MAKE LEARNING VISIBLE IN MATHEMATICS

1

2 + ***2*** = ***4*.**

It just adds up, right? But think about how you know that two plus two equals four. Did you memorize the answer from a flashcard? Did someone tell you that and then expect that you accept it as truth? Did you discover the answer while engaged in a relevant task? Were you asked to explore a concept, and when you grasped the concept, someone provided you with labels for the ideas? In all likelihood, it was a combination of these things that led you to come to understand the concept of the number two, the possibility of combining like items, and the idea that the sum is a result of these combinations. Over time, you were able to consider an unknown term such as x in the equation $2 + x = 4$ and master increasingly complex ideas that are based on algebraic thinking. Your learning became visible to you, your teachers, and your family.

And that's what this book is about—making learning visible. By visible learning, we mean several things. First and foremost, students and teachers should be able to see and document learning. Teachers should understand the impact that they, and their actions, have on students. Students should also see evidence of their own progress toward their learning goals. Visible learning helps teachers identify attributes and influences that work. Visible learning also helps teachers better understand their impact on student learning, and helps students become their own teachers. In this way, both teachers and students become lifelong learners and develop a love for learning. Importantly, this is not a book about visible teaching. We do, of course, provide evidence for various teacher moves, but our goal is not to make teaching visible but rather the *learning* visible. Before we explore the research behind visible learning, let's consider the ways in which you may have been taught mathematics. We need to accept and understand that high-quality learning may require that we discard ineffective pedagogy that we may have experienced as learners of mathematics.

Forgetting the Past

Do you remember the *Men in Black* movies? The agents who are protecting the universe have neuralyzers, which erase memories. They use them to erase encounters with intergalactic aliens so that people on planet Earth are kept in the dark about threats to their world. We wish we had that little flashy thing. If we did, we'd erase teachers' memories of some of the ways they were taught mathematics when they were younger. And we'd replace those memories with intentional instruction, punctuated with collaborative learning opportunities, rich discussions about mathematical concepts, excitement over persisting through complex problem solving, and the application of ideas to situations and problems that matter. We don't mean to offend anyone, but we have all suffered through some pretty bad mathematics instruction in our lives. Nancy remembers piles of worksheets. Her third-grade teacher had math packets that she distributed the first of each month. Students had specific calculation-driven problems that they had to do every night, page after page of practicing computation with little or no context. A significant amount of class time was spent reviewing the homework, irrespective of whether or not students got the problem wrong or right. In fact, when she asked if they could skip the problems everyone completed correctly, she was invited to have a meeting with the teacher and the principal.

In algebra, Doug's teacher required that specifically assigned students write out one of their completed homework problems on the chalkboard while the teacher publicly commended or criticized people. Doug wasn't academically prepared for entry-level algebra, so he hid outside the classroom until the teacher ran out of problems each day. (He took the tardies rather than show everyone he didn't understand the homework.) When this ritual was completed, the teacher explained the next section of the textbook while students took notes. The teacher wrote on an overhead projector with rollers on each side, winding away, page after page. Doug learned to copy quickly into his Cornell notes since the teacher often accidentally erased much of what he wrote because of his left-hand hook writing style. When finished with this, students were directed to complete the assigned odd-numbered problems from the back of the book in a silent classroom. Any problems not completed during class time automatically became homework. Doug copied from his friend Rob on the bus ride home each day but failed every test. This spectator sport version of algebra did not work for students who did not already know the content. Doug's learning wasn't visible to himself, or to his teacher.

If you're worrying about Doug, after failing algebra in ninth grade, he then had a teacher who was passionate about her students' learning. She modeled her thinking every day. She structured collaborative group tasks and assigned problems that were relevant and interesting. Doug eventually went on to earn a master's degree in bio-statistics.

John did okay in mathematics and enjoyed the routines, but if offered, he would have dropped mathematics at the first chance given. But his school made all students enroll in mathematics right to the last year of high school. It was in this last year that he met Mr. Tomlinson—rather strict, a little forbidding, but dedicated to the notion that every one of his students should share his passion for mathematics. He gave his students the end-of-the-year high-stakes exam at the start of the year to show them where they needed to learn. Though the whole class failed, Mr. Tomlinson was able to say, "This is the standard required, and I am going to get you all to this bar." Throughout the year, Mr. Tomlinson persistently engaged his students in how to think in mathematics, working on spotting similarities and differences in mathematical problems so they did not automatically make the same mistakes every time. This teacher certainly saw something in John that John did not see in himself. John ended up with a minor in statistics and major in psychometrics as part of his doctoral program.

Learning isn't linear; it's recursive.

These memories of unfortunate mathematics instruction need to be erased by *Men in Black* Agent K using his neuralyzer, as we know that one of the significant impacts on the way teachers teach is how they were taught. We want to focus on the good examples—the teachers we remember who guided our understanding and love of mathematics.

We've already asked you to forget the less-than-effective learning experiences you've had, so we feel comfortable asking you one more thing. Forget about prescriptive curricula, scripted lesson plans, and worksheets. Learning isn't linear; it's recursive. Prescriptive curriculum isn't matched to students' instructional needs. Sometimes students know more than the curriculum allows for, and other times they need a lot of scaffolding and support to develop deep understanding and skills. As we will discuss later in this book, it's really about determining the impact that teachers have on students and making adjustments to ensure that the impact is as significant as possible.

A major flaw of highly scripted lessons is that they don't allow teachers to respond with joy to the errors students make. Yes, joy. Errors help teachers understand students' thinking and address it. Errors should be celebrated because they provide an opportunity for instruction, and thus learning. As Michael Jordan noted in his Nike ad, "I've missed more than 9,000 shots in my career. I've lost almost 300 games. 26 times, I've been trusted to take the game winning shot and missed. I've failed over and over and over again in my life. And that is why I succeed."

