**Learning to Call Bullsh*t via Induction:
Categorization Training Improves Critical Thinking Performance**

Benjamin A. Motz^{1,2}, Emily R. Fyfe¹, Taylor P. Guba¹

¹ Department of Psychological and Brain Sciences and Cognitive Science Program

² eLearning Research and Practice Lab

Indiana University Bloomington

Author Note

Benjamin A. Motz <https://orcid.org/0000-0002-0379-2184>

Emily R. Fyfe <https://orcid.org/0000-0002-5420-360X>

Taylor P. Guba <https://orcid.org/0000-0002-6131-2417>

Study materials, data, and analysis scripts are publicly available at <https://osf.io/vcgzf/>.

This research was made possible by support from the Reboot Foundation, and portions of this research have been reported on the Reboot Foundation's blog, but the funder had no role in research design nor analysis and there are no conflicts of interest related to the conduct of this study. We appreciate the contributions of Alexandria Hartburg, Ulrich Boser, Michelle Sublette, Emily Cheshire Brown, Helen Lee Bouygues, and Steve Sloman. Author T.P.G. is now with University of Delaware.

Correspondence concerning this article should be addressed to Benjamin A. Motz, eLearning Research and Practice Lab, 1320 East 10th Street, Bloomington, IN, 47405. Email: bmotz@indiana.edu

Abstract

Participants performed a categorization training task, where each trial presented a claim based on an observation, and participants marked which fallacy or bias, if any, applied to the claim. In two studies, we measure the effect of this training task on critical thinking, measured using an open-ended critical thinking assessment. In Study 1, we pilot these materials in an online college course and observe credible improvements in critical thinking performance. In Study 2, we conduct a randomized controlled experiment and observe credible improvements in critical thinking relative to no training, and relative to comparable learning activities focused on conventional curricular content. We infer that the categorization training task facilitated inductive learning of patterns of biased and flawed reasoning, which improved participants' ability to identify such patterns in a delayed critical thinking assessment. Such categorization training shows promise as an effective and practical method for improving learners' resistance to online disinformation.

Keywords: critical thinking, inductive learning, causal reasoning, categorization, education, disinformation

General Audience Summary

Helping students learn to think critically is an important priority across education levels, but there is currently no clear consensus on how to teach critical thinking. Many proposed teaching strategies are very elaborate and involve extensive discussion and analysis, which makes it difficult for teachers to incorporate critical thinking instruction into their existing curriculum. In our current research, we explore whether a more focused approach might improve students' critical thinking performance. Specifically, we examine an inductive approach to critical thinking instruction. Inductive learning, commonly known as learning by example, happens when a person learns a general pattern from multiple examples of the pattern. In this case, we ask whether participants can learn to identify illogical or biased claims, a common measure of critical thinking ability, by presenting them with multiple examples of different kinds of fallacies and biases. We presented participants with sets of example scenarios in which an individual makes a claim based on an observation, and participants categorized which fallacy or bias, if any, the individual in the scenario was committing. In two studies, we find that this critical thinking categorization practice improves performance on a delayed open-ended critical thinking assessment. These exciting results show that critical thinking skills can be taught using a well-established, specific, psychologically-grounded method, and these results come at a critical juncture in our society. In a time of fake news, disinformation and a lack of critical thinking skills among the populace can have dire consequences. We believe our current results are of broad interest to educators and applied cognitive psychologists, and show promise for improving peoples' defenses against disinformation.

Learning to Call Bullsh*t via Induction:

Categorization Training Improves Critical Thinking Performance

The importance of critical thinking is a consensus view across education. Both postsecondary institutions (Association of American Colleges and Universities, 2005) and now the current generation of K-12 common standards place heavy emphasis, explicitly, on developing higher-order thinking and analytical skills (Porter et al., 2011). But underneath this consensus, there are mixed views on how to improve critical thinking, with similarly mixed outcomes (Abrami et al., 2015; Niu et al., 2013; Tiruneh et al., 2014). In this study, we examine a specific operationalization of critical thinking, and evaluate a corresponding training regimen, using an inductive learning paradigm, to improve critical thinking performance.

What is critical thinking? Although controversy persists (Johnson & Hamby, 2015), an increasingly common view defines critical thinking as the ability to detect and identify fallacies (such as logical errors or biases) in peoples' claims (Bergstrom & West, 2021; Lawson et al., 2015; Schmaltz et al., 2017; Schmaltz & Lilienfeld, 2014; Sternberg, 1999). Essentially, when presented with a statement about the world, someone who is skilled at critical thinking should be able to discriminate whether the claim is justified based on the information provided. Further, when claims are unjustified, they should be able to describe the specific issues that call into question their validity. This view of critical thinking lacks the scope and grandeur of broader definitions (which may include elements of creativity, self-efficacy, problem solving, and attitudinal dispositions; Ennis, 1987; Facione, 1990; Halpern, 2013), but has the converse benefits of discriminant validity and specificity. Moreover, this proposed view of critical thinking has a clear operationalization in a real-world task of increasing importance.

Individuals urgently need to be able to counteract growing volumes of disinformation (Lazer et al., 2018; Machete & Turpin, 2020). The prevalence of misleading online news and information presents real threats to democracy, public discourse, and public health (Lewandowsky et al., 2017). For example, online disinformation (that has been thoroughly

debunked) about side effects from vaccines has resulted in large-scale vaccine hesitancy, creating secondary public health crises caused not by illness, but by a contagion of fallacious claims (Horton, 2020; Poland & Spier, 2010). These claims, and our susceptibility to them, follow well-worn psychological paths, such as our human tendency to view correlational findings, or even just rare coincidences, as providing evidence of causal mechanisms. When exposed to these types of claims, one who is skilled at critical thinking should be able to recognize that mere associations do not warrant claims about one factor causing another, and to challenge the validity of claims.

Individuals also urgently need to detect and correct biased claims. Another example of our susceptibility is *confirmation bias*, the tendency to privilege evidence that supports one's views while ignoring evidence that contradicts them, a common feature of ideological extremism (Lilienfeld et al., 2009; Tavris & Aronson, 2020). In social media, the construction of polarized echo chambers is attributed to confirmation bias, as users promote posts and articles that only reflect one side of a debate (Garimella et al., 2018). Users who share one-sided posts do not believe themselves to be spreading false information, they simply neglect evidence to the contrary (Pennycook et al., 2019; Pennycook & Rand, 2019, 2021). Again, when exposed to such a scenario, one who is skilled at critical thinking should be able to recognize the single sidedness of the arguments, and to challenge the validity of claims.

Increasingly, many people look to our education system to provide students with the critical thinking skills necessary to mitigate their susceptibility to disinformation and bias. However, developing a learner's ability to think critically — to identify fallacies such as confirmation bias or a causal inference from rare coincidences — requires dedicated effort. Many students simply show no measurable improvement in general critical thinking ability during college (Arum & Roksa, 2011). Moreover, improving a student's understanding of domain knowledge (e.g., neuroscience) does not necessarily improve a student's ability to identify fallacies within that domain (e.g., neuromyths; Im et al., 2018). Incidentally, most

college-age learners maintain views that have no evidentiary basis whatsoever, such as the existence of people who can move objects with their minds (70% endorsement), or that certain objects such as rabbits' feet genuinely bring good luck (69% endorsement; Lobato et al., 2014), providing indirect but complementary evidence of low critical thinking ability (Lantian et al., 2021). If education systems are to improve students' critical thinking skills, a more focused educational strategy is necessary.

Reviews and meta-analyses of instructional strategies for improving critical thinking have examined a wide range of published approaches, with varying degrees of efficacy (Abrami et al., 2015; Niu et al., 2013; Tiruneh et al., 2014). Across these reviews, three general patterns emerge: effective instructional strategies (1) tend to be embedded within the discipline, (2) tend to expose students to authentic critical thinking problems in real contexts, and (3) tend to occur over extended periods of time. Situating or contextualizing instruction within a discipline, rather than presenting instruction in abstraction, is a well-known feature of effective educational programs in critical thinking (Willingham, 2008). And the importance of exposure to authentic problems over time mirrors recent findings that emphasize the importance of deliberate practice (Heijltjes et al., 2014; van Gelder, 2005) for improving critical thinking skills.

