

Covering the Costs

Caroline M. Hoxby

Most Americans, whether employers or parents or people who do business internationally, recognize that our students' achievement is mismatched with our economy. The growing sectors of our economy are highly skill-intensive, and only the shrinking sectors require unskilled laborers. Yet, as evinced by scores on the National Assessment of Educational Progress (NAEP), the share of our population that is capable of performing highly skilled jobs is no greater than it was forty years ago. Our students' achievement is mediocre compared to the achievement of the people worldwide with whom they will have to compete for jobs in the future. For instance, American fifteen-year-olds scored below the average in mathematics in 2009 among students in member nations of the Organisation for Economic Co-operation and Development (OECD). They merely scored at the average in reading.¹ Moreover, in the future, US students will not compete only with OECD students. They will compete with millions of people from countries like India and China where the number of well-educated young adults is growing very rapidly.

The recognition that American students must improve is not enough, however. The United States needs to find the methods and the resources to make the improvements. In this paper, I examine some of the methods that hold the greatest promise and argue that they are affordable with the resources we *already* have.

If the past few decades are anything to go by, we might scoff at the idea of improving American students' skills at no additional cost. From 1970 to 2010, per-pupil expenditure on public elementary and secondary education rose by 327 percent in dollars adjusted for inflation.² Over the same period, high school students' scores on the NAEP rose not at all. In reading, they scored 285.2 in 1971 and 296 in 2008. In mathematics, they scored 304 in 1973 and 306 in 2008.³ In other words, achievement stubbornly failed to improve even when we "improved" schools in ways that required enormous increases in resources. Why, then, should we think that we might now be able to raise achievement without additional resources?

The answer is that, now, we are contemplating very different methods of improving schools. These methods are not only much more promising than the "improvements" conducted in the past. They are also more likely to pay for themselves because they move money toward the people who, and schools that, actually raise achievement and—crucially—move it *away* from those that do not. Put another way, these new methods tend to *align* resources with effectiveness, much as private markets do. Because these methods push schools into greater alignment with the private sector, the realignments are not only feasible but sustainable over the long term. This makes them differ fundamentally from previous "improvements" that required across-the-board increases in spending or that required omniscient regulators who somehow command-and-control in opposition to market forces.

Before we get to the new methods that can potentially pay for themselves, however, one might ask why containing cost growth in K–12 education is more important now than in the past. First and most important, skills are the crucial input for our growth industries. If the United States can only produce skills at much greater expense than other countries can, it will lose its comparative advantage in high-skill industries. Jobs in those industries will drain away to other countries with lower costs of producing skills. This is basic economics, and it matters much more than in decades past when few other countries produced educated people. Second, many Americans are disillusioned and exhausted by calls for further increases in education spending. They refuse to keep “throwing money at the problem.” Third, for the last forty years, US schools have not been accurately reporting the true growth in their education spending. They have promised increasingly generous pensions to teachers and other staff members and put insufficient money into the trust funds intended to pay those pensions. The pensions are now due and will continue to be paid for decades. As a result, *even if spending on current students does not rise at all*, spending on public education will rise substantially simply to meet pension obligations to retired staff. Finally, rising health care costs have been crowding out education spending in government budgets for more than twenty years. They will continue to do so because of the aging of the population—even if health care inflation abates. All in all, it would be foolhardy to suggest that we improve US education via methods that require further, substantial increases in spending.

What are these new methods of improving American education that are both more promising than past methods and that can plausibly pay for themselves? They are:

- i. school choice and competition,

- ii. rewarding teachers based on their value-added, and
- iii. using technology to substitute for routine instruction and to customize instruction

What This Paper Is and Is Not About

It is the goal of this paper to outline what is *feasible* given the resources we have and what we know about methods of raising achievement. My goal is to describe financial models whereby schools could realize improvements without needing additional resources. In that sense, it is a realistic paper.

However, it is not the goal of this paper to predict or be constrained by the bastardized policies that fallible politicians often enact when they bow to pressures from lobbyists, public sector unions, fundraisers, and other interest groups. Politicians often misuse resources to pay for support, bribe opponents to be silent, or “grandfather” ineffective programs. They create programs with elements that undermine the intended policy—sometimes through mere incompetence, other times because some interest group prefers that the program self-destruct. While papers on the politics of education and the misuses of government funds are important for understanding the forces opposed to true reform, relentless attention to such topics obscures the possibilities for raising American students’ skills.

Furthermore, this paper does not attempt a comprehensive review of the evidence on the effectiveness of the three improvement methods listed above. Nor does this paper attempt to “get into the weeds” of policies such as exactly how value-added can best be estimated. There are many papers that cover such topics.

Rather, the focus is novel: do logic and the available evidence suggest that new, promising methods of improving achievement can be implemented within current education budgets? The phrase “available evidence” is important. This paper does not purport

to offer definitive evidence. Indeed, one of the goals of this paper is to suggest what evidence we need to gather via future policy experiments and future research.

