22 Videorecording as Theory

Rogers Hall
University of California–Berkeley

> Carl: They're not tattoos, they're skin ILLUSTRATIONS! Don't you ever call them tattoos. Let me tell you. Don't you look at those illustrations too long, because they'll come alive and they'll tell stories.
>
> Rod Steiger as Carl, in a film adaptation of Ray Bradbury's *The Illustrated Man* (1951, directed by Jack Smight, 1969)

We can learn a great deal about the detailed structure of learning and teaching by watching videotape of people in action. But, we still don't know much about our own activities of collecting, watching, or interpreting video as a stable source of data for research and presentation purposes. That one can get into trouble or, worse still, get others into trouble, by showing films of study participants is apparent in the kinds of "war stories" that researchers swap outside conference meeting rooms. Relatively little has been written about this. In this chapter, I use difficulties encountered in my own research on teaching and learning mathematics to illustrate what I take to be a set of core problems with appropriate uses of video as data. These problems include:

- Processes of collecting video data and making selections from it are usually (and too easily) deleted from research accounts.
- Production values that are preserved in technical arrangements for collecting video become a permanent part of the data one is recording.

- Video databases may be public resources by virtue of accepted arrangements for scholarly review and research sponsorship, yet, we cannot fully anticipate the public use of database materials.

These problems persist because there is no single community of practice surrounding appropriate use of video records as data in research with human (and technical or organizational) participants. I present several local suggestions for overcoming some of these problems, but I expect that no single set of standard practices is possible. Instead, a heterogeneous set of agreements and conventional practices will need to emerge over time. The chapters collected in this volume are a step in that direction.

DATA ARE TECHNOLOGY (AS WELL AS THEORY) LADEN

In this section, I focus on how video and audio records are used to capture stretches of doing, learning, and teaching. My point is only to show that inclusions/exclusions of what is "primary" about the resulting data follow directly from technical and theoretical expectations about what a record of these activities should look like. These expectations create data that are both technology and theory laden. In a later section, I consider what happens when these records "travel" outside the contexts of use in which they are collected.

Technology and Research Practices

There is a widespread though controversial belief in the human sciences that new technologies shape the specific cultural practices we call human thinking (Eisenstein, 1979; Goody, 1977, 1987; Latour, 1986; Olson, 1994; Street, 1984; Tulviste, 1991). Given this family of conjectures, we might learn something by applying the same logic of development to the specific cultural practices we call cognitive and educational research.

To what extent do technologies for getting access to phenomena of human action matter when doing research? For example, without the widespread availability of audiotape recording and playback machines, two of the most radical innovations in close studies of human action would be completely unworkable: studies of thinking as symbolic information processing using verbal protocol analysis (Ericsson & H. A. Simon, 1984) and studies of the social order using analyses of talk-in-interaction (Garfinkel, 1967; Sacks, 1992; Schegloff, 1984). Both are scientific movements that have reorganized how we think about thinking and human action. But, starting with a common (and apparently neutral) technology to capture people talking, they have created accounts of human action that may turn out to be

incommensurable (Costall & Still, 1987; Coulter, 1991; Suchman, 1987; Vera & H. A. Simon, 1993).

By considering how new combinations of text, sound, image, and interactivity can be used to report research findings about learning and teaching, we raise an interesting methodological question. How should we expect new technical media for recording human action, along with a growing suite of tools for their analysis, to shape the specific cultural practices of cognitive and educational researchers?

The first five figures (and following sections) in this chapter examine different sorts of trouble that can arise when we treat videorecordings as objective or theory neutral data. Each shows a type of activity that is important in research on learning or teaching. These activities include: making inferences, solving problems, teaching, telling, and explaining that (with one exception) are drawn from my own work. The figures show scenes from these studies, collected in ways that make aspects of the phenomena of learning or teaching more or less visible. In each case, the technical arrangements for recording data reflect and help to reproduce specific cultural practices of analysis, inference, and publication in educational research.

Scenes on the left in each figure show a narrow framing of each activity (deleted), whereas scenes on the right show a wider framing (restored).[1] In each contrasting pair of scenes, my aim is to show that any activity of interest (e.g., solving problems) must be rendered in ways that selectively delete or foreground aspects of the original setting. This is not usually an oversight on the part of researchers (Star, 1983), but rather, is driven by technical arrangements for capturing human activity and by theoretical expectations about the boundaries of that activity. How one approaches the "primary data" in each case matters a great deal to what is later available for analysis.

