Developing writing skills through students giving instructional explanations

Kwangsu Cho
University of Missouri
Sungkyunkwan University, Seoul

Christian Schunn
University of Pittsburgh

1. Introduction

Writing skills are considered to be critical for academic and professional success (National Commission on Writing, 2004). However, a large number of students are not writing well. According to National Assessment of Educational Progress (2002), 69% of 8th graders and 77% of 12th graders have only basic writing skills. Moreover, 50% of college students cannot produce texts that are relatively free of errors (ICAS, 2002). This unfortunate situation also permeates government and industry sectors. State government employees are found to have weak writing skills (NCW, 2004). Salaried employees in major US firms also lack writing skills (NCW, 2004).

Writing is a very difficult skill for students to master. Writing is an ill-structured and complex task that requires a number of cognitive processes such as planning, translating, reviewing, and monitoring (Hayes *et al.*, 1987). High quality writing further requires a large amount of situation-specific adaptation, with the features requiring adaptation being relatively abstract and complex constructs like the writing goal, the genre, and the audience. In order to teach these complex and abstract elements of writing effectively, rich instructional explanations are highly likely to be required.

While rich instruction is required, limited instructional attention is the actual situation. There are generally too few writing activities taking place in typical classrooms. While part of the problem may be the significant amount of student time

required for each piece of writing, the largest bottleneck standing in the way of more writing in the classroom is instructor time. The complex, open nature of writing makes instruction challenging: reading, commenting on, and grading on student writing easily overwhelms instructors. Together these two challenges lead to near total neglect of writing especially in subject matter courses (NCW, 2003).

To illustrate, we conducted an analysis of writing in undergraduate psychology courses at twelve different universities, covering a range of university ranking, university type (research, teaching), university size (large, medium, small), university funding source (public, private), and tuition costs (high, medium, and low). We obtained all course syllabi (for at least two semesters in most cases), and then coded for the presence of required papers. We also coded for the presence of required multiple drafts, because learning to write is likely to be much weaker without attending to instructor feedback.

Although one might have expected that university size or rank would strongly predict the inclusion of writing, in fact annual tuition was the best predictor. Those with the highest tuition levels were the ones to most commonly require writing in their courses (see Figure 1). Presumably tuition predicts class size (and teaching assistant resources), which in turn predicts inclusion of writing. Interestingly, required drafts were very rare across the board (see Figure 1), and almost non-existent in the highest tuition category. If psychology instructors primarily see writing in their classes from a writing-to-learn perspective rather than from a learning-to-write perspective, then that would explain why required drafting is so rare, regardless of institutional resources.

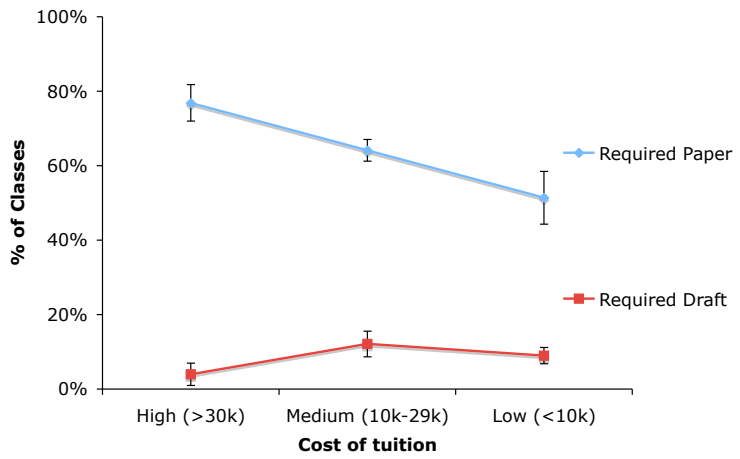

Figure 1. The mean percentage (and standard error) of psychology undergraduate courses that include a required paper or a required draft as a function of annual tuition (in state).

Responding to Leinhardt's call for research on the development of higher levels of academic literacy in students (Young & Leinhardt, 1998), one possible solution to including more writing (in the resource poor settings) with revisions (in all settings) is use of reciprocal peer reviewing (RPR). In RPR, students review each other's writing rather than the traditional model of only receiving feedback from the instructor or a teaching assistant. Through RPR, students are potentially learning from both the feedback they receive and from the task of giving feedback.

Over the past twenty-five years, scholars of rhetoric and composition have continually emphasized the importance of assigning multiple drafts for improving students' writing skills (Beach & Friedrich, 2006) and participating in peer review. These pedagogical principles grow out of a long research tradition showing that gains in writing occur in classes that adopt a process-based approach and less teacher-centered classroom (Hillocks, 1984). Not surprisingly, these approaches are consistent with learning theories

that promote active learning, including collaborative and cooperative learning, provision of feedback, repeated opportunities to practice, and relevant domain-specific tasks (Ashbaugh, Johnstone, & Warfield, 2002; Cornelius-White, 2007; Palincsar & Brown, 1984; Vygotsky, 1978).

In theory, RPR addresses the bottleneck problem. Regardless of class size, with every increase in the number of authors, a corresponding increase in the number of potential reviewers occurs. Indeed, we have seen RPR used with substantial papers in classes with as many as 300 students.