Rewarding and Employing Teachers Based on Their Value-Added

In recent years, researchers have demonstrated that individual teachers differ substantially in their value-added—their propensity to raise students' achievement. For our purposes, the key take-aways from the research are as follows:

1. A teacher who is in the top 10 percent of the current distribution of value-added raises student achievement by several times what a teacher in the bottom 10 percent does.⁴
2. If all US teachers had value-added equal to what the current top 10 percent has, the average American student would achieve at the level of students whose parents have incomes in the top 10 percent of the family income distribution. This is approximately equivalent to the level at which the average student in Singapore achieves.⁵
3. Even after a single year of teaching, we can predict a teacher's value-added sufficiently well that the retention decision can be made at that point. Two years of teaching experience adds information, but subsequent years of teaching add little to our predictions.⁶
4. After the first few years of teaching (during which most teachers' value-added rises, apparently through on-the-job learning), a teacher's value-added remains fairly stable.⁷
5. Teachers' measured value-added based on students' test scores is highly correlated with their measured value-added based on long-term outcomes such as students' earnings, employment, and college attainment.⁸

6. Teachers' value-added on students' non-cognitive outcomes, such as truancy, attendance, and disciplinary incidents, can be measured and gives therefore a better fix on how teachers affect long-term outcomes.⁹
7. Teachers with master's degrees do not have systemically higher value-added. If anything, the evidence suggests that master's degrees are associated with lower value-added.¹⁰
8. There are almost no credentials that predict a teacher's value-added, with the exception that individuals who attended a more selective college tend to have higher value-added. This is presumably due not only to their higher aptitude and better education but also to the fact that selective colleges admit students in a fairly holistic way that rewards leadership, motivation, and similar traits.¹¹
9. There is evidence that teachers improve their value-added when exposed to other teachers who have high value-added and when they are offered pay based on their value-added.¹²

Suppose that we wanted to have a pay system that attained the goal that all teachers (except the brand new, untried ones) would have value-added in the range we currently see among the top 10 percent of teachers (hereafter: "high value-added teachers"). A well-designed pay system would achieve this both through selection (getting the right individuals to become and stay teachers) and through incentives (giving teachers incentives to raise their value-added). What would such a system look like? First, within teaching, pay would be aligned with value-added. This would provide the right incentives for improvement. Second, to ensure that high value-added individuals select into teaching, their pay would have to be competitive with alternative jobs that they could obtain. Since the evidence suggests that people who can be high value-added teachers are also better at other jobs, we should assume that their alternative jobs will be those that a baccalaureate degree holder with above-average aptitude, college quality, and motivation could

obtain. Note that these are *not* the alternative jobs that most current teachers now take if they leave teaching. This is because most current teachers have aptitude and college quality well below average among baccalaureate degree holders.

Consider some private sector occupations that are filled by people with baccalaureate or more education (but not a professional degree or PhD): accountants, compensation and benefits managers, computer programmers, editors, landscape architects and surveyors, property managers, occupational therapists, regional planners, public relations specialists, and buyers for major retail stores. These people are all paid based on their productivity and rewarded by private sector employers for their intelligence, motivation, and ability to work with other people. Benjamin Scafidi, David Sjoquist, and Todd Stinebrickner show that only about 5 percent of people who leave teaching ultimately take a job along these lines.¹³ Sara Champion, Annalisa Mastri, and Kathryn Shaw show that the teachers who leave for such jobs are those whose value-added in teaching is unusually high.¹⁴

Let us say, then, that any high value-added teacher's compensation—pay plus benefits—must be equal to the average for full-time workers in these occupations.¹⁵ Their average annual compensation was \$89,989 in the 2009–10 school year (the most recent year for which full financial data on US schools are available).¹⁶ Suppose that, by employing value-added-based pay and creating the right selection and incentives, US schools were able to attain a teaching workforce almost entirely composed of high value-added teachers. Would they have sufficient money to pay compensation of \$89,989 to all teachers except untried, new ones? The answer is “yes” with current budgets. In the 2009–10 school year, public schools paid \$275.3 billion in compensation (wages, salaries, payments for benefits) to instructors and employed 3,123,957 full-time equivalent classroom teachers.¹⁷ This is \$88,132 per teacher—enough to pay 95 percent of teachers the competitive compensation of \$89,989 and pay the remaining 5 percent of teachers (presumably

the untried ones) compensation of about \$53,000. If, by employing only high value-added teachers, schools could eliminate instructional support staff (people who are not certified instructors but who help with instructional improvement, curriculum development, staff training, and the like), the average teacher might be paid \$94,177 out of current budgets.¹⁸

Keep in mind that no adjustment has been made for the fact that teachers' contract hours are only about 1,200 hours per year (about 0.6 of full-time) and that teaching jobs have work hours, work locations, and vacation timing that allow them to avoid child care costs that the typical full-time worker must bear. Thus, it is generous to assume that they need to be compensated like someone who works full time. I have made such generous assumptions because *high* value-added teachers may actually work a good many more hours than their contract hours.

