THE JOURNAL OF THE LEARNING SCIENCES, *10*(3), 223–264

Knowing When You've Brought Them in: Scientific Genre Knowledge and Communities of Practice

D. Kevin O'Neill
Faculty of Education
Simon Fraser University

Increasingly, researchers in the learning sciences are appealing to notions of community to shape the design of learning technologies and curricular innovations. Many of these designs, including those in the area of project-based science, show strong promise; but, it is a challenging matter to understand the influences of these innovations in a detailed enough fashion to refine them over time. This work demands sensitive, theoretically grounded ways to assess the depth to which particular facets of innovations help enculturate students into communities of discourse and practice.

Taking genre theory and the sociology of science as points of departure, I demonstrate a unique approach to the problems of developing and assessing students' understanding of persuasive practices in the scientific community. The research I discuss revolves around students' use of a professional scientific genre of scientific writing, the Research Article or Introduction, Methods, Results, Discission (IMRD) report (Swales, 1990), as they compose reports about their own original research. Using data from an innovative project-based high school science class, I demonstrate how genre use provides a window on the effectiveness of a learning environment in helping use discipline-specific tools of persuasion.

In the classroom studied here, students developed e-mail mentoring relationships with volunteer scientists across the United States and Canada. Working in partnership with the teacher, these "telementors" served not only as inquiry guides for students, but also as a critical audience that helped shape the arguments they made about their research. Detailed analysis of the final reports produced by teams of students in the class revealed a significant relation between their fulfillment of the customary persuasive functions of a scientific research article and sustained correspondence with their telementors. A significant relation was also observed between sustained dialogue

Correspondence and requests for reprints should be sent to D. Kevin O'Neill, Faculty of Education, Simon Fraser University, 8888 University Drive, Burnaby, British Columbia, Canada V5A 1S6. E-mail: dko@acm.org

with telementors and careful hedging of knowledge claims. I situate these findings within a body of theory that suggests the value of telementoring relationships consists not only the ongoing advice and guidance they furnish, but in the ways that a professional audience shapes students' ideas about the sorts of arguments that are called for in science class.

Because the analysis of genre use is a relatively noninvasive way to examine students' understandings of scientific persuasion (as compared with survey instruments or pull-out interviews), this method can serve as a useful tool for reformers wishing to compare the outcomes from iterations or conditions of design experiments that aim to develop students' understanding of persuasive practices in the scientific community. It may also make a useful transfer measure for a wide range of classroom innovations that aim to cultivate scientific reasoning and persuasion, such as science-oriented tools for computer-supported collaborative learning.

> The page is no more than a score is to a Scarlatti sonata performed in a Santa Barbara living room … an archive mediating between an imagined event and a distant realization. To help people write more effectively we need to unpack the entire transaction and identify what the words are doing in the middle. (Bazerman 1988, pp. 9–10)

Increasingly, researchers in the learning sciences are appealing to notions of community to shape the design of classroom interventions and learning technologies (Brown & Campione, 1994; Pea & Gomez, 1992; Scardamalia & Bereiter, 1994). References to Lave's concept of legitimate peripheral participation (Lave & Wenger, 1991) and Vygotsky's concept of the zone of proximal development (Vygotsky, 1978) are rife in our professional literature. However, it is difficult to assess the depth to which particular innovations help enculturate students into the intellectual customs and norms of specific communities of discourse and practice. This matter is of genuine concern to learning scientists because it limits our ability to tune the design of classroom innovations over time.

Here, I offer a new approach to this problem, based on students' use of the textual genres produced and used within a professional community. As an illustration of the approach, I present an evaluation of a classroom design experiment involving "telementoring," or long-term online relationships between volunteer scientists and high school students conducting independent science projects. In particular, I test the hypothesis that students' involvement with their online mentors influences the ways in which they use the customary "tools of argument" in a genre of scientific writing. This genre is known as the Research Article or "IMRD" report.

TEXTUAL GENRES IN COMMUNITIES OF PRACTICE

Implicit in the idea of community are elements of stability and dynamism, continuity and change, shared interests, and ongoing contention of interests. Every community is held together by common interests and practices, but there is rarely perfect consensus over how to achieve common goals. For these and other reasons, participants in communities of discourse and practice find themselves arguing over similar issues again and again. Genre theorists (Bazerman, 1988; Miller, 1994) observed that over time these recurring rhetorical problems can give rise to customary forms of communication. For example, because academic researchers repeatedly find themselves competing for jobs and funding, they have developed the genres of the curriculum vitae, the letter of recommendation, and the grant proposal. These genres are easily recognizable to initiates in the field, and have such well-wrought expectations surrounding them that they form a constraint on acceptable practice. To write a curriculum vita in an uncustomary way is, therefore, to risk unemployment.

Genres are an important sociocognitive phenomenon because they embody the intellectual norms and values of the communities that produce and reproduce them. They give newcomers to a community a sense of its public priorities—its collective judgment of what is most relevant and important to its ambitions. From an individual's perspective, genres also structure the problems of persuasion associated with contributing to a community's discourse. As Bazerman (1988) put it,

> A genre is a socially recognized, repeated strategy for achieving similar goals in situations socially perceived as being similar. A genre provides a writer with a way of formulating responses in certain circumstances and a reader a way of recognizing the kind of message being transmitted. … Thus the formal features that are shared by the corpus of texts in a genre and by which we usually recognize a text's inclusion in a genre, are the linguistic/symbolic solution to a problem in social interaction. (p. 62)

In recent years there has been an explosion of scholarly interest in the nature and variety of textual genres, and what they reveal about the communities of discourse and practice that produce them (e.g., Bazerman, 1988; Cope & Kalantzis, 1993; Freedman, 1993). Much of this interest stems from a belief that a deeper understanding of how particular genres have developed would help educators to better design instructional approaches and materials. In particular, educators might be better equipped to foster discipline-specific literacies by "enculturating" students into disciplinary practice and communication, in the sense used by Perkins (1993).

This objective should be carefully distinguished from that of teaching students about the mere surface features of genres, such as the names and order of sections in a scientific research article. Genre should also be distinguished from writing style because it is not about writers' whimsy or use of ornament. Writing in a disciplinary genre entails the deliberate use of a community's customs to serve one's own goals of persuasion in a particular situation (Flower & Hayes, 1980). As I discuss at length, teaching students to use disciplinary genres has long been an objective of science educators, though not always a clearly stated one. More recently, tools for computer-supported collaborative learning in science (e.g., Bell, 1997; O'Neill & Gomez, 1994; Suthers, 1998) have continued in this tradition, attempting to use the affordances of new media to scaffold students' thinking in discipline-appropriate ways.

Whether writing takes place on a computer or on paper, a well-structured problem of persuasion is easier to solve than an ill-structured one. By structuring rhetorical problems, genres do a part of the writer's work. In this way, a genre is a brand of distributed intelligence (Pea, 1992) that represents the collective accomplishment of the community of discourse and practice that creates it. Because genre knowledge is a communal form of intellectual property, mastering it is an important part of joining a new intellectual community.

BLUFFING THEIR WAY INTO SCIENCE: STUDENTS USING PROFESSIONAL GENRES IN THE CLASSROOM

As researchers have previously illustrated, school and work settings develop their own unique genres of oral and written discourse (Berkenkotter & Huckin, 1995; Yates & Orlikowski, 1992). To deal with some of their recurring rhetorical problems, for instance, schools have developed the genres of the report card, the teacher's disciplinary note to parents, and the five-paragraph theme. This is natural enough, given that the school is a community in its own right and necessarily develops unique norms. Nevertheless, the school's mandate to help students understand and participate in communities of discourse and practice beyond its walls runs against this separation of genres. Helping students learn to write in the customary genres of adult work settings—from business memos in typing class to lab reports and research articles in science—has been a developing part of practice in schools since the early 1900s (Russell, 1991).

Under common K–12 teaching conditions (the lab report, science fair project), students are called on to do some difficult play-acting. Despite knowing little about the academic disciplines and not feeling invested in its discourses, they must attempt to produce pieces of writing that imitate those of initiates:

> The student has to appropriate (or be appropriated by) a specialized discourse, and he has to do this as though he were easily and comfortably one with his audience, as though he were a member of the academy. … He must learn to speak our language. Or he must dare to speak it or to carry off the bluff, since speaking and writing will most certainly be required long before the skill is "learned." (Bartholomae, 1985, p. 135)

The "bluffing" that Bartholomae refers to is a necessary stage in entering a discourse community, because language is only acquired through use (Bruner, 1990). It may be precisely because they know this that many of the students and teachers whom I have interviewed seem to feel that imitating professional scientific writing in the classroom makes sense. The kinds of sense it makes to them vary considerably, however. Figure 1 summarizes the responses of 12 high school students in a project-based science class to the question, “Why does your teacher have you write project reports in the particular format that he does?” This question was posed in an open-ended fashion during a series of focus groups, and the responses categorized post hoc. Some students offered more than one reason.

The activity in this particular project-based high school science class, which I discuss in greater detail later, was largely structured around the production of a scientific research article in a quasi-professional format (Polman, 2000). “Milestone” assignments submitted by each project team in the course of their work were designed by the teacher to correspond directly with sections of the required final report. Ongoing advice and feedback provided to project teams by both the teacher and volunteer scientist mentors (discussed extensively later) were also driven by these milestones. Thus, in this setting, there was stronger integration of the writing and research tasks than one might find in some other high school science classrooms. This is hinted at by the fact that an equal number of students cited guidance in their research as a reason for writing in an authentic genre as mentioned guidance in their writing.

Despite the relatively high frequency of these two answers, however, there was not complete consensus among the students about their teacher’s reasons for having them write in this genre. Although of less salience to them than guidance in research and writing, the 12 students mentioned a variety of other possible reasons for their teacher’s choice. Five students suggested that part of the purpose of writing in an authentic scientific genre was to help them accurately play the role of “real scientists” in the class. In this conception, the genre is an important prop in a collective game of academic make-believe, consistent with Bartholomae’s earlier characterization. What may be most surprising about the

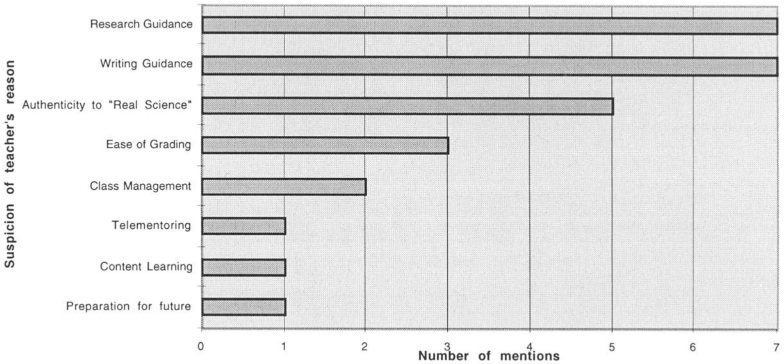

FIGURE 1 Students’ impressions of the purposes of imitating scientific genres.

responses of these dozen students is that so few of them felt the format of the report was chosen for pragmatic reasons, such as helping the teacher to grade the paper (by making the absence of valued elements more obvious) or helping him or her to manage the class by communicating his or her expectations more clearly. It appears that imitating "real scientists" was a sensible idea to these students.