---

[1] The contrasting terms, *deleted* and *restored*, reflect my own interests and assumptions as a researcher. Rather than masking these interests, I use material from my own research to show how different approaches to recording can give access to different kinds of phenomena.

Deleted

...without moving

Restored

...with moving

FIG. 22.1. Making inferences with and without moving.

Thinking With and Without Moving

To the left in FIG. 22.1, a subject sits in an experimental device that records eye movements as she does things like make inferences during a reading comprehension task. In order for us to tell "What is she looking at?",[2] the subject in this situation needs to keep her body relatively still and attend to what she is shown. To the right in FIG. 22.1, two civil engineers lean into a table layered with what they call "paper space" views of a design project (Hall, 1994). Without getting their bodies into this position, they would have nothing coherent to look at or read, and so would have difficulty making inferences about their ongoing design problem. The activity that researchers hope to study across scenes in FIG. 22.1 might be the same (e.g., making inferences about complex informational displays), but in one approach the body and task are immobilized to get at inference, whereas in the other, inferences and tasks appear only when people and materials move around in coordinated ways.

---

[2]This image comes from an advertisement in the *Proceedings of the Fourth European Conference on Eye Movements* (Luer & Lass, 1987, p. 398). It is the only graphical rendering of an intact person I could find in the entire conference proceedings. The few graphical images used to report research results showed an eyeball in plan view, oriented toward a fixed, planar display for a researcher-presented task.

This is an extreme contrast, but it is a good place to start thinking about how technologies for recording primary data can influence the phenomena being studied. Though studies that use eye movements to make inferences about attention and the contents of working memory can tell us very subtle things about how people use text and graphics while working on experimental tasks (e.g., Hegarty, 1992, reported that subjects "animate" connected parts of mechanical displays while trying to explain how they work), these technical arrangements for recording data strip away many of the resources that people have in situations where they find and solve complex problems. These resources include other people, changes in perspective made possible by reorienting one's body and/or the representational forms in use, and use of one's own and others' hands to assemble and index parts of the forms displayed.

In the studies I have been doing of "math at work" (Hall & R. Stevens, 1995, 1996; R. Stevens & Hall, in press), these resources appear to be necessary for people to get their work done, to teach, and to learn. Briefly, our approach (see also C. Goodwin, 1993; Hutchins, 1995) has been to gain access to ongoing work activities, to film a "wide" view of activity with a stationary camera that can be operated by remote control, to use hand-held cameras with a steady cam feature to "follow" detailed work with representational forms, and to locate wireless audio microphones on people and in places where we have found that talk routinely occurs. By starting with multiple video and audio perspectives on activities that occur naturally in the workplace, we can do work at a theoretical level that stretches cognition (e.g., inference, calculation, teaching, learning) over people, things, and space in new and interesting ways (Lave, 1988). Again, the point here is simply that our technical arrangements for collecting primary data greatly influence the kinds of inferences we will be able to make later.

FIG. 22.2. Solving problems with and without heads or bodies.

Solving Problems With/Without Whole People

FIG. 22.2 shows two scenes from studies of design "problem solving" in elementary and middle school mathematics (Berg, Chiu, & Hall, 1994; Goldman, 1994; Goldman, Moschkovich, & The Middle School Mathematics Through Applications Project, 1995; Hall, 1995). In the first scene, hands are busy at the computer screen of a prototype software environment. With the computer in full frame, we can recover relatively little about whose hands are in play, how this activity at the interface got organized, or where it is going. The second scene shows a combined perspective on these same people and their activities. At the lower right, a picture-in-picture image shows a group of students and their teacher clustered around a computer. Using a digital mixing board, this wide perspective on their activity can be inserted and moved around inside a simultaneously recorded image that follows the details of work at the interface. Using the combined images, for example, we might find that a student has been systematically "carpentered" (R. P. McDermott, Gospodinoff, & Aron, 1978) out of access to a design-in-progress at the computer screen. Both scenes are records of problem-solving activity, but each gives a different sense of who is solving what problems. From the first to the second scene, what gets restored is the people who are manipulating objects on the screen (or not) and the way in which they are organizing their attempt to find a solution.

