

WHAT THE
BEST COLLEGE
TEACHERS
DO

Ken Bain

HARVARD UNIVERSITY PRESS

Cambridge, Massachusetts

London, England

2004

Copyright © 2004 by the President and Fellows of Harvard College
All rights reserved
Printed in the United States of America

Library of Congress Cataloging-in-Publication Data

Bain, Ken.

What the best college teachers do / Ken Bain.

p. cm.

Includes bibliographical references and index.

ISBN 0-674-01325-5 (alk. paper)

1. College teaching—United States. 2. Effective teaching—United States.

I. Title.

LB2331.B34 2004

378.1'2—dc22 2003062869

CONTENTS

1. Introduction: Defining the Best	1
2. What Do They Know about How We Learn?	22
3. How Do They Prepare to Teach?	48
4. What Do They Expect of Their Students?	68
5. How Do They Conduct Class?	98
6. How Do They Treat Their Students?	135
7. How Do They Evaluate Their Students and Themselves?	150
Epilogue: What Can We Learn from Them?	173
Appendix: How the Study Was Conducted	181
Notes	191
Acknowledgments	201
Index	203



INTRODUCTION: DEFINING THE BEST

When Ralph Lynn graduated from college in 1932, decked out in a variety of academic honors, he began doing other people's laundry to survive the depression. Ten years later, he acquired a correspondence-course teaching certificate and taught high school history classes for six months before entering the army in late 1942. He spent most of World War II in London looking at other people's dirty laundry—censoring soldiers' letters to keep them from revealing too much about troop movements to the folks back home—and reading history. When he came home in 1945, he asked his alma mater, Baylor University, to let him teach. Later, he went north to the University of Wisconsin to acquire a Ph.D. in European history. In 1953 he returned to Texas, where he taught for the next twenty-one years.

When Lynn retired in 1974, more than one hundred of his former students who now held academic posts paid him tribute. One of them, Robert Fulghum, who later wrote a much celebrated book claiming that he learned everything he needed to know about life in kindergarten, confessed that Ralph Lynn was the “best teacher in the world.” Another student, Ann Richards, who became the governor of Texas in 1991, wrote that Lynn's classes “offered us a window to the world, and for a young girl from Waco, his classes were great adventures.” They were, she explained some years after leaving the governor's mansion, like “magical tours into the great minds and movements of history.” Hal Wingo, who took classes from Lynn long before he became the editor of *People* magazine, concluded that Lynn offered the best argument he knew for human cloning. “Nothing would give me more hope for the future,” the

editor explained, “than to think that Ralph Lynn, in all his wisdom and wit, will be around educating new generations from here to eternity.”¹

What did Lynn do to have such a sustained and substantial influence on the intellectual and moral development of his students? What do any of the best college and university teachers do to help and encourage students to achieve remarkable learning results? What does Jeanette Norden, a professor of cell biology who teaches the brain to medical students at Vanderbilt University, do that enables her students to learn so deeply? How does Ann Woodworth, a professor of theater at Northwestern University, lift her acting students to heights of thespian brilliance? Given that human cloning is not an option, is it possible to do some intellectual cloning, to capture the thinking of people like Don Saari from the University of California at Irvine, whose calculus students have sometimes claimed 90 percent of the A’s on departmental examinations? Can we capture the magic of Paul Travis and Suhail Hanna, who taught history and literature in a small freshwater college in Oklahoma in the 1970s and later at other institutions from Pennsylvania to Kansas, inspiring their students to new intellectual levels?

What makes some teachers successful with students of diverse backgrounds? Consider the case of Paul Baker, a teacher who spent nearly fifty years empowering his students to find their own creativity. In the 1940s Baker developed for an undergraduate theater program a course he called “Integration of Abilities,” a mind-charging exploration of the creative process that attracted as many future engineers, scientists, and historians as it did actors and other artists. By the late 1950s, he used the course to build the graduate program in theater at the Dallas Theater Center and later at Trinity University, revolutionizing theater productions around the world. By the 1970s he was employing the integrations method as head of the new performing arts magnet high school in Dallas, changing the lives of

many students whom others had dismissed as failures. In the early 1990s, now retired on a small ranch in East Texas, he took the same approach in creating a program for the local elementary school that pushed standardized test scores in that rural community to historic highs. How did he do it?

