Chapter 1

Predicting

INTRODUCTION

My wife, Anne, teaches kindergarten, and during the 2020–2021 school year she taught it remotely from our dining room. Whenever you find it challenging to maintain the focus of your students, you should envision attempting to corral the attention of a couple of dozen five-year-olds through your computer screen and feel just a little bit better about your prospects. I was working from home as well during the pandemic, and so I had the pleasure (I think) of overhearing almost everything she did with her students all day long. I hummed along to the songs she sang about the weather and days of the week, counted my way to one hundred by ones and fives and tens, and listened to her patiently introduce letters and syllables and words to her students.

At some point every day she would read a picture book to the faces gathered on her screen, and this process always started in the same way. She would display the cover of the book, read aloud the title, and then ask her students some variation of the following question: "What do you think this book is going to be about?" She would then pose additional questions based on the specifics of the picture and the title. "What do you think will happen if you give a mouse a cookie?" "What would it be like to take a polar bear for a walk?" A final set of questions was designed to surface whatever students already knew about the subject matter of the book: "Where do polar bears live? Are polar bears pets?" The students would sing out their answers, some of which were on target and some of which were not. (Polar bears do NOT make good pets.) But almost always the process of reading a story aloud to the students began with a few minutes of these efforts by the students to make some predictions about what would unfold in the story they were about to hear.

In higher education we tend to go into class with our content guns blazing: I HAVE SOME STUFF TO TEACH YOU, LET'S GET STARTED! But the research on human learning that we will consider in this chapter suggests that the first step in the learning process should be to follow Anne's lead—not necessarily through her specific activity of asking pre-reading questions–but by asking students to engage in predictive activities of some kind or another *before* we expose them to new course content. Such activities could include inviting students to answer questions about what they are about to learn or experience, but could also include asking students to solve problems that are beyond their current ability level or to try their hand at a new skill before they have been given any formal instruction. This approach can seem counterintuitive since it turns the normal teaching sequence on its head. The operating assumption of many teachers runs like this: *First* I teach them the material, *then* I have them answer questions or solve problems with it. Read the story about polar bears, and then ask them where polar bears live. Give students a lecture on poetry, and then have them interpret a poem. Show students how to solve a particular type of math problem, and then give them one to solve on their own. But learning research suggests something quite different. It tells us that asking students to use their existing knowledge and skills to struggle with the material *before* we teach it can provide a robust foundation for deeper learning.

Providing opportunities for students to undertake this struggle doesn't take much time or effort, which is why it makes an ideal place to start a book on small teaching: just a few minutes at the beginning of a class period or a unit or even a course has enormous potential to improve both student learning and your teaching.

IN THEORY

To understand how and why predictive activities support learning, consider first an elegant series of experiments conducted by three researchers at UCLA (Kornell, Jenson Hayes, and Bjork 2009). The authors asked subjects in one of these experiments to memorize a series of loosely connected word pairs, such as *whale–mammal*. One group of the subjects was given 13 seconds to study each word pair; the other group was given 8 seconds to see only the first word and to make a prediction about the second, after which they had 5 seconds to see the full word pair. Since the second word of the pair was linked to the first, but only one of many possibilities (since you might well guess *sea* or *ocean* or *large* if you saw only the word *whale*), participants typically guessed the second half of the pair incorrectly. Note also that the subjects in the second (prediction) group had only 5 seconds to view the correct answer, so 8 seconds less than those in the first (nonprediction) group. Yet in spite of this shorter study time, and in spite of the fact that subjects in the second group frequently predicted the second half of the word pair incorrectly, the subjects in the second group performed significantly better than those in the first group when they were asked to recollect the word pairs on a subsequent exam: 67% accuracy in the second group versus 55% in the first. The unsuccessful prediction attempts in the experimental group, the authors explain, "were, remarkably, more effective than was spending the same time studying the answer to be recalled later" (p. 994). In other words, taking a few seconds to predict the answer before learning it, *even when the prediction was incorrect*, seemed to increase subsequent retention of learned material. This was true even when that prediction time substituted for—rather than supplemented—more conventional forms of studying.

