord, C.G., Ross, L., and Lepper, M.R. (1979). Biased assimilation and attitude polarization: the effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology*, v37. pp. 2098-2109.
- Nisbett, R. E., Krantz, D. H., Jepson, C., & Kunda, Z. (1983). The Use of Statistical Heuristics in Everyday Inductive Reasoning. *Psychological Review*, 90, 339-363.
- Nisbett, R.E. & Ross, L. (1980). *Human inference: Strategies and shortcomings of social judgement*. Englewood Cliffs, N.J.: Prentice-Hall.
- Stanovich, K. E. (1999). *Who is Rational? Studies of Individual Differences in Reasoning*. Mahway, NJ: Lawrence Erlbaum Associates.