Behavior Research Methods
2009, 41 (3), 691-698
doi:10.3758/BRM.41.3.691

The efficacy of a Web-based counterargument tutor

Christopher R. Wolfe
Miami University, Oxford, Ohio

M. Anne Britt
Northern Illinois University, DeKalb, Illinois

Melina Petrovic and Michael Albrecht
Miami University, Oxford, Ohio

and

Kristopher Kopp
Northern Illinois University, DeKalb, Illinois

In two experiments, we developed and tested an interactive Web-based tutor to help students identify and evaluate counterarguments. In Experiment 1, we determined the extent to which high- and low-argumentation-ability participants were able to identify counterarguments. We tested the effectiveness of having participants read didactic text regarding counterarguments and highlight claims. Both preparations had some positive effects that were often limited to high-ability participants. The Web-based intervention included interactive exercises on identifying and using counterarguments. Web-based presentation was state driven, using a Java Server Pages page. As participants progressively identified argument elements, the page changed display state and presented feedback by checking what the user clicked against elements that we had coded in XML beforehand. Instructions and feedback strings were indexed by state, so that changing state selected new text to display. In Experiment 2, the tutor was effective in teaching participants to identify counterarguments, recognize responses, and determine whether counterarguments were rebutted, dismissed, or conceded.

A few years ago, the first author was using *The Blank Slate* by Steven Pinker (2003) in a sophomore seminar. The discussion of the first chapter was confusing until it became apparent that a number of the students erroneously thought that Pinker was arguing for the blank slate rather than against it. As an author, Pinker makes extensive use of counterargument in several of his books. Unfortunately, many readers have real trouble comprehending the uses of counterargument, which is the focus of this research and development campaign.

Consider the following written argument that includes a counterargument:

> We should prohibit the use of cell phones while driving because drivers talking on cell phones are more likely to get into accidents. Some people say that it is unwise to create laws to protect people from themselves, but people who talk on the phone while driving are a threat to pedestrians and other drivers as well as themselves.

We define an argument as one or more claims supported by one or more reasons (Toulmin, 1958; Voss, 2005; Wolfe & Britt, 2008). In this case, "we should prohibit the use of cell phones while driving" is the claim, which is supported by the reason "because drivers talking on cell phones are more likely to get into accidents." A counterargument is an argument against the author's claim or supporting reasons. In this example, the author raises the counterargument that "some people say that it is unwise to create laws to protect people from themselves," which attacks the supporting reason. Finally, the author provides a response to that counterargument: "but people who talk on the phone while driving are a threat to pedestrians and other drivers as well as themselves." Generally speaking, the claim of one argument commonly serves as the reason for a second argument, and the backing of a given argument can itself be another argument (Voss, 2005; Voss, Fincher-Kiefer, Wiley, & Silfes, 1993). Effectively generating and evaluating counterarguments are fundamental skills of argumentation, but research suggests that even bright students often have difficulty processing counterarguments (Wolfe & Britt, 2008; Wolfe, Britt, & Butler, 2009).

The ability to effectively use counterarguments develops relatively late (Kuhn & Udell, 2003). When asked to

C. R. Wolfe, wolfecr@muohio.edu

generate arguments, adults typically generate far more reasons in favor of a position they support than they do reasons on the other side (Macpherson & Stanovich, 2007; Perkins, 1985; Perkins, Farady, & Bushey, 1991; Toplak & Stanovich, 2003). Although talented writers often include counterarguments in their published "authentic arguments" (Wolfe & Britt, 2005), many undergraduate authors fail to include counterarguments in their written essays (Wolfe & Britt, 2008). Moreover, in unpublished pilot research, we have found that undergraduate readers have difficulty even identifying counterarguments in argumentative texts.

There are, however, positive consequences of including counterarguments in argumentative essays. Wolfe et al. (2009) found that brief written arguments that include counterarguments and rebuttals produced significantly higher levels of agreement, higher ratings of argument strength, and better impressions of the author than the same arguments presented without counterarguments. In short, facility with counterarguments is important to both readers and writers of argumentative texts. Even bright undergraduates, however, have difficulty understanding argumentative texts with counterarguments.

We hypothesize that part of the difficulty readers have in identifying and evaluating counterarguments stems from difficulties in comprehending claims made by the author. There are three aspects of a claim: the theme, the side, and the specific predicate (Britt, Kurby, Dandotkar, & Wolfe, 2008; Wolfe et al., 2009). Consider two similar claims:

1. Talking on a cell phone while driving should be discouraged.
2. Talking on a cell phone while driving should be prohibited.