Linda remembers playing a logic game using attribute blocks with her students. The beginning of the game required that students listen carefully to the ideas of others and draw some conclusions as to whether those ideas were correct or accurate. At one point, she commented to an incorrect response, "That's a really important mistake. I hope you all heard it!" The reaction of almost every student was a look of surprise. It was as if the students were thinking, "Have you lost your mind? The goal in math is to get it right!" That response made a real impact on Linda's teaching moves in terms of recognizing how important it is for students to understand they learn and develop understanding from making mistakes (and, in fact, she still says that to this day!). The very best mathematicians wallow in the enjoyment of struggling with mathematical ideas, and this should be among the aims of math teachers—to help students enjoy the struggle of mathematics.

When students don't make errors, it's probably because they already know the content and didn't really need the lesson. We didn't say throw away textbooks. They are a resource that can be useful. Use them wisely, and make adjustments as you deem necessary to respond to the needs of your students. Remember, it is your students, not the curriculum writers, who direct the learning in your classroom.

What Makes for Good Instruction?

When we talk about high-quality instruction, we're always asked the chicken-and-egg question: "Which comes first?" Should a mathematics lesson start with teacher-led instruction or with students attempting to solve problems on their own? Our answer: it depends. It depends on the learning intention. It depends on the expectations. It depends on students' background knowledge. It depends on students' cognitive, social, and emotional development and readiness. It depends where you are going next (and there needs to be a next). And it depends on the day. Some days, lessons start with collaborative tasks. Other days, lessons are more effective when students have an opportunity to talk about their thinking with the entire class or see worked examples. And still other days, it's more effective to ask students to work individually. Much of teaching is dependent on responding to student data in real time, and each teacher has his or her own strengths and personality that shine through in the best lessons. Great teachers are much like jazz musicians, both deliberately setting the stage and then improvising. Great teachers have plans yet respond to student learning and needs in real time.

Errors help teachers understand students' thinking and address it. . . . They provide an opportunity for instruction, and thus learning.

But even the most recognized performers had to learn techniques before applying them. Jazz musicians have to understand standards of music, even if they choose to break the rules. Similarly, great teachers need to know the tools of their craft before they can create the most effective lessons. Enter *Visible Learning*.

The Evidence Base

The starting point for our exploration of learning mathematics is John's books, *Visible Learning* (2009) and *Visible Learning for Teachers* (2012). At the time these books were published, his work was based on more than 800 meta-analyses conducted by researchers all over the world, which included more than 50,000 individual studies that included more than 250 million students. It has been claimed to be the most comprehensive review of educational research ever conducted. And the thing is, it's still going on. At the time of this writing, the database included 1200 meta-analyses, with more than 70,000 studies and 300 million students. A lot of data, right? But the story underlying the data is the critical matter; and it has not changed since the first book in 2009.

Meta-Analyses

A **meta-analysis** is a statistical tool for combining findings from different studies with the goal of identifying patterns that can inform practice.

Before we explore the findings, we should discuss the idea of a meta-analysis because it is the basic building block for the recommendations in this book. At its root, a **meta-analysis** is a statistical tool for combining findings from different studies with the goal of identifying patterns that can inform practice. It's the old preponderance of evidence that we're looking for, because individual studies have a hard time making a compelling case for change. But a meta-analysis synthesizes what is currently known about a given topic and can result in strong recommendations about the impact or effect of a specific practice. For example, there was competing evidence about periodontitis (inflammation of the tissue around the teeth) and whether or not it is associated with increased risk of coronary heart disease. The published evidence contained some conflicts, and recommendations about treatment were piecemeal. A meta-analysis of five prospective studies with 86,092 patients suggested that individuals with periodontitis had a 1.14 times higher risk of developing coronary heart disease than the controls (Bahekar, Singh, Saha, Molnar, & Arora, 2007). The result of the meta-analysis was a set of clear recommendations for treatment of periodontitis, such as the use of scaling and root planing (SRP), or deep cleaning of the teeth, as initial treatment. The evidence suggests that this has the potential of significantly reducing the incidence of heart disease. While this book is not about health care or business, we hope that the value of meta-analyses in changing practice is clear.

The statistical approach for conducting meta-analyses is beyond the scope of this book, but it is important to note that this tool allows researchers to identify trends across many different studies and their participants.

Video 1.1
What Is Visible Learning for Mathematics?

http://resources.corwin.com/VL-mathematics

To read a QR code, you must have a smartphone or tablet with a camera. We recommend that you download a QR code reader app that is made specifically for your phone or tablet brand.

Effect Sizes

The meta-analyses were used to calculate effect sizes for each practice. You might remember from your statistics class that studies report statistical significance. Researchers make the case that something "worked" when chance is reduced to 5 percent (as in $p < 0.05$) or 1 percent (as in $p < 0.01$). What they really mean is that the probability of seeing the outcome found as the result of chance events is very small, less than 5 percent or less than 1 percent. One way to increase the likelihood that statistical significance is reached is to increase the number of people in the study, also known as sample size. We're not saying that researchers inflate the size of the research group to obtain significant findings. We are saying that simply because something is statistically significant doesn't mean that it's worth implementing. For example, if the sample size was 1,000 participants, then a correlation only needs to exceed 0.044 to be considered "statistically significant," meaning the results are due to factors other than chance; if 10,000 are sampled, then a correlation of 0.014 is needed, or if 100,000 are sampled, then a correlation of 0.004 is sufficient to show a nonchance relationship. Yes, you can be confident that these values are greater than zero, but are they of any practical value? That's where effect size comes in.

