THE JOURNAL OF THE LEARNING SCIENCES, *13*(1), 77–103

Ontological Innovation and the Role of Theory in Design Experiments

Andrea A. diSessa
Graduate School of Education
University of California, Berkeley

Paul Cobb
Peabody College
Vanderbilt University

The motivation for this article is our belief that theory is critically important but currently underplayed in design research studies. We seek to characterize and illustrate a genre of theorizing that seems to us strongly synergistic with design-based research. We begin by drawing contrasts with kinds of theory that are relevant but, we contend, by themselves inadequate. A central element of the type of productive design-based theorizing on which we focus is "ontological innovation," hypothesizing and developing explanatory constructs, new categories of things in the world that help explain how it works. A key criterion to which we adhere when discussing ontological innovations is that theory must do real design work in generating, selecting and validating design alternatives at the level at which they are consequential for learning. Developing and refining an ontological innovation is challenging and requires the kind of extensive, iterative work that characterizes design experiments more generally. However, the pay-off in terms of clarity of focus and explanatory power can be great. We present two case studies that illustrate the development, refinement, extension, and instructional application of ontological innovations.

The nature of theory and its role in educational studies and in design research in particular is a deep topic deserving of extended discussion. It has frequently been argued that the development of theory should be one of the primary goals of design research (e.g., Collins, Joseph, & Bielaczyc, this issue; Confrey & Lachance,

Correspondence and requests for reprints should be sent to Andrea A. diSessa, University of California at Berkeley, Graduate School of Education, 4533 Tolman Hall, #1670, Berkeley, CA 94720–1670. E-mail: disessa@soe.berkeley.edu

2000; Design-Based Research Collaborative, 2003; Edelson, 2002; Gravemeijer, 1994a; Labato, 2003; Simon, 2000). Despite this recognition of the potential contributions of theory, we contend that design experiments have generally been underdeveloped as contexts for the development of theory. Our primary purpose in this article is to enrich ongoing discussion of the role of theory in design research with some new avenues of consideration.

On the vast background of the nature and function of theories, the goals of this article are modest. We cannot hope for definitive arguments in favor of theory, nor for an extensive characterization of how theories are built (or should be built), tested, and how they work. Instead, we aim only to accomplish the following: We will argue briefly that theory is important. Then, our principal thrust will be to characterize roughly elements of a genre of theory that has been underplayed in design experiments but that can, we believe, be productively developed, tested, revised, and elaborated while conducting such studies. We initially characterize this type of theory negatively, by contrasting other meanings for "theory." On the positive side, we concentrate on one distinguishing process in developing the kinds of theories in which we are interested, that of "ontological innovation"—the invention of new scientific categories, specifically categories that do useful work in generating, selecting among, and assessing design alternatives. Examples of ontological innovations developed and refined while conducting programs of design research include elements of the theory of quantitative reasoning proposed by Thompson and colleagues (Thompson & Thompson, 1994, 1996) and the constructs of model-of and model-for formulated by Gravemeijer (1994b). Examples of ontological innovations that were developed in other disciplines but that have been used widely in educational research include the constellation of constructs that comprise Lave's (1988) and Wenger's (1998) situated theory of learning and the notion of an activity system as formulated by Engestrom (1998, 1999).

Discussions of the nature of science, and of theories in particular, tend to suffer from a high degree of abstractness. Meta-scientific frameworks, which might facilitate description and comparison of theories, are not common currency. This difficulty is exacerbated in design research, where investigators tend to follow their noses, doing the work of science as they understand it, without extended rationale or public explication. Compounding the problem is the fact that design research is a relatively new and still evolving methodology. The community of researchers has not had time to filter and share reflections on the phenomenology of "doing science" in this way. As a consequence, the delineation of agreed upon practices and the development of a grounded language for describing them is very much in progress.

To combat abstractness, we trim our coverage and depth goals further to make room for two case study descriptions of our own practices of theory development and use. Our intent in presenting these case studies is to exemplify and concretize the discussion of ontological innovations in the context of design research.

THE IMPORTANCE OF THEORY

In many corners of science, the importance of theory is completely uncontested. The history of physics, for example, may be viewed profitably as a series of broader and deeper theories, each of which was complemented, refined, or subsumed in successor theories. Newton's theory was substantially modified by Einstein's special relativity. Quantum mechanics located Newtonian mechanics as an approximation relevant to everyday sized objects and everyday motions. Later, field theory integrated special relativity and quantum mechanics.

Theories have always displayed a principal part of the power and elegance of science. They embody generalization, bringing order to a vast array of seemingly disparate phenomena that come to be seen as special cases of some theory. They encapsulate the most secure of our knowledge claims at any stage of scientific advance. They enable us to discriminate between relations that are necessary and those that are contingent. They delineate classes of phenomena that are worthy of inquiry and specify how to look and what to see in order to understand them. This last characteristic—epigrammatically, "teaching us how to see"—is particularly evident in the type of theory on which we focus in this article.

Theory even provides a suggestive framework for understanding sciences' evolution; once a theory is initially established, scientists typically (1) follow up details and special cases, (2) adapt theories to application, and (3) seek to account for anomalies that, sometimes, result in the creation of new alternative, subsuming, or complementary theories.

Education seems a different matter. No psychological or sociocultural theory has had the breadth, precision and detailed validation of any of a number of theories in physical or biological sciences. Theories concerning educational matters seem to replace one another, rather than subsume, extend, or complement other theories. Although the state of the art constantly changes, it is often difficult to tell that progress is being made. Indeed, some have been drawn to argue that human sciences are simply not amenable to theory building as we know it in other sciences (e.g., see the discussion of "theoretical frameworks" and references in Eisenhart, 1991).

We see no reason to believe that the human sciences are essentially different from other sciences with respect to the value of theory. All well-developed current sciences went through early stages where progress was more chaotic and uncertain. When things settled down, theories and theoretical development have always proved central to progress. Indeed, the complexity of extended real-world interventions that characterize design studies highlights the pressing need for theory while simultaneously making the development of useful theories more difficult. When working in the multifaceted or even seemingly chaotic settings in which design studies are conducted, one must have some orientation on central versus peripheral concerns, and one must be very clear on what general results are intended.

In our view, the apparent complexity and messiness of human action and learning in these settings indicate our collective limited theoretical sophistication.

One thing is certain: If we do not pursue theories that enable us to come to grips with the complexity of design experiments, we will not achieve them. The central problem, as we see it, is that of how best to pursue theoretical agendas, befitting what we believe we know so far and acknowledging whatever special circumstances are our lot in education-related sciences.

CHARACTERIZING "THEORY" FOR DESIGN STUDIES

Design studies, or design experiments, are iterative, situated, and theory-based attempts simultaneously to understand and improve educational processes (Brown, 1992; Cobb, Confrey, diSessa, Lehrer, & Schauble, 2003; Collins, 1992; Edelson, 2002). Theory can, however, mean different things to different people. What kinds of theories can be produced by and can serve to further the aims of design experiments? Among the many possible criteria for types of theories, one stands out as critical for design studies. Theory must do real design work in generating, selecting and validating design alternatives at the level at which they are consequential for learning.[1] With this criterion in mind, we identify several types of theory that can be of value when conducting design studies. However, we argue that, by themselves, they constitute an inadequate theoretical basis for design research in the long term.