How Students Encounter Scientific Genres and Scientific Thinking

Even if the imitation of professional science makes sense to students themselves, one may ask whether it is necessary for students to use the genres of professional science to appreciate or understand science. Here, I argue that the imitation of professional practice is to some extent inescapable, because it is implicit in the very notion of teaching science. The more important questions for researchers and educators, I believe, are what purposes this imitation is intended to serve, how well it serves them, and how well it might serve them under different conditions.

To get a firm grip on these questions, it is necessary to understand the variety of influences that students come under in their acquisition of scientific genre knowledge, and what messages these send about the practice of science. These influences include textbooks, encyclopedias, the style guides that may be prepared by school or school board personnel, and the handouts or class displays used by individual teachers. Most important, of course, are the pieces of scientific writing that students are asked to produce themselves. These include the presentation boards that students involved in science fairs commonly prepare (which usually have a strict prescribed format) and the lab reports they may be asked to prepare after a hands-on lab.

Sutton (1989) provided a comparison of guidelines provided to students for writing up labs as far back as 1898:

> Looking into the origins of this pattern of writing … it is very interesting to see the variation in the flexibility allowed, and in how much emphasis is placed on the preliminary statement of ideas. One extreme may be represented by C. B. Owen of Stowe School in his Methods for Science Masters (1956). He offered the mnemonic: High Powered Motors Often Crash, to trigger recall of the need for Heading, Picture, Method, Observations and Conclusion. … MacNair (1904) suggested: "The Object Aimed At," "What Was Done," "What Was Seen," and "What the Result Proved." … A. G. Hughes (1933) advocated the headings "Purpose," "Apparatus," "Observation," "Inference." He stressed the importance of discussion before practical work to clarify its purpose. (p. 139)

Each of these sets of guidelines reflects different ideas about the purposes behind students' imitation of scientific practice. In the course of his review, Sutton characterized two general classes of guidelines: those that depict science as a regimen of careful recording (Science as "Describing What Happens") and those that depict science as a regimen of withholding judgment until all the data are in (Science as "Data First and Theory Later"). Both classes of guidelines send particular messages to students about the nature of scientific practice, as indeed such guidelines do in the world of adult professional practice (e.g., Bazerman, 1988, chapter 9 discusses the messages of the *American Psychological Association* guidelines). In effect, these guidelines emphasize different problems of persuasion for authors and identify different genres through which solutions to those problems can be developed.

Today, educators continue to develop and use a variety of guidelines for science writing. For example, the poster illustrated in Figure 2 was observed in a middle school science classroom in 1996. In it, the reader will notice some similarities to the guidelines Sutton describes earlier, though I would argue that this poster presents a more inclusive view of scientific practice. Unlike most of Sutton's examples, it does not focus narrowly on the act of observation, and actually encourages

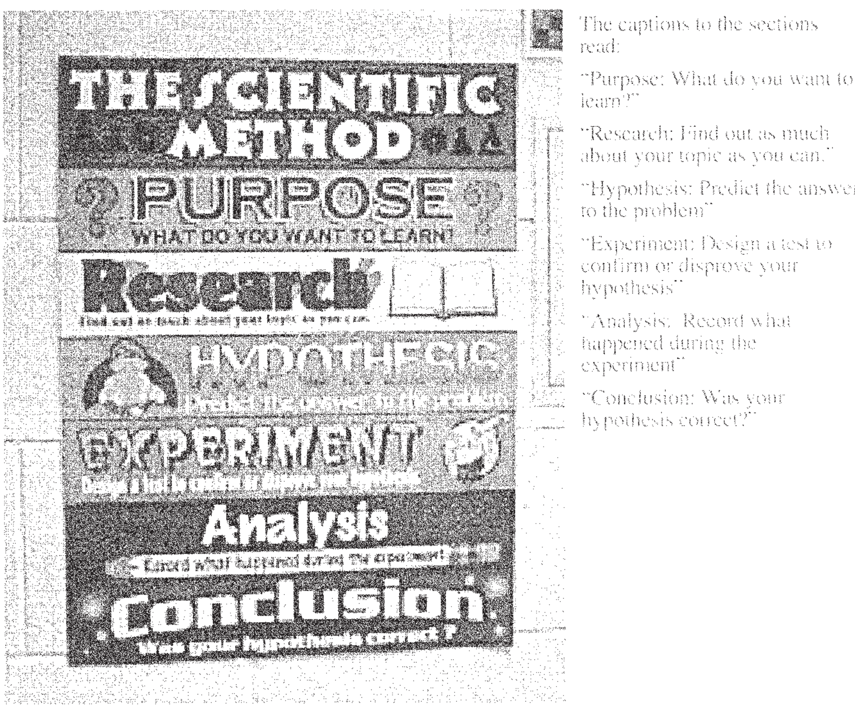

FIGURE 2 "Scientific Method" poster observed in a middle school science classroom.

students to generate hypotheses before the outcome of an experiment is known. It has its own limitations, however. Through its illustrations, the poster actually mystifies the process of hypothesis generation (which it pictures as a child gazing into a crystal ball) and encourages the idea that "research" is something bookish, done in the library alone. Finally, like Sutton's examples, this poster continues to give preferential place to experimental protocol in the development of scientific knowledge. In fact, a great deal of scientific practice does not involve much laboratory experimentation (e.g., Astronomy, Atmospheric Science, Botany, or Ecology).

I do not point out these limitations to be harsh, but to emphasize how difficult a job any such set of guidelines has to do. If we accept that school science must imitate professional science in some respects, we must ask which ones. Each of the representations of scientific research and reporting described earlier attempts one answer to this difficult question, reducing an extremely varied, complex, and large-scale set of practices to something small and simple enough to be useful for teachers and students in a host of local contexts.

Using and Misusing Scientific Genres

As I showed earlier, many of the genres we present to students for reporting their work are derivative of ones invented by and for participants in a radically different community of discourse and practice (that of professional science) in response the unique rhetorical problems that have recurred there over hundreds of years (Bazerman, 1988). We cannot expect to simply drop these genres into the classroom and have them fit the native activity there. In fact, professional genres often do not fit the native activity in classrooms. As an example, following is a Method section from a research report produced by a student in a project-based high school science classroom in the 1994–1995 school year. This project, on photochemical smog, was somewhat unique among high school projects in that the student was given several weeks to complete it and was required, in that time, to come up with a question that could be addressed with numerical data. The following excerpt, from a students' article on smog, is indicative of the misuse of scientific genres that takes place in this and other classrooms:

> Method: My problem with this topic was that *all I found was the temperature and precipitation data.* I sat at Mosaic and Netscape for hours just cruising through the information endlessly. I even tried Lycos and all of the other searching mechanisms to find *the rate of photochemical smog. Nobody had it.* This time period was quite frustrating. Finally, I posted on a newsgroup. For awhile, I did not hear anything, but finally a very nice person wrote me back. A man on the California Air Resources Board sent me quite a bit of information. As a result, *I had to change my topic. I decided to try and find a*

> *correlation between the precipitation and temperature and ozone statistics between 1970 and 1979.* That is when I could get down to business. [italics added]

The reader may ask what is wrong with this passage. Nothing, I would answer, unless the goal of having the student write it was to teach a lesson about the role a Method section plays in a research article. The italicized portions are those that appear to serve part of the customary function of a Method section. The first two italicized parts explain a practical constraint on the investigation: The desired data could not be located. The next part explains a strategic response to this problem: The question is changed to fit the available data. The balance of the section, however, is simply an adventure story, told by the student to the teacher about the difficulty of completing the assignment. This story makes appeals to considerations traditionally associated not with evaluating knowledge claims, but with grading: the student's use of all of the resources made available by the teacher, the long hours dedicated the project, and resulting frustration. The student's message is clear: "I worked very hard, I used what you gave me, and I'm fed up. Any weaknesses in this paper are not my fault." What little the student says about the strategic decisions she made in her investigation is almost lost amid this argument over grades.

I do not intend to belittle this student's concerns. This kind of frank communication has a legitimate place in the relationship between student and teacher. Unfortunately, it does not help students to understand the kinds of argumentation peculiar to scientific discourse. In fact, it is arguable that the kind of writing quoted earlier can make the imitation of scientific genres in the classroom worse than useless. Having students routinely produce papers that center on the grades they feel they deserve rather than the knowledge claims justified by their work puts educators at risk of teaching and reinforcing a caricature of scientific writing and research. The problem for learning scientists is how to practically change the teaching–learning context so that the focus of students' writing is on knowledge claims, rather than grades.

UNDERSTANDING GENRE MISUSE: STUDENTS' RHETORICAL SITUATIONS

> The situation controls the rhetorical response in the same sense that the question controls the answer and the problem controls the solution. Not the rhetor and not persuasive intent, but the situation is the source and ground of rhetorical activity. (Bitzer, 1968, p. 6)

The research discussed here was driven by the suspicion, shared by my teacher collaborators, that the misuse of scientific genres illustrated earlier stems largely from the structure of the situations in which students are asked to undertake scientific research and writing. In particular, we suspected that the audience (or lack of it) that students were writing for distorted the ways in which they used the disciplinary genres to which they were exposed.

To help explain this suspicion, I present Figures 3 and 4, which illustrate some of the important differences between the situations presented by school science and professional science. Figure 3, from Berkenkotter and Huckin (1995, p. 62), depicts the role of research articles in the ongoing professional practice of the sciences. To secure funding to carry out research ("Lab Activity," center top), a researcher must be able to cite previous accomplishments in grant proposals ("Citations," bottom left). But these accomplishments are largely recognized through the process of peer review ("Manuscript Under Consideration," middle right); thus, published articles are a critically important form of currency in the social and professional credit system of the sciences—what Latour and Woolgar (1979) called the "cycle of credit."

Although scientists may be judged by their peers on a myriad of both professional and personal criteria, writing for a critical audience clearly has a privileged influence on a researcher's professional fortunes in science (Myers, 1990). Contrast this situation with a student's situation at school. Figure 4 depicts my own impression of the place of paper writing in the life of a student. Here there are many more ways to gain credit than by writing. A significant portion of a student's

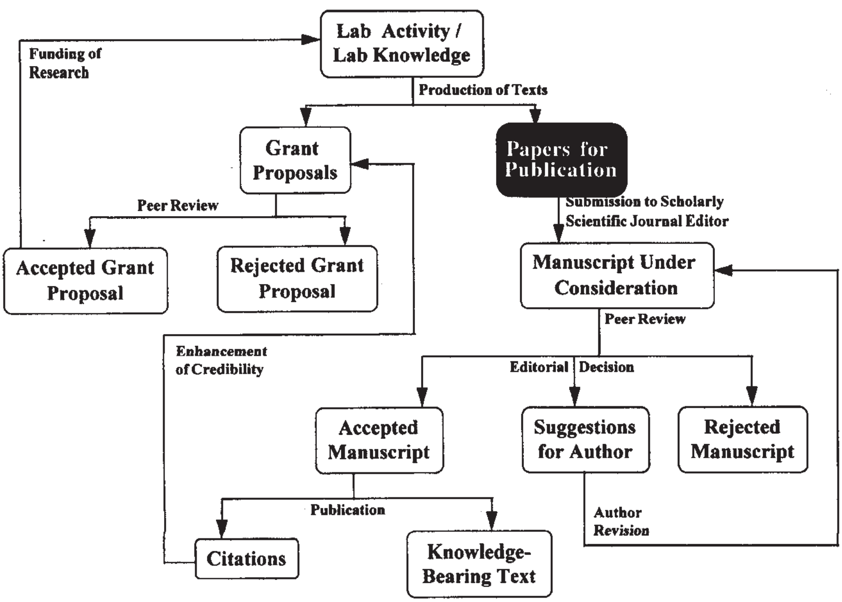

FIGURE 3 Place of scientific research articles in professional practice [shading added]. From *Genre Knowledge in Disciplinary Communications: Cognition/Culture/Power* (p. 62), by C. Berkenkotter and T. N. Huckin, 1996, Hillsdale, NJ: Lawrence Erlbaum Associates, Inc. Copyright 1996 by Lawrence Erlbaum Associates. Used with permission.

grades may actually derive from the teacher's first-hand observations of how he or she behaves in class: attendance, on-task behavior, contribution to discussions, cooperation with classmates, and so on. This variety of credit mechanisms necessarily makes writing less important for the student than for the adult professional, whose stock of credit can go up or down tremendously from writing (or lack of it) alone. And, when students do write, they are frequently assessed on the accuracy with which they can summarize or the work of others, rather than on the competence with which they can state and defend their own knowledge claims.