However, the shift to RPR also marks a very interesting shift from instructor-centric to student-centric writing instruction. This shift involves instructional explanations in two different ways. The obvious shift involves the nature of instructional explanations that students (as authors) receive from their peers. As it turns out, systematic comparisons of instructor-generated vs. peer-generated feedback have found that the feedback that students receive from their peers has many similarities with the feedback that they receive from instructors. While there are some systematic differences, overall the similarities outweigh the differences (K. Cho, Schunn, & Charney, 2006; Patchan, Charney, & Schunn, 2009).

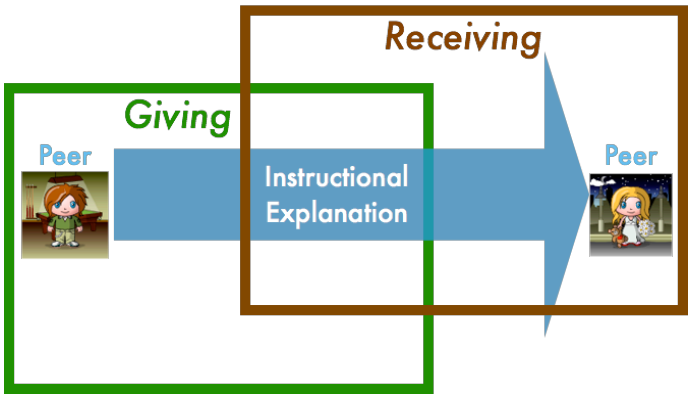

Figure 2. The two sides of instructional explanations in peer review.

The less obvious shift involves the change from students being the *receivers* of instructional explanations to students being the *generators* of instructional explanations. It is this kind of instructional explanation that is the focus herein. A student participating in RPR generates explanations on peer writing through a rich process. In the role of reviewer, a student engages in reading, text analysis, and writing. He or she must carefully read a draft, interpret the evaluative criteria, detect and prioritize problems, make a holistic assessment, and draw on writing skills to formulate comments. Coming to understand the criteria well enough to apply them to another student's paper provides students with the opportunity to improve their own writing and revision activities. More learning gains are possible through a second round of review, in which the same reviewers evaluate the writers' revised drafts and receive feedback on the helpfulness of their comments. We summarize evidence that, indeed, students are benefiting from giving these instructional explanations.

In the next section, we describe a technological support tool that we have created for broadly implementing peer reviewing in small and large classrooms. Then, in later

sections we describe research that we have done on the benefits that this system provides for students in learning how to write, both looking at the benefits of *receiving* peer feedback and the benefits of *giving* peer feedback.

2. A Technological Solution

To fully realize the benefits of reciprocal peer review in many settings, an automatic administration mechanism is required due to the logistical complexity of keeping tracking of scores—if not hundreds—of papers, reviewers, and reviews. A web-based platform for reciprocal peer review offers unique opportunities for improving writing instruction across the board.

We have developed such a system called SWoRD (K. Cho & Schunn, 2007), for Scaffolded Writing and Rewriting in the Disciplines. SWoRD is a Web-based, hybrid intelligent system that implements reciprocal peer reviewing of writing. It was initially developed for use in large undergraduate courses in academic disciplines in which writing is rarely assigned. Since 2002, SWoRD has been used by about 6,000 students from 120 courses at fifteen universities and for 5 courses in secondary schools in the U.S.. Interestingly, the largest user setting of SWoRD is smaller disciplinary courses in which writing may have happened previously, but in weaker form (e.g., only a single draft or with relatively little feedback).

SWoRD enables instructors to implement a wide range of reciprocal peer review activities. First, students as authors write first drafts in response to any task given by the instructor and submit them online. Then, students are randomly assigned a set of three to six of these first drafts to peer review. As reviewers, they analyze the written texts in detail along several evaluative dimensions in response to prompts that incorporate

explicit rubrics. Students submit written comments and ratings online. Figure 3 presents an example of the prompts used to guide the ratings and comments students provide on their peers' writing.

