

The Real Value of Teachers

If good teachers matter, why don't we act like it?

Thinking K-16, Winter 2004, Volume 8, Issue 1

Presented by Russlyn Ali

The Real Value of Teachers

Using New Information about Teacher Effectiveness to Close the Achievement Gap

By Kevin Carey

We spend hundreds of billions of dollars a year on public education, more than we spend on just about anything, even national defense.¹ Most of that money goes to pay teachers. We depend on them to educate our children, to produce the intellectual capital that is the foundation of our future as a society.

But do we really know how important teachers are?

Incredibly, you can walk into almost any school in America, go down the hall to the first couple of classrooms you find, look at the teachers inside, and realize this: nobody, not the principal, not the parents, not the students, not even the teachers themselves, actually knows how effective those teachers are in helping their students learn. They probably have an opinion, maybe even some anecdotal evidence. But in terms of accurate, verifiable information about how effective individual teachers are at helping each of their students learn and make progress from the beginning of the school year to the end? In the vast majority of schools, nobody knows.

This is a big problem.

This void of information compromises our public schools in a large number of ways. If we don't know which teachers are the most effective and in which areas, we can't match them up with the students who need the most help. We can't find out what makes good teachers good and bad teachers bad. We're limited in the ways we prepare them to enter the field, and the way we train them once they get there. We can't give them a chance to help one another, and we can't give them a way to help themselves. Ultimately, not

knowing how good teachers are prevents them from being as good as they can be. As a result, students suffer.

But this is a problem with a solution. The reason we don't have this information is not because it doesn't exist. It's mostly because we're not looking for it. Fortunately, some people have started to do just that, and we can profit from their example. If we follow their lead, there are many ways we can use this information to help our schools and teachers meet the challenges ahead.

In these pages, we will look at data that cuts through what is outside the control of schools and teachers to reveal the tremendous impact teachers make on student learning. We will also show how states, districts and schools are using this information to accomplish two primary goals:

- Increase the overall number of effective teachers, which includes improving the effectiveness of teachers currently in the classroom;
- Get more effective teachers into the classrooms of the low-income children who rely on them the most for their learning.

This isn't something that can wait. Our new national goals for education are centered on closing the achievement gap for low-income and minority students. This gap exists because, in addition to more stress in their lives, these children have been, and continue to be, unjustly denied a fair share of those things that assure their success in school—our expectations, our resources, and our best teachers. Justice alone demands we redress these inequities.

The Real Value of Teachers, Kevin Carey, Thinking K-16, Volume 8, Issue 1, Winter 2004, pp. 3-42. Copyright 2003 The Education Trust. All rights reserved. Reprinted by Permission.

But the maldistribution of these resources is not just morally wrong, it's also economically wrong. The future of the American economy depends on having many more knowledgeable, skilled workers throughout a range of jobs and industries. We will need all of the well-educated and well-trained workers we can get. To secure this future and meet our goals will mean providing all the things students need to excel, beginning with effective teachers.

What we didn't know then that we know now—Good teachers matter most

Why haven't we known how effective teachers are?

Historically, it's been a data problem. The only way to really gauge teachers' effectiveness is to measure what each of their students knows when they arrive in the classroom, and then see how that changes during the school year. To perform that calculation in a way that is fair, accurate, and comparable to other teachers, you must give every student in different classrooms the same test, and the test needs to be consistent with what the students were supposed to be taught during the year. Then all that information must be stored somewhere in a way that allows you to follow student's scores from one year to the next, between different schools, keeping track of which teachers they had along the way.

In other words, you need common academic standards, yearly standardized tests, and computers. Until fairly recently, few states had any of those things, and some still don't have all of them. Thus, it wasn't possible to find out how effective teachers are. As a result, many people came to think that since they didn't have the information, it therefore must not be very important.