The results of laboratory experiments like this one prompted another group of researchers, led by cognitive psychologist Elizabeth Bjork, to see whether they could reproduce the positive learning effect of prediction in an actual classroom (Carey 2014b). The researchers gave students in Bjork's introductory psychology class short multiple-choice pretests before some of the lectures in her course. Since the pretest questions asked them about material that had not yet been covered, the students performed about as well on the pretests as they would have from guessing randomly—again, as in the laboratory experiment, they made plenty of wrong predictions. Lectures on the subject matter followed immediately after the pretests, so the students received quick feedback on their answers. At the end of the term, the students took a final exam that contained multiple-choice questions similar to the ones on the pretests. The results paralleled the results of the laboratory experiment almost exactly: students performed around 10% better on questions from the subject areas in which they had been pretested than on those on which they had not. Bjork concluded from this experiment that "giving students a pretest on topics to be covered in a lecture improves their ability to answer related questions about those topics on a later final exam" (Carey 2014b). Note, of course, that even though the vocabulary has changed slightly here—from *prediction* to *pretesting*—the cognitive activity is similar: asking learners to give answers to questions or anticipate outcomes about which they do not yet have sufficient information or understanding. They are trying before they are ready.

Before we explore the reasons that prediction boosts learning, consider one final example of prediction in higher education, this one from an online environment (Ogan, Aleven, and Jones 2009).

Three researchers from Carnegie Mellon University developed an online tutoring program that demonstrated the power of predicting in helping students improve their intercultural understanding in hybrid language courses. The two French courses described in the experiment each met once a week in a face-to-face environment, but otherwise the students did their course work online. Part of the goal for these courses was to help students develop what the authors called "intercultural competence, that is, the ability to think and act in culturally appropriate ways" (p. 268). This can be an extremely difficult skill to develop, as anyone who has ever traveled in a foreign country can likely attest. The ability to speak the language of a foreign country does not necessarily guarantee your ability to understand how to hail a cab, tip in a restaurant, or approach a stranger in the Paris Metro to ask which train will take you to the airport in time to catch your flight home (as I once discovered, to my great sorrow). So, in this experiment we are moving beyond the realm of simply knowing and retaining information into the broader realm of comprehension—that is, understanding how to use and apply in other contexts the information you have learned. The intercultural competence sought by these instructors requires learners to think and act with their knowledge, not just report it back.

To help students acquire this type of deeper comprehension, two experts in computer-assisted learning worked with a language professor to develop an online tutoring program based on the use of film clips. In the control condition of this experiment, students were shown film clips highlighting cultural attitudes or behaviors that are normally taught in introductory French classes. As the students watched the short film clips, they had the opportunity to take notes on what they saw. The students in the experimental group, by contrast, were given the opportunity to use the power of prediction to improve their learning. Their film clips would pause at key moments, ask them to make a prediction about what was about to unfold, and then require them to ponder what actually happened once the clip had finish: the authors described their three-part sequence with the catchy phrase *pause–predict–ponder*. The prediction the students had to make actually came from a drop-down menu of choices, but then they had open text boxes to explain why they made that prediction. After they had watched the remainder of the clip, they had to answer a simple question about whether or not their prediction was correct and then respond to prompts to help them reflect on their prediction, such as: “If so [i.e., if your prediction was correct], did you see anything you didn’t expect about the French culture? If not, what happened that you didn’t predict?” Students in both conditions concluded their viewing of the film clips with required postings to a discussion board to allow them to process and review what they had seen.