Grand Theories

"Grand" theories of cognitive or social processes are the educational equivalent of "evolution" or "Newtonian mechanics." Examples include Piaget's theory of intellectual development (Gruber & Voneche, 1977), Skinner's behaviorist theory of learning (Skinner, 1974), or Newell's "unified theories of cognition" (Newell, 1990). The difficulties with such theories lie typically in some combination of being, as yet, immature (e.g., false as categorical prescriptions of cognitive or social processes), imprecise (so that implications at the level of design decisions are unsure), or simply too high-level to inform the vast majority of consequential decisions in creating good instruction. To take a specific case, Piaget developed his theory to address epistemological issues that concern the nature and growth of knowledge. Nonetheless, his work had a strong educational influence from the 1960s at least through the early 90s. We feel, as many others do, that Piaget's ideas

[1]Learning, per se, is a canonical goal for educational-related design studies. However, alternative goals certainly exist, including engaged participation, community development, development of a personal scientific identity, and so on. To simplify discussion, we take learning here as a paradigmatic goal, and support for learning as the relevant design focus.

were overextended into education. Indeed, Piaget felt that way himself. Attempts to apply his ideas have often resulted in unwarranted generalizations of what were actually more focused empirical results (Metz, 1995). Furthermore as Walkerdine (1988) documents, readings of Piaget's ideas were primarily used to rule out approaches to learning and instruction ("young children cannot consider scientific explanation because they are concrete thinkers"), rather than to manage design decisions among possibilities in ultimately effective instruction.[2] Piaget's theory is powerful and continues to be an important source of insight. However, it was not developed with the intention of informing design and is inadequate, by itself, to do so deeply and effectively.

Orienting Frameworks

"Constructivist theory" or "cultural-historical theory" are often appealed to as the basis for instructional design. On the one hand, commitments at this level are highly consequential in identifying family resemblances and perhaps core values among those who do design research. They define important social affiliations among scientists and name widespread thrusts into thinking about and carrying out instruction. Yet, much that is involved in developing specific designs does not follow from these orientations. What happens when "building on student ideas" seems to fail to work? How, exactly, should one build on student ideas? Communities of learners are sometimes not optimal learning cultures when they are patterned on a limited range of exemplars that emphasize only induction into participation. Some cultural artifacts are much more powerful for learning than others; how do we measure limitations and improve elements that are weak?

Orienting frameworks seldom provide strong constraints or detailed prescriptions. Their value instead resides in the general perspectives that they provide for conceptualizing issues of learning, teaching, and instructional design. Principles of refinement and the rectification of failure, the meat of much design, typically require detailed specification within such a perspective. Orienting frameworks are probably best viewed as meta-theories, presumed general constraints that define general aspects of hoped for and needed more specific theoretical frameworks.

Frameworks for Action

Another common usage of "theory" refers to more or less general prescriptions of pedagogical strategies. Examples include Papert's "constructionism" (Papert &

[2]A fuller treatment would require attention to recent attempts at instructional design that appeal to grand theory in selecting many specific features of instructional interventions, e.g., Anderson (1990). Without meaning to appear dismissive, for lack of space, we say only that there is no wide-spread acceptance of such theories and their consequences in educational design.

Harel, 1991), Ann Brown and Joe Campione's "fostering a community of learners" (Brown & Campione, 1996) and Ann-Marie Palincsar's "reciprocal teaching" (Palincsar & Brown, 1984). Marcia Linn talks about "scaffolded knowledge integration" (Linn & Hsi, 2000). In some of our own work (diSessa, 1992) we advocate "learning by designing" as a powerful framework for instruction. These frameworks certainly provide some focus and direction to the design of learning environments, and they are often effective heuristics. However, such frameworks themselves typically fail to serve the role of theory in the sense that we believe to be most important for one central reason: They do not cleanly separate their scientific claims and validation from their suggested actions. That is, the theory or theories behind frameworks for action are relatively inexplicit, complex, and often involve multiple very diverse elements that cannot plausibly be brought under a single umbrella. In the case of students' learning through design, we argue that it is attractive because it combines good affective properties with good cognitive ones. The details of the separate theories of cognition or affect are not really at stake.

We introduce the phrase "managing the gap" to name the issue that is behind the failure of most frameworks for action to achieve what we would like accomplished. The "gap" arises from the fact that instruction is the result of many sorts of complex, interacting elements. Instruction depends on the values of the participants; it depends on technological infrastructure; it depends on the nature of classroom discourse; it depends on practicalities such as available time. We also want to *make* instruction both depend on and serve to test theory. And yet, in order to see and assess the impact of underlying theory, we must cleanly separate it from the myriad of other issues that we handle, as best we can, in the management of trade-offs among the multiple constraints impinging on instruction. In the ideal case, then, pedagogical strategies and conjectures are separated by a *carefully considered and articulated gap* from the theory or theories that explain or motivate them. A well-managed gap separates the implications of a particular theoretical claim from other claims and also from atheoretical aspects of design. Attention and effort are necessary to perform this management.

To take a noneducational example, there can be no doubt that there is science in the design of airplanes. However, the shape of a Boeing 747 aircraft does not follow in a direct and simple way from any of this science. Neither does the shape of the aircraft, as a whole, directly test elements of the underlying theory. With sufficient care (corresponding to managing the gap between design and theory), however, a failure attributable to the shape of the aircraft might implicate a failure of a theory of strength of materials, not just to a careless mistake, a failure to anticipate transient loads, or a poor choice of materials.

As with grand theories and orienting frameworks, our point is not to denigrate frameworks for action. They play a critical role in organizing research about instruction given the current state of the art, and they will probably continue to do so for the foreseeable future. Yet, their purpose is more directly practical and less directly connected to core scientific values, such as precision, exhaustive and explicit formulation, and falsifiability. Frameworks for action specify goals and actions in instruction; the type of theory with which we are concerned typically lies "on the other side of the gap," where entities and relations are specified, but actions, goals, and instrumentality are much less so.

Domain Specific Instructional Theories

When people refer to theories in the context of design experiments, they are often referring to *domain specific instructional theories*. Theories of this type entail the conceptual analysis of a significant disciplinary idea (e.g., density, sampling distribution) together with the specification of both successive patterns of reasoning and the means of supporting their emergence. Means of support are usually construed broadly and can include the classroom activity structure and discourse as well as instructional tasks and tools. In contrast to frameworks for action, theories of this type are of more than heuristic value and embody testable conjectures about both learning processes and the means of "engineering" them. Furthermore, although they are typically developed within an orienting framework, they reflect the detailed work involved in formulating, testing, and revising a hypothetical learning trajectory.

As we indicated, design studies involve iterative cycles of design and analysis. Domain specific instructional theories represent an advance in the design aspect of design research. In fact, we can view them as constituting the rationale for pedagogical designs; generalization across a number of such theories can result in the development of a framework for action such as that of Realistic Mathematics Education developed at the Freudenthal Institute in the Netherlands (cf. Streefland 1991; Treffers, 1987). However, the fact that domain specific instructional theories are tied to specific designs reflects the emphasis on relatively situated and pragmatic concerns at the expense of managing the gap between explanatory theory and pedagogical conjecture.

Our intent in making these observations is, again, not to disparage theories of this type, particularly as their development constitutes one of the primary reasons why we conduct design studies. Instead, it is to highlight that design research will not be particularly progressive in the long run if the motivation for conducting experiments is restricted to that of producing domain specific instructional theories. It bears noting, for example, that although such theories are informed by repeated cycles of design and analysis, the analyses are typically conducted by using established theoretical constructs. As a consequence of the concerns that motivate their development—including communicating design rationale to teachers—domain specific instructional theories will not by themselves enable us to scrutinize and reconceptualize aspects of competence in mathematics or science, or to better handle the complexity and diversity of settings in which we conduct experiments.

Beyond narrowing our consideration by contrasting the cases of grand theory, orienting frameworks, frameworks for action, and domain specific instructional theories, can we be more affirmative in characterizing a type of theory that can be important for design work? Of course, a huge amount could and should be said, but here we concentrate on one important element, ontological innovation.

ONTOLOGICAL INNOVATION

The idea behind ontological innovation is deceptively simple. Science needs its own set of terms or categories to pursue its work. Again, this has always been true of developed sciences: "force," "gene," "natural selection," "molecule," "element," "catalyst." The process of creating such categories, however, is far more complicated than writing down a definition, or finding a relevant meaning in a dictionary. Instead, defining the technical terms of science is more like finding and validating a new category of existence in the world; hence we use the term *ontological innovation*.