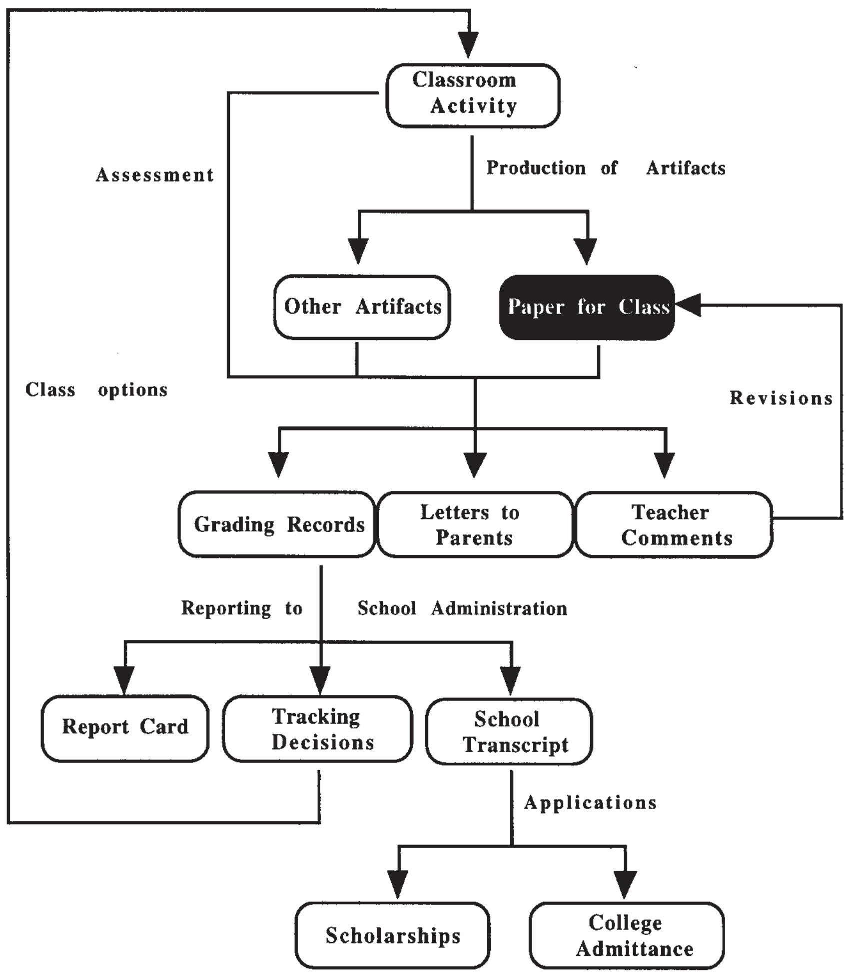

FIGURE 4 Place of writing in school life.

Finally, and most important for my argument here, there is usually only one formal giver of credit for writing in a student's life—the teacher. Because students write in such different "rhetorical situations" (Bitzer, 1968) from scientists, we should not be surprised if they misuse scientific genres, and misunderstand the nature of scientific research.

Teaching students about how knowledge claims are credentialed in the scientific community means far more than teaching them about logic, or about the formal features (e.g., heading structure) of professional genres alone. Such trivial conceptions of what it means to teach scientific persuasion are, to borrow an image from Russell (1991), as absurd as teaching someone the rules of movement for chess pieces without teaching them the objective of the game. I would claim that in its deepest conception, teaching students about making written arguments in science should mean providing them with tasks and situations in which they can faithfully use the tools of scientific argument to pursue their own persuasive goals in writing. To the greatest practical extent, those situations and those goals should be a match for those of adult scientists.

Providing a Critical Audience Through Telementoring

It is well enough to assert this, of course; but how does one do it? One way that may be practical on a large scale is a combination of project-based science (Krajcik, Blumenfeld, Marx & Soloway, 1994; Ruopp, Gal, Drayton, & Pfister, 1993) and telementoring (O'Neill & Gomez; 1998; O'Neill, Wagner, & Gomez, 1996) that I have explored with teachers involved in the CoVis project (Pea, 1993). As I describe in greater detail later, this approach involves giving students greater agency to determine not only what they study, but how they study it. Then, to help them manage this responsibility, the teacher arranges for students to get ongoing advice and guidance from knowledgeable adult volunteers on a routine basis, via the Internet. Because they offer a wide range of advice and guidance from a distance, they are referred to as telementors.

As I discuss later, telementors turn out to be important not only as sources of knowledgeable advice, but also as a responsive critical audience for students' work. Genre theory and sociocultural theory (e.g., Wertsch, 1991) would both suggest that routine and purposeful discourse with a critical audience that is not involved in grading students' work has the potential to alter their understandings of scientific reporting in positive ways. If students have the chance to internalize a dialogue that concerns itself largely with knowledge claims, it might change the rhetorical situations (Bitzer, 1968) in which they understand themselves to be working. By providing a context in which the dialogic nature of students' writing is shaped by a wider audience than their teacher alone, telementoring should enable students to appropriate the genres of scientific writing more authentically. I call this the “rhetorical situation hypothesis,” and set out to test it later.

ILLUSTRATION: TELEMENTORING AND GENRE USE

Research Setting

The data discussed following were collected in a project-based Earth Science classroom at a suburban Chicago high school. While the school and this class are attended primarily by White students from affluent families, most of the students under study were not traditionally high-achieving and might be characterized as “science avoiders” seeking to fulfill a minimum science requirement for graduation. The teacher, Rory Wagner,[1] is a veteran of the CoVis project who began developing his project-based teaching style in 1992. He is now at an advanced stage in the development of his project-based teaching style, and dedicates three contiguous quarters of the school year to students’ independent project work.

In the first academic quarter of each year, Wagner gives his students a grounding in the phenomena they might choose to investigate in the remainder of the year, mostly through traditional lectures and demonstrations. This quarter ends with an open-book content test, which is mentioned again in later sections. From that point forward, students are evaluated largely on their performance on project work, which they conduct in self-selected teams of between one and three students. Each team pursuing a research agenda of its own design and the teams have tremendous freedom in this regard. When asked to describe what phenomena his students are permitted to do research on, Wagner responds, “anything that isn’t living—that’s Biology.”

Each team project lasts for roughly 7 weeks, and the only strict requirements are that students pose a clear research question that they can address with some form of numerical data analysis, and submit a quasi-professional research report on their work. Wagner provides very little class-wide training in data analysis or scientific argumentation, choosing instead to do this on a team-by-team basis as needed. As students’ project work progresses, they submit portions of their final research articles as “milestones.” Each of these is carefully assessed and followed up on in class. For example, each team’s research questions and the data they plan to analyze are submitted weeks before the final report comes due, at which time Wagner offers additional support if they appear weak.

---

[1]His real name, used here with permission. All other names mentioned in the text, for students and volunteer mentors, are pseudonyms.

Wagner has some help in providing this support. One of the most unusual aspects of the students' working conditions is that each team is assigned a volunteer on-line mentor, or telementor. These volunteers are recruited by Mr. Wagner via the Usenet news ".sci" hierarchy, and other venues, to advise and assist students' research through exchanging periodic e-mail. Most are graduate students or professionals in the geosciences, such as environmental engineers. Each volunteer is assigned to work with one specific team of students, after the students have declared their research agendas for the project. Although some of the volunteer mentors have teaching experience, they are not specially trained. In fact, their orientation to telementoring consists entirely of a series of e-mail messages describing the nature of the class the students are enrolled in, the work they are expected to do, and the kinds of support they are likely to need.

The volunteer mentors have, on many occasions, been of great help to students in locating data sources for their research and/or suggesting manageable lines of investigation; but equally important, they provide a critical sounding-board for students' ideas. In Wagner's words, the purpose of getting his students involved with scientist mentors is,

> that the kids can see how scientists think, how they work. Not only get the information, because they're going to get that—because they could get that from anybody who's knowledgeable—but also the process of doing it. … So I'm just asking the mentors to get kids to think like scientists think in dealing with them.

When Wagner and I began our work together, we hoped that purposeful discussion between the research teams and their mentors would enable students to internalize critical voices in a Vygotskyan sense, and produce more thoughtful and discipline-appropriate writing. Volunteers were made aware that their mentees were to produce written reports of their work, and sometimes saw these reports; but they were not complicit in Wagner's agenda to teach students about genres of scientific reporting, and were not instructed to coach particular argument strategies. Given the wide variation in the design of the students' projects, we viewed this as impractical.

Telementoring Relationships

Because the telementoring relationships in Wagner's class did not follow a rigid protocol, no two were alike; however, they were governed by common dynamics that are discussed in detail elsewhere (O'Neill, 1998; O'Neill & Gomez, 1998). For the sake of illustration, I provide one detailed example of how mentors served as a supportive and critical audience for students' research. This is followed by a broader characterization of the mentor–mentee discourse across the entire class.