I have come to think of the leftmost video record as a "primal screen scene," because it closely reflects the interests of an interface designer (myself, along with others, at the time) in how interface objects are being used by children who are (we hope) learning and solving problems as they design things. A stationary camera sits above and behind a group working at the computer, and a single wireless microphone is taped to the side of the computer display. As a recording of action at an interface, this technical arrangement has a nice feature: We can see which objects people are pointing at as they speak, so we have a chance to recover the meaning of indexical talk[3] like, "No! Move *that* over *there*!" This would not be true of a

---

[3]Indexicality refers to the way utterances depend on their context for meaning and the fact that the meaning of words can shift as the context changes (Duranti & C. Goodwin, 1992; Hanks, 1996). It is important to realize that this is a *commonplace* feature of language use in multiparty interaction.

(continued)

record of the screen recorded by taking the video signal directly from the back of the computer display (a number of companies sell devices for doing this) or by replaying a "dribble file" that records a chronologically ordered sequence of interface events. Ironically, taking only a wide view of this interaction would show us that people are pointing at or animating objects on the screen, but we could not tell which objects were in play or what their important features were. So *both* wide and follow perspectives are necessary for understanding what people are doing in this case.

...from a head with no bodies (teacher as exhibit)

...between heads and bodies (teaching dioramas)

FIG. 22.3. Teaching with and without acting or listening.

Teaching With and Without Acting or Listening

FIG. 22.3 shows two different approaches to capturing "teaching" in a middle school mathematics classroom. In the first scene, the camera frames a conventional view of teaching as talking to students, so we get a talking head without bodies and utterances without hearers. In the second scene, collected in the same classroom after thinking a little harder about the routine structure of teaching interventions, we get a frame that includes more aspects of teaching as an activity done in interaction with learners. The teacher stops in at a local group,

---

including situations where people have shared access to physical or computational artifacts (e.g., a graphical user interface). Goodwin and Goodwin (1996) described this particular kind of indexical talk as "articulating the surface" of complex representational displays.

and (most of) the participants lean in to have a conversation about the day's activity. From the first to the second scene, what gets restored is the teacher's activity and in progress relations with the students she teaches. If we thought of these two scenes as the construction of museum exhibits (i.e., images that we and our readers would later take to be representative of something), the first treats the teacher herself as the exhibit whereas the second frames the teacher's interactions with local groups as a collection of dioramas[4] that unfold over time.

Thus far we have another, relatively simple story about trying to include more of what people are doing in the records that we collect to study them. It is important to point out two things, however. First, one does not get a usable "teaching diorama" by recording a fixed, wide-angle view of an entire classroom with an open audio source. Instead, the scene shown to the right in FIG. 22.3 was recorded with a stationary camera that was operated (i.e., panning and zooming) by a person whose sole job it was to follow teaching events. The teacher wore a wireless lapel microphone, clipped at the middle of her blouse so that it would pick up local student voices without being drowned out by her own talk. As the record was made, the camera operator wore headphones that shut out ambient noise in the rest of the classroom, allowing him (me, in this case) to follow the teacher's local conversations with students quite closely. Second, by combining a record of these teaching dioramas with focused records of group work like those described in the preceding section, it is possible to treat classroom videography as a sampling scheme that weaves together the multiple, local perspectives of teachers and students as they work together across settings. This is a labor- and equipment-intensive approach,[5] of course, but it has the nice feature that subsequent

---

[4] In *Scenes from Deep Time: Early Pictorial Representations of the Prehistoric World*, Martin Rudwick (1992) analyzed the history of different technical approaches to constructing museum exhibits that were intended to "show" Cretaceous life in paleontology. His point, and the point of comparing "teachers as exhibits" with "teaching dioramas," is that the ways in which we frame actors (whether dinosaurs, students, or teachers) has a great deal to do with how we then come to see and understand their activity.

[5] In this case, three cameras and two or more microphones would be required: (a) one camera and one microphone to follow the teacher and (b) two cameras and one or more microphones to follow the local work of a group. For a less equipment-intensive alternative, see Stigler's (1995) approach to capturing the perspective of an "ideal" student as part of the Third International Mathematics and Science Study.