One might wonder how it is possible that US public schools could, within their current budgets, pay teachers in a manner that is so competitive with private sector rewards. The main explanation is that although high value-added teachers are currently underpaid, low value-added teachers who have high seniority, master's degrees, and other paper credentials are systemically overpaid relative to their alternative jobs. They have no incentive to leave teaching, therefore. They also have no incentive to improve their value-added. Low value-added teachers absorb so much of the total compensation budget that little is left for high value-added starting teachers, who are not only underpaid if they do teach but who tend to leave teaching as a result. A second, less important explanation is that a non-trivial share of instructional compensation goes to people who are not classroom teachers but who provide some form of instructional "assistance." Randomized controlled trials suggest that such assistants add little or no value in terms of student achievement.¹⁹ As a logical matter, this may be because principals often steer incompetent teachers into non-classroom instructional jobs. That is, such jobs act as "holding tanks" for ineffective teach-

ers who have not been dismissed—sometimes because of formal barriers to dismissal but more often because there is no norm of dismissal for mere incompetence.

Importantly, a system in which teachers are compensated in a manner that is fully competitive with their private sector alternatives would be highly sustainable. This is a major but often underappreciated benefit of education reforms that move schools in the direction of managing themselves as private organizations do. This is in contrast to well-meaning but starry-eyed educational methods that depend on school personnel being so altruistic or specially attuned to teaching that they are rare or ignore the fact that they are undercompensated.

Employing Technology to Enrich Instruction and Substitute for Mundane Instruction

The combination of modern computing, software, and the Internet has led various researchers and commentators to suggest that technology could substantially improve American education. The evidence on this front is still very slight, owing mainly to the absence of evidence from randomized controlled trials or similarly robust policy experiments.²⁰ However, there is fairly compelling logic to the claim. The logic is that current technology allows students to interact with course materials in a highly individuated way (as opposed to a non-individuated textbook, say), to obtain instant feedback and diagnosis in some subject areas (for instance, on mathematics problems), and to gain access to a rich array of auxiliary information via hyperlinks. Technology should also allow teachers and administrators to track students' progress—on some types of assessments, at least—fairly easily.

Suppose that we accept the idea that modern technology can enhance instruction and improve achievement. We must still grapple with the question of whether it can pay for itself. We ought not to be overly sanguine because previous waves of

technology—television, instructional films, telephony, audio-visual machines, photocopying, central computers, and computers on teachers' desktops—appear simply to have added to schools' costs. When they were introduced, it was often argued that they would substitute for mundane instructional or logistical tasks and thus pay for themselves via reductions in school staff. It is nearly impossible to find evidence of this, given the unrelenting increase in the ratio of school staff to pupils over the years in which such technology was introduced.

What makes current technology more likely to generate cost savings than, say, television or instructional films? The key distinction is the degree of interactivity—which boils down to the claim that students can be “hooked” by software that responds to them in a way that they cannot be “hooked” by technology that allows them to be passive. Thus, the technology can keep them productive and engaged for non-trivial periods of time—periods that staff can allocate to other students and other tasks. Indeed, since almost no one claims that technology can substitute for some instructional tasks—motivating children, organizing projects, teaching higher level writing skills, diagnosing learning problems in a manner that requires empathy—the logic of modern technology is that staff should focus on the non-substitutable activities. They can then be spread more thinly among students because each student only needs to engage in non-substitutable activities for part of each day. This is known as the hybrid model—that is, a school in which students work with teachers part of the time and work on individual computers the remaining time.

As a logical matter, a hybrid school could improve achievement within current educational budgets. It just has to be the case that the necessary and substantial investments in technology (computers, networking, software, technology maintenance staff) are offset by reductions in other staff costs. This seems possible, but can it occur in practice?

Unfortunately, there is a paucity of hybrid schools in the United States owing to the newness of the hybrid model.²¹ Since it is impossible to describe a large, representative sample of the hybrid schools that will probably exist in the future (since they do not exist now), it makes sense to take a case study approach and analyze one network of schools on which detailed, accurate financial information is available: the Rocketship schools located in Santa Clara County, California.

The Rocketship schools are hybrid schools that serve students who are largely poor and Hispanic or black. They attain some of the highest scores for students from such backgrounds among California schools.²² Their students spend part of each day in a “learning lab” in which they work on computers. The schools also use computer-based technology for curricular enrichment, diagnosis, and tracking progress. (Although the Rocketship schools are charter schools and largely admit students via lottery, they have not yet been evaluated using lottery-based methods. Thus, some of their high performance may be due to motivated or able students’ self-selecting into them. However, what really interests us here is their financial model.)