Say, for example, that a digital app was found to be statistically significant in changing students' learning in geometry. Sounds good, you say to yourself, and you consider purchasing or adopting it. But then you learn that it increased students' performance by only three right answers for every twenty-five choices (and the research team had data from 9,000 students). If it were free and easy to implement this change, it might be worth it to have students get a tiny bit better as users of geometric knowledge. But if it were time-consuming, difficult, or expensive, you

Effect size represents the magnitude of the impact that a given approach has.

EFFECT SIZE FOR SELF-VERBALIZATION AND SELF-QUESTIONING = **0.64**

should ask yourself if it's worth it to go to all of this trouble for such a small gain. That's **effect size**—it represents the magnitude of the impact that a given approach has.

Visible Learning provides readers with effect sizes for many influences under investigation. As an example, self-verbalization and self-questioning—students thinking and talking about their own learning progress—has a reasonably strong effect size at 0.64 (we'll talk more about what the effect size number tells us in the next section). The effect sizes can be ranked from those with the highest impact to those with the lowest. But that doesn't mean that teachers should just take the top ten or twenty and try to implement them immediately. Rather, as we will discuss later in this book, some of the highly useful practices are more effective when focused on surface learning (initial acquisition of knowledge) while others work better for deep learning (consolidation of knowledge) and still others work to encourage transfer (application to new and novel situations).

Noticing What Does and Does Not Work

If you attend any conference or read just about any professional journal, not to mention subscribe to blogs or visit Pinterest, you'll get the sense that everything works. Yet educators have much to learn from practices that do not work. In fact, we would argue that learning from what doesn't work, and not repeating those mistakes, is a valuable use of time. To determine what doesn't work, we turn our attention to effect sizes again. Effect sizes can be negative or positive, and they scale from low to high. Intuitively, an effect size of 0.60 is better than an effect size of 0.20. Intuitively, we should welcome any effect that is greater than zero, as zero means "no growth," and clearly any negative effect size means a negative growth. If only it was this simple.

An **influence** is an instructional strategy, idea, or tool we use in schools.

It turns out that about 95 percent or more of the **influences** (instructional strategies, ideas, or tools) that we use in schools have a positive effect; that is, the effect size of nearly everything we do is greater than zero. This helps explain why so many people can argue "with evidence" that their pet project works. If you set the bar at showing any growth above zero, it is indeed hard to find programs and practices that don't work. As described in *Visible Learning* (2009), we have to reject the starting point of zero. Students naturally mature and develop over the course of a year, and thus actions, activities, and interventions that teachers use should *extend learning beyond what a student can achieve by simply attending school for a year*.

This is why John set the bar of acceptability higher—at the average of all the influences he compiled—from the home, parents, schools, teachers, curricula, and teaching strategies. This average was 0.40, and John called it the "**hinge point**." He then undertook studying the underlying attributes that would explain why those influences higher than 0.40 had such a positive impact compared with those lower than 0.40. His findings were the impetus for the *Visible Learning* story. We expect, at minimum, students' learning to progress a full year for every year that they are in school. And we hope that students gain more than that. Ensuring this level of growth requires a relentless focus on learning rather than on teaching.

Hinge point is the average point at which we can consider that something is working enough for a student to gain one year's growth for a year of schooling.

Borrowing from *Visible Learning*, the **barometer of influence** and hinge point are effective in explaining what we focus on in this book and why. Here's an example of how this might play out in learning mathematics. Let's focus on volunteer tutors, which some have argued could be used to address the basic skills needs that some students have in mathematics. In essence, students are taught by volunteers, often parents or university students, and this instruction focuses on topics such as adding fractions, long division, or some other skill. Importantly, we are not advocating for skills-based instruction, but rather using this example to highlight the use of effect sizes. As with much of the educational research, there are studies that contradict other studies. For example, Scott (2007) described an experiment in engaging parents as volunteers to boost mathematics learning. She suggests that the effort was worthwhile but does not provide information on the impact it had in terms of learning that exceeded one year. Similarly, Carmody and Wood (2009) describe a volunteer tutoring program, this time with college seniors tutoring their younger peers in college mathematics classes. They report that their effort was generally well received, but do not provide information about the impact that it had on students' learning. That's where the meta-analyses and effect size data can teach us. The barometer and hinge point for volunteer tutors are presented in Figure 1.1. Note that this approach rests in the zone of "teacher effects," which is below the level of desired effects but better than reverse effects. Our focus in *Visible Learning for Mathematics* is on actions that fall inside the *zone of desired effects*, which is 0.40 and above. When actions are in the range of 0.40 and above, the data suggest that the effort extends beyond that which was expected from attending school for a year.

The **barometer of influence** is a visual scale that can help us understand where an influence falls in terms of relative effect size.

EFFECT SIZE FOR VOLUNTEER TUTORS = **0.26**

THE BAROMETER FOR THE INFLUENCE OF VOLUNTEER TUTORS

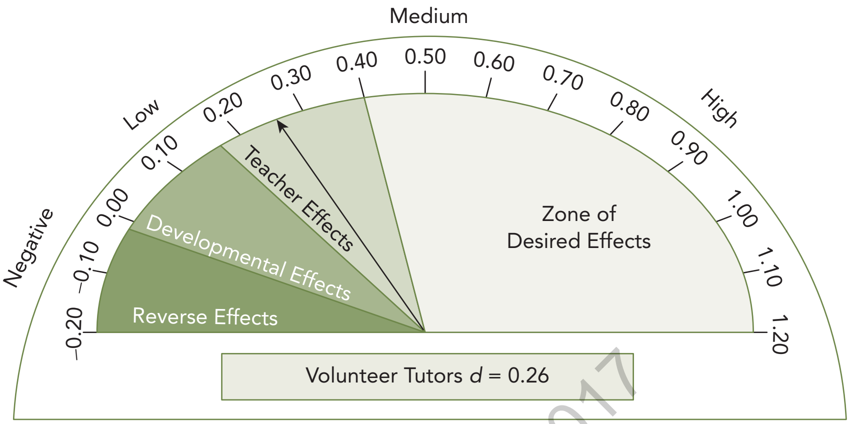

Source: Adapted from Hattie (2012).