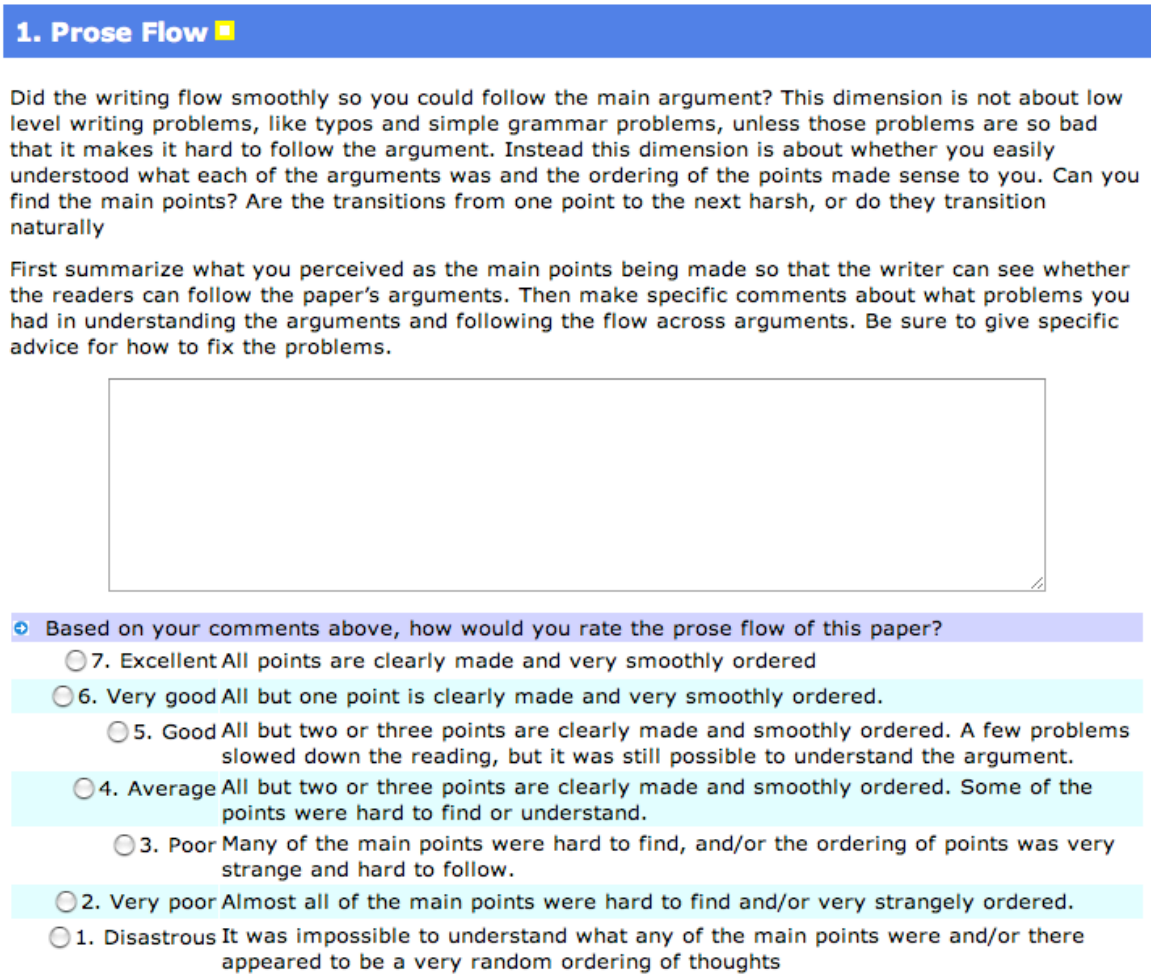

Figure 3. Example writing prompt and evaluation rubrics used in SWoRD.

After students as authors receive the feedback from their peers, they revise their drafts and re-submit them. SWoRD also asks authors provide comments to the reviewers on the helpfulness of the suggestions. Then the revised drafts are made available to the same set of reviewers who evaluated the first drafts. The reviewers then observe how the

revised drafts have changed (or not) in response to the full set of reviews. The reviewers rate and comment along with the same dimensions, providing data on gains from the first draft. In this explicit, step-by-step fashion, students are required to go through iterations of writing, reviewing, and revision. We hypothesize that this iterative writing and reviewing help students to develop good models of writing practice which they may begin to apply to other writing contexts. Once these behaviors become automatized over multiple writing tasks, students can become better self-regulated writers (Zimmerman & Kitsantas, 1999, 2002).

SWoRD's design makes it well suited to overcoming the core obstacles to implementing peer review described above. First, it relieves the logistical burden by automatically administering the collection and exchange of drafts and reviews, monitoring completion of each step, and providing summative statistics on both writing and reviewing activities. Instructors may increase the number of reviews, the number of reviewers, and the rounds of review without additional effort (Rada, Michailidis, & Wang, 1994).

Second, the SWoRD system helps to establish and maintain conducive pedagogical practices with peer review and makes it easier for instructors to design writing assignments and rubrics. SWoRD incorporates a case-based reasoning (CBR) module, storing instances of writing genres and writing evaluation rubrics. The CBR module currently stores 25 types of writing evaluation rubrics that are designed for particular types of writing. For example, there are rubrics for scientific research papers and there are different rubrics for papers that ask students to explain a scientific theory and then apply it to an everyday situation. The former involves longer papers and has

rubrics specific to typical sections of a research report (e.g., abstract, introduction, methods) whereas the latter has rubrics for the (shorter) paper as a whole (e.g., the flow, argument transparency, and insight provided in a paper).

Each dimension of the rubric has clear explanations and rating points have a clear anchor—what features a paper should have to deserve the given rating in order to maximize consistency across reviewers. When instructors create their courses in SWoRD, instructors may select existing rubrics for their writing assignments, modify existing rubrics for their purpose, or create new ones. Modified and new rubrics are also stored for the use of other instructors.

Third, to support effective feedback generation, SWoRD helps to provide an atmosphere conducive to collaboration and guides students to produce high quality reviews. Student papers are randomly assigned to reviewers, with authors and reviewers double-blinded to each other's identity, thereby fostering peer's willingness to provide reviewer that are appropriately critical. In addition, by receiving comments from multiple students (rather than just one or two as in typical classroom practice), the issue of diverse audience becomes very salient. Writing is difficult not just because one must write clearly and persuasively to an individual starting from a different position than the writer, but rather it is especially difficult because one must write clearly and persuasively to different groups of individuals each starting in a different position. Feedback from just one person does not make this diversity of audience problem salient, but feedback from multiple reviewers can. Indeed, our research shows that students benefit from receiving comments from both stronger and weaker writers (Nelson, Melot, Stevens, & Schunn, 2008).