The Rocketship schools have current per-pupil expenditures equal to 79 percent of that of traditional public schools in their county: \$7,492 for Rocketship and \$9,463 per school year for the other schools.²³ How do they manage this? First, their ratio of pupils to classroom teachers is 30.5 while the traditional public schools’ is 21.6. Thus, Rocketship schools need only two teachers for every three teachers whom the traditional public schools need. According to their accounts, this entire reduction is attained by means of computers being used for mundane instruction and practice of skills. Second, Rocketship schools spend a much lower share of their budget on the wages and salaries of *non*-teachers: 12.7 percent as opposed to 32.6 percent. This is largely because they have approximately one non-teaching staff member for every

three such people at traditional public schools. The schools' explanation is that they have less need for administrators and support staff because technology performs many of the tracking and paperwork tasks that such people perform in traditional public schools.²⁴

So far we have compared Rocketship to the traditional public schools only on the basis of current expenditures which include some spending for technology (software licenses, network access, staff for maintaining technology, and the like) but which do not include spending on computer hardware. However, in Santa Clara County, total capital outlay expenditure was \$1,322 per pupil in the 2009–10 school year. It is simply not plausible that Rocketship spent a similar amount on capital that it would need regardless of whether it adopted a hybrid model (that is, the same sort of buildings and equipment that the traditional public schools have) and *then* spent an additional \$1,971 per pupil on just the equipment needed for its hybrid model. With that additional expenditure, Rocketship could purchase two computers and associated networking equipment for every student every year. (In fact, if it did do this, it would be unable to count the computers as capital expenditures because capital must have a working life of more than a year.)

In other words, it is clear that Rocketship is able to use a hybrid model of instruction while staying well within the total per-pupil budget of local, traditional public schools.

The hybrid model is probably not the only reason why Rocketship schools have lower per-pupil current expenditure. They are charter schools so they have stronger incentives to use their funds efficiently than traditional public schools. They also have greater management autonomy than most traditional public schools.²⁵ These are points that I take up in the next section. However, the Rocketship financial model is a proof-of-concept. That is, it is proof that it is possible to reduce staff sufficiently by means of technology that the technology pays for itself—without apparent harm to (indeed, with apparent benefits for!) achievement. It is this question

that was the main source of doubt—not the question of whether technology could be used to enhance or individuate instruction.

Choice and Competition

The best available evidence, which is based on comparing students who are randomly “lotteried-in” and attend choice schools to students who are randomly “lotteried-out” and attend traditional public schools, indicates that charter schools and vouchers raise student achievement.²⁶ Moreover, the typical choice school in the United States—a charter school—has per-pupil spending less than half that of the average traditional public school.²⁷ Private schools involved in voucher systems also spend far less than the average traditional public school.²⁸ Thus, the most straightforward look at the facts suggests that choice schools have both higher value-added in achievement and lower costs.

Why should we be unsurprised by this dual attainment, as a logical matter? The answer is that *choice* creates a greater demand for schools that students prefer. *Competition* ensures that the supply side will generate more seats in schools that can efficiently produce what students prefer and fewer seats in schools that produce it inefficiently. That choice must be *combined* with competition is a point missed by many commentators, but it is crucial to the discussion of costs in this paper because it is supply-side forces that ultimately determine schools’ productivity and the affordability of education.

Owing to the numerous programs that use the word “choice” without adhering to the basic tenets of the idea, it is worth recalling that all true choice and competition reforms:

- i. Attempt to make schools more productive by:
 - A. Forcing them to compete for students and the funding that follows them

- B. Giving them autonomy to manage how they add value—that is, affect students' outcomes
2. Confine government intervention in K–12 education to:
 - A. Ensuring that schools compete on an even playing field
 - B. Ensuring that families have accurate information on how schools affect students' outcomes (information on *effects*, not merely on outcomes)
 - C. Ensuring that individual students can make the investments in their own human capital that will pay off

When a choice and competition reform has a design that implements these ideals (something that is very feasible but requires detail beyond the scope of this paper), it gives schools strong incentives to raise their benefit-to-cost ratios and unbinds the constraints that prevent some schools from doing this now.²⁹ With better information and funding that follows students, families are more aware of schools' benefits *and* costs. Since families do not reward schools for wasting resources on programs that do not benefit their children, schools are under pressure to eliminate unproductive programs, teachers, and work rules. Schools are also under pressure to adopt innovations that create efficiencies or raise achievement per dollar spent. Schools can be more productive when families are able to match their children's educational needs to schools' pedagogical strengths.

Because choice and competition represent such a thoroughgoing change in schools' incentives and constraints, we expect to see schools change on numerous dimensions that raise their benefit-to-cost ratios. For instance, it should come as no surprise that charter schools have been pioneers in paying teachers based on their performance and in using the hybrid model of technology in instruction.³⁰

Consider a few other ways in which choice schools often manage their budgets differently to produce higher value-added than traditional public schools do.

First, recognizing that most teachers do not improve their value-added once they have experience of about four years, choice schools frequently employ numerous less-experienced teachers in combination with a “leaven” of highly skilled master teachers who are not only more experienced but who are good at conveying the fruits of their wisdom to the less-experienced teachers. This system allows a choice school to pay teachers in a manner that is highly competitive with traditional public schools and the private sector: there is more money for less-experienced teachers because money is not wasted on giving seniority pay to teachers who are no more productive. This system also allows a choice school to keep highly effective teachers involved in instruction, rather than inducing them to choose an administrative job in which they earn higher salaries but leave teaching.

