

Introduction

I BEGAN WORKING on school improvement some thirty years ago in Chicago. Back then, dropout rates in the Chicago Public Schools (CPS) hovered near 50 percent, less than 30 percent of students were at grade level in reading and math, school buildings were in decay, and basic concerns about issues of order and safety were widespread.¹ The district served 90 percent children of color and 90 percent were from low-income families. In this context, improving schools meant a direct attack on institutionalized inequity by race and poverty.

Fast-forward to today: we now know that dramatic improvements have occurred in the city's schools. High school graduation rates are approaching 80 percent.² CPS students now achieve some of largest learning gains per year of instruction in the nation.³ While these improvements may be exceptional, there are indications that public schools all across the country are gradually getting better. Nationally, high school graduation rates, for example, are at their highest level ever.⁴ Academic standards have risen substantially.⁵ Educators are working hard at trying to get better, and yet dissatisfaction with the state of public education persists.

The problem is that America's aspirations for what we want schools to accomplish today are increasing at a very fast rate. Educators now seek to embrace the goals of the Common Core State Standards and Next Generation Science Standards. They are exhorted to personalize instruction, to promote student social-emotional development, to make all students

technology-savvy. Educators are pressed to prepare every student for a future economy that is evolving rapidly and where the speed of technological change is outstripping the capacity of most US social and political systems to keep pace. And accompanying these larger macroeconomic disruptions, schools are often on the front line to redress an ever-expanding list of social concerns.

Given this backdrop, it is not surprising that a chasm has been widening for some time between our rising aspirations for what we would like schools to accomplish and what they are able to routinely achieve. And this gap is greatest for our most disadvantaged students and in our most troubled community contexts. This has come to be a major social justice issue for our time.

Moreover, there are indications that the economic, technological, and social forces that have been driving the growing gap between aspirations and delivery will continue to compound. This means that the expectations we hold for our educational institutions will continue to change, and likely at an increasing rate. So not only do our schools need to keep getting better, educators need to learn how to accomplish improvements faster in the years ahead.

This challenge is the organizing perspective that my colleagues and I brought to *Learning to Improve—How is it that America's schools can get better at getting better?*⁶ Put simply, in that volume we argued that we need answers to two big questions:

- How can our educational organizations *continuously get better at what they do?*
And recognizing the quickening pace of change,
- How do we *accelerate learning to improve?*

In this volume, I, with contributing authors, offer six stories of educational organizations that have stepped up to the challenges of “getting better at getting better.” These organizations span the nation from north to south and east to west; their work is alive in urban and rural contexts. *Improvement in Action* brings forward the visions, aspirations, and concrete moves that these educators have made to vitalize improvement.

These stories *inform how* other educators might embark on similar activities. They also *inspire why* working in these pioneering ways really matter for the children, the families, and the communities that educators aim to serve.

Before sharing the six stories, I recount below a bit of recent educational reform history. It sets the context for where the discipline of improvement science and working in structured improvement networks takes root.

AN ASPIRATION IN SEARCH OF A METHOD

In response to the forces of change noted above, policy leaders and reform advocates have been pressing an increasing number of new ideas into America's educational systems. Unfortunately, these ideas often arrive without the expertise and practical know-how to make them work. Reformers and policy makers then get frustrated that educators cannot seem to turn their "great" ideas into effective execution; and the commitment of educators gets called into question.⁷

W. Edwards Deming, the improvement guru whose writings and research have inspired productive change across many different industries and sectors, said that education is a field characterized by "miracle goals and no methods."⁸ He made this observation in 1991, as policy leaders touted national education goals that pledged, among other things, for the US to be first in the world in math and science by 2000. Deming's characterization of "miracle goals" was apt then, is equally true looking back now on No Child Left Behind (NCLB), and remains salient as we move into the era of Common Core Standards.

Performance Management

Important initiatives have been developing over the last two decades that seek to tackle Deming's methods question. One such effort is the performance management strategy embraced by NCLB and now found in many school and teacher evaluation systems. A strength of this approach is that it has focused educators' attention on data. As a result, there is much greater transparency today around student outcomes, including disparities in

