

## The Instructional Core

*The Kendall School, a K–8 school in a suburban neighborhood, is struggling with its mathematics curriculum. The district has adopted a new math curriculum, keyed to the state content and performance standards, and the Kendall School is into its second full year of implementation. It is clear not only that the curriculum is not improving math performance—scores have actually declined in a few classrooms—but also that teachers are struggling with the new content. Parents are beginning to express frustration because they don’t understand the homework that students are bringing home and the math problems in the homework don’t look anything like the math that they studied when they were in school. Pat Granger, the principal of Kendall, is starting to field complaints from both teachers and parents about the state of math instruction.*


*The district math coordinator and the professional development consultant whom the district has hired have a different view. “The problem, as far as I’m concerned,” the math coordinator says, “is that this is not just a change in content—it is a change in the whole way we think about math in the classroom. It requires teachers to put much more control in the hands of students; it requires students to think about the mathematical ideas, not just the procedures; and it requires a significantly higher level of mathematical knowledge than teachers have been expected to have.”*

*Granger says, “We thought we were adopting a curriculum, and it turns out, we’ve adopted a monster. This thing has set off a whole range of issues we’re not very well prepared to deal with.”*

What's happening at the Kendall School is a version of what happens whenever a school or system undertakes an instructional change that is well outside the existing knowledge and skill set of students, teachers, and administrators. Students are unfamiliar with the new classroom demands and, for the most part, don't understand why content that was familiar to them in the past now seems strange and unfamiliar. Teachers find their established ways of teaching disrupted and are confronted with classroom management and pedagogy issues for which they may have only minimal preparation. Administrators are faced not just with responding to teachers' expressions of uncertainty, but also the reverberations of this uncertainty through students and their parents. What's going on here is that the system has adopted a *disruptive technology*—a curriculum and a set of pedagogical practices that require people to think and act differently than they have in the past.<sup>1</sup> The sources of uncertainty in this disruptive technology are rooted in the relationship between the teacher and the student in the presence of content—the instructional core (figure 1.1). Making meaningful and productive changes in instructional practice requires us to confront how they upset and, in some sense, reprogram our past ways of doing things. The success of the new math curriculum at Kendall depends on understanding exactly what has to change in the instructional core for the new curriculum to deliver on its promise of higher learning of mathematics. The instructional core anchors the practice of rounds and any other school- or district-level instructional improvement process.

In its simplest terms, the instructional core is composed of the teacher and the student in the presence of content. In the work of the philosopher of education,

**FIGURE 1.1 THE INSTRUCTIONAL CORE**



**FIGURE 1.2 SEVEN PRINCIPLES OF THE INSTRUCTIONAL CORE**

1. Increases in student learning occur only as a consequence of improvements in the level of content, teachers' knowledge and skill, and student engagement.
2. If you change any single element of the instructional core, you have to change the other two.
3. If you can't see it in the core, it's not there.
4. Task predicts performance.
5. The real accountability system is in the tasks that students are asked to do.
6. We learn to do the work by doing the work, *not* by telling other people to do the work, *not* by having done the work at some time in the past, and *not* by hiring experts who can act as proxies for our knowledge about how to do the work.
7. Description before analysis, analysis before prediction, prediction before evaluation.

David Hawkins, this is the “I” (the teacher), the “thou” (the student), and the “it” (the content).<sup>2</sup> As the Hawkins framework was developed and elaborated on by David Cohen and Deborah Ball, it is the *relationship* between the teacher, the student, and the content—*not* the qualities of any one of them by themselves—that determines the nature of instructional practice, and each corner of the instructional core has its own particular role and resources to bring to the instructional process.<sup>3</sup> In his seminal work on instructional practice, Walter Doyle locates the *instructional task* at the center of the instructional core.<sup>4</sup> Simply stated, the instructional task is the actual work that students are asked to do in the process of instruction—*not* what teachers *think* they are asking students to do, or what the official curriculum *says* that students are asked to do, but what they are *actually* asked to do. So, for example, in an “advanced” science class, if students are asked to memorize elements and their atomic structures, the actual task that the students are being asked to do is a memorization task, even though the teacher might think that because the material is difficult and the work is beyond what students in the “regular” science class are being asked to do, it is a higher-level task.