After the class, the researchers looked at two different measures to see whether the pause–predict–ponder exercises had improved student mastery of intercultural competence: student scores on tests of cultural knowledge, and their more general cultural thinking or reasoning skills through their discussion board posts. The students who had the opportunity to make predictions outscored their peers on the first exam by about that same 10% margin that we saw in the earlier experiments, with some diminishing returns on the subsequent assignments—which might tell us that prediction, like many of the active learning interventions we will consider in this book, especially helps new learners. Ratings of the posts in the discussion board also showed the students in the experimental condition performing significantly higher on assessments of intercultural competence than those in the control condition. In the discussion of their results, the researchers note that students made correct predictions only about 40% of the time, another point in favor of the notion (to be qualified shortly) that wrong predictions do no harm. They also point to an interesting side benefit they witnessed: students in the experimental condition posted more frequently on the discussion boards, and more frequently on target, than those in the control condition: "The experimental group showed a better ability to maintain a productive discussion compared to the control group (p. 283)." It seems that in this case the prediction activity also helped engage the students more thoroughly in the material.

Researchers who study the brain can help clarify the mechanics that underpin the results of all of these experiments. Neuroscientists are increasingly demonstrating that our brains are prediction-making machines, and that our learning stems most fundamentally from the cycle of making predictions and then adjusting our thinking in light of the accuracy of those predictions. Stanislas Dehaene is a professor of experimental cognitive psychology at the College de France, and the Director of the Neurospin Brain Imaging Center. In his book *How We Learn: Why Brains Learn Better Than Any Machine . . . for Now*, he offers a detailed but accessible tour of learning and the brain, and concludes that "Generating a prediction, detecting one's error, and correcting oneself are the very foundation of effective learning" (Dehaene 2020, p. 209). Our brains continuously create models of the world around us, use those models to predict how our experiences will unfold, and engage in corrective re-modeling in light of what actually happens. That corrective re-modeling is what we call learning. A new driver's brain makes continuous predictions about how the car will behave as she drives—she expects the car to slow down and come to a stop when she brakes. The first time she drives in the snow, applies the brake, and finds herself skidding, her brain notes the failure of her existing mental model—applying the brakes doesn't *always* stop the car with the same efficacy. She has to take weather conditions into account. Her mental model of driving expands, her predictive abilities as a driver improve. She has learned.

In a classroom setting, predictive activities reveal to students the gaps and problems in their existing knowledge of the course subject matter and provoke them to fill and repair their understanding. When I am asked to make a prediction or try a new skill, I am forced to surface whatever knowledge and skills I currently possess and use them in service of the task in front of me. In some cases, I will see immediately that I don't have what I need—*I have no idea how to do this*, I might think to myself. In a classroom setting, that recognition should push me toward learning: *If I want to succeed in here, I need more information*. But I might equally well think I *do* have enough information or skill to get the job done, in which case I will make a confident prediction, assuming that I already know everything I need to know. When I discover that my prediction was incorrect, I am put back on my heels again. What went wrong? What information was I missing? What would I have needed to know in order to get it right? Since very few students, if any, will come into our classes knowing everything they need to know already, the predictive activities we design are likely to reveal at least some errors in their thinking, showing them the gaps in their knowledge and skills. Ideally, your course material then gives them precisely what they need in order to fill those gaps.

Predictive activities can also give students a clearer understanding of what and how they need to study and learn in any particular course. As Elizabeth Bjork points out in relation to the practice tests in her experiments, "Taking a practice test and getting answers wrong seems to improve subsequent study, because the test adjusts our thinking in some way to the kind of material we need to know" (quoted in Carey 2014b). Well-planned pretests or predictive activities alert the students to essential course content and the testing style of the instructor. Envision a student sitting in a chemistry course in the fall semester of her first year of college. She had a year of chemistry in high school, and her high school teacher focused entirely on having students memorize facts and formulae, knowledge of which he assessed exclusively through multiple-choice tests. She has been conditioned by that teacher to think about and learn chemistry through rote learning, a memorization technique she will carry into her new course—unless, of course, her college professor opens the semester by asking the class to try to answer some conceptual questions, and mentions that these problems are similar to ones that will appear on the final exam. When that first-year student finds herself scratching her head and unable to come up with the answers, she will immediately see that she has to approach this course, and her learning, in a new way. She has to focus on conceptual understanding instead of the memorization of facts. If she hadn't encountered that first-day predictive activity, she might have spent the first five weeks of the semester, before the first major exam, focused entirely on repeating her high school study practices.