Figure 1.1

We expect, at minimum, students' learning to progress a full year for every year that they are in school.

Caution: That doesn't mean that everything below 0.40 effect size is not worthy of attention. Hattie (2012) points out that the hinge point of 0.40 is not absolute. In actuality, each influence does have its own hinge point; therefore the hinge point of 0.40 is simply a good starting point for discussion about the nuances, variability, quality of the studies, and other factors that give an influence a particular effect size. It's just not black-and-white, and there are likely some useful approaches for teaching and learning that are not above this average. For example, drama and arts programs have an effect size of 0.35, almost ensuring that students gain a year's worth of achievement for a year of education. We are not suggesting that drama and art be removed from the curriculum. In fact, artistic expression and aesthetic understanding may be valuable in and of themselves.

EFFECT SIZE FOR SIMULATIONS = **0.33**

It is also important to realize that some of the aggregate scores mask situations in which specific actions can be strategically used to improve students' understanding. Simulations are a good case. The effect size of simulations is 0.33, below the threshold that we established. But what if simulations were really effective in deepening understanding, but not as useful when used for surface learning? (See Chapters 4 and 5 for more on surface and deep learning.) In this case, the strategic deployment of simulations could be important. There are situations like this that we will review in this book as we focus on the balance and sequencing of surface learning compared with deep learning or transfer learning. For now, let's turn our attention to actions that teachers can take to improve student learning. We'll start by directly addressing a major debate in mathematics education: direct instruction compared with dialogic approaches.

Direct and Dialogic Approaches to Teaching and Learning

Debates about the teaching of mathematics have raged for decades. In general, the debate centers on the role of direct instruction versus dialogic instruction, with some teachers and researchers advocating for one or the other. Proponents of both models of instruction have similar goals—student mastery of mathematics. But they differ in the ways in which learning opportunities are organized within the context of a lesson. According to Munter, Stein, and Smith (2015b):

> In the direct instruction model, when students have the prerequisite conceptual and procedural knowledge, they will learn from (a) watching clear, complete demonstrations of how to solve problems, with accompanying explanations and accurate definitions; (b) practicing similar problems sequenced according to difficulty; and (c) receiving immediate, corrective feedback. Whereas in the dialogic model, students must (a) actively engage in new mathematics, persevering to solve novel problems; (b) participate in a discourse of conjecture, explanation, and argumentation; (c) engage in generalization and abstraction, developing efficient problem-solving strategies and relating their ideas to conventional procedures; and to achieve fluency with these skills, (d) engage in some amount of practice. (p. 6)

As the authors note, there are several similarities and some important differences between these two competing models. In terms of similarities, both focus on students' conceptual understanding and procedural fluency. In other words, students have to know the *why* and *how* of mathematics.

Neither model advocates that students simply memorize formulas and procedures. As the National Council of Teachers of Mathematics (2014) states, procedural fluency is built on a foundation of conceptual understanding. Students need to develop strategic reasoning and problem solving. To accomplish this, both models suggest that (1) mathematics instruction be carefully designed around rigorous mathematical tasks, (2) students' reasoning is monitored, and (3) students are provided ample opportunities for skill- and application-based practice.

Direct instruction is when the teacher decides the learning intentions and success criteria, makes them transparent to the students, demonstrates them by modeling, evaluates if they understand what they have been told by checking for understanding, and re-tells them what they have been told by tying it all together with closure.

Munter, Stein, and Smith (2015b) also identify a number of differences between the two models, namely in the types of tasks students are invited to complete, the role of classroom discourse, collaborative learning, and the role of feedback. Figure 1.2 contains their list of similarities and differences. Importantly, these researchers also recognize that teachers use aspects of each model. As they note, "teachers in dialogic classrooms may very well demonstrate some procedures, just as students in a direct instruction classroom may very well engage in project-based activities" (p. 9). They argue that the purposes for using different aspects of each model may vary, and the outcomes may be different, but note that "high-quality instruction must include the identification of both instructional practices and the underlying rationales for employing those practices" (p. 9).

We agree that direct instruction should not be thought of as "spray-and-pray" didactic show-and-tell transmission of knowledge. Neither direct nor dialogical instruction should be confused with "lots of talking" or didactic approaches. John (Hattie, 2009) defines **direct instruction** in a way that conveys an intentional, well-planned, and student-centered guided approach to teaching. "In a nutshell, the teacher decides the learning intentions and success criteria, makes them transparent to the students, demonstrates them by modeling, evaluates if they understand what they have been told by checking for understanding, and re-tells them what they have been told by tying it all together with closure" (p. 206).

EFFECT SIZE FOR DIRECT INSTRUCTION = **0.59**

EFFECT SIZE FOR CLASSROOM DISCUSSION = **0.82**

When thinking of direct instruction in this way, the effect size is 0.59. Dialogic instruction also has a high effect size of 0.82. This doesn't mean that teachers should always choose one approach over another. It should never be an either/or situation. The bigger conversation, and purpose of this book, is to show how teachers can choose the right approach at the right time to ensure learning, and how both dialogic and direct approaches have a role to play throughout the learning process, but in different ways.

