

Let History Not Repeat Itself: Overcoming Obstacles to the Common Core's Success

BY JOHN E. CHUBB

» At least twice in recent history, the United States has come together around bold ideas to reform its schools. And twice the ideas failed to move the nation's students substantially forward. Today, the country has seized upon another big idea. Through the Common Core State Standards project, 45 states and the District of Columbia have endorsed academic standards in English language arts and mathematics that raise the bar for student achievement and redefine what it means in the 21st century.

The standards not only establish best-in-the-world expectations for traditional knowledge and skills. They specify the capacities that students need for success in college and careers. They aim for a deeper form of learning that requires analytical, creative, and communication skills, as well as others not well-represented in standards today, either here or abroad. The Common Core project is a potential watershed in U.S. education, uniting most of the country around an ambitious yet practically grounded vision of what schools and students need to accomplish.

But the Common Core is hardly the first such inflection point in recent education reform. In 1983 the federal report *A Nation at Risk* set off a wave of reform in which states raised high school graduation requirements, introduced academic standards, and took other widely recommended measures. Within a decade, policymakers were disappointed in the results and looking for new ideas. In 2002, those ideas came together in another initiative with historic promise. The federal No Child Left Behind Act (NCLB), adopted by a bipartisan coalition, hoped that testing, accountability, and a bit of educational choice would catalyze innovation and

improvement. A decade later, policymakers of both political persuasions are again looking for new ideas.

The Common Core, despite some controversy, is the idea of the day. And it deserves to be. Higher achievement begins with higher expectations; if we can't specify what more we want students to know and be able to do, we shouldn't expect much in the way of progress. Until now, states have made a hash of setting academic standards. Experts rate the standards of very few states an A—and those of most states a C or worse. Once standards are set, few states have been willing to ask students to master a high percentage of them to be recognized as "proficient." Most states have declared two to three times as many students proficient as has the National Assessment of Educational Progress (NAEP), the nation's "report card." The Common Core promises to overturn the status quo: one set of standards that demands serious levels of proficiency, is aligned with modern expectations, and is embraced by most states.

But will it? History is not on its side. In the past, the government's reach has regularly exceeded its grasp. But history need not be repeated. The disappointments of the reforms prompted by *A Nation at Risk* and embodied in NCLB should provide lessons for those working today to make the Common Core a success. Policymakers should be looking hard at past efforts to raise standards and achievement to understand what is likely to go wrong. Unfortunately, that is easier said than done. The architects of NCLB looked at decades of experience implementing the Elementary and Secondary Education Act (ESEA) as they crafted the original's successor. The problem is that the necessary lessons are not always obvious. They sometimes require research and analysis to understand what they really are. Careful thought must be given to the road that this latest initiative must travel before students should be expected to learn any more.



LESSON 1: THE ASSESSMENTS

If there is one thing we have learned about academic standards, it is that they are a necessary but insufficient condition for improving student learning. States with exemplary standards do not have consistently exemplary achievement. Massachusetts has developed some of the strongest academic standards in the nation—high, clear, and deep—and emerged as the highest-achieving state over the last decade, according to NAEP. By contrast, California also developed strong standards, but its achievement has remained well below the national average and improved very little. Brookings Institution Scholar Tom Loveless recently completed a study of the effects of state academic standards and concluded that the Common Core was likely to make little difference for student learning. Massachusetts and California merely illustrate the pattern found across the 50 states.

The reason standards are not consistently associated with achievement is that standards are only the first step in a long, interdependent chain that ends with student learning, but includes many links in between. The next link in the chain is the assessments that measure progress against the standards. Even the clearest standards—and the Common Core standards are clear—do not become concrete until they are translated into tests of student knowledge, assessments that specify what the standards actually require students to know and be able to do. Over the last decade or so, states have not always done the best job of matching assessments to standards. Some states, out of expedience, have adopted assessments off the shelves of test publishers and ended up with tests that measure skills more basic than the aspirations of their standards. When this occurs, standards are effectively lowered: what matters to teachers and students is not what standards say will be assessed; what matters is what *is* assessed.