The model of the instructional core provides the basic framework for how to intervene in the instructional process so as to improve the quality and level of student learning. Seven principles guide our work with the instructional core (figure 1.2).

**FIRST PRINCIPLE: *Increases in student learning occur only as a consequence of improvements in the level of content, teachers' knowledge and skill, and student engagement.***

There are only three ways to improve student learning at scale. The first is to increase the level of knowledge and skill that the teacher brings to the instructional process. The second is to increase the level and complexity of the content that students are asked to learn. And the third is to change the role of the student in the instructional process. That's it. If you are not doing one of these three things, you are not improving instruction and learning. Everything else is instrumental. That is, everything that's *not* in the instructional core can only affect student learning and performance by somehow influencing what goes on *inside* the core.

When educators think about "changing" instruction, they typically focus not on the instructional core, but on the various structures and processes that surround the core. They might choose, for example, to group students in a particular way because of a theory about how grouping will affect the relationship of the student and the teacher in the presence of content. But it is not the grouping practice that produces student learning. Rather, it is the change in the knowledge and skill that teachers bring to the practice, the type of content to which students gain access, and the role that students play in their own learning that determine what students will know and be able to do. If changes in grouping practices don't alter the core, then the likelihood they will affect student learning is remote.

What about content and performance standards? Standards only operate by influencing the level of the content that's actually being taught. Their effect in actual classrooms depends on whether there are materials that reflect the standards, whether teachers know how to teach what the materials and standards require, and whether students find the work that they are being asked to do worthwhile and engaging.

What about professional development? Professional development works, if it works at all, by influencing what teachers *do*, not by influencing what they think they ought to do or what the professional developers think teachers ought to do. The quality and impact of professional development depends on what teachers are being asked to learn, how they are learning it, and whether they can make the practices they are being asked to try work in their classrooms.

What about supervision, evaluation, and strong instructional leadership? Administrators' influence on the quality and effectiveness of classroom instruction

is determined *not* by the leadership practices they manifest, but by the way those practices influence the knowledge and skill of teachers, the level of work in classrooms, and level of active learning by students. Much of what well-intentioned policy makers and administrators do in the name of school improvement never actually reaches the instructional core. Much of it doesn't even reach the classroom, much less inside the classroom. Our best ideas about policy and management don't *cause* student learning to increase. At the very best, when they are working well, they *create conditions* that influence what goes on inside the instructional core. The primary work of schooling occurs inside classrooms, *not* in the organizations and institutions that surround the classroom. Schools don't improve through political and managerial incantation; they improve through the complex and demanding work of teaching and learning.

**SECOND PRINCIPLE: *If you change any single element of the instructional core, you have to change the other two to affect student learning.***

The second principle follows from the first. So, for example, if your improvement strategy begins with a curriculum solution—say, as at the Kendall School, the adoption of a new math curriculum—then you have to invest in the new knowledge and skill required of teachers to teach that curriculum if you expect it to contribute to new student learning. A failure to address teachers' knowledge and skill as part of a curriculum-based improvement strategy typically produces low-level teaching of high-level content, a situation we see with considerable frequency in American classrooms. We call this “knocking the corners off the grand piano to get it through the door.” Teachers assign high-level text or complex problems and then structure student learning around familiar fill-in-the-blank worksheets. Or teachers walk students through a straight procedural explanation of how to find the answer, leaving the students in the role of recording what the teacher says, rather than actively thinking through the problems for themselves. If you invest in teacher professional development without a clear understanding of where you expect it to lead in terms of the actual content that students are expected to master, then you get random innovation across classrooms and the innovation has no systemwide or schoolwide impact on student learning.