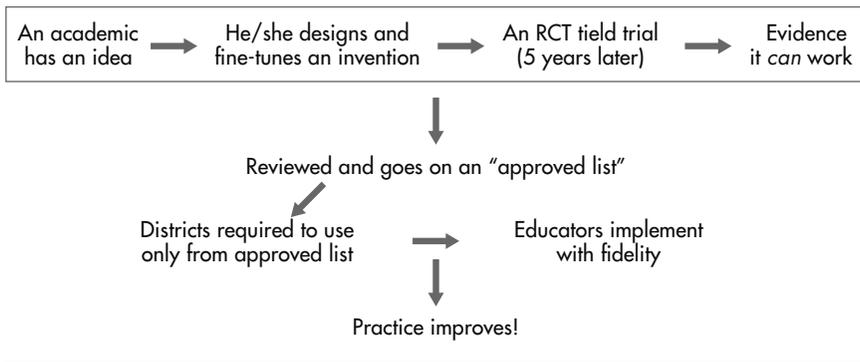
these outcomes among groups. This transparency is a new, important phenomenon. I think back to one of my first efforts in Chicago in the late 1980s to analyze systemwide student achievement data broken down by racial-ethnic subgroups. Everyone knew there were large disparities in student attainment, but back then there was little published data on this point. As our tabulations began to shed light on this issue, we initiated conversations with system and community leaders about making these data public. Some of them feared, however, that our findings would be read as blaming the victim and cautioned against releasing them.⁹ In the end, the need for data transparency won out, but it took an argument to accomplish this. Today, the commitment to transparency about results is settled, and this development is central in allowing improvement science to spring forward in the education context.

In somewhat simplified form, the basic logic operating in performance management systems is as follows: policy makers set the targets, create incentives, provide schools with data dashboards, and then hold individuals accountable for reaching the established targets. This is a plausible strategy if you think that the missing ingredient is effort, or the problem is that educators are just not focused on the right objectives. However, if you think that improving outcomes is principally a problem of educator learning—of not knowing how to do better—then performance management as the primary strategy is lacking in a critical regard. Specifically, there is no detailed working theory embedded in this strategy about *how* to actually improve practice. Instead, educators are told the goals; they are told to work harder and figure out how to reach them; and they are put on notice that failing has consequences for them and their schools.

The Evidence-Based Practice Movement

Evidence-based practice represents a second response to Deming's methods problem (see figure I.1). This movement has been prevalent in medicine for over fifty years, and it began to take hold in education in the early 2000s. In the ideal scenario, a researcher in an academic institution—or perhaps a commercial educational products firm—has some research-based ideas about new educational materials, programs, or tools. Systematic design and

FIGURE I.1 A process map for evidence-based practice



initial testing occur in a controlled context orchestrated by the researcher, and a proof of concept emerges. A small field trial, which is also directed by the researcher, takes place; this is followed up with a larger, often multiyear randomized control trial (RCT). The final results are eventually reviewed by the What Works Clearinghouse (WWC—part of the US Department of Education’s Institute of Education Sciences (IES)). If the intervention is judged successful, then typically seven or eight years after the initial research inquiries began, it is added to a list of effective programs that districts are encouraged to buy and implement with fidelity.

The strength of this evidence-based movement is that it has brought increased conceptual and analytic discipline into education. This is a much-needed antidote in a field where supporting evidence has often been an afterthought. This attention to bringing empirical evidence to bear, and to drawing on more basic social and psychological research findings to undergird change, is another grounding tenet of improvement science.

But the primary focus in the evidence-based practice movement is on estimating a program’s *standard effect size*. Basically this is just the *average difference* between a group of students, classrooms, or schools that received the intervention and another group that did not. At present, the focus on average outcomes is so deeply embedded in policy and practice contexts that educators tend to lose sight of what this perspective obscures. Even

the large-scale RCTs, which sit at the apex of the four tiers of evidence for educational interventions, as defined in the Every Student Succeeds Act (ESSA), are about estimating average differences in outcomes.¹⁰ Advocates use these results as evidence that a program “works,” but it is more precise to say that the results are evidence that a program *can* work. That is, assuming a positive effect was found, then the program presumably had to work somewhere for some students for this average difference to emerge. But it is generally not possible to tease out for whom the program worked, nor in what contexts and under what conditions.

In addition, the technical term of a *standard effect size* is easily misunderstood. In broader public use, it is often construed into a more general claim, leading lay audiences to think that the standard effect size is what routinely occurs and what should reliably result—that is, the standard—but it does not mean that.

There is now a growing body of evidence that almost all interventions actually have highly variable effects, working well in some places and not at all in others.¹¹ Researchers report that the standard deviation in the size of an intervention’s effects is as large as, and sometimes even larger than, the average effect. This means, again, that while these studies provide rigorous causal inferences that the intervention *can* work (i.e., the statistical evidence about the average difference), these interventions often have not worked in a quarter to a third or more of the sites where they have been attempted. Moreover, there are reasons to believe that the results, which occurred under the controlled conditions of an RCT, likely underestimate the propensity for the null effects that will occur if the intervention were to scale more broadly.¹²

In essence, the education field is relearning observations that date back to the 1970s—and are documented in the evaluations of programs such as Planned Variation Head Start, Follow Through models, and the television series *Sesame Street*. Interventions produce wide variability in outcomes. Moreover, this variability in effects is especially important when the aim is to address long-standing disparities in educational outcomes, because these null effects tend not to be randomly distributed. Rather, the most disadvantaged students and contexts are often the ones left behind.¹³