Strangely enough, this was actually conventional wisdom for a long time—that teachers

aren't very important. Researchers and sociologists interested in education performed all kinds of complicated statistical modeling and analysis trying to figure out what makes the modern school system tick and what really determines whether or not students succeed. Lacking any good information about teacher effectiveness, they found little to indicate that teachers had an effect. And so some people chose to interpret these results to mean that teachers don't matter very much.²

But by the 1990s some states had successfully put all the elements in place: standards, assessments, and computers. These allowed researchers to track the yearly progress of many thousands of individual students, match them up with individual teachers, drop all that data into computer databanks, and really take a look at it. What did they find? That teachers mattered after all.³

Not only that, they found that teacher effectiveness varies tremendously—some are much more effective than others. Some of the earliest and best analysis has been done in Tennessee, where researchers found that all else being equal, students assigned to the most effective teachers for three years in a row performed 50 percentile points higher—that's on a 100-point scale—than comparable students assigned to the least effective teachers for three years in a row.⁴

So large was the impact of teachers on student learning that it exceeded any one thing about the students themselves. The authors of the study concluded that teacher effectiveness is the “the single biggest factor influencing gains in achievement,” an influence bigger than race, poverty, parent's education, or any of the other factors that are often thought to doom children to failure.⁵

The central importance of teachers was confirmed by recent findings from Texas. Using a massive database of student test scores from thousands of schools and hundreds of thousands of students, researchers analyzed the math performance of individual students over time, calculat-

ing the effect of individual teachers on how much students learn. The conclusion: teacher effectiveness varied dramatically and had a major impact on student performance, so much so that "...having a high quality teacher throughout elementary school can substantially offset or even eliminate the disadvantage of low socio-economic background."⁶

So the advent of academic standards, consistent yearly testing, and computers has given us the ability to do what we could never do before—measure the effectiveness of individual teachers in helping students learn. And that, in turn, has thrown a bright light on the critical importance of teachers in closing the academic achievement gap. The specter of futility that has haunted discussions of education policy for decades—the idea that schools and teachers have a limited ability to help students, particularly “disadvantaged” students—has been decisively put to rest. Students of color and students in poverty can learn up to high standards just as much as anyone else, when they have effective teachers.

Some states still have some logistic issues to work out to get their standards, assessments, and computer systems fully in place. What none of the states has are good excuses, not any more. The data that states and districts need to find out how effective their teachers are and how to help them be even more effective is quite literally sitting at their doorstep. All that remains is the will and the foresight to find that information, and use it wisely.

How to get fair and honest data about teacher effectiveness

Step one in using teacher effectiveness information to improve public education is getting the information itself. This is a fairly complicated process, involving the collection and analysis of lots of data about students, achievement, and

teachers. But it's also quite doable, which we know because some people are already doing it, right now, with great success. Easily the best example is the system that's currently up and running in Tennessee.

Created by law in the early 1990s, the system is called TVAAS, the Tennessee Value Added Assessment System. TVAAS is a “value-added” system because it's based on measuring the amount of additional learning—the “value”—that a given district, school, or teacher adds to their students during a given school year, as measured by annual tests in five different subject areas. This means that teacher effectiveness measures aren't based on the overall performance of students; they're based on the *improvement* of students from the beginning of the school year to the end. A teacher who makes better than expected progress with previously low-achieving students would be rated more effective than a teacher who made only moderate progress with previously high-achieving students, even if the high-achieving students still finished the year with better overall scores.