Even the clearest standards—and the Common Core standards are clear—do not become concrete until they are translated into tests of student knowledge, assessments that specify what the standards actually require students to know and be able to do.

This point is obvious enough. But what is not obvious is exactly *how* assessments should correspond to standards. The conventional wisdom is that the 50 states, scrambling independently to write their own standards and develop their own tests, have consistently chosen lower-level tests, only loosely aligned with standards, simply to get the job done. Common Core states are doing things differently. They have banded together in two coalitions, the Partnership for Assessment of Readiness for College and Careers (PARCC) and Smarter Balanced Assessment Consortium (SBAC), each supported with hundreds of millions of federal dollars, to do the job right—to take the time and go to the expense necessary to write the high-quality assessments that high-quality standards deserve.

Yet there is no systematic research demonstrating that state tests have consistently failed to measure up to their respective standards. There is no systematic research demonstrating exactly what test attributes are most closely associated with higher achievement. In fact, over the last decade many states revised their initial assessments to add constructed response questions, demanding that students do more than answer the much-maligned multiple choice variety. The constructed response questions, which require both long and short answers, asked students to explain their work in math problems, analyze documents, write persuasive essays, and perform other open-ended tasks. The new Common Core tests promise much more of this open-ended form of questioning. But it is



far from clear how important this kind of questioning is for raising achievement.

Open-ended questions are only one attribute that may distinguish assessments that demand deeper learning from those that do not. Another is the balance between fiction and non-fiction in assessing reading comprehension. Here, too, there is a popular critique that state assessments have fallen short of state standards by relying too much on literary text. The Common Core assessments promise to ask students to interpret scientific, historical, technical, and other forms of writing better aligned with modern requirements. This makes a great deal of sense. But, again, there is little evidence that state tests have been skewed toward literary text. Indeed, the evidence is that state tests over the last decade shifted substantially toward the assessment of non-fiction comprehension and analysis skills.

The point here is that research has not yet shown clearly what assessment attributes are associated with gains in student learning. This is true of U.S. and international assessments. Consider math assessments. U.S. assessments in mathematics over the last decade (both state tests and NAEP) have placed more emphasis on problem-solving skills. The NAEP for grade eight includes a large percentage of items that require students to set up a math problem from a real-world situation and then use the proper skills to solve it. The new Common Core assessments appear to promise more of the same. But do such assessments promote deeper learning? International assessments at the same grade level appear to demand less open-ended problem-solving but higher-level math skills, especially algebra. Again, the point is that we do not know what assessment attributes drive the learning that our standards aim to promote.

Lacking answers to these fundamental questions, the nation should be promoting research to understand much better the role that assessments play in

stimulating learning. While we can safely say that assessments that employ multiple choice questions to measure basic skills will not encourage the attainment of high academic standards, we are not equipped to say beyond that what exactly assessments should do. Sophisticated closed-ended response questions can certainly evaluate students' ability to analyze and reason in nuanced ways. This is the stock in trade of the SAT and LSAT, for example. The closed-ended response section of the Advanced Placement Calculus exam also demands that students solve challenging mathematics problems and not just excel at filling in bubble sheets. But standardized exams that require formulaic five-paragraph essays—staples of many state assessments today—are not necessarily a friend of good writing, as rote writing programs help students crank out predictable “proficient” writing samples.

It is essential that the new assessments be supported by research carefully designed to show when and how assessment is really working.

The two Common Core assessment coalitions are staffed with experts who undoubtedly recognize the outstanding questions about assessments. The coalitions are also led by states with their own judgments about how assessment should look; the assessments that each coalition produces will be a product of political compromises and best guesses. They are likely to represent real progress. But we would be kidding ourselves to think that we already know how best to assess the new standards. It is essential that the new assessments be supported by research carefully designed to show when and how assessment is really working.