For more than fifteen years I have raised such questions in looking at the practices and thinking of the best teachers, those people who have remarkable success in helping their students achieve exceptional learning results. Much of the inspiration for the inquiry came from the extraordinarily successful teachers I have encountered in my own life. It has occurred to me that teaching is one of those human endeavors that seldom benefits from its past. Great teachers emerge, they touch the lives of their students, and perhaps only through some of those students do they have any influence on the broad art of teaching. For the most part, their insights die with them, and subsequent generations must discover anew the wisdom that drove their practices. At best, some small fragment of their talent endures, broken pieces on which later generations perch without realizing the full measure of the ancient wealth beneath them.

A decade ago, I confronted the tragedy of losing some of that wealth in the death of a talented teacher whom I never formally met. When I was a graduate student at the University of Texas in the early 1970s, I heard about a young professor, fresh from his own studies at the University of Chicago, who had students sitting in the aisles for the chance to take his class. Nearly every day, I saw a small army of people follow Tom Philpott from class to the departmental lounge, where they continued the conversations his teaching had started. In the late 1980s my son and daughter-in-law took Philpott's class in U.S. urban history, and I watched as it provoked new questions and perspectives. I listened with renewed interest to their stories of students—even many who were not registered for the

class—who crowded into the legendary teacher’s classroom for a charge to their intellectual batteries. I wanted to interview Philpott about his teaching and possibly videotape some classes, but that chance never came. A short while later he took his own life. His colleagues eulogized him, his students remembered his classes, and perhaps a few of them who became teachers carried some pieces of his talent into their own careers. But for the most part his library of teaching talents and practices burned to the ground when he died. His scholarship on the development of neighborhoods in Chicago remains, but he never captured his own scholarship of teaching, and no one else did it for him.

In this book I have tried to capture the collective scholarship of some of the best teachers in the United States, to record not just what they do but also how they think, and most of all, to begin to conceptualize their practices. The study initially included only a handful of teachers at two universities, but eventually it encompassed professors at two dozen institutions—from open admissions colleges to highly selective research universities. Some taught primarily students with the best academic credentials; others worked with students who had substandard school records. Altogether, my colleagues and I looked at the thinking and practices of between sixty and seventy teachers. We studied nearly three dozen of them extensively, the others, less exhaustively. A few of the latter subjects were speakers in one of the annual series I organized at Vanderbilt and Northwestern that featured professors from other institutions who had achieved impressive teaching results. The subjects came from both medical school faculties and undergraduate departments in a variety of disciplines, including the natural and social sciences, the humanities, and the performing arts. A few came from graduate programs in management, and two came from law schools. We wanted to know what outstanding professors do and think that might explain their accomplishments. Most important, we wanted to know if the lessons they taught us could inform other people’s

teaching. I have directed this book to people who teach, but its conclusions should also be of interest to students and their parents.

DEFINING EXCELLENCE

To begin this study we had to define what we meant by outstanding teachers. That turned out to be a fairly simple matter. All the professors we chose to put under our pedagogical microscope had achieved remarkable success in helping their students learn in ways that made a sustained, substantial, and positive influence on how those students think, act, and feel. The actual classroom performance of the teachers did not matter to us; so long as the teachers did not do their students (or anyone else) any harm in the process, we cared little about how they achieved their results. Dazzling lecture styles, lively classroom discussions, problem-based exercises, and popular field research or projects might or might not contribute to the *telos* of good teaching. Their presence or absence, however, never dictated which people we investigated. We chose teachers because they produced important educational results.

What counted as evidence that a professor profoundly helped and encouraged students to learn deeply and remarkably? That question proved to be more complex. No one type of evidence would do in every case. We simply looked for proof of an educator's excellence, and if we found it, we used that person in the study. In some cases the evidence came in clearly labeled packages; in others, we had to collect it from unmarked jars and piece it together like anthropologists in search of a lost civilization. The types of evidence available depended on both the individual and the discipline.