An Example of Research Guidance in a Telementoring Relationship

One case from the 1995–1996 school year revolved around Dan, a PhD student in Physics, and three students (Andy, Cori, and Bill) pursuing a project on how astronomers identify black holes. This relationship began with a brief exchange in which the students told Dan a little about their idea for the project, and he probed them about what relevant course work they had completed in school. From this point, the conversation turned directly to the design of an investigation that would enable the students to pursue their curiosity about black holes, while satisfying Mr. Wagner's stringent requirement for a nontrivial empirical claim in the final report. Following, the students use the word "infomercial" to describe the sorts of projects that Mr. Wagner routinely rejects [italics added]:

> Dear Dan Jeffries,
>
> I'm glad to know you'd be interested in helping us. I hope you got our information on our educational past. We need to come up with a thesis proposal. We submitted several proposals that weren't accepted because they didn't fully meet the thesis requirements. Here's the bind: *We must have a topic question that will not turn our project into an infomercial. However, we need specific data (however much data on Black Holes is complex Physics that we can't use/understand).* We're thinking of the following type of project proposal. Finding several reports on Black Holes that may exist. Using the data we know (in simplified terms) we evaluate the data known on this these supposed Black Holes. We then conclude whether any of these can truly be black holes (dependent on whether they meet our "requirements for black holes"). This case study will require transferring a great deal of complex material into simplified, workable terms. Perhaps this is an area that you may be helpful in. Please write us and tell us if such information we're looking for exists and is workable, or if you have any ideas for our project. Your help is greatly appreciated. Keep in touch.
>
> Thank You,
> Andy, Cori, and Bill

In his reply to this message, Dan offers advice on three fronts. In the first paragraph of the next message, he points Andy, Cori, and Bill to publications where they are likely to find reading materials on black holes that they will be able to understand with their limited theoretical knowledge. In the second paragraph, he cautiously suggests that with the right limitations, their black holes agenda could satisfy Mr. Wagner's project criteria. Finally, he reminds them that it is not too late to choose something simpler to do, and suggests that they may find some inspiration in a publication *called Sky and Telescope:*

> Andy et al.
>
> I don’t think I gave you the month on that *Physics Today* article, I’m not sure of the month but it’s number 8. I would assume that’s August but I’m not sure if *PhT* puts out 12 issues a year. Once again, I’d look at *Sky and Telescope* or *Astronomy* first. I talked it over with a few of my colleagues, and they suggested that *Physics Today* might be too advanced.
>
> As I understand it, your problem is that you cannot just say “This is what a black hole is, isn’t it cool,” but you have to come up with some verifiable, answerable question. The process used by astronomers to determine if a source is a black hole isn’t all that trivial. If you do not have to go into how they measure the velocity of the surrounding material, that might be do-able. Still, there are only three sources that have been conclusively identified as black holes. (And one of them was just announced days ago.)
>
> In addition to looking at the idea of how the existence of black holes is determined, you might also ask yourself if there are any other astronomical questions that you are interested in. You might look through a few back issues of *Sky and Telescope,* and see what interests you.
>
> Let me know what you think.
>
> Dan

Later in the exchange, when it became clear to Dan that his young mentees were not going to give up on the black holes idea, he began to problematize for them the nature of the “proof” they should be looking for with a series of heuristic questions:

> The basic questions you would want to ask are: What is a black hole? If it exists, where would we expect to look for it? What would we see? Has anyone looked for this, and what did they find?

This is an example of the type of discourse that is likely to influence some of the ways that students approach tasks of persuasion in a research article. Dan’s questions make clear that establishing knowledge claims in the domain the students have chosen is hard. His questions also encourage the students to do particular things to address this challenge, including conducting targeted background research (“where would we expect to look for a black hole?”), describing their hypotheses carefully (“what would we see?”), and hedging their knowledge claims carefully. As we see later, each of these facets of scientific persuasion is targeted in the Science Argument Strategies coding scheme.

With Dan’s suggestions and further consultation with their teacher, Andy, Cori, and Bill finally managed to settle on the following plan for their data gathering and analysis:

Thu 7 Dec 1995

Dear Dan,

Thanks a great deal for your continual responses and input. Your help is much appreciated. We think we've found a nifty idea for our project.

First we'll briefly explain the features and dynamics of a Black Hole, talking about simple physics, formation of, etc. Next we'll research the three known Black Holes and find information about how and why these areas were positively identified as a black hole. Next we'll do a "case study" on the areas that scientists think may be Black Holes. By comparing observations between the Black Holes and the "possible Black Holes" we can conclude which of these "possible Black Holes" are most likely to exist.

Some of the following information may help.

1. What are the names of the three Known Black Holes, where might information be found on them (we'll find it).

2. Where might information be found on the unknown Black Holes.

3. What "traits" in these Black Holes and supposed Black Holes would be the most simplistic and beneficial to helping us compare.

Once again Thank You Greatly,

Andy, Cori, and Bill

Following this message was a lengthy exchange about the theory surrounding black holes, and how astronomers collect and interpret data about them. Much of this text shows Dan translating what he knows about the phenomenon of black holes into simpler language for his mentees. When the students completed their report, Dan also provided extensive, paragraph-by-paragraph comments on the work. Following, for example, is one of Dan's minute reactions to the team's Data Analysis section, included here in italics after his quote from the students' paper:

> Cygnus X–1 however cannot be conclusively proven a black hole. Few common traits could be found between Cygnus X–1 and NGC4261/M 87. Although thought for years to be a black hole because it's density outmatched a neutron star's maximum capability, it clearly is not. However, Cygnus's density may have been miscalculated slightly. Assuming Cygnus X–1 is actually 20% less dense (this is quite possible given it's distance), then it is low enough in density to be a neutron star yet too large to be classified as anything else. Therefore, Cygnus is most likely not a black hole, but a neutron star instead.
>
> *Given its distance, or uncertainty in the determination of its distance?*

Along with these paragraph-by-paragraph comments, Dan provided summary comments on the entire paper that emphasized how well the students limited and defended their knowledge claims:

> Interesting paper. You did a good job of making what conclusions you can from a limited data set. This is a must for anyone in astronomy. In the field, because there is almost always limited data, error analysis is crucial, as well as a healthy sense of scepticism. I think you did a good job of showing why the nature of Cygnus X–1 is in doubt. More information on what you would expect to observe from typical neutron stars would help to strengthen your claim that the Cygnus X–1 source is a neutron star. Good Work! I hope you had fun with this.

Returning to my argument about the rhetorical situations in which we place students, I want to make clear that the educational value of telementoring relationships like this one does not reside chiefly in students tapping the authoritative knowledge of a "real" scientist. This is valuable, but they could do it more easily through a Web-based "ask-an-expert" service. What is most valuable with regard to the educational goals at issue here is the relationship that Andy, Cori, and Bill formed with an audience that they respected, that demanded disciplinary rigor in their work, and that supported them in achieving it. When students build a relationship with this sort of audience, they can better understand disciplinary concerns and feel more responsible to serve them in their work. Andy articulated this sense of responsibility well when I asked how he felt about sending his paper to Dan:

> It can be … scary … because I'm sending a pretty amateur paper to a physics professor, asking him to put comments on it! I don't know, it might be quite laughable to him. But on another level, [the mentors are] probably anxious to see your paper. … You know, they put some element of work into this whole mentoring deal, they want to see what the student has got out of it, and how successful their mentoring was.

A Broader View of the Content of Telementoring Dialogues

As the relationship between Dan, Andy, Cori, and Bill would suggest, the telementoring relationships in Wagner's class were often friendly, but they also were functional. This is due, in part, to the fact that Wagner orchestrated the relationships explicitly to facilitate project work, and carefully matched students' research interests with their mentors' expertise. In the words of one student, this allowed research teams to get "straight down to business" with their mentors.

Although Dan's case is typical in some respects, no single example can adequately illustrate the range and diversity of telementoring relationships. To provide a wider view of what telementoring discourse was like in Wagner's class, Table 1 shows some of the results from a coding of the dialogue "moves" observed

TABLE 1

Proportion of Total Message Text Occupied by Dialogue Moves in 26 Mentor-Student E-Mail Exchanges

Mentors			*Students*			
Topic of Dialogue	*Words*	*%*	*Topic of Dialogue*	*Words*	*%*	*Description*
Conceptual			*Conceptual*			
Project ideas	1356	11	Project ideas	2795	23	Discussion of possible directions for an investigation
Domain concepts	3883	32	Domain concepts	955	8	Discussion of domain phenomena, terms, methods, practices or standards of evidence related to the investigation
Questions	706	5.8	Questions	1004	8.3	Direct question of any kind
Procedural/Advice			*Procedural/Advice*			
Requests	731	6	Requests	2383	20	A direct request for information or advice
Project status	379	3.1	Project status	1663	14	Discussion of the progress made on the investigation, challenges that have come up, etc.
Hints	567	4.7	Hints	0	0	Indirect suggestion of a course of action or inquiry
Suggestions	1284	10.6	Suggestions	0	0	Direct suggestion of a course of action or inquiry
Offers	613	5.1	Offers	0	0	Offer of further aid, information, etc.
Advice	1432	12	Advice	0	0	Recommendation for action
Relationship			*Relationship*			
Hello/goodbye	542	4.5	Hello/goodbye	1886	16	Students and mentors introducing themselves to one another or drawing their relationships to a close
Thanks	0	0	Thanks	837	7	Expression of gratitude for past action
Complaints	183	3.5	Complaints	0	0	Complaints about past behavior (e.g., lack of communication, vagueness of request or advice, etc.)
Apologies	431	3.5	Apologies	320	2.7	Apology for past behavior
Compliments	230	1.9	Compliments	71	.6	Recognition of work well done

Total words coded as moves: 12,016

in transcripts of the e-mail between students and their mentors during the final project cycle of 1995–1996.[2]

As this overview shows, mentors wrote approximately 1.8 words of telementoring dialogue for each word written by students (including message headers, etc.). This ratio is not surprising, given that mentors normally had more convenient access to e-mail, and probably wrote much more quickly than their mentees as well. What is most important, of course, is how this text was used. Most of it was taken up not in friendly chat (though that occurred), but in seeking or offering project-related advice and guidance of the kinds shown in the sample relationship earlier. A smaller proportion of the message text was dedicated to relationship maintenance moves, such as hellos, goodbyes, thanks, and apologies. More will be said in a later section about the content of the telementoring dialogues and the relation this bears to the scientific argument strategies that students use.

Research Design

Having briefly discussed the nature of the dialogue that takes place between student research teams and their telementors, we can begin to focus on the question of whether this dialogue had any apparent influence over students' use of the tools of argument in a disciplinary genre. I take the latter as a meaningful barometer of the disciplinary "enculturation" that project-based teaching and telementoring are intended to provide.

Science Argument Strategies Coding

The main coding instrument for this study, the Science Argument Strategies (SAS) coding scheme, was designed to capture differences in students' use of the customary tools of argument in research articles. The version of the coding scheme used in this analysis (see Appendix) represents a synthesis of observations about the research article genre from literature in the sociology of science (Bazerman, 1988; Myers, 1990) and linguistics (Swales, 1990) with observations made through close reading of a corpus of roughly 150 research articles written by middle and high school students in two project science classrooms. In addition to Wagner's class, reports were obtained from a middle school class in an inner-city setting. Using these data, and in response to feedback from collaborating teachers and colleagues, the coding scheme went through several stages of refinement over the course of approximately 3 years.

It is important to note that despite deriving its focus from studies of professional scientific writing, the development of this coding scheme was largely inductive.

[2]All coding was done by myself alone. Percentages sum to more than 100 because dialogue moves often overlapped in the text.

That is, no behaviors are included in the scheme that were not observed in samples of actual students' writing, or that are not logical complements to observed behaviors. Thus, the coding scheme does not unfairly apply the standards of professional science to the work of students.

The portions of the scheme discussed here code for customary argument moves in research articles, and the use of hedges and qualifications to strategically soften claims. Using the SAS scheme, I coded a sample of 31 reports written by teams of students in Wagner's Earth Science class over the course of three project cycles in the 1995–1996 school year. (Many more papers were coded in the course of developing the instrument, but the results are not discussed here because matching data on the writers' telementoring experiences were not available.) The time to code each report varied between 15 and 45 min, depending on the length and clarity of the text. To the extent possible, coding was conducted blind of students' telementoring experiences. The results of students' surveys and focus groups about telementoring were not consulted during the coding process, and although in a few cases students' reports specifically cited their telementors, they did not describe the depth of telementoring relationships in sufficient detail to bias the coding.