The essential challenge can be expressed simply enough. Scientific terms must "cut nature at its joints." That is, they must make distinctions that really make a difference, ignore the ones that prove to be inconsequential, and enable us to deepen our explanations of the phenomena of interest. We must develop theoretical constructs that empower us to see order, pattern, and regularity in the complex settings in which we conduct design experiments. Ontological innovations are attributions we make to the world that necessarily participate in our deepest explanatory frameworks.

We use two examples to illustrate the notion of ontological innovation and some intrinsic difficulties ("bootstrapping") in developing one. A common laboratory experiment has students measure a varying voltage and consequent current in a circuit. The students "discover" Ohm's law, that current (I) is proportional to voltage (V), "I = (1/R) V" (the proportionality constant is the inverse of resistance). The task of a real scientist is much different. The terms current, voltage and resistance must first be conceptualized and operationalized. The bootstrapping process is complex. It turns out that meter mechanisms mainly measure current. But we cannot even know this until we have operationalized current in a way other than "that which proportionally drives meter mechanics"! Further, how do we measure voltage? A modern scientist could use Ohm's law and a known resistance together with a current-sensitive meter mechanism. The problem is obvious: knowing Ohm's law is by far the easiest way to know how to measure (in this case, voltage), in order to "discover" Ohm's law.

Measurement is a key practice in education as well as in the physical sciences, so it is worth pursuing our metaphor just a bit further. Voltage must be measured with a two-wire probe. Why not a single wire? Again, conceptualization appears necessary prior to measurement, which seems prior to validating the conceptualization. Measuring current requires one to break a circuit and insert the meter, and it requires the meter not add significant resistance to the circuit. How would one know that? The full theory appears necessary in order to measure, in order to test the theory.

All of these bootstrapping problems apply to design studies. The last of our observations—constrained and careful breaking of a circuit in order to tap theoretical properties necessary to understand it—is particularly evocative. Not only do theoretical terms provide us with explanatory constructs that we use to understand phenomena, but they guide us in the tricky interventions necessary even to "see" theory-relevant things. In the student exercise of "discovering Ohm's law" that started this example, students just insert a probe as directed, and black-box mechanisms read out relevant things that go together exactly as the theory says they should. Scientists in the process of developing theory somehow perform operations to see things that they are just beginning to conceptualize.

Our second example undoes an accidental salience in the first, that of precise, quantitative measures in accomplishing science. This example makes the point that scientific categories do not always need to have quantitative properties to provide for strong implications and falsifiability.

Consider the case of the germ theory of disease: tiny, invisible, biological elements cause human illness. No precise measurements or quantitative laws are at stake, and the exact nature of those biological elements, although certainly within the scientific program, is not relevant to some powerful theory-based predictions, both practical and intrascientific. To wit, if we isolate sick people from well people, the well people will not get sick. If illness truly develops spontaneously, the theory is refuted. Turning to practical implications, we might be able to "kill" those invisible elements, say with heat or some kind of poison. We might Pasteurize, or conceptualize what a mold can do for some diseases.

The germ theory of disease provides a clear *locus for refinement*. What exactly are those tiny elements, and how do they work? Some of them (bacteria), it turns out, are full-fledged organisms susceptible to many killing methods. Some (viruses) have quite peculiar ontological characteristics, and "killing them" or preventing their spread becomes subtle and theoretically intensive. The concept of locus for refinement is important, as, unlike theories that are treated as *faits accomplis*, the development of ontogocial innovations in design-based research seems inevitably a long process of refinement and extension. This was very much the case for the innovations developed by Thompson and colleagues (Thompson & Thompson, 1994, 1996) as well as in our own work.

In the following two sections, we provide two brief case studies of design experiments that led to, and then developed, ontological innovations. The main elements of our analysis are to sketch (a) the origins and nature of ontological innovation, (b) their impact on the conceptualization of learning and on subsequent educational design, and (c) the processes of validation and refinement, and what that further precipitated in design. When possible, we describe the need for and our efforts to manage the gap between theory and the details of instruction, and, in addition, we bring out a few relevant and interesting properties of design experiments, per se.

CASE STUDY I: META-REPRESENTATIONAL COMPETENCE

Opportunistic Research, Retrospective Analysis, and Discovery

In 1989 and 1990 the Boxer group at UC Berkeley was in the midst of a design experiment aimed at teaching elementary school students some high school and university physics (diSessa, 1995). We won't recount the motivations or results of that experiment because the focus of our attention here is something new that emerged from that context. In fact, what we observed was quite unexpected. It concerned what we subsequently came to refer to as the students' *meta-representational competence*: their substantial expertise in inventing, evaluating, and refining a variety of representational forms. Although these observations did not fit cleanly within our existing scientific framework, they did fit beautifully within our orientating constructivist framework, and also remarkably well in our framework for action, students' design. The fact that what we discovered was entirely consistent with our orienting and action frameworks, but quite distinct scientifically, exemplifies the contrasts we made in introducing theory-for-design.

Discovery of new foci of attention is actually quite common in design experiments, if it is not absolutely necessary. Failures or surprising successes not infrequently push toward, and sometimes enable, new lines of inquiry, possibly involving new ontologies. In some other cases, typically in initial failures, we manage the gap by patching enough to get by, without pulling the surprising occurrence into the core, scientific program. For the discovery relevant here, we had no means to pursue the new theoretical issues by any particular actions during the project in which it emerged.

What we could do within the initial design experiment was to conduct a *retrospective analysis* of our "surprising occurrence," and, from that, get a head start on follow-up work aimed specifically at clarifying what we had discovered. Retrospective analyses more frequently occur closer to the main scientific line of an experiment, to get at microstructure related to conjectures that might have been roughly validated, or to study unanticipated, theory-relevant failure. For example, a post test might show that students learned something, but one might want to backtrack to see in more detail how that happened, say, in classroom discussion.

How did we come to have the video that allowed a retrospective analysis of an unanticipated event? It was part luck and part typical design-experiment discipline. Design experiments usually encompass extended periods of time and much activity that is not highly scripted and/or it might be outside the focus of principal conjectures (constituting a somewhat uncontrolled part of the "gap" referred to previously). During many of these times, nothing particularly interesting or relevant occurs. It is manifestly impossible to study everything that happens in a design experiment, and trying to do so exhibits lack of scientific focus. Systematic analysis has its virtues, but researchers have finite time and hopefully the wisdom to study mostly issues that are well prepared and tractable.

On the other hand, surprising things essentially always happen, and some data collection beyond core focus is always a good idea, pending trade-offs of time and effort with other tasks. In our experience, much of that data never gets analyzed. Many times we decide in advance on principles of data selection. For example, we have often decided to tape every full-class discussion, even if nothing critical is anticipated.

We did not have that particular principle in action during the case in point—which would have led to a corpus of hundreds of videos. The lucky part of having video of our surprising event was that an alert graduate student noticed some interesting things going on in a warm-up activity, so he suggested we video the full activity.

The task as set was an introduction to graphing. Instead of jumping into instruction, we had time for and decided to implement a more exploratory exercise of seeing what our students knew about graphing. At least, that is what we thought we were about. We asked students to design their own representations of motion, expecting that some might produce graphs and show us how they made them. The orienting framework suggested studying student ideas; one major course-specific frame for action, student design, suggested posing a design task.

An explosion of ideas resulted. Students produced and critiqued dozens of representational forms. None of the first many representations was graphing. Instead, graphing emerged clearly only very late in the activity, about two-thirds through the 4 to 5 day activity (40 minutes per day). What was most salient for us was the almost uniformly intense engagement on the part of the students. Sometimes the teacher couldn't even find her way into the conversation. Students came to class with ideas, and they went out continuing to discuss possibilities that had arisen.