Improving Critical Thinking via Category Induction

These features of effective critical thinking instruction (contextualized materials, deliberate practice, and repeated exposure) bear striking resemblance to principles of inductive learning. Inductive learning, broadly defined as learning by example, involves extracting a general rule, pattern, or schema based on key structural features (and not the variable surface features) present across multiple examples. Like critical thinking instruction, inductive learning is improved by exposure to more concrete examples (Fong et al., 1986; Fyfe et al., 2014), by spacing exposure over time (Carvalho & Goldstone, 2015; Wahlheim et al., 2011), and by practice testing (Izawa, 1967; Tulving, 1967). And while inductive learning might commonly be associated with contrived laboratory-based experiments, the general framework has

constructively been applied to more naturalistic learning tasks, such as learning to discriminate artists' painting styles (Kang & Pashler, 2012; Kornell & Bjork, 2008), learning to medically diagnose skin disorders (Brooks et al., 1991), learning to sort natural rock specimens with complex category structures (Nosofsky et al., 2018), and learning to classify concepts in real education settings (Carvalho et al., 2016; Day et al., 2015; Dunlosky et al., 2013; Quilici & Mayer, 2002). For all these tasks, the inductive approach is to present an example of a problem, have the learner respond to the problem, provide feedback, then repeat.

But by contrast, learning activities currently proposed in the critical thinking literature are generally much more elaborate. For example, Renaud and Murray's (2008) critical thinking intervention involved essay questions that required multiple pages of written text response. Zohar and Tamir (1993) infused critical thinking scenarios into a biology curriculum, where activities extended for hours, sometimes covering multiple class periods with lively student debates. Self and Self (2017), who also propose exposing students to regular practical critical thinking scenarios, advocate that this exposure should methodically cover each of the ten steps of Kolb's experiential learning process (Kolb, 2015). While such training regimens are impressive in their detail and scope, they also create practical barriers for implementation, such as heavy incursion into existing curricular activities and steep grading demands. It is also unclear if these laborious efforts cause better critical thinking outcomes than a more basic inductive approach.

If we define critical thinking as the ability to detect and identify fallacies, an inductive approach should simply present examples of potentially fallacious claims, situated in a course's content domain and distributed across the semester, and invite students to identify which fallacy (if any, among a set of fallacies) are evident in the claim. These might be presented as multiple-choice problems, where students categorize claims according to basic fallacies, and receive prompt feedback on their responses.

This approach is strong because it relies on categorization, which is fundamental to human cognition (Harnad, 2017), and category learning builds structure in the organization of semantic knowledge (Goldstone et al., 2017). Showing real-world examples of a category improves one's ability to correctly identify new (unstudied) examples of the category, even when the category is very high-level, such as decision making principles (Rawson et al., 2015), intuitions about statistics (Fong et al., 1986), or the deep meaning of proverbs (Tullis & Goldstone, 2016). Along these lines, one might expect improvement in our ability to identify critical thinking fallacies when people are shown examples and tasked to categorize the fallacy present. In effect, learners are inducing mental models of the category structure of these fallacies (Cheng et al., 1986; Goldstone et al., 2013; Nosofsky, 1986; Smith et al., 1992). Research suggests that knowledge of these categories, in particular, has general utility: the specific ability to detect and reject fallacious claims has been shown to predict academic performance, resistance to misconceptions, and the ability to weigh and interpret evidence (Bensley et al., 2014; Bensley & Lilienfeld, 2017; Kowalski & Taylor, 2004; McCutcheon et al., 1992).

Preliminary evidence of these kinds of improvements has already been observed in a critical thinking study using a multiple choice format for induction. Hitchcock (2004) saw better critical thinking performance among students who performed computer-guided practice with multiple choice questions than students who engaged in small group discussion, but this study lacked random assignment. As Hitchcock concluded, "more work needs to be done," to build confidence that critical thinking gains can be achieved by this approach.

Researchers have recently begun to build this evidence base, using a tactic presently referred to as "prebunking" (Basol et al., 2021; Lewandowsky & van der Linden, 2021). The prebunking approach involves inductive exposure to small innocuous doses of misinformation, which progressively build an individual's ability to respond appropriately to misleading claims in the wild. Like the idea that one might become inoculated from persuasion (Papageorgis &

McGuire, 1961), prebunking has shown promise in building an individual's cognitive defenses against common forms of disinformation.

The Current Research

Here we explore whether an inductive training regimen, where learners categorize example claims according to different types of fallacies and biases, improves general critical thinking performance. Considering that critical thinking instruction should be situated within disciplinary content, we developed materials for implementation in a college Introductory Psychology course (although this general approach could easily be adapted to other courses). Introductory Psychology is one of the most popular college courses, taken by roughly 1.5 million students each year in the US (Gurung et al., 2016), with programs placing heavy emphasis on the importance of critical thinking (Homa et al., 2013), and where critical thinking outcomes similarly include the ability to recognize flaws in explanations (Halonen et al., 2013).

In the first study, we pilot a set of critical thinking practice problems in a live Introductory Psychology class. In the second study, we randomly assign online research participants to receive the same critical thinking practice problems, to receive non-critical-thinking Introductory Psychology practice questions, or to a no-training control condition. In both studies, we measure critical thinking performance pre- and post-training using the Psychology Critical Thinking Exam (PCTE, updated version; Lawson et al., 2015). The PCTE is an open-ended assessment where each item describes a fallacious claim drawn from some evidence, and examinees are asked to describe the problem (if one exists) with the conclusion. The open-ended nature of the PCTE provides a reasonably authentic measure of whether participants can accurately identify and describe these fallacies from memory, in the absence of any cues.

Study 1 (Classroom Pilot)

Study 1 was an exploratory pilot study, designed to test for possible benefits of inductive learning using multiple choice questions about critical thinking scenarios in a live education setting. All enrolled students are exposed to critical thinking categorization practice, and we measure improvements in critical thinking performance by comparing scores on pretest (at the start of the semester) and posttest (at the end of the semester). All materials are publicly available at <https://osf.io/vcgzf/>.

Method

Participants. Students enrolled in online sections of Introductory Psychology at Indiana University Bloomington during the Spring 2017 semester constituted a convenience sample for examining the effect of critical thinking training. Students who withdrew from the course are not included, resulting in 504 students enrolled. Aggregate institutional records from these sections indicate that 50.2% are female, with an average age of 19.9 years at the end of the semester. However, 122 students did not access either the pretest or posttest assessment (which were ungraded; any response earned a small amount of completion credit toward the course score), and another 22 students accessed neither of them, so these students are also excluded, leaving 360 students who all submitted a pretest and a posttest.

Materials. The Psychology Critical Thinking Exam (PCTE, updated version; Lawson et al., 2015) was used to assess critical thinking. Each item in the PCTE is a short description of a fictional scenario where someone makes a fallacious or biased inference based on some observation, and examinees are instructed to “State whether or not there is a problem with the person’s conclusions and explain the problem.” There are seven such problems in the full PCTE, which we label here as: Random Chance, Lack of Control, Correlation is not Causation, Overgeneralization, Experimenter Bias, Confirmation Bias, and Oversimplification. However, in the current study, we excluded items related to Oversimplification, because these were in direct conflict with the students’ textbook (Schacter et al., 2016), which explicitly promoted the

importance of parsimony, stating that the “simplest theory that explains all the evidence is the best one” (p. 41). Moreover, the original authors of the PCTE found group differences specifically on the Oversimplification items, and suggested that specialized training may be necessary for students to understand this principle in particular. So in the current study, we use a reduced form of the PCTE, including only the first six fallacies, assessed across 12 items, with two items per fallacy. See Table 1.

Table 1. Fallacies and Biases Included in Training and Tests

Fallacy	Description
Confirmation Bias	Inferences based only on positive evidence and ignoring disconfirming evidence
Correlation is not Causation	Inferring a causal relationship on the basis of correlational data
Experimenter Bias	Inferences based on questions that were biased, loaded, or leading
Lack of Control	Inferring that improvement was due to an experimental intervention without comparison to a control group
Overgeneralization	Inferring that findings generalize to a larger group based on a biased sample
Random Chance	Inferring systematic relationships from chance occurrences

We developed four problem sets (training materials) for students to practice categorizing scenarios according to these fallacies. Each problem set contained seven multiple choice questions, where each question included a scenario where someone drew a potentially fallacious conclusion from their own observations, and the options corresponded to the 6 fallacies in the PCTE, plus a seventh, where there was no problem with the claim. The scenarios in these multiple choice practice tests were designed to be related to the content

areas of the four Introductory Psychology units in which they were embedded: Neuroscience, Sensation & Perception, Memory, and Learning. In these practice exercises, we carefully controlled the properties of the individuals described in each scenario. Half the items described a male and the other half described a female (randomly assigned across scenarios), and the names of these individuals were designed to be representative of the US population (forenames drawn from US Social Security Administration records, and surnames drawn from US Census records). Because it was not known whether the individuals in these training exercises should be presented as authoritative researchers or as lay public (the PCTE includes a mix of both), we randomly assigned students to receive either of two versions of the training scenarios throughout the semester, one describing inferences seemingly drawn by experts (e.g., Professor Denise Crane), and the other where the person drawing inferences was described by a first name only, with no implication of expertise (e.g., Denise). All training materials are available at <https://osf.io/cg7k6/>.