Jeanette Norden from Vanderbilt University's Medical School and Ann Woodworth from the Theatre Department at Northwestern illustrate two different patterns of evidence. Norden's medical students face a standardized test of their learning in the form of the National Board of Medical Examiners and the United States

Medical Licensing Examination. Their group performance on sections of the exam that cover Norden's field provides a strong indication of her students' learning. So does the students' testimony about how well her class prepared them for the rotation in neurology, the National Boards, and careers in medicine. So do the examinations she uses in her classes, carefully and rigorously constructed instruments that take students through specific cases that require extensive knowledge, advanced understanding, and sophisticated clinical reasoning skills. And so do her colleagues' statements about how well her students are prepared for subsequent work. Norden has won every award for teaching granted by the medical school and selected by the students—some of the awards more times than the university will now allow. When Vanderbilt's chancellor established endowed chairs of teaching excellence in 1993, Norden was the first recipient of that honor. In late 2000, the American Association of Medical Colleges presented her with its Robert Glaser award for teaching excellence.

Ann Woodworth also came with a plethora of teaching awards—including appointment to an endowed chair of teaching excellence at Northwestern. But those recognitions, while important and substantial, gave us no direct evidence about student learning. Woodworth's field certainly emphasizes student performance, but it has no standardized measure of dramatic accomplishments. What convinced us that her teaching was worthy of careful study? First, we had a large body of testimony from her students, not just that she was entertaining or witty, but that she helped them achieve substantial results. We were impressed with the consistency of the testimony, with the kinds of praise the students offered (“you’ll learn more from her class than from any other at this school”; “this class changed my life”), and with the perfect marks they gave her in response to questions about stimulating intellectual interest and helping students learn. Second, we had considerable evidence about

what Woodworth taught, information we gathered from her students, from her account of her courses, and from a term-long observation of one of her classes. Finally, we saw the performances of her students, both in final productions and in classroom work, in which her assistance often transformed a stale rendition into something magical.

Glowing reviews from students and colleagues alone were insufficient, however. We wanted indications from a variety of sources that a particular teacher was worthy of study. Although we did not insist that every instructor present exactly the same kinds of support, we did have two acid tests that all instructors had to meet before we included them in our final results.

First, we insisted on evidence that most of their students were highly satisfied with the teaching and inspired by it to continue to learn. This was no mere popularity contest; we were not interested in people because they were well liked by their students. Rather, we wanted indications from the students that the teacher had “reached them” intellectually and educationally, and had left them wanting more. We rejected the standards of a former dean who used to say, “I don’t care if the students liked the class or not as long as they learned the material,” which meant “I just want to see how they performed on the final.” We too were concerned with how students performed on the final, but we had to weigh the growing body of evidence that students can “perform” on many types of examinations without changing their understanding or the way they subsequently think, act, or feel.² We were equally concerned with how they performed after the final. We were convinced that if students emerged from the class hating the experience, they were less likely to continue learning, or even to retain what they had supposedly gained from the class. A teacher might scare students into memorizing material for short-term recall by threatening punishment or imposing excessively burdensome workloads, but those tactics

might also leave students traumatized by the experience and disliking the subject matter. Any teacher who causes students to hate the subject has certainly violated our principle of “do no harm.”

We recognize that some professors might be enormously successful in helping a few students learn but far less so with most of them. Colleagues have told us about former professors who stimulated their intellectual development but left most students flat. These people obviously valued those mentors and sometimes even modeled their own careers after them, taking pride in what they saw as the elite cadre of their satisfied students, and perhaps even believing that alienation of the masses set them on a higher plane. Such professors may have great value for the academy, but they did not make our cut. We sought people who can make a silk purse out of what others might regard as a sow’s ears, who constantly help their students do far better than anyone else expects.

Our second acid test concerned what students learned. This is tricky because it involved judgments about a variety of disciplines. We sought evidence that colleagues in the field or in closely related fields would regard the learning objectives as worthy and substantial. Yet we remained open to the possibility that some remarkable teachers developed highly valuable learning objectives that ignored the boundaries of the discipline and even, on occasion, offended many disciplinary purists—the medical school professor, for example, who integrated issues of personal and emotional development into a basic science class, helping to redefine the study of medicine. Indeed, most of the highly successful teachers in the study broke traditional definitions of courses, convincing us that success in helping students learn even some core material benefits from the teacher’s willingness to recognize that human learning is a complex process. Thus we had to apply a sweeping sense of educational worth that stemmed not from any one discipline but rather from a broad educational tradition that values the liberal arts (including the natural sciences), critical thinking, problem solving,

creativity, curiosity, concern with ethical issues, and both a breadth and depth of specific knowledge and of the various methodologies and standards of evidence used to create that knowledge.