After determining the overall gain in student achievement, TVAAS then produces a measurement of teacher effectiveness by comparing the *actual* growth in student learning to the *expected* growth. The expected growth level is created by starting with the normal amount of academic progress that a typical student is expected to make in a given subject and grade, and then using statistical controls to adjust that anticipated progress up or down, based on the previous achievement history of each student. If a teacher has a student that has previously struggled to make academic progress over a number of years—because of motivation, aptitude, family life, or whatever the

The data that states and districts need to find out how effective their teachers are and how to help them be even more effective is sitting at their doorstep.

reason may be—then the amount of growth that a teacher is expected to help that student achieve is adjusted down accordingly. This has the effect of screening out whatever non-teacher factors affect student learning and isolating the individual teacher's contribution.⁷

TVAAS data are compiled by the state every year, and then distributed to districts, schools and teachers. In grades 2 through 8, the ratings are calculated by comparing the results of each student's test scores in a given year to the previous year. In high school, TVAAS ratings are calculated using state-developed end-of-course exams, comparing actual student pass rates to predicted pass rates based on the students' prior perfor-

mance. In addition to measuring the effectiveness of each teacher, TVAAS provides a wealth of other information about student learning, showing the amount of progress made by different kinds of students in different subjects, grade levels, and school buildings.

More than a decade of results from TVAAS and other value-added systems has shown some remarkable things about teacher effectiveness. Perhaps most importantly, it shows that some teachers are simply much more effective than others.⁸ A teacher who gets exactly the expected growth in student achievement would have a rating of 100 under TVAAS, a teacher who gets 20% more growth than expected would have a

Beyond the Basics: Incorporating more information to create richer teacher effectiveness data

The TVAAS system is based on the results of annual standardized tests. This is also true of the systems used in the other states and school districts that have begun to create value-added information.

It makes sense to use these tests. Soon every state will have yearly assessments for grades 3-8 in at least reading and math, given to all students statewide and aligned with state standards. The results will be used to determine adequate yearly progress for schools and districts under the No Child Left Behind Act. Some states also use the yearly tests to determine whether students advance or graduate. Because many of these encompass an entire year of learning, the tests naturally match up with a measure of teachers' impact on student achievement.

These tests, like all assessments, have their limitations. But those limitations are not so great that they compromise the tremendous utility of value-added information. The teachers and principals that are actually using the data in the field make that clear. Still, we'd be better off if we knew more.

Some districts are solving this problem by gathering more comparable student data throughout the year. For example, many districts are now administering standard benchmark assessments to students every 6-9 weeks, to monitor progress on an ongoing basis and make mid-course adjustments if necessary. A number of California districts use a periodic assessment called "Results" to assess growth in reading competence throughout the year. The Northwest Evaluation Association (NWEA) uses a system of quick, periodic computer-based assessments to track student progress.

Each of these processes, and many others like them, contain some form of common, progress-based student assessment data. Therefore, all could potentially be used to create value-added measures and provide more detailed information about student progress that will help teachers better understand where they're having the most success and where they may need to refocus instruction so they can better help students before they move on to someone else's classroom.

In the end, we have to keep in mind that creating teacher effectiveness information is really still in its infancy. As more schools create and use the data, as we know more about student achievement itself, we'll be better able to translate that data into rich, nuanced, meaningful information about teacher success in helping students learn.

rating of 120, etc. We find that actual TVAAS ratings don't all bunch within a few percentage points of 100. Instead, some teachers get ratings of 200 points or more, some get less than 50.

It's important to acknowledge that TVAAS effectiveness measurements aren't perfect. Since student test scores themselves only provide an estimate of the breadth and depth of a student's knowledge, measurements of teacher effectiveness based on those tests can never attain 100% precision. That said, the designers of TVAAS have gone to great lengths to be fair. Effectiveness ratings are based on multiple years of data, and are adjusted to account for the different learning history of each student. The system also adjusts its rating for the amount of student achievement information available for each teacher. If insufficient data is available to provide a reliable rating—as with newly-hired teachers, for example—the system gives them the benefit of the doubt and assumes that their performance is equal to the system average.