In addition, because authors assess the quality of the reviews, students have incentives to make comments that are constructive and helpful. SWoRD also provides reviewers an interface for self-monitoring, making it easy for reviewers to compare their comments with other reviewers on the same papers. By making it easy for students to view and re-view an explicit evaluative rubric, SWoRD helps students articulate their evaluations and expand their knowledge of discourse structures. Assuming that students review five drafts, they are exposed to the rubric about 14 times per writing assignment—10 as reviewer across both drafts, twice as self-assessor across both drafts, and twice as receiver of feedback organized around the rubric.

SWoRD also incorporates a calibration exercise in which students practice reviewing three writing samples (good, mid-level, and poor) with given evaluation rubrics before actually reviewing peer drafts. The exercise module provides feedback on how the instructor would rate each sample on each dimension. SWoRD provides students with a multimedia module allowing them to hear real student-authors explaining their experiences with helpful and unhelpful feedback.

3. Empirical Research

In the early research on SWoRD, studies focused on the effectiveness of peers' instructional explanations relative to instructor explanations. We found that peers, on average, generated roughly similar kinds of feedback, although with some differences in relative length, focus on problems, inclusion of solutions, and inclusion of praise (Patchan et al., 2009). Multiple peers produce much more feedback in total than a single instructor; the peers find the comments just as useful as instructor comments (K. Cho et

al., 2006); and more revision behavior results from the multiple peer feedback, producing better final drafts (K. Cho & Schunn, 2007).

This chapter examines the less-explored aspect of generating instructional explanations: the role of providing peer reviews. In particular, we focus on the question of whether providing peer reviews improves the reviewer's own writing skills. Peer reviewing is an active process that can help the reviewer learn which features of writing are desirable and which features are undesirable. Thus reviewers are engaged in exercising important skills for writing (Bereiter & Scardamalia, 1987; Fitzgerald, 1987; Flower *et al.*, 1986). These activities may improve the reviewers' own writing skills, by finding strengths and weaknesses, reinforcing successful strategies, and calling attention to unsuccessful strategies that the reviewers have already used in their own writing.

On the other hand, while providing instructional explanations might have opportunities for learning, they might not always involve actual learning. The reviewer may be exposed to problems they themselves do not have or may not be able to generate methods for effective revision. Further, the literature on transfer of skills finds that even when there is clear overlap of underlying skills between the learning environment and the testing environment, transfer is not always observed (Bransford & Schwartz, 2001). For example, the time between generating a possible solution to the time at which it is required in later writing may be so great that the insight is lost. Further, the proposed solution may be framed in the mind of the reviewer in such a specific way that it does not seem applicable to other writing tasks.

Further, assuming there are benefits of reviewing for one's own writing, there is the question about what aspects of generating peer reviews is actually useful for

improving ones own learning. Critically for this book, there are three propositions worth considering: 1) participating in generating of instructional explanations *per se* is useful, 2) simply the task of reading others' papers is sufficient in terms of learning from models (Charney & Carlson, 1995), or 3) practicing detection skills by evaluating others' papers is most useful for learning to write.

Recently, Cho (in prep) examined the hypothesis that reviewing (evaluating AND generating explanations) is be more helpful than simple reading (i.e., learning from models) in supporting students' development of high quality writing. Students in a college physics class were randomly assigned to one of the three conditions: After-reviewing, After-Reading, and No-reading-or-reviewing. All the students were given three sample lab papers that had the same structure as the papers the students were asked to generate. Students in the After-reviewing condition were asked to review the papers with the given rubric and to generate written comments on and rate the papers. Students in the After-reading condition were asked to read the sample papers twice but were not told to comment on or rate them. The students in the No-reading-or-reviewing condition were not asked to read or review the papers. Instead these students were simply notified that the sample papers were available in the SWoRD system. All the students were then asked to generate a paper of their own. Three physics PhD students evaluated the papers to see which condition produced the strongest quality papers according the given evaluation rubric. That is, the study tested whether students could apply something from the reading or reviewing tasks to improve their own writing.

Figure 4 presents the mean 1st draft quality (averaged across evaluation dimensions). The writing quality in the After-Reviewing condition was significantly

higher than in either the After-Reading or in No-reading-or-reviewing conditions. Furthermore, the After-Reading condition papers were no better than the No-reading-or-reviewing condition papers. These results suggest that reviewing activities do benefit the writer's own writing. Further, these results suggest that the locus of the benefit does not stem (at least in this context) from simply reading the papers of others. However, it is unclear from these results whether evaluating papers or generating comments both contribute the writer's learning.