(continued)

analyses can question relations between students' perspectives on instruction and those of the teacher, between the curriculum as taught and the curriculum as enacted by students, etc. In addition, if the teacher visits most of the groups in a classroom, a record of teaching dioramas will also be a record, at some level of granularity, of what these groups are doing.

Deleted

...without asking (semi-structured interview)

Restored

...with asking, showing, and negotiated understanding (participatory design)

FIG. 22.4. Telling with and without asking.

Telling With and Without Asking or Interacting

In FIG. 22.4, two interview records taken with teachers are shown (Hall, Knudsen, & Greeno, 1996). In the first, we see only the "respondent" (a teacher) telling about her work without any evident access to an interviewer. That is, the teacher (it could be anyone in a typical, semistructured interview) is technically framed as a respondent: She tells us things when asked, she asks relatively few questions of her own, and we can see nothing of the interactional work of the interviewer in the resulting record. In the second scene, the same teacher (and a coworker) sit at a table with a group of researchers. Everyone is talking, not

---

Using a single camera and (I presume) an audio source attached to the teacher, camera operators were trained to "point the camera toward that which should be the focus of the ideal student at any given time" and to "be certain to capture everything that the teacher is doing to instruct the class" (p. 15). It is probably the case that stronger theoretical commitments allow one to record less, rather than more. In any case, you get what you invest in (i.e., investing in operator time, in equipment, and in later indexing and analysis of the resulting records).

only to each other, but also about records of interaction drawn from these teachers' classrooms. From the first to the second scene, what is restored is the relation between research participants (teachers), people doing the research (us), inspectable records of the phenomena up for discussion (e.g., students working on a design problem), and the negotiated character of arriving at an analysis of those records.

This comparison shifts our attention from collecting video records of people doing things (e.g., using a computer to solve a problem or teaching) to collecting records of people telling us about what they do. Interviewing people raises all the usual problems of distinguishing what they say from what they actually do (H. S. Becker & Geer, 1969; B. Jordan & Henderson, 1995), but my point here is that technical and organizational arrangements for "telling" make a difference in sorting these things out. First, as is shown in the scene to the right in FIG. 22.4, we can ask people questions in the context of prior records of what they and others have actually been doing together. How selected aspects of a prior video record can be used in an interview is a complex issue. In the scene to the right, for example, video excerpts of local student work were selected both by teachers and academic researchers, and the contrasts between what we selected turned out to be interesting for our attempts at a joint analysis of how assessment systems operate in elementary school mathematics classrooms. With these records in view, however, telling about what happened is held accountable to aspects of what did happen, at least as these happenings were captured in the video record and to the extent that discussions systematically revisit this record of interaction.[6] Second, and as implied already by this example, one can invite the people usually treated as respondents to participate in the analysis. In this case, we were working over the summer with a fifth- and sixth-grade teacher, using video records we had recorded in their classrooms during the preceding school year. We framed the collaboration as an instance of the "participatory

---

[6] This is a crowning feature of video-based interaction analysis (B. Jordan & Henderson, 1995), where the video record up for discussion is played until someone asks to "stop" (technically, to "pause") and discuss something they have seen. The discipline of returning to the tape, ironically, is supported by a feature of consumer video playback mechanisms: They resume play after they have been placed on "pause" for more than several minutes, so analysts' tendencies to pontificate are faithfully interrupted by automatic resumption of the video record.

design" (Schuler & Namioka, 1993; Suchman & Trigg, 1991a) of assessment systems, with the intention of designing and field testing new forms of assessment in their classrooms the following year.

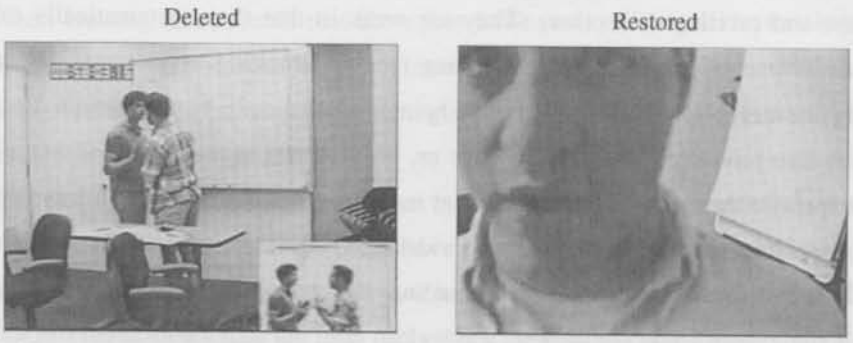

FIG. 22.5. Explaining with and without seeing.