Hypotheses

In analyzing the coding of students' papers, two rival hypotheses were considered. The first hypothesis, mentioned earlier, follows from the theoretically driven suspicion that the lack of an audience concerned chiefly with knowledge claims distorts the opportunities that projects provide for students to learn why science is reported as it is. If this suspicion is true, one would expect that teams who invested greater effort in engaging their mentors as a critical audience would be more apt to write about what they thought they knew on the basis of their research than about the effort they invested or the grades they deserved. Their efforts to sustain dialogue with practitioners of science while carrying out their projects would enable them to build more sophisticated models of the rhetorical situation (Bitzer, 1968) surrounding the reporting task, and this should be evident in their writing. This is the "rhetorical situation" hypothesis.

A second possibility, equally important to consider, is that telementoring would have a relatively small influence on students' approach to the reporting task. Under this hypothesis, any variation in the sophistication of students' arguments would more likely be a function of general academic ability, or interest in science. Under this hypothesis, traditional high achievers would be expected to write the best reports, regardless of the depth of their telementoring relationships. Perhaps they would have both sophisticated arguments and sustained mentoring relationships, but only because they were generally "good students."

Rather than disrupting the unique work being carried out in the research setting to establish artificial conditions for hypothesis testing, this research followed the model of design experiments (Hawkins & Collins, in press). The correlational analysis presented later focuses on exploiting the natural variation in classroom activity and its outcomes to explore the relations between interacting variables. In line with the hypotheses stated earlier, the following analysis explored the relations between the genre use measures and

- The volume of correspondence between the team and its mentor. Is there a reliable relation between the argument strategies that a student research team sees fit to use and the frequency with which they corresponded with their mentors?
- Team members' average and maximum grades on a typical Earth Science content test. Are the students who use scientific genres in an authentic manner the same ones who perform well on a traditional academic task?
- Volunteer scientists' holistic ratings of the quality of students' reports. Does the Argument Strategies coding appear to reflect what practitioners of science value in research articles?

The relations among these variables and the sophistication of the argument strategies used by students were assessed using statistical correlation. Because the data from the argument coding are noncontinuous and the sample sizes concerned are sometimes small, Spearman's rank correlation (ρ) was used. This minimized the influence of outliers on the levels of significance reported.

Content Test Scores

The content test scores used here to test the "good students" hypothesis were from an open-book Earth Science content test, administered in the first quarter of 1995–1996 as part of students' routine work for the course. This test contained a variety of long-answer questions related to both conceptual knowledge in Earth Science (e.g., "What is the big bang, and what happened then?") and concepts related to research methods (e.g., "What is the difference between data and information?"). All of the completed tests were graded by the teacher, who was not aware at the time that the grades would be used in this research. Here, I take the test scores as a reasonable indicator of students' academic ability and commitment to success at the outset of the course.

Because the test scores are for individual students (whereas the research articles were composed by teams), the analysis discussed later inspected the relations between genre use and both the mean test score for teach team and the high score for each team. High scores were considered to address the possibility that the highest scoring member of a team had "carried" his or her partners through the composition of the research article.

Scientists' Holistic Ratings of Student Reports

To test the consistency of the Science Argument Strategies coding with the standards of practicing scientists, holistic ratings of a sample of the 31 previously coded papers were solicited from a small group of graduate students, professionals, and university faculty in the sciences who had previously volunteered to serve as mentors. None of the volunteers were friends or associates of the author. Each was asked to rate students' research articles with respect to (a) their overall quality, (b) the quality of the research discussed in the paper, and (c) the quality of the arguments presented about the research.

Each volunteer used a prepared form to rate two sets of three student articles. First, as an introduction, each volunteer read and rated an identical set of three papers chosen to familiarize them with the range of quality in the students' work. The three papers in this training set were chosen to represent the range of scores received by the 31 papers under the IMRD dimension of the Science Argument Strategies scheme (see later).

After rating the three articles in the training set, each volunteer coded 31 unique articles chosen at random from the 31 already coded. This resulted in 18 sets of holistic ratings for 18 unique articles, in addition to the 3 in the training set. All articles in both the training set and the unique set were presented in a uniform word-processing format to minimize the influence of surface attributes (such as page layout or font size and style) on raters' judgments. Spelling, grammar, and paragraph structure were, however, left unchanged from the original submissions.

Quantity of E-Mail Exchanged by Students and Mentors

Direct logs of the e-mail correspondence between Wagner's students and their mentors were available only for the final project cycle of the 1995–1996 school year. However, 17 of the project reports analyzed later were composed in the first and second project cycles of that year. For the analysis conducted here, the quantity of correspondence between student teams and their mentors was estimated on the basis of survey data.

In a brief survey administered at the completion of each project, students filled a six-cell table with estimates of the number of messages their team sent to and received from their telementors at the beginning, middle, and end of the project. The variable used in the following analysis is the mean of the estimates of message traffic provided by the members of each team. Students' estimates were generally quite consistent: Just 5 of the 31 teams' estimates (8.3%) had standard deviations higher than 3.5 messages. Figure 5 shows the distribution of estimated message traffic for the 31 project teams over the 1995–1996 school year.

As Figure 10 makes clear, the volume of correspondence between student teams and their mentors varied widely, from no correspondence at all over the

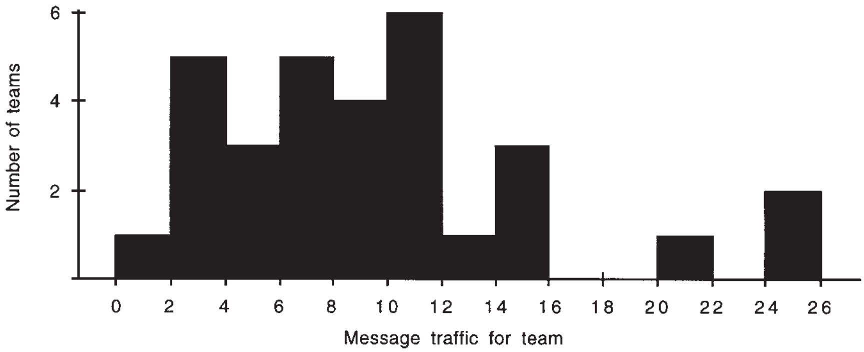

FIGURE 5 Students' estimates of total message traffic with mentors.

course of a project in one case to virtually one message per day in two others. The majority of teams exchanged between 4 and 16 messages with their mentors. Although this number may seem small to some readers, it can be put in perspective by mentioning that one of the richest relations that occurred in this class (described at length in O'Neill, 1998) involved the exchange of just 15 messages over a period of 9 weeks.

Given this observation, one might wonder what the number of messages exchanged by student teams and their mentors means. This issue is worth considering carefully. Two important possible influences on the frequency of students' correspondence with their mentors are grading and technology access. If, for instance, students were awarded a mark for each message they sent to their mentors, message traffic might simply reflect students' efforts to boost their grades. In Wagner's class, however, no grades were assigned for simply exchanging e-mail with mentors. In a similar vein, one might suspect that the frequency of students' correspondence with mentors could merely reflect their ease of access to e-mail. However, in Wagner's classroom, all students had equal access to email via six computers with direct connections to the Internet. Very few students had Internet access at home, and focus group data suggest that those few who did rarely used it to correspond with their mentors. One team in which a student was known to have corresponded with his mentor from home was excluded from the analysis presented here.

The number of messages exchanged by students and their mentors would mean very little if many of the messages had been merely sociable ("How are you? I am fine."). As the content coding presented in Table 1 shows, however, students and their mentors invested little of their total message text in sociable chat. As noted elsewhere, students tended to invest effort in sustaining their mentor relationships only if they felt it could make them more productive in their work (O'Neill & Gomez, 1998). Thus in this class, the number of email messages exchanged over the course of a project can be regarded as a reasonable measure of the effort that students chose to invest in sustaining their mentor relationships.

Argument Coding Details

Customary Argument Functions (IMRD) Scoring

Following Swales (1990), I refer to the genre in which Wagner's students wrote their articles as the "IMRD genre." As mentioned earlier, this is the clearly dominant genre for reporting scientific research. The name "IMRD" itself refers to the customary names of the major sections in such articles: Introduction, Method, Results, and Discussion. Despite some minor variations, these are very consistent across fields of inquiry in science.

In pilot work with Wagner during the 1994–1995 school year, I identified a small set of argument functions performed by his students in each section of their IMRD reports. Several of these functions roughly parallel those discussed by Bazerman (1988) and Swales (1990) in their discussion of IMRD reports written by scientific professionals. For each of the 31 reports in my sample from 1995–1996, I coded the presence or absence of these customary argument functions. I then made a simple count of the number of customary argument functions performed by the student research teams in each section of their reports. The functions coded for, and the sections in which they usually appear, are described later. Beside each list of argument functions is a bar chart showing the distribution of scores on this variable for the 31 reports coded. A quote from a high-scoring report is also provided, by way of illustration, with tags (e.g., "[a]") to mark the beginning of passages that perform specific persuasive functions.

A paper earned one point for fulfilling each previously identified argument function anywhere in the body of the paper (the argument functions did not have to appear in the proper section in order to count). Because six argument functions had been identified for Introductions, 6 was the highest possible score on the Introduction dimension. Eighteen argument functions had been identified for all four sections, so 18 could be considered a "perfect" score. The highest score among the 31 papers discussed here was 14.

Introduction. To illustrate the Introduction functions described in Figure 6, a brief quote is provided from the report on Black Holes prepared by Andy, Cori, and Bill in connection with the case discussed earlier. They scored a 6 on their Introduction, though for space considerations, not all coded moves are included here:

> [b & e] A black hole is a region in space with such gravitational pull that light and all other matter are compacted. No amount of energy could ever remove

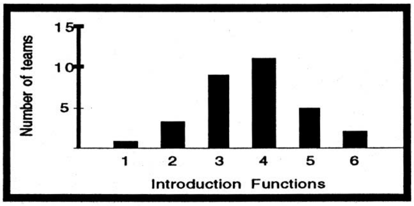

a) States a purpose in the form of a problem, question, or issue to be resolved
b) Explains the significance of this purpose to the audience (i.e. why should we care?)
c) Summarizes Method
d) Summarizes Results
e) Provides background research into the broad topic area
f) Summarizes important findings from earlier work on the problem by others (names names and gives citations)

FIGURE 6 Introduction functions and distribution of scores.

> such matter from the black hole. It is as if an entire galaxies have disappeared into an unidentifiable region. [a] Only three black holes have been conclusively identified; NGC4261, center of M87 and NGC 4258 (discovered just weeks after our project began). Nonetheless, several regions of space have been theorized as being black hole regions. [c] We decided to select two of these "non proven" black holes, Cygnus X–1 and Sagittarius A, and compare their characteristics to that of the two proven black holes, NGC 4261 and M87. We would then predict which of these two "unidentified" regions of space are most likely black holes. [d] By looking for similarities between the two proven and "unproven" black hole regions we were able to conclude that Sagittarius A is most likely a black hole yet Cygnus X–1 is beyond a reasonable doubt not.

Here we see students doing an admirable job of contextualizing their work by stating a clear problem (identifying black holes), mentioning previous findings ("proven" black holes), and summarizing their investigative methods and results. In other paragraphs not shown here for space considerations, the students also provide an overview of the theory of black hole formation and the significance of black holes as a phenomenon.