Had we designed such a rich and enthusiastic exploration, we would have been extraordinarily pleased. Having it presented to us with little preplanning, we were stunned. Although not in our game plan, we felt we had to follow up. We set our agenda to understand how it was possible for this activity to arise. More particularly:

1. What knowledge accounted for the contributions of students?
2. What contributions did the teacher make? To what extent were her actions critical to the outcomes?
3. How did ideas build on one another: For example, did they flow logically, or accumulate in an ad hoc manner? How did students actually get to graphing?

4. What conceptual knowledge of motion (the designated learning target of the course) was implicated in what the students did?

Briefly, we found (diSessa, Hammer, Sherin, & Kolpakowski, 1991):

1. Students seemed to have a wealth of ideas about how to create and judge representations. The logic of the representational forms that were created was, for the most part, scientifically appropriate. Students saw many good and bad properties of the representations they produced, and they were creative in patching problems that were uncovered.
2. We note that teacher interventions constitute a potential competitor to "student knowledge" as an answer to the question "how did students manage this?" It became quickly evident that many student creations were not and could not have been contributed by, or instructed by the teacher. In fact, she did not understand several of the representations initially when they were produced. With respect to the narrower question, did the teacher make graphing appear, our analysis suggested that while she pushed criteria of judgment that would favor graphing, she did not either introduce key ideas, nor coerce student judgments that "graphing is best" in the end. The primary outcome concerning teacher strategies was to note the extraordinary flexibility of this teacher's strategies, which were applied in a highly contextualized manner.
3. Students combined ideas from other students fluidly and sensibly. With respect to graphing, per se, the trajectory had both elements of inevitability, and elements of idiosyncrasy. Students seemed to have or develop all the basic resources needed to understand graphing, and it matched their criteria of judgment as a "good representation" well. Yet, the particular path of invention involved probably irreproducible specificity and timing of conceptual contributions by students.
4. The study of conceptual knowledge of motion was treated separately in an article that set conceptual development within another framework for action, "benchmark lessons" (diSessa & Minstrell, 1998). The analysis led to some interesting conclusions and speculations about new paths of conceptual development (possibly contributing to a domain specific instructional theory), but, in fact, conceptual development concerning motion seemed nearly orthogonal to the main line of representational design.

Ontological Innovation: Meta-Representational Competence

The major result of our retrospective analysis was the "discovery" of *meta-representational competence* or simply MRC. MRC, as an ontological innovation, implicates a specific body of knowledge that students have, and which can be developed, lying behind students' abilities to *create, critique,* and *adapt* a very wide range of effective scientific representations. Elements of our retrospective analysis, cited above, suggested that students' prior knowledge, as distinct from "what the teacher taught," lay behind MRC (item 2), and, furthermore, that that knowledge was also distinct from other bodies of knowledge, such as conceptual knowledge of motion (item 4). "Discovering MRC," of course, is a very substantial simplification of the process, which seemed to us consistent with the bootstrapping of conjectures on the nature of the entity, and methods of "seeing" and measuring it, which we introduced with the "discovery" of Ohm's law.

Three points are worth emphasizing about this ontological innovation. First, we feel MRC is a genuinely new construction. Although several bodies of literature relate to it, our formulation is different in some important respects from any of them. Developmental psychology has studied children's representational competence, but treated it as slowly changing and not nearly as differentiated (see below), richly generative, and capable of supporting midduration (several days or weeks) development (contrast both a "lab task" on the short time-scale end, or a "developmental level" on the long time-scale end). Prior instructional work that has highlighted student representational creativity (e.g., Cadwell, 2003) has not led to a fine-grained analysis of the specific knowledge underlying that competence, and it has usually been disjoint from consideration of the acquisition of competence with scientific representations.

Second, we emphasize the analogy with the germ theory of disease. The existence of MRC can be a substantial discovery at a relatively undifferentiated level. So, for example, we believe that essentially any teacher can reproduce at least some aspects of the creativity and energy from students that drew our attention in the first place, given a little help in understanding what MRC is and how one acts to evoke it. But, in addition, MRC sets a locus for refinement that can bring additional science and practical impact. See continuing discussion, later in this article.

Third, like essentially all scientific discoveries, instrumental impact is not easy to lay out. Indeed, we expect a lot of creative design will be needed to add to MRC's purely scientific properties in the service of good instruction. The gap between theory and practical action is complex, and it is not bridged merely with more science. In our case, we conjecture two important functional niches for MRC. First, it is a valid target of instruction in its own right, worthy of its own slot in the curriculum. Making such a judgment implicates values (what should be taught) not just science (what students know or can come to know). Our second functional niche is as a frame and background for the study of the many specific scientific representations that one finds now in conventional curriculum. Simply put, particular representations may be more learnable and sensible to students in the context of understanding what representations do in general, how they work, and what alternatives might exist for given ones (like graphing). We have taken small steps in both of these directions, but the agenda is far from finished.

Refinement, Iteration, and Extension

This section emphasizes MRC as a locus for refinement. Over a roughly 5-year period, we followed up on our initial discovery in several directions. First, we developed a preliminary taxonomy of subsystems of MRC, including the two most obvious parts mentioned above, inventive resources (how do students make new representations?), and critique (how are representations judged?). Inventive resources were amenable to study in additional examples of full-class design, very much paralleling our original miniexperiment; one sees direct evidence of inventive resources in the physical representations students produce (Sherin, 2000). For example, students regularly and easily apply a particular schema for time-based representations, where time is displayed as a sequence of separate icons or images. In contrast, the breadth and nature of students' criteria of judgment were difficult to study in these instructional contexts, so we did a lab experiment with a small number of students, seeking to draw out and study criteria (diSessa, 2002). (We had to cut the fabric of in-class learning to see student criteria, as one needs to cut a circuit to measure current.) Design experiments seem intrinsically to call for multiple methods and naturalistic study is not *sine qua non* (Brown, 1992).

Iteration of the "inventing graphing" classroom activity has been critical in separating accidental from intrinsic properties of MRC. In the years since, we have undertaken about a half-dozen "replications" in various contexts. Typical of design experiment methodology, replication does not mean literally trying to do exactly the same thing again, but exploring variations and what stays the same (and what changes) on redoing an activity. Briefly, we have found:

1. MRC has shown through strongly in every iteration, with a broad range of students across multiple age levels. Indeed, the kinds of representations our first group of students did not nearly exhaust the possibilities.
2. While our first try generated truly exceptional interest and engagement, every iteration also showed substantial student involvement.
3. While our first teacher continued to serve as a benchmark for scaffolding expertise, teachers with very different styles can still find their way effectively in conducting these sessions.

We initiated a study of MRC in a neighboring domain, not concerning representing motion, but with respect to representing spatially distributed data, for example, altitude across some terrain, or brightness across a photograph. This constituted a sort of "transfer" test for the concept of MRC. While we never believed MRC was completely domain-independent, it was important to get some calibration. Is motion a special ground of representational competence?

Briefly, we discovered that MRC does basically carry over to spatially distributed data, with some important caveats. Students were quite competent in producing a fairly wide range of cogent representation; they were relatively good at judging their value; and they even spontaneously developed approximations of prototypical scientific representations (topographic maps, roughly the equivalent of "graphing" with respect to spatially distributed data). The caveats, however, are illuminating. First, students' MRC does not support the same kind of enthusiastic, extended engagement with design that we encountered with motion. In recent work, we feel we uncovered some reasons for this (Sherin, Azevedo, & diSessa, in press). Second, students' "representational misconceptions" seem to be more prominent with respect to spatially distributed data (Friedman, 1999; Elby, 2000). As a result, we developed and tested instructional strategies for supporting better competence for dealing with certain representational technologies (image processing systems; see Friedman & diSessa, 1999).

