

Chapter 3

STRATEGY 1: CLARIFYING, SHARING, AND UNDERSTANDING LEARNING INTENTIONS AND SUCCESS CRITERIA

OVERVIEW

When we ask teachers about their learning intentions for a particular lesson, it is common for them to reply by saying, "I'm going to have the students do this." In other words, teachers respond in terms of the activities in which they plan to engage their students. This is, of course, entirely understandable. The only way that teachers get students to learn is by engaging them in activities. However, if the teachers select activities without a clear view of the learning that is intended, it is far less likely that the students will learn what they need to learn. By engaging in an activity, students usually learn something, but not all outcomes of learning are equally valuable. By being clear about what it is we want students to learn—in Stephen Covey's words, by "starting with the end in mind" (Covey, 1989)—it is more likely that our students will learn what we need them to learn.

WHY USE THIS STRATEGY

Before we can find out what our learners are learning, before we can give feedback, before we can engage our learners as resources for one another and as owners of their own learning, we have to be clear about where we are going.

GETTING TO UNDERSTAND THE STRATEGY

Ralph Tyler's *Basic Principles of Curriculum and Instruction* was first published in 1949. It is a testament to its influence that more than six decades later, it is still in print. Tyler identifies "four fundamental questions which must be answered in developing any curriculum and plan of instruction":

1. What educational purposes should the school seek to attain?
2. What educational experiences can be provided that are likely to attain these purposes?

3. How can these educational experiences be effectively organized?
4. How can we determine whether these purposes are being attained?

(Tyler, 1949, p. 1)

In many ways, such a formulation, and in particular the idea of starting with the purposes of education, is obvious. As Yogi Berra, the former New York Yankees baseball player, once said, "If you don't know where you are going, you might wind up somewhere else." The fundamental questions are also empowering for students. As one tenth-grade student said,

Something like English, for every topic we start we get success criteria with different grades on it and we have to highlight the key bits and the things we can include in our essay or coursework to get that grade. (William & Leahy, 2014)

It is therefore hardly surprising that designing educational experiences backward from the intended outcomes has been extremely influential (see, for example, Wiggins & McTighe, 2000). However, while it seems straightforward to specify aims for our students' learning—so much so that many teachers and schools begin their implementation of formative assessment by working on learning intentions—it is important to understand that developing good learning intentions is not easy, and our work with teachers suggests that this first strategy may actually be the most difficult of the five formative assessment strategies to implement well.

In particular, the apparent straightforwardness of specifying learning intentions and success criteria has led to simplistic and rigid implementation that undermines their effective use. Many frameworks for the evaluation of teaching provide rubrics of teacher behaviors, and it is common to find that to attain the highest level of performance, teachers must, among other things, begin each instructional sequence with a clear statement of the intended learning outcomes. As a result, many districts mandate that every lesson must begin with a learning outcome. The consequences are a predictable "culture of compliance"; teachers post learning outcomes on a board, students copy these into their notebooks, and the lesson then often proceeds without further reference to the learning outcomes, unless, of course, an administrator happens to stop by to observe the lesson. One student, who was asked to describe a learning intention, said, "It's what we copy into our books at the beginning of the lesson, while we're talking to our friends."

Now, of course, it is, in general, a good idea that students know where they are going—we are both embarrassed about how long we spent teaching without telling our students where we were going. After all, it is really rather obvious that if our students know where we are headed, we are more likely to arrive at the correct destination. But there are at least three reasons why starting every single lesson with a learning outcome is likely to be a bad idea.

First, so teaching single, v to the s student track. H it is enti In such : equally : Sometin direction to the ta

This illu things n because dents in rather th to experi they are tainty wi in the pa it is appi

Second, a teache ing that the base calculate problem problem

Third, a recipe fc son with not all cl any learn classroom happy w by any te ate to lea of conte

First, sometimes you don't always know exactly where the lesson is going. If you are teaching high school chemistry students how to balance chemical equations, there is a single, well-defined goal for the whole class. You are trying to get every single student to the same point, so if the student's solution differs from the standard solution, the student is probably heading off in the wrong direction and needs to be brought back on track. However, if you are asking the students to respond to a poem by Emily Dickinson, it is entirely appropriate that some students head off in different directions from others. In such a situation, there isn't a single goal for the whole class but a whole "horizon" of equally acceptable goals. Indeed, it is okay if different students learn different things. Sometimes, of course, it seems as if a student is headed in a completely inappropriate direction, but this can often be the result of the student finding a more creative approach to the task than those envisaged by the teacher.

This illustrates a more general point, which is that sometimes it is appropriate to do things not because they are guaranteed to result in specific learning outcomes but because they are important experiences for students. The national curriculum for students in Scotland explicitly recognizes this by talking about "experiences and outcomes" rather than relying solely on "outcomes." There are some things that we want students to experience not because they necessarily lead to specific learning outcomes but because they are valuable things to do together. It may not be possible to predict with any certainty what the learning outcomes will be, but if your experience is that a particular task in the past resulted in valuable and significant learning for your students, our view is that it is appropriate to use the task.

Second, sometimes telling students where you are going spoils the journey. For example, a teacher might set a class a problem in which the key to solving the problem is realizing that triangles with the same base and same height have the same area even though the base and height of the triangles are not given in the problem, so it is not possible to calculate the area (see Wiliam, 2011a, pp. 56–57). If you tell the students they are doing problems that involve areas of triangles, what could be an interesting and challenging problem is reduced to a rather more routine exercise.

Third, and perhaps most important, starting every lesson with a learning outcome is a recipe for uninspired and uninspiring teaching. Those who mandate starting every lesson with a learning outcome forget one important, and universal, fact about schools: not all children are motivated. If all students were indeed motivated, we could announce any learning outcome, and our students would be delighted. They would come into our classrooms and tell us that they were desperate to learn something, and they would be happy with any learning outcome we choose. Unfortunately, this is not the reality faced by any teachers we know. Our students do not always come into our classrooms desperate to learn. They need to be engaged and enthused, and, again unfortunately, statements of content standards excite very few students. They get excited by ideas, arguments, dis-

cussions, and debates. And this is the real skill of good teachers: to find a way to engage students in rigorous content matter by making it come alive.

For example, a teacher begins a ninth-grade science lesson by asking the class, "Why does it take longer to cook potatoes at 400°F in an oven than it does at 212°F in a saucepan of boiling water?" Most of the students know enough about cooking potatoes to realize that the premise of the question is correct—it does indeed take longer to cook potatoes by broiling than by boiling—but none of the students notice the potential paradox. The important point here is that the teacher did have a clear learning intention in mind at the outset, and in fact shared it with the students after the discussion, but the teacher judged that the best way to start the lesson was not with a learning intention but with an engaging question.

So, in general, teachers do need to plan their classroom activities backward from the learning goals, but these cannot always be specified in advance. And it is not always advisable to tell the students at the outset either, because it spoils the journey and there are better ways to engage and motivate students. In the remainder of this chapter, we look at a number of techniques that you can use to develop, clarify, and share learning intentions with your students, and ensure that they understand where they are going.

TECHNIQUES

Keep the Context of Learning Out of the Learning Intention

Before we discuss learning intentions and success criteria, it seems sensible to discuss three assumptions that we believe should drive what teachers do. The first assumption is that your main aim as a teacher is to change your students. Most of the time, we are trying to get our students to be able to do things that they could not do before, although there are also times when we are trying to get them not to do something that they might have done before, like get pregnant or take drugs. But the teacher's main aim is always to change his students. The second assumption is that this change is directed and planned. We are not interested in random changes in what students can do, but we are interested in particular changes in our students that we believe have value. The third assumption is that these changes are more likely to come about if we have a clear plan for what we want our students to learn.

Ralph Tyler's model for the development of curriculum and instruction outlined earlier is sometimes called the "objectives" model for the obvious reason that the model assumes objectives should drive curriculum and instruction; and in some ways, the idea of learning intentions and success criteria is just a modern take on the same idea. Indeed, we could just carry on using the terms *aims* and *objectives*. However, we believe that a change in vocabulary is justified, because people often use *objectives* to mean "behavioral objectives," and one major criticism of Tyler's objectives model is there are many things we want our students to learn that are not easily described in behavioral terms.