Procedure.

Pretests and posttests were created by splitting the PCTE into two forms with six items each (one item per fallacy), and students were randomly assigned either form for pretest, and the other was used for posttest (as in Lawson et al., 2015). The pretest was administered in the first week of the semester, and the posttest was administered in the final (16th) week of the semester. These were ungraded, but students received a small amount of course credit for completion.

Training. Students received direct instruction on critical thinking during the third week of the semester (in a content unit about research methods). This included a statement on the importance of critical thinking, having students read definitions and descriptions of each of the six fallacies included in the PCTE, and a short graded multiple-choice practice test (which was presented on the same page as the definitions) that students could attempt as many times as they wanted prior to the deadline. The practice test included six scenarios (similar to the PCTE

items) illustrating each of the six fallacies, as well as a seventh item where there was nothing wrong with the inference in the scenario, and students were instructed to identify the problem (if any) in multiple-choice format (with 7 labeled options: Random chance, Lack of control, Correlation is not causation, Overgeneralization, Experimenter Bias, Confirmation Bias, and There is no fallacy in this conclusion).

Then, in the subsequent four units of the course (Neuroscience, Sensation & Perception, Memory, and Learning), training activities were embedded amidst typical online Introductory Psychology lessons and content. Each training activity involved the completion of a graded multiple choice problem set, as described above (see *Materials*). There were roughly 2 weeks between the administration of each problem set, which could be completed multiple times prior to the due date, allowing students to maximize their scores.

Data Analysis.

Response scoring. Items on the PCTE (Lawson et al., 2015), which were used for the pretest and posttest, are open-ended questions which ask participants to explain the problem in a scenario. We manually scored participants' responses using the original PCTE scoring rubric, which was modeled on a scale used by Williams et al. (2003): 0 (no problem identified), 1 (a problem recognized but misidentified), 2 (identified main problem, but also mentioned less relevant problems), and 3 (identified only the main problem). We did not require participants to respond with specific words, terms, or labels, rather to describe the problem in their own words — for example, the response “Although the weight loss supplement may have actually caused their weight-loss, there are other factors that could have,” was scored as 3, the maximum score, for a “Lack of Control” item (PCTE item #3). Blank responses were scored as 0. In two cases when non-English responses were provided, we used Google Translate to translate the students' responses to English and then scored the responses accordingly. Responses were scored by a trained coder who was not able to see whether the responses were from a pretest or posttest.

Statistical analysis. The goal of statistical analysis in this study is to examine differences in PCTE scores between pretest and posttest. Toward this goal, we carried out two different analytical approaches. Because PCTE scores are ordinal, we first analyzed the data using a hierarchical ordered probit model, adapted from the model described by Liddell & Kruschke (2018); but due to the complexity of this approach, we have opted to make it available as an online supplement (see Supplementary Materials). Additionally, we also implemented a more intuitive and coarse analysis, simply measuring individual test performance as the percent of items that participants got precisely correct (i.e., scored 3 out of 3 on a particular item), and subtracting pretest performance from posttest performance, creating a metric improvement score for each individual. Only this second analysis is described in the current article's main text, however, both yielded identical inferences.

We use Bayesian estimation for our statistical analyses. Bayesian estimation provides a statistical framework for directly estimating parameter values (such as the typical PCTE improvement score) and quantifying the uncertainty in these estimates. We estimated the tendency of participants' improvements pretest-to-posttest using a robust, hierarchical, Bayesian version of the t-test, described in Kruschke (2013). The prior for group-level estimates of this improvement score had a mean of 0, and a standard deviation that was 100-times the empirical standard deviation of the sample, so that priors were vague and uninformed. Considering that our goal is to examine differences between pretest and posttest, and not to test the improbability of our results under a null hypothesis that such differences do not exist, Bayesian methods are particularly well suited to the current study.

We estimated parameter values in JAGS (Plummer, 2003), using the runjags package (Denwood, 2016) for R. The full model specifications are available at <https://osf.io/vcgzf/>. Our estimation runs used 500 steps of adaptation, 1,000 steps of burn-in, and then 250,000 Markov chain Monte Carlo (MCMC) steps across 4 chains, thinned to every 5th step for a total of 50,000

saved steps total. The effective sample size (ESS) was at least 20,000 for all model estimates, well above the 10,000 recommended by (Kruschke, 2014).

Results and Discussion

Prior to any critical thinking training, participants received a score of 3 (the maximum score) on 30.3% of the pretest questions, while 3.4% of responses received a score of 2, 40.0% received a score of 1, and 26.7% received a score of 0.

Following critical thinking training (at the end of the semester), participants received a score of 3 on 40.3% of the posttest questions, while 4.2% of responses received a score of 2, 32.2% received a score of 1, and 23.3% received a score of 0.

Lawson et al. (2015) did not report details of the distribution of PCTE scores, and instead reported the average cumulative score across all items for different groups of students in their study, which ranged from 35.5% of points possible (senior art majors) to 77.5% (senior psychology majors who recently received critical thinking instruction). If we similarly aggregate the current results (maximum score of 18; students can earn up to 3 points on each of the 6 items), we observe cumulative percent scores of 45.8% at pretest and 53.8% at posttest for students in the training group, roughly in line with Lawson's observations.

At pretest, participants were precisely correct (receiving a score of 3-out-of-3) on 1.82 items (out of 6 items total, 30.3%), and at posttest, were precisely correct on 2.42 items (40.3%). We estimated the mean tendency of this improvement to be 0.596 more items (95% HDI: 0.410 to 0.782), a credibly non-zero improvement in the number of correct responses on the PCTE following training. See Figure 1.

Figure 1. Performance outcomes in Study 1. Histograms on the top panels illustrate the number of items participants got precisely correct (score of 3-out-of-3) on pretest (left panel) and posttest (right panel). Dashed lines indicate group means. The plot on the bottom panel shows individual student-level improvement from pretest to posttest. Each dot illustrates a single student's improvement, with minor horizontal jitter added for visualization (improvement scores are integers). Semi-transparent blue lines show posterior estimates of the mean and standard deviation, modeled as a normal distribution at the group level, which is credibly centered above zero. (Low resolution copy)

For participants whose scenarios included titled full names (e.g., Professor Denise Crane), these metric improvement scores were slightly smaller than scores from participants whose scenarios only included first names (e.g., Denise), but the difference between these groups was not credibly different from zero (difference mode = -0.172; 95% HDI: -0.56 to 0.17). Nevertheless, for our subsequent work (see Study 2 Procedure, below) we opted to only use first names in the training materials.

Figure 2. Pretest and posttest performance for each fallacy in Study 1. Horizontal bars indicate the percent of all participants in the training group who scored 3-out-of-3 for that fallacy, on pretest and on posttest. Error bars show the 95% HDI of each point estimate, as modeled with a simple binomial distribution. (Low resolution copy)

There were modest differences in critical thinking performance across the different fallacies present in the PCTE, shown in Figure 2. Students were least likely to respond correctly on items presenting the “Correlation is not Causation” fallacy, both at pretest and at posttest. This is not surprising, considering that people routinely consider patterns of correlation to provide evidence of causality (Sloman, 2005). But paradoxically, items presenting causal claims from uncontrolled experiments (labeled “Lack of Control”) were most likely to be correct, both at pretest and at posttest, suggesting that respondents may have a reasonable grasp of some features of causality.

We investigated further, examining student performance on the training task, which had multiple-choice rather than open-ended items. Looking only at students’ first attempt at each quiz, we calculated a confusion matrix, shown in Figure 3. Again, “Correlation is not Causation” had the lowest percent correct. We find that the most common incorrect response, when presented with a “Correlation is not Causation” claim, was “Overgeneralization” (15% of

responses). There is an intuitive explanation of this pattern, that participants may be skeptical of whether a correlational pattern would *generalize* to a new sample or setting, which is not inconsistent with concerns of whether a general causal relationship is present (personal communication, S. Sloman). This may also explain the proportionately high frequency of “Correlation is not Causation” responses on “Random Chance” items, also possibly stemming from concerns about a finding’s generalizability. On the other side of the spectrum, items about “Experimenter Bias” had the highest percent correct, consistent with the observation that identification of this fallacy had the largest improvement from pretest to posttest (as shown Figure 2).