Second, choice schools often rethink the school year and day. Traditional public schools spend considerable effort ensuring that the number of hours that a teacher is in the classroom is below some amount, that her hours for preparation are above some amount, that the days in the school year are below some amount, and that professional development days are above some amount. In contrast, many choice schools recognize that students’ achievement can be *directly* affected by the hours and days they spend on school grounds, in the school’s custodial care (not necessarily in instruction), and on fundamental tasks like reading. Thus, it is not unusual to see choice schools experiment with year-round calendars; school days that start early and end late; and school days that contain substantial periods for meals, homework, and play. Choice schools often make these changes pay for themselves by substituting non-teachers for teachers efficiently (when instruction is not going on), by reducing losses associated with students taking books and materials home, and by reducing the need for remediation and disabled instruction.

Third, recognizing that diagnostic assessment is inherently inexpensive (owing to the massive economies of scale associated

with modern testing, especially computer-adaptive testing), choice schools often give their students frequent, short assessments that provide near-instant diagnostic information about students' learning gaps. This allows teachers to modify their instruction, revisit confusing topics, and provide remediation in real time. This is in contrast to many traditional public schools that feel so burdened by assessment that they "save it" for the annual state mandatory exam that generates results with such a long lag that they are useless for instructional adjustments in real time. It also appears that traditional public schools are unwilling to move even modest parts of their budgets to assessment (which is cheap but has no natural lobby) and from—say—staff compensation (expensive but has a lobby).

So far I have emphasized the tendency of choice schools to save money or, at least, get more with the same budget. However, there are some items that might cost *more* in a system with true choice and competition. First, we might expect school facilities to need some excess capacity if schools are to shrink and expand with students' demand for them. Second, we might expect additional transportation costs if students are not constrained to attend schools based purely on geographic attendance zones. Third, we might expect some costs associated with making information on schools' effects available to families in a user-friendly way.

Regarding the capacity issue, the areas in which choice programs play the largest role (Washington, DC, New York City, Chicago, and so on) have facilities expenditures that are no higher per pupil than areas in which choice plays no role.³¹ As a logical matter, this is because they have re-purposed (often as charter schools) school buildings from which enrollment was draining because families disliked the school in question. However, deliberate systems of re-purposing, which make a great deal of sense, are only needed when existing buildings have exhausted their excess capacity. The most recent, representative data indicate that 59 percent of US schools have excess capacity of at least 5 percent. About

40 percent of schools have excess capacity greater than 20 percent. Thus, the system could easily allow much more flexibility in enrollment than it does currently. Indeed, even among the 18 percent of schools that house more students than they were designed to house, about half are only temporarily “over-enrolled” owing to renovation. And, portable buildings (which often house administrative offices and storage rather than classrooms) have proven to be a highly efficient way of allowing schools to change size smoothly.³²

School transportation costs in cities—like Washington, DC, New York City, and Chicago—with prevalent choice are not higher on a per-pupil basis than in cities with no choice.³³ However, such cities are not representative of the United States: they are densely populated and have high-coverage systems of public transportation. In order to consider more typical areas, let us compare Pennsylvania and Ohio. These states have comparable shares of their students in private schools and charter schools but Pennsylvania provides school bus transportation for private and charter school students whereas Ohio does not. Yet, Pennsylvania does not spend more per student on school transportation than Ohio does, regardless of whether we adjust for factors like population density.³⁴ Logically, this is because the vast majority of stops on a school bus route are at points where students are picked up or dropped off, not at schools (which are relatively few, even in an area where choice is prevalent).

Computing schools’ value-added and other effects on outcomes may be somewhat beyond the typical person, but making the computations costs a trivial amount relative to education budgets. This is because the computations rely on administrative data that are gathered already. Even the latest value-added evidence, which shows that teachers’ value-added can be measured in terms of students’ non-academic outcomes (such as attendance and discipline) and long-run outcomes (such as college-going), can be compiled at so little cost that no more than 0.1 percent (*one-tenth* of 1 percent) of a state’s total education budget could plausibly be absorbed.³⁵

Moreover, providing parents with information on schools costs very little. New York City, in which all prospective ninth graders must submit a preference-ordered list of high schools, is widely recognized for its successful system of informing families about their options, partly through written or online communications and partly through local “choice” centers where families consult individual counselors. Other cities, such as San Francisco, Boston, Charlotte-Mecklenberg, North Carolina, and London, have similar programs to disseminate information. In no case does such a program account for even 0.1 percent of the local education budget. Simply put, disseminating information is very inexpensive. This is because, logically, it is an activity in which there are enormous economies of scale.

Summing up, choice and competition reforms do impose some costs, but they are minor relative to the sources of cost saving—which arise through thoroughgoing incentives to improve productivity.

Virtuous Circles

We have already seen that choice and competition give incentives to schools to be productive—that is, to adopt educational methods that deliver benefits disproportionate to their costs. As a result, we see in practice that choice schools are more likely to reward teachers based on their performance, use technology in hybrid instruction, and so on. But the complementarity between the new methods of improving education goes deeper than such joint adoption.