In this chapter, we use *learning outcomes* to describe what it is that students can do as a result of engaging in educational activities (literally, the outcomes of the learning process). Of course, as every teacher knows, what students learn as a result of engaging with a particular sequence of educational activities is not easy to predict. Sometimes students learn what we want them to learn, and sometimes they learn something else. We use *learning intentions* to describe the things that we want our students to learn and *success criteria* to describe the criteria that we use to judge whether the learning activities in which we engaged our students were successful. One fourth grader describes being given success criteria thus: "It's like knowing the teacher's secret" (Spendlove, 2009, p. 18).

Elementary school teachers often communicate learning intentions and success criteria to their students as WALT (We are learning to . . .) and WILF (What I'm looking for). Others have used a puppet called Timmy the teaching tiger or Ollie (OLI or "our learning intention") to convey the learning intentions. Such acronyms can be useful in engaging students in the process, but as Shirley Clarke (2001) points out, it is important to ensure that the focus is on what these terms mean, and how they relate to the work, otherwise they just become "gimmicks."

In a later work, Shirley Clarke (2003) shows that devising learning intentions and success criteria is actually very difficult. Or, to be more precise, devising bad learning intentions and success criteria is actually very easy, but devising good ones is a significant challenge—one, like developing good questions (see next chapter), that is so challenging that no teacher ever truly masters it. However, there are some very simple and straightforward things that teachers can do to substantially improve their learning intentions.

Perhaps the most straightforward is to address the fact that, as Shirley Clarke points out, learning intentions are far less useful if they specifically include the context of the learning. She describes an elementary school teacher who tells her students that the learning intention is "To understand the effects of banana production on the banana producers" (Clarke, 2003, p. 23). The reason for this particular activity is almost certainly not to ensure that students come to understand the specific difficulties faced by banana producers. This particular task has more likely been selected because the teacher is using this specific context to illustrate the more general impact of production on producers, and it would be better if the learning intention reflected this. The learning intention would therefore become "to understand the impact of production on producers" with banana production being the context of the learning.

This example illustrates a rather important principle at the heart of teaching, which is that when we teach our students how to do something, we are rarely interested in their ability to do the things we have taught them. That sounds paradoxical but it is, in our experience, true. If we teach our students to add two numbers, such as 278 and 345, we are no longer interested in those two numbers; our students can almost

certainly add those two numbers because we have shown them how to do it. We are interested in whether they can use what they learned to add two different numbers. If we teach our students how to punctuate a sentence, we are no longer interested in that particular sentence; we are interested in whether our students can apply what they have learned to a different sentence, or, even better, that they can punctuate when writing a story as opposed to doing a punctuation exercise. Another way of stating this principle is that generalizability is at the heart of good teaching. If the only things our students can do are what we have taught them to do, they can't do very much at all, because those exact circumstances are unlikely ever to arise again. The only useful learning is that which the student can apply beyond the context of the learning.

This might seem straightforward to accomplish, but it turns out our brains aren't very good at disentangling the context of the learning from the learning itself. It has long been known that adults rarely use what they have been taught in school to solve problems in the "real world." For example, Nunes, Carraher, and Schliemann (1993) find that Brazilian market traders, fishermen, building site foremen, carpenters, and farmers are able to solve a variety of mathematical problems more effectively than students with more years of schooling, but they do not use the formal school-taught methods.

Another example comes from the work of Jean Lave (1988), who describes a group of adults in a Weight Watchers class learning to prepare meals:

In this case they are to fix a serving of cottage cheese, supposing the amount laid out for the meal is three-quarters of the two-thirds cup the program allowed. The problem solver in this example begins the task muttering that he took a calculus course in college. . . . After a pause he suddenly announces that he "got it!" From then on he appears certain he is correct, even before carrying out the procedure. He fills a measuring cup two-thirds full of cottage cheese, dumps it out on the cutting board, pats it into a circle, marks a cross on it, scoops away one quadrant, and serves the rest.

Therefore, "take three-quarters of two-thirds of a cup of cottage cheese" is not just the problem statement but also the solution to the problem and the procedure for solving it. The setting is part of the calculating process, and the solution is simply the problem statement, enacted with the setting. At no time does the Weight Watcher check his procedure against a paper and pencil algorithm, which would produce $\frac{3}{4}$ cup $\times \frac{2}{3}$ cup = $\frac{1}{2}$ cup. Instead, the coincidence of the problem, setting, and enactment are the means by which checking took place.

As these and many other examples show, our brains do not appear to be very good at transferring knowledge from one context to another. Many people imagine that our brains are somehow like computers that either have particular knowledge or don't, but it turns out our brains don't work like this.

For example, the results of one experiment shows that scuba divers who were taught things at the surface remembered 55 percent more of what they had been taught when they were tested at the surface than when they were tested underwater. But when they had been taught things underwater, they remembered 35 percent more when they were tested underwater than when they were tested at the surface (Godden & Baddeley, 1975).

In another experiment, a group of adults were each given three gin and tonics and then taught a nineteen-item set of instructions for a route on a map. The following day, half of the students were tested on their recall of the instructions sober, and half were tested after three more gin and tonics. The students tested while intoxicated recalled almost 40 percent more of the details than those tested while sober (Lowe, 1980).

Finally, students who were tested in the room in which they were taught something got better scores than if they were tested in a different room (S. M. Smith, Glenberg, & Bjork, 1978).

We see this failure of our brains to transfer from one context to another played out in middle and high schools every day. The science teacher asks the math teacher, "Why can't these students do equations and graphs?" and the math teacher lamely replies, "Well, they can do equations and graphs in the math classroom."

This research has led some to conclude that teaching abstract ideas in the hope that students are able to transfer learning from one context to another is impossible (Lave, 1988, p. 19). However, as Anderson, Reder, and Simon (1996) point out,

Contrary to Lave's opinion, a large body of empirical research on transfer in psychology . . . demonstrates that there can be either large amounts of transfer, a modest amount of transfer, no transfer at all, or even negative transfer. How much there is and whether transfer is positive depends in reliable ways on the experimental situation and the relation of the material originally learned to the transfer material. (p. 7)

In particular, we know that transfer from one context to another is much more effective when students encounter the same idea in more than one context (Bjork & Richardson-Klavehn, 1989), and this is why the idea of keeping the context of the learning separate from the learning intention is so important. Returning to the example of banana production discussed earlier, the learning intention is "to understand the impact of production on producers"; the context of the learning is banana production; and potential success criteria would describe the extent to which students could apply what they had learned in the context of banana production to a different context, such as sugar production.

Tips

Get Students to Apply What They Have Learned in Other Contexts

Designing learning activities backward from the learning intentions and success criteria that they are intended to engender feels odd at first, but with practice, it becomes natural. To get you started, look at the following list of potential learning intentions. Decide whether each one is really a learning intention or just an activity, and for those you think are learning intentions, decide whether they are context free. Then, for those that relate to your own teaching responsibilities, spend some time thinking about how they might be improved.

We are learning to:

- know ways of controlling floods in Bangladesh
- express ideas, words, thoughts, and feelings through creative work in dance
- recognize triangular numbers
- paint a landscape
- make a plastic key fob
- create a cartoon animal using descriptive phrases and sentences
- complete all ten questions on pages 23 to 25
- use science equipment and kitchen compounds to carry out purposeful experiments
- identify the causes of plantation slavery
- count up to ten in French

As an example, for the first of these, the important observation is that it is unlikely that the teacher chose the topic primarily for the students to understand how to control floods in Bangladesh. It is far more likely the teacher intends that the students understand the factors that affect flooding in general, with Bangladesh as a special case. So the learning intention becomes "Know ways of controlling floods," with Bangladesh being the context of the learning. Possible success criteria might then relate to being able to list different causes of flooding, knowing how each of these might be mitigated, and then applying the knowledge in a different context.