If you raise the level of content *and* the knowledge and skill of teachers without changing the role of the student in the instructional process, you get another common situation in American classrooms: Teachers are doing all, or most, of the work,

exercising considerable flair and control in the classroom, and students are sitting passively, watching the teacher perform. A common student question in these classrooms is, “Teacher, should I write this down?” If you raise the level of teachers’ knowledge and skill in general pedagogy without anchoring it in content, you get high-level practice disconnected from a clear understanding of what students are actually learning, and from the specific issues that students have with specific cognitive tasks. This is what David Hawkins means when he says, “Without an it, there is no content for the context, no figure, no heat, but only an affair of mirrors confronting each other.”<sup>5</sup>

We frequently hear educators talk about how well the lesson went, without reference to what students were actually doing and the visible evidence of what students actually knew as a consequence of the teaching. Mostly, the lesson has “gone well” when it has gone according to plan, without any specific reference to what students do or don’t know as a consequence of the teaching. Intervening on any single axis of the instructional core means that you have to intervene on the other two to have a predictable effect on student learning.

If you invest in higher-level content *and* teacher knowledge and skill, but you neglect the role of the student in the instructional process, you get students (and parents) who, as at the Kendall School, don’t understand the new roles and demands that they are expected to meet. Americans are much more comfortable talking about changing content and teaching than they are about changing the role of the student in instruction. We focus much more attention on textbook adoptions and curriculum alignment, for example, than we do on analyzing students’ actual responses to the content, what motivates them to high levels of engagement with the content, and their actual role in the instructional process. In the more advanced strategies of improvement, we focus attention on helping teachers get familiar with new content and pedagogy, but we focus relatively little attention on what students are doing when they are actively engaged in learning what we think they should learn.

This is one big difference between American schools and schools in other countries. Here we spend a great deal of time worrying about *what* we’re teaching and *how* it is being taught. In other places, people also spend a great deal of time worrying about *whether* students are actually interested in, actively engaged in, and able to explain how they the students think about what adults are trying to teach them. There are differences between elementary schools and secondary schools in the United States on this score. It is much more common, although still not the dominant practice, in U.S. elementary schools for teachers to pay attention

to whether students are actually interested and engaged in learning. Most of the instruction we observe in secondary schools is about “delivering” the content and, most importantly, about deciding which students are smart and which are “deserving” of further attainment. The culture of American schools, in its deep structure, is very teacher-centric. You only see the magnitude of this when you step outside the culture. We tend to focus more on what the teacher is doing in front of the classroom than we do on the work that is actually on top of the student’s desk. More about this later.

The instructional core provides a heuristic for assessing the likelihood that any systemic improvement strategy, or any particular change in policy or practice, will result in any real improvement in student learning:

- How will this affect teachers’ knowledge and skills?
- How will this affect the level of content in classrooms?
- How will this affect the role of the student in the instructional process?
- How will this affect the relationship between the teacher, the student, and content?

In a more specific context, the questions might sound something like “We’re doing formative assessment”—yes, but how will your investment in the technology of assessment influence teachers’ knowledge and skill, the level of content you expect to see in the classroom, and the role of the student in the instructional process? “We’re focusing on developing strong instructional leaders”—yes, but what actual practice that will lead to improvements in content, knowledge and skill, and student engagement are you asking leaders to engage in? “We’re adopting a new, more challenging math curriculum”—yes, but how would you know whether the instructional practice on which the curriculum is predicated is actually occurring in classrooms, and with what level of depth and consistency?

**THIRD PRINCIPLE: *If you can’t see it in the core, it’s not there.***

The third principle is, in general, a good rule for the design of large-scale improvement strategies. It doesn’t matter how much money you’ve spent. Nor does it even really matter whether everyone thinks it’s a fantastic idea (since many people like best the changes that are the least disruptive). And, above all, it doesn’t matter whether everyone else is doing it. What matters is whether you can see it in the core. If you can’t, it’s not there.

The instructional core also helps us predict what we would expect to see happening to student learning over time. Here the central idea is *the academic task*. As an example, one of our superintendents' networks was visiting a school in a network member's district. This particularly thoughtful and active host superintendent had managed to make quite a lot happen instructionally in his district in a relatively short period. In our visit, we broke into groups of three or four and did a series of rotations through classrooms, with two groups seeing each of four classrooms at a given grade level for a period of time. We then observed the team meeting of the teachers in the grade level whose classrooms we had observed. So, essentially, we saw the instruction in each classroom and then we saw the teachers talking about the instruction in their team meeting. Because the district and the school had worked hard on curriculum alignment, the teachers were able to talk about a common lesson sequence they were teaching and about the work that students were producing in that sequence, according to a common assessment that all the teachers were using. This is a rather sophisticated system.