In each case, the theme of the claim is using cell phones while driving, and in both cases the side is against this practice. In the case of the first claim, the specific predicate is that the practice should be discouraged, whereas in the second case the predicate is that the practice should be prohibited (i.e., outlawed). We have found that people are generally quite good at remembering the theme and side of a claim but significantly worse at recalling specific claim predicates (Britt et al., 2008). Indeed, participants often incorrectly recall an argument predicate like "should be outlawed" as "should be discouraged," even though there are clear differences that necessitate the use of different reasons and prompt different counterarguments.

Counterarguments generally attack the claim and, sometimes, the supporting reasons. In the case of the two claims above, "the police have enough to do without ticketing people for talking on their cell phones" is a counterargument to Claim 2 but not to Claim 1. Many people have difficulty identifying the claim predicate and are thus unsure about what is being argued for, so they have difficulty identifying the counterargument predicate and may not fully understand the argument against that claim. In both the classroom and the laboratory, we have seen university students confuse a counterargument with the author's main claim.

As part of a broader research-and-development campaign, we set out to develop an interactive Web-based tutor to help students identify and evaluate counterarguments. Our research and development model may be described as theoretically informed, empirical, and recursive. It is theoretical in that it is informed by psychological theory and research (as well as by work in other disciplines). We have conducted basic research (e.g., Wolfe & Britt, 2008; Wolfe et al., 2009) on the use of counterarguments, and we have conducted applied research in identifying critical issues, developing the tutor, and testing the effectiveness of the tutor (once built) in controlled laboratory experiments. Thus, our process has been recursive, in that the results of one set of experiments are used to generate materials that are tested in the next set of experiments. In the research reported here, the results of Experiment 1 were used to identify the key components of a Web-based tutor and to test their efficacy.

Our tutor is based on the notion that people make use of an argumentation schema to comprehend written arguments. The schema organizes knowledge according to the manner in which it is used. In reading argumentative texts, the argumentation schema may be evoked by a provocative claim, although reasons may occasionally induce the schema. Perhaps the most important expectation generated by the schema is for a reason, an answer to the question, "Why should I accept the claim?" Readers who are generally sympathetic to the claim and reason may readily accept an argument. However, if a reader is less inclined to accept the claim, the schema provokes a consideration of counterarguments. Readers who can readily generate counterarguments may be less likely to "buy" an argument. However, the tendency of readers to minimize cognitive load (Ericsson & Kintsch, 1995) suggests that if the author can fill a counterargument slot in the reader's schema and then adequately address the counterargument, readers are more likely to be satisfied.

Unfortunately, for many individuals, the argumentation schema is underdeveloped. We have found that, as a precursor to deeper comprehension, many people could benefit from practice in recognizing different aspects of arguments and in parsing arguments into their constituent parts. Our approach is to provide brief didactic texts to help people understand various aspects of argumentation and then to provide opportunities for guided practice with scaffolding and simple feedback. Early modules in the argumentation tutor, such as the counterargument tutorial presented here, are designed to be followed by subsequent modules that present more sophisticated skills that build on a more fully developed schema. To illustrate, we are developing a module to help students recognize when university assignments require argumentation in different fields of study across the curriculum and to help them develop arguments in different writing genres.

Another dimension of this research that may be of interest to readers of *Behavior Research Methods* is the computer interface we developed for having students interact with written arguments. Sophisticated uses of text provide special challenges that are keenly felt in Web-based tutor-

ing environments lacking direct interventions by a human teacher. Programming techniques used for addressing such challenges may prove useful in other research and in educational settings that make extensive use of different kinds of text. One such example is the drop-and-drag technique pioneered by Britt and Gabrys (2001), through which users highlight and drag segments of text to labeled "bins" on the screen. By demonstrating that the technique could be used for collecting data and providing feedback, Britt and Gabrys (2004) confirmed the usefulness of the drop-and-drag interface for both research and education. Similarly, the state-based click-on-text interface we developed for this counterargument tutor may have broader utility for other text-intensive contexts.