What might that research look like? First, research might look systematically at current state assessments to understand, all things being equal, what attributes appear to be associated with higher and more rapidly improving student achievement. Second, similar research should be conducted internationally: are there testing characteristics consistently associated with higher and/or lower achieving nations? Third, as the coalitions develop the Common Core assessments, the differences between assessments should be analyzed to understand what, if any, difference they may make for student achievement. Fourth, the coalitions should be encouraged to create alternate forms of their assessments, to see what consequences different approaches might have for gauging high standards. What learning is gleaned from an extended-response question, say, versus sophisticated closed-ended responses?

We want to understand what kinds of assessments most consistently gauge the higher-order thinking skills and important knowledge that the Common Core is all about. And we want assessments that are reliable and valid. Validity, in particular, is about more than accurately measuring the new standards; it is also about what assessments predict. Colleges and law schools use the SAT and LSAT to predict student success in higher education. The new Common Core assessments are supposed to gauge readiness for college and careers. We need to know if they actually do. Research should examine the relationship between scores and proficiency on the new assessments and success rates for entering and *completing* college or becoming gainfully employed. This research will take time, but data on new high school graduates can inform testing programs year by year as experience unfolds. If the Common Core is really to be different, evidence of its predictive validity will be crucial.

LESSON 2: PERFORMANCE LEVELS

Once assessments are developed, the coalitions will establish performance criteria. They will specify what score is required for a student to be considered proficient. They will also specify higher and lower levels of achievement, from below basic and basic to advanced. Experts have crafted reasonably objective procedures for matching scores on assessment items with overall performance categories. Experts can also compare scores on new assessments to scores on established assessments for benchmarks. The Common Core assessments are to be benchmarked against NAEP and international tests. These objective procedures are then likely to come up against subjective judgments: Will the new performance standards place too many students in categories below proficient to be politically acceptable? Will they make previous state standards appear unacceptably low? Will they exacerbate achievement gaps between rich and poor or black and white?

One can only hope that given the spirit of the Common Core—to raise and unify state standards—that the subjective phase of standard-setting will not undermine standards that are objectively correct. Research, coupled with communication, can help keep the standards high. The state of Tennessee successfully raised its performance standards several years ago by waging a pre-emptive communication campaign to help citizens understand that the new higher standards might cut existing estimates of student proficiency in *half*. When the first year of testing found scores

One can only hope that given the spirit of the Common Core—to raise and unify state standards—that the subjective phase of standard-setting will not undermine standards that are objectively correct.

were indeed down, parents and citizens had been well-prepared, and policymakers and schools faced relatively little backlash. Contrast that with states that have tried to raise standards for high school diplomas and then given in to resistance from parents whose kids would not graduate.

We need research aimed at verifying or gauging, first, how well new performance levels actually do compare with national and international benchmarks. Second, and more important, we need research to help states predict how new assessments will change their existing performance measures. And, third, states and support organizations must undertake educational campaigns to help the public appreciate how test scores will change under the Common Core and why these changes should be embraced.

LESSON 3: ACCOUNTABILITY

If proficiency levels are set high, the *next* challenge might seem to be helping students achieve them. But history suggests it is not. While teaching and learning are the ultimate targets of higher standards, they will not become the focus unless students, teachers, and schools understand how they will be responsible for achieving them—and believe their responsibility is fair. All of these details will be spelled out in state accountability frameworks. Unfortunately, this presents an immediate problem. State accountability frameworks are now being overhauled state by state. The Obama administration has granted states waivers from the federal accountability framework that was established

Research conducted before and during the implementation of NCLB provides some guidance for how state policymakers ought to provide accountability for meeting academic standards.

by NCLB. Now the Common Core standards are likely to be introduced into a variety of state accountability systems, which means that the new standards will have differing effects.