Because of the adjustments to ensure accuracy, and because teachers who are skilled and effective in one year tend to be effective in the next year, the value-added ratings reported in Tennessee and elsewhere are generally consistent over time. The highest-rated teachers in one year tend to be highly rated the next year and the lowest-rated teachers tend to remain low. While differences in tests, students, and teacher performance can cause some year-to-year variation, TVAAS provides reliable information about which teachers are most effective in helping their students improve.⁹

Value-Added in Practice

Across the country, states and districts are finding creative and helpful uses for value-added data. For example, over the last five years, educators, community groups, and city leaders in Chattanooga have spearheaded a comprehensive initiative to improve the quality of education

in struggling, high-poverty urban elementary schools. One recently implemented component is offering significant financial incentives to highly effective teachers to teach in those schools. Teachers who consistently show the highest TVAAS gains were guaranteed an extra \$5,000 per year in salary for three years, starting in 2002-2003, if they taught in one of nine low-performing elementary schools. Local businesses pitched in by providing the teachers with housing benefits and free graduate education. A second cohort of high-performing teachers was added for the 2003-2004 school year.

The schools in Chattanooga have used multiple strategies to give students the best possible instruction, including the incentives described above to bring effective teachers in, and a staff reconstitution process to move less successful teachers out. They are analyzing and using their data to create and maintain the best possible match-up of schools, subjects, and students. For example, elementary schools in Chattanooga reformulate their classes on a *weekly basis* depending on the particular students and subjects being taught, to ensure that the students who have specific instructional needs are assigned the most effective teacher for that skill or concept.

The net result of these policies and other efforts in Chattanooga to improve funding, management, leadership, and community support has been a significant increase in student achievement. Each of the nine schools targeted by this initiative was previously among the 20 lowest-scoring in the state on 3rd grade reading tests. All nine schools have shown statistically significant gains in all five tested subject areas in each of

The designers of TVAAS have gone to great lengths to be fair. Effectiveness ratings are based on multiple years of data, and are adjusted to account for the different learning history of each student.

the last two years, growing faster than the other schools in the district. During that time, the proportion of students reading at grade level in grade three increased from 22.6% to 36.9%. (For more examples of how value-added teacher effectiveness data is being used across the country see pages 38 and 39.)

Value-added data can transform the way teachers and school leaders understand their schools and educate schoolchildren. We've spoken to many of these local educators (see page 17), and they all emphasize how important the information is to understanding and improving student learning. The information has proved to be useful in a number of different and often unexpected ways. One district administrator told us this story:

"The test scores in one of our schools were so low that under state law we were allowed to reconstitute it—to evaluate the entire staff and replace them if necessary. So I sat down with the principal and said, 'Which teachers should be replaced?' The principal said, 'First off, I want to replace Mrs. Jones.'¹⁰ I said, 'Why Mrs. Jones?' The principal said, 'She's very negative. She's always complaining to me about her students, how they don't perform well enough, how they never live up to her expectations.' Well, after the meeting I went back and looked up the value-added data for that school. It turns out that Mrs. Jones was consistently the most effective teacher in the entire building. She was 'negative' only in the sense that she had very high expectations for her students and so was never really satisfied—and the students benefited tremendously as a result."

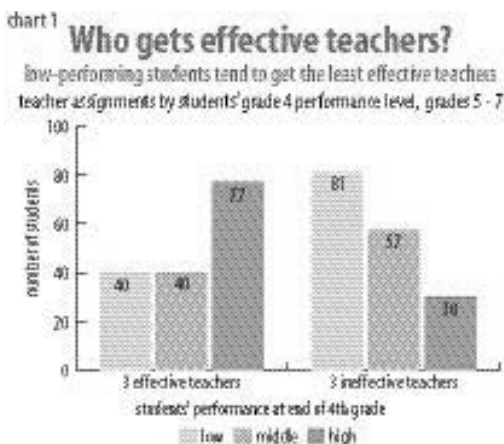
This is just one example of how measuring teacher effectiveness brings crucial information out into the open. Without it, Mrs. Jones would have been out of a job. With it, she was recognized for her real achievements. It is hard to overestimate the importance of this, and frightening

to think of how many teachers like Mrs. Jones are out there, setting high expectations for their students without the benefit of real, objective information that validates their success. Our public schools need all the Mrs. Joneses they can get.