Explaining With and Without Seeing

In FIG. 22.5, we see two different perspectives on a single explanation for why rates can be added together in a difficult algebra problem (Hall, 1996). In the first scene (an observer perspective), two teachers meet in a simulated act of collision, while one of them narrates a simple explanation for why combining rates would mean that two drivers are "goin' faster." In the second scene (a participant perspective), recorded later with the original explainer (wearing a plaid shirt in the first scene), we see a re-creation of the simulated collision that was filmed to examine the recipient's experience during this explanation. To do this, I carried a camera and acted like the recipient (i.e., followed instructions, marched toward the explainer, and eventually collided with him), while the original explainer read from a transcript of his earlier directives and explanatory comments. What gets restored from the first to the second scene is the fact that both of these teachers (i.e., participants in this situation) had an experience that we (i.e., observers of the situation) did not have and might easily have overlooked.

This kind of constructed video record helps raise some important questions about what we count as primary data in research on teaching and learning. First, it helps us to see that every type of video record described in this chapter encodes a particular perspective on ongoing human activity. What is more, in almost every one of these cases, the perspective that goes "on the record" is one that no participant in the recorded activity could have had. If

nothing else, this should encourage us to think carefully about claims that videotape provides "objective" or "realistic" records of human action.[7]

Second, this case should help us to see that video records of human activity are weak in some ways and privileged in others. They are weak in that they systematically miss the experience of participants in the events being recorded, putting us as observers into the seemingly inescapable situation of always being perpendicular to the action. But video records are also privileged because they give us, as observers to the action, access to things that participants either miss or cannot do. That is, video records are plastic in ways that real-time experience is not: we can slow down a videotape, we can watch multiparty interactions repeatedly, we can make video and audio recordings that contain more or different things than any single participant could see or hear, and we can hold our judgments about the actions of participants in the record accountable to other observers.

Summary and a Note About Production Values

Across these brief comparisons, I have argued that technical arrangements for collecting video records can:

- Reorganize the tasks and experiences of research participants.
- Serve different research interests by selectively attending to different aspects of human activity.
- Reinforce or break open traditional boundaries between researchers and their study participants.
- Provide both limited and privileged access to aspects of human interaction.

Across different types of activity in learning and teaching, the way technology is used to capture primary data and our expectations about what these data need to show create records of phenomena that are both technology and theory laden. Primary data are technology laden in the sense that video and audio devices are deployed to record human activity in ways that

---

[7] C. Goodwin (1994) gave an extended discussion of how the same videotape can be used to "see" entirely different things, contrasting expert legal testimony given in the trial of Los Angeles police officers accused of beating an African American motorist (Rodney King).

make selections from ongoing interaction. Primary data are also theory laden in the sense that theoretical interests focus researchers on which parts of ongoing interaction are relevant, reliable, or usable given existing methods of analysis. Technical and theoretical constraints drive toward creating data records that show just those parts of interaction we already find interesting and little more. Although they can always be questioned, the boundaries of phenomena in research on learning and teaching are largely fixed at the time primary data are constructed.

Before closing this section, it is important to consider what production values we expect of video recordings, given that most of us have grown up with television and commercial film. By production values, I mean the way we judge the quality of a filmed record of human activity (e.g., Is it well lit? Can we hear it? Is there too much or too little activity? Are we close enough or too close to the action? Can we follow what people are saying?) and the way that these judgments have a tendency to creep into both what we choose to record and how we watch it. These values are implicit but pervasive. For example, up until the last several years, almost no one seen on television was ever shown watching television (Ehrenreich, 1988). It was, ironically, as if the fictional worlds we most wanted to watch were places where everyone had something better to do than waste their time watching television. As another example (also something that has started to change recently), we rarely see any embedded evidence that television or movie images have been constructed (e.g., images of recording equipment, rapid panning or zooming, etc.), as if the viewer had unmediated access to the events unfolding on film. We should be careful about how our expectations of film records of activity, shaped in these other genre, come into play when we construct and analyze records of learning and teaching.