Research conducted before and during the implementation of NCLB provides some guidance for how state policymakers ought to provide accountability for meeting academic standards. Several themes are clear:

1. Accountability works. States that adopted accountability measures with consequences for inadequate student outcomes made greater gains and larger reductions in their achievement gaps on NAEP than states that did not implement such measures. Once NCLB was implemented, students in the lowest 10 percent of NAEP scores made larger gains than they had during the period prior to NCLB—an indication that NCLB likely had a positive effect on the achievement of students at the lowest end of the achievement spectrum. These and similar findings suggest that states should adopt some regime of clear outcome expectations, differentiated by subgroup and reinforced with sanctions or consequences for failing to meet them.
2. Research also indicated major flaws in the NCLB accountability regime. The first was the exclusive focus on proficiency. The system did not recognize student achievement above or below the proficiency bar; if students made yearly gains that still left them below proficiency, schools received no credit. If students above the proficiency bar made gains—became “advanced,” for instance—schools also received no credit. And if top students slid backward, schools were not sanctioned, provided these students remained proficient. This system was not only unfair to schools and students—not recognizing progress when it was clearly occurring—it also provided undesirable and unintended incentives for teachers: it encouraged



them to pay less attention to high-achieving students and to students with little chance of reaching proficiency. Research is mixed on exactly how the incentives altered teacher behavior, but there is no doubt about the perversity of the incentives.

3. Research has also pointed a way out of the proficiency obsession. A number of states, with federal approval, have been implementing “growth models” that gauge student achievement year over year and provide accountability for progress along the entire achievement scale, not just at the proficiency bar. A consensus now exists that growth modeling of some sort is a better way to provide accountability for student achievement. But with ESEA not yet reauthorized, there is no official agreement on how growth models should be used. Controversies surround questions such as how much growth should be considered adequate, and how schools can be given credit for above-proficiency progress by top students without compromising attention to low-achieving students.
4. Research has also exposed great differences in the thousands of schools that fail each year to make adequate yearly progress (or AYP) under NCLB. Because NCLB requires progress not only by schools but by student subgroups (low-income, limited English proficient, students with disabilities, etc.) schools have numerous performance measures to hit in order to make AYP overall. Some schools perform poorly for all or most of their students. These are the lowest 5 to 10 percent of public schools, the schools that the Obama administration has consistently targeted for the intervention known as “turnaround.” At least half of all schools that do not make AYP fail to do so only because one or two subgroups have fallen short; relatively few schools are overall failures. In recognition of this, the Department of Education under George W. Bush began giving

states approval to implement differentiated accountability models. Tough measures, like school closure and reconstitution, were reserved for schoolwide failures. Schools that missed AYP for a subgroup or two were given much lighter sanctions. The Obama administration has likewise asked states to propose their own definitions of AYP and sanctions, if they want an NCLB waiver. The idea of differentiated accountability has logical merit. The rub is that research does not point to any principles that might govern a model system with differentiation; state policymakers do not know what is likely to work best. But policymakers do know that schools must feel that they are being treated fairly, and a one-size-fits-all sanctions regime is not going to be effective.

5. While policymakers seem to agree that students should not be stuck in truly poor schools, waiting patiently while authorities try to improve them, there is no agreement on what relief should be provided. Research has not yet pointed the way. For schools that chronically fail to perform—those beyond six or seven years of not making AYP—NCLB calls for restructuring. But with nearly a decade of experience, all we really know for sure is that most schools, given a choice by the federal government, have “restructured” by doing as little as possible. Yet where tougher measures have been adopted—converting the school to a charter or reconstituting the whole staff, for example—there is little evidence that they have helped. For schools with fewer years of failure, NCLB offers two innovative options—students can choose to attend a different, non-failing public school or receive private tutoring. Neither of these options has been shown to effect positive change. Failing schools are seldom surrounded with successful educational options, and tutoring is often so short-lived that it lacks the intensity to matter. The Obama administration has tacitly said that while NCLB does not seem to offer solutions

itself, the federal government may never come up with fixes for students and schools. That may be best left to the states. The administration asks states in their waiver applications to propose remedies for schools and students that actually work. This makes sense. But research does not yet indicate what seems to work best for chronically poor-performing schools, or even for those with lesser issues.