Ultimately and perhaps most simply, predictive activities mimic something we normally ask of learners who are attempting to master a skill: requiring them to *try before they are ready*. We can all likely draw from our experiences with attempts to master skills of one sort or another, and we know full well that however much one might read in advance about throwing a football or painting a portrait or giving a speech, the real learning happens after we have thrown ourselves into the situation and made that first (unsuccessful) attempt. When I took a class to become licensed in scuba diving, we spent the first half of every session in a classroom taking notes on some skill we would have to practice in the pool. I typically jumped in the pool for the second half of class thinking I had that skill mastered, but within a few minutes the gaps in my knowledge were revealed, and I floundered around for a while, doing it completely wrong until the instructor swam over and gave me the help I needed, at which point the real learning began.

We facilitate this type of learning in many academic contexts, asking students to try out cognitive skills before they are ready. I don't spend the entire semester lecturing to my freshman composition students about all of the writing techniques they will need to write a perfect academic essay and then give them one final assignment to show me how they have mastered those skills. I assign essays from the beginning of the semester, even though some of what they need to write great academic essays won't be covered for another 4 or 8 or 12 weeks. Asking students to make predictions before learning new material just represents another version of this common teaching approach.

MODELS

The ideal grounds for small teaching activities related to prediction are the openings of a learning experience—a course, a learning unit, a class period. Consider the following models for leveraging predicting into the beginning of anything you might be teaching.

Activating Prior Knowledge

Engaging in an act of prediction begins with my rooting around in my brain for anything which I already know that will help me make a successful prediction. It requires me to tap into my existing knowledge and skills, whatever they might be. I use my existing knowledge of French culture to predict what will happen in the video, or my prior knowledge of chemistry to predict what will happen in an experiment. This *activation of prior knowledge*, as learning researchers call it, has been demonstrated as having plenty of power to boost learning. "Prior knowledge plays a critical role in learning," explain Susan Ambrose and Marsha Lovett, "which means that. . .faculty members need to assess the content, beliefs and skills students bring with them into courses and. . .use that information as both a foundation for new learning as well as an opportunity to intervene when content knowledge is inaccurate or insufficient" (Ambrose and Lovett 2014, p. 16). The easiest way to activate the prior knowledge of your students is to ask them to make individual and collective knowledge dumps, telling you everything they know—or think they know—about your subject before you begin teaching them. You might think about this as a pre-prediction activity, or a half-step toward prediction. Instead of drawing up their prior knowledge and using it to solve a new problem, students are simply drawing up that prior knowledge and showing it to you—and the research suggests that this step on its own, *even if they don't use their prior knowledge to solve a new problem,* can prepare them for deeper learning.

Our first small teaching strategy might therefore be the easiest one to implement in the book: before you teach something new to your students, ask them what they already know about it. You could do this in multiple ways that fit the frame of small teaching:

- Prior to introducing new content in a course, ask students to take a pre-quiz or respond to two or three questions about the subject matter on the course's learning management system and then summarize those results briefly at the start of the first or second class.
- At the start of any individual class period, ask students to write down what they think they already know about the subject for that day. Tell them to list three to five things they have learned in previous classes or from their life experiences, and have them pair up to compare notes. Solicit a few responses and build on them in the opening of your presentation of the new material.

- At the start of the semester, devote part of one class period to assessing students' current state of knowledge, either through whole-class or group activities or through a written pretest. Once you have heard what students have to offer in such an exercise and have gained a glimpse into their existing knowledge, you can strategize how to build upon it most effectively in the course.