Finally, at the most basic level, the overall corpus of work in the WWC remains modest. This contrasts with health care, where clinical trials have been supported at multiple orders of magnitude greater than in education, and where the knowledge base has been building for over half a century.¹⁴ To put this difference in perspective, in 2010 alone, some twenty-seven thousand clinical trials were published in medicine. Contrast this with education, where after more than a decade of effort, there are about one thousand entries in the WWC. In short, the overall body of program effectiveness evidence available to educators is very limited. These studies may afford a useful starting point for improvement for some educators on some problems, and in some contexts; but realistically, if our aim is to close the aspirations chasm described earlier, this evidence base will not take us very far anytime soon. We need to look beyond this for improvement guidance.

School-Based Learning Communities

Forming school-based learning communities represents a third strategy for advancing improvement. In some ways, this approach represents the polar opposite of the evidence-based practice movement. As noted above, field trials tend to focus on evaluating programs such as a new curriculum, a pedagogic practice, or a technology. Since the emphasis is placed on average results, the influences of local contexts fade into the background. Communities of practice, in contrast, take local context very seriously. They direct attention to the day-to-day problems that educators experience in their own classrooms, schools, and districts.

A strength of this strategy is that it directly engages practitioners as problem solvers—identifying issues they care about addressing—and recognizes them as active agents of change. Learning communities break down the work silos and norms of private practice that have long characterized teaching.¹⁵ In their stead, collaborative coaching, walkthrough processes, and collective inquiry cycles have been building in many schools and districts since the 1990s. These developments, like performance management and evidence-based practice, represent real progress. And they reinforce another core theme in improvement science—that those engaged in the work are central to its improvement.

But the research-base and empirical grounding of communities of practice is often unclear and sometimes absent. Most important, these communities have weak mechanisms for articulating what they are learning, and then testing and refining it under more diverse conditions. As a result, what gets learned in many such communities tends to live and die with the individuals who have learned it. There is no established mechanism for developing a practical knowledge base that might more readily support improvements at larger scale.¹⁶

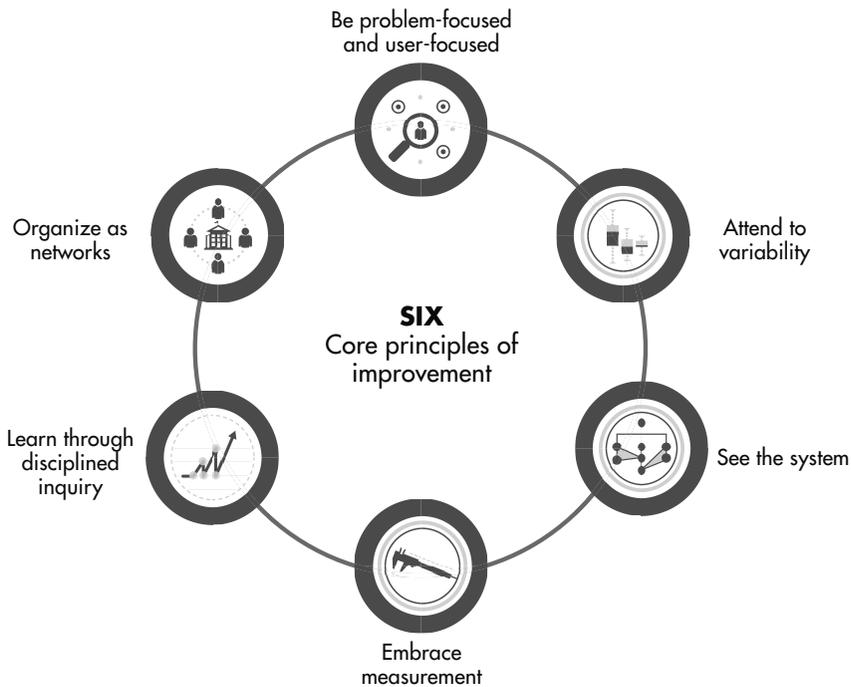
To sum up, then, each of these three approaches has distinctive strengths, but none alone or even in combination will help to close the aspirations chasm. Instead, as educators, we need to draw on the best of what these strategies offer and join to them the best of what has been learned about improvement in other sectors and by pioneers in our own field, some of whose stories are recounted here. Moving toward such a synthesis is where the improvement paradigm, introduced in *Learning to Improve*, took its inspiration.

A BRIEF REVIEW OF THE IMPROVEMENT PARADIGM

The paradigm shared in *Learning to Improve* is organized around the six principles highlighted in figure I.2. The first principle—*to be problem-centered and user-specific*—is anchored in a deceptively simple question: “What is the specific problem or problems you are trying to solve?” The critical word here is *specific*. Educators typically know what outcomes they want, but often do not know exactly what they need to change to achieve them.