Well-designed projects could evaluate and improve the design of the following: growth models, full-spectrum achievement incentives and consequences, differentiated accountability plans, school turnaround models, remedies for students such as choice and tutoring, and other innovations.

When the Common Core is launched into America's schools in 2014-2015, these accountability questions will either be receiving varied answers as states create their own accountability systems or they will have received a single answer through the reauthorization of ESEA. Either way, the answers would benefit from timely and rigorous research. Because many states now have waivers from NCLB, there is considerable experimentation taking place. That experimentation should be the target of formal evaluations to determine what, if any, practice is working. In addition, the NCLB experience now provides years of evidence of how different accountability provisions may affect student achievement. Well-designed projects could evaluate and improve the design of the following: growth models, full-spectrum achievement incentives and consequences, differentiated accountability plans, school turnaround models, remedies for students such as choice and tutoring, and other innovations. Let us not forget the new ideas that states are implementing with Race to the Top and i3

grants. Those touch on accountability as well, but they have not been seriously evaluated.

LESSON 4: TEACHERS

Research has helped advance understanding of the teacher's role in student achievement more than any other element of the achievement puzzle. Over the last decade, we have come to widely accept the idea that teachers are the single most important determinant of what students learn in school. Top teachers working with students over just a few years can literally change lives; they can help students move up the norm-referenced achievement scale by 50 national percentiles within three years. Low-performing teachers can have opposite effects, causing students to slip rapidly backward in achievement. Policymakers, who believe these findings are true, are now seeking ways to put top teachers in every classroom.

Research has also shown that this is exceedingly hard to do. Top teaching prospects are hard to identify right out of college. Verbal aptitude and subject matter expertise predict a bit. But certification, training programs, Praxis scores, and other attributes on teacher resumes predict absolutely nothing about success. A teacher's performance, for the most part, cannot be known until it is observed in the classroom. Experience surely helps, but some novices grow into superstars, and others never really succeed. Research has not shown us the conditions under which teachers prosper or fail.

This may be changing. With the advent of value-added measures of teacher success in raising test scores and formally validated qualitative teacher observation tools, we can now measure much more reliably how well teachers are helping students learn. We can also gauge much more reliably how teachers are improving their performance over time. What should be done with this information? Some suggest it be used to reward successful teachers and remove unsuccessful ones:



manage teachers by merit and build a more successful teaching force via smart retention and dismissal.

While better management surely makes sense, it may not be the greatest opportunity for improving teaching. Better measurement provides the chance to understand the process of teacher improvement. What professional development structures facilitate teacher improvement? What training programs have the most effect? What pre-service experiences offer the greatest payoff? What effect do principals have on teacher development, and how can more effective leaders be developed?

Right now teacher quality is a black box. We know what successful teachers look like, but we have little idea how to produce them. Research, supported by vastly better measures of teacher success, could open the box. If policymakers could understand how the best teachers are “made,” they could support the programs and conditions that help great teachers develop.

With the introduction of the Common Core, the role of teachers will become even more important. The new standards demand that students develop levels of understanding and skills of analysis, synthesis, and evaluation that will not be acquired through rote or didactic instruction. That means that teachers will need higher-level skills themselves, and they will have to teach much more interactively. Teachers are going to need much assistance reaching these new heights. But research currently provides almost no guideposts on how to help them do that, and the literature on professional development is almost uniformly negative. Massive studies sponsored by the Institute of Education Sciences have found little evidence that professional development ever works.

It is not an exaggeration to say that just as we need to develop our teachers more, we could not know less how to do so. Yet the opportunity to gain that knowledge is now with us. We should seize that opportunity by undertaking well-designed projects aimed at gauging what forms of training and

Right now teacher quality is a black box. We know what successful teachers look like, but we have little idea how to produce them. Research, supported by vastly better measures of teacher success, could open the box.

professional development produce more successful classroom teachers. A few states have begun to pave the way. Tennessee, for example, has measured the effects of four-year teacher training programs over the last several years and found stunning differences in the performance of the top-ranked and bottom-ranked programs. Teacher-level data now make it possible to conduct nuanced evaluations of all kinds of training and professional development programs. Now, during the brief window before the Common Core launch, teacher preparation—traditional and alternative—and professional development should be rigorously evaluated. So, too, should training programs focused on the Common Core.