As an example of the third suggestion, when I teach a seminar on twenty-first-century British literature and culture, I set aside 30 minutes on the first day of the semester for a knowledge dump in which students respond to the following question: "When I say the word *British*, what are the primary impressions that form in your mind?" I pose this question and give them a few minutes to make a list on paper, and then we fill up the board with their impressions. I try to categorize their impressions so we can see patterns—and the same basic patterns always emerge. *British* to them means the royal family, traditional activities like having tea and biscuits or taking long walks in the country, green landscapes, and ruined castles. None of them ever mention the ethnic diversity of Britain, or the fact that *Britishness* encompasses the peoples of Scotland and Wales and Northern Ireland, or its recent political history. The second class period, in which we begin considering the political and cultural context of British literature, then gives me an opportunity not to dismiss their impressions but to build upon them as I introduce features of Britishness that they have never encountered. Throughout the semester I will occasionally return to the impressions we laid out on that first day and remind them about how our current text or discussion connects back to them. Such an opening-class activity can obviously work in almost any type of course, modified in ways that best fit your inclinations. For an upper-level course you might want to narrow down the questions, since they might have more general extensive knowledge than would fit into a paragraph. The overall strategy here remains happily simple: prior to the introduction of any major new chunk of course content, spend a class period (for a whole course) or a few minutes (for a new unit) asking them what they already know.

Polling Predictions

The use of classroom polling—whether you go high-tech with programs like Poll Everywhere or low-tech with colored index cards or even just raised hands—presents a very simple route to making prediction part of your course lectures, as Derek Bruff points out in his book *Teaching with Classroom Response Systems: Creating Active Learning Environments* (Bruff 2009). His chapter "A Taxonomy of Clicker Questions" points to the power of prediction to increase comprehension in addition to the benefits it should provide in boosting memory of individual facts and concepts. For example, Bruff gives an example of a math instructor at a small college who "shows his students a graphing program that allows him to vary a parameter in a function, such as the parameter ω in the function sin (ωt), and asks his students to predict what will happen to the graph of a function when he changes that parameter. After the students vote with their clickers, he demonstrates the correct answer using his graphing program" (p. 85). Students cannot answer questions like this with simple plug-and-chug–type knowledge; they have to possess a conceptual understanding of the problem to make an accurate prediction. The failure or success of their predictions enables them to re-model that conceptual understanding, which is of course the most important learning the course should induce.

Bruff also notes that classroom polling can support the opportunity for instructors to ask students for predictions behind the screen of anonymity, which can be useful in certain contexts. He provides an example of a health and wellness course at another university in which instructors want to draw attention to student perceptions of drinking on campus. The instructors first ask the students how many alcoholic drinks they consumed at their last social occasion, but then they also ask students to predict what they think the responses of their peers will look like. “The differences between the predicted votes and the actual votes,” explains Bruff, “are often surprising to students because it turns out that students are not always as risky as they think they are” (p. 86). The benefit of such a quick prediction exercise is the rich discussion that follows: “This activity can lead to a productive class-wide discussion of social perceptions of risky behavior and the role that marketing, in particular, plays in those perceptions” (p. 86). Such discussions, in other words, can encourage the students to reflect on why their predictions were incorrect—and the role that social media or beer commercials might play in driving their perceptions of their peers’ consumption of alcohol. The potential screen of anonymity provided by clickers obviously can serve a useful purpose when asking students to make predictions based on their own personal behaviors. It also could prove useful anytime you feel students might not want their predictions shared publicly, either because they want something kept private or because they feel they might be embarrassed by making a wildly incorrect prediction in front of their peers (or in front of you).

Prediction–Exposure–Feedback

Even without formal testing or the use of clickers, you can always ask students to make informal, in-class predictions about any course material to which they are about to be exposed. This could happen in almost any discipline, in any type of class.

Scientists know full well how prediction plays a role in the scientific method—in the form of the hypothesis—and likely already ask students to engage in predictive activities in their use of laboratory experiments and reports. But outside of the laboratory, and in other disciplines, instructors can still follow this same basic approach. Get into the habit of asking students to make predictions about new content based on their knowledge from earlier in the semester, from their previous courses, or from their own general knowledge. *How Learning Works* gaves two quick examples of this: "Before asking students to read an article from the 1970s, you might ask them what was going on historically at the time that might have informed the author's perspective. Or when presenting students with a design problem, you might ask them how a famous designer, whose work they know, might have approached the problem." In these kinds of questions, again, you are requiring students "not only to draw on prior knowledge but also to use it to reason about new knowledge" (Ambrose, Bridges, DiPietro, Lovett, and Norman 2010, p. 33). Ideally, you will both ask for the prediction and give them the opportunity to explain why they made it; doing so will require them to examine their thinking and might help them recognize their knowledge and skill gaps. Even more ideally, after you give them the answer you might ask them to explain why their predictions did or did not hold true.