The Web-based intervention tested in Experiment 2 uses a Java Server Page to present interactive exercises on identifying and using counterarguments. The presentation is state-driven, meaning that the tutor has a small number of states it can be in at any given time and operates on input to change states and produce outputs. As participants progressively identify argument elements, the page changes display state and presents feedback by checking what the user clicks against elements coded in XML. Instructions and feedback strings are indexed by state, so that changing state selects new text to display. The program pulls coded arguments from a file, allows the user to click on portions of text on the screen, checks selected text against argument categories, and provides appropriate feedback on the screen.

EXPERIMENT 1
Didactic Text and Claim Identification

The purpose of the first experiment was to determine the extent to which participants of high and low argumentation ability would be able to identify and evaluate counterarguments in brief written essays and to test the effectiveness of two potential interventions. The first intervention was to have participants read about 1,300 words of didactic text on the constituents of an argument, with particular emphasis on counterarguments (the *tutorial-text* condition). The second intervention was to require participants to actively identify the argument claim while reading argumentative texts (the *identify-claim-only* condition). The rationale for the latter intervention was to reduce confusion between the claim and the counterargument and to teach participants to attend to the specific claim predicate, allowing them to better determine the degree to which counterarguments actually address the claim advanced by the author.

Method

Participants. One hundred twenty native English-speaking students from an introductory psychology class at Miami University participated for partial course credit.

Materials. All participants received seven brief written arguments—six with counterarguments and one without—in two randomly selected orders. The arguments were about a proposed highway, the death penalty, a university math requirement, television and the family, hypnosis, driving with cell phones, and whether the founding fathers of the United States were deists. The arguments averaged about 570 words in length. Participants in the tutorial-text condition first read a 1,300-word document called "Counterarguments and Other-Side Information."

Design. In a between-participants experimental design with an additional quasi-experimental measure, participants were randomly assigned to one of three groups. The 40 participants in the tutorial-text group received the didactic text, "Counterarguments and Other-Side Information," and instructions to first identify the main argument claim. The 40 participants in the identify-claim-only group also were instructed to first identify the main argument claim, but they did not receive any didactic text. Finally, the remaining 40 participants received no additional instruction. To study the consequences of these manipulations for people of high and low argumentation ability, the participants were first tested on an instrument consisting of 18 items from previously published Law School Admissions Tests (LSATs), and they were divided into high- and low-LSAT groups by a median split.

Procedure. After providing their informed consent, the participants were randomly assigned in blocks of 30 to one of the three experimental conditions. The participants in the tutorial-text condition read the didactic text and were instructed to first identify the main argument claim. The participants in the identify-claim-only condition were instructed to first identify the main claim being made in each essay, and the control participants received no intervention. The participants then read the seven argumentative essays in one of two randomly assigned orders. For each essay, the participants were asked to write down any counterarguments and rebuttals. All of the participants were given an unrelated task before completing the LSAT. The participants were then thanked and debriefed.

Results

The random assignment of participants yielded no differences in LSAT scores ($F < 1$). The participants had three tasks: to identify in writing the claim (tutorial-text and identify-claim-only groups), identify the counterargument, and identify the response to the counterargument. The participants' responses were scored for gist using different words but capturing the same substantial meaning. The identification of the claim was also analyzed verbatim (i.e., when the responses used the same words as those used in the target text). In each case, responses were blind-coded to experimental condition and LSAT score. Two raters at Miami University and Northern Illinois University were trained on a subset of responses and then rated 20% of the materials separately, with an interrater reliability of .91. These analyses yielded five independent measures: identifying the claim predicate using verbatim scoring, identifying the claim predicate using gist scoring, identifying the counterargument theme using gist scoring, identifying the response predicate using gist scoring, and false alarms on any measure.

Participants who had read the tutorial text correctly identified the claim predicate verbatim significantly more often than did those who had not read the tutorial text; the control group was excluded by definition. Over seven arguments, the mean for the tutorial-text group was 3.28 (SD = 1.62), and the mean for the identify-claim-only group was 2.58 (SD = 1.53) [$F(1,78) = 3.95$, $p < .05$]. However, there were no significant differences for gist measures [$F(1,78) = 2.41$, $p = .12$]. This suggests that the didactic text helped participants represent the claim predicate more precisely. However, performance was far from perfect.

For two of the dependent measures, there was an interaction between tutorial (tutorial-text + identify-claim vs. identify-claim-only and control groups) and LSAT score.