COMPARING DIRECT AND DIALOGIC INSTRUCTION

Dialogic Instruction	Distinction	Direct Instruction
Fundamental to both knowing and learning mathematics. Students need opportunities in both small-group and whole-class settings to talk about their thinking, questions, and arguments.	The importance and role of **talk**	Most important during the guided practice phase, when students are required to explain to the teacher how they have solved problems in order to ensure they are encoding new knowledge.
Provides a venue for more talking and listening than is available in a totally teacher-led lesson. Students should have regular opportunities to work on and talk about solving problems in collaboration with peers.	The importance of and role of **group work**	An optional component of a lesson; when employed, it should follow guided practice on problem solving, focus primarily on verifying that the procedures that have just been demonstrated work, and provide additional practice opportunities.
Dictated by both disciplinary and developmental (i.e., building new knowledge from prior knowledge) progressions.	The **sequencing of topics**	Dictated primarily by a disciplinary progression (i.e., prerequisites determined by the structure of mathematics).
Two main types of tasks are important: (1) tasks that initiate students to new ideas and deepen their understanding of concepts (and to which they do not have an immediate solution), and (2) tasks that help them become more competent with what they already know (with type 2 generally not preceding type 1 and both engaging students in reasoning).	The nature and ordering of **instructional tasks**	Students should be given opportunities to use and build on what they have just seen the teacher demonstrate by practicing similar problems, sequenced by difficulty. Tasks afford opportunities to develop the ability to adapt a procedure to fit a novel situation as well as to discriminate between classes of problems (the more varied practice students do, the more adaptability they will develop).
Students should be given time to wrestle with tasks that involve big ideas, without teachers interfering to correct their work. After this, feedback can come in small-group or whole-class settings; the purpose is not merely correcting misconceptions, but advancing students' growing intellectual authority about how to judge the correctness of one's own and others' reasoning.	The nature, timing, source, and purpose of **feedback**	Students should receive immediate feedback from the teacher regarding how their strategies need to be corrected (rather than emphasizing that mistakes have been made). In addition to one-to-one feedback, when multiple students have a particular misconception, teachers should bring the issue to the entire class's attention in order to correct the misconception for all.

(Continued)

(Continued)

Dialogic Instruction	Distinction	Direct Instruction
Students' learning pathways are emergent. Students should make, refine, and explore conjectures on the basis of evidence and use a variety of reasoning and proof techniques to confirm or disprove those conjectures (CCSS-M-SMP 3), asking questions that drive instruction and lead to new investigations.	The emphasis on **creativity**	Students' learning pathways are predetermined and carefully designed for. To "make conjectures and build a logical progression of statements to explore the truth of their conjectures" (CCSS-M-SMP 3) is limited to trying solution strategies for solving a problem posed to them.
Students' thinking and activity are consistent sources of ideas of which to make deliberate use: by flexibly following students' reasoning, the teacher can build on their initial thinking to move toward important ideas of the discipline.	The purpose of **diagnosing student thinking**	Through efficient instructional design and close monitoring (or interviewing), the teacher should diagnose the cause of errors (often a missing prerequisite skill) and intervene on exactly the component of the strategy that likely caused the error.
Students participate in the defining process, with the teacher ensuring that definitions are mathematically sound and formalized at the appropriate time for students' current understanding.	The introduction and role of **definitions**	At the outset of learning a new topic, students should be provided an accurate definition of relevant concepts.
Representations are used not just for illustrating mathematical ideas, but also for thinking with. Representations are created in the moment to support/afford shared attention to specific pieces of the problem space and how they interconnect.	The nature and role of **representations**	Representations are used to illustrate mathematical ideas (e.g., introducing an area model for multi-digit multiplication after teaching the algorithm), not to think with or to anchor problem-solving conversations.

Source: Munter, Stein, and Smith (2015b). Used with permission.

Figure 1.2

Precision teaching is about knowing *what* strategies to implement *when* for maximum impact.

Many readers of *Visible Learning* (Hattie, 2009) attend to the details about effect sizes and measuring one's impact (important, to be sure), but fewer may notice that this body of research points to *when* it works as well as *what* works. Knowing *what* strategies to implement *when* for maximum impact is what we think of as **precision teaching**.

The Balance of Surface, Deep, and Transfer Learning

As mentioned in the preface, it's useful when planning for precision teaching to think of the nature of learning in the categories of surface, deep, and transfer. It is a framing device for making decisions about *how* and *when* you engage in certain tasks, questioning techniques, and teaching strategies. The most powerful model for understanding these three categories is the SOLO (structure of observed learning outcomes) model developed by Biggs and Collis (1982). In this model, there are four levels, termed "unistructural," "multi-structural," "relational," and "extended abstract." Simply put, this means "an idea" and "many ideas" (which together are surface), and "relating ideas" and "extending ideas" (which together signify deep). Transfer is when students take their learning and use it in new situations. Figure 1.3 shows two examples of the SOLO model for mathematics.

Teaching Takeaway

The issue should not be direct versus dialogic but rather the right approach at the right time to ensure learning.

One key to effective teaching is to design clear learning intentions and success criteria (which we'll discuss in Chapter 2), which include a combination of surface, deep, and transfer learning, with the exact combination depending on the decision of the teacher, based on how the lesson fits into the curriculum, how long- or short-term the learning intentions are, and the complexity of the desired learning. Also, we recognize that learning is not an event, it is a process. It would be convenient to say that surface, deep, and transfer learning always occur in that order, or that surface learning should happen at the beginning of a unit and transfer at the end. In truth, these three kinds of learning spiral around one another across an ever-widening plane. Also, we want to be clear that because learning does not fall into a linear and repeating pattern—and is different for different students—we are in no way suggesting a specific order or scaffold of methods. In education, we spend a great deal of time debating particular methods of teaching and the pros and cons of certain strategies and their progression as applied to different content areas. The bottom line is that there are many phases to learning, and there is no one way or one set of understandings that unravels the processes of learning. Our attention is better placed on the effect we, as teachers, have on student learning. Sometimes that means we need multiple strategies, and, more often, many students need different teaching strategies from those that they have been getting (Hattie, 2012). Before further discussing how these phases of learning interweave, let's dive into what each one means in terms of mathematics.