Earlier in this chapter, I described for you three experiments on learning designed to demonstrate the power of predictions. Imagine that I was describing those experiments for students in a psychology or education course. I describe the setup of the first experiment—with one group of students trying to memorize the word pairs for 13 seconds, and a second group making their guess and then seeing the correct answer for five seconds—and then stop and ask the students to predict the results. I could do this in a generic way, and just invite some students to raise their hands and speculate, or I could offer specific options and invite students to select the one they think would be correct:

- Memorizers performed 50% better on the test.
- Memorizers performed 10% better on the test.
- The two groups performed about equally.
- Predicters performed 10% better on the test.
- Predicters performed 50% better on the test.

(If all is going well in this chapter, you will remember that the answer is number four.) Students could select the answer they think was correct using polling or simply by raising their hands. After I record the student votes, I might invite a few students to explain their answers, and then show the results. After those results have been revealed, I can use them to make the two points that matter the most: prediction enhanced learning, and it did so by a small but significant margin. Once students have seen the results of this initial experiment, the next two experiments I tell them about will help confirm the knowledge that they have now lodged more firmly in their brains.

If your teaching routine includes describing surveys or experiments for your students, consider varying the typical pattern of describing the setup and then showing the results. Instead describe the setup and ask students to predict the results. After you have shown them the results, invite them to reflect upon the accuracy or inaccuracy of their predictions: Why did they get it wrong or right? What did they learn when the results were revealed?

Closing Predictions

Predictions can close a class as easily as they can open them, but in the case of closing predictions you are pointing students toward the material that they will be reading or studying for homework. Many textbooks include prediction-style questions at the beginning of a chapter; I suspect that few students read or think about those questions unless they are specifically required to do so. If you have spent a class period finishing up Chapter Five in your textbook, you might end class by asking students to answer one of the prediction questions at the beginning of Chapter Six. When they are sitting down to read Chapter Six that evening or the next day, and spot that question again, they will recall the predictions they made in class, which should (re)-activate their prior knowledge.

You can push this activity one step further by asking students to revisit a closing prediction question in the opening minutes of the next class and reflect on whether or not they got it right and why they did (or didn't). To continue with the example from the previous section, let's assume I had presented the three experiments on prediction and learning in my class period on Monday. But I had reserved a fourth case, one that introduces a new wrinkle to the research, for the end of class. In the closing minutes of the period I describe the setup of this last experiment and ask them to predict the results. But this time I send them away without revealing what happened, and leave them with a teaser: "You'll find the results of this experiment in your reading for Wednesday's class." When students return on Wednesday, I can ask a few students to remind me what they had predicted on Monday, to describe the results they read about, and to account for the difference: What principle or new idea explains the discrepancy between the first three experiments and the fourth one?

PRINCIPLES

Asking students to make predictions requires a very small investment of time, which makes predicting an ideal small teaching activity. The following principles can help guide the creation of prediction activities in your classroom.

Stay Conceptual. Remember that part of the reason predictions work is that they require students to draw up whatever knowledge they might have that will assist them in making their prediction. If you ask them questions that are so specific that they have no prior knowledge to activate, you won't see this benefit. It seems unlikely, for example, that asking students to predict the meaning of a word in a language with which they are totally unfamiliar, in a different alphabet, will offer much learning benefit. Focus prediction activities on the major conceptual material that will maximize their learning in the course.

Provide Fast Feedback. Close the loop on every prediction your students make by providing feedback as immediately as possible. Predictions made at the opening of a class session should be addressed within that class session. Those made at the end, even if they will be answered by the reading or studying that they will do for the next session, should still be reviewed at the opening of that next session. Predictions made in online environments should provide feedback within the same session they are made. Remember that you don't want wrong predictions hanging around in students' heads for very long; the more immediate the correction, the better.