Over six arguments with counterarguments, the tutorial text worked significantly better with high-LSAT participants than with low-LSAT participants in helping them identify counterargument themes using a criteria of gist scoring (see Table 1). There was not a main effect of tutorial ($F < 1$); the LSAT was of borderline significance [$F(1,116) = 3.65$, $p = .059$] and the tutorial × LSAT interaction was significant [$F(1,116) = 4.34$, $p < .04$]. This indicates that the high-LSAT participants who had read the tutorial text and identified the main claim missed significantly fewer counterargument themes than did the identify-claim-only and control participants.

The tutorial also worked significantly better with high-LSAT participants than with low-LSAT participants in helping them identify the predicate of the response to the counterargument using a criteria of gist scoring (see Table 1). There was no main effect of tutorial [$F(2,114) = 1.01$, $p = .37$]; there was a borderline effect of high-LSAT participants doing better [$F(1,114) = 4.43$, $p = .06$] and the tutorial × LSAT interaction was significant [$F(2,114) = 4.44$, $p = .01$]. This suggests that the tutorial text was ineffective for low-LSAT participants and that the tutorial text and instructions to identify the main claim helped high-LSAT participants identify the predicate of the response to the counterargument.

There was one additional main effect of tutorial that did not interact with LSAT. The tutorial significantly reduced the number of false alarms, with a mean of 4.90 in the tutorial-text group, as compared with a mean of 7.08 for participants in the identify-claim-only group and 8.38 for the control group [$F(1,116) = 4.93$, $p < .03$]. In our scoring system, the theoretical maximum was 21 false alarms: 1 for each claim, counterargument, and response to counterargument for all seven problems. LSAT was not significant [$F(1,116) = 2.05$, $p = .15$], and the tutorial × LSAT interaction was not significant ($F < 1$).

Discussion

The results of Experiment 1 underscore the need for tutoring on counterarguments; participants in all three groups performed poorly. Presenting the tutorial text and asking participants to first identify the main claim had some positive effects. However, the results were often limited to high-ability participants, as measured by LSAT score. This suggests that more and better didactic text would be required of the Web-based tutor to make it effective with students of all ability levels.

Developing the Web-Based Counterargument Tutor

Following the results of Experiment 1, we determined that the Web-based tutor would have to improve upon the offline interventions in several ways. First, we expanded the tutorial text to about six screens, totaling over 1,500 words, and we made the text more accessible. Figure 1 shows the first screen of "The Role of Counterarguments in Argumentation."

The second and more important change was the creation of interactive exercises for giving each user practice at parsing arguments into their constituents. The users received brief practice arguments one at a time. Figure 2 presents 1 of the 22 brief arguments. The argument reads as follows:

> You should not believe what you read on the Internet because a lot of uneducated people can make Web pages. Of course, some would say that uneducated people often produce quality work. But let's face it, that is really wishful thinking.

For each argument, users had to highlight different aspects of the argument, including the main claim, counterargument, and response to the counterargument. For each element of the argument, users were given written feedback, and they had to identify each element correctly before proceeding to the next argument.

The participants' first task was to click on the main claim. If they got it wrong, they had to go back and try again. Then they were to click on the supporting reason and then the counterargument. Next, they had to determine whether the counterargument attacked the claim, the supporting reason, or something else. Finally, they had to determine whether the response to the counterargument was a rebuttal, a concession, or simply a dismissal of the counterargument, as was the case in the example in Figure 2.

Table 1
Means and Standard Deviations of Correct Identifications of Counterargument Themes and of Response Predicates, by Group and LSAT Level

	Identified Counterargument Theme		Identified Response Predicate	
Group	*M*	*SD*	*M*	*SD*
Tutorial-Text and Identify-Claim				
High LSAT	1.36	1.14	1.64	1.36
Low LSAT	0.56	0.98	0.94	0.94
Identify-Claim-Only				
High LSAT	0.91	1.07	1.73	1.49
Low LSAT	0.83	0.92	0.61	0.92
Control				
High LSAT	0.75	0.91	0.65	0.67
Low LSAT	0.90	1.21	1.15	1.50

Note—LSAT, Law School Admissions Tests.