As soon as you begin to think of what you are teaching as special cases of something else, you will begin thinking about the more general category of which this particular example is a special case, and then you can start generating other examples in the same category. Table 3.1 shows examples of some confused learning intentions based on the work of Shirley Clarke (2005) but extended to show how removing the context of the learning from the learning itself makes it easier to identify appropriate success criteria.

Table 3.1: Confused and Clarified Learning Intentions and Success Criteria

Confused learning intention	Clarified learning intention	Context of learning	Potential new context
To be able to write instructions on how to change a bicycle tire	To be able to write instructions	Changing a bicycle tire	Writing instructions on how to make a sandwich
To be able to construct arguments for or against assisted suicide	To be able to construct arguments for or against emotionally charged propositions	Assisted suicide	Constructing arguments for or against abortion
To create a written description of a friend	To be able to create an effective characterization	Describing a friend	Creating an effective characterization of a relative

Mix It Up


The ability of students to apply things in other contexts is increased if their learning activities are varied. Students generally prefer to do one thing at a time, and get frustrated if the activities involve “chopping and changing” between different activities, but they tend to learn more. This was rather powerfully illustrated by a twelve-week experiment that involved sixty-four eight-year-old students who practiced throwing beanbags into buckets in gym class (Kerr & Booth, 1978). Half of the students practiced tossing the beanbags into a bucket three feet away, while for the other half the practice was varied by having the bucket either two feet away or four feet away. At the end of the experiment, all the students were tested on their ability to throw the beanbag into a bucket three feet away. The students who had practiced with buckets placed two and four feet away, and who therefore never practiced throwing a beanbag into a bucket three feet away, actually outscored the students who did nothing but practice throwing beanbags into a bucket three feet away. Of course, this is just a single study, but there is now substantial research evidence that despite their preference to work on one idea or context at a time, students learn more when they are required to vary their practice (for an authoritative yet readable summary, see P. C. Brown et al., 2014).

Cautions


Differentiate Success Criteria, Not Learning Intentions

Many districts have adopted “differentiated instruction” as a policy for their classrooms, although this often appears to be done with little clear idea about what the term actually means. Indeed, most definitions of differentiated instruction are rather vague, in effect saying little more than students are different, and we should recognize this in our teaching. Nevertheless, in many districts, teachers are encouraged to cater to differences in their students by having different learning intentions for different students in their class. This seems to us unfortunate for at least two reasons: first, it assumes that the teacher knows which students are capable of which work, therefore creating a danger of self-fulfilling prophecies, and, second, it is a lot of work for the teacher. Now, of course, different students will learn different things in a lesson, particularly one where teachers give students independence over their work, so it is unlikely that a single success criterion is appropriate for all students. However, we do think that all the students in the group should have reasonably similar goals. In other words, we think that differentiation should take place in terms of the success criteria rather than the learning intentions.


As a concrete example, consider a lesson on the area of a trapezoid. In a typical middle school classroom in the United States, the teacher explains the rule to students, performs one worked example for the whole class on the board, and then students practice the technique individually, in their seats, while the teacher walks up and down the aisles, providing additional support where needed. In Japan, it is more likely that the teacher shows one way of finding the area of a trapezoid to the class, and then invites students, working in pairs or groups, to see how many conceptually distinct ways they can generate to find the area of a trapezoid. For example, if the teacher introduces the “duplicate and rotate” method:



then students might find the “parallel cut and rotate” method:



There are exactly thirteen conceptually distinct ways of finding the area of a trapezoid. Many are simple to find, but others are much harder, either because they are not easy to spot, or because the algebraic manipulation required to show that the method still results in the standard formula of $\frac{1}{2}(a + b) \times h$ is much more demanding, such as the method that involves completing the triangle and then using properties of similar triangles to subtract the area of the smaller triangle from that of the larger triangle:



The important point about such a task is that the learning intention is the same for all members of the class—to find as many conceptually distinct ways of calculating the area of a trapezoid given the length of the two parallel sides and distance between them. The differentiation comes as a result of the way that the teacher implements the activity, supports some students, and pushes others to find specific methods, such as the example using similar triangles just described.

Enhancements

Process versus Product Success Criteria

Some authors argue that success criteria are more useful if they address the process by which students are to produce the work than if they merely describe the features of the final product. Clarke (2005) defines *process success criteria* as those which “summarise the key steps or ingredients the student needs in order to fulfil the learning objective—the main things to do, include or focus on. These give a framework for formative assessment to take place” (p. 29).

Process success criteria can be particularly helpful to students if they break the process that students are expected to follow into a number of smaller, more easily managed steps. This idea is nicely encapsulated by Liane-Paixao (2014), who wrote and recorded a song that summarizes the benefits of process success criteria from the student’s point of view:

Break It Down

You want me to follow you, steady
But I don’t think we’re really ready
You gimme no clue, came out of the blue
I’m feeling so lost
I cannot pack and go

So, where are you leading me to?
 That, you don't say . . .
 I gave up on guessing, and then you complain
 It feels like a game
 I zone out for the day
 (You gotta break it)

CHORUS:

Break it down, break it down
Break it down for me
Learning can be so easy
When every step is there
And I can see

Don't go 'round turning my rocks over
 Pinpointing all that's error
 Without showing how to learn from them
 And make it all better
 You wait until I fail to tell me
 That I really could
 I could've done much better
 It is water under the bridge
 It doesn't matter (You gotta break it)

(Chorus)

BRIDGE:

I gotta learn on my own, guide me
 Warn me before a sharp turn, mind me
 Life is a slippery slope that we can't cruise on blindly
 Gotta learn how to get to where I gotta go (yeah)

Scaffold my learning
 and coach me along the way
 Empower me so that one day
 I can stand on my own
Break it down, break it down . . .

The danger with process success criteria, of course, is that they can be too specific. If the process success criteria lay out, step-by-step, exactly what the student has to do to be successful, little learning is likely to take place. Moreover, students are less likely to generate insightful solutions.

For example
 A suitable le
 success crite
 do this is to
 common de
 of the denor

Example
 Example

A typical set
 identify the l
 all the fractio
 the fractions
 how to iden
 and others n
 which fractio
 in our view,
 problem. Th
 trajectory of
 ferent ways a

Similar cons:
 that leg brac
 mous hero o
 weaker stude
 strain all stu
 This does no
 they can be a
 to make sure
 they are doir
 haps even fre
 when we dis
 challenging t
 failure to pro
 just found 10

To Be Forr

The foregoi
 ing to the ide
 learning inter
 it is to help th

For example, a teacher is instructing a class how to identify the largest fraction from a set. A suitable learning intention is “to identify the largest fraction in a set,” and the product success criteria are differentiated to include examples of differing difficulty. One way to do this is to have a criterion relating to those cases where one of the denominators is a common denominator for all the fractions in the set and another criterion where none of the denominators is a common denominator of the set:

Example 1	$\frac{4}{5}$	$\frac{3}{4}$	$\frac{7}{10}$	$\frac{17}{20}$
Example 2	$\frac{4}{5}$	$\frac{3}{4}$	$\frac{7}{10}$	$\frac{5}{8}$

A typical set of process success criteria for such a task would remind the students to first identify the lowest common denominator for the four fractions being compared, convert all the fractions to equivalent fractions with a common denominator, and then compare the fractions. However, if students are not given process success criteria that tell them how to identify the largest fraction, some might convert the fractions to percentages, and others might realize that identifying the largest fraction is equivalent to identifying which fraction needs the smallest amount added to it to make into one whole, which are, in our view, eminently sensible, and perhaps even superior, approaches to solving the problem. The important point here is that while rubrics may identify the most common trajectory of improvement for most students, there will be students who improve in different ways and follow different tracks.