Why identifying effective teachers is so important to under-served students

The more we know and understand about teachers, schools, and students, the more we come to realize that good information about teacher quality can be leveraged to improve almost every important aspect of our education system. And it won't be a moment too soon, because there's a disturbing side to all this new information about teacher quality: The more we know about effective teachers and how important they are, the more we also know *where* they are—mainly in the schools and classrooms of America's better-off white students. Where they are *not* is teaching low-income and minority students, the students who have traditionally been short-changed by the education system, the students who are most dependent on our public schools.

Because only a few states and districts are collecting value-added teacher effectiveness data, we have limited information about the distribution of effective teachers. But we do know quite a lot about the distribution of teachers with characteristics that are, to various degrees, correlated with effectiveness. And what we know is extremely disturbing. No matter which study you examine, no matter which measure of teacher qualities you use, the pattern is always the same—poor students, low-performing students, and students of color are far more likely than other students to have teachers who are inexperienced, uncertified, poorly educated, and under-performing. Many of those teachers demonstrate most or all those unfortunate qualities all at the same time.¹¹ (see "The Opportunity Gap" on page 36)



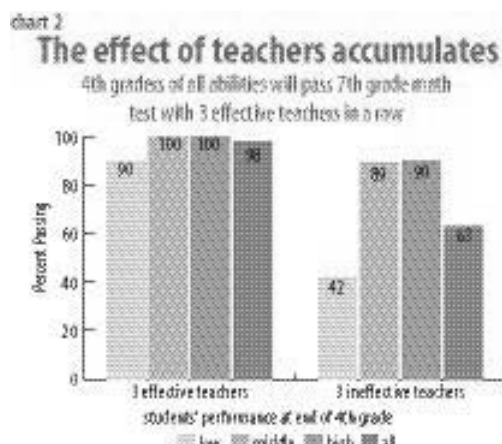
Source: Sifha Babu and Robert Wendt, *Teacher Accountability: ICAP-Based Teacher Effectiveness Judges in the Investigation of Teacher Effect on Student Achievement for a State Assessment Program*, AERA annual meeting, 2003.

Is it any wonder that the achievement gap grows wider the longer students remain in school? Given what we know about the importance of good teaching, what else could we possibly expect?

The crushing impact of the maldistribution of effective teachers to low-income and minority students is hard to overestimate. This is illustrated by recent analyses using value-added data from Dallas.¹² The charts above show the performance of different middle school students assigned to different teachers, and how those assignments affected their performance in math.

Above, we show the performance of two groups of students on the state's 2000 7th grade mathematics test. Both groups entered the 5th grade three years earlier. One group was only assigned to effective math teachers, as measured by the Dallas value-added system, during the 5th, 6th, and 7th grades. The other group was only assigned to ineffective teachers during those three years. The students are further broken out by their previous math performance—low, middle or high—prior to beginning the 5th grade.¹³

We can immediately see how the system is already starting to give up on many of these nine-year-olds. Almost twice as many of the previously



Source: Sifha Babu and Robert Wendt, *Teacher Accountability: ICAP-Based Teacher Effectiveness Judges in the Investigation of Teacher Effect on Student Achievement for a State Assessment Program*, AERA annual meeting, 2003.

high-achieving students were assigned to a string of effective teachers as were low-achieving students—77 to 40. Conversely, more than twice as many low-achieving students were assigned to a string of ineffective teachers as high-achieving students—81 to 30. Low-achieving students, those who need help the most, are being systematically sorted into the classrooms of the least effective teachers.

Chart 2 also shows just how important and potentially damaging these decisions can be. *Every one* of the mid- and high-achieving students who got effective teachers passed the test. By contrast, only 42% of the previously low-achieving students who got ineffective teachers passed the test.