THE SOLO MODEL APPLIED TO MATHEMATICS

Learning Intentions		Success Criteria
SOLO 1: Represent and solve problems involving addition and subtraction.		
Uni-/Multi-Structural	Know basic facts for addition and subtraction. Represent addition and subtraction using multiple models (manipulatives, number lines, bar diagrams, etc.).	I know my sums to twenty in both addition and subtraction. I can show my thinking using manipulatives and pictures.
Relational	Understand the meaning of addition or subtraction by modeling what is happening in a contextual situation (Carpenter, Fennema, Franke, Levi, & Empson, 2014). Recognize when either addition or subtraction is used to solve problems in different situations.	When I read a word problem, I can describe what is happening and use addition or subtraction to find a solution.
Extended Abstract	Use addition and subtraction to solve problems in a variety of situations.	I can use what I know about addition and subtraction contexts to figure out how to use addition and subtraction to solve problems beyond those I solve in class.
SOLO 2: Reason with shapes and their attributes.		
Uni-/Multi-Structural	Know the definitions and key attributes for shapes.	I can identify and name the attributes of shapes.
Relational	Recognize relationships among shapes.	I can explain how two shapes are related to each other.
Extended Abstract	Classify two-dimensional shapes based on properties.	I can create a diagram to show how different quadrilaterals are related to each other.

Source: Adapted from Biggs and Collis (1982).

This figure and a blank template are available for download at **http://resources.corwin.com/VL-mathematics**

Figure 1.3

Surface Learning

In mathematics, we can think of **surface learning** as having two parts. First, it is initial learning of concepts and skills. When content is new, all of us have a limited understanding. That doesn't mean we're not working on complex problems; it's just that the depth of thinking isn't there yet. Whether a student is exposed to a new idea or information through an initial exploration or some form of structured teacher-led instruction (or perhaps a combination of the two), it is the introductory level of learning—the initiation to, and early understanding of, new ideas that begins with developing conceptual understanding—and at the right time, the explicit introduction of the labels and procedures that help give the concepts some structure. Let us be clear: surface learning is not shallow learning. It is not about rote skills and meaningless algorithms. It is not prioritizing "superficial" learning or low-level skills over higher order skills. It should not be mistaken for engaging in procedures that have no grounding in conceptual understanding. Second, surface learning of concepts and skills goes beyond just an introductory point; students need the time and space to begin to consolidate their new learning. It is through this early consolidation that they can begin to retrieve information efficiently, so that they make room for more complex problem solving. For example, counting is an early skill, and one that necessarily relies initially on memorization and rehearsal. Very young children learn how to recite numbers in the correct order, and in the same developmental space are also learning the one-to-one correspondence needed to count objects. In formal algebra, surface learning may focus on notation and conventions. While the operations students are using are familiar, the notation is different. Multiplication between a coefficient and a variable is noted as $3x$, which means 3 times x. Throughout schooling, there are introductions to new skills, concepts, and procedures that, over time, should become increasingly easier for the learner to retrieve.

Surface learning is the initiation to new ideas. It begins with development of conceptual understanding, and then, at the right time, labels and procedures are explicitly introduced to give structure to concepts.

Surface learning is not shallow learning. It is not about rote skills and meaningless algorithms.

Importantly, through developing surface learning, students can take action to develop initial conceptual understanding, build mathematical habits of mind, hone their strategic thinking, and begin to develop fluency in skills. For example, surface learning strategies can be used to help students begin developing their metacognitive skills (thinking about their thinking). Alternatively, surface learning strategies can be used to provide students with labels (vocabulary) for the concepts they have discovered or explored. In addition, surface learning strategies can be used to address students' misconceptions and errors.

One challenge with surface learning is that there is often an overreliance on it, and we must think of the goal of mathematics instruction as being much more than surface learning. When learning stalls at the surface level, students do not have opportunities to connect conceptual understandings about one topic to other topics, and then to apply their understandings to more complex or real-world situations. That is, after all, one of the goals of learning and doing mathematics. Surface learning gives students the toolbox they need to build something. In mathematics, this toolbox includes a variety of representations (e.g., knowing about various manipulatives and visuals like number lines or bar diagrams) and problem-solving strategies (e.g., how to create an organized list or work with a simpler case), as well as mastering the notation and conventions of mathematics. But a true craftsman has not only a repertoire of tools, but also the knowledge of which tools are best suited for the task at hand. Making those decisions is where **deep learning** comes to the forefront, and, as teachers, we should always focus on moving students forward from surface to deep learning.

Deep learning is about consolidating understanding of mathematical concepts and procedures and making connections among ideas.

Deep Learning

The deep phase of learning provides students with opportunities to consolidate their understanding of mathematical concepts and procedures and make deeper connections among ideas. Often, this is accomplished when students work collaboratively with their peers, use academic language, and interact in richer ways with ideas and information.

Mrs. Graham started the school year for her fourth graders working with factors and multiples, connecting this work to previous third-grade experiences with arrays as models for multiplication, and extending these ideas to understanding prime and composite numbers. Students started by building and describing rectangular arrays for numbers from 1 to 50 (some students continued on to 100) and then discussed their answers to a variety of questions that developed the idea of prime and composite numbers. Class discussion incorporated mathematical vocabulary so it became a natural part of the student conversations (surface learning). The next day, students played a game called Factor Game (http://www.tc.pbs.org/teachers/mathline/lessonplans/pdf/msmp/factor.pdf) in which an understanding of primes and composites was crucial to developing strategies to win (deep learning is now occurring). However, the story doesn't end there. In March, students were beginning to study area and perimeter of rectangles. Following an initial exploration, several students approached Mrs. Graham to comment, "This is just like what we did last September when we were building arrays and finding primes and composites!" Talk about making connections!