Similar considerations apply to planning frames and writing templates. In the same way that leg braces can be both a support and constraint (think, for example, of the eponymous hero of the movie *Forrest Gump*), process specifications and structures can help weaker students do better work than they might without those supports, but they constrain all students from finding imaginative and creative solutions to tasks they are set. This does not mean that we should not use process success criteria—as noted earlier, they can be an important element in supporting our students’ learning—but we do need to make sure our students understand that if they feel confident that they know what they are doing, they can throw the leg braces away. This means that occasionally, perhaps even frequently, our students will fall flat on their faces, but, as will become clear when we discuss eliciting evidence, failures are indications that the work is sufficiently challenging to make it worthwhile. As Thomas Edison once said, when asked about his failure to produce a workable filament for the electric light bulb, “I have not failed. I’ve just found 10,000 ways that won’t work.”

To Be Formative, Learning Intentions and Success Criteria Need to Be Generic

The foregoing discussion of process success criteria illustrates a more general point relating to the idea of generalizability in learning that we discussed earlier. The aim of any learning intention or success criterion is not to help the students complete the activity—it is to help them learn. Checklists that provide step-by-step instructions for students to

follow do not help students complete tasks on their own. This is because what psychologists call “executive control”—deciding what to do next in a task—is being done by the checklist, not the student. Too often, the result is students who can complete a task when directed but cannot work independently.

This point often gets lost when teachers are devising learning intentions and success criteria. If students are to learn something in doing one task that they can carry over to another task, the same learning intentions and success criteria must apply to more than a single task. Specific learning intentions and success criteria do help different raters agree on whether students have achieved a particular level of attainment, so they are useful for summative assessment. However, where the assessment is intended to function formatively, a degree of generality has to be built into the learning intentions and success criteria (Arter & McTighe, 2001). Of course, if the success criteria are too general, they become so meaningless that they provide little or no guidance for students. The art (and it is an art rather than a science) is to find an appropriate balancing point between these two extremes, and, of course, the location of that balancing point will be different for different students, different for different aspects of the subject matter, and different for students at different stages of expertise.

Be Aware of Strengths and Weaknesses of Analytic and Holistic Learning Intentions and Success Criteria

Similar considerations apply to the debate about whether learning intentions and success criteria should be analytic or holistic. It might be assumed that analytic approaches to describing quality, in which different aspects of quality are identified, might be superior for both summative and formative purposes. Specifying the dimensions of quality that are important increases the likelihood that different raters look at the work in the same way, therefore improving the usefulness of descriptions for summative assessment. Moreover, when the learning intentions and success criteria highlight the important dimensions of quality, students are more likely to pay attention to the aspects of the work that matter, and therefore improve the quality of their work.

These arguments certainly have merit, but it is also important to remember that analytic approaches to describing quality often miss out on the “glue” that holds a piece of work together. Sometimes a piece of work seems to check all the right boxes, but somehow the whole is less than the sum of the parts. And at other times, a piece of work can get all the pieces wrong but somehow cohere as a whole. As with so many aspects of formative assessment, there are no simple answers here, but we hope that the principles we lay out help you resolve—in ways that better meet the learning needs of your students—the various tensions we discussed.

Start With Samples of Work, Rather Than Rubrics, to Communicate Quality

The extraordinary increase in the use of scoring rubrics for all kinds of assessments in the United States in recent years is perfectly understandable. If, as this chapter argues, students do better if they know what they are meant to be doing, laying out the characteristics of different standards of work must surely be helpful. However, it is our view that the use of rubrics has, in many cases, spread to a point where the benefit may well be less than the harm.

Before we point out the shortcomings of rubrics, however, we want to make it clear we are not saying that rubrics cannot be useful in education—indeed, we believe that they have an important role to play—but we think the intelligent use of rubrics requires a deep understanding of their potential shortcomings.

First, rubrics are essentially unidimensional. While most rubrics identify multiple aspects of a piece of work, the improvement in each aspect is almost always along a single dimension. In other words, there is only one way in which that aspect of the work can be improved. This may be appropriate for some kinds of work, but wherever a measure of creativity is involved, a piece of work can be improved in a number of different ways, but only one of these is usually reflected in the rubric. Those using the rubric are therefore channeled into producing improvements only along the dimension identified in the rubric.

Second, rubrics are often more focused on evaluating student achievement than improving. Alfie Kohn suggests that there are three main reasons why we want to assess the quality of our students' efforts: to rank students, induce them to try harder, or provide feedback that helps them improve (Kohn, 2006, p. 14). Proponents of rubrics point out that rubrics can give students feedback to enable them to improve their work, which is certainly true, but the fact that rubrics almost always identify a number of different levels of quality in student responses suggests that the evaluative function is also vital. And it is important to remember that the intent behind the use of rubrics is irrelevant; what matters is what rubrics do to students.

Third, rubrics rarely have the same meanings for students that they do for teachers. The highest level of a rubric for a piece of persuasive writing might require that the essay make an effective argument, but if the student does not know what constitutes an effective argument, the rubric is of limited use.

While rubrics have a role to play, we believe they are best regarded as the *culmination* of a developmental process that begins with examination of samples of students' work. So, before your students do a laboratory report or write a ghost story, spend some time getting them to look at other students' attempts at similar tasks. Many teachers believe that it is wasteful to take time that students could use to generate their own work to look

at the work of others, but there are two immediate benefits of getting students to look at samples of student work: first, students are better at spotting mistakes in the work of others than they are in their own work. Assessing one's own work, as well as assessing the work of one's peers in the classroom, is emotionally charged, and the emotional resonances can often interfere with engaging in the demands of the task. However, assessing the work of anonymous others is emotionally neutral, so students are able to focus more effectively on the task. Second, when students notice mistakes in the work of others, they are less likely to make the same mistakes in their own work.

For the youngest students, we recommend starting with just two pieces of work where one is relatively strong and one is relatively weak. For example, the teacher might present a kindergarten class with two photographs of something children have built such as a sandcastle or something made of Lego® bricks and ask them which is better, and why. In general, students seem to find it easier to identify qualities of good work when two pieces of work are contrasted than when they are given a single piece of work, regardless of whether it is strong or weak.

Where the two pieces of work are relatively large—for example, watercolor paintings or drawings—the teacher can display them side by side to the class. For smaller pieces of work, a document camera (sometimes called a visualizer or visual presenter) is very useful. The teacher can then ask students whether one of the two is better than the other, and if they respond that one is better, the teacher can ask which one it is and why. Think-pair-share (see next chapter) is a particularly useful technique here for giving students time to think about their perceptions of quality and begin to develop some vocabulary to put this into words.

As students become more sophisticated in their appreciation of the differences in quality between different pieces of work, you can use samples of work that are much closer in quality, and introduce more gradations of quality by including more examples. The language that students use to describe their perceptions of the differences in quality between different pieces of work will then naturally lend itself to the construction of success criteria for the work. More importantly, the success criteria will be anchored in actual samples of work rather than vague and decontextualized statements that can be interpreted in a number of ways.

Success criteria can be simple checklists of features that need to be present for the work to be regarded as good, such as the following checklist for early writing, which can be accompanied by pictures or diagrams to cue less fluent readers:

- ☐ All my sentences start with a capital letters.
- ☐ Every sentence ends with a "stop sign" (".", "?", or "!").
- ☐ I used finger spacing between words.
- ☐ I used the word wall to spell my words correctly.

☐ My handwriting is neat.

☐ I read through my work to make sure it sounds right.

Before showing the work to the teacher, students can then be asked to check through the list. Some teachers ask students to use the checklist to complete a self-assessment by choosing an emoticon such as a happy or sad face. Asking students how happy they are with their work does have a role to play, even in pre-K (most adults are surprised how well even four- and five-year-olds are able to evaluate their own work). However, when the students are using a checklist, it probably makes more sense for them to make sure each of the required features is present. After all, we don't want students handing in writing that they know isn't good enough and they know how to improve!