In the team meeting, a problem emerged. The student work was obviously quite variable from classroom to classroom. In one classroom in particular, there were a number of students whose assessment results suggested that they apparently did not understand the content. The team leader asked the teachers what they thought explained the differences among classrooms. Each teacher offered an explanation. The explanations had mainly to do with the teachers' interpretations of the students' skill levels at the beginning of the unit. That is, the teachers felt that students who were struggling with the content had weak prior learning. So the discussion quickly shifted to what kind of remedial strategies one might use to bring those students up to the desired level.

What the teachers didn't know—because they had never observed each other teaching—was that the *actual work* that we observed students doing, within a nominally common curriculum framework, was quite different in each of the four classrooms. And the level of the student work that was presented at the grade-level meeting was quite close to the actual work that students were being asked to do in each classroom. In other words, the variability in student performance was a result of the teaching that was going on and the actual tasks that students were asked to do, *not*, as the teachers hypothesized, a result of the students' prior knowledge. This was yet more evidence for a simple, but powerful lesson—hold on to your hats—*teaching causes learning*. In the absence of direct evidence on what her colleagues were doing, the team leader, whose students produced the most consistently high-



level work, was projecting her own practice onto the practice of the other teachers on the team. This led her to suggest that the variability couldn't be the result of differences in teaching, since "we're all teaching the same thing." In fact, they weren't.

What was different in the four classrooms was what exactly the students were being asked to do and the degree to which the teacher engaged students in the work by scaffolding their learning up to the complexity of the task. The curriculum was the same; the tasks were different.

In one classroom, the teacher took twenty minutes of the fifty-five-minute period explaining the task and directing students through a detailed procedural drill on what to do. The instructions were so complex that most students (and observers) couldn't repeat them when the children were released to work on their own. In another classroom, the teacher focused very little time on setting up the task, passed out the materials, and asked students to work individually on the task and to consult other students in their group if they got stuck. In yet another classroom, the teacher passed out the task, assigned roles to students at tables, and then circulated through the room answering individual students' questions.

In the team leader's classroom, the teacher spent less than five minutes reminding the students of how the task they were about to do was connected to the previous day's work, asked students what they had learned from that work, and then spent about five minutes walking students through a discussion of a model task that was similar to the one they were being asked to do. She then put students in groups, assigned roles, and circulated through the room. When we asked students in the first three classrooms what they were working on, none of them could reliably describe the task. When we asked students in the fourth classroom, they could reliably tell us what they were expected to do and tell us how it was connected to what they had done earlier.

It is important to add here that students in all four classrooms were "engaged," by conventional definitions—that is, they were attentive, nondisruptive, and compliant. If you were doing a windshield survey of classroom climate in this school, with the typical supervisory checklist, you would see, without exception, classrooms that were quiet and orderly and in which the teachers had done everything that the external environment expected of them. The "Do Now" was in the upper left-hand corner of the whiteboard; the specific objective of the day was prominently displayed, referenced to the appropriate state standard; the "Students Will Be Able to Do" was adjacent to the standard. If you stayed at the surface-level character-

istics of the classroom, you would predict that students were all getting access to the same work.

But in reality, students were engaged in very different levels of work in different classrooms around a common curriculum unit. In the classroom where students were explicitly drawing on prior knowledge about how to address the task and where they had experience working individually and in groups, not surprisingly, they were relatively competent at doing what the teacher expected them to do, and they did it at a relatively high level. The teacher was free to work with individual students who were struggling with the task. Not surprisingly, things were different when the teacher was the main source of information on the task and the teacher's practice at setting up the task was disconnected from the students' understanding of it. Students were confused about the task and variable in their engagement with it. In our experience, the latter situation is much more common than the former in American schools. One of our favorite questions to ask students during an observation is "What's going on here?" The most frequent response is, "I don't know," or "Ask the teacher—she knows."