In short, we included in our study only those teachers who showed strong evidence of helping and encouraging their students to learn in ways that would usually win praise and respect from both disciplinary colleagues and the broader academic community. But we also tried to include some educators who were operating on the fringes of current norms, defining learning wealth in important new ways. We also studied a few people who were highly successful with some classes and less so with others. For example, some teachers achieved wonderful results with large or small classes, advanced or beginning courses, but not with both. Such cases allowed us to make some comparisons between what worked and what did not.

We wanted to study teachers who had a sustained influence on their students, but the evidence for that proved difficult to obtain, especially in the early phases of our research. We talked with some students years after they had taken a particular professor and heard their testimonies about the way the class touched their minds and influenced their lives. We did not, however, systematically follow students; nor did we rely on those interviews alone to decide that someone deserved attention. Instead, we looked for something that would tell us more immediately that the impact was lasting. The concept of deep learners, first developed by Swedish theorists in the 1970s, helped us spot indications of sustained influence.³

We assumed that deep learning was likely to last, and so we listened closely for evidence of it in the language students used to describe their experiences. Did they speak about “learning the material” or about developing an understanding, making something their own, “getting into it,” and “making sense of it all?” We were drawn to classes in which students talked not about how much they had to remember but about how much they came to understand

(and as a result remembered). Some students talked about courses that “transformed their lives,” “changed everything,” and even “messed with their heads.” We looked for signs that students developed multiple perspectives and the ability to think about their own thinking; that they tried to understand ideas for themselves; that they attempted to reason with the concepts and information they encountered, to use the material widely, and to relate it to previous experience and learning. Did they think about assumptions, evidence, and conclusions?

Consider, for example, two sets of comments. One came from students who told us that the class “required a lot of work,” that the professor motivated them to “get it done,” and was thorough and fair, “covering,” as one student put it, “all the stuff that would be on the exam” and “never surprising us with problems we hadn’t seen.” The students dwelled on being successful “in the course” and offered high praise because the instructor helped them achieve that goal. While these comments were all quite favorable, they did not necessarily point to deep learning. In contrast, the second set of students talked about how they could “put a lot of things together now” or “get inside” their own heads. They stressed that they wanted to learn more, sometimes spoke about changing majors to study under a particular professor, and seemed in awe of and fascinated with how much they didn’t know. “I thought it was all cut and dried before I took this course,” one student explained. “It’s pretty exciting stuff.” They talked about issues that the course had raised, how they learned to think differently, how the course had changed their lives, and what they planned to do with what they had learned. They easily discussed arguments they had encountered, questioned assumptions, and distinguished between evidence and conclusions. Students mentioned books they had subsequently read because the course raised their interest, projects they had undertaken, or changes in plans. In commenting about a math class one student explained, “He didn’t just show us how to solve the

problem but helped us think about it so we could do it on our own. I can think through problems better now.” In reference to a history class, that reflection became, “I don’t just memorize stuff in here. I have to think about arguments and evidence.” The second set of comments suggested sustained influences while the first didn’t tell us enough.

As our inquiry developed, it generated enormous interest from colleagues, who often suggested that we consider particular people. All potential subjects entered the study on probation while we examined their learning objectives and pressed them for evidence about success in fostering meaningful results. Sometimes we quietly dropped people from consideration, not because we came to believe that they were ineffective teachers, but because we just did not have enough data to know, one way or the other. My objective in this book is not to notify these colleagues who were not included in the study but to learn as much as possible from the most successful teachers. Consequently, though I mention the names of many people we analyzed, I do not provide a complete list.

CONDUCTING THE STUDY

Once we had identified our subjects, we studied them. Some we observed in the classroom, laboratory, or studio; others, we videotaped. For still others, we did both. We had long conversations with many of the teachers and their students; looked at course materials, including syllabi, examinations, assignment sheets, and even some lecture notes; considered examples of students’ work; conducted what we called “small group analyses,” in which we interviewed entire classes in small groups; asked some people to analyze and describe their own teaching practices and philosophies in more formal reflections; and in a few cases actually sat in on an entire course. The methods of collection and analysis varied, but they all came from approaches common in history, literary analysis, investigative

journalism, and anthropology. The talks we heard, the interviews we conducted, the class materials and other writings we read, and the notes we took while observing a class formed the texts that we subsequently scrutinized (see the appendix for details on the study).