Fundamentally, all of the new methods we consider in this paper move schools closer to the incentives and rewards that arise in the private sector. Once a school starts down this path, the productivity of various people and inputs becomes much more obvious and many decisions may be reconsidered. In contrast, the command-and-control methods and the arbitrary rewards that currently prevail in the traditional public schools tend to obscure

the productivity differences among different policies. As a result, many policies in the traditional public schools do not hold up under benefits-versus-costs scrutiny.

To see this, consider a concrete example. Suppose that a school breaks free from the current system of paying teachers solely on their credentials and seniority and begins to pay them based on their productivity (their effects on students' outcomes, broadly construed). Then, a school may realize that teachers who can work smoothly with technology are able to instruct more students because hybrid classrooms work well for them. Therefore, the school may offer teachers more pay if they show themselves to be apt at hybrid instruction. As a result, the school will tend to fill up with teachers who are good with technology. Those who are not will go elsewhere. But, once the school is filled with technology-enabled teachers, it may realize that it would be productive to invest in even better technology—perhaps technology that allows students to submit homework via computer or that allows teachers to communicate efficiently with parents. With such technology in hand, the school may find it productive to rethink its school day or alter its evaluation system. Perhaps students complete more homework at the school. Perhaps evaluations should be more frequent or problem-based. The point is that once every policy's benefits and costs are considered before it is adopted (or maintained), new methods that improve productivity will be routinely adopted.

Schools that not only increase achievement but do so in the most productive manner create the foundation for skill-based economic growth.

The Need for New Research

The first generation of research on teachers' value-added, technology, and choice focused—for good reason—on simple demonstrations of efficacy. Do teachers actually differ in their value-added? Do they respond at all to rewards? Does some technology raise

achievement? Does a student who attends a choice school learn more than one who is lotteried-out? And so on.

Such research plays a crucial role. Without it, policymakers would have little basis for suggesting the expansion of new methods. Nevertheless, the first generation of research typically left the financial side of the methods very vague—partly because the policy trials being analyzed often had arbitrary, even bizarre, financing. For instance, in the few policy trials that exist, the size of rewards for teachers was selected quite arbitrarily—not on a scientific basis designed to maximize the benefits for the costs. Similarly, choice schools are funded in different ways and to different degrees in different cities and states. Often, we can make no sense of why similar states fund their choice schools at such different levels. There are also choice schools that are so poorly funded—for instance, vouchers equal to one-sixth of local per-pupil funding—that it is impossible to extrapolate from their experience to system-wide choice and competition. Schools' technology adoption has often been funded by large federal grants that give schools strong incentives to adopt the technology that is “free” under the grant rather than the technology that would most improve their students' learning.

Indeed, in this paper, we have seen that there is often little research on exactly how pay should change with performance, on how hybrid instruction should be funded, and on exactly what amount of funding should follow each student in a choice and competition environment. Future research needs to fill these gaps for several reasons.

First, reforms that are fundamentally intended to move schools closer to the incentives and constraints of the private sector need to get the finances right. It would be bizarre to say that financial incentives should play an important role, yet the amounts of the incentives should be determined by guessing. Second, policymakers who are fortunate enough to have the support to enact important school reforms often find themselves in a quandary when setting actual amounts. Researchers often help policymakers set up

the framework of a new system—for instance, by estimating teachers’ value-added—but then beg ignorance when asked for money amounts (such as how much more a teacher should be paid if her value-added is at the ninetieth percentile). The result is that a policy that might work well ends up with a bad financial design and it fails—even though everyone’s intentions were good. This often leaves a bad taste, making it harder to enact future reforms. Third, as emphasized in this paper, the new methods of improving schools are all *systemic* methods at heart—that is, they rely on changing incentives and constraints system-wide. To see this, contrast a policy that alters the entire way in which teachers are paid to a policy that introduces a new textbook (a non-systemic method). The new textbook policy will interact in limited ways with other school policies, but the new way of paying teachers will interact with every other important school policy. Because such interactions are important, we need much more information about the financial implications of each policy.

In short, this paper attempts to demonstrate that, with current education budgets, we can afford new methods of improving schools—specifically those that move schools closer to private sector rewards and productivity. However, this demonstration also indicates that we need more research to get those rewards right.

Notes

1. See US Department of Education, “Highlights from PISA 2009: Performance of US 15-Year-Old Students in Reading, Mathematics, and Science Literacy in an International Context,” December 2010, <http://nces.ed.gov/pubs2011/2011004.pdf>.
2. The inflation adjustment is done using the GDP deflator (US Department of Commerce, Bureau of Economic Analysis, 2012). The GDP deflator is the most appropriate price index for the purpose of comparing education investments to the

US economy as a whole. The other data source is the “Digest of Education Statistics” (US Department of Education, 2012) for total expenditure and average daily attendance in the 1969–70 and 2009–10 school years.