A final note: Because MRC was not anything that we expected to run into, a significant part of our work in following up was locating and positioning our work with respect to various bodies of literature concerning general representational competence. See, particularly, Sherin (2000), which situated our work with respect to developmental studies of representational competence. Scientific work requires such accountability, and it is one reason design researchers need to limit their attention to a limited number of theoretical foci in any given experiment, relegating many of the other issues that arise to a practical managing of the gap.

Synthesis and Extension

Ontological innovation often provides at least as much of a basis for further work as it does "results," per se. MRC proved a solid basis for a progressing paradigm of scientific research: It encompasses a wide and varying domain of phenomenology; it suggested the possibility of new instructional goals (elements of MRC, itself; representational design with respect to different classes of representations), and new measures of success (high-level understanding of representations, not simply learning "additional particular scientific representations"); finally, it interpreted and helped us deal with educational difficulties such as representational misconceptions (Elby, 2000) and how to conceptualize problems we discovered in students' use of technology (Friedman & diSessa, 1999). A recent article, diSessa (inpress), attempts to synthesize the major lessons learned in our studies of MRC, and it sets a further agenda for research.

In this section, we briefly highlight productive connections between MRC and other research agendas. Ontological innovation "bleeds over" to related studies, either demanding complementary study (which is more prominent in our second case study, below), or contextualizing and contributing to other research foci.

MRC has proved a good context in which to pursue the study of teacher strategies. Recall, this was one focus of the initial MRC retrospective analysis. MRC is highly generative knowledge, in contrast to strategic or "structural" understanding such as understanding particular scientific concepts. Getting students to design and critique novel representations is a prototypical task for learning and also a test of advancement. Yet, such tasks have had comparatively little study, in contrast to, for example, lessons focusing on structural scientific understanding (e.g., jigsaw lessons to illuminate function/structure relations in ecology). How does one guide student design without stifling it? The Boxer group produced some relevant studies. Madanes (1997) developed the concept of critical pedagogical moves, moves that provide an essential leap in the focus or flow of inquiry. MRC strongly contextualizes such moves. What counts as an "essential leap" must be calibrated by understanding the development of specific competence, in this case, representational competence. Granados (2000) looked at the problem of establishing an appropriate intersubjective agreement among teacher and students on the range and limits of design tasks. diSessa and Minstrell (1998) looked at MRC-related data to elaborate the meaning of a particular genre of somewhat more teacher-driven lesson (compared to open design), *benchmark lessons*.

Finally, the Boxer group has produced a small set of articles trying to understand interest and engagement, the overt reasons we started our study of "inventing graphing" in the first place. In Sherin, Azevedo, and diSessa (in press), we identify structural differences in MRC related to different domains (e.g., MRC related to motion and MRC related to spatially distributed data), which account, in part, for differences in student engagement and enthusiasm. Azevedo (submitted) attends to student personalization of their efforts in MRC-based learning tasks, implicating "personal excursions" as highly functional in allowing students strong personal links, while maintaining a focus on "what the students are supposed to be learning." Again, the point of mentioning these is to illustrate how ontological innovation can open up new avenues of pursuing older questions, or how it provides synergistic contextualization of complementary studies.

CASE STUDY 2: SOCIOMATHEMATICAL NORMS

Opportunistic Research, Reorientation, and Retrospective Analysis

In 1986, the research team of Paul Cobb, Erna Yackel, and Terry Wood began a series of year-long design experiments in collaboration with first, second, and third grade teachers. The second classroom session of the very first of these experiments involved the observation of an unexpected phenomenon that precipitated a new strand to our research program that has continued for 17 years. Briefly, we intended to focus on young children's mathematical learning primarily from a cognitive perspective. However, the teacher with whom we collaborated initiated a shift in discourse during the second whole class discussion from talking with the students about their solutions to explicitly talking with them about how they should talk about their solutions. Our investigation of shifts of this type, which we subsequently came to call the *renegotiation of classroom social norms*, led to an explicit focus on classroom discourse and communication. Thus, a single surprising observation led us to reorient our research program on the fly and led to a series of investigations, each of which spawned new issues that required us to extend our prior analyses. Our decision to video-record all classroom sessions made it possible for us to modify our research focus and to pursue issues that arose opportunistically by conducting retrospective analyses, similar to what happened in the meta-representational competence case study.

As part of this research program, Erna Yackel observed that the teachers with whom we worked frequently asked if anyone had solved a task a different way and began analyzing teachers' and students' interactions in these situations. This investigation turned out to be particularly fruitful in that Yackel clarified and illustrated how teachers and their students negotiated the norm of what counts as a different mathematical solution. This analysis also involved a significant advance over our initial work on classroom social norms in that the latter are not specific to mathematics. The social norm that students should make different contributions to discussions is applicable to subject matter areas other than mathematics. In contrast, the norm of what counts as a *mathematically* different solution is specific to mathematics. The analysis of mathematical difference therefore precipitated an ontological innovation—sociomathematical norms, of which mathematical difference is one—that overcame a limitation in our analyses of classroom social norms.

In explicating the norm of mathematical difference, we made a number of observations (Yackel & Cobb, 1996):

1. The norm of mathematical difference established in a particular classroom is interactively constituted by the teacher and students, and is thus a joint accomplishment. This theoretical claim holds for all classrooms including those in which that teacher acts as an overt mathematical authority and has a predetermined view of what will count as mathematically different. Even in these cases, the students have to play their part by coming to act in accord with the teacher's expectations. Consequently, this norm does not correspond to an overlap in the teacher's and students' individual understandings. Instead, it is located within the ontological domain of social interaction and communication.

2. What counts as a mathematical difference can differ markedly from one classroom to another, and this can significantly influence the nature of the learning opportunities that arise for students. For example, in some first grade classrooms, a number of different types of counting methods for finding sums and differences count as different, whereas in other classrooms in which we have worked, counting methods are treated as the same and are differentiated from various thinking or derived-fact solutions.

3. The establishment of the norm of mathematical difference can give rise to learning opportunities for students during whole class discussions that otherwise would not have arisen. To make contributions that count as different, students have to make their own solutions objects of reflection in order to assess similarities and differences with other contributions.

4. The norm of mathematical difference can also give rise to learning opportunities for teachers. This is particularly the case when they do not have prior experience asking students to explain their reasoning and therefore have little basis for anticipating either students' solutions or the mathematically significant distinctions between them.

Ontological Innovation: Sociomathematical Norms

We came to realize that in analyzing the norm of mathematical difference, we were drawing attention to an aspect of the hidden curriculum of mathematics classrooms. This realization led us to focus explicitly on what students had to know in action in order to participate effectively in whole class discussions conducted by the collaborating teachers. In doing so, we identified several other specifically mathematical norms whose negotiation is typically less explicit than that of mathematical difference. These include the norms of what counts as a sophisticated mathematical solution, an acceptable mathematical explanation, and a clear mathematical explanation (Yackel & Cobb, 1996; Bowers, Cobb, & McClain, 1999). To differentiate these norms that constitute part of the hidden curriculum of mathematics classes from general social norms, we proposed to call them sociomathematical norms.

We continued this line of research because of its potential usefulness, a key consideration for design researchers. For example, our analyses pointed to a range of understandings in action that students appear to develop as they contribute to the initial establishment and ongoing regeneration of sociomathematical norms. These include judging what counts as a different solution, a sophisticated solution, an acceptable argument, and a clear argument. In the course of conducting these analyses, we came to reconceptualize one of our primary instructional goals, that of supporting students' development of intellectual autonomy in mathematics. We had previously viewed intellectual autonomy as a characteristic of particular students but now recast it as a characteristic of a student's way of participating in the practices of a classroom community. In particular, students who are intellectually autonomous in mathematics draw on their own intellectual capabilities when making mathematical decisions and judgments as they participate in these activities. The link between the growth of intellectual autonomy and development of classrooms in which mathematics is realized as a form of inquiry is readily apparent given that in such classrooms, the teacher and students together constitute a community of validators. However, students can take over some of the teacher's traditional responsibilities only to the extent that they can judge in action both when it is appropriate to make a mathematical contribution and what counts as an acceptable mathematical contribution. These are the types of judgments that the teacher and students negotiate when establishing sociomathematical norms. Furthermore, these types of judgments and the specifically mathematical beliefs and values that they entail capture and specify much of what is implied by the notion of mathematical disposition, a major focus of reform recommendations (National Council of Teachers of Mathematics, 1991, 2000). We therefore conjectured in initiating and guiding the negotiation of sociomathematical norms, teachers are simultaneously supporting their students' development of particular mathematical dispositions.