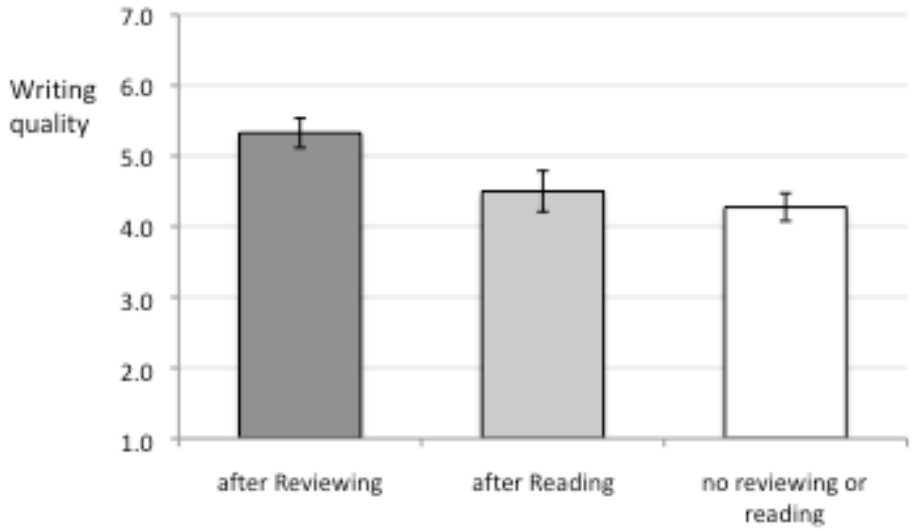

Figure 4. Mean 1st draft quality (and SE bars) as a function of condition.

Wooley, Was, Schunn, and Dalton (2008) followed up on this work to examine the relative benefits of evaluating examples versus providing revision suggestions. In order to populate a complex nested experimental design, this study took place in a large educational psychology undergraduate class (over 180 students). A third of the students wrote a paper early in the semester. For the purposes of the experiment, they were simply generating materials for the other conditions. Another third of the students (Evaluate-First

condition) were asked to evaluate five peer papers, and then write their own paper. The final third of the students (Write-First condition), wrote for the same deadline as the Evaluate-First condition, but did not have to evaluate papers before writing. Students were randomly assigned to one of these conditions.

The Evaluate-First condition was further subdivided into two conditions critical to the question of the efficacy of instructional explanations. The Rate-Only sub-condition had students evaluate five papers using a rubric (three dimensions with 1-to-7 Likert ratings like those shown in Figure 2). They did not have to generate any comments. The Rate+Comment sub-condition had students rate the papers as well as generate helpful comments. By examining the quality of the students' own writing across these two conditions, we can examine whether the comment generation process confers benefits to the reviewer above-and-beyond the benefits of simply carefully evaluating models (i.e., practicing error detection skills). Three PhD students blinded to condition evaluated all of the 1st draft papers in the Evaluate-First (Rate-Only and Rate+Comment) and Write-First conditions.

Figure 5 presents the mean 1st draft quality scores in the Rate-Only and Rate+Comment conditions on each of the three writing dimensions. We see that the commenting task did indeed help students in their own writing—generating explanations for others in this instructional setting were indeed instructional, for the author AND for the reviewer. In addition, only the Rate+Comment condition generated papers that were of higher quality than in the Write-First condition, suggesting that simply evaluating papers did not convey much writing benefit to the students.

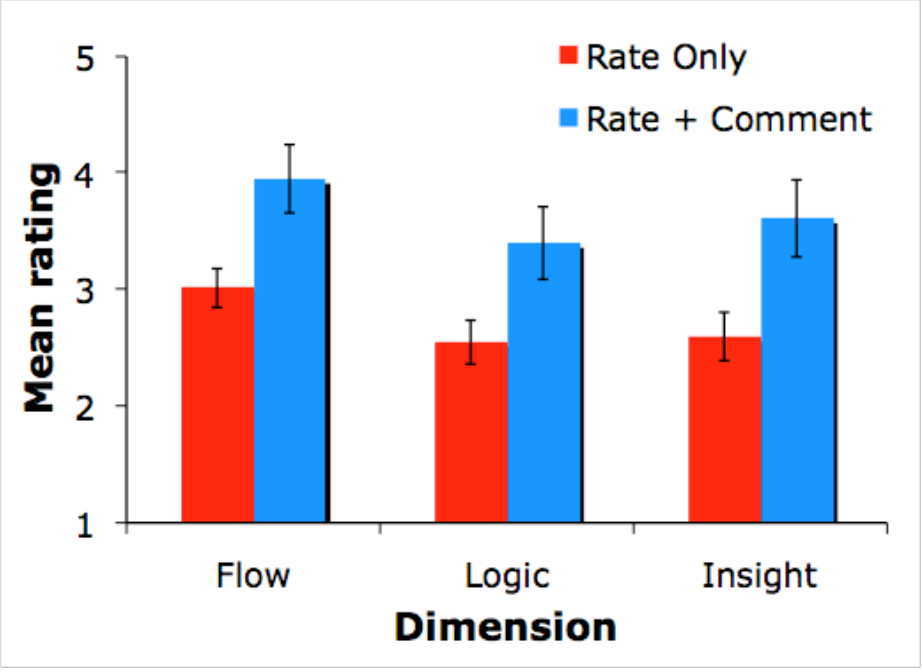

Figure 5. Mean 1st draft quality (and SE bars) as a function of condition.

These two studies suggest that the generation of instructional explanations is critical for learning, rather than just reading or evaluating examples. However, we do not mean to suggest that quality or features of the examples being evaluated play no role at all in student learning. We do not think that generating instructional explanations is a magical learning activity that conveys learning benefits regardless of content, as some kind of exercise for the mind. Rather we think that generating explanations conveys learning benefits because typically generating explanations for peers will produce explanations about problems that the students themselves are struggling with in their own writing. And thus, we predict that the quality of the examples serving as objects of review and commenting will influence what the reviewer takes away from the experience.