Was this dismal performance the result of their low aptitude, or their ineffective teachers? To find the answer, look at previously low-achieving students who were lucky enough to get three effective teachers in a row—90% passed the test.

This one grade isn't an anomaly. Chart 3 on page 10 shows the pass rate of previously low-achieving students in every grade from 3 to 8 in both reading and math, comparing those who had effective teachers to those who did not.

In every grade and subject, the pattern is exactly the same. Low-achieving students with

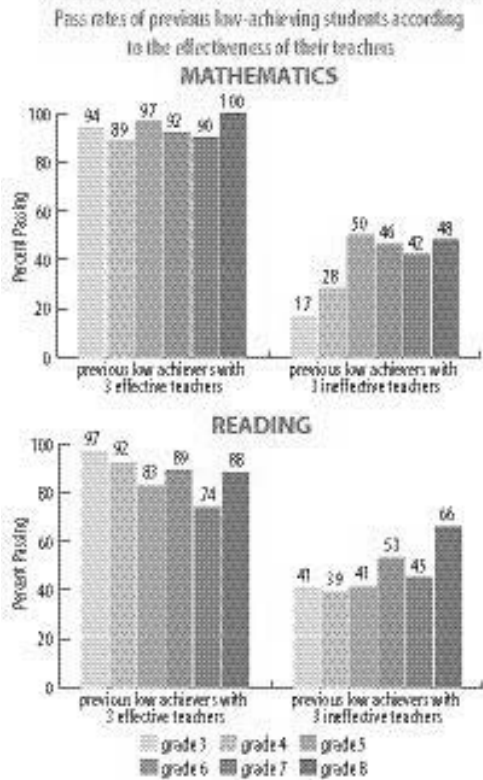
effective teachers are perfectly capable of catching up and meeting the standard. Low-achieving students with ineffective teachers are getting crushed.

But sadly, in every case there are also far more of the latter than the former. Chart 4 shows the exact same group of students and teachers, this time showing the *distribution* of different kinds of students to different kinds of teachers within the Dallas School District. Once again, in every grade and subject, the pattern is exactly the same, but this pattern is far less hopeful. Previously low-achieving students are more likely—in many cases *far more likely*—to be assigned to ineffective

teachers than to effective teachers. This indefensible practice starts at the earliest ages, literally as the students walk in the elementary school door. It continues, year after year, eventually growing so large and so pervasive that one wonders if the small handful of previously low-performing students who were actually assigned to highly effective math teachers through the 8th grade got there through some kind of bureaucratic foul-up or case of misfiled paperwork.


And what happened to those students? As the chart shows, every single one of them passed the math test and almost nine in ten passed the read-

chart 3
Low-achievers become high achievers with effective teachers



Source: Sika Bakula and Robert Wendt, *Teacher Accountability: K-16 Based Teacher Effectiveness Indicators and Investigation of Teacher Effects on Student Achievement for a State Assessment Program*, ABRA annual meeting, 2005.

chart 4
Low-achievers are more likely to be assigned ineffective teachers



Source: Sika Bakula and Robert Wendt, *Teacher Accountability: K-16 Based Teacher Effectiveness Indicators and Investigation of Teacher Effects on Student Achievement for a State Assessment Program*, ABRA annual meeting, 2005.

ing test. Low-achieving students can learn—if *they have good teachers*.

In these two charts we see the foundations of the achievement gap in American public education laid bare. Young children who arrive at school behind and struggle in school, many of whom are poor and/or children of color, are being abandoned by the education system at an early age and systematically shuffled off into the classrooms of our least effective teachers. The result is a vicious cycle of further low performance and further excuses to not give them the teachers they need.