		Response							
		Confirmation Bias	Correlation is not Causation	Experimenter Bias	Lack of Control	Nothing Wrong	Overgeneralization	Random Chance	[Blank]
Correct Response	Confirmation Bias	0.48	0.08	0.13	0.08	0.05	0.07	0.08	0.03
	Correlation is not Causation	0.07	0.42	0.06	0.08	0.11	0.15	0.09	0.03
	Experimenter Bias	0.13	0.06	0.53	0.09	0.04	0.06	0.06	0.03
	Lack of Control	0.07	0.1	0.07	0.49	0.1	0.07	0.07	0.03
	Nothing Wrong	0.08	0.09	0.06	0.08	0.52	0.08	0.07	0.03
	Overgeneralization	0.07	0.08	0.06	0.1	0.09	0.51	0.07	0.02
	Random Chance	0.06	0.18	0.05	0.09	0.04	0.09	0.47	0.03

Figure 3. Confusion matrix for training task in Study 1. Horizontal rows indicate the correct response for each training item, and columns indicate participants’ actual responses. Values (and color codings) indicate row percentages. Only includes data from the participant’s first attempt at each training set. Light blue squares along the diagonal indicate a high proportion of correct responses. Proportionately frequent errors include “Correlation is not Causation” responses to “Random Chance” items, and “Overgeneralization” responses to “Correlation is not Causation” items, discussed in the main text. (Low resolution copy)

Taking stock, we observe credible improvements in critical thinking scores from pretest to posttest among those students who were exposed to explicit, direct, critical thinking instruction and practice during their semester of enrollment in Introductory Psychology. However, these students were *also* exposed to routine instruction in Psychology (and in their other courses) so the improvement pre-to-post might also be attributable to other curricular content. Another confound is that the students received both direct instruction on critical thinking fallacies *and* received opportunities to practice categorizing claims according to these fallacies. Thus, this present study conflates the possible benefits of training with the benefits of categorization practice (this was also the case in Bensley & Spero, 2014), limiting our ability to explain the students' improvements by category induction. If we were to draw causal inferences from this group alone, we would certainly be guilty of the "Lack of Control" fallacy.

While the present results provide modest correlational evidence in support of the benefits of critical thinking training in practice, they lack the controls that would be characteristic of a randomized controlled experiment. Ideally, one would administer a pretest to a full cohort of students, then randomize these students to receive critical thinking training using categorization practice or no practice, and finally administer a posttest. However, given the present suggestive evidence of the benefits of training, it would be potentially unethical to withhold this training from enrolled students. Instead, we conduct a controlled experiment with online research volunteers using MTurk.

Study 2 (Randomized Controlled Experiment)

Study 2 was confirmatory, designed to examine the benefits of categorization practice with critical thinking scenarios in a randomized controlled experiment. Prior to data collection, we publicly registered all study methods at <https://osf.io/9j583>, which included fully specified

analysis scripts based on simulated data. This study, as described below, contained no deviations from this pre-registration, and all data are available at <https://osf.io/vcgzf/>.

Method

Participants. Participants were recruited on MTurk (<https://mturk.com>), with the requirements that they were based in the US, were at least 18 years old, and had an MTurk approval rating above 50%. The study's sample size was determined by a financial constraint — we continued to recruit participants in batches until funding was exhausted. Initially, 601 potential participants accessed the screening session, and of these, 497 participants completed the screener and correctly entered their MTurk ID in a validation question. Sixty of those who completed the screener were excluded because they did not write coherent responses to simple questions, leaving 437 participants who passed the screener and were invited to complete the full study. Following exclusions and attrition over the subsequent multi-day online study, a total of 253 participants completed the full sequence. All results described hereafter are limited to these 253 participants with complete datasets. According to responses provided in the initial screening session, 68.4% of these participants had completed a Bachelor's degree or higher, and 85.0% agreed or strongly agreed with the statement "I am skilled at critical thinking." We collected no other demographic data about participants.

Procedure. The current study had four sessions: (1) Screener, (2) Pretest and Training, (3) Intervention, and (4) Posttest (see Figure 4), all conducted online using Qualtrics (Provo, UT). Each day we monitored who had satisfactorily completed each session, and sent invitations for subsequent sessions accordingly (using MTurk's messaging system). In effect, the average time between participants' completion of each session was about 1.4 days.

Figure 4. Overview of Study 2 procedure. Participants who satisfactorily completed each session were invited to the subsequent session. Participants were randomly assigned to one of three conditions for session 3; those who were assigned to the No Treatment Control condition were invited directly to Session 4 after completing Session 2. The average time between each session was 1.4 days. All participants received some critical thinking training in Session 2; the key manipulation is the form of practice (or lack thereof) in Session 3. See main text for details of each session. (Low resolution copy)

Session 1: Screener. The purpose of the Screener was to identify human participants (not automated bots) who could follow instructions and provide coherent responses to open-ended questions. Participants were paid \$0.20 this session, and the median completion time for participants in the final sample was 3.2 minutes.

Session 2: Pretest and Training. This session included a pretest assessment of participants' critical thinking, immediately followed by training about critical thinking fallacies. Just like the training group in Study 1, the pretest was half the items on the PCTE, which was counterbalanced: half the participants got odd numbered items, and the other half got even numbered items for the pretest. We used the same reduced version of the PCTE (omitting items related to Oversimplification), with 12 items total (6 item pretest, 6 item posttest).

After the pretest, participants completed a training round that included two phases: a study phase and a test phase. During the study phase, participants were asked to read and

study instructional text about the 6 fallacies described in Table 1. During the subsequent test phase, participants were asked to respond to 7 multiple choice items, corresponding to each of the 6 fallacies plus an additional item where there was no fallacy, all available at <https://osf.io/n2hu8>. Each item included a scenario similar to the PCTE items, and participants were asked to mark which fallacy (if any) the person committed. The training text from the study phase was included on the same page as each test item, so participants could refer back to definitions of these fallacies when selecting their responses. Participants received feedback indicating the correct answer after each selection.

Participants were paid \$2 for completing session 2, and the median completion time for participants in the final sample was 20.5 minutes.

Session 3: Intervention. Participants were randomly assigned to one of three intervention conditions: (1) critical thinking categorization practice ($n=81$); (2) non-critical thinking categorization practice ($n=91$); or (3) no intervention control ($n=81$). Participants in the no intervention control group were not invited to complete session 3; after completing session 2 they were invited directly to session 4. Participants in the other two conditions were paid \$3 for completing session 3.

For the critical thinking practice group, we administered the same 4 practice sets as in Study 1, each containing 7 multiple-choice critical thinking scenarios. However, whereas practice sets in study 1 randomized the individual's name in each scenario (some participants saw "Professor Denise Crane" while others simply saw "Denise"), in Study 2 all practice sets only used the individual's first name ("Denise") for simplicity, and because we observed directionally better performance when participants were shown first name only. The sets were still organized around typical content units in an Introductory Psychology course (neuroscience, sensation and perception, memory, and learning). After participants selected their answer, verification feedback (correct/incorrect) was shown, and participants had to get at least 5 of 7

questions correct in each set; if they got fewer than 5 items correct, they had to repeat the set a maximum of three times. The median completion time was 36.6 minutes for the final sample of participants in the critical thinking practice group.

Table 2. Examples of items from Session 3 (Intervention) in Study 2

Critical Thinking Categorization Practice	Non-Critical Thinking Categorization Practice
<p>Directions: Read the scenario below and select which fallacy (if any) the person made. You can use the descriptions of the fallacies in the table to help. After you have made a selection, click the arrow to proceed. You will get feedback on your answer, which you can look over.</p> <p>Many people have suggested that some of the symptoms of ADHD might be caused by hyperactivity of the mirror neuron system. In order to treat these symptoms, Douglas has designed a therapy that should improve cognitive control of mirror neurons. This therapy, which involves watching cartoons of animals while keeping perfectly still, was administered to one hundred 12-year-olds who had previously been diagnosed with ADHD. After 6 months of therapy, these patients showed an overall decline in symptoms of ADHD. Douglas concluded that the therapy reduces symptoms of ADHD.</p> <ul style="list-style-type: none"> ☐ Random chance ☒ Lack of control ☐ Correlation is not causation ☐ Overgeneralization ☐ Confirmation bias ☐ Experimenter bias ☐ There is no fallacy in this conclusion 	<p>Directions: Read the scenario below and select which concept applies to that scenario. You can use the definitions of the concepts in the table to help. After you have made a selection, click the arrow to proceed. You will get feedback on your answer, which you can look over.</p> <p>A group of people suffering from depression recently began a new experimental drug therapy trial. Like most pharmacological treatments for depression, this experimental drug affects the action of a certain neurotransmitter. However, this new drug is designed to avoid some of the adverse effects of previous drug therapies affecting this neurotransmitter, such as insomnia, drowsiness, and food cravings. Therapy for people suffering from depression often involves drugs that affect the action of which neurotransmitter?</p> <ul style="list-style-type: none"> ☐ Cerebellum ☐ Substantia nigra ☒ Serotonin ☐ Prefrontal cortex ☐ Temporal lobe ☐ Amygdala ☐ Hippocampus

For the non-critical thinking practice group, we modified existing multiple-choice Introductory Psychology exam questions corresponding to each content unit (neuroscience, sensation and perception, memory, and learning), so that they were length-matched with the critical thinking practice items, and so that there were 7 response options, available at <https://osf.io/bq5a8> (also see Table 2). Each item asked participants to categorize a scenario according to a structure or concept in psychology, with no mention of logical fallacies. Like the

critical thinking questions, these non-critical thinking questions were shown alongside full definitions of all the relevant structures and concepts, verification feedback was shown after each response, and participants had to get 5 of 7 correct to proceed to the next set. The median completion time was 32.6 minutes for the final sample of participants in the non-critical thinking practice group.