Tips

Remember That Rubrics Are Really Nothing More Than Collections of Success Criteria

Heidi Andrade (2000) defines a rubric as follows:

An instructional rubric is usually a one- or two-page document that describes varying levels of quality, from excellent to poor, for a specific assignment. It is usually used with a relatively complex assignment, such as a long-term project, an essay, or a research paper. Its purposes are to give students informative feedback about their works in progress and to give detailed evaluations of their final products. (p. 13)

Before we go on, one feature of the definition needs discussion. It is true that rubrics are usually given to students to help them complete a specific assignment. However, as noted above, if the rubric is focused solely on what a student needs to do to complete just that task, it is not much use in helping the students learn. Such a rubric will be useful for summative assessment because it can ensure that students are clear about the requirements of the task so that "If they know it they show it." But to support learning, rubrics must be relevant to multiple tasks, ideally over an extended period of time, so that students can come to understand the deep principles that underpin quality work across a range of settings and tasks. Although we have not conducted a definitive study, it seems to us that most of the rubrics given to students are much better suited to summative functions of assessment rather than guiding student learning in the future.

In North America, rubrics are almost always presented in terms of a rectangular grid, with the vertical dimension used to identify different aspects of the task and the horizontal dimension used to identify different degrees or gradations of quality. Some authors have degrees of quality increasing from left to right, which makes sense given that in English, we read from left to right, and in mathematics and science, graphs are almost always drawn with values on the horizontal axis increasing from left to right. Other authors prefer to have degrees of quality decreasing from left to right so that the rubric headlines high-quality rather than low-quality work. We can see advantages of both and

no compelling reason to prefer one over the other. Whichever works best for you and your students is the one to use.

Another issue that needs to be settled in devising a rubric is how the levels of quality are described. Some rubrics simply use numbers (e.g., "1," "2," "3," and "4") while others use labels such as "below basic," "basic," "proficient," and "advanced," or "novice," "apprentice," "practitioner," and "expert." The problem with such labels is that they can send the message to students that once they get to the top level, they are done, which, of course, in learning is never the case—there is always a next level. Some teachers have asked students to choose for themselves how the different levels of achievement are described. One group of students chose a bicycle to represent the lowest level of achievement, a car for the second, and a space rocket for the third. However, perhaps the best idea of this kind we have seen from students, which really highlights the developmental nature of the best rubrics, was a group who chose frogspawn for the lowest level, tadpoles for the second, and a frog for the third.

The rubric format is now so common we often forget that the format makes us do things that may, or may not, be helpful. For example, we have yet to see a rubric that has a blank cell. There is something in every cell because, for some reason, there has to be, and so we contrive entries by using adjectives such as "some" and "many" to create distinctions that allow us to put something in every cell. However, for some aspects of work, there isn't really a continuum of development, but just two levels, equivalent to understanding something and not understanding something. Think of riding a bicycle. At level 1, there would be "Can't do it" and at level 4, "Can do it." There really isn't anything at levels 2 and 3. Children move so quickly from level 1 to level 4, it would be artificial and misleading to put anything between them.

Another problem with the rubric format is it creates the impression that the degrees of quality are somehow equivalent across the different components of the rubric. In other words, giving a 3 to a piece of work on two different components of a rubric sends the message that these two aspects of the work are equally good, and, perhaps more problematically, that if a student gets a 2 on one component of the rubric and a 3 on all the others, the work is somehow weaker on the component that was scored a 2. This may, of course, be the case, but in our experience this is unlikely unless the rubric has been specifically constructed in this way, which makes even greater demands on how the rubric is constructed, and which in turn leads to the problems we discussed in the previous paragraph.

As we said earlier, none of this means that rubrics are not useful; they can be highly effective ways of communicating standards of achievement to our students, not least because the format is extremely simple and compact. But all teachers need to be aware that sometimes we pay a high price for that simplicity and compactness. If it makes sense to arrange your success criteria in the form of a scoring rubric, then by all means do so.

But also
a differ

Cautio

Use E

When q
to deter:
want to
of work
another
Similarl
into the
intentio
about th
work in
pieces a
ing on k
the shor

Be Av

Perhaps
be effec
the wor
in Burg
what Gt
ing have
then eve
high-qu
without
agreeme
when, it
work, w
of this c

Don't

Many p
with the
success
we thin
is the ex
teacher

But also consider what you are giving up by using this format, and also consider whether a different format, such as a checklist, might be more effective for your students.

Cautions

Use Examples Where Deep Features Are Not Aligned With Surface Features

When quality has multiple dimensions, care needs to be taken that students are not able to determine which pieces are better by focusing on surface features. For example, if you want to get students to focus on characterization, it might be useful to have one piece of work with deep characterization but poor grammar, spelling, and handwriting, and another piece with good grammar, spelling, and handwriting but weak characterization. Similarly, where pieces of work involve extended pieces of writing, be careful not to fall into the trap of "longest is best." For example: A social studies teacher is sharing learning intentions with an eighth-grade class by giving them a source document, a key question about the document, and four different responses to the question. Students are asked to work in pairs to put the four responses in order of quality. However, because the longer pieces are better, some students are able to generate the correct order of quality by focusing on length. The teacher learns his lesson, and in the future makes sure that sometimes the shortest answer is the best, and the longest answer is the worst.

Be Aware That Sometimes Quality Cannot Be Put Into Words

Perhaps the greatest danger with rubrics is that they are used where quality really cannot be effectively described in words. As Roger Shattuck once said, "Words do not reflect the world, not because there is no world, but because words are not mirrors" (quoted in Burgess, 1992, p. 119). Sometimes the best we can do is help our students develop what Guy Claxton calls a "nose for quality" (Claxton, 1995). If the teachers in a building have a shared sense of quality, and they are able to share this with their students, then even if this understanding cannot be put into words, there is an effective basis for high-quality assessment. When teachers have a shared understanding of quality, even without rubrics, different teachers will give similar ratings to students' work and it is this agreement between teachers that provides a secure basis for summative assessment. And when, in addition, students understand quality, they are better able to guide their own work, which is the key aspect of the first strategy of formative assessment and the focus of this chapter.

Don't Abdicate Responsibility for Quality

Many people describe the process of constructing learning intentions and success criteria with the students as a process of "negotiating" or "co-creating" learning intentions and success criteria with students. We have no problem with the process, nor the terms, but we think it is important that it is clear the teacher is in charge of the process. The teacher is the expert in the subject and the students are novices, so while it is appropriate for the teacher to listen to the students' views, it is essential that the final learning intentions

and success criteria, however they are developed, are consistent with the teacher's understanding of quality in the subject.

Enhancements

Use Rubrics as the Starting Point for a Dialogue Between You and Your Students

One of the advantages of rubrics is that they allow teachers to give really quick feedback to students on their work. For example, a teacher gives her students a rubric for a writing task that identifies six aspects of quality: focus, organization, grammar, effectiveness, length, and completeness. For each of the six dimensions of quality, the rubric identifies four levels of achievement (1, 2, 3, and 4) and for each level, provides a description of the typical features of work at this level. Because the teacher knows that the students have a copy of the full rubric, rather than detailing the features of the work, the teacher can provide feedback simply by giving a rating for each of the dimensions of quality, as shown in Figure 3.1. The students can then refer to the rubric to see what each of the ratings mean.

Figure 3.1: Use of a Rubric Summary to Provide Rapid Feedback

Focus	1	2	3	④
Organization	1	2	3	④
Grammar	1	2	③	4
Effectiveness	1	2	3	④
Length	1	2	③	4
Completeness	1	②	3	4

However, where the goal is to help students understand the learning intentions and success criteria, a more powerful activity is to get the students to do a self-assessment for each of the dimensions specified in the rubric and then for the teacher to look at the students' self-assessment and report back to the student, "I would have given you a higher rating in two of the categories and a lower rating in one of the categories. See if you can figure out which categories these are, and why." This is an example of making feedback into detective work, discussed in further detail in the chapter on feedback.

Find Out What Your Students Think They Are Learning

No matter how much the teacher involves the students in the "co-creation" of the learning intentions and success criteria, and no matter how well those learning intentions and success criteria are designed, there is always the danger that students' understandings are different from what was intended. It is therefore useful periodically to ask students what they think they are learning.