A TERRAIN FOR (THINKING ABOUT) USING VIDEOTAPE

Most of the research projects that I work on are supported with public funds, so the data we collect in those projects could be considered a public resource. This is true both in the traditional sense of academics making their data available to one another for peer review and purposes of replication, but also in the sense that resources produced with public funds should be available to the public that underwrites them. This second sense of research as a public good is particularly important when the development of curricula, assessments, or software is

part of the funded activity. What parts of a research project should be a public resource, and for what purposes, are deeply contested questions.

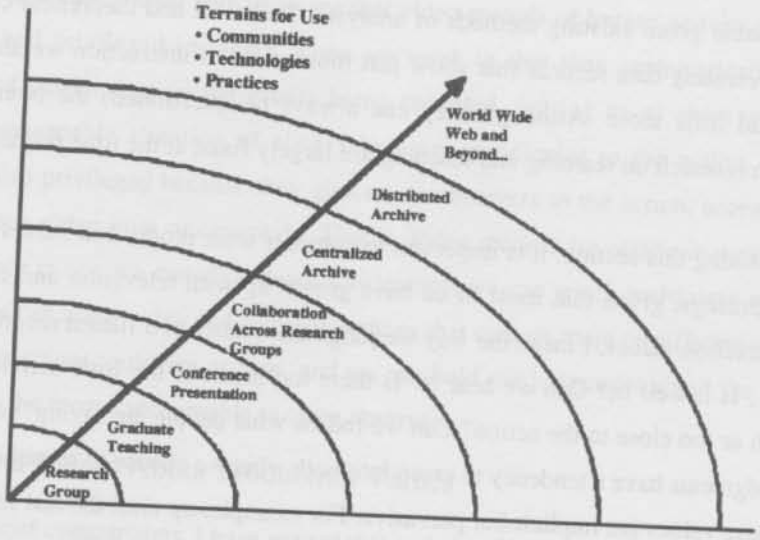

FIG. 22.6. A terrain showing contexts of use that are increasingly distant (and unpredictable) from a research project that produces video records.

FIG. 22.6 shows one way to think about records of primary data as public resources by following them as they travel inside and outside the research projects where they are collected and used.[8] The terrain for use is organized in concentric circles to show that, as one moves outward, the types of use (and users) expand, whereas the amount of information available about the actual activities of participants in a project usually contracts. Places in this terrain consist of user communities, their practices for manipulating and interpreting video records, and the technologies that (in some cases) support these practices. The cast is plural at each level out from the center, so the terrain is complicated. A small move across a boundary can bring about a very large change in the hypothetical (or actual) scenario of use.

---

[8] This view was constructed jointly by a subgroup discussing ethical uses of field data at a 1995 National Design Experiments Consortium meeting. The group included Maryl Gearhart, Janet Schofeld, Raul Zaritsky, Tammy Berman, and Eric Baumgarten. They should not, of course, be held responsible for what I have to say here.

In the previous section, I argued that collections of primary data about learning and teaching necessarily reflect deletions (or selections) brought on by technical arrangements and theoretical interests. Here, the question is: What happens when these constructed records move out from the communities that have collected them? We should immediately recognize this as a volatile mix—more people looking, with varied and unpredictable purposes, with less surrounding material to inform what they might find. The mix is particularly volatile when we talk about something as value laden as the reorganization of public schooling.

I visit three places across this terrain briefly:

- Use of primary data records within a research project.
- Using these records in a conference setting.
- Using these records to illustrate research issues and analytic categories in a public forum like the World Wide Web.

At each stop, I focus on what "informed consent" given by study participants might mean, and how new technical arrangements for recording and reporting the phenomena of learning and teaching create possible problems.