STUDENT RATINGS

Before turning to a summary of the major findings of our study, we should consider one more methodological issue: What role can student ratings play in helping identify outstanding teaching? How did they influence our decisions?

In meeting new faculty members, I have discovered that many teachers have a vague knowledge of the famous Dr. Fox experiments, a knowledge just blurry enough to produce skepticism about any attempt to identify and define teaching excellence. In that study, originally published in the 1970s, three researchers hired an actor to deliver a lecture to a group of educators. They instructed him to make his delivery highly expressive and entertaining but to offer little content in a talk riddled with logical confusions and repetitions. The experimenters gave their “professor” a fictional curriculum vitae, complete with a list of publications, and called him Dr. Fox. When they asked listeners to rate the lecture, the numbers appeared quite favorable, and one of the respondents even claimed to have read some of Dr. Fox’s work.⁴

Many faculty members familiar with this experiment have concluded that student ratings are useless because lectures filled with junk can “seduce” students if the teacher is entertaining. But on closer examination, the original Dr. Fox study had one major flaw: it asked the wrong questions. Many of the questions simply asked if the actor did what he was instructed to do. For example, he had been told to display expressiveness and enthusiasm, and one of the survey questions then asked, “Did he seem interested in his subject?”⁵ No wonder the ratings were so high. Not a single one of the

eight questions asked the audience members if they had learned anything—the element we regarded as so crucial in spotting excellent teaching. Researchers made no effort to test the listeners on the knowledge they had gained from the lectures (although subsequent experiments with Dr. Fox did so), or even to ask them whether they believed they had in fact learned anything.

Far less well known and publicized were the subsequent studies done on what came to be known as the “Dr. Fox effect,” which pointed out these methodological flaws in the original study and drew far more conservative conclusions from the investigations. All told, what we can learn about identifying teaching excellence from the Dr. Fox experiments seems pretty meager. At best, they may help us understand what questions we should and shouldn’t be asking on the student rating forms. Rather than asking if professors were expressive or used a particular technique, we should ask if they helped students learn or stimulated their interest in the subject. Indeed, research has found high positive correlations between student ratings and external measures of student learning when such questions are used.⁶ Most important, student ratings can, as one observer put it, “report the extent to which the students have been reached [educationally].”⁷ If we want to know if students think that something has helped and encouraged them to learn, what better way to find out than to ask them. As for expressiveness, Herbert Marsh, an Australian researcher, and others found in subsequent Dr. Fox experiments that students usually perform better on examinations after hearing exciting lecturers than they do after dull ones, but that should surprise no one.⁸

Students do not always have sophisticated definitions of what it means to learn in a particular discipline. Thus we could not rely on the numbers alone to tell us whether someone had been helping people learn at the high level expected in this study. That information came only from looking at course materials, including syllabi and methods of evaluation, or from interviewing both instructors

and their students. Student ratings could help supplement these more qualitative inquiries, especially the numbers that emerge from questions like the two that appeared on both the Northwestern and the Vanderbilt student rating forms: Rate how much the teaching helped you learn, and rate how well the course stimulated you intellectually.

Yet many people remain highly suspicious of any study of teaching quality that draws even part of its evidence from student ratings. Educators not familiar with the Dr. Fox experiments may have a headline acquaintance with a more recent study. In 1993 Nalini Ambady and Robert Rosenthal showed students short film clips of teachers and asked them to rate those professors using the same instrument others had used after taking classes with the same instructors.⁹ The researchers wanted to know how small the exposure could be and still generate ratings that were substantially the same as those that came after an entire semester of viewing the professor. When *Lingua Franca* and other publications reported that high positive correlations began to appear after the experimental group saw only a few seconds of the professor, some academics came to believe that all student ratings arise from superficial observations and amount to little more than the most primitive of popularity tests. These critics failed to consider, however, that the Ambady and Rosenthal study could point to a much different conclusion: students, with long histories of dealing with both highly stimulating and discouraging teachers, may develop an ability to guess quite accurately, even after only a few seconds of exposure, which professors will ultimately advance their education and which will not. In short, the instant judgments may stem from concerns about who can help them learn and grow rather than from a focus on amorphous qualities of personality and friendship. Ambady and Rosenthal made this point in their article: “Not only do we possess the remarkable ability to form impressions of others . . . but, perhaps more remarkably, the impressions that we form can be quite accurate!”