It is important to stress that in developing the theoretical construct of sociomathematical norm, we proposed a way of making sense of what is going on in mathematics classrooms. This proposal involves an ontological innovation in that it is concerned not merely with how to look at things in general, but with particular things to look at in classrooms. As this innovation contributes more directly to analysis than to the design aspect of design research, it does not give rise unproblematically to instructional recommendations or prescriptions. Nonetheless, the gap between theory and practice is relatively narrow because the proposal involves empirically testable conjectures about the relation between social process in which students participate and the understandings that they develop while doing so. We in fact viewed the construct of sociomathematical norm as a conceptual tool that can be used to orient ongoing interpretations of classroom events. These interpretations and the suppositions and assumptions that underpin them have real consequences in that they inform (but do not determine in detail) instructional design and pedagogical decisions that are made both in action and while reflecting on prior classroom events and planning for future instructional sessions.

Intervention, Iteration, and Refinement

Sociomathematical norms provided an excellent locus for refinement. The initial analyses that we conducted when formulating the construct of sociomathematical norm were descriptive and involved the retrospective analysis of video-recordings and other data sources including copies of the students' written work. In subsequent design experiments conducted by Paul Cobb, Kay McClain, Koeno Gravemeijer, Janet Bowers, and Michelle Stephan, we monitored classroom sociomathematical norms on-line and intervened on the basis of these analyses. Our intent in doing so was to improve the quality of whole class discussions as social settings for the students' mathematical learning. For the most part, the conjectures that motivated these interventions proved to be viable. However, we made three unanticipated observations that suggest directions for further inquiry and refinement.

The first observation concerns the way in which the emergence of one sociomathematical norm (e.g., mathematical difference) can provide a basis for the emergence of subsequent sociomathematical norms (e.g., mathematical sophistication; see McClain & Cobb, 2001). In retrospect, this finding is almost banal. For example, if one solution is judged to be more sophisticated than another, that clearly implicates a judgment of difference. Valuing and being able to judge sophistication automatically carries over into judgments of difference, and likely vice versa. This finding orients us to investigate evolutionary patterns in the development of sociomathematical norms. The second observation concerns the important role that the development of notation schemes played in supporting the emergence of particular sociomathematical norms. This observation indicates the need to investigate the process by which sociomathematical norms emerge in more detail and, in particular, to attend to the material supports involved. The third observation concerns the manner in which the establishment of particular sociomathematical norms could directly influence the mathematical agenda. We found, for example, that the establishment of certain sociomathematical norms could give rise to learning opportunities for students that extended beyond those that we conjectured when intervening (McClain & Cobb, 2001). As an illustration, the distinctions between the sophistication of various types of reasoning made in particular task settings are sometimes adapted and extended spontaneously to other settings by either the teacher or the students. This observation indicates that there is much yet to be learned about the establishment and subsequent evolution of sociomathematical norms on the one hand, and the mathematical learning opportunities that arise for students on the other.

Several researchers have conducted related investigations and have identified a number of additional sociomathematical norms (e.g., Hershkowitz & Schwartz, 1999; Sfard, 2000; Simon & Blume, 1996). For example, Hershkowitz and Schwartz (1999) document the norms of what constitutes a good mathematical hypothesis and what constitutes mathematical evidence when solving open-ended instructional activities. One can imagine how a listing of sociomathematical norms might be extended almost indefinitely. In keeping with the basic tenet that theoretical constructs do useful work, the most pressing concern in our view is not to develop an exhaustive taxonomy of norms. Instead, it is to identify and investigate in greater detail those norms that appear to be particularly relevant to the agenda of supporting students' learning of significant mathematical ideas.

Synthesis and Extension

In the process of developing the notion of sociomathematical norms, we became aware of a limitation of this construct that was significant, given our concerns and interests as design researchers. Although sociomathematical norms are specific to mathematics, they are relatively broad in that they cut across mathematical domains. We therefore found it necessary to develop an additional construct, that of a *classroom mathematical practice*. This construct is concerned with the normative ways of reasoning and arguing that the teacher and students establish while discussing particular mathematical ideas. In analyzing the evolution of classroom mathematical practices in a particular classroom, we can in effect document the collective mathematical learning of the classroom community. Taken together, the three constructs of social norms, sociomathematical norms, and classroom mathematical practices constitute the core elements of the interpretive framework that we currently use to analyze the classroom microculture and thus the social situation of students' learning (Cobb, 2000). Viewed within the setting of this framework, sociomathematical norms are integral features of a proposed ontology of the mathematics classroom.

Sociomathematical norms and, more generally, the development of an interpretive framework that enables us to come to grips with the messiness and complexity of classrooms has been one of several strands of our research program. The analysis of sociomathematical norms has in fact proved relevant in several other strands that are surprisingly parallel to the case of meta-representational competence. For example, Kay McClain has focused on teachers' decision making and learning. Her analyses point to both the challenges involved for the teacher in initiating and guiding the renegotiation of sociomathematical norms, and the importance of doing so if tasks, tools and inscription are to function as means of supporting an instructional agenda that aims at significant mathematical ideas (McClain, 2002). In contrast to the relatively passive role that is frequently implied by the metaphor of the teacher as a facilitator, her analyses of sociomathematical norms contribute to the development of an empowering vision of the teacher's role in the mathematics classroom.

The construct of sociomathematical norm has also proved useful in understanding both the process by which students' interest in engaging in mathematical activity can be cultivated and the personal identities that students develop in particular classrooms. In the latter case, we draw on the constructs of social and sociomathematical norms to operationalize the idea of the normative identity as a doer of mathematics that students are, in effect, invited to develop in particular classrooms[3] (Cobb & Hodge, 2003). Although this strand of our work is still in a formative phase, our hope is that it will contribute to the goal of ensuring that *all* students have access to significant mathematical ideas.

Reflection

The case of sociomathematical norms, parallel to the meta-representational competence case, serves to illustrate that design research can serve as the context for the development of ontological innovations that enable us to specify with some

[3]Speaking of students being *invited* to develop a particular type of identity as a doer of mathematics is a sloppy way of talking in that it implies that the normative identity is established solely by the teacher, and that the students then accept or reject it. In keeping with the view that norms are joint accomplishments, it is more appropriate to say that students develop their personal classroom identities as they contribute to (or oppose) the ongoing regeneration of the normative identity as a doer of mathematics.

precision the phenomena that we seek to study. Theoretical work of this type can be demanding in that it involves "building the plane while flying it." However, the payoff is that conceptual tools are developed in the context in which they are designed to be used, that of developing analyses that can feed back to inform the ongoing instructional design and teaching effort.

The case of sociomathematical norms seems particularly analogous to that of germ theory. Just as germs are tiny, invisible, biological elements, so we can no more observe sociomathematical norms directly than we can directly perceive students' conceptual process. Instead, we have to infer the norms established in a classroom by identifying patterns and regularities in teachers' and students' classroom interactions. Just as there are standards of evidence for making inferences about the presence of particular types of germs, so much of our work has involved explicating and scrutinizing the criteria that we use when identifying the sociomathematical norms established in particular classrooms. Further, just as germ theory is an interpretive perspective that enables us to make sense of certain biological phenomena, so sociomathematical norms are constituents of an interpretive framework that enables us to make sense of certain social phenomena. In both cases, the primary motive was to understand while remaining vigilant that proposed theoretical constructs do useful work (Stokes, 1997). Together with Toulmin (1963), we contend that the quest to understand is a primary hallmark of science, with prediction and control as necessary by-products.