3. These differences are *not* statistically significantly different from zero. Moreover, they are negligible in comparison to a standard deviation on the NAEP (about 30 points) or in comparison to the white-black score gap (29 points). Students who are seventeen years old take the high school versions of the NAEP. The source is the National Assessment of Educational Progress Long Term Trend data tool (US Department of Education, National Center for Education Statistics, 2013).
4. This calculation is based on the persistent teacher effects computed by Douglas Staiger and Jonah Rockoff, “Searching for Effective Teachers with Imperfect Information,” *Journal of Economic Perspectives* 24, no. 3 (Summer 2010), 97–118. As they point out, however, their value-added calculations are very similar to those of numerous other researchers who use different data but like methodology. It is important to use shrinkage methods to compute the persistent teacher effect, as opposed to the annual teacher effect which also contains noise from student composition.
5. This calculation is based on National Assessment of Educational Progress individual student data, equated to TIMSS (Trends in International Mathematics and Science Study) data for the latter statement.
6. See Staiger and Rockoff, “Searching for Effective Teachers.”
7. See *ibid.* for a summary of the numerous studies that have demonstrated this point.
8. See Raj Chetty, John Friedman, and Jonah Rockoff, “The Long-term Impacts of Teachers: Teacher Value-added and Student Outcomes in Adulthood,” NBER Working Paper 17699, 2011.
9. See Clement Kirabo Jackson, “Non-Cognitive Ability, Test Scores, and Teacher Quality: Evidence from 9th Grade Teachers in North Carolina,” NBER Working Paper 18624, 2012.

10. See Thomas J. Kane, Jonah E. Rockoff, and Douglas O. Staiger, “What Does Certification Tell Us About Teacher Effectiveness? Evidence from New York City,” *Economics of Education Review* 27, no. 6 (December 2008), 615–631.
11. See Kane, Rockoff, and Staiger, “What Does Certification Tell Us?” and Staiger and Rockoff, “Searching for Effective Teachers.”
12. On the former point, see C. Kirabo Jackson and Elias Bruegmann, “Teaching Students and Teaching Each Other: The Importance of Peer Learning for Teachers,” *American Economic Journal: Applied Economics* 1, no. 4 (2009), 85–108. On the latter point, see Sally Hudson, “The Effects of Performance-Based Teacher Pay on Student Achievement,” SIEPR Discussion Paper 09-023, 2010, Stanford Institute for Economic Policy Research. In addition, there is some evidence that a very well-designed evaluation and feedback system may help teachers improve their value-added: see Eric Taylor and John Tyler, “The Effect of Evaluation on Teacher Performance,” *American Economic Review* 102, no. 7 (2012), 3628–51.
13. See Benjamin Scafidi, David Sjoquist, and Todd Stinebrickner, “Do Teachers Really Leave for Higher Paying Jobs in Alternative Occupations?” *The B.E. Journal of Economic Analysis & Policy* 6, no. 1 (2006), 1–44.
14. See Sara Champion, Annalisa Mastro, and Kathryn Shaw, “The Teachers Who Leave: Pulled by Opportunity or Pushed by Accountability?” Stanford University working paper, 2011, available at <http://www.rand.org/content/dam/rand/www/external/labor/seminars/allca/Champion-Mastro-Shaw-2011.pdf>.
15. Full-time is defined by the US Bureau of Labor Statistics as a nominal 2,080 hours per year, but in fact few of the workers in question are paid on a truly hourly basis or “punch the clock,” so their exact hours of work are not relevant or even accurately recorded. Full-time is thus what is usual among workers who define themselves as “full-time” on the jobs in question. They

are at the office more than forty hours in a typical non-vacation week, but they have paid time off in various forms.

16. Author's calculations based on Occupational Employment Statistics, US Bureau of Labor Statistics, 2013.
17. Author's calculations based on the 2009–10 Common Core of Data (US Department of Education, 2013) and the 2009–10 (2010 fiscal year) Public Elementary-Secondary Education Finance Data (US Department of Commerce, 2013).
18. Instructional support staff is defined in Exhibit 1A: School-Level Personnel by the United States Department of Education, Institute of Education Sciences, National Center for Education Statistics, 2013, <http://nces.ed.gov/pubs/web/9619ex1a.asp>.
19. See Alan Krueger, "Experimental Estimates of Education Production Functions," *Quarterly Journal of Economics* 114, no. 2 (May 1999), 497–532.
20. The best randomized controlled trial of hybrid instruction is currently Linda Cavalluzzo, Deborah Lowther, Christine Mokher, and Xitao Fan, "Effects of the Kentucky Virtual Schools' Hybrid Program for Algebra I on Grade 9 Student Math Achievement" (NCEE 2012-4020) (Washington, DC: National Center for Education Evaluation and Regional Assistance, Institute of Education Sciences, US Department of Education, 2012). However, given the way the study was conducted, it is only useful for assessing the effects on achievement, not the effects on schools' finances. Other studies that examine somewhat similar models of technology use include Lisa Barrow, Lisa Markman, and Cecilia Rouse, "Technology's Edge: The Educational Benefits of Computer-Aided Instruction," *American Economic Journal: Economic Policy* 1, no. 1 (2009), 52–74; and Larisa Campuzano, Mark Dynarski, Roberto Agodini, and Kristina Rall, "Effectiveness Of Reading And Mathematics Software Products: Findings From Two Student Cohorts" (NCEE 2009-4041), (Washington, DC: US Department of Education, Institute of Education Sciences, National Center for Education Evaluation and Regional Assistance, 2009).