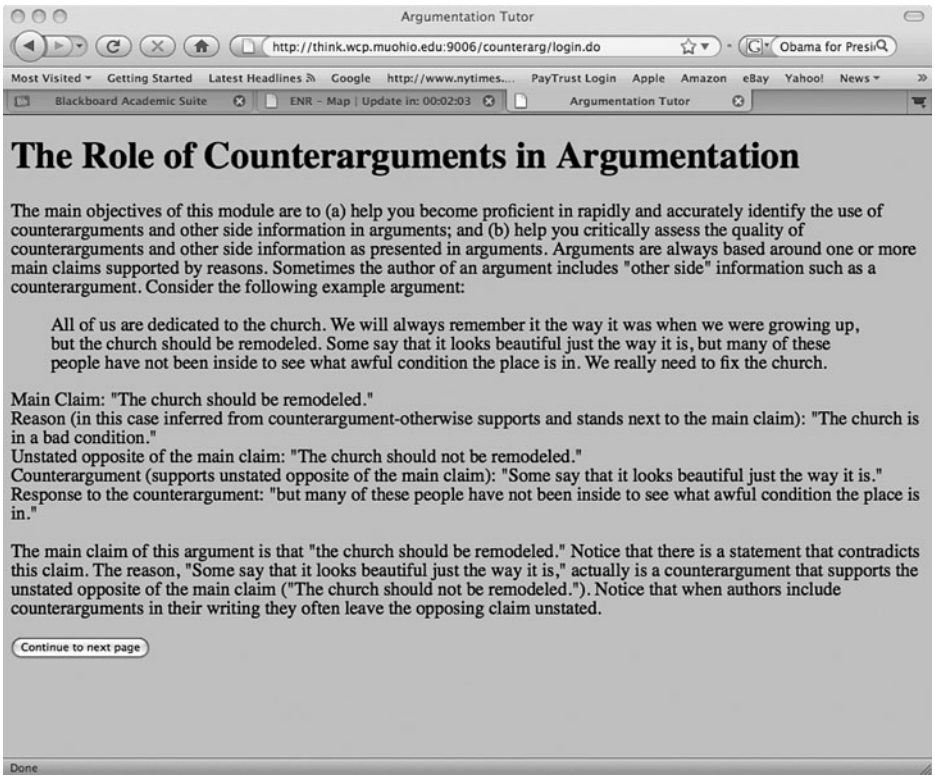

Figure 1. The first of six screens of didactic text.

The Web-based tutor was implemented on a Macintosh OS X server running the Apache Tomcat Java Servlet. The application consists of a set of Java Server Pages (JSP) pages. Argument-file parsing and data logging are handled on the server; presentation and feedback are handled on the client side. We divided the arguments into elements by hand and stored them in XML format. Java classes were written to parse the argument file into passages and to handle randomization. Each of the instruction screens is a separate JSP page. A single JSP page handles the presentation of the task. This page gets a passage object from the passage collection and constructs JavaScript arrays and functions to handle the sequencing of the task on the client side.

The task presentation is state driven. As the user progressively identifies argument elements, the page changes its display state and presents feedback by checking what the user clicks against the element identified in the passage object (i.e., as identified by the experimenters in the XML document). Instructions and feedback strings are assigned to arrays that are indexed by state, so that changing state selects a new set of strings to be displayed on the screen.

The application pulls hand-parsed, XML-coded arguments from a prepared file. It then allows the user to click on portions of text on the screen. The application checks selected text against preestablished argument categories. Finally, the Web-based tutor provides appropriate feedback as developed beforehand by the authors. Users must exhibit the appropriate behavior before going on to the next sample argument.

EXPERIMENT 2
Testing the Web-Based Tutor

We tested the counterargument tutor in a controlled laboratory experiment using LSAT for additional quasi-experimental comparisons on the effect of the interactive tutor between users of high and low argumentation ability.

Method

Participants. One hundred eleven native English-speaking students from an introductory psychology class at Miami University participated for partial course credit.

Materials. The participants in the counterargument-tutor condition received the Web-based tutor consisting of over 1,500 words of didactic text and interactive exercises using 22 brief arguments and 1 long argument. The participants in the control group received a Web-based tutorial on postmodernism and deconstruction. The materials for the control tutor were significantly edited texts adapted from Hedges (2007) and Wikipedia (*Deconstruction*, 2007). The control tutor was designed to address sophisticated uses of language, be demanding for undergraduate readers, and have some degree of face validity in relation to subsequent tasks. All of the participants received a counterargument identification task.

Design. In a between-participants experimental design with an additional quasi-experimental measure, the participants were assigned in random blocks to either the counterargument tutor (n = 56) or the control tutor (n = 55). The participants had been given the 18-item LSAT instrument and were assigned to the high or low group according to a median split.