The authors of the Dallas study conclude:

“Clearly, there is a tremendous interaction effect between longitudinal exposure to ineffective teachers and effective teachers when crossed with prior student achievement level. A sequence of ineffective teachers with a student already low achieving is educationally deadly.”¹⁴

From Information To Action

The maldistribution of good teachers means that our new ability to create value-added measures of teacher effectiveness has immediate utility. But we can’t just assume that simply giving schools and policymakers better information about teacher effectiveness will guarantee that they’ll use that information wisely. We have to *make* it happen. Those actions can take many different forms, but in the end they boil down to two related things: We need to use value-added data to improve the effectiveness of teachers generally, and we need to use it to get many more effective teachers into the classrooms of many more low-income and minority students.

To do that, we have to start by making some pretty fundamental changes to the way we think about teachers. Teachers are the fundamental resource of public education. They do the work, they make the difference.

But we tend not to treat teachers as a resource. If you have a vital resource, something that whatever you’re doing crucially depends on, the rational thing to do is have an ongoing process of:

1. Trying to get as much of the resource as possible.
2. Thinking about how best to distribute and deploy that resource to the various parts of your organization, in a way that is strategic, effective, and efficient.
3. Continually evaluating the policies you adopt to accomplish 1 and 2, revising them, improving them, and always focusing on how to maximize the resource and use it as effectively as possible.

Broadly speaking, public education uses this approach for the other really important resource it has—money. As it happens, poor and minority kids have historically been short-changed on that front as well, a problem that continues to this day.¹⁵ But at least policymakers actually sit down every year or two and explicitly decide how to divide dollars up among different states, regions, and school districts. As a result, the nature of those decisions is highly public and thoroughly debated, and there’s actually been some recent progress in making school funding more fair.¹⁶ If state policymakers do a really bad job of dividing up dollars, or don’t provide enough, you can sue them. You might even win.¹⁷

By contrast, there is no such public process of distributing effective teachers, no evaluation of how many we have, no discussion of the best way to deploy them. The reason for this is simple: dollars are easy to measure whereas, until recently, there was no way to accurately measure the value of teachers. So instead of being the result of actual choices and policy decisions, the ultimate amount and distribution of effective teachers has always been the residual effect of various other decisions and circumstances.

Children have just as much of a right to enough good teachers as they do to enough dollars, and we need an attitude and approach to public policy that reflects that. We can't pretend that because teacher value is harder to measure and control than dollar value, we have no responsibility to care about teacher value, no responsibility to take steps to distribute that value fairly and effectively, no responsibility to increase the overall amount of value that teachers provide.

To meet those obligations, we need to think about the whole process of employing teachers for what it is: a market, in the classic economic sense of the word. Teachers do the selling in this market, schools do the buying, and teacher salaries are the price of the service being bought and sold. In theory, markets are efficient. In an efficient market, the price of the service being bought and sold naturally reaches a point where both the buyers and the sellers are satisfied with the quantity and quality of the service being provided and the amount of money it costs to buy.

But a lot of different things have to happen in order for a market to be efficient, and a lot of

those things are missing from the job market for teachers. The market is ignorant, in the sense that it lacks crucial information. And it is in some ways badly over-regulated; there are numerous laws and practices in place that prevent the buyers and sellers from acting in a reasonable, rational way. As a result, the teacher job market is dysfunctional and inefficient, and its flaws manifest themselves in ways that are particularly bad for low-income and minority students.

So, one thing that needs to be done is improve the efficiency of the teacher job market. But that's not the *only* thing that needs to be

done. Efficient markets are good for some things, such as ensuring an adequate supply of what's needed at a reasonable price. But they're bad at other things, such as reflecting broader social values that transcend the specific self-interest of the buyers and the sellers. In this sense, the teacher job market is very much *under*-regulated, in a way that is also bad for low-income and minority students.

To maximize the value of teacher effectiveness information in helping all students and closing the achievement gap, to get more effective teachers and match teachers' strengths with the students in greatest need, we need to do five things.

1. *Make the teacher job market smarter.*
2. *Stop interfering with the market by eliminating old regulations that make things worse.*
3. *Start interfering with the market by implementing new regulations that make things better.*
4. *Level—and then un-level—the playing field for the buyers.*
5. *Increase the overall effectiveness of teachers.*

If states can take these five steps, the result will be a system of training, hiring, compensating, and distributing teachers that is more rational, more fair, more effective, and more focused on helping students who have been traditionally left out in the cold.