For our part, we have relied not on instant impressions but rather on the kind of detailed and sustained study outlined above and discussed more fully in the pages to come. I will return in the last chapter to the process of evaluating teaching, but for now it is worth emphasizing that this study follows the criteria of outcomes. We identify teaching excellence when we see evidence about remarkable feats of student learning and indications that the teaching helped and encouraged those results; we learn something about developing teaching excellence when we try to discover what fostered that educational success. Ratings from students of how much they learned and whether the professor stimulated their interests and intellectual development often told us a good deal about the quality of teaching, but we looked at far more evidence before concluding that it was exceptional.

MAJOR CONCLUSIONS

Let's begin with the major conclusions of this study, the broad patterns of thinking and practice we found among our subjects. One word of caution, however: anyone who expects a simple list of do's and don'ts may be greatly disappointed. The ideas here require careful and sophisticated thinking, deep professional learning, and often fundamental conceptual shifts. They do not lend themselves to teaching by the numbers.¹⁰

Our conclusions emerge from six broad questions we asked about the teachers we examined.

1. What Do the Best Teachers Know and Understand?

Without exception, outstanding teachers know their subjects extremely well. They are all active and accomplished scholars, artists, or scientists. Some have long and impressive publication lists, the kind the academy has long valued. Others have more modest records; or in a few cases, virtually none at all. But whether well

WHAT DO THEY KNOW ABOUT HOW WE LEARN?



In the early 1980s, two physicists at Arizona State University wanted to know whether a typical introductory physics course, with its traditional emphasis on Newton's laws of motion, changed the way students thought about motion. As you read this account, you might substitute for the line "think about motion" any other phrase that fits your subject. Do the students in any class change the way they think?

To find out, Ibrahim Abou Halloun and David Hestenes devised and validated an examination to determine how students understand motion. They gave the test to people entering the classes of four different physics professors, all good teachers according to both colleagues and their students. On the front side, the results surprised no one. Most students entered the course with an elementary, intuitive theory about the physical world, what the physicists called "a cross between Aristotelian and 14th-century impetus ideas." In short, they did not think about motion the way Isaac Newton did, let alone like Richard Feynman. But that was before the students took introductory physics.

Did the course change student thinking? Not really. After the term was over, the two physicists gave their examination once more and discovered that the course had made comparatively small changes in the way students thought.¹ Even many "A" students continued to think like Aristotle rather than like Newton. They had memorized formulae and learned to plug the right numbers into them, but they did not change their basic conceptions. Instead, they had interpreted everything they heard about motion in terms

of the intuitive framework they had brought with them to the course.

Halloun and Hestenes wanted to probe this disturbing result a little further. They conducted individual interviews with some of the people who continued to reject Newton's perspectives to see if they could dissuade them from their misguided assumptions. During those interviews, they asked the students questions about some elementary motion problems, questions that required them to rely on their theories about motion to predict what would happen in a simple physics experiment. The students made their projections, and then the researchers performed the experiment in front of them so they could see whether they got it right. Obviously, those who relied on inadequate theories about motion had faulty predictions. At that point, the physicists asked the students to explain the discrepancy between their ideas and the experiment.

What they heard astonished them: many of the students still refused to give up their mistaken ideas about motion. Instead, they argued that the experiment they had just witnessed did not exactly apply to the law of motion in question; it was a special case, or it didn't quite fit the mistaken theory or law that they held as true. "As a rule," Halloun and Hestenes wrote, "students held firm to mistaken beliefs even when confronted with phenomena that contradicted those beliefs." If the researchers pointed out a contradiction or the students recognized one, "they tended at first not to question their own beliefs, but to argue that the observed instance was governed by some other law or principle and the principle they were using applied to a slightly different case."² The students performed all kinds of mental gymnastics to avoid confronting and revising the fundamental underlying principles that guided their understanding of the physical universe. Perhaps most disturbing, some of these students had received high grades in the class.