LESSON 5: TECHNOLOGY

The launch of the Common Core coincides with what appears to be an inflection point in the adoption of education technology. Schools have been quite slow incorporating information technology into their core instructional programs. Students today, by and large, experience school as they have for generations: in classrooms of 25 students led by a teacher. Curriculum is still largely in print, and teachers use technology mainly to aid traditional lectures. (Think PowerPoint.) Students depend on their teachers very directly to learn; they tend to learn more directly through technology (with online courses or computer-driven instruction) only if they have a special need. Elementary schools, for instance, use computer programs for remediation in reading and math, and



high schools deliver content online to students who need to make up classes or for dropouts seeking an alternative to a degree. High schools also use online delivery for AP courses if enrollment is insufficient to justify a teacher-led class. But core instruction remains largely traditional.

The persistence of the status quo owes partly to political resistance: more technological instruction could mean fewer teachers, an unpopular idea with teachers and unions. But resistance to these new modes of delivery is due also to simple discomfort with technology on the part of teachers and administrators and to technology that has yet to live up to its promise. Yes, students today are digital natives, and they spend countless hours with technology. But technology may still be better for communicating with friends than mastering mathematics.

All of this is quickly changing. Economic pressures, unrelenting since the great recession, are forcing public schools to look for efficiencies, and technology is finding a widening audience. Standardized testing has raised the value of technology that can help teachers and students keep track of ongoing achievement against new standards—the process known as formative assessment. Schools want the capability to manage and interpret volumes of student data.

More than anything, however, schools are beginning to see how technology can be “blended” with traditional instruction to give students a learning experience that makes the best use of both teachers and technology. For example, students might learn a subject by acquiring core content online and practicing skills with an intelligent computer-based tutoring program. Beyond the foundational skills, they would also spend time working with a teacher on research projects that require creativity, analysis, and higher-level thinking. Students might also work with teachers in small groups for extra assistance. All of this work might be done under the

The Common Core is being developed without serious attention to new ways of doing schooling and new ways of incorporating technology.

watchful eye of a teacher managing a class that is substantially larger than what is now the norm.

In the early days of online or computer-based instruction, technology and traditional instruction seemed to be caught up in a zero-sum game. Indeed, full-time online schools took students away from regular public schools—prompting pitched political battles. The future will surely have a place for full-time, K–12 online instruction. But with most parents wanting their children to have experience and supervision outside of the home, and with education having an important social dimension, place-based schooling is likely to dominate indefinitely. Schools and technology firms are now experimenting as never before. What they are seeing is that technology can potentially improve place-based schooling, not replace it.

This should therefore be an opportune time for the launch of the Common Core. Information technology can help the effort in so many ways. It can provide formative assessments and track student progress. It can support the rapid development of effective curricula. As textbooks are replaced by online resources, instructional programs can be evaluated continuously, with ineffective lessons replaced by effective ones. Research by Grover (Russ) Whitehurst, former head of the Institute of Education Sciences, has shown the very substantial difference that curriculum can make for student achievement. Technology can facilitate the measurement of curriculum effects and support its ongoing improvement. Technology can blend with face-to-face teaching to provide more efficient and effective instruction in demanding new standards. It can perhaps



provide the scaffolding as teachers themselves provide the higher-order skills.