Procedure. The participants were run individually or in small groups under laboratory conditions. After providing their informed consent, the participants first completed the LSAT task. When they

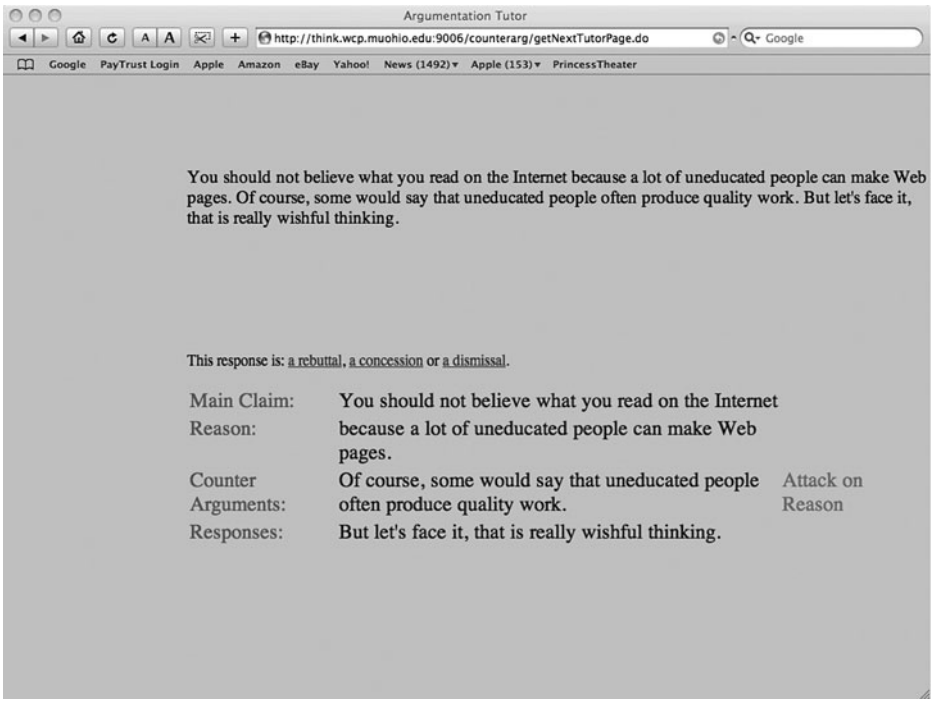

Figure 2. Interactive exercise with 1 of the 22 brief arguments used in Experiment 2.

were done, they were given either the Web-based counterargument tutor or the deconstruction control tutor at individual workstations.

To ensure that they were familiar with the terms, all of the participants first read a description of counterarguments and examples using an abbreviated version of the tutorial text used in Experiment 1. We wanted to rule out any effect of unfamiliarity with the basic concepts. Thus, everyone in both the control and the experimental groups read a one-page description of counterarguments and a correctly worked example, all of which totaled 269 words. The participants then read 11 brief arguments and 1 longer argument. These arguments were not the same as those included in the tutorial. For each argument, the participants underlined the text that presented a counterargument and indicated whether it presented an alternative claim or attacked the main claim or the supporting reason.

The participants identified the responses to the counterarguments by underlining the relevant text and indicated whether the author conceded, rebutted, or dismissed the counterargument. The pencil-and-paper underlining task replaced the written responses and gist scoring used in Experiment 1; thus, the dependent measures were more straightforward. The participants had to highlight the main claims, counterarguments, and so forth on paper copies using different colored markers. After they were done with the online tutor, the participants received a posttest on identifying counterarguments and their function and identifying responses to the counterargument and their function. The participants were then thanked and debriefed.

Results

As in Experiment 1, the random assignment of participants yielded no differences in LSAT scores ($F < 1$). Thus, we are confident that any differences between experimental conditions are attributable to the tutor rather than to ability level, as measured by the LSAT instrument. There were five independent measures in this experiment: For the 11 short arguments, we assessed the mean correct in identifying the counterargument, identifying the response, and characterizing the response to the counterargument. On the long argument, we assessed the number of participants who correctly identified the counterargument and the response to the counterargument.