Session 4: Posttest. In this session, participants completed the remaining half of the PCTE items. As with the pretest, no information about the fallacies were displayed during the posttest. Participants were paid \$5 for completing session 4, and the median completion time was 11.1 minutes.

Data Analysis

Response Scoring. Similar to Study 1, responses to PCTE items were scored by a trained coder who was not able to see participants' condition assignments, nor whether the responses were from a pretest or posttest. Scoring again used the PCTE's original coding scheme: 0 (no problem identified), 1 (a problem recognized but misidentified), 2 (identified main problem, but also mentioned less relevant problems), and 3 (identified only the main problem). Our primary analytical goal is to assess whether improvement from pretest to posttest differs between the three intervention conditions: (1) critical thinking categorization practice; (2) non-critical thinking categorization practice; or (3) no intervention control.

Statistical Analysis. We registered two separate statistical analyses prior to data collection, and both were implemented without deviation from their registered forms. These analyses were identical to the analyses described in Study 1 (ordered probit model, and metric improvement model), but extended to include pairwise contrasts in improvement pre-to-post between groups (Study 1 only had pre-to-post contrasts for one group). Additionally, the prior for group-level estimates of improvement scores for each treatment group used the mean improvement for the full study sample. As with Study 1, the ordered probit analysis is described

in supplementary material, and the metric analysis is described in the main text below; both yielded identical findings.

We estimated model parameters with 100 adaptation steps, 500 burn-in steps, and 100,000 samples thinned to every 5th step (20,000 saved samples) across 4 MCMC chains, again using JAGS and the `runjags` package for R, and the full registered model specification is available at <https://osf.io/vju8w/>. The effective sample size (ESS) was at least 20,000 for all model estimates, well above the 10,000 recommended by (Kruschke, 2014).

Results and Discussion

Averaging across all treatment conditions, prior to any training, participants got 2.4 items precisely correct (scored 3 out of 3; 40.3%), received a score of 2 on 0.4 items (6%), received a score of 1 on 2.0 items (33%), and received a score of 0 on 1.2 items (21%), out of 6 items total on the pretest. These pretest scores are modestly higher than pretest scores observed in our previous college sample (Study 1), which may be due to differences in age or participants' education level at time of the study; for example, 68.4% of participants in Study 2 reported having completed a Bachelor's degree or higher, compared with 0% in Study 1, who were early in college enrollment. Moreover, attrition during Study 2's four sessions may have produced a more selective final sample than Study 1.

On the posttest, still averaging across all treatment conditions, participants got a score of 3 on 3.1 items (52.3%), a score of 2 on 0.4 items (7.2%), a score of 1 on 1.4 items (23.7%), and a score of 0 on 1 item (16.8%), out of 6 items total.

According to Lawson et al.'s (2015) point tallying method, these results correspond to an average pretest score of 55.3% (9.96 points out of 18 possible) and an average posttest score of 65% (11.7 points).

The number of items participants got precisely correct (3-out-of-3), again aggregating across all treatment conditions, was 2.4 at pretest and 3.1 at posttest. Posterior estimates of

average improvement (posttest-pretest; full sample) had a mode of 0.72, and a 95% HDI of 0.55 to 0.93. It makes sense that we see credibly non-zero improvement across the full study sample, considering that all participants received some training on critical thinking fallacies immediately after taking the pretest.

In our hierarchical model, we also estimated improvement from pretest to posttest for each treatment condition separately at the group-level. The modal improvement estimate in the critical thinking practice condition was 1.33 items (95% HDI: 0.96 to 1.67), in the non-critical thinking practice condition was 0.46 items (95% HDI: 0.13 to 0.81), and in the no treatment control was 0.41 items (95% HDI: 0.13 to 0.70). These improvement scores, with group-level distribution estimates, are shown in Figure 5, visibly demonstrating that scores are larger in the critical thinking practice condition than in the other two conditions. Pre-registered pairwise contrasts find that improvement in the critical thinking practice condition was credibly larger than the non-critical thinking practice condition (difference mode = 0.82; 95% HDI: 0.33 to 1.32), and larger than the no treatment control (difference mode = 0.89; 95% HDI: 0.44 to 1.35). There was no evidence of any credible difference between the non-critical thinking practice condition and the no treatment control condition (difference mode = 0.035; 95% HDI: -0.395 to 0.491).

Figure 5. Improvement in the number of items correct between pretest and posttest in Study 2. Each dot illustrates a single participant's improvement, with minor horizontal jitter added for visualization (improvement scores are integers). Semi-transparent blue lines show credible estimates of the mean and standard deviation, modeled as a normal distribution at the group level. Results indicate larger improvements in the critical thinking practice condition (top row) compared with the other two conditions. (Low resolution copy)

The magnitude of these differences are notably larger than observed in Study 1, which might again be attributable to sample differences between these studies, or to the relative recency of practice and training at posttest in Study 2. On this second point, the posttest was administered 1.4 days after practice in Study 2, but roughly 14 days after the most recent practice session in Study 1.

Considering that all three treatment groups (critical thinking practice, non-critical thinking practice, and no practice) took both pretest and posttest, Study 2 affords the opportunity to

investigate differences in the patterns of improvement for the six different fallacies, for the different treatment groups. In an exploratory (not pre-registered) analysis we estimated parameters for a metric improvement model similar to the one described above, but with an additional hierarchical level for the different fallacies (see <https://osf.io/bfvg3/>). Consistent with Study 1, we observe the smallest improvement for PCTE items related to “Correlation is not Causation,” and largest improvement for “Experimenter Bias.” Unlike Study 1, we observe robust improvements for items related to “Random Chance” in the group that received critical thinking practice. However, interpreting random chance coincidences as evidence of a causal mechanism is, in a sense, misattributing a correlational pattern for causation, in the special case involving few observations.

Figure 6. Improvement from pretest and posttest in Study 2, separated by fallacy. Horizontal location indicates the magnitude of improvement from pretest to posttest, estimated on a latent metric scale. Points show the most credible estimate, error bars show the 95% HDI.

In Study 2, all participants received training about logical fallacies, and in all three conditions, participants demonstrated improved performance on an open-ended critical thinking assessment. However, participants who had a separate session practicing categorizing scenarios according to these logical fallacies had significantly higher gains. Participants assigned to this condition had roughly three times larger improvement pre-to-post (1.33 more items correct) than participants who practiced categorizing basic concepts (0.46 items) and participants who had no practice (0.41 items).

General Discussion

The current research extends a robust research lineage demonstrating that the strategic use of abstract “rules for reasoning” is a skill, and that this skill can be trained through induction (reviewed in Holland et al., 1989; Nisbett, 2013). In the present study, we show that such an approach can be constructively applied to train participants to categorize claims according to different logical fallacies and biases, and that this training improves performance on a delayed, open-ended critical thinking assessment. In Study 1, we observed this improvement over an academic semester in an authentic college education setting, and in Study 2, we conducted a pre-registered randomized controlled experiment, demonstrating appreciable improvements caused by induction with critical thinking practice problems, beyond those observed by a single training session, and by practice with traditional learning activities.

By practicing the categorization of fallacies and biases, we believe that participants induce mental models of the abstract patterns exemplified in the training materials. Armed with such mental models, a participant later exposed to a logical fallacy or a biased claim will be better-equipped to recognize and label the claim as fallacious, a skill commonly associated with critical thinking ability, and one that is urgently in demand, considering the volume of disinformation online.

The central feature of the current work is our inductive approach to critical thinking training, presenting multiple sets of contextualized example claims which participants categorized using a simple multiple-choice response format. In this regard, our research was opposite Renaud & Murray's (2008) approach; rather than designing a critical thinking intervention involving open-form essay questions, and evaluating critical thinking improvement using a multiple-choice assessment instrument (Watson-Glaser Critical Thinking Appraisal; WGCTA), we created an intervention using multiple-choice practice questions, and assessed improvement on an open-form critical thinking assessment. The current results add to existing evidence that multiple choice practice questions can develop higher-order understanding (Hitchcock, 2004; Jensen et al., 2014; Scully, 2017), and we believe that our proposed intervention is more portable and easy to integrate in education settings, addressing calls for incorporating modular critical thinking practice throughout a course (Stevens et al., 2016).