Method. To illustrate the Methods functions described in Figure 7, following is a quote from a report that investigated the empirical relation between the volume of snowfall in mountainous areas and the number of avalanches that occur there. The method amounted to plotting the two variables together on a graph, and visually inspecting them; but the research is described with uncommon precision and care to explain its significance. Because the students did not fully explain the rationale behind their research methods [c], they only scored a 2 on this scale:

> At the start of our project we planned to find out whether there is a pattern of avalanches in North America but we ran into some roadblocks along the way. There is not an organization that documents the time and date of each ava-

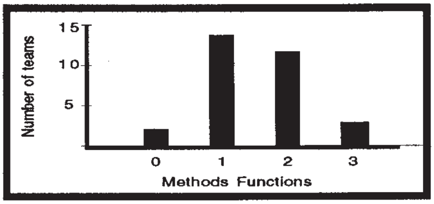

a) Describes what was done by the investigators (built a physical model, collected samples, gathered sources from libraries or electronic archives, etc.)
b) Does this description with precision appropriate for others to reproduce the results
c) Explains why the method followed could be expected to lead to a resolution of the problem, or an answer to the question

FIGURE 7 Methods functions and distribution of scores.

> lanche along with its location. … [a] We began by calling different ski areas to find if they published there [*sic*] own information about avalanche records and snowfall. We found that all of the Colorado ski areas report their data to the Colorado Avalanche Information Center. In calling the center we met Dan Altman, who turned out to be an incredible help. Dan provided is with the information we needed for analyzing annual snowfall and number of avalanches. … [b] We immediately produced spreadsheets and graphs to analyze the data received (see Spreadsheet 1, Graph 1) We examined the line graph and found that there was definitely not a direct parallel between annual snowfall and total avalanches.

Readers may note some similarity between this quote and the “adventure story” used as a negative example in the section titled Using and Misusing Scientific Genres. These students, too, also have a story to tell; but it is different from the story discussed earlier in an important way. The Method section in the previously discussed smog story dwelled on the time, hard work, and emotional turmoil involved in a search for information, without actually describing what motivated it or any of the search criteria applied. In contrast, the avalanches group tells a story about a systematic search for data, and how their interim findings re-shaped the direction of their research. A further important difference is that, after discovering which agencies routinely collect avalanche data, which ones aggregate it, and in what form, the team actually carried out some data analysis that brought them to formulate a less naive research agenda. This set of contrasting examples highlights the fact that a good Method section does not avoid storytelling altogether. Rather, it tells a story about a set of strategic actions leading to the production of knowledge claims.

Results. To illustrate the Results functions described in Figure 8, a quote is provided from a report on volcanic activity in the Pacific rim. This report’s authors scored 4 points for Results functions because they foreshadow their research results [c], present and gloss their data, and provide an interpretation of these data in a very concise manner:

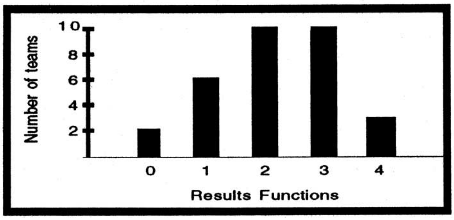

a) Foreshadows results briefly (in a sentence or two)
b) Presents the data collected or found (in tabular or graphical form, a set of images, etc.)
c) Characterizes or "glosses" the data for the non-specialist
d) Provides an interpretation of the data with respect to the original question or problem (in the form of calculations and/or prose that refers specifically to the data)

FIGURE 8 Results functions and distribution of scores.

> [a & d] The amount of volcanic activity directly corresponds to the rate of [continental plate] movement as suspected. [b & c] Of the three island arcs studied, Japan had the greatest incidence of volcanic activity; at that location, the Pacific Plate has the greatest measured velocity: 9.29 centimeters per year. The Aleutian Islands had the second greatest incident of volcanic activity; at that location, the velocity of the Pacific Plate measured 8.02 centimeters per year. The Philippines had the least incident of volcanic activity; at that point, the velocity of the Pacific Plate measured 7.27 centimeters per year. See graph number 3.

Discussion. Finally, following is a quote from a report that investigated why the coastal waters in California are so much colder than the surrounding land and ocean. The authors attributed this phenomenon, in part, to the upwelling of cold water from the ocean depths. This report scored a 4 in the Discussion dimension (Figure 9) because it offers an explanation of the phenomenon (function a) and the lack of available data [b]. In later paragraphs, the authors also describe the importance of the results and offer suggestions for further study [d]:

> The data on upwelling are indirect. [b] This is because it is really hard to measure vertical velocities (up and down flow) in the ocean, because they are very small. We infer that upwelling occurs from a combination of theory and measurements. [a] From theory, we know that wind blowing toward the Equator along a coast will tend to cause upwelling. The upwelled water comes from about 200 meters depth, so it is colder than the surrounding [sea surface temperatures] when it reaches the surface. From measurements of wind direction and SST we can say that cold coastal SSTs during periods of upwelling favorable winds (i.e., equator ward along the coast) are consistent with the theory of upwelling.

Relation of IMRD Score To Scientists' Holistic Ratings

This scheme for coding students' argument strategies would be of little value if it did not reflect what scientists find persuasive; but encouragingly, results from coding with the IMRD portion of the Science Argument Strategies scheme appear largely consistent with the holistic ratings solicited from volunteer scientists. As Table 2 shows, IMRD scores for the 18 randomly chosen papers in this subset were found to correlate significantly with volunteer scientists' ratings of each report as a whole, the quality of the argument presented in it, and the research on which it was based.

These findings, although based on a small quantity of data, suggest that the persuasive moves coded for under the IMRD score did capture important aspects of what the six volunteer scientists valued in students' reports.

Hedging Score

Myers' analysis of professional articles in the field of Biology reveals the critical importance of appropriately "hedging" arguments (softening them or limiting their scope), in order to gain credence from a professional audience (Myers, 1990). In the spirit of Myers' observations, the Hedging score (Figure 10) was developed to capture students' attempts to strategically soften the claims they made about their research. In the class studied here as in many others, it is common for students to over-

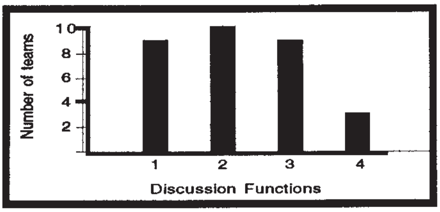

a) States the conclusions that can be made about the original question or problem from the data or information collected
b) Attempts to explain how the data support, refute, or are unrelated to hypotheses about the question or problem mentioned earlier
c) Discusses the importance of the results with reference to the significance of the original question or problem
d) Makes suggestions for further study

FIGURE 9 Discussion functions and distribution of scores.

TABLE 2
Correlations Between Introduction, Method, Results, and Discussion (IMRD) Score and Scientists' Holistic Ratings

	Whole	*Argument*	*Research*
IMRD score	.562*	.632*	.535*

*$p \leq .05$.

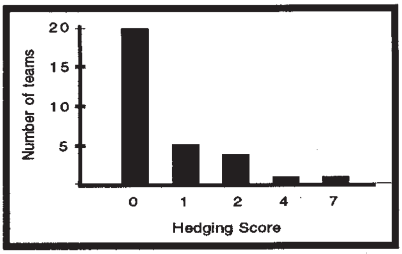

a) Acknowledges possible flaws in the method or calculations performed by the writers
b) Acknowledges the limits of the writers' own experience or data
c) Emphasizes the provisional nature of the conclusion or argument put forward (e.g. "More data are needed for a definite answer to this question")

FIGURE 10 Hedging functions and distribution of scores.

state the case for their preferred interpretations of research results. This may be done under the naive assumption that stronger arguments are inherently more persuasive. In the context of the class under study, it was hoped that ongoing conversations with scientist mentors would shed light on the ways in which they might be tempted to overinterpret their data, teaching them a more cautious and mature approach to arguing about science, and helping them to avoid the pitfall of overly bold claims.

It should be noted that unlike the IMRD function scores discussed earlier, the Hedging score has no predetermined upper limit. Although a paper can earn only one point on the IMRD scales by performing each of the previously identified argument functions, an article can earn multiple points on the Hedging scale by acknowledging a number of different possible weaknesses in the argument it presents. In principle it is possible for a team to hedge excessively, but my corpus did not contain any examples of this behavior. As the distribution of scores for the 31 papers in my sample suggests, it is far more common for students to overstate claims than to overqualify them.

To illustrate what writing looks like on the upper end of the Hedging scale, following is a quote from the Discussion section of a paper that evaluated the supporting evidence for two rival theories about the locomotion of an aquatic dinosaur known as the Plesiosaur. This paper was not only unique in the question it chose to address, but in the care with which the evidence was considered. The following passage shows some of the thoughtful provisos that earned its authors a Hedging score of 6:

> Alexander [a researcher in the field] argues that underwater flight is a more efficient means of propelling large reptiles in water, but there are no hard data to support this conclusion. … [b] Given the paucity of fossil evidence regarding the Plesiosaur it is currently impossible to reach a firm conclusion regarding its method of locomotion. [c] There is a significant amount of circumstantial evidence that seems to favor the underwater flight hypothesis, but other forms of locomotion cannot be ruled out at the present time. Until further fossil data are found or a living Plesiosaur is located (perhaps at Loch Ness be-

cause “Nessie” is thought to be a Plesiosaur-like creature), the true mode of locomotion of this extinct reptile must remain unproved.

Relation of Hedging Score to Scientists’ Holistic Ratings

Notwithstanding Swales’ findings, is the sort of claim hedging illustrated earlier really valued by practicing scientists? Apparently so. The Hedging scores for the 18 reports in my random sample correlated significantly with scientists’ holistic ratings, though not as strongly as was the case with the IMRD scores. A significant relation was found between the Hedging scores and scientists’ ratings of both the argument and the quality of the research; however, no significant relation was revealed between the Hedging scores and scientists’ ratings of the paper as a whole (see Table 3). Thus, the Hedging score appears to capture some part of what the scientists valued in students’ reports, if not as much as was captured by the IMRD score.

Relation of Argument Coding to Frequency of Correspondence

IMRD Scores

The findings of the previous section suggest that the SAS coding scheme does capture some of the argument strategies that scientists value in research reports. With the SAS coding as a gauge, I then asked, Does students’ engagement in telementoring relationships make a difference in the quality of the arguments they make about their research? The “rhetorical situation” hypothesis would suggest that through extended discussions with telementors, student research teams would become more informed about the difficulty of supporting knowledge claims based on their research. Knowing also that their mentors would be hoping and expecting to read their finished research articles, students’ persuasive goals might shift away from an exclusive focus on the time and effort they invested in their projects, placing greater emphasis on the defense of knowledge claims. But did this actually happen?

TABLE 3
Correlations Between IMRD Score and Scientists’ Holistic Ratings

	Whole	*Argument*	*Research*
Hedging score	.248	.306*	.316*

*$p \leq .05$.

The data suggest it did. The Spearman correlation between the 31 teams' frequency of correspondence with their telementors and their total IMRD function scores was found to be positive and statistically significant ($r = .357, p < .05$; see Figure 11). Furthermore, there was no significant correlation between the IMRD function scores and team members' top or average scores on the open book content test. It is arguable, then, that students' telementoring experiences had an influence on their argument strategies that would not have been predicted from their performance on traditional academic tasks.