As you can see, students move to deep learning when they plan, investigate, and elaborate on their conceptual understandings, and then begin to make generalizations. This is not about rote learning of rules or procedures. It is about students taking the surface knowledge (which includes conceptual understanding) and, through the intentional instruction designed by the teacher, seeing how their conceptual understanding links to more efficient and flexible ways of thinking about the concept. In Mrs. Graham's class, students began by developing surface knowledge of factors and multiples using concrete models and connected that to primes and composites. Mrs. Graham's use of the Factor Game provided students a way to apply their surface knowledge to developing strategies to win a game . . . deep knowledge. A teacher who nurtures strategic thinking and action throughout the year will nurture students who know when to use surface knowledge and when deep knowledge is needed.

Students move to deep learning when they plan, investigate, and elaborate on their conceptual understandings, and then begin to make generalizations.

We need to balance our expectations with our reality. This means more explicit alignment between what teachers claim success looks like, how the tasks students are assigned align with these claims about success, and how success is measured by end-of-course assessments or assignments. It is not a matter of all surface or all deep. It is a matter of being clear about when surface and when deep is truly required.

Consider this example from algebra. A deep learning aspect of algebra comes when students explore functions—in particular, the meaning of the slope of a line. Surface knowledge focuses on understanding the term mx in the slope-intercept ($y = mx + b$) form to mean m copies of the variable x. Deep learning requires students to understand and show that this term represented visually is the steepness or flatness of the slope of a line and the rate of change of the variables. Such learning might come from working collaboratively to explore a group of functions represented in multiple ways (equations, tables of values, and graphs) and make inferences about the slope in each representation. At this point, students are connecting their conceptual knowledge of ratio to their surface knowledge of algebraic notation and the process of graphing. This is deep learning in action.

Transfer Learning

The ultimate goal, and one that is hard to realize, is transfer. Learning demands that students be able to apply—or transfer—their knowledge, skills, and strategies to new tasks and new situations. That transfer is so difficult to attain is one of our closely kept secrets—so often we pronounce that students can transfer, but the processes of teaching them this skill are too often not discussed, and we'll visit that in Chapter 6.

Transfer is the phase of learning in which students take the reins of their own learning and are able to apply their thinking to new contexts and situations.

Transfer is both a goal of learning and also a mechanism for propelling learning. Transfer as a goal means that teachers want students to begin to take the reins of their own learning, think metacognitively, and apply what they know to a variety of real-world contexts. When students reach this level, learning has been accomplished.

Nancy once heard a mathematics teacher say that transfer is what happens when students do math without someone telling them to do math. It's when they reach into their toolbox and decide what tools to employ to solve new and complex problems on their own.

For example, transfer learning happens when students look at data from a science or engineering task that requires them to make sense of a linear function and its slope. They will use their surface knowledge of notation and convention, along with their deep understanding of slope as a ratio, to solve a challenge around designing an electrical circuit using materials with a variety of properties. Ohm's law ($V = iR$, where V represents voltage, i represents the current, and R represents resistance) is the linear function that relates the relevant aspects of the circuit, and students will use their mathematics knowledge in finding their solution.

One of the concerns is that students (often those who struggle) attempt to transfer *without* detecting similarities and differences between concepts and situations, and the transfer does not work (and they see this as evidence that they are dumb). Memorizing facts, passing tests, and moving on to the next grade level or course is not the true purpose of school, although sadly, many students think it is. School is a time to apprentice students into the act of becoming their own teachers. We want them to be self-directed, have the dispositions needed to formulate their own questions, and possess the tools to pursue them. In other words, as students' learning becomes visible to them, we want it to become the catalyst for continued learning, whether the teacher is present or not. However, we don't leave these things to chance. Close association between a previously learned task and a novel situation is necessary for promoting transfer of learning. Therefore, we teach with intention, making sure that students acquire and consolidate the needed skills, processes, and metacognitive awareness that make self-directed learning possible.

One of the struggles in teaching mathematics is to determine how much to tell students versus how to support students as they engage in productive struggle on their own, and when to know which is the right step to take. Let's take a look at helping elementary-age children build a toolbox of problem-solving strategies. Linda once attended a workshop for teachers that opened a whole new world of problem-solving strategies to use when solving nonroutine or open-ended problems. She was excited to take these problems back to her students and give them the opportunities to solve rich problems that involved some higher order thinking—that is, solving problems that involve much more than simple calculations. After some careful planning, she started a Monday class with her fifth graders by presenting the following problem.

> Mrs. Thompson, the school cook, is making pancakes for the special fifth-grade breakfast. She needs 49 pounds of flour. She can buy flour in 3-pound bags and 5-pound bags. She only uses full bags of flour. How can she get the exact amount of flour she needs?

Having never solved this type of problem before, the students rebelled. Choruses of "I don't know what they want me to do!" rang out across the classroom. "But they said in this workshop that kids could do this!" Linda thought.

Refusing to give in to the students' lament that the work was too hard, Linda decided that she needed to go about this differently. She resolved to spend each Monday introducing a specific strategy, presenting a problem to employ that strategy for students to solve together, and discuss their thinking. This was followed by an independent "problem of the week" for students to solve. After introducing all of the strategies (surface learning) and following up with independent applications of those strategies for students (deep learning), students continued to work independently or in small groups to solve a variety of open-ended problems on their own using strategies of their choice (transfer learning). Later that year, a group of girls approached Linda asking why she had saved all of the easy problems for the end of the year. That's transfer!