However—and this is a big however—the Common Core is being developed without serious attention to new ways of doing schooling and new ways of incorporating technology. Other than electronic formative assessment, technology does not appear to be on the radar screen. Worse, the Common Core may actually reinforce existing impediments to technological innovation. The new standards are conceived as annual markers that students should hit before proceeding to the next grade. Assessments are conceived as they are today, as annual events. Instruction, to the extent it has yet been imagined, is assumed to be teacher-led. Yet, technology makes it possible for students to move at their own pace and show mastery when they are ready—some much quicker than annually, some a bit slower. Students might also proceed at different paces in different subjects. Some students might learn better through teachers, others through technology. The standards are the same for all; the pace and manner in which they're attained can vary.

The Common Core faces one very high hurdle that technology could help solve. As standards are raised, the chance is great that disadvantaged students will struggle even more than they do today. The nation faces a huge challenge in expecting more from every student, when even lower standards already are not being taught effectively or mastered widely now. There is no single solution to this challenge; better teachers, fairer accountability, and other measures will all help. But ultimately, there is no reason to believe that these measures taken together will be enough. Technology, however, offers a new way. It provides additional modalities for instruction. It permits teachers to be more engaged with some students and less with others, as individual needs dictate. And it provides much more opportunity for targeted re-teaching and practice with skills.

The difficulty here, as with other challenges, is that we simply do not know enough to prescribe a path forward. As in the other areas discussed, research is sorely needed—into effective uses of formative assessment data, into more effective curricula in all subject areas, into models of blended learning that more effectively mix teachers and technology, and into instructional models that are especially successful in helping reduce the achievement gap that the Common Core may exacerbate. Meta-analyses have established the overall merit of online learning. But this research provides little guidance for creating technology-based programs that will be effective with the new standards.

The most difficult work is just beginning—helping schools and teachers figure out how high standards will be taught and learned. There is a real danger that the traditional structure of the new standards could discourage promising new approaches to teaching and learning through technology, just as interest in new approaches is taking off. It would be unfortunate in the extreme if the Common Core became less a force for deeper learning and more a reinforcement of the status quo. That has been the history of school reform. It need not be here. Hard evidence of success is a great antidote to the frustrations of major change. >>

REFERENCES

- Carey, Kevin and Robert Manwaring, *Growth Models and Accountability: A Recipe for Remaking ESEA* (Washington, DC: Education Sector 2011) http://www.educationsector.org/sites/default/files/publications/GrowthModelsAndAccountability_Release%20.pdf
- Chetty, Raj, John N. Friedman, Jonah E. Rockoff, "The Long-Term Impacts of Teachers: Teacher Value-Added and Student Outcomes in Adulthood" Working Paper 17699 (Cambridge, MA: National Bureau of Economic Research, 2011). <http://www.nber.org/papers/w17699.pdf>
- Chubb, John E., *Learning from No Child Left Behind* (Stanford, CA: Hoover Institution, 2009).
- Dee, Thomas and Brian Jacob, "The Impact of No Child Left Behind on Student Achievement" Working Paper 15531 (Cambridge, MA: National Bureau of Economic Research, 2009) <http://www.nber.org/papers/w15531.pdf>

ESSELECT

Our experts' take on the issues that matter today . . . or will tomorrow.

Finn, Chester E. Jr. and Daniela R. Fairchild, eds., *Education Reform for the Digital Era* (Washington, DC: Thomas B. Fordham Institute, 2012) <http://www.edexcellence.net/publications/education-reform-for-the-digital-era.html>

Getting to 2014 (And Beyond): The Choices and Challenges Ahead (Washington, DC: Education Sector, 2012) http://www.educationsector.org/sites/default/files/publications/Gettingto2014_BEYOND_Release.pdf

Gordon, Robert, Thomas J. Kane, and Douglas O. Staiger, *Identifying Effective Teachers Using Performance on the Job* (Washington, DC: The Brookings Institution, 2006).

Hanushek, Eric A. "Impacts and Implications of State Accountability Systems," in *Within Our Reach: How America Can Educate Every Child*, ed. John E. Chubb (Stanford, CA: Hoover Institution Press, 2005): 95-112.