In the past, as concerns arose about some educational problem, educators often moved quickly to implement reforms. Typically they did so by drawing on a standard set of solutions such as adding a new curriculum, offering more professional development, hiring more staff, or introducing another new program. In contrast to just jumping on solutions like this, improving organizations take time to analyze and understand better the actual problems they have to solve. And they go about this analysis of root causes in a very specific way—by being user-centered. Improvers try to

FIGURE 1.2 The six principles that form the improvement paradigm



see the problems through the eyes of the students whom they serve and the adults who work with and for them: “What are those people actually experiencing and how do they make sense of the environment that embeds their work?”

This examination of the root causes of specific problems leads to the second principle: *Attend to the variability* in the outcomes regularly observed. Education is complex work. Studies demonstrate that whenever organizational complexity exists, wide variability in performance typically accompanies it.¹⁷ As noted above, addressing this concern requires moving beyond knowing that something can work on average to learning *how* to achieve improved outcomes *reliably* for different subgroups of students and their teachers, and in the many different organizational and community contexts in which they work.

This is part of why the improvement paradigm is especially well suited for attacking disparities in educational outcomes. The focus on variability in performance, and the sources creating it, is at the center of the work. This stands in stark contrast to all that remains hidden in standard reports about mean differences and average trends, even when this information is presented for separate racial, ethnic, or linguistic subgroups. These reports may well document the existence of a broad educational concern, but they generally offer little insight as to how one might actually solve it.

This consideration drives improvers to zoom in on seeing how their educational systems create these unsatisfactory outcomes. This brings us to the third principle: *See the system*. This principle directs attention to understanding more fully the underlying processes, norms, and organizational arrangements that produce unacceptable outcomes. Improvement teams then use this systems knowledge to develop a working theory of improvement.

Principles 4 and 5 turn to the empirical heart of the improvement paradigm. First in this regard is to *embrace measurement* (principle 4). As noted earlier, educators today have a lot of student outcome data and consequently a much better handle on the variability in these outcomes than ever before. But these measures tend to be global in content and are generally available only after the work is done. Standardized test scores, for example, arrive after the academic year is over. For this reason, they are referred to as *lagging indicators*. Improvers, in contrast, typically need more timely and finer-grained information that gets down into the actual work processes, the roles people enact and the prevailing norms that shape their work, because, as noted above, this is where the observed outcomes actually take root. Such *leading* indicators provide signals as to what is and isn't working on the ground and help educators to identify where the next round of changes should focus.

Now, on balance, the measurable aims targeted for improvement are often expressed in terms of standard lagging indicators. But even here, to inform improvement efforts, teams may need to organize these data in new ways to help them see more clearly the problem they are trying to solve. For example, rather than only looking at overall year-by-year outcome results for high school graduation, several of the teams whose stories

we tell tracked longitudinally the progress of individual students and subgroups of students toward graduation over time. They also created new measures explicitly focused on the specific changes in practice that they were trying to effect.¹⁸ Developing these practice-based measures is one of the ways that improvement science builds on the spirit of performance management, but goes well beyond it.

Likewise, as is characteristic of evidence-based practice, improvement teams *learn through disciplined inquiry* (principle 5). At the core of improvement research are rapid iterative cycles of testing possible changes against data, and then typically revising, retesting, and refining these change ideas to work well across diverse conditions. Throughout these cycles, improvers are constantly questioning: “Are the proposed changes actually happening as envisioned? Are we seeing the immediate process outcomes expected from these changes? And where process improvements are occurring, is there evidence that this is actually moving us toward the measurable aims we seek?” The kind of causal thinking embedded here will often lead improvers to stepping back a bit to ask still other, often more fundamental, questions: “What are we assuming that we may need to revisit? What else do these results suggest we might need to work on?”

Improvers are opportunistic and pragmatic in the design of their inquiries. The plan for improvement cycles may take many different forms, ranging from lesson study protocols, to small RCTs, to interrupted time series analysis (often visualized as run charts), to qualitative field observations, and more. At base is the question, “How can I efficiently learn about the problem at hand and whether the change I am introducing is actually an improvement?” In this regard, the method of inquiry is simply a servant to improvement learning.

As will be seen in several of the narratives that follow, improvement cycles often aim to create good routines that scaffold how educators might better do their work. Importantly, these routines are being developed with, by, and for educators. In this regard, improvement activity mirrors the social dynamic of school-based learning communities. In moving beyond this local learning, as new protocols take on the status of tested routines—things we know *can* work—improvement teams recognize that

these routines are likely to need some adaptation in order to be taken up and used well by other colleagues and in other settings. So the process of spreading an effective change idea is itself a “learning to improve” problem.