This story is part of a small but growing body of literature that

questions whether students always learn as much as we have traditionally thought they did. The scholarly work on this issue asks not if students can pass our examinations but whether their education has a sustained, substantial, and positive influence on the way they think, act, and feel. Researchers have found that even some “good” students may not progress as much intellectually as we once thought. They have discovered that some people make A’s by learning to “plug and chug,” memorizing formulae, sticking numbers in the right equation or the right vocabulary into a paper, but understanding little. When the class is over, they quickly forget much of what they have “learned.”³ Participants at a 1987 conference on science education, for example, saw this problem in math. “Those who successfully complete calculus,” they concluded, “frequently fail to gain a conceptual understanding of the subject or an appreciation of its importance” because instructors rely on “‘plug and chug’ exercises that have little connection with the real world.”⁴ Even when learners have acquired some conceptual understanding of a discipline or field, they are often unable to link that knowledge to real-world situations or problem-solving contexts.

LEARNING FROM THE BEST

What do the best teachers know that helps them overcome—at least partially and sometimes fully—these problems?

We discovered that they know their disciplines well and are active and accomplished scholars, artists, or scientists—even if they do not always have long publication records. But that necessary knowledge alone can’t account for their teaching success. If it did, then any expert in the field would become an outstanding educator, but that clearly doesn’t happen. Nor is it the case that experts just need more time to become better teachers. We encountered many professors, all eminent scholars in their fields, who spent hours

crafting lectures that reflected the latest and most advanced scholarly and scientific knowledge only to produce students who understood little of that sophistication. One of those people, a medical school professor who was not part of the study, once told us with both pride and some measure of frustration that he didn't worry about whether students "got it" as long as every line of his lectures reflected the "highest standards of scientific quality and cutting-edge knowledge in the field."

What else do the best teachers know that might explain their successes in helping students learn deeply? We found two other kinds of knowledge that seem to be at play. First, they have an unusually keen sense of the histories of their disciplines, including the controversies that have swirled within them, and that understanding seems to help them reflect deeply on the nature of thinking within their fields. They can then use that ability to think about their own thinking—what we call "metacognition"—and their understanding of the discipline *qua* discipline to grasp how other people might learn. They know what has to come first, and they can distinguish between foundational concepts and elaborations or illustrations of those ideas. They realize where people are likely to face difficulties developing their own comprehension, and they can use that understanding to simplify and clarify complex topics for others, tell the right story, or raise a powerfully provocative question. There's a catch to all this, however. That kind of understanding is obviously rooted in each individual field of study and defies generalization.

Yet something else seems to be at work that transcends the various disciplines and therefore is more useful to our general study. To put it simply, the people we analyzed have generally cobbled together from their own experiences working with students conceptions of human learning that are remarkably similar to some ideas that have emerged in the research and theoretical literature on

cognition, motivation, and human development. Those ideas help them understand and cope with situations like the physics story and myriad other learning problems.

Here are the key concepts we found.

1. Knowledge Is Constructed, Not Received

Perhaps the best way to understand this notion is to contrast it with an older idea. According to the traditional view, memory is a great storage bin. We put knowledge in it and then later pick out what we need. Thus you often hear people say, “My students must learn the material before they can think about it,” presumably meaning that they must store it somewhere for later use.

The best teachers don’t think of memory that way, and neither do a lot of learning scientists. Instead, they say that we construct our sense of reality out of all the sensory input we receive, and that process begins in the crib. We see, hear, feel, smell, and taste, and we begin connecting all those sensations in our brains to build patterns of the way we think the world works. So our brains are both storage and processing units. At some point, we begin using those existing patterns to understand new sensory input. By the time we reach college, we have thousands of mental models, or schemas, that we use to try to understand the lectures we hear, the texts we read, and so forth.

For example, I have a mental model of something called a classroom. When I enter a room and receive some sensory input through the lens in my eyes, I understand the input in terms of that previously existing model, and I know I’m not in a train station. But this enormously useful ability can also present problems for learners. When we encounter new material, we try to comprehend it in terms of something we think we already know. We use our existing mental models to shape the sensory inputs we receive. That means that when we talk to students, our thoughts do not travel seamlessly from our brains to theirs. The students bring paradigms to the class

that shape how they construct meaning. Even if they know nothing about our subjects, they still use an existing mental model of something to build their knowledge of what we tell them, often leading to an understanding that is quite different from what we intend to convey. “The trouble with people,” Josh Billings once remarked, “is not that they don’t know but that they know so much that ain’t so!”