21. This is not to say that the use of technology to enhance instruction is rare in public schools. According to the latest nationally representative study of teachers' use of technology in instruction, at least half of teachers claim to use technology on a frequent basis for activities such as problem-solving, drills, practice of basic skills, and preparing text. However, it appears that almost none of these teachers are in schools that deliberately substitute technology for teacher time. See Lucinda Gray, Nina Thomas, and Laurie Lewis, "Teachers' Use of Educational Technology in US Public Schools: 2009" (NCES 2010-040) (Washington, DC, National Center for Education Statistics, Institute of Education Sciences, US Department of Education, 2010).
22. See Rocketship Education, "Annual Report to the Santa Clara County Office and Board of Education," January 2013, <http://www.sccoe.org/supoffice/innovative-schools/Charter%20School%20Annual%20Reports/Rocketship%20Education.pdf>.
23. These financial data are for the 2009–10 school year. For sources, see the next note.
24. Rocketship 2009–10 school year financial data are from Rocketship Education, "Rocketship Education and its Affiliates: Consolidated Audited Financial Statements for the Year Ended June 30, 2010," <https://rsed.box.com/shared/z3kxe9iu00>. All enrollment data are based on the 2009–10 Common Core of Data. Traditional public schools' financial data are from the 2009–10 (2010 fiscal year) Public Elementary-Secondary Education Finance Data.
25. It is extremely difficult to compare charter and traditional public schools' capital expenditures because they are financed so differently, especially for charter management organizations that are starting up new schools.
26. See, for instance, Caroline Hoxby and Jonah Rockoff, "Findings from the City of Big Shoulders: Students Learn More in Chicago Charter Schools," *Education Next* 5, no. 4 (Fall 2005); Caroline Hoxby and Sonali Murarka, "How New York

- City's Charter Schools Affect Achievement," *Education Next* 8, no. 3 (Summer 2008); Joshua Angrist, Sarah Cohodes, Susan Dynarski, Jon B. Fullerton, Thomas J. Kane, Parag Pathak, and Christopher Walters, "Student Achievement in Massachusetts' Charter Schools," Center for Education Policy Research at Harvard University, 2011; Joshua Angrist, Sarah Cohodes, Susan Dynarski, Parag Pathak, and Christopher Walters, "Stand and Deliver: Post-Secondary Outcomes at Boston's Charter High Schools," MIT manuscript, 2013.
27. Author's calculations based on matching the 2009–10 Common Core of Data (US Department of Education, 2013) to the 2009–10 (2010 fiscal year) Public Elementary-Secondary Education Finance Data (US Department of Commerce, 2013). The former dataset is the source of student enrollment data. The latter dataset is the source of expenditure data.
 28. Author's calculations based on Internal Revenue Service, Statistics of Income Division, "Exempt Organizations Business Master File Extract," 2013; and on US Department of Education, National Center for Education Statistics, "Schools and Staffing Survey, 2007–08."
 29. I do not emphasize the liberating effect of choice and competition reforms on the constraints that apparently bind many traditional public schools: work rules, union contracts, state and federal mandates, and the like. The reason is that many of these constraints are essentially self-imposed because schools have had weak or no incentives to oppose them. That is, any given principal can legitimately complain that he or she is bound by current work rules, but the system of traditional public education *chose* these work rules in negotiations with staff over a long period in which incentives to be productive were extremely weak.
 30. See Hoxby and Murarka, "How New York City's Charter Schools Affect Achievement."
 31. Author's calculations based on the 2009–10 (2010 fiscal year) "Public Elementary-Secondary Education Finance Data" (US Department of Commerce, 2013).

32. The source of the facilities statistics in this paragraph is Bradford Chaney and Laurie Lewis, “Public School Principals Report on Their School Facilities: Fall 2005” (NCES 2007-007), US Department of Education (Washington, DC: National Center for Education Statistics, 2007).
33. Author’s calculations based on the 2009–10 (2010 fiscal year) “Public Elementary-Secondary Education Finance Data” (US Department of Commerce, 2013).
34. The adjustments make little difference in any case, owing to the similarity between Ohio’s and Pennsylvania’s districts in terms of student density, school size, and the like. Actually, Pennsylvania spends slightly less per student than Ohio does on school transportation, but the difference is too small to be interpretable.
35. These numbers are based on the *actual* costs of such computations as made by researchers working with districts’ and states’ administrative data, including data on non-test-score outcomes such as studied by Chetty et al., “Long-term Impacts of Teachers,” and by Clement Kirabo Jackson, “Non-Cognitive Ability, Test Scores, and Teacher Quality.”