When we taught middle and high school math, we would both get our students to undertake “mathematical investigations”—open-ended (and often open-begun!) tasks to get students to understand how mathematics generates knowledge (as opposed to what knowledge mathematics generates). In one such lesson, Dylan was getting a class to investigate the number of “polyominoes” that could be made from a given number of squares, given that each square had to join edge to edge and corner to corner with other squares. The class quickly established that there were only five distinct tetrominoes (i.e., five different ways of arranging four squares, not counting rotations and reflections as different). He then asked students to find out how many pentominoes there were. When he asked students why they thought they were doing this particular investigation, they replied that they thought what was important was finding out how many pentominoes there were. It turns out there are exactly twelve pentominoes, but there is nothing interesting or important about this fact. This particular task had been chosen to prove a challenging (and somewhat messy) context for students to practice working systematically and logically, but the students did not realize this.

One particularly powerful way to find out what your students think they are learning is to ask them to make up test items, with correct answers, about what they have been studying. You will often find out what you think they are learning is not the same as what they think they are learning, which is useful, but this particular task of generating items is beneficial to students in itself.

It has long been known that learners remember responses they generate themselves better than those responses that are given to them, and this is now often called the *generation effect* (Slamecka & Graf, 1978). In particular, one hour students spend writing test questions on what they have been studying results in more learning for them than one hour spent working with a study guide, answering practice tests, or leaving the students to their own devices (Foos, Mora, & Tkacz, 1994).

Big Ideas, Learning Progressions, and Staging Posts

So far in this chapter, we have treated learning intentions and success criteria for one lesson as relatively isolated from those for another lesson. The problem with such an approach is that each lesson or unit stands or falls on its own, few connections are made, and the student’s experience is fragmented. In this final section, we lay out an approach that you can use to begin to look for the “big pictures” in your curriculum. The idea is that instead of dealing with pine needles and twigs, you identify the trunk of the tree and the main branches.

In terms of the logic of Ralph Tyler’s approach to curriculum and instruction we presented at the beginning of this chapter, you might assume that we should have started the book with this idea, but in our experience, it can take many years of teaching to be clear about the big picture of what we are teaching. Moreover, the idea that all our teaching should be focused on big ideas is likely to be unrealistic. Our state and national

standards simply aren't well-enough designed for this to be possible, at least in any reasonable period of time. Our recommendation is therefore that you should work both bottom-up and top-down. Begin by clarifying learning intentions and success criteria lesson by lesson and unit by unit, and share them with your students. Over time, however, we encourage you to give some attention to a top-down process of starting from the big ideas in what you are teaching, mapping out the routes your students are likely to take in developing their understanding of the big ideas and figuring out what the most important staging points are along the way. As well as making the curriculum simpler and more manageable for teachers, focusing on the big picture in this way helps students see how what they are learning builds on what they have done before, and how their current learning prepares them for what is to come. This idea of progression in learning is particularly important in helping students develop what Carol Dweck calls a *growth mindset*, which is explored in more detail in Chapter 5.

We can, of course, define a big idea in any way we please. To us, the most powerful way involves more than using a label. To be useful as a big idea, a concept should provide a way of making sense of a number of related phenomena in an effective and efficient way. Most importantly, a big idea is *generative* in the sense that you can apply it in areas that the student has not yet encountered. It does not just make sense of what the student already has experienced, but it also provides a way of making sense of what the student encounters in the future. For this reason, the best big ideas form part of student learning over many years and can even bridge from elementary to secondary school.

Just how big an idea needs to be to qualify as big is, of course, a question that does not have a simple answer. When there are relatively few big ideas, obviously you can explore each idea to some depth, but the risk is that the ideas become so vague as to be difficult to use. A relatively large number of big ideas allows for more comprehensive coverage, but the danger is then that the ideas become small atoms of content, and their relationship to the other content is lost. One group of researchers that looked at this recommends that for each school subject, the number of big ideas should be between five and ten (Popham, Keller, Moulding, Pellegrino, & Sandifer, 2005), but there could be many good reasons for more or fewer.

A group of science education researchers from around the world led by Wynne Harlen proposes a list of ten big ideas of science that could provide a focus for the learning of science from kindergarten to high school (Harlen, 2010, p. iv):

1. All material in the Universe is made of very small particles.
2. Objects can affect other objects at a distance.
3. Changing the movement of an object requires a net force acting on it.
4. The total amount of energy in the Universe is always the same but energy can be transformed when things change or are made to happen.

5. The composition of the Earth and its atmosphere and the processes occurring within them.
6. The solar system is a very small part of one of millions of galaxies in the Universe.
7. Organisms are organized on a cellular basis.
8. Organisms require a supply of energy and materials for which they are often dependent on or in competition with other organisms.
9. Genetic information is passed from one generation of organisms to another.
10. The diversity of organisms, living and extinct, is the result of evolution.

In addition, the group looks at how scientific knowledge is created and suggests four big ideas about science (Harlen, 2010, p. v):

1. Science assumes that for every effect there is one or more causes.
2. Scientific explanations, theories, and models are those that best fit the facts known at a particular time.
3. The knowledge produced by science is used in some technologies to create products to serve human ends.
4. Applications of science often have ethical, social, economic, and political implications.

Because big ideas are meant to have traction over a number of years, they aren't of much use as learning intentions, so it is generally helpful to break big ideas down into smaller elements. This may seem paradoxical. The whole purpose of a big idea is for it to be a single concept, and breaking it down into smaller concepts seems to undermine the entire purpose of having big ideas. However, students' understanding of big ideas is meant to develop over time, so there is nothing inherently contradictory in having a big idea that can be broken down into a number of developmental steps.

In recent years in the United States, these developmental sequences have been called *learning progressions*—indications of the stages through which students typically progress in reaching an understanding of a big idea. For example, learning progressions in science is defined as “empirically grounded and testable hypotheses about how students' understanding of, and ability to use, core scientific concepts, explanations, and related scientific practices grow and become more sophisticated over time, with appropriate instruction” (Corcoran, Mosher, & Rogat, 2009, p. 20).

Perhaps not surprisingly, most of the early work on developing learning progressions focuses on science and mathematics, because the developmental trajectories of learners in these subjects have been studied in some detail in recent years (see, for example, Corcoran et al., 2009; Daro, Mosher, & Corcoran, 2011).

However, as far as we are aware, the idea of learning progressions applies equally well in all school subjects. For example, the following is a learning progression proposed by a working group in England looking at students' developing conceptions of chronology and cause and effect in history (National Curriculum History Working Group, 1990, pp. 121–125):

- Recognize everyday time conventions.
- Place a few straightforward events in chronological sequence; demonstrate, by reference to the past, an awareness that actions have consequences.
- Demonstrate awareness of a variety of changes within a short time span; demonstrate an awareness of human motivation illustrated by reference to events of the past.
- Employ appropriate chronological conventions by using time lines or other diagrammatic representation of historical issues; understand that historical events usually have more than one cause or consequence.
- Demonstrate a clear understanding of change over varied time periods; understand that historical events have different types of causes and consequences.
- Recognize some of the complexities inherent in the idea of change, when explaining historical issues; when explaining historical issues, place some causes and consequences in a sensible order of importance.
- When explaining historical issues, show a detailed awareness of the idea of change; when examining historical issues, draw the distinction between causes, intentions, motives, and reasons.
- Apply extensive understanding of change to complex historical issues; produce a well-argued hierarchy of causes for complex, historical issues.
- Demonstrate an awareness of the problems inherent in the idea of change; demonstrate an awareness of the problems inherent in the idea of causation.
- Demonstrate a clear understanding of the complexities of the relationship between cause, consequence, and change.

Now, different social studies teachers might have other views about the appropriateness of this particular learning progression, but the important point is that this does attempt to lay out how the student improves when he gets better at history.

It might be assumed that, to save teachers all over the country from “reinventing the wheel,” learning progressions could be developed at a national or at least state level. However, because the sequence in which topics are taught vary from state to state, district to district, and even building to building within the same district, progressions developed nationally or at state level are likely to need adaptation if they are to faithfully

trace the learning journeys students take in any particular school. For some ideas about generating learning progressions that meet the needs of your students, see Leahy and Wiliam (2011).