Turning first to the identification of counterarguments, the participants in the counterargument-tutor condition identified significantly more counterarguments than did the participants in the control condition (Table 2) [$F(1,107) = 26.02, p < .0001$]. There was also a main effect of LSAT, with high-LSAT participants having a mean of 7.58 and low-LSAT participants having a mean of 6.62 [$F(1,107) = 4.93, p < .03$]. The LSAT $\times$ condition interaction was not significant ($F < 1$), and there were no differences between the groups in determining which part of the argument was being attacked (Fs < 1).

On the task of identifying responses to the counterarguments, participants in the experimental group performed significantly better than those in the control group (Table 2) [$F(1,107) = 8.62, p < .004$]. There was also a main effect of LSAT, with a high-LSAT mean of 6.79 and a low-LSAT mean of 5.18 [$F(1,107) = 8.85, p < .003$]. The LSAT $\times$ condition interaction was not significant [$F(1,107) = 2.13, p = .15$].

In characterizing the responses to the counterarguments, participants in the counterargument tutor group were significantly better (Table 2) [$F(1,107) = 12.02, p < .0008$]. There was also a significant main effect of LSAT [$F(1,107) = 14.48, p < .0002$], with high-LSAT participants having a mean of 6.31 and low-LSAT participants having a mean of 4.97. The condition $\times$ LSAT interaction was not significant [$F(1,107) = 1.42, p = .24$].

Table 2
Means and Standard Deviations of Correct Identifications of Counterarguments, of Correct Identifications of Responses, and of Correct Characterizations of Responses

	Identification of Counterarguments		Identification of Responses		Characterization of Responses	
Group	*M*	*SD*	*M*	*SD*	*M*	*SD*
Counterargument tutor	8.25	1.74	7.10	2.53	7.13	1.72
Control tutor	6.22	2.41	5.31	2.97	5.31	2.00

On the single long argument, participants in the counterargument-tutor condition were slightly but not significantly better at identifying the counterargument [$\chi^2(109) = 2.36, p < .12$]. However, participants who had received the counterargument tutor were significantly better at identifying the response to the counterargument on the long argument (37 of 54, 69%) than were their counterparts in the control group (21 of 55, 38%) [$\chi^2(109) = 10.24, p < .001$].

Discussion

These results suggest that the Web-based counterargument tutor was generally effective in teaching participants to identify counterarguments and responses to counterarguments and to determine whether the counterarguments were rebutted, dismissed, or conceded. Unlike in Experiment 1, argumentation ability, as measured by the 18-item LSAT task, did not interact with the presence or absence of the counterargument tutor. We attribute this to the additional text and interactive experiences that were designed to be effective with participants of high and low argumentation ability; the data seem to support our interpretation. The task instructions also provided a good deal of didactic text, as well as a correctly worked example, for participants in both the control and experimental conditions. This may have reduced some of the variability; however, it increases our confidence that the results represent deeper conceptual change and not just superficial familiarity with the key concepts.

GENERAL DISCUSSION

These experiments show that people have difficulties with counterarguments. Even when instructed in the identification of counterarguments and of the responses to them, our participants performed poorly. Recognizing and assessing arguments and counterarguments is an important, advanced literacy skill that even good college students often lack. Our recursive research-and-development model was labor intensive. However, the use of data from controlled experiments increases our confidence in the efficacy of the tutor. The Web-based counterargument tutor with interactive exercises and didactic text demonstrably improved performance. This suggests that the counterargument tutor may be useful as part of a broader effort to improve argument comprehension and production skills. It also indicates that educators at all levels may benefit from additional attention to counterarguments and, more generally, to "other-side" information.

The state-based click-on-text interface used here may have some broader utility for education, research, and other text-intensive contexts. Other forms of writing, such as narratives, expository texts, and even prose and poetry, could be hand coded along a number of dimensions, including identifying the topic sentences of paragraphs, the elements of story grammar (e.g., the protagonist), and rhyme scheme. Such an interface could be used to shed light on research issues, such as what readers consider the most interesting or important portions of a text. For example, do all readers identify the same person as the main character in a work of fiction? Perhaps boys and girls identify different characters in a children's book as "the protagonist." An interface like the one described here may help researchers determine when and how readers make such determinations online, in "real time," while reading. Of course, parsing text and coding it by hand in XML format is somewhat time consuming. However, reliably coding texts automatically for complex constructs, such as argument constituents, remains elusive.