Indeed it is here that the instructor can play a critical role in shaping learning that happens in the peer review context: by influencing which papers each peer is asked to

review. Instructors often use examples to help students develop understandings on concepts and definitions. Consistently, much research on examples in mathematics and science reveals that examples may enhance knowledge acquisition and problem solving (Atkinson, Derry, Renkl, & Wortham, 2000; Sweller & Cooper, 1985). Examining examples may help learners to understand general rules and apply the learned rules to given problems whose structural features are similar to the examples (Anderson, Fincham, & Douglass, 1997). Leinhardt (2001) argued "The generation or selection of examples is a fundamental part of constructing a good explanation. But developing, recognizing, or selecting an appropriate example or counterexample is difficult" (p. 347).

Recently, Cho and Cho (2009) examined the role of instructional examples in improving the effect of participating in instructional explanation activities, focusing on the mean quality and diversity of quality of the examples in learning to write from reviewing. In reviewing, peer writing of low quality may have a different role for the reviewer's learning from that of high quality writing. In mathematics, explaining why correct solutions are correct and why incorrect solutions are incorrect was more effective in learning than explaining only why correct solutions are correct (Siegler, 2002). This prior work suggests that learning from low quality examples may be different from learning from high quality examples. An obvious difference between the two types of examples is that low quality examples include more errors than high quality examples. There are different perspectives on how errors in low quality examples influence learning.

High quality examples may be more beneficial in acquiring knowledge of high quality writing than low quality examples. Students can model their writing on the superior features of high quality examples, whereas errors in low quality examples can be harmful if students mimic the errors without caution. In addition, revising errors may impose a heavy

cognitive load on limited working memory (Paas, Renkl, & Sweller, 2004), which may interfere with abstracting writing principles from examples. In other words, students who review high quality examples can better concentrate on features of high quality writing than students who review low quality examples, which requires paying attention to detecting the many errors contained in the document. Thus, students may understand more about what writing should be and how to write well from high quality examples.

However, low quality examples can help students develop knowledge about various constraints on writing and apply the learned constraints to their own writing. Low quality examples can also provide more opportunities to detect, diagnose and revise errors than high quality examples. While revising errors, students may think deeply what errors are critical, why they are harmful, and how to fix them. For instance, Zimmerman and Kitsantas (2002) found that coping models that showed how to revise errors were more beneficial in acquiring writing skill than mastery models that did not include any error. Thus, low quality examples can help students to improve writing by reducing errors.

Whether low or high quality examples are better could further depend upon the writing abilities of the reviewer. For example, stronger writers may already know to avoid common errors shown by weaker writers, whereas weaker writers might need to practice detecting and repairing those more basic writing issues found in the writing of weaker writing peers.

Cho and Cho (2009) analyzed data from a class used SWoRD, placing students analytically into one of four conditions in a 2x2 design. First students were divided into poor and good writers on the basis of first draft scores. Crossed with that dimension was whether or not the papers the students reviewed were low or high in quality. Median splits were used to create these categories. The critical question concerns the average final draft paper scores

of students in each reviewing category: that is, who learned the most that could be applied to their own revision work?

Cho and Cho found an interaction between the two conditions (see Figure 6). Poor writers benefited most from seeing low quality examples. By contrast, good writers benefited equally from reviewing high and low quality drafts.

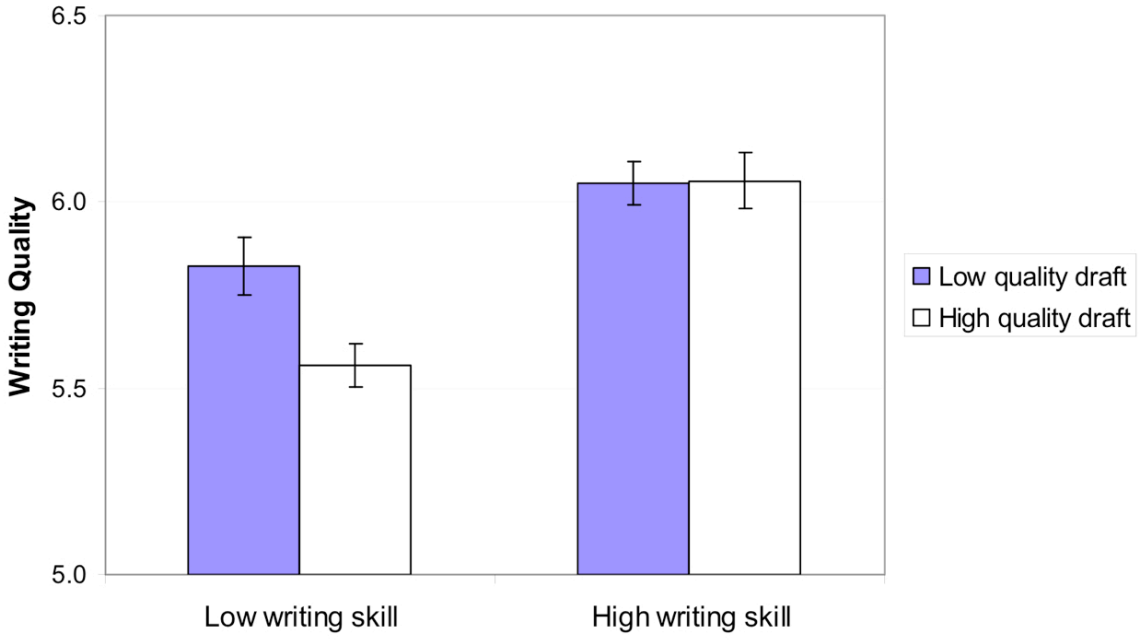

Figure 6. Mean final draft scores (and SE bars) as function of writer skill and whether the papers reviewed were of poor or good average quality.