After learning progressions are developed, it is then possible to select particular points along the “learning journey” as learning intentions. Most often the learning intentions relate to what a teacher can cover in a particular instructional period, such as a lesson or unit of work, but it is also sometimes useful to define learning intentions as what David Perkins calls *troublesome knowledge* (Perkins, 1999), which is discussed in more detail in the next chapter. The important idea here is if we know there are particular things that are difficult for students, then we need to find out whether they understand them.

The goals of education in play, of course, strongly influence the nature of big ideas, and what might be a big idea for one person might be trivial or even irrelevant for another. Big ideas may not even be “true” in fields where knowledge is contested, but they always provide ways of making sense of experience.

Tips

Get the Grain Size of Big Ideas Right

When developing big ideas for their curriculum, it is difficult, but important, to get the *grain size* of big ideas right for the task at hand. When the grain size is too large then the idea tends to be so loosely framed that it is not helpful in guiding learning. For example, one group of teachers adopted the following list of big ideas for social studies: government and civics, cultures and societies, economics, geography, and historical perspective (NTI Social Studies, 2014). These are useful ways of dividing up the content matter, but they provide little guidance about what the students should be learning in this subject.

On the other hand, when the grain size is too small, the number of big ideas becomes large, and ultimately unmanageable. Think of this: in the context of teaching about the American Revolution, we could adopt as potential big ideas the nature of patriotism, whether the bravery of the founders was an admirable quality, or whether the British government was fair to the colonists (Propel, no date). The problem with this list is the grain size of these ideas is so fine that dozens, and even perhaps hundreds, would be needed to cover just one year’s work in social studies.

Table 3.2 shows examples of suggestions for big ideas that seem to us to be sufficiently general to be powerful over a number of years of a child’s development but also sufficiently precise for determining whether students are improving.

Table 3.2: Suggestions for Big Ideas

Subject	Big idea
English language arts	The “hero’s journey” (Campbell, 1949/2004) as a structure for understanding myths and legends.
Social studies (geography)	Patterns of human development are influenced by, and in turn influence, physical features of the environment.
Social studies (history)	All historical sources are products of their time, but understanding the circumstances of their creation helps us resolve conflicts of evidence.
Mathematics	Fractions, decimals, percentages, and ratios are all ways of expressing numbers, and every real number can be represented as a point on a number line.
Science	All matter is made of very small particles.

Cautions

Not All Useful Learning Intentions Will Be Big Ideas

Even with a clear focus on big ideas, there is always a need for some “mopping up” of individual pieces of content that don’t fit neatly anywhere else. For example, it is useful for students to know how to read Roman numerals, but understanding Roman numerals is not a big idea in the sense we have described it here. Moreover, the Roman numeral system is highly idiosyncratic, using bases of both 5 and 10, so learning about the Roman numeral system doesn’t help you do anything apart from using Roman numerals; the topic of Roman numerals is probably best treated as a “one-off.”

In our view, it is generally possible, over a period of several years, to organize the curriculum so that at least half of the content is in the form of big ideas. This may seem modest, but the clarity brought to the curriculum, for both the teachers and students, is a huge advantage. If you can get two-thirds, you’re doing well, not least because state standards were never designed to be coherent, so imposing some order on them is challenging.

Learning Progressions Need Both an Empirical and a Theoretical Basis

The fact that students happen to learn something in a particular sequence does not mean that the sequence is an appropriate learning progression. After all, the particular learning sequence might just be an accident of the particular way that students had been taught. To be useful as a learning progression, there has to be some underlying rationale that students are likely to follow in using this particular sequence in advancing their work. However, as every teacher knows, by itself, a compelling logical rationale that students should learn something in a particular sequence does not mean that they will do so.

We have all taught students who, for whatever reason, seem to find easy things difficult and difficult things easy. That is why Denvir and Brown (1986) suggest there should be both a logical and an empirical basis for the learning progression: there should be clear reasons why the elements of the progression are arranged in a particular sequence, and it should be rare (if not impossible) for students to reach the later parts of the progression without going through the earlier parts.

Enhancements

Involve Students in Clarifying the Big Ideas

As students get more involved in their own learning, they become clearer about their own learning, and over time, develop views about what the subjects they are studying are really all about. In the same way that we recommended that students are involved in the co-construction of the learning intentions and success criteria, it can be useful to ask them for their views about what the big ideas are of the subjects they are studying. This helps students develop a more strategic view of the subject and also provides useful information for you about the students' experiences. For example, if your students tell you that mathematics is all about getting the correct answer quickly, or English language arts is all about reading and writing, that is a cue for you to change what you are doing. Perhaps most importantly, our students can sometimes help us think about the subjects that we teach in a way that had not occurred to us and help us move our own thinking forward as well.

Case Studies

In this section, we present two case studies of teachers sharing learning intentions and success criteria with students. The extracts are transcripts of videos taken from Wiliam and Leahy (2014). We encourage you to read the transcripts, and then think about what you have read by responding to the reflection questions at the end of the transcript before moving on. The two case studies in this chapter, and the two in the chapter on "activating students as owners of their own learning," are both taken from secondary schools. You may be surprised by the extent to which the conversations between teachers and students, and among the students, are dominated by what they need to do to get a good grade on an examination. Of course doing well on tests or examinations is not the only—or even the most important—reason to educate young people, but they are, in many schools and districts, a feature of classroom life. We have included these examples to show that using formative assessment to improve the quality of instruction is entirely compatible with success on tests and examinations. Indeed, we know of nothing else that you can do to have a bigger impact on your students' test and examination performance. The point is that if high-stakes tests and examinations are a part of the system in which you work, then formative assessment can be used both to deepen student learning and increase test scores at the same time. You don't have to choose between teaching well and increasing test scores. You can do both.

English

A group of tenth-grade students is preparing for a formal written examination that will assess their understanding of English literature. The examination consists of two examination papers, each seventy-five to ninety minutes long, during each of which the students will attempt two or three questions, so they will have thirty to forty-five minutes for their response to each question.

Chris Flack, the teacher, shares with the students the following handout, which includes the four lists of success criteria associated with the four highest grades available—C, B, A, and A* (equivalent to an A+ grade)—in the examination they are about to take.

Handout

Match the four lists of criteria below to the correct grade. Which one is C, B, A, and A*?

- Shows originality of analysis and interpretation
- Sustained evidence of enthusiastic and personal response
- Shows concise textual analysis and understanding

- Demonstrates analytic skill
- Shows independent understanding
- Shows layers of meaning—how meaning is conveyed through appropriate detail
- Shows relevant comparisons

- Demonstrates insight
- Explores connections between texts
- Explores how meaning is conveyed through language, structure, and form
- Explores connections between texts—refers to some detail

- Demonstrates analytical and interpretive skill
- Explores alternative approaches and interpretations
- Considers and evaluates the way that meaning is conveyed through language, structure, and form

With a partner, the students discuss the key differences between each grade. Here is a transcript of the video:

- Teacher: Exams are creeping up on us now. We are going to try to work out how you are going to be assessed in this thing. You've got a sheet in front of you. It's got the GCSE criteria for C, B, A, and A*. Have a read, and try and work out which one is which first of all. So where's the B, which one is the A, which one's the dirty C? Obviously once you have done this, I want you to work out why, because anyone can write a grade next to a sentence. Try to work out why that grade is what you are thinking it is.