Induce Reflection. As Daniel Willingham has argued, "Memory is the residue of thought" (Willingham 2014, p. 54). In other words, we remember what we spend a little time thinking about. Prediction provides an excellent spur for thought, in that you can ask students to think about why they made their prediction, what actually happened (if the prediction leads to direct observation), and why their prediction was right or wrong. If you are asking students at the beginning of a class to write down a prediction and having a few of them read their predictions aloud, return to those students at the end of class and ask them to explain why they made those predictions. Students who made correct predictions can be asked to articulate the principle or concept that helped them get it right; students who made incorrect predictions can repair their understanding by articulating the correct ideas.

SMALL TEACHING QUICK TIPS: PREDICTING

Predictive activities are the ideal starting point for a small teaching approach, because they are so easy to slot into the opening and closing of a learning unit. These reliable prediction activities give you some practical starting points.

- Open the course with a class brainstorming activity in which you ask students to surface their prior knowledge. Give them time individually or in small groups, and then work as a class to organize that knowledge in ways that will set up the learning to come.
- At the beginning of the class, unit, or course, give students a brief pretest on the material. For example, give an opening-week pretest that is similar in format to the final exam.
- Use classroom polling to break up your course lectures and ask students questions about the next topic you will cover. Pose the question, have them respond with whatever poll technology you are using, and then invite some paired or whole-class discussion of their responses before you move into your explanation.
- When presenting cases, problems, examples, or histories, stop before the conclusion and ask students to predict the outcome. Invite them to reflect afterward on why they got it wrong or right. In other words: pause, predict, ponder.
- Close class by asking students to make predictions about material that will be covered in the next class session.

CONCLUSION

As I was finishing the second edition of this book, an article appeared in a psychology journal which provided another plank of support for the power of prediction, as well as one more reason why prediction has such a positive impact (Brod 2021). The article describes an experiment in which the researchers asked students to predict the scores of a soccer game, then revealed the correct scores to them, and afterward tested their memory of those scores. But they compared the students who engaged in these predictive activities to another group of students who did something slightly different: they were shown the correct scores, and then asked to explain whether those scores matched their expectations about them. In other words, the students in that second group were essentially being asked: If you *had* made a prediction, would it have been correct? As you might expect from the research I have provided throughout this chapter, the students who made the actual predictions had better memory for the soccer games than did the students who only reflected afterward upon how well the scores matched their expectations. The concrete act of making the prediction was the crucial differentiator: they had to use their prior knowledge of soccer to hazard a prediction, and then see the results afterward.

The researchers in this study added one interesting element to their experiment, though. They looked at the pupil dilation of the subjects as the answer was revealed to them. When the subjects who had made predictions saw the correct answers, their pupils dilated, indicating surprise. The researchers thus argue that one of the driving mechanisms behind the learning power of prediction is the quick emotional burst we get from seeing how our prediction turns out. As we shall discuss further in Chapter Eight, emotions can play a key role in enhancing our learning. In this experiment, the predictions that the subjects made about the soccer games heightened their emotional investment in finding out the scores, and this emotional investment paid dividends in their learning. This finding on the emotional impact of prediction led the researchers to conclude that random guessing doesn't have the same impact as prediction. When I make a random guess, I'm aware of its randomness, and don't make an emotional commitment to it. When I make a prediction, I'm drawing on my prior knowledge to try to understand a novel situation, and as a result I become more emotionally invested in the outcome: I want to know if I was right. It's that emotional investment, the researchers argue, that makes the difference, and that helps lodge the new knowledge firmly in my brain.

Keep this finding in mind as you are engineering predictive activities for your students. If you give them a pretest, encourage them to draw on their previous knowledge to try and get as many correct as possible. If you are asking students to predict something in class with electronic polling, actually give them the time they would need to think about it and make that prediction. The more they commit themselves to their predictions, the more emotion they will feel at the revelation of the correct answer—and the more they are likely to remember it in the future.