Meta analyses of critical thinking instruction in education settings have inferred that contextualized materials, deliberate practice, and repeated exposure are common elements of successful strategies for improving critical thinking (Abrami et al., 2015; Niu et al., 2013; Tiruneh et al., 2014; Willingham, 2008). However, with rare exceptions (Heijltjes et al., 2014; Hitchcock, 2004), we have found no scholarly review that directly considered, let alone examined, the use of an inductive strategy for building learners' critical thinking skills. Instead, education scholars tend to advocate for deep analysis of open-ended cases and in-depth group discussion, strategies which are laudable, but time-intensive and threaten incursion into other curricular instruction. Cáceres (2020) observed that many teachers do make ad-hoc efforts to infuse critical thinking into subject matter instruction, albeit implemented informally, without dedicated modules or lessons. These efforts suggest that many educators are willing to incorporate critical thinking instruction into their curriculum, and the instructional strategy we propose herein offers a flexible and effective method for doing so.

Our current implementation takes the form of an online quiz, a routine learning device in contemporary education settings. However, alternative formats might also confer similar benefits, such as a recent web browser-based game for detecting fake news (Basol et al., 2020; Roozenbeek & van der Linden, 2019). This game's prebunking approach extends a metaphor where people become inoculated against persuasive techniques (McGuire, 1964; Papageorgis & McGuire, 1961), by exposing them to weakened forms of misinformation, progressively training their resistance to fake news. If resistance to misinformation is analogous to being able to accurately identify fallacies in claims, we see the current results as providing converging evidence of the benefits of induction to building such cognitive defenses.

In the current study, we found that inductive training, and these resulting defenses, were not equally resistant to all fallacies. The single most difficult fallacy to identify, as evidenced by it having the lowest scores during pretest, training, posttest, and improvement from pretest to posttest, involved causal claims drawn from correlational data (labelled "Correlation is not Causation"). Indeed, in the course of natural human experience, correlations typically provide evidence that a causal mechanism is at play, and humans are generally disposed to infer causal structure from correlational patterns (Sloman, 2005). Nevertheless, special cases of fallacious correlational inferences, where change over time is attributed to an antecedent action (labelled "Lack of Control") and where rare coincidences are attributed to a causal relationship (labelled "Random Chance"), showed improvements during training in comparison to "Correlation is not Causation" items. Evidently the current training develops sensitivities to specific features of correlational inferences, but general causal claims from correlational findings are more resistant to induction, a topic that is ripe for future research. In contrast, the items that did not involve correlational patterns ("Confirmation Bias" and "Experimenter Bias") showed credible improvements attributable to training.

Still, we acknowledge some limitations of the present results. For example, the difference in the magnitude of improvement (from pretest to posttest) between studies 1 and 2,

which differed substantially in the delay between training and posttest, raise questions about the persistence of these effects. Even so, it may yet be possible to strengthen the inductive training regimen, and perhaps to see more enduring effects over longer delays, by presenting more examples and/or repeating the examples to increase the number of training items. Expanding the set of practice problems, perhaps further to align with other disciplinary content (beyond Introductory Psychology) would also test the generalizability of this approach. Further, the question of whether inductive training has isolated benefits for our narrow operationalization of critical thinking ability, or whether it facilitates broader benefits (such as a tendency to approach claims with skepticism, or other attitudinal dispositions) also begs for further examination.

Democracy depends on a well-informed society. Considering the volumes of disinformation that presently circulate online, it is critically important to equip individuals to identify and reject fallacious and biased claims, or more bluntly, to be able to call bullsh*t. The current work advances a novel method for doing so, which has potential for practical impact given its flexibility and effectiveness.

References

- Abrami, P. C., Bernard, R. M., Borokhovski, E., Waddington, D. I., Wade, C. A., & Persson, T. (2015). Strategies for Teaching Students to Think Critically: A Meta-Analysis. *Review of Educational Research*, 85(2), 275–314. <https://doi.org/10.3102/0034654314551063>
- Arum, R., & Roksa, J. (2011). *Academically Adrift: Limited Learning on College Campuses*. University of Chicago Press.
- Association of American Colleges and Universities. (2005). *Liberal Education Outcomes: A Preliminary Report on Student Achievement in College*. Association of American Colleges & Universities.
- Basol, M., Roozenbeek, J., Berriche, M., Uenal, F., McClanahan, W. P., & Linden, S. van der. (2021). Towards psychological herd immunity: Cross-cultural evidence for two prebunking interventions against COVID-19 misinformation. *Big Data & Society*, 8(1), 20539517211013868. <https://doi.org/10.1177/20539517211013868>
- Basol, M., Roozenbeek, J., & van der Linden, S. (2020). Good News about Bad News: Gamified Inoculation Boosts Confidence and Cognitive Immunity Against Fake News. *Journal of Cognition*, 3(1), 2. <https://doi.org/10.5334/joc.91>
- Bensley, D. A., & Lilienfeld, S. O. (2017). Psychological Misconceptions: Recent Scientific Advances and Unresolved Issues. *Current Directions in Psychological Science*, 26(4), 377–382. <https://doi.org/10.1177/0963721417699026>
- Bensley, D. A., Lilienfeld, S. O., & Powell, L. A. (2014). A new measure of psychological misconceptions: Relations with academic background, critical thinking, and acceptance of paranormal and pseudoscientific claims. In *Learning and Individual Differences* (Vol. 36, pp. 9–18). <https://doi.org/10.1016/j.lindif.2014.07.009>
- Bensley, D. A., & Spero, R. A. (2014). Improving critical thinking skills and metacognitive monitoring through direct infusion. *Thinking Skills and Creativity*, 12, 55–68.

<https://doi.org/10.1016/j.tsc.2014.02.001>

Bergstrom, C. T., & West, J. D. (2021). *Calling Bullshit: The Art of Skepticism in a Data-driven World*. Random House.

Brooks, L. R., Norman, G. R., & Allen, S. W. (1991). Role of specific similarity in a medical diagnostic task. *Journal of Experimental Psychology. General*, 120(3), 278–287.

<https://doi.org/10.1037/0096-3445.120.3.278>

Cáceres, M., Nussbaum, M., & Ortiz, J. (2020). Integrating critical thinking into the classroom: A teacher's perspective. *Thinking Skills and Creativity*, 37, 100674.

<https://doi.org/10.1016/j.tsc.2020.100674>

Carvalho, P. F., Braithwaite, D. W., de Leeuw, J. R., Motz, B. A., & Goldstone, R. L. (2016). An In Vivo Study of Self-Regulated Study Sequencing in Introductory Psychology Courses.

PloS One, 11(3), e0152115. <https://doi.org/10.1371/journal.pone.0152115>

Carvalho, P. F., & Goldstone, R. L. (2015). What you learn is more than what you see: what can sequencing effects tell us about inductive category learning? *Frontiers in Psychology*, 6,

505. <https://doi.org/10.3389/fpsyg.2015.00505>

Cheng, P. W., Holyoak, K. J., Nisbett, R. E., & Oliver, L. M. (1986). Pragmatic versus syntactic approaches to training deductive reasoning. *Cognitive Psychology*, 18(3), 293–328.

[https://doi.org/10.1016/0010-0285\(86\)90002-2](https://doi.org/10.1016/0010-0285(86)90002-2)

Day, S. B., Motz, B. A., & Goldstone, R. L. (2015). The Cognitive Costs of Context: The Effects of Concreteness and Immersiveness in Instructional Examples. *Frontiers in Psychology*, 6,

1876. <https://doi.org/10.3389/fpsyg.2015.01876>

Denwood, M. (2016). runjags: An R Package Providing Interface Utilities, Model Templates, Parallel Computing Methods and Additional Distributions for MCMC Models in JAGS.