Equally important, research-based evidence gets woven throughout the processes of continuous improvement. This is where the conceptual frameworks and statistical methods that develop from more basic research add value to how these practice-based efforts are carried out. Cast as improvers, practitioners and researchers bring distinctive knowledge and expertise to the table and collaborate in practical problem solving. This is very different from the traditional models that see researchers as the expert and educators as the implementers of their external knowledge.

Finally, as educators tackle larger and more difficult problems, they *organize as improvement networks* (principle 6). These are necessary because contemporary educational systems, and many of the problems embedded within them, are now so complex that few educators can solve them alone. Addressing these concerns requires coordinated, collective action. Whether this networked activity occurs within a single organization or in a deliberately structured cross-organizational arrangement, certain commonalities are evident:

- Each network is organized around a specific problem.
- Participants’ efforts are disciplined by a shared working theory of improvement that is anchored in seeing how their system generates the unsatisfactory outcomes targeted for change.
- The network develops and uses common measures and inquiry tools so that members can learn together and quickly share whether the changes they are testing are actually moving their system in the right direction.
- These intentionally formed networks draw together diverse sources of expertise from practitioners, school leaders, researchers, designers, technologists, and, depending on the problem, often families and students as well.
- Participants are deliberate in consolidating the practical knowledge developing from their efforts and making it quickly accessible to still

others who may further refine it as they seek to integrate it effectively in their particular contexts.

These are the core mechanisms through which a networked improvement community accelerates learning to improve.

Embedded in the improvement paradigm is a key shift from telling educators *what* to do to fostering a social learning environment where everyone is learning about the *how* and *why* of making schools work better. Creating broad agency for improvement is a core idea in Deming's work on how organizations continuously improve. Writing seventy years ago, Deming understood the dysfunction that was rampant in industry—where a small number of people thought about the design of the work while the vast majority labored as replaceable parts. In Deming's mind, when corporations disregarded their workers, they were throwing away most of the social intelligence and goodwill resident in their organizations that could advance improvement. Such behavior was both morally unconscionable and would ultimately undermine corporate goals. Deming's insights remain equally salient today as we seek to improve our nation's schools.

SINCE THE PUBLICATION OF *LEARNING TO IMPROVE*

In writing *Learning to Improve*, my coauthors and I sought to bring together a compelling set of ideas, with accompanying tools and practices, about how educators might engage in new and different ways to make real progress. The volume drew heavily on our early experiences at the Carnegie Foundation in initiating two networked improvement communities (NICs). Three big questions motivated the launch of these NICs:

- What would it take to move from a compelling set of ideas—the six improvement principles—to execution?
- Would anyone, anywhere, agree to work with us in these new and different ways?
- If they did, could we assemble a body of evidence that such work truly advanced measurable improvements?

We sought to learn fast about how, and whether, we might be able to move this set of promising ideas into practical action.

The first, and by far the larger NIC, was the Carnegie Math Pathways (CMP); the second was the Building a Teaching Effectiveness Network (BTEN). In each case, we tackled a problem that had previously seemed insurmountable: the abysmal failure rates in developmental math instruction in community colleges and the low rate of retention among new and early career teachers, respectively. Both were high-leverage problems. Developmental math courses in community colleges operated as the gatekeeper to opportunity for large numbers of students, especially students of color and from poverty. As for new teachers, they often encountered an incoherent and largely dysfunctional human resource development system that produced the same outcomes year after year—predictable failure rates and high churn, cohort after cohort. Getting new teachers better, faster and holding on to them could change student outcomes for a generation to come.

In terms of our own internal learning goals, the Carnegie Math Pathways provided us with an entry into working with ideas about a deliberately structured improvement community. In a complementary fashion, BTEN was our first test bed for engaging improvement science on the ground in schools.

These two initiatives created the base of practical experience that we shared in *Learning to Improve*. As we concluded that book, we noted that the idea of continuous improvement might be easy to say but harder to do (well). Subsequently, we have now heard this same sentiment expressed by others who have become excited about these principles, tools, and methods and have shared with us their earlier experiences in trying to make them work. So the rationale for *Improvement in Action* emerged—to make more visible how those seriously engaged in improvement actually do the work.

When *Learning to Improve* was first published, staff at the Carnegie Foundation for the Advancement of Teaching knew that there was interest in the field to learn more, but we did not anticipate the intensity of this interest, nor how quickly it would grow. Five years later, we are excited that efforts at continuous improvement and working in improvement networks

have moved to the center of educational reform. The Every Student Succeeds Act, for example, has devolved significant resources and authorities to states and districts with strong guidance to be evidence-based and to engage in continuous improvement. Philanthropic institutions are committing increasing resources to initiating and growing improvement networks around targeted educational problems. Participation in Carnegie's annual Improvement Summit on Education expands each year, and the project proposals submitted for consideration continue to strengthen in depth of detail.