I’m not just saying that students bring misconceptions to class, as a philosophy professor concluded a few years ago when he heard these ideas in a workshop. Actually, I’m arguing something much more fundamental: the teachers we encountered believe everybody constructs knowledge and that we all use existing constructions to understand any new sensory input. When these highly effective educators try to teach the basic facts in their disciplines, they want students to see a portion of reality the way the latest research and scholarship in the discipline has come to see it. They don’t think of it as just getting students to “absorb some knowledge,” as many other people put it. Because they believe that students must use their existing mental models to interpret what they encounter, they think about what they do as stimulating construction, not “transmitting knowledge.” Furthermore, because they recognize that the higher-order concepts of their disciplines often run counter to the models of reality that everyday experience has encouraged most people to construct, they often want students to do something that human beings don’t do very well: build new mental models of reality.

But that’s the problem.

2. Mental Models Change Slowly

How can we stimulate students to build new models, to engage in what some call “deep” learning as opposed to “surface” learning in which they merely remember something long enough to pass the examination? Our subjects generally believe that to accomplish that feat, learners must (1) face a situation in which their mental model

will not work (that is, will not help them explain or do something); (2) care that it does not work strongly enough to stop and grapple with the issue at hand; and (3) be able to handle the emotional trauma that sometimes accompanies challenges to longstanding beliefs.

The teachers in our study often talked about “challenging students intellectually.” That meant they wanted to create what some of the literature calls an “expectation failure,” a situation in which existing mental models will lead to faulty expectations, causing their students to realize the problems they face in believing whatever they believe. Yet these highly effective teachers realized that human beings face too many expectation failures in life to care about all of them, so students may not engage in the deep thinking required to build completely new models. Furthermore, they understood that people have so many paradigms of reality that they may not know which of their schemas has led to the faulty predictions, so they may correct the wrong ones. That’s partly where the physics students went wrong when they encountered experiments in which their conceptions of motion did not work. Finally, the best teachers understood that their students may find so much emotional comfort in some existing model of reality that they cling to it even in the face of repeated expectation failures.

Such ideas have important implications for the teachers. They conduct class and craft assignments in a way that allows students to try their own thinking, come up short, receive feedback, and try again. They give students a safe space in which to construct ideas, and they often spend a great deal of time creating a kind of scaffolding to help students engage in that construction (which is different from the popular notion of “covering” the material, but in ways that are sometimes difficult to grasp). Because they attempt to place students in situations in which some of their mental models will not work, they try to understand those models and the emotional baggage attached to them. They listen to student conceptions before

challenging them. Rather than telling students they are wrong and then providing the “correct” answers, they often ask questions to help students see their own mistakes.

Perhaps this general approach is most apparent in the way the teachers in the study approached a controversy that still rages in many disciplines, from the sciences to the humanities. On one side of that debate, teachers have argued that students cannot learn to think, to analyze, to synthesize, and to make judgments until they “know” the “basic facts” of the discipline. People in this school of thought have tended to emphasize the delivery of information to the exclusion of all other teaching activities. They seldom expect their students to reason (that will supposedly come after they have “learned the material”). On their examinations, these professors often test for recall, or simple recognition of information (on a multiple-choice examination, for example).

Teachers in our study come down on the other side of that controversy. They believe that students must learn the facts *while* learning to use them to make decisions about what they understand or what they should do. To them, “learning” makes little sense unless it has some sustained influence on the way the learner subsequently thinks, acts, or feels. So they teach the “facts” in a rich context of problems, issues, and questions.

Consider the approaches of two anatomy professors, one who has been enormously successful and was included in the study and the other, outside the study, who has, to put it gently, had difficulty fostering learning. The latter insisted that students must simply “learn the facts.” There “isn’t much to discuss,” he told us. “The structure of the human body is well known by scientists, and students must simply absorb a lot of facts. There isn’t any other way to teach except to stand in front of them and give them those facts. We can’t discuss the way you might in a literature class.” He talked about “transmitting” knowledge and insisted that the primary objective of the course was to “memorize large chunks of information.”