Close inspection of the relation between the IMRD scores and frequency of correspondence on a project-by-project basis (see Table 4) revealed that Project 3 was the only one for which the correlation was individually significant ($r = .552, p < .01$). This finding may be noteworthy from the standpoint of implementation, because it suggests that telementoring had no measurable influence on IMRD scores until the students had developed a substantial collective experience with project-based sci-

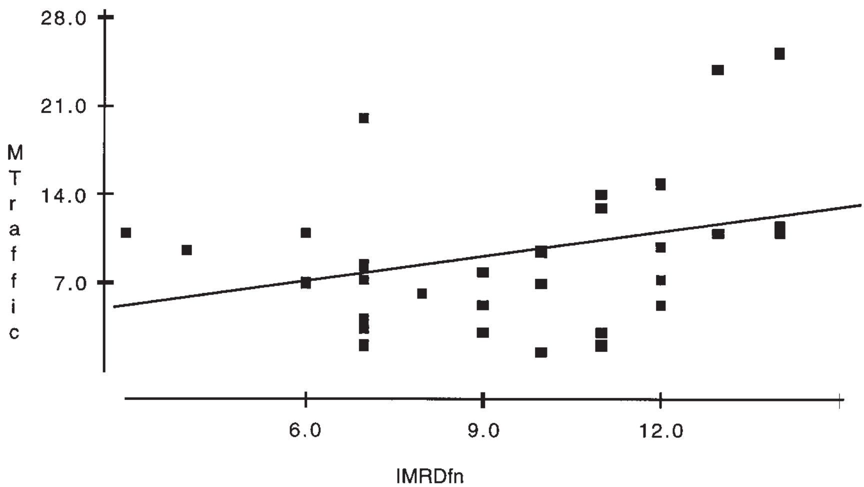

FIGURE 11 Scatterplot of IMRD function scores and message traffic for Projects 1 through 3.

TABLE 4
Correlations Between IMRD Scores and Message Traffic, by Project

	Projects 1–3	*Project 1*	*Project 2*	*Project 3*
Research articles coded	31	9	8	14
Value of ρ between Message Traffic and IMRD Score	.357*	.230	.118	.552**

$*p \leq 0.05$. $**p \leq .06$.

ence, telementoring, and/or composing reports in the IMRD genre. The small number of groups for whom both papers and matching telementoring data were available from Projects 1 and 2 leave some doubt, however, as to how early the apparent effect began.

Hedging Scores

The data from the Hedging coding also lend support to the rhetorical situation hypothesis. As the figures in Table 5 show, more extended dialogue with telementors appeared to increase students' caution in stating claims. As with the IMRD function scores discussed earlier, project-by-project inspection of the data shows that the expected correlations are most significant in the third project of the year. Hedging scores had a strong, statistically significant correlation with the number of messages exchanged by students and their mentors over the course of the third project ($r = .627$, $p < .01$).

Giving the "good students are good students" hypothesis its due consideration once again, the teams' Hedging scores were not found to correlate significantly with either the top team member's score or the teams' average scores on the Earth Science content test. Thus, the apparent influence of telementoring activity on students' written arguments cannot easily be explained away as the product of differing levels of ability between teams.

Relation of Argument Coding to E-Mail Content

According to the theory I reviewed at the beginning of this article, we should not be surprised to see a significant relation between genre use and telementoring activity. Nonetheless, the relation observed here raises an important question: What exactly was it in the telementoring dialogues that changed students' approach to their writing? This question is a practical one, for if it turns out that specific dialogue moves or subjects are linked to the production of more sophisticated arguments, we might

TABLE 5
Correlations Between Hedging Scores and Message Traffic, by Project

	Projects 1–3	*Project 1*	*Project 2*	*Project 3*
Research articles coded	31	9	8	14
Value of ρ between Message Traffic and Hedging Score	.332*	.120	.195	.627**

$^*p \leq .05$. $^{**}p \leq .01$.

be better able to advise students and mentors about the sorts of conversations they should strive to have.

Though full e-mail logs were available only for a subset of the student teams in the dataset (n = 12), I offer some preliminary observations on the basis of these data. On the students' side of the telementoring dialogues, several relations with argument strategies stand out. To begin with, there were statistically significant ($p < .01$) relations between the teams' IMRD scores and the quantity of message text they invested in both discussing specific concepts in their domains of inquiry, and posing direct questions to their mentors (see Table 6). Dan's dialogue with Andi, Cori, and Bill in the section titled "An Example of Research Guidance in a Telementoring Relationship" provides good illustrations of these features, which provide important avenues for mentors to observe and support students' thinking. Students' reports on the status of their investigations serve this purpose even more directly, which may explain the equally significant relation observed between the teams' IMRD scores and the quantity of message text invested in status reports. Less significant relations ($p < .05$) were observed between the Hedging scores and the quantity of message text invested in asking questions and discussing domain concepts.

On the mentors' side of the dialogue, the quantity of message text invested in posing questions to students was significantly related to both the IMRD ($p < .01$) and Hedging ($p < .05$) scores. Although mentors' questions were not categorized in detail, this coding category included both questions about the status of the investigation (e.g., "Haven't heard from you in a while. How is your project coming?") and heuristic questions of the sort illustrated in Dan's dialogue with Andi, Cori, and Bill (e.g., "What is a black hole? If it exists, where would we expect to look for it?"). If taken seriously, one would expect these sorts of questions to influence students' thinking about the nature of the evidence required to firmly establish their knowledge claims. This in turn would influence their use of hedging and other argument strategies.

These preliminary findings should not be taken to indicate that the effort students and their mentors invested in subjects or facets of dialogue other than questions, domain concepts, and status reports were in any way wasted. Students and their mentors can hardly get to the point of asking one another pertinent questions, or frankly discussing the progress of an investigation without first building a rapport. And they sometimes cannot sustain this rapport without resorting to repair strategies, such as apologizing for past behavior. What these findings tentatively suggest, however, is that as students and mentors carry out their work together, they should strive to ask one another informed questions about the subject matter domain and the course of the investigation itself, rather than merely requesting information and passing hints and suggestions.

TABLE 6
Correlaztions Between IMRD and Hedging Scores and Moves in Telementoring Dialogue

Mentor					*Students*				
Topic of Dialogue	*Words*	*%*	*IMRD* ρ	*Hedge* ρ	*Topic of Dialogue*	*Words*	*%*	*IMRD* ρ	*Hedge* ρ
Conceptual					*Conceptual*				
Project ideas	1,356	11	.143	–.004	Project ideas	2,795	23	0.97	–.043
Domain concepts	3,883	32	.322	.274	Domain concepts	955	8	.725**	.596*
Questions	706	5.8	.712**	.638*	Questions	1,004	8.3	.702**	.596*
Procedural/Advice					*Procedural/Advice*				
Requests	731	6	.254	–.057	Requests	2,383	20	–.040	.050
Project status	379	3.1	.070	–.039	Project status	1,663	14	.711**	.501
Hints	567	4.7	.145	–.009	Hints	0	0	.010	0
Suggestions	1,284	10.6	.302	.069	Suggestions	0	0	0	0
Offers	613	5.1	.023	.111	Offers	0	0	0	0
Advice	1,432	12	.094	–.057	Advice	0	0	0	0
Relationship					*Relationship*				
Hello/goodbye	542	4.5	0	0	Hello/goodbye	1,886	16	–.113	–.002
Thanks	0	0	0	0	Thanks	837	7	.334	–.201
Complaints	183	3.5	.445	.196	Complaints	0	0	0	0
Apologies	431	3.5	–.490	–.245	Apologies	320	2.7	–.401	–.245
Compliments	230	1.9	.445	.196	Compliments	71	.6	–.401	–.245
Total words coded as moves: 12,016									

**p* ≤ .05. ***p* ≤ .01.

Summary of Results

This analysis of 31 research articles composed by teams of students in a project-based high school science class showed a significant relation between the use of sophisticated argument strategies in a professional genre of writing and sustained e-mail relationships with volunteer mentors. In particular, teams of students who sustained their telementoring relationships were likely to fulfill the customary persuasive functions of an IMRD report, as well as to soften their arguments in strategic ways. Because all students in the class had access to ongoing advice and guidance from the classroom teacher as well as their telementors, the influence of guidance from these two sources cannot be fully disentangled. However qualitative evidence suggests that routine discourse with a critical audience besides the teacher played an important role in helping traditional science avoiders put forward more sophisticated scientific arguments.

The most likely alternative explanation of these findings, the "good students are good students" hypothesis, found no support in the available data. Under this hypothesis, "good" students would be expected both to produce thoughtful reports and sustain their telementoring relationships, simply because both were expected of them and because they were capable of both. However, the lack of any significant correlation between the genre use scores and team members' average or maximum grades on a content test appears to disconfirm this hypothesis. One would have expected the most dedicated or compliant students to score high on a traditional academic task like a content test, but these high scores did not predict the quality of the written arguments about project work.

Finally, we may speculate that both the best research articles and the richest telementoring relationships were produced by teams whose members possessed the strongest prior writing ability. Hypothetically, this ability would be beneficial to both the composition of research articles and the maintenance of telementoring relationships, but might not correlate significantly with high scores on the content test. Data to definitively address this hypothesis were not collected in the study, but the fact that the test used in this study was composed largely of long-answer questions casts this explanation in doubt.

CONCLUSION

The notion of communities of practice is now a powerful driving force behind design work in the learning sciences. It is, however, a slippery idea to get a grip on; and if we cannot judge our success in initiating students into new intellectual communities, there is little hope that we can improve our designs systematically. Here, I proposed genre analysis as a theoretically grounded way to assess students' understanding of persuasive practices in the scientific community. I then illustrated the approach using data from a project-based Earth Science class in which students engaged directly with volunteer scientists in lengthy on-line mentoring relationships.

Genre theory is useful as more than a basis for assessment, however. It also brings pedagogical problems into focus that may be ignored in other frameworks. In particular, genre theory clarifies the importance of the social situations in which we ask students to apply new knowledge. Even advanced project-based learning environments such as the one I discussed here present some risk of students learning to misuse the new disciplinary forms of expression (genres) to which they are exposed. In science for example, if students lack an audience that attends critically to the knowledge claims justified by their research, they may fail to use and master the appropriate tools of persuasion. This will be particularly true if other sorts of tools appear to work equally well in gaining credit. After all, it is easier to marshal arguments about the amount of time invested in an investigation than about the merits of the data collected and the analysis performed on them.

Understood in this frame, the unique value of innovations like telementoring lies in the routine involvement they foster between students and an audience of knowledgeable adults who are more attuned to scientific persuasion than teachers can often afford to be, or than student peers are often able to be. When students engage deeply with such an audience, and understand that scientific argument strategies are valued by it, they can better understand the task of mastering these strategies, and take it more seriously. As a teacher involved in my research explained,

> I'm hoping that by being involved with [their telementors], that they're putting [their work] out there for someone else, [and] that they're going to be a little more critical of themselves. Because it's not me that they've known for three years. … They'll have somebody new that they're presenting [their arguments] to, and somebody who will give them a different kind of feedback than the feedback I've given them.