Journal of Statistical Software, Articles, 71(9), 1–25. <https://doi.org/10.18637/jss.v071.i09>

Dunlosky, J., Rawson, K. A., Marsh, E. J., Nathan, M. J., & Willingham, D. T. (2013). Improving Students' Learning With Effective Learning Techniques: Promising Directions From

- Cognitive and Educational Psychology. *Psychological Science in the Public Interest: A Journal of the American Psychological Society*, 14(1), 4–58.
<https://doi.org/10.1177/1529100612453266>
- Ennis, R. H. (1987). A taxonomy of critical thinking dispositions and abilities. In J. B. Baron & R. J. Sternberg (Eds.), *Teaching Thinking Skills: Theory and Practice.*, 275, 9–26. W H Freeman/Times Books/ Henry Holt & Co.
- Facione, P. (1990). *Critical Thinking: A Statement of Expert Consensus for Purposes of Educational Asesment and Instruction. Research Findings and Recommendations* (ERIC Document ED315423). American Philosophical Association.
- Fong, G. T., Krantz, D. H., & Nisbett, R. E. (1986). The effects of statistical training on thinking about everyday problems. *Cognitive Psychology*, 18(3), 253–292.
[https://doi.org/10.1016/0010-0285\(86\)90001-0](https://doi.org/10.1016/0010-0285(86)90001-0)
- Fyfe, E. R., McNeil, N. M., Son, J. Y., & Goldstone, R. L. (2014). Concreteness Fading in Mathematics and Science Instruction: a Systematic Review. *Educational Psychology Review*, 26(1), 9–25. <https://doi.org/10.1007/s10648-014-9249-3>
- Garimella, K., De Francisci Morales, G., Gionis, A., & Mathioudakis, M. (2018). Political Discourse on Social Media: Echo Chambers, Gatekeepers, and the Price of Bipartisanship. *Proceedings of the 2018 World Wide Web Conference*, 913–922.
<https://doi.org/10.1145/3178876.3186139>
- Goldstone, R. L., Kersten, A., & Carvalho, P. F. (2013). Concepts and categorization. In A. F. Healy (Ed.), *Handbook of psychology: Experimental psychology, Vol* (Vol. 4, pp. 607–630). John Wiley & Sons, Inc., xix.
- Goldstone, R. L., Rogosky, B. J., Pevtzw, R., & Blair, M. (2017). Chapter 34 - The Construction of Perceptual and Semantic Features During Category Learning. In H. Cohen & C. Lefebvre (Eds.), *Handbook of Categorization in Cognitive Science (Second Edition)* (pp. 851–882). Elsevier. <https://doi.org/10.1016/B978-0-08-101107-2.00034-8>

- Gurung, R. A. R., Hackathorn, J., Enns, C., Frantz, S., Cacioppo, J. T., Loop, T., & Freeman, J. E. (2016). Strengthening introductory psychology: A new model for teaching the introductory course. *The American Psychologist*, 71(2), 112–124.
<https://doi.org/10.1037/a0040012>
- Halonen, J. S., Buskist, W., Dunn, D. S., Freeman, J., Hill, G. W., Enns, C., & Others. (2013). APA guidelines for the undergraduate psychology major (version 2.0). *Washington, DC: APA*.
- Halpern, D. F. (2013). *Thought and Knowledge: An Introduction to Critical Thinking*. Psychology Press.
- Harnad, S. (2017). Chapter 2 - To Cognize is to Categorize: Cognition is Categorization. In H. Cohen & C. Lefebvre (Eds.), *Handbook of Categorization in Cognitive Science (Second Edition)* (pp. 21–54). Elsevier. <https://doi.org/10.1016/B978-0-08-101107-2.00002-6>
- Heijltjes, A., van Gog, T., & Paas, F. (2014). Improving Students' Critical Thinking: Empirical Support for Explicit Instructions Combined with Practice. *Applied Cognitive Psychology*, 28(4), 518–530. <https://doi.org/10.1002/acp.3025>
- Hitchcock, D. (2004). The Effectiveness of Computer Assisted Instruction in Critical Thinking. *Informal Logic*, 24(3). <https://doi.org/10.22329/il.v24i3.2145>
- Holland, J. H., Holyoak, K. J., Nisbett, R. E., & Thagard, P. R. (1989). *Induction: Processes of Inference, Learning, and Discovery*. MIT Press.
- Homa, N., Hackathorn, J., Brown, C. M., Garczynski, A., Solomon, E. D., Tennial, R., Sanborn, U. A., & Gurung, R. A. R. (2013). An Analysis of Learning Objectives and Content Coverage in Introductory Psychology Syllabi. *Teaching of Psychology*, 40(3), 169–174.
<https://doi.org/10.1177/0098628313487456>
- Horton, R. (2020). Offline: Managing the COVID-19 vaccine infodemic. *The Lancet*, 396(10261), 1474. [https://doi.org/10.1016/S0140-6736\(20\)32315-1](https://doi.org/10.1016/S0140-6736(20)32315-1)
- Im, S.-H., Cho, J.-Y., Dubinsky, J. M., & Varma, S. (2018). Taking an educational psychology

- course improves neuroscience literacy but does not reduce belief in neuromyths. *PloS One*, 13(2), e0192163. <https://doi.org/10.1371/journal.pone.0192163>
- Izawa, C. (1967). Function of test trials in paired-associate learning. *Journal of Experimental Psychology*, 75(2), 194–209. <https://doi.org/10.1037/h0024971>
- Jensen, J. L., McDaniel, M. A., Woodard, S. M., & Kummer, T. A. (2014). Teaching to the Test...or Testing to Teach: Exams Requiring Higher Order Thinking Skills Encourage Greater Conceptual Understanding. *Educational Psychology Review*, 26(2), 307–329. <https://doi.org/10.1007/s10648-013-9248-9>
- Johnson, R. H., & Hamby, B. (2015). A Meta-Level Approach to the Problem of Defining “Critical Thinking.” In *Argumentation* (Vol. 29, Issue 4, pp. 417–430). <https://doi.org/10.1007/s10503-015-9356-4>
- Kang, S. H. K., & Pashler, H. (2012). Learning Painting Styles: Spacing is Advantageous when it Promotes Discriminative Contrast. *Applied Cognitive Psychology*, 26(1), 97–103. <https://doi.org/10.1002/acp.1801>
- Kim, E. S., & Willson, V. L. (2010). Evaluating Pretest Effects in Pre–Post Studies. *Educational and Psychological Measurement*, 70(5), 744–759. <https://doi.org/10.1177/0013164410366687>
- Kolb, D. A. (2015). *Experiential Learning: Experience as the Source of Learning and Development*. Pearson Education.
- Kornell, N., & Bjork, R. A. (2008). Learning concepts and categories: is spacing the “enemy of induction”? *Psychological Science*, 19(6), 585–592. <https://doi.org/10.1111/j.1467-9280.2008.02127.x>
- Kowalski, P., & Taylor, A. K. (2004). Ability and Critical Thinking as Predictors of Change in Students’ Psychological Misconceptions. *Journal of Instructional Psychology*, 31(4).
- Kruschke, J. K. (2013). Bayesian estimation supersedes the t test. *Journal of Experimental Psychology. General*, 142(2), 573–603. <https://doi.org/10.1037/a0029146>

- Kruschke, J. K. (2014). *Doing Bayesian Data Analysis: A Tutorial with R, JAGS, and Stan*. Academic Press.
- Lantian, A., Bagneux, V., Delouvée, S., & Gauvrit, N. (2021). Maybe a free thinker but not a critical one: High conspiracy belief is associated with low critical thinking ability. *Applied Cognitive Psychology*, 35(3), 674–684. <https://doi.org/10.1002/acp.3790>
- Lawson, T. J., Jordan-Fleming, M. K., & Bodle, J. H. (2015). Measuring Psychological Critical Thinking: An Update. *Teaching of Psychology*, 42(3), 248–253. <https://doi.org/10.1177/0098628315587624>
- Lazer, D. M. J., Baum, M. A., Benkler, Y., Berinsky, A. J., Greenhill, K. M., Menczer, F., Metzger, M. J., Nyhan, B., Pennycook, G., Rothschild, D., Schudson, M., Sloman, S. A., Sunstein, C. R., Thorson, E. A., Watts, D. J., & Zittrain, J. L. (2018). The science of fake news. *Science*, 359(6380), 1094–1096. <https://doi.org/10.1126/science.aao2998>
- Lewandowsky, S., Ecker, U. K. H., & Cook, J. (2017). Beyond Misinformation: Understanding and Coping with the “Post-Truth” Era. *Journal of Applied Research in Memory and Cognition*, 6(4), 353–369. <https://doi.org/10.1016/j.jarmac.2017.07.008>
- Lewandowsky, S., & van der Linden, S. (2021). Countering Misinformation and Fake News Through Inoculation and Prebunking. *European Review of Social Psychology*, 1–38. <https://doi.org/10.1080/10463283.2021.1876983>
- Liddell, T. M., & Kruschke, J. K. (2018). Analyzing ordinal data with metric models: What could possibly go wrong? *Journal of Experimental Social Psychology*, 79, 328–348. <https://doi.org/10.1016/j.jesp.2018.08.009>
- Lilienfeld, S. O., Ammirati, R., & Landfield, K. (2009). Giving Debiasing Away: Can Psychological Research on Correcting Cognitive Errors Promote Human Welfare? *Perspectives on Psychological Science: A Journal of the Association for Psychological Science*, 4(4), 390–398. <https://doi.org/10.1111/j.1745-6924.2009.01144.x>
- Lobato, E., Mendoza, J., Sims, V., & Chin, M. (2014). Examining the relationship between

conspiracy theories, paranormal beliefs, and pseudoscience acceptance among a university population. *Applied Cognitive Psychology*, 28(5), 617–625.