#### **FOURTH PRINCIPLE: *The task predicts performance.***

What determines what students know and are able to do is not what the curriculum says they are supposed to do, or even what the teacher thinks he or she is asking students to do. What predicts performance is *what students are actually doing*. Memorization tasks produce fluency in memorization and recall, not necessarily understanding. Memorizing the elements of the periodic table is not the same as understanding the properties of the elements. The single biggest observational discipline we have to teach people in our networks is to look on top of the students' desks rather than at the teacher in front of the room. The only way to find out what students are actually doing is to observe what they are doing—not, unfortunately, to ask teachers what students have done after the fact or to look at the results of student work after they have engaged in the task. What was interesting about our observation with the superintendents' network was that for a brief moment, for this particular task, we, the observers, actually knew more about what was going on in these classrooms than the teachers did. This is an unsettling commentary on the instructional culture of American schooling.

Walter Doyle, from whom we have drawn most of our understanding of the nature of academic work, makes an interesting point about accountability:

Accountability drives the task system in the classroom. As a result, students are especially sensitive to cues that signal accountability or define how tasks are to be accomplished. In addition, students tend to take seriously only that work for which they are held accountable.<sup>6</sup>

The accountability problem in the classroom is a microcosm of the accountability problem in the broader system. Other things being equal, people tend to want to do what they are expected to do in complex social systems with interlocking expectations. But to do what they are expected to do, they must know not only *what* they are expected to do but also *how* they are expected to do it, and what *knowledge and skill* they need to learn how. This is the distinction that Nobel economist Thomas Schelling makes between doing the right thing and knowing the right thing to do. When we put teachers and students in situations where the task is vague and unspecified, but the expectations for performance are specific and high, we are expecting them to do the right thing without knowing the right thing to do. Students in three of the classrooms we observed that day were dutifully doing what they thought the teacher expected them to do, without knowing either what they were actually supposed to do or, more importantly *why* they should *want* to do it. We frequently see the same pattern with teachers—they are doing their best to do precisely what they think is expected of them without the *what* or the *why* being clear. This is a failure of the system, not of the teachers. Students in the fourth classroom had discussed how the task was related to the previous day's work and what they had learned from that work; they had seen and discussed a version of the task with the teacher *before* they were asked to work independently and in groups on the task. It was also clear from the way they worked that they were familiar with this routine. Notice also that the practice of the team leader did not trickle into the classrooms of the other teachers at her grade level—the culture of autonomous practice guaranteed that.

**FIFTH PRINCIPLE: *The real accountability system is in the tasks that students are asked to do.***

This connection between doing the right thing and knowing the right thing to do leads to the fifth principle. From a policy and managerial perspective, we tend to think of accountability as a systemic issue. Accountability, in this view, is the way we steer the system toward a good collective result, using performance measures, standards, rewards, and sanctions. From this perspective, we tend to think that

if we just get the incentives and structures right, good things will follow. In fact, this view of accountability rests on a heroic, largely unfounded assumption that students and teachers actually know what to do, that they know how to do it, and, most importantly, that they are able to derive some personal meaning and satisfaction from having done it. If you can't solve this problem of accountability at the classroom level, then the system-level work on accountability is mostly about the manipulation of political and managerial symbols, not about the improvement of learning.

In our experience working with teachers, principals, and system-level administrators around problems of large-scale improvement, people tend to be much more specific about what they expect by way of student performance than they are about what in classrooms would lead to the performance they desire. American schools have traditionally had an extraordinarily weak instructional culture, which has led, in turn, to extremely high variability in student performance among classrooms within schools, and to an extremely low capacity to affect instructional practice and student learning at scale. Trying to move performance in a system with a weak instructional culture is like pushing on a string. It doesn't do any good to know that there *is* an instructional core and that the tasks that students are asked to do within that core actually drive student learning, if the core itself differs from one classroom to another and if people aren't used to thinking about instruction as a collective practice.