Two students working together, followed by a whole-class discussion:

- Polly: For a C you don't actually have to analyze anything.
- Sofia: You just have to show that you understand it.
- Polly: Yes, basically that you understand. And for B it's more.
- Teacher: Thirty seconds.
- Ana: And it says sustained evidence so it's not just like you . . .
- Teacher: Right, then. Which one is the A then? Steph.
- Steph: Demonstrate an analytical and interpretive skill. Explore alternative approaches and interpretations. Consider and evaluate the way meaning is conveyed through language, structure, and form.
- Teacher: Good. It is that one. How did you know that it was an A? Someone else that isn't Steph.
- Briony: Because for the A* it says "originality of analysis," and you have been telling us that we need an original idea to get an A*.
- Teacher: Excellent. Good. So the difference between an A and an A* is that for an A* you really have to try to pin down that originality of analysis. What other differences did you spot between an A and an A*?
- Polly: For an A* you really have to have an opinion on it and be enthusiastic about it.
- Teacher: Yes, you do. We have had this discussion before. Enthusiasm doesn't mean you have to jump up and down and pretend you love the poem. It means you have to have an enthusiastic response. Good, anything else that you spotted? Prin.
- Prin: It uses words like *sustain* and *concise*.
- Teacher: Absolutely, you've got to get straight to the point for an A*. You've got to know what you are saying and get there as accurately as you can. Good, hopefully by the end of the lesson you will be another step nearer that A*.

The video ends with a student and the teacher reflecting on the process after the lesson:

- Anika: If they give a list of, like, A*, what you need to do for an A* and what you need to do for a C, you'll make sure you do more A* work than a C, because you know what you are doing. But if you don't have that success criteria you don't know what you are doing, what you need to do to get that grade.

Teacher: It is slightly disingenuous because some of them aren't an A yet. They are not a miracle class. But by talking only about A and A*, the Cs and the Bs, you'll find they reach further toward an A. So if I keep talking about an A*, even those that aren't an A will still make progress.

Reflection

Before reading further, take a few moments to reflect on this transcript, and think about what how effective this kind of activity was in helping the students be clear about what they had to do to get a good grade.

German

Ian Storey is teaching a group of eighth-grade students in an accelerated class. Shortly, these students will be taking an examination that most students take at the end of tenth grade. The examination includes three pieces of school-based assessment—extended pieces of writing that the students complete individually, under controlled conditions—and before they redraft a piece of writing they have just completed and been given feedback on, the teacher wants the students to have a good understanding of the criteria against which their work will be assessed.

Teacher: Going from a C/B to an A/A* for our coursework is what we are all aiming for. So, I thought the best thing to do would be to moderate some coursework, not yours but other pieces of coursework that have already gone through the process of GCSE examination and they've been moderated by a team of teachers. So remember, give your examiner a break and not a breakdown (class joins in).

Thank you. Okay. So,

Ihr habt zehn minuten. Ihr must jetzt diese drei Kurzarbeit Blätter lesen und... wie gut oder schlecht es ist. You have ten minutes. You must now read these three sheets of coursework, and say how good or bad it is.

So, in English, we have communication, range of language, and accuracy mark. You will need your coursework descriptor sheets that we used a few weeks ago. It's this sheet that you are going to talk about in eight minutes (there is a countdown timer on the whiteboard).

Students in pairs looking at work, followed by a whole-class discussion:

Liam: He spelt [*sic*] "hotel" wrong.

Ben: It's got a wide range of vocabulary.

- Liam: But not used correctly.
- Ben: Language range. He has tried to use big words and complicated.
- Liam: But he hasn't used them successfully, so about a three or a four I guess.
- Ben: I'd say a three because there is some variety.
- Teacher: Oscar (choosing a student from small name cards held in hand).
- Oscar: Communication. For marks, I gave him between a seven and eight because there was a lot of detail, and he did lots of descriptions. He gave a lot of opinion.
- Aria: I think it would be a seven, but it were [sic] really good but he could have included a bit more information.
- Teacher: Yes, it was lacking a bit of information, but it definitely was a seven, and you are absolutely spot-on with that. Oscar, what did you give it for range of language?
- Oscar: Five. He was quite strong in the past tense, but he could have used more complex sentences.
- Teacher: Interesting, and you gave it six for accuracy?
- Oscar: Yes, because there was [sic] only a few minor mistakes.
- Teacher: Yes, right, okay. Rosie, tell me why you disagreed with Oscar and said that it should have a six rather than a five for range.
- Rosie: Because he used a large variety of sentence types and vocabulary. He expanded all the sentences a lot more.
- Teacher: Interesting. Christian, talk to me a bit about the work you and Rhiannah did.
- Christian: We gave it, for communication, a four because it had quite a fair bit of useful information but not enough to give it a five. Language range, a four, because again it used a lot but just not enough. And for language accuracy we gave it a four, as well, because although he used quite a lot of different words not all of them were spelt right or used in the right context.
- Teacher: Who agrees that this coursework is not worthy of a C?
- Krishna: All the stuff was not explained in detail.
- Teacher: Detail. That's the key word, isn't it?—Detail. If there is no detail you can't get anything above a C. Now we need to consider any feedback that we could give this student. What feedback could we give this student? David, tell me please.

- David: Widen his range of sentences a bit more.
- Teacher: Well, you are really on the right lines there. Well done. Zoe.
- Zoe: Use more opinions and a wider range of opinions, because he uses the same ones.
- Teacher: Good. Really important what Zoe was saying. Just bouncing off Zoe there. What she was saying about opinions. Opinions are really, really important.

After lesson ends:

- Teacher: From the lesson they have a clear understanding of what it is the exam board is looking for. They now know the marking criteria. They were given student-friendly assessment criteria grids as well. I did change it slightly because sometimes the exam board language isn't that student-friendly. So from that point of view they now know what to expect in a piece of coursework. They also had a chance to talk to their peers, they got ideas, they exchanged ideas about what they could include in their coursework drafts, and, all in all, they demonstrated to me that they can develop their coursework and take it forward so that they can gain not just C and D grades but looking at As and A*s.

Reflection

As with the previous case study, take a few moments to reflect on this transcript, and think about how effective this activity was in helping the students understand the criteria against which their work would be judged. In particular, note any similarities and differences between this example and the previous case study.

We end the chapter with a recap, followed by a reflection checklist you might find useful to complete, a planning sheet for using a technique with a class, and a peer observation sheet (so that you can ask a colleague to give you feedback).

Each of the sheets are at the back of this book for you to copy if needed, so feel free to write all over the following sheets!

RECAP

The aim of any learning intention or success criterion is not to help the students complete the activity—it is to help them learn.

- Sometimes you should not set a learning intention at the start of the lesson.
 - You may not know exactly where the lesson is going—it is the experience rather than the outcome.
 - It may spoil the journey.
 - It may not inspire students.
- Keep the context of learning out of the learning intention.
 - Get students to apply in other contexts what they have learned.
 - Mix it up.
 - Differentiate success criteria, not learning intentions.
 - Make learning intentions and success criteria generic.
 - Be aware of the strengths and weaknesses of analytic and holistic learning intentions and success criteria.
- Start with samples of work rather than rubrics, to communicate quality.
 - Use anonymous work.
 - Compare two or more pieces of work.
 - Use a document camera.
 - Use examples where deep features are not aligned with surface features.
 - Use rubrics as the starting point for a dialogue between you and your students.
 - Recognize that sometimes, quality cannot be put into words.
 - Help students develop a nose for quality.
 - Don't abdicate responsibility for quality.
 - Find out what your students think they are learning.
- Use big ideas, learning progressions, and staging posts.
 - Over time move to a top-down process starting from the big ideas.
 - Develop learning progressions.
 - Get the grain size of big ideas right.
 - Understand that not all useful learning intentions will be big ideas.
 - Ensure that learning progressions have both an empirical and a theoretical basis.

REFLECTION CHECKLIST FOR STRATEGY 1: Clarifying, Sharing, and Understanding Learning Intentions and Success Criteria

	I don't do this	I do this sometimes	This is embedded in my practice	I could support someone else
I know what the learning intention of the lesson is, although sometimes I do not tell the students at the start of the lesson.				
I keep the learning intention and success criteria for a lesson context free.				
I communicate quality by using at least two pieces of anonymous work.				
At the end of a lesson I sometimes ask my students what they have learned.				
I use rubrics to discuss quality with my students.				
Other techniques for this strategy that I use to improve student learning:				