In short, a movement is gaining force, but it has also surfaced a critical new question: "What does it actually take to move this promising set of ideas into quality execution?" As noted above, the results of our own first efforts were encouraging, but this just raised a next question: "Could other educators working on different problems and in varied contexts make the improvement principles work for them as well?" Realistically, we knew that variation in the use of the improvement principles, tools, and methods was likely to manifest as engagements in continuous improvement spread. Some groups might make good use of these resources but others would take them in less effective directions. So a concern for *quality* in the enactment of improvement came into view.

A FOCUS ON *QUALITY* IN CONTINUOUS IMPROVEMENT

Over the three decades that I have been involved in school improvement efforts, I have seen a lot of good ideas, and a few not-so-good ones, come and go. Throughout, I have remained hopeful about the possibilities for improving our nation's schools. But I am also a critical realist, and my realist voice reminds me that the history of educational reform is replete with good ideas that got taken up in superficial ways.¹⁹ Leaders often embrace these new ideas precisely because they are new and doing so casts them as reformers. A frenzied spate of activity typically follows, but productive improvements prove harder to find. As the next leader takes the reins of the organization, that person champions other new ideas, and the process

just continues to churn. The academic literature refers to this phenomenon as a form of *ritualized rationality*.²⁰ As a very experienced and somewhat haggard Chicago teacher once told me as she rolled her eyes, “This, too, shall pass.”

I worry that this same phenomenon might torpedo the current embrace of continuous improvement. Educators are increasingly using terms like *improvement science* and *networked improvement communities*. They tell me they are implementing new routines like root cause analysis and conducting Plan-Do-Study-Act (PDSA) cycles, but sometimes it is not clear that the fundamental changes in thinking and action that the improvement paradigm demands are actually taking place. Let me offer a couple of examples to make this more concrete.

Over the last few years, I have had opportunities to interact with a number of different groups who have heard about improvement science and networked improvement communities. Sometimes I am told, “Oh, we’re already doing that.”

“Really?” I ask.

“Oh yes, at the end of every year, we look at our accountability data and we talk together as a faculty about it and then we figure out some new things that we might try next year.”

In short, the annual school improvement plan has become “continuous improvement”—a new label for an old practice.

In other places, I have witnessed constructive first steps taken toward continuous improvement. Professional development programs focused on improvement have been added either at the district level or through some regional intermediary, such as a county office. These sessions introduce educators to the basic tools and methods of improvement science. Teachers then are encouraged to do PDSA cycles, but then often left largely on their own to figure out exactly how to design and carry out these cycles. More problematic, the source of change ideas may be parochial and the research base doubtful. Additionally, the data necessary to evaluate whether a changes is actually an improvement may be missing altogether, or take the form of, “Well, what does your gut say?” Other key shifts in practice, such

as bringing the users' perspective into the work and identifying the specific sources of variation in performance that need to be addressed, may also be weak or absent.

Likewise, organizational leaders supporting these efforts may not recognize how the shift into improvement affects their own work. They continue to view their responsibilities as providing services to schools and districts rather than operating as full learning partners in the improvement efforts. I have seen administrators counting the number of PDSAs that have been completed in schools and encouraging more activity. How all of this local activity consolidates into collective progress, however, remains unaddressed. Likewise, the role of improvement leadership, as orchestrating a set of social learning processes, has not yet been activated. Instead, improvement efforts look a bit like performance management: administrators offer training, monitor activity levels, and then hold others accountable for change. The focus isn't on what is being learned, but rather on how much is being done.

On balance, the transformation for any organization into a continuously improving one is no small endeavor, and the tendency to fall back on old ways of thinking and doing is understandable, and even to be expected. The six improvement principles represent a fundamental change and a major challenge to the ways of thinking and acting for both practicing educators and for those who partner with them, including educational researchers and commercial vendors. Helping to illuminate this shift into *quality* in continuous improvement is the central purpose of the cases that follow in this volume. In these stories, I use the term *continuous improvement* in reference to a broad range of activity aiming to advance educational outcomes through iterative cycles of change over time. I distinguish this term from the *improvement paradigm*, which constitutes a transformation in the ways educators think and act as they advance such activity. The improvement paradigm is about living the six principles, about bringing agency to frontline educators as improvers while also fusing research knowledge, and the expertise of researchers and others, more productively into these processes.