This is why we have invested a good deal of our professional energy in building the competence of leaders in schools to observe, analyze, and affect instructional practice. We have deliberately drawn on the medical model in this work, not because we think educators ought to act more like physicians, but because medicine has, in our view, the most powerful social practice for analyzing and understanding its own work—the medical rounds model. In most instances, principals, lead teachers, and system-level administrators are trying to improve the performance of their schools without knowing what the actual practice would have to look like to get the results they want at the classroom and school level. We work with educators on the observation and analysis of teaching practice not because we think it's good for their souls (although it may be), but because we think you cannot change learning and performance at scale without creating a strong, visible, transparent common culture of instructional practice. And you can't create a common culture of practice without actually *engaging* in the practice yourself. We know this is heresy, since

most administrators and support staff in schools choose to do what they are doing precisely because they see work in classrooms as too limiting. But this heresy leads to the sixth principle.

**SIXTH PRINCIPLE:** *We learn to do the work by doing the work, not by telling other people to do the work, not by having done the work at some time in the past, and not by hiring experts who can act as proxies for our knowledge about how to do the work.*

The genius of the medical rounds model is that the profession reproduces its practice and the surrounding culture through direct, face-to-face interactions around the work. To be sure, there is an ample supply of knowledge that comes from outside sources into the practice of physicians. And certainly, there are strong external controls and incentives that drive the practice in a given direction. But nested within these external structures and incentives is a social process for inducting people into the practice, for sustaining and developing norms of practice, and for making face-to-face evaluations of practice. The education sector, which is no less knowledge-intensive than medicine at its core, has no such culture-building practice. It should not surprise us, then, that the enterprise is atomized at its core. Education is essentially an occupation trying to be a profession without a professional practice.

When we work with people to develop their knowledge and practice around the instructional core, they typically ask two questions in the earliest stages of the work: “Can you tell us what high-level instruction looks like?” And, “How do I get people to do it?” People want an immediate framework for judging whether teachers are “doing it,” and they want us to tell them how to get people who are not currently “doing it” to “do it.” Our role is to disappoint people. We are fairly adamant in resisting answering these questions, to the point that it has become a standing joke in our practice. Why? Because we think people have to engage in sustained description and analysis of instructional practice before they can acquire either the expertise or the authority to judge it, much less to evaluate other people doing it. Most of the educators we work with—understandably, given the pressure they are under—want an immediate shortcut to the answer. You don’t build a culture by taking shortcuts. It took over a hundred years to build the current dysfunctional instructional culture of American schools; it won’t be transformed by taking a three-day or six-week course in supervision and evaluation.

**SEVENTH PRINCIPLE: *Description before analysis, analysis before prediction, prediction before evaluation.***

You build a common culture of instruction by focusing on the language that people use to describe what they see and by essentially forcing people to develop a common language over time. Language *is* culture, and vice versa. When you jump straight from observation to evaluation, you short-circuit the difficult process of developing a common language to use in describing what you see going on in classrooms. In the absence of such a language, what you mean by some key term—*student engagement*, for example—might be completely different from what your colleague means by it, and you end up agreeing to disagree because it's too hard to figure out how to negotiate your differences. In our work, we insist that people develop a strong descriptive language and that they go through several iterations of a process for developing a common language *before* we move on to the tasks of analyzing, predicting, and evaluating (see chapter 4 for a deeper discussion of this process).

*Analysis* is getting people to work at grouping what they see into mutually agreed-upon categories and to start to make some judgments about how the categories are related to each other. *Prediction* is learning to use the evidence of observation and the analysis to make causal arguments about what kind of student learning we would expect to see as a consequence of the instruction we have observed. Typically, we ask people, “If you were a student in this classroom and you did exactly what the teacher expected you to do, what would you know how to do?” This question stems directly from the fourth principle—task predicts performance.

Only after people have developed the disciplines of description, analysis, and prediction do we raise the issue of evaluation, and then, we don't raise it in the typical form of “was this good teaching or not?” We ask people to address the question “What is the next level of work in this classroom, school, or system?” We pose the evaluative question in this way specifically to avoid the superficial classification of practice into “good,” “mediocre,” and “bad” because we want practitioners to think about the process of improvement as a *clinical practice*. That is, our job is to make the practice better over time, *not* to mete out rewards and punishments.

There is also an issue of humility involved here. Most of the people who, by virtue of their positional authority, are evaluating teachers could not themselves do what they are asking teachers to do. Teachers know this. The escalating demands of teaching practice are such that the knowledge and skill required to do the work