CONCLUSION

We have attempted to illustrate the role of a particular class of theories in design experiments—those involving ontological innovation.[4] We argued that although grand theory, orienting frameworks, frameworks for pedagogical action, and domain specific instructional theories are useful to the conduct of design studies, they do not provide an adequate theoretical basis for design research in the long term. Ontological innovation illustrates a kind of theory building that has been underplayed in design research but that is, we believe, well suited to being developed in and to contributing to design studies. Of course, we do not mean to imply that all theories involving ontological innovation are "good" or useful, nor that useful theories that do not involve ontological innovation are impossible.

[4]As a somewhat ironic historical observation, we note that Collins' well-known 1992 paper on design experiments seemed to present "determining all the variables in a design experiment" as a first, and possibly not-that-difficult task. He said that, "Because [educational innovations] have been so varied, they should have uncovered most of the variables needed for a design theory." (p. 20) Whether they count as variables or not, we believe that ontological innovations obviously are critical to consider, difficult to define, and not nearly all already discovered.

In characterizing ontological innovations, it is more feasible to focus on their function in cycles of design and analysis than to try and identify inherent properties apart from the contexts in which they are developed, refined, and used. With regard to design, ontological innovations can delineate new competencies that become a focus on instruction (e.g., meta-representational competence) or can provide new goals for instructional environments (e.g., a productive ecology of sociomathematical norms). In addition, their formulation can lead to the reformulation of previously adopted instructional goals (e.g., intellectual autonomy as a characteristic of a student's way of participating in the practices of a classroom community rather than as a characteristic of a student per se). Furthermore, an ontological innovation can orient the designer to scrutinize previously implicit aspects of a design. For example, students' use of inscriptions is a primary means of support in almost all designs. However, different designs support students' developing competence in inventing, evaluating, and refining a variety of representational forms more or less well. The theoretical category of meta-representational competence orients the designer to examine this aspect of a proposed design. Similarly, all designs involve assumptions about the nature of classroom interactions and discourse. The theoretical category of sociomathematical norms orients the designer to both explicate these assumptions, thereby further specifying the design, and to examine the types of mathematical dispositions that the design appears to support. Finally, an ontological innovation can inform the formulation of new benchmarks for assessing either the consequences of instruction (e.g., meta-representational competence) or the classroom learning environment (e.g., sociomathematical norms).

With regard to the analysis aspect of the design research cycle, ontological innovations provide new lenses for making sense of what is happening in the complex, more-or-less real world instructional setting in which a design study is conducted. An ontological innovation is of value to the extent that it enables us to see and account for patterns and regularities that can inform pedagogical and design decisions. In this regard, we note that interpretations of classroom events reflect suppositions and assumptions about learning, teaching, and mathematics as well as a range of issues typically subsumed under the heading of affective factors such as students' dispositions towards and interest in engaging in mathematical activity. These suppositions and assumptions are, in part, embodied in ontological assumptions, and thus ontology is reflected in the interpretations we make of classroom events. Although ontological innovations may not determine particular pedagogical and design decisions in any detail, they can profoundly shape the things we see, and thus the data on which those decisions are based.

Previously, and in this article generally, we emphasized the contributions ontological innovation can make to design experiments. But we also believe design experiments are excellent contexts in which to do the work of ontological innovation. In particular, design experiments

1. Promote *grounding* of theoretical constructs in real-world experiences.
2. Foster the development of *useful* constructs.
3. Provide multiple exposures to empirical test that aid in the difficult and extended work of *refinement*.
4. By the same token, help develop constructs that are *robust* in their application across variations in context.

In presenting the two case studies, we illustrated that the complexity of instructional settings often drives the formulation of new theoretical entities. As the development of the categories of meta-representational competence and of sociomathematical norms were both occasioned by surprising observations, we should clarify that this is a common but by no means necessary story. Work on an ontological innovation might, for example, be initiated by the anticipation of an interpretive challenge in an upcoming design experiment. For example, can we operationalize, know how to see, "a student's concept" in classroom discussions? In such a case, preparation for the experiment and for the development of an innovation might include exploring potentially relevant theoretical constructs in other disciplines. One might think that one could just import those other-disciplinary constructs, but in our experience that is rare. Instead, they serve as resources and scholarly contextualization for our own work of innovation.

Initial invention or discovery in a narrow sense is a small part of ontological innovation. Both case studies emphasize new theoretical entities as *loci for refinement* of many sorts. Typically, our conceptualization of a new theoretical category needs to be revised and refined as it is used and proves to be unduly simplistic, overly restrictive, or inadequate. There are almost always contextual differences that are instructionally important (e.g., between motion representations and representations of spatially distributed data). Sometimes, new constructs wind up doing part of, but not all of what we had initially hoped (e.g., complementing sociomathematical norms with the idea of mathematical practices). The need for further study is almost a foregone conclusion precisely because the posited existence of new theoretical entities is consequential. While retrospective studies may advance our understanding, perhaps even by beginning an ontological innovation, many further cycles of conjecture and implementation are typically needed both to specify the nature of new entities, and also to develop a rich understanding of details that affect instruction differently in different contexts. It is in the course of this process that the innovation is operationalized and standards of evidence are clarified so that other researchers can both use the category with some fidelity and can monitor analyses in which it has been employed.

In addition to refinement and operationalization, we briefly illustrated how collateral studies can be affected by ontological innovation. Study of teacher strategies and student engagement were stimulated and contextualized by the relevant ontological innovations in both our case studies. These developments serve to further emphasize the programmatic nature of design research that involves the formulation of ontological innovations.

Making ontological innovations is plainly difficult. We have refrained from narrow prescriptions about how to develop or validate them in part because that encroaches on some of the deepest issues of "how to do (theoretical) science well." Even if we had unlimited expository space, that is a task for a larger community.

We have taken pains to acknowledge that an ontological innovation stands at some distance from many elements of instructional design. "Managing the gap" between theory and what we must do to teach is an essential part of the task. It does not advance either science or instruction to pretend that theories are encompassing and highly prescriptive of instruction. As we have argued, theory development in general, and ontological innovations in particular, are nonetheless critical for the long-term scientific health and practical power of design-related educational research.

ACKNOWLEDGMENT

We are grateful to Sasha Barab for his constructive critique of a previous draft of this article.

REFERENCES

Anderson, J. R. (1990). *The adaptive character of thought.* Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Azevedo, F. (2000). Designing representations of terrains. *Journal of Mathematical Behavior, 19*(4), 423–480.

Azevedo, F. (submitted). *Personal excursions: Investigating the dynamics of student engagement.*

Bowers, J. S., Cobb, P., & McClain, K. (1999). The evolution of mathematical practices: A case study. *Cognition and Instruction, 17*, 25–64.

Brown, A. L. (1992). Design experiments: Theoretical and methodological challenges in creating complex interventions in classroom settings. *Journal of the Learning Sciences, 2,* 141–178.

Brown, A. L., & Campione, J. C. (1996). Psychological theory and the design of innovative learning environments: On procedures, principles, and systems. In R. Glaser (Ed.), *Innovations in learning: New environments for education* (pp. 289–325). Mahwah, NJ: Lawrence Erlbaum Associates.

Cadwell, L. B. (2003). *Bringing learning to life: The Reggio approach to early childhood education.* New York: Teachers College Press.

Cobb, P. (2000). The importance of a situated view of learning to the design of research and instruction. In J. Boaler (Ed.), *Multiple perspectives on mathematics teaching and learning* (pp. 45–82). Stamford, CT: Ablex.