INTRODUCING THE SIX STORIES

Learning to Improve offered a vision of practitioners and researchers fusing their domains of expertise to make a difference for all students. This volume, offered five years later, is testimony to moving beyond a vision. The chapters that follow provide accounts of a diverse set of improvement efforts in action. The stories open up some of the many ways that the six improvement principles are being used by different educators in different organizational settings and with different resources, but in each instance to pursue ambitious change efforts. Although each of the organizations used somewhat different labels for the roles, routines, and structures that they put in place to orchestrate their change efforts, for simplicity, I refer to all of these developments as the work of an *improvement hub*.

At one level, the chapters are a collage of improvement narratives. Each describes impressive work in a different educational context. Two detail the efforts of traditional school districts, Fresno Unified School District and Menomonee Falls. Two other chapters tell the stories of highly innovative charter management organizations, Summit Public Schools and High Tech High. And two more are accounts of intermediate organizations that work with large networks of schools, New Visions for Public Schools and the National Writing Project.

Chapter 1 chronicles the improvement efforts of Fresno Unified School District, located in California's Central Valley. This chapter introduces the metaphor of a leaky pipeline as a heuristic device, not only for the work in Fresno, but also for the efforts of New Vision for Public Schools, Summit Public Schools, and High Tech High (chapters 2, 3, and 4, respectively). At multiple steps in the journey from high school entry to matriculation into college, Fresno students were falling through the cracks. Improving high school graduation and college-going rates required Fresno educators to carefully orchestrate a series of efforts at every juncture where the pipeline leaked. The Fresno team kept their eyes on seeing the system—on understanding how a set of processes that students experience as they move through high school created the overall outcomes they were observing. Their problem-solving took them into thinking about data in new ways,

creating new data tools and processes for their use, and putting in place the staffing and professional development supports necessary for practitioners to turn useable evidence into productive action. The Fresno story also introduces a theme that will be repeated in subsequent accounts: how insights from relevant research and the inclusion of researchers as partners can inform productive change.

The New Vision for Public Schools case presented in chapter 2 introduces the organization and role of an analytic hub of a networked improvement community. New Visions' efforts took the organization deeply into how basic administrative systems around student absenteeism, course-taking, grades, credits earned, and, in New York, the number of Regents exams passed, contributed to the graduation problem they wanted to solve. On the surface, the functioning of routine administrative data systems like these might seem mundane and uninteresting. But the breakthrough at New Visions was to realize that embedded within the operations of these systems was a complex array of taken-for-granted ways of working that needed to be understood more deeply, and then challenged, to advance better outcomes. Like Fresno's leaders, the hub at New Visions focused on building good data systems and tools that aimed to provide the right information to the right people and at the right time. Hub staff collaborated with educators in their network of schools to field-test and refine a key change idea, that they eventually came to call the Strategic Data Check-Ins process. The breakthroughs affected by New Visions are now moving out across the entirety of the New York City Public School System.

The portrait of Summit Public Schools in chapter 3 opens up a window into what it means for a system of schools to function as a learning organization. Started some fifteen years ago as a single school, Summit has grown into a charter management organization that now operates a network of schools in California and Washington. As part of this organizational growth, Summit embraced an extraordinarily innovative agenda for high school reform. It aimed to give students more choice in the work they do and to move at their own pace toward well-articulated learning outcomes. In Summit's vision of a good school, teachers function as mentors and advisers supporting students to take charge of their own learning. Figuring

out how to actualize these aspirations—every day, for every student, and in every learning context—proved a huge challenge. And it was this challenge that led Summit to develop its own internal capacity to undertake and advance continuous improvement. The organization invested in building improvement capabilities among school staff, in assembling expertise in a central hub to facilitate and support these activities, and in building analytic capacity to inform and guide these efforts. Summit also drew on research and nurtured relationships with applied researchers whose knowledge and skills could help them on their journey to improve.

High Tech High (HTH), the focus of chapter 4, offers a story of improvement science going deep in a most unexpected place. The progressive educators who founded HTH harbored a fundamental skepticism about data, and this sat alongside a strong commitment to individual teacher autonomy. These are not conditions in which one might expect to find a ready embrace of either improvement science or organizing as improvement networks. HTH educators, however, recognized that even though everyone was working hard, their organization was coming up short on its core aspiration: preparing every one of its students to graduate high school and succeed in college. This recognition propelled them to explore new ways of trying to get better. Successes in a modest initial improvement project expanded interest among HTH educators to learn more about improvement science and generated momentum for its increased use. An organic process of organizational development emerged with transformative effects. Today, High Tech High operates multiple networked improvement communities across its system of schools. It has developed a strong cadre of teacher leaders who support these improvement efforts, and has woven improvement science into its new graduate school of education. Having reached these milestones, it has also moved the goal line. The aim now is to make this approach normative in the way their educators do their own work and to enable the spread of these practices to a next generation of teachers and other educational professionals.