<https://doi.org/10.1002/acp.3042>

Machete, P., & Turpin, M. (2020). The Use of Critical Thinking to Identify Fake News: A Systematic Literature Review. *Responsible Design, Implementation and Use of Information and Communication Technology*, 12067, 235. https://doi.org/10.1007/978-3-030-45002-1_20

McCutcheon, L. E., Apperson, J. M., Hanson, E., & Wynn, V. (1992). Relationships among critical thinking skills, academic achievement, and misconceptions about psychology. *Psychological Reports*, 71(2), 635–639. <https://doi.org/10.2466/pr0.1992.71.2.635>

McGuire, W. J. (1964). Some Contemporary Approaches¹¹The research by the author which is reported in this chapter was greatly facilitated by two successive grants from the National Science Foundation, Division of Social Sciences. He wishes to acknowledge the aid itself and the enlightened manner in which it was administered by the Foundation. In L. Berkowitz (Ed.), *Advances in Experimental Social Psychology* (Vol. 1, pp. 191–229). Academic Press. [https://doi.org/10.1016/S0065-2601\(08\)60052-0](https://doi.org/10.1016/S0065-2601(08)60052-0)

Nisbett, R. E. (2013). *Rules for reasoning*. Psychology Press.

Niu, L., Behar-Horenstein, L. S., & Garvan, C. W. (2013). Do instructional interventions influence college students' critical thinking skills? A meta-analysis. *Educational Research Review*, 9, 114–128. <https://doi.org/10.1016/j.edurev.2012.12.002>

Nosofsky, R. M. (1986). Attention, similarity, and the identification–categorization relationship. *Journal of Experimental Psychology. General*, 115(1), 39–57. <https://doi.org/10.1037/0096-3445.115.1.39>

Nosofsky, R. M., Sanders, C. A., & McDaniel, M. A. (2018). Tests of an exemplar-memory model of classification learning in a high-dimensional natural-science category domain. *Journal of Experimental Psychology. General*, 147(3), 328–353.

<https://doi.org/10.1037/xge0000369>

Papageorgis, D., & McGuire, W. J. (1961). The generality of immunity to persuasion produced by pre-exposure to weakened counterarguments. *Journal of Abnormal and Social Psychology*, 62, 475–481. <https://doi.org/10.1037/h0048430>

Pennycook, G., Epstein, Z., Mosleh, M., Arechar, A. A., Eckles, D., & Rand, D. G. (2019). Shifting attention to accuracy can reduce misinformation online. In *PsyArXiv*. <https://doi.org/10.31234/osf.io/3n9u8>

Pennycook, G., & Rand, D. G. (2019). Lazy, not biased: Susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning. *Cognition*, 188, 39–50. <https://doi.org/10.1016/j.cognition.2018.06.011>

Pennycook, G., & Rand, D. G. (2021). The psychology of fake news. *Trends in Cognitive Sciences*, 25(5), 388-402. <https://doi.org/10.1016/j.tics.2021.02.007>

Plummer, M. (2003). *JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling*. <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.13.3406>

Poland, G. A., & Spier, R. (2010). Fear, misinformation, and innumerates: how the Wakefield paper, the press, and advocacy groups damaged the public health. *Vaccine*, 28(12), 2361–2362. <https://doi.org/10.1016/j.vaccine.2010.02.052>

Porter, A., McMaken, J., Hwang, J., & Yang, R. (2011). Common Core standards. *Educational Researcher*, 40(3), 103–116. <https://doi.org/10.3102/0013189x11405038>

Quilici, J. L., & Mayer, R. E. (2002). Teaching students to recognize structural similarities between statistics word problems. *Applied Cognitive Psychology*, 16(3), 325–342. <https://doi.org/10.1002/acp.796>

Rawson, K. A., Thomas, R. C., & Jacoby, L. L. (2015). The Power of Examples: Illustrative Examples Enhance Conceptual Learning of Declarative Concepts. *Educational Psychology Review*, 27(3), 483–504. <https://doi.org/10.1007/s10648-014-9273-3>

Renaud, R. D., & Murray, H. G. (2008). A comparison of a subject-specific and a general

- measure of critical thinking. *Thinking Skills and Creativity*, 3(2), 85–93.
<https://doi.org/10.1016/j.tsc.2008.03.005>
- Roozenbeek, J., & van der Linden, S. (2019). Fake news game confers psychological resistance against online misinformation. *Palgrave Communications*, 5(1), 1–10.
<https://doi.org/10.1057/s41599-019-0279-9>
- Schacter, D., Gilbert, D., Wegner, D., & Hood, B. (2016). *Psychology*.
<https://doi.org/10.1007/978-1-137-40673-6>
- Schmaltz, R., Jansen, E., & Wenckowski, N. (2017). Redefining Critical Thinking: Teaching Students to Think like Scientists. *Frontiers in Psychology*, 8, 459.
<https://doi.org/10.3389/fpsyg.2017.00459>
- Schmaltz, R., & Lilienfeld, S. O. (2014). Hauntings, homeopathy, and the Hopkinsville Goblins: using pseudoscience to teach scientific thinking. *Frontiers in Psychology*, 5, 336.
<https://doi.org/10.3389/fpsyg.2014.00336>
- Scully, D. (2017). Constructing multiple-choice items to measure higher-order thinking. *Practical Assessment, Research, and Evaluation*, 22(4), 1–13. <https://doi.org/10.7275/SWGT-RJ52>
- Self, T. T., & Self, J. T. (2017). Reality in the Classroom: Teaching Critical Thinking with Scenarios. *Journal of Hospitality & Tourism Education*, 29(4), 187–193.
<https://doi.org/10.1080/10963758.2017.1382367>
- Sloman, S. (2005). *Causal Models: How People Think About the World and Its Alternatives*. Oxford University Press.
- Smith, E. E., Langston, C., & Nisbett, R. E. (1992). The case for rules in reasoning. *Cognitive Science*, 16(1), 1–40. https://doi.org/10.1207/s15516709cog1601_1
- Sternberg, R. J. (1999). Teaching psychology students to be savvy consumers and producers of research questions. *Teaching of Psychology*, 26(3), 211–213.
<https://doi.org/10.1207/S15328023TOP260311>
- Stevens, C., Witkow, M. R., & Smelt, B. (2016). Strengthening scientific reasoning skills in

introductory psychology: Evidence from community college and liberal arts classrooms.

Scholarship of Teaching and Learning in Psychology, 2(4), 245–260.

<https://doi.org/10.1037/stl0000070>

Tavris, C., & Aronson, E. (2020). *Mistakes Were Made (but Not by Me) Third Edition: Why We Justify Foolish Beliefs, Bad Decisions, and Hurtful Acts*. Houghton Mifflin Harcourt.

Tiruneh, D. T., Verburgh, A., & Elen, J. (2014). Effectiveness of Critical Thinking Instruction in Higher Education: A Systematic Review of Intervention Studies. *Higher Education Studies*, 4(1), 1. <https://doi.org/10.5539/hes.v4n1p1>

Tullis, J. G., & Goldstone, R. L. (2016). Comparison versus reminding. *Cognitive Research: Principles and Implications*, 1(1), 20. <https://doi.org/10.1186/s41235-016-0028-1>

Tulving, E. (1967). The effects of presentation and recall of material in free-recall learning. *Journal of Verbal Learning and Verbal Behavior*, 6(2), 175–184.

[https://doi.org/10.1016/S0022-5371\(67\)80092-6](https://doi.org/10.1016/S0022-5371(67)80092-6)

van Gelder, T. (2005). Teaching Critical Thinking: Some Lessons From Cognitive Science. *College Teaching*, 53(1), 41–48. <https://doi.org/10.3200/CTCH.53.1.41-48>

Wahlheim, C. N., Dunlosky, J., & Jacoby, L. L. (2011). Spacing enhances the learning of natural concepts: an investigation of mechanisms, metacognition, and aging. *Memory & Cognition*, 39(5), 750–763. <https://doi.org/10.3758/s13421-010-0063-y>

Williams, R. L., Oliver, R., Allin, J. L., Winn, B., & Booher, C. S. (2003). Psychological Critical Thinking As a Course Predictor and Outcome Variable. *Teaching of Psychology*, 30(3), 220–223. https://doi.org/10.1207/S15328023TOP3003_04

Willingham, D. T. (2008). Critical Thinking: Why Is It So Hard to Teach? *Arts Education Policy Review*, 109(4), 21–32. <https://doi.org/10.3200/AEPR.109.4.21-32>

Zohar, A., & Tamir, P. (1993). Incorporating critical thinking into a regular high school biology curriculum. *School Science and Mathematics*, 93(3), 136–140.