4. Discussion

In this research, we examined the role of giving reviews as an example of providing instructional explanations. We investigated the gains for reviewers-as-writers rather than the more traditional focus of reviewers-as-surrogate-feedback source. Thus, this research speaks to the pedagogical value of providing the reviews in addition to its practical advantage of making rich feedback available more often to students. Although

this research did not focus on the nature of the comments being generated, it is worth noting that, like with peer tutoring, reviewers are expected to provide peer writers with coherent explanations or suggestions for improvement. By writing comments to others, therefore, reviewers may be more engaged in constructing a coherent understanding of writing as a result of developing these coherent explanations (Bargh & Schul, 1980).

Reviewing is a process of problem solving in which reviewers are engaged in exercising important skills for writing (Bereiter & Scardamalia, 1987; Fitzgerald, 1987; Flower et al., 1986) such as detection, diagnosis, and solution generation along with reading and commenting. These activities may improve reviewers' own writing and revising skills by reinforcing successful strategies and by calling attention to unsuccessful strategies that the reviewers have already used in their own writing.

A structured interview in a large undergraduate course about their experiences with the scaffolded peer review process support the findings. Here are typical responses from the undergraduate students to the question, "How did giving feedback help your own writing?" Students often mentioned they learned what they should not do in their own writing as well as what they should do. For example, a student said, "Well, I reviewed before I wrote so that definitely helped because I was able to see what other people were doing too and know what to expect and or what was expected and so that helped and then just again seeing, just more writing styles and different ideas." Another student expressed a consistent opinion by saying, "Yeah, I think it helped me in writing; it allows me to see how other people are writing and allows me to see their mistakes and that can only help me better write my papers…. It just shows me what not to do and what to do; gives me a better understanding of how to write a paper." Another student also

mentioned, “I think it helped me look out for things that I didn’t want to repeat; I didn’t want to waste other people’s time in effect, so that’s what I made sure. I really thought about what I had done to other people and told them what to do, not to use their mom as a reference, that’s just point blank common sense that some people just didn’t have that so.”

In addition, some students mentioned learning about taking an audience perspective to their own writing, especially by learning what others might want to see in writing. For example, a student said, “I guess, maybe, just looking at the flaws other people have in their paper I would make sure I didn’t do the same thing. Like, because I had to look for smooth transitions I, you know, I knew that’s what they’d be looking for in my paper so I made sure my transitions were ok. *So, I guess, by giving feedback on their papers I knew what they’d be looking at for mine*.” Also another student commented on the audience issue, “I think it gave me an idea of what other people look for, it helped me recognize what other problems there could be in writing. It also, I think, helped me learn how to write for a more general audience because I could see what other people think is important and that kind of thing. *It mostly helped me see how other people see writing*.” Another student expressed a consistent experience by saying “Oh yeah, *it gave me a better perspective of how my audience is going to be perceiving me* because I was the audience and tried to perceive how other people, saw what to look for, saw to look for things that usually when I’m writing I don’t catch. So I tried to reread my own writing as I was one of the students who was going to be grading it which I usually didn’t try to do before.”

Young and Leinhardt (1998) argued that disciplinary writing skills develop

mainly along two dimensions: the content and rhetoric of the discipline. Along the content dimension, students develop more detailed and integrated understandings of disciplinary knowledge. Along with the rhetorical dimension, students develop "disciplinary acts of argument and interpretation, evaluating and qualifying claims and evidence, and using rhetorical and linguistic conventions to support these acts of analysis and synthesis" (p. 56).

In this chapter, we argue that peer reviewing may provide students with valuable opportunities for understanding disciplinary writing skills. In the process of peer reviewing, students read, analyze, and assess peer papers and explain how to improve them in reference to the disciplinary rubric. Unlike trouble-shooting their own papers, peer reviewers are asked to construct instructional explanations that are coherent and accurate according to disciplinary rhetoric and content (Leinhardt, 2001).

The findings from the SWoRD research tentatively make a number of suggestions for classroom practice regarding instructional explanations. First, the SWoRD work suggests that having students engage in constructing *instructional* explanations can be useful for learning, rather than just the more typical student activities of self-explanations or explanations to the teacher. Second, the SWoRD work suggests that students can benefit from providing instructional explanations to both weak and strong examples. Further, there is some suggestion that the best examples to provide instructional explanations for may depend upon the skill level of the student.

References

Anderson, J. R., Fincham, J., & Douglass, S. (1997). The role of examples and rules in the acquisition of a cognitive skill. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 23*(4), 932-945.

Ashbaugh, H., Johnstone, K. M., & Warfield, T. D. (2002). Outcome assessment of a writing-skill improvement initiative: Results and methodological implications. *Issues in Accounting Education, 17*(2), 124-148.