In the course of learning how to do improvement, each of the four organizations mentioned above developed expertise and considerable practical know-how to advance real progress on improving high school graduation

and college-going rates. These problems have functioned for too long as gatekeepers to opportunity, and especially for students of color and low income. The lessons that we can learn from their stories opens up possibilities for scaling these successes more broadly.

Chapter 5 ventures into the domain of improving classroom instruction. Efforts to engage the principles, tools, and methods of improvement science and organize as instructional improvement communities are more recent than the work cited above. This chapter offers a first look at a distinctive set of challenges that improvers confront as they turn their attention in this direction. The problem that the National Writing Project (NWP) took on was motivated by a new emphasis on argument writing that surfaced as part of the Common Core standards. The NWP story documents how the organization drew on the expertise of consulting teachers and the established relationships of trust that were resident in its national network of local research-practice sites to mobilize and quickly advance measureable improvements in students' writing. To form this networked improvement community, NWP focused on: *what to teach* (i.e., the community developed text sets to enable instruction around the writing of arguments), *how to teach it* (i.e., consulting teachers offered local professional development on core pedagogical practices), and *how to know* whether the teaching and learning was actually improving (i.e., for this purpose a new instrument, the Using Sources Tool, was developed to inform NWP's continuous improvement efforts). The case offers a glimpse into what educators can accomplish in improving instruction when they organize as a professional scientific community solving a problem together.

Finally, the work of the School District of Menomonee Falls (Wisconsin), taken up in chapter 6, is quite different from the others. It is not a story that begins with a specific problem to be addressed, but rather with an ambitious organizational aim—the transformation of a whole school district into a continuous improvement organization. Menomonee's district leaders aimed to reorient the work of everyone in the school system in this direction. In support of this goal, literally every person in the district—teachers, auxiliary staff, operations personnel, board members, the leadership team, and students—was trained in continuous improvement methods, and

improvement projects were launched all across the system. The chapter offers a detailed account of an evolving set of leadership actions carried out over several years—first to catalyze these changes and then to institutionalize them in how the school system does its work.

Interestingly, chapter 6 can be read comparatively with the narratives of Summit Public Schools and High Tech High. Like Menomonee Falls, these other two cases can also be read as stories of transformative leadership, although the courses they followed were quite different. Both Summit and HTH began with a modest improvement project that grew into a gradually expanding set of projects. The success of the initial projects, and the know-how developed and energy generated among those first participants, quickly brought more colleagues into the work. Sustained support from organizational leaders enabled improvement efforts to expand and eventually transform their organizations over time. Taken together, these three cases offer considerable fodder for conversations about the role and actions of leaders in advancing continuous improvement deeply into how educational systems operate.

While each of the six cases tells a different story, important themes repeat across the cases that are worthy of note at the outset. To launch their efforts, educators in all six contexts examined extant data to better understand how their systems operated to produce their current unsatisfactory outcomes. All turned to extant research evidence and drew on external research expertise to both deepen their understandings about how their systems worked, and to identify possible change ideas. And all continuously monitored the progress of their improvement efforts as the basis for fueling subsequent tests of change.

Additionally, each organization committed to developing its hub capabilities to support improvement. They created specialized staff positions and new organizational roles. Hub staff facilitated and supported local improvement teams. They invested broadly in developing improvement knowledge, skills, and dispositions among staff within their respective organizations. Along the way, hubs developed new data resources and processes that sought to inform local change efforts in a more timely fashion.

Lastly, all of these contexts also benefited from stable leadership and supportive institutional governance environments that made it possible for efforts to germinate and grow. In some cases, the organizing leaders were new to their contexts (chapters 1 and 6); in the others, the leadership for transformation had been present for some time and only the ideas were new.

For this book, my colleagues and I worked closely with leaders in each organization to learn about their distinctive improvement journeys. We provide our take on their respective stories as a way to illustrate, amplify, and teach key improvement ideas in action. In some instances, we studied these organizations from outside; in others, we offered direct support as they launched and advanced their improvement efforts. Regardless, we acknowledge that the full range of work occurring in each of these organizations is much more extensive than our accounts, constrained to chapter length and a specific focus, permit.

Although each case stands on its own as an account of action for possible discussion among a group of readers, the chapters together show how the same principle, concept, or tool was applied by different teams and in different educational contexts. Seeing it at work across multiple and diverse contexts creates opportunities for these resources to be understood more flexibly and deeply than any one telling would allow.

Finally, each of the organizations has made real progress toward its improvement aims. It is notable that even given their documented successes, their organizational leaders remain humble about their improvement journeys. They acknowledge that they still have much more to learn and many miles still to go. In this sense, the chapters are dynamic portraits of improvement in action. I am deeply grateful that these six organizations have allowed us to share some of the *what*, *how*, and *why* of their efforts. They afford all of us an opportunity to see the extraordinary in the ordinary day-to-day work of getting better. Now, on to their stories.