Trouble/Heterogeneity at Home

At the center of the terrain are research or development projects that collect, analyze, and preserve the kinds of primary data contrasted in the preceding section. Even within research projects, interests and commitments are heterogeneous. This is particularly true of research projects that seek the active participation of teachers and learners in developing, documenting, or analyzing instructional interventions. As these projects mature, the people who show up in the records treated as data are still present and sometimes are actively involved in the analysis of those data, in resulting design decisions, and in report writing. In this context for use, informed consent obtained at the outset of a project can be challenged by these (traditionally absent) participants. One response is to hold the ongoing analysis accountable to the concerns of study participants who live in both worlds (observed and observer), treating their views of the data and its interpretation as privileged accounts. Another response is to give study participants the right to selectively withdraw their consent as research subjects when their informed views of use change (e.g., to remove records that include them from a corpus under study), yet still to continue collaborating in other ways with the research project.

Trouble/Heterogeneity in the Research Community

As records of primary data travel into the research community, we need to make decisions about:

- What records to choose.
- How to provide surrounding context for these records so that they can be used in the ways we expect.
- How to limit or respond to unanticipated uses effectively.

Particularly with video records of teaching, the range of responses one gets at a public meeting seems to be enormous.[9] What one person sees as an artful interaction between a teacher and a student, another person takes as the basis for confirming generalizations about all of their worst experiences in public school. Under these circumstances, informed consent, taken at the outset of a study or renegotiated as that study unfolds, is probably inadequate to deal with challenges to reasonable use that arise outside a research project. One response is to obtain signed, restrictive agreements with viewers or users of video records in addition to releases from study participants.

Trouble/Heterogeneity Online

Perhaps all bets are off when records of primary data are allowed to circulate freely around the Internet, as are the text and still images in this chapter.[10] In the worst-case scenario, we should expect these records to be repurposed in ways that undermine the entire research undertaking, regardless of the kind of surrounding details we attach to the records. In the average case, particularly if "reality close" media (e.g., image and sound) are included as data to illustrate research settings or analytic categories, it is likely that only people who are already interested in these issues will seek out and use these illustrations. But informed consent and the ethical treatment of research participants isn't decided on the average. Copyright protections are often suggested as a solution to this problem. But in this context of use, if we consider the expense of pursuing cases of infringement, legal protections like

---

[9] Wendy Mackay (1994) described a variety of similarly challenging situations that appear regularly among researchers in human computer interaction.

[10] A draft of this chapter (along with others in this volume) was available through a password-protected Web site for a graduate course on research methods, organized by Dick Lesh at Purdue University.

copyright are probably more about protecting the rights and limiting the liabilities of research institutions and authors than they are about protecting research participants.

DISCUSSION

In this chapter, I have examined and then combined two issues. The first is how we should think about records of primary data we collect in research on learning and teaching. Most of us would agree that what we take as data are strongly theory laden, although the ways in which recording technologies implement these theories may be less obvious. New media for recording, analyzing, and reporting primary data necessarily delete or reorganize aspects of the original phenomena, even as they add new dimensions to what we can learn about human activity. I say this as a caution against taking these new media as being relatively complete, direct, or veridical.

The second issue carries this caution into the wider question raised by research that seeks to reform or reorganize public schooling: How can new technical arrangements for collecting, analyzing, and reporting research on teaching and learning be managed in ways that:

- Support collaboration across research projects.
- Convey vivid illustrations of complex instructional interventions to widely different audiences.
- Protect the rights and privacy of a heterogeneous group of study participants.

I have argued that our conventional understanding of informed consent breaks down as primary data move across wider contexts of use. One reason for this breakdown is that we cannot yet anticipate how different communities, technologies, and practices will take up and use these records. We need new practices as much as we need new media for doing and talking about research on learning and teaching.

ACKNOWLEDGMENTS

An early version of this chapter was given at a symposium on "A Discussion of the 'Multimedia Journal Article' Format as a Technology for Reporting Research and Evaluation Projects" at the annual meeting of the American Educational Research Association (San Francisco, 1995). I would like to thank members of that panel, particularly Lawrence Erlbaum

and Jan Hawkins, for helpful conversations about these ideas. I also would like to thank Michael Sipusic and Jeremy Roschelle for more general discussions about using videotape in educational research. The title and focus of this chapter borrows from Elinor Ochs' account of transcription as theory (Ochs, 1979). Support for this work came from a National Academy of Education/Spencer Foundation Postdoctoral Fellowship and a National Science Foundation grant (ESI-94552771) to the Math@Work Project conducted at the University of California, Berkeley.