Hanushek, Eric A. and Steven G. Rivkin, "Generalizations about Using Value-Added Measures of Teacher Quality," *American Economic Review: Papers & Proceedings* 100 (May 2010): 267-271. <http://hanushek.stanford.edu/sites/default/files/publications/Hanushek%2BRivkin%202010%20AER%20100%282%29.pdf>

Hassel, Emily Ayscue and Bryan C. Hassel, *Seizing Opportunity at the Top: How the U.S. Can Reach Every Student with an Excellent Teacher* (Chapel Hill, NC: Public Impact, 2011) http://opportunityculture.org/seizing_opportunity_policybrief-public_impact.pdf

Headden, Susan, "A Test Worth Teaching To," *Washington Monthly*, May/June 2012. http://www.washingtonmonthly.com/magazine/mayjune_2012/special_report/a_test_worth_teaching_to037194.php

Lerner, Lawrence S., Ursula Goodenough, John Lynch, Martha Schwartz, and Richard Schwartz, *The State of State Science Standards 2012* (Washington, DC: Thomas B. Fordham Institute, 2012) <http://www.edexcellencemedia.net/publications/2012/2012-State-of-State-Science-Standards/2012-State-of-State-Science-Standards-FINAL.pdf>

Manwaring, Robert, *Restructuring 'Restructuring': Improving Interventions for Low-Performing Schools and Districts* (Washington, D.C.: Education Sector, 2010) <http://www.educationsector.org/sites/default/files/publications/Restructuring.pdf>

Middle School Mathematics Professional Development Impact Study (Washington, DC: Institute of Education Sciences, April 2010) <http://ies.ed.gov/ncee/pubs/20104009/pdf/20104010.pdf>

Peterson, Paul E. and Frederick Hess, "Few States Set World-Class Standards," *Education Next* Vol. 8, no. 1 (Summer 2008): 70-73. <http://educationnext.org/few-states-set-worldclass-standards/>

Loveless, Tom, *The 2012 Brown Center Report on American Education: How Well Are American Students Learning?* (Washington, DC: The Brookings Institution, 2012) http://www.brookings.edu/~media/newsletters/0216_brown_education_loveless.pdf

Loveless, Tom, *The 2007 Brown Center Report on American Education: How Well Are American Students Learning?* (Washington, DC: The Brookings Institution, 2007) http://www.brookings.edu/research/reports/2007/12/~media/Research/Files/Reports/2007/12/11%20education%20loveless/1211_education_loveless.PDF

Loveless, Tom, Steve Farkas, and Ann Duffett, *High Achieving Students in the Era of NCLB* (Washington, DC: Thomas B. Fordham Institute, 2008) http://www.edexcellencemedia.net/publications/2008/200806_highachievingstudentsintheeraofnochild-leftbehind/20080618_high_achievers.pdf

Sanders, William L. and June C. Rivers, *Cumulative and Residual Effects of Teachers on Future Student Academic Achievement* (Knoxville, TN: University of Tennessee Value-Added Research and Assessment Center, 1996) http://www.cgp.upenn.edu/pdf/Sanders_Rivers-TVASS_teacher%20effects.pdf

Whitehurst, Grover J., *Don't Forget Curriculum* (Washington, DC: The Brookings Institution, 2009). <http://www.brookings.edu/research/papers/2009/10/14-curriculum-whitehurst>.

ABOUT THE AUTHOR

JOHN E. CHUBB is interim CEO of Education Sector. He can be reached at jchubb@educationsector.org

ACKNOWLEDGEMENTS

This piece was funded by The William and Flora Hewlett Foundation. We thank the foundation for its support. The findings and conclusions presented here are those of the author alone and do not necessarily represent the opinions of the foundation.

ABOUT EDUCATION SECTOR

Education Sector is an independent think tank that challenges conventional thinking in education policy. We are a nonprofit, nonpartisan organization committed to achieving measurable impact in education, both by improving existing reform initiatives and by developing new, innovative solutions to our nation's most pressing education problems.



EDUCATIONSECTOR

www.educationsector.org