Atkinson, R. K., Derry, S. J., Renkl, A., & Wortham, D. W. (2000). Learning from examples: Instructional principles from the worked examples research. *Review of Educational Research, 70*, 181–214.

Bargh, J. A., & Schul, Y. (1980). On the cognitive benefits of teaching. *Journal of Educational Psychology, 72*, 593-604.

Beach, R., & Friedrich, T. (2006). Response to writing. In C. A. MacArthur, S. Graham & J. Fitzgerald (Eds.), *Handbook of writing research*: Guilford.

Bereiter, C., & Scardamalia, M. (1987). *The psychology of written composition*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Bransford, J. D., & Schwartz, D. L. (2001). Rethinking transfer: A simple proposal with multiple implications. *Review of Research in Education, 24*, 61-100.

Charney, D., & Carlson, R. A. (1995). Learning to write in a genre: What student writers take from model texts. *Research in the Teaching of English, 29*, 88-125.

Cho, K., & Schunn, C. D. (2007). Scaffolded writing and rewriting in the discipline: A web-based reciprocal peer review system. *Computers & Education, 48*(3), 409-426.

Cho, K., Schunn, C. D., & Charney, D. (2006). Commenting on writing: Typology and perceived helpfulness of comments from novice peer reviewers and subject matter experts. *Written Communication, 23*(3), 260-294.

Cho, Y.-H., & Cho, K. (2009). The role of examples in learning to write in physics, *Annual Conference of the American Educational Research Association*. San Diego, CA.

Cornelius-White, J. (2007). Learner-centered teacher-student relationships are effective: A meta-analysis. *Review of Educational Research, 77*(1), 113-143.

Fitzgerald, J. (1987). Research on revision in writing. *Review of Educational Research, 57*(4), 481-506.

Flower, L., Hayes, J. R., Carey, L., Schriver, K., & Stratman, J. (1986). Detection, diagnosis, and the strategies of revision. *College Composition and Communication, 37*(1), 16-55.

Hayes, J. R., Flower, L., Schriver, K., Stratman, J., & Carey, L. (1987). Cognitive processes in revision. In S. Rosenberg (Ed.), *Advances in applied psycholinguistics vol. 2. Reading, writing and language processing*. Cambridge, England: Cambridge.

Hillocks, G. (1984). What works in teaching composition: A meta-analysis of experimental treament studies. *American Journal of Education, 93*, 133-170.

Leinhardt, G. (2001). Instructional explanations: A commonplace for teaching and location for contrast. In V. Richardson (Ed.), *Handbook of research on teaching* (4th ed., pp. 333-357). Washington, DC: AERA.

National Assessment of Educational Progress. (2002). *Writing report card for the nation and the states*: National Center for Education Statistics, U. S. Department of Education.

National Commission on Writing. (2003). *The neglected "r" the need for a writing revolution*: The College Board.

National Commission on Writing. (2004). *Writing: A ticket to work…or a ticket out: A survey of*
business leaders: The College Board.

Nelson, M. M., Melot, B., Stevens, C., & Schunn, C. D. (2008). *The effects of skill diversity in peer*
feedback: It's what. You don't know. Paper presented at the 30th Annual Meeting of the Cognitive Science Society.

Paas, F. G. W. C., Renkl, A., & Sweller, J. (2004). Cognitive load theory: Instructional implications of the interaction between information structures and cognitive architecture. *Instructional Science, 32*, 1–8.

Palincsar, A. S., & Brown, A. L. (1984). Reciprocal teaching of comprehension-fostering and comprehension-monitoring activities. *Cognition and Instruction, 1*(2), 117-175.

Patchan, M. M., Charney, D., & Schunn, C. D. (2009). A validation study of students' end comments: Comparing comments by students, a writing instructor, and a content instructor. *Journal of Writing Research, 1*(2), 124-152.

Rada, R., Michailidis, A., & Wang, W. (1994). Collaborative hypermedia in a classroom setting. *Journal of Educational Multimedia and Hypermedia, 3*, 21–36.

Siegler, R. S. (2002). Microgenetic studies of self-explanations. In N. Granott & J. Parziale (Eds.), *Microdevelopment: Transition processes in development and learning* (pp. 31-58). New York: Cambridge University.

Sweller, J., & Cooper, G. A. (1985). The use of worked examples as a substitute for problem solving in learning algebra. *Cognition and Instruction, 2*(1), 59-89.

Vygotsky, L. S. (1978). *Mind in society*. Cambridge: Harvard University Press.

Wooley, R., Was, C., Schunn, C., & Dalton, D. (2008). *The effects of feedback elaboration on the giver of feedback.* Paper presented at the 30th Annual Meeting of the Cognitive Science Society.

Young, K. M., & Leinhardt, G. (1998). Writing from primary documents: A way of knowing in history. *Written Communication, 15*, 25-68.

Zimmerman, B. J., & Kitsantas, A. (1999). Acquiring writing revision skill: Shifting from process to outcome self-regulatory goals. *Journal of Educational Psychology, 91*, 241-250.

Zimmerman, B. J., & Kitsantas, A. (2002). Acquiring writing revision and self-regulatory skill through observation and emulation. *Journal of Educational Psychology, 94*, 660-668.