We undertook this project because we understood that argumentation skills are broadly applicable—yet underdeveloped—even among bright university students. However, we did not initially expect performance with counterarguments to be as poor as was demonstrated in these experiments. We expected to begin teaching more sophisticated skills of the sort needed to critically evaluate the way authors, such as Steven Pinker, use counterarguments in their work. We discovered that more fundamental lessons are the first order of business. Some colleagues in the humanities give lip service to subtle uses of argumentation in the classroom. But we found that, even with argumentative texts of only 40 words, asking untutored students to underline the counterargument yielded success rates just above 50%. As researchers, we know that one of the lessons of this research-and-development project is not to assume the presence of skills that have not been demonstrated. As citizens with an investment in higher education, we find the lessons learned about student abilities to be even more sobering.

AUTHOR NOTE

This research was supported by a grant from the U.S. Department of Education Institute of Education Sciences. Correspondence concerning this article should be addressed to C. R. Wolfe, Department of Psychology, Miami University, Oxford, OH 45056 (e-mail: wolfecr@muohio.edu).

REFERENCES

BRITT, M. A., & GABRYS, G. L. (2001). Teaching advanced literacy

skills for the World Wide Web. In C. R. Wolfe (Ed.), *Learning and teaching on the World Wide Web* (pp. 74-90). San Diego: Academic Press.

BRITT, M. A., & GABRYS, G. [L.] (2004). Collecting responses through Web page drag and drop. *Behavior Research Methods, Instruments, & Computers*, **36**, 52-68.

BRITT, M. A., KURBY, C. A., DANDOTKAR, S., & WOLFE, C. R. (2008). I agreed with what? Memory for simple argument claims. *Discourse Processes*, **45**, 52-84. doi:10.1080/01638530701739207

DECONSTRUCTION (2007). Retrieved September 18, 2007, from Wikipedia: http://en.wikipedia.org/wiki/Deconstruction.

ERICSSON, K. A., & KINTSCH, W. (1995). Long-term working memory. *Psychological Review*, **102**, 211-245.

HEDGES, W. (2007). *Using deconstruction to astonish friends and confound enemies (in 2 easy steps)*. Retrieved September 18, 2007, from www.sou.edu/English/Hedges/Sodashop/RCenter/Theory/Howto/decon.htm.

KUHN, D., & UDELL, W. (2003). The development of argument skills. *Child Development*, **74**, 1245-1260.

MACPHERSON, R., & STANOVICH, K. E. (2007). Cognitive ability, thinking dispositions, and instructional set as predictors of critical thinking. *Learning & Individual Differences*, **17**, 115-127.

PERKINS, D. N. (1985). Postprimary education has little impact on informal reasoning. *Journal of Educational Psychology*, **77**, 562-571.

PERKINS, D. N., FARADY, M., & BUSHEY, B. (1991). Everyday reasoning and the roots of intelligence. In J. F. Voss, D. N. Perkins, & J. W. Segal (Eds.), *Informal reasoning and education* (pp. 83-105). Hillsdale, NJ: Erlbaum.

PINKER, S. (2003). *The blank slate: The modern denial of human nature*. New York: Penguin.

TOPLAK, M. E., & STANOVICH, K. E. (2003). Associations between myside bias on an informal reasoning task and amount of post-secondary education. *Applied Cognitive Psychology*, **17**, 851-860.

TOULMIN, S. E. (1958). *The uses of argument*. Cambridge: Cambridge University Press.

VOSS, J. F. (2005). Toulmin's model and the solving of ill-structured problems. *Argumentation*, **19**, 321-329.

VOSS, J. F., FINCHER-KIEFER, R., WILEY, J., & SILFES, L. N. (1993). On the processing of arguments. *Argumentation*, **7**, 165-181.

WOLFE, C. R., & BRITT, M. A. (2005, November). *The use of other side information: Explaining the my side bias in argumentation*. Paper presented at the 46th Annual Meeting of the Psychonomic Society, Toronto, Canada.

WOLFE, C. R., & BRITT, M. A. (2008). Locus of the myside bias in written argumentation. *Thinking & Reasoning*, **14**, 1-27. doi:10.1080/13546780701527674

WOLFE, C. R., BRITT, M. A., & BUTLER, J. A. (2009). Argumentation schema and the myside bias in written argumentation. *Written Communication*, **26**, 183-209. doi:10.1177/0741088309333019

(Manuscript received November 19, 2008;
revision accepted for publication March 12, 2009.)