In ongoing research, I have been working to understand why not all students are able to sustain their telementoring relationships (O'Neill & Gomez, 1998), and how new technologies may enable them to do so (O'Neill & Harris, 2000; O'Neill & Scardamalia, 2000). I have also been developing a Web-based service that I hope will enable telementoring to be support teaching on a large scale (O'Neill & Gomez, 2000). I am hopeful about the number of volunteers that could be engaged in this work, because current census data on volunteering and Internet use among adults with postsecondary degrees suggest that as many as 2.7 million people in the United States and Canada may be willing and able to serve as telementors (O'Neill & Harris, 2000).

The telementoring study illustrates three ways in which genre analysis can be of value to learning scientists. These are its sensitivity, noninvasiveness, and connection with traditionally valued educational outcomes. The sensitivity of genre use measures is particularly valuable to designers who test their work in complex classroom settings, for the success of design experiments must be gauged not only summatively, but in ways that support continuous tuning and refinement. To the extent that coding schemes like the SAS help researchers evaluate the effects of a new design on students' use of particular kinds of argument strategies, they can help to strengthen that design in selective ways (e.g., targeting Methods functions or Hedging specifically).

Also important in design experiments research is the need to avoid fatigue effects. Because design experiments researchers rely on students and teachers in so many ways—from co-planning and designing innovations to critiquing software and more—it is easy for assessment-related tasks to become the last straw. By allowing researchers to take the traditional products of students' classroom research as rich sources of data, genre analysis largely avoids the complications of surveys and performance tasks, which may not be taken seriously.

Finally, although measures of genre use are a new technique for learning scientists, they are tied to traditional educational values in useful ways. In recent years, we have seen the development of exciting computing tools that afford the development of whole new forms of literacy in science (e.g., Bell, 1997; Loh et al., 1997; O'Neill & Gomez, 1994; Suthers, 1998). The promise of these tools is great; but if we are to engage teaching professionals in large-scale work with them, it will be important to show how the new literacies they shape are related to those that science educators and practitioners know and value. In addition to understanding the benefits of new tools on an empirical basis, we must be able to assure educators that they do not forsake traditional goals. As the history of writing in science shows, these have often become traditional for good reason.

ACKNOWLEDGMENTS

I would like to acknowledge the generous support of the National Science Foundation under Grant RED–945729 and the Illinois State Board of Education during the data collection and analysis phases of this research.

I also offer my sincere thanks to Rory Wagner of New Trier High, Winnetka, Illinois, Judith Lachance-Whitcomb of Jordan Community School, Chicago, and their students for the invaluable contributions of time, effort, and encouragement they offered me. This research also benefited from my collaboration with Marlene Scardamalia and Carl Bereiter, and the CSILE/Knowledge-Building team at the Ontario Institute for Studies in Education of the University of Toronto (OISE/UT). That collaboration was generously supported by the James S. McDonnell Foundation, the Natural Sciences and Engineering Research Council of Canada, and the Office of Learning Technologies, Human Resources Development Canada. Finally, I offer my thanks to two anonymous reviewers, and to Laura D'Amico and Mary Lamon for their thoughtful responses to previous drafts of this article.

REFERENCES

Bartholomae, D. (1985). Inventing the University. In M. Rose (Ed.), *When a writer can't write: Studies in writer's block and other composing-process problems* (pp. 134–165). New York: Guilford.

Bazerman, C. (1988). *Shaping written knowledge: The genre and activity of the experimental article in science*. Madison: University of Wisconsin Press.

Bell, P. (1997, December). *Using argument representations to make thinking visible for individuals and groups*. Paper presented at the conference on Computer Supported Collaborative Learning, Toronto, Canada.

Berkenkotter, C., & Huckin, T. N. (1995). *Genre knowledge in disciplinary communication: Cognition/culture/power*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Bitzer, L. (1968). The rhetorical situation. *Philosophy and Rhetoric*, *1*(1), 1–14.

Brown, A. L., & Campione, J. C. (1994). Guided discovery in a community of learners. In K. McGilly (Ed.), *Classroom lessons: Integrating cognitive theory and classroom practice* (pp. 229–272). Cambridge, MA: MIT Press/Bradford.

Bruner, J. S. (1990). *Acts of meaning*. Cambridge, MA: Harvard University Press.

Cope, B., & Kalanzis, M. (1993). *The powers of literacy: A genre approach to teaching writing*. Pittsburgh, PA: University of Pittsburgh Press.

Flower, L., & Hayes, J. R. (1980). The cognition of discovery: Defining a rhetorical problem. *College Composition and Communication*, *31,* 21–32.

Freedman, A. (1993). Show and tell? The role of explicit teaching in the learning of new genres. *Research in the Teaching of English*, *27,* 222–251.

Hawkins, J., & Collins, A. (in press). Design experiments: Evaluating the role of technology in supporting school change. In J. C. Hawkins (Ed.), *Design experiments: Using technology to restructure schools*. New York: Cambridge University Press.

Krajcik, J. S., Blumenfeld, P. C., Marx, R. W., & Soloway, E. (1994). A collaborative model for helping teachers learn project-based instruction. *Elementary School Journal*, *94,* 483–497.

Latour, B., & Woolgar, S. (1979). *Laboratory life: The construction of scientific facts*. Princeton, NJ: Princeton University Press.

Lave, J., & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge, MA: Cambridge University Press.

Loh, B., Radinsky, J., Reiser, B. J., Gomez, L. M., Edelson, D. C., & Russell, E. (1997, December). *The Progress Portfolio: Promoting reflective inquiry in complex investigation environments.* Paper presented at the conference on Computer Supported Collaborative Learning, Toronto, Canada.

Miller, C. R. (1994). Genre as social action. In A. Freedman & P. Medway (Eds.), *Genre and the new rhetoric* (pp. 23–42). London: Taylor and Francis.

Myers, G. (1990). *Writing biology: Texts in the social construction of scientific knowledge*. Madison: University of Wisconsin Press.

O'Neill, D. K. (1998). *Engaging science practice through science practitioners: Design experiments in K–12 telementoring*. Unpublished doctoral dissertation, Northwestern University.

O'Neill, D. K., & Gomez, L. M. (1994, June). *The collaboratory notebook: A networked knowledge-building environment for project learning*. Paper presented at Ed-Media '94: World Conference on Educational Multimedia and Hypermedia, Vancouver, BC, Canada.

O'Neill, D. K., & Gomez, L. M. (1998, November). *Sustaining mentoring relationships on-line*. Paper presented at the ACM Conference on Computer-Supported Cooperative Work, Seattle, WA.

O'Neill, D. K., & Gomez, L. M. (2000). *Unfinished business: Human relationships in the design of CSCW systems*. Unpublished manuscript.

O'Neill, D. K., & Harris, J. B. (2000). *Is everybody happy? Bridging the perspectives and developmental needs of participants in telementoring initiatives*. Paper presented at the annual meeting of the American Educational Research Association, New Orleans, LA.

O'Neill, D. K., & Scardamalia, M. (2000, June). *Mentoring in the open: A strategy for supporting human development in the knowledge society*. Paper presented at ICLS 2000: International Conference on the Learning Sciences, Ann Arbor, MI.

O'Neill, D. K., Wagner, R., & Gomez, L. M. (1996, November). Online mentors: Experimenting in science class. *Educational Leadership, 54,* 39–42.

Pea, R. D. (1992). Practices of distributed intelligence and designs for education. In G. Salomon (Ed.), *Distributed cognitions* (pp. 89–132). New York: Cambridge University Press.

Pea, R. D. (1993, May). Distributed multimedia learning environments: The collaborative visualization project. *Communications of the ACM*, 60–63.

Pea, R. D., & Gomez, L. M. (1992). *Distributed multimedia learning environments: Why and how?* [Tech. Rep. 25]. Chicago: Northwestern University, Institute for the Learning Sciences.

Perkins, D. N. (1993). Creating a culture of thinking. *Educational Leadership, 51*(3), 98–99.

Polman, J. L. (2000). *Designing project-based science: Connecting learners through guided inquiry*. New York: Teachers College Press.

Ruopp, R., Gal, S., Drayton, B., & Pfister, M. (Eds.). (1993). *LabNet: Toward a community of practice*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Russell, D. (1991). *Writing in the academic disciplines, 1870–1990: A curricular history*. Carbondale: Southern Illinois University Press.

Scardamalia, M., & Bereiter, C. (1994). Computer support for knowledge-building communities. *Journal of the Learning Sciences, 3,* 265–283.

Suthers, D. (1998, April). *Representations for scaffolding collaborative inquiry on ill-structured problems*. Paper presented at the annual meeting of the American Educational Research Association, San Diego, CA.

Sutton, C. (1989). Writing and reading in science: The hidden messages. In R. Millar (Ed.), *Doing science: Images of science in science education* (pp. 137–159). London: Falmer.

Swales, J. (1990). *Genre analysis: English in academic and research settings*. New York: Cambridge University Press.

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

Wertsch, J. V. (1991). *Voices of the mind: A sociocultural approach to mediated action*. Cambridge, MA: Harvard University Press.

Yates, J., & Orlikowski, W. J. (1992). Genres of organizational communication: A structurational approach to studying communication and media. *Academy of Management Review, 17,* 299–326.

APPENDIX
Select Categories From the Science Argument Strategies Coding Scheme

Paper Title: ____________________

Team Number: ________

Date Submitted: ____________________

Writers & IDs ____________________ ID __________

____________________ ID __________

____________________ ID __________

____________________ ID __________

Teacher, Period: ____________________

Date Coded: ________ Coder: ____________________

Section-Specific Persuasive Functions

RA Sections Present (i.e. marked with a labeled heading)

☐ Introduction (sometimes labeled "Background")	☐ Method (sometimes labeled "Procedure")	☐ Results (sometimes labeled "Data")	☐ Discussion (sometimes labeled "Conclusions")	☐ References (sometimes labeled “Bibliography” or “Literature Cited”)

Introduction

☐ States a purpose in the form of a problem, question, or issue to be resolved	☐ Explains the significance of this purpose to the general audience (i.e. why do we care?)	☐ Provides background research into the broad topic area
☐ Summarizes important findings from earlier work on the problem by others (names names and gives citations)	☐ Summarizes Method	☐ Summarizes Results

Methods

☐ Describes what was done by the investigators (built a physical model, collected samples, gathered research sources from libraries or electronic archives, etc.)	☐ Does this description in terms and with precision appropriate to others who might want to reproduce the results	☐ Explains why the method followed by the writers could be expected to lead to a resolution of the problem, or to an answer to the question	☐ Mentions specific search goals and/or criteria employed at library or in Internet searches (e.g. ‘we searched with the keywords "sea surface temperature’")

(continued)

APPENDIX (*Continued*)

Results

☐ Foreshadows results briefly (in a sentence or two)	☐ Presents the data collected or found (in tabular or graphical form, a set of images, etc.)	☐ Characterizes or "glosses" the data for the non specialist	☐ Provides an interpretation of data with respect to the original question (in the form of calculations and/or prose that refers specifically to the data)

Discussion

☐ States the conclusions that can be made about the original question or problem from the data or information collected	☐ Attempts to explain how the data support, refute, or are unrelated to a question or problem mentioned earlier	☐ Discusses the importance of the results with reference to the significance of the original question	☐ Makes suggestions for further study

Other Persuasive Features

Hedging (record line numbers and count)

Acknowledges possible flaws in method or calculations performed by the writers	Acknowledges the limits of the writer's own experience or data	Emphasizes the provisional nature of the conclusion or argument put forward (e.g. "More data are needed for a definite answer to this question")
[ ] #	[ ] #	[ ] #