Cobb, P., Confrey, J., diSessa, A., Lehrer, R., & Schauble, L. (2003). Design experiments in education research. *Educational Researcher, 32*(1), 9–13.

Cobb, P., & Hodge, L. L. (2003, April). An interpretive scheme for analyzing the identities that students develop in mathematics classrooms. Paper presented at the annual meeting of the American Educational Research Association, Chicago.

Collins, A. (1992). Toward a design science of education. In E. Scanlon & T. O'Shea (Eds.), *New directions in educational technology* (pp. 15–22). New York: Springer-Verlag.

Confrey, J., & Lachance, A. (2000). Transformative teaching experiments through conjecture-driven research design. In E. Kelly & R. Lesh (Eds.), *Handbook of research design in mathematics and science education* (pp. 231–266). Mahwah, NJ: Lawrence Erlbaum Associates.

Design-Based Research Collaborative (2003). Design-based research: An emerging paradigm for educational inquiry, *Educational Researcher, 32*(1), 5–8.

diSessa, A. A. (1992). Images of learning. In E. De Corte, M. C. Linn, H. Mandl, and L. Verschaffel (Eds.), *Computer-based learning environments and problem solving* (pp. 19–40). Berlin: Springer.

diSessa, A. A., (1995). The many faces of a computational medium. In A. diSessa, C. Hoyles, R. Noss, with L. Edwards (Eds.), *Computers and exploratory learning* (pp. 337–359). Berlin: Springer Verlag.

diSessa, A. A. (2002). Students' criteria for representational adequacy. In K. Gravemeijer, R. Lehrer, B. van Oers, & L. Verschaffel (Eds.), *Symbolizing, modeling and tool use in mathematics education* (pp. 105–129). Dortrecht: Kluwer.

diSessa, A. A. (in press). Meta-representation: Native competence and targets for instruction. *Cognition and Instruction.*

diSessa, A. A., Hammer, D., Sherin, B., & Kolpakowski, T. (1991). Inventing graphing: Meta-representational expertise in children. *Journal of Mathematical Behavior, 10*(2), 117–160.

diSessa, A. A., & Minstrell, J. (1998). Cultivating conceptual change with benchmark lessons. In J. G. Greeno & S. V. Goldman (Eds.), *Thinking practices in mathematics & science learning* (pp. 155–187). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Edelson, D. C. (2002). Design research: What we learn when we engage in design. *Journal of the Learning Sciences, 11*, 105–121.

Eisenhart, M. (1991). Conceptual frameworks for research. In R. G. Underhill (Ed.), *Proceedings of the thirteenth annual meeting of the North American Chapter of the Psychology of Mathematics Education, V. 1* (pp. 202–219). Blackburg, VA: Virginia Tech.

Elby, A. (2000). What students' learning about representations tells us about constructivism. *Journal of Mathematical Behavior, 19*(4), 481–502.

Engestrom, Y. (1998). Reorganizing the motivational sphere of classroom culture: An activity - theoretical analysis of planning in a teacher team. In F. Seeger, J. Voigt & U. Waschescio (Eds.), *The culture of the mathematics classroom*. New York: Cambridge University Press.

Engestrom, Y. (1999). Expansive visibilization of work: An activity-theoretical perspective. *Computer Supported Cooperative Work, 8*, 63–93.

Friedman, J. S. (1999). *Image processing in a science classroom: The role of students' understanding of the technology*. Unpublished doctoral dissertation. Berkeley, CA: Graduate School of Education University of California.

Friedman, J. S., & diSessa, A. A. (1999). What should students know about technology? The case of scientific visualization. *International Journal of Technology and Science Education, 9*(3), 175–196.

Granados, R. (2000). Constructing intersubjectivity in representational design. *Journal of Mathematical Behavior, 19*(4), 503–530.

Gravemeijer, K. (1994a). Educational development and developmental research. *Journal for Research in Mathematics Education, 25*, 443–471.

Gravemeijer, K. (1994b). *Developing realistic mathematics education*. Utrecht, The Netherlands: CD-β Press.

Gruber, H., & Vonèche, J. J. (1977). *The essential Piaget*. New York: Basic Books.

Hershkowitz, R., & Schwarz, B. (1999). The emergent perspective in rich learning environments: Some roles of tools and activities in the construction of sociomathematical norms. *Educational Studies in Mathematics, 39*, 149–166.

Labato, J. (2003). How design experiments can inform a rethinking of transfer and vice versa. *Educational Researcher, 32*(1), 17–20.

Lave, J. (1988). *Cognition in practice: Mind, mathematics and culture in everyday life*. New York: Cambridge University Press.

Linn, M. C., & Hsi, S. (2000). *Computers, teachers, peers: Science learning partners*. Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Madanes, R. (1997). *Teaching through discussion: Using critical moves and support moves*. Paper presented at the Annual Meeting of the American Educational Research Association, Chicago.

McClain, K. (2002). Teacher's and students' understanding: The role of tools and inscriptions in supporting effective communication. *Journal of the Learning Sciences, 11*, 217–249.

McClain, K., & Cobb, P. (2001). The development of sociomathematical norms in one first-grade classroom. *Journal for Research in Mathematics Education, 32*, 234–266.

Metz, K. (1995). Reassessment of developmental constraints on children's science instruction. *Review of Educational Research, 65*(2), 93–127.

National Council of Teachers of Mathematics (1991). *Professional standards for teaching mathematics*. Reston, VA: National Council of Teachers of Mathematics.

National Council of Teachers of Mathematics (2000). *Principles and standards for school mathematics*. National Council of Teachers of Mathematics.

Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.

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

Papert, S, & Harel, I. (1991). *Constructionism: Research reports and essays, 1985–1990*. Norwood, NJ: Ablex.

Sfard, A. (2000). On the reform movement and the limits of mathematical discourse. *Mathematical thinking and learning, 2*, 157–189.

Sherin, B. (2000). How students invent representations of motion. *Journal of Mathematical Behavior, 19*(4), 399–441.

Sherin, B., Azevedo, F., & diSessa, A. (in press). Exploration zones: A framework for describing the emergent structure of learning activities. In A. Rosebery, B. Warren, R. Nemirovsky, & J. Solomon (Eds.), *Everyday matters in science and mathematics*. Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Simon, M.A. (2000). Research on the development of mathematics teachers: The teacher development experiment. In A.E. Kelly & R.A. Lesh (Eds.), *Handbook of research design in mathematics and science education* (pp. 335–359). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Simon, M. A., & Blume, G. W. (1996). Justification in the mathematics classroom: A study of prospective elementary teachers. *Journal of Mathematical Behavior, 15*, 3–31.

Skinner, B. F. (1974). *About behaviorism*. New York: Knopf.

Stokes, D. (1997). *Pasteur's quadrant: Basic science and technological innovation*. Washington, DC: Brookings Institute.

Streefland, L. (1991). *Fractions in realistic mathematics education. A paradigm of developmental research*. Dordrecht, Netherlands: Kluwer.

Thompson, P. W., & Thompson, A. G. (1994). Talking about rates conceptually, part I: A teacher's struggle. *Journal for Research in Mathematics Education, 25*, 279–303.

Thompson, A.G., & Thompson, P. W. (1996). Talking about rates conceptually, part II: Mathematical knowledge for teaching. *Journal for Research in Mathematics Education, 27*, 2–24.

Toulmin, S. (1963). *Foresight and understanding*. New York: Harper Torchbook.

Treffers, A. (1987). *Three dimensions: A model of goal and theory description in mathematics instruction - The Wiskobas Project*. Dordrecht, Netherlands: Reidel.

Walkerdine, V. (1988). *The mastery of reason*. London: Routledge.

Wenger, E. (1998). *Communities of practice*. New York: Cambridge University Press.

Yackel, E., & Cobb, P. (1996). Sociomathematical norms, argumentation, and autonomy in mathematics. *Journal for Research in Mathematics Education, 27*, 458–477.