It's important to note that within the context of a year, a unit, or even a single lesson, there can be evidence of all three types of learning, and

THE RELATIONSHIP BETWEEN SURFACE, DEEP, AND TRANSFER LEARNING IN MATHEMATICS

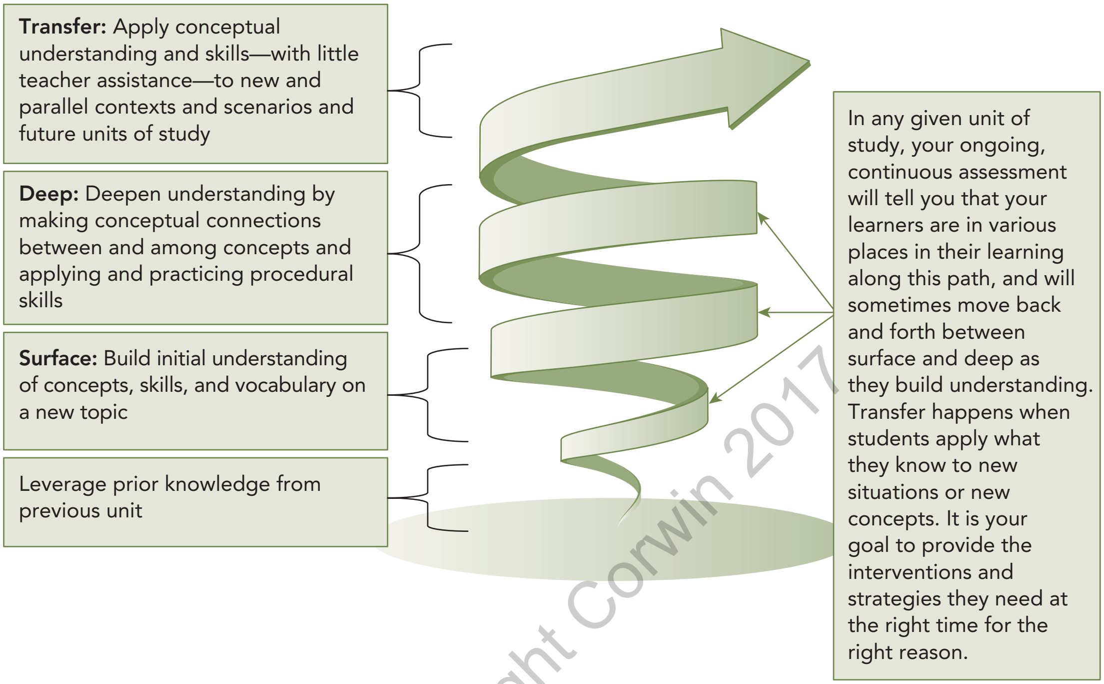

Source: Spiral image copyright © iStock/EssentialsCollection/83158933.

Figure 1.4

that students can sometimes move among various kinds of learning depending on where they are as individual learners. Figure 1.4 describes the relationship between surface, deep, and transfer learning.

Surface, Deep, and Transfer Learning Working in Concert

As mentioned before, when it comes to the surface, deep, and transfer phases of learning, knowing *what* strategies to implement *when* for maximum impact on learning is key. How, then, should we define learning, since learning is our goal? John defines it as

> the process of developing sufficient surface knowledge to then move to deeper understanding such that one can appropriately transfer this learning to new tasks and situations.

Video 1.2
Balancing Surface, Deep, and Transfer Learning

http://resources.corwin.com/VL-mathematics

Learning has to start with fundamental conceptual understanding, skills, and vocabulary. You have to know *something* before you can do something with it. Then, with appropriate instruction about how to relate and extend ideas, surface learning transforms into deep learning. Deep learning is an important foundation for students to then apply what they've learned in new and novel situations, which happens at the transfer phase. And tying all of this together is clarity about learning outcomes and success criteria, on the part of both teachers and students. If students know where they are going and how they'll know when they get there, they are better able to set their own expectations, self-monitor, and predict or self-report their own achievement. All of these phases can be present within the body of a single lesson or multiday or multiweek unit, as well as extend across the course of a school year.

EFFECT SIZE FOR STUDENT SELF-MONITORING = **0.45**

EFFECT SIZE FOR SELF-REPORTED GRADES/STUDENT EXPECTATIONS = **1.44**

Conclusion

Teachers have choices. As a teacher, you can unintentionally use instructional routines and procedures that don't work, or don't work for the intended purpose. Or you can choose to focus on *learning*, embrace the evidence, update your classrooms, and impact student learning in wildly positive ways. You can consider the nature of the phases of surface, deep, and transfer learning and concentrate on more precisely and strategically organizing your lessons and orchestrating your classrooms by harnessing the power of activities that are in the zone of desired effects—above a 0.40 hinge point. Understanding the phases of learning by examining the evidence will help you to make instructional choices that positively impact student learning in your classroom.

The following chapters are meant to help you design lessons and appropriately employ instructional moves that honor students' need to develop their surface understanding of a topic, help you extend the depth of their mathematical learning, and help them transfer their learning to new tasks and projects. The journey starts when you turn the page and delve into the topic that matters most—establishing learning intentions and success criteria. Let's start the journey of making mathematics learning visible for students.

Reflection and Discussion Questions

1. Think about the instructional strategies you use most often. Which do you believe are most effective? What evidence do you have for their impact? Save these notes so you can see how the evidence in this book supports or challenges your thinking about effective practices.
2. Identify one important mathematics topic that you teach. Think about your goals for this topic in terms of the SOLO model discussed in this chapter. Do your learning intentions and success criteria lean more toward surface (uni- and multi-structural) or deep (relational and extended abstract)? Are they balanced across the two?
3. A key element of transfer learning is thinking about opportunities for students to move their learning from math class, to use their knowledge to solve their own problems. Think about the important mathematical ideas you teach. For each one, begin to list situations that might encourage transfer of learning. These might be applications in another subject area or situations in real life where the mathematics is important.