1. Make the teacher job market smarter.

It takes a lot of things happening simultaneously to make an efficient market. One of those things is information, knowledge of the price and quality of what's being bought and sold. The price part of the equation in the teacher job market isn't a secret. It's easy enough to find out what salaries are in a given school district, and if a district wants to know how much a teacher wants to get paid, all they have to do is ask. But the job

The teacher job market is dysfunctional and inefficient, and its flaws manifest themselves in ways that are particularly bad for low-income and minority students.

market for teachers suffers from a massive shortfall of information about quality.

Other labor markets are different. For example, the job market for lawyers is quite dynamic in the way it sorts attorneys of varying quality and reacts accordingly. This is partially because crucial aspects of their labor market are less regulated, a subject we'll discuss in more detail in the next section. But it's also because lawyers have the freedom to provide powerful evidence of their value to prospective employers—satisfied clients, significant cases won, well-written briefs, deals negotiated, etc. If a lawyer is effective, the market will respond.

Good teachers, by contrast, are hamstrung in their ability to enter the job market and ask for a price that reflects the quality of their services. Teachers aren't able to gather and disseminate accurate information about their ability to help students learn. Only schools, districts, and states can do that, by creating value-added teacher effectiveness information, and right now most of them aren't doing it. As a result, effective teachers are probably among the most under-valued professionals in the labor market today.

We need to remove this veil of ignorance that we've thrown over the teacher job market. Good teachers deserve to have the ability to demonstrate just how valuable they are when they're looking for a job, or when they're negotiating to stay in the job they already have. Schools need the ability to go into the market and find teachers who are most effective in teaching the particular students those schools serve and the particular subjects where they need the most help. Creating and using value-added teacher effectiveness data would help fill this information void.

In a more open teacher job market, highly effective teachers that are willing and able to help the most challenging students would command more money, because that combination of skills is scarce and will soon be in high demand.

2. Stop interfering with the market by eliminating old regulations that make things worse.

But even if we made the market smarter, schools still wouldn't be able to use that information to their fullest advantage, because the job market for teachers suffers from three different kinds of regulatory interference, all of which are bad for under-served students: price controls, barriers to exit, and barriers to entry.

Price Controls: Not Paying Teachers What They're Really Worth

In the vast majority of public schools the price of teachers is predetermined by a "single salary schedule."¹⁸ All teachers get paid the same single salary, regardless of how well they teach. Adjustments to the schedule are made based on only two factors. One is experience, measured in years of service. The other is education, either through getting a Master's degree or accumulating units of continuing education through college or professional development. The problem is that while teachers get paid more for accumulating classroom time—either years in the K-12 classroom or hours in the college classroom—the connection between increasing these things and increasing effectiveness is at best inconsistent and at worse nonexistent.¹⁹ By contrast, school administrators are prohibited from using salaries to recruit, hire, or retain teachers on the basis of a proven ability to help students learn.

In a more open teacher job market, highly effective teachers that are willing and able to help the most challenging students would command more money, because that combination of skills is scarce and will soon be in high demand.²⁰ But

when you artificially limit the price of something that's in demand, people will provide less of that something because they can't get paid what it's really worth. Thus, price controls on teacher salaries artificially limit the supply of a precious national resource—teachers who can help close the achievement gap.

The idea of paying teachers more money based on the quality or the difficulty of their work is not new. In addition to the incentives we previously described in Chattanooga, other states and districts have experimented with bonus programs, some better than others, based on various measures of teacher quality. They've also experimented with less direct compensation, providing things like loan forgiveness, housing assistance and subsidized tuition for graduate school to teachers willing to work in hard-to-staff schools.

But programs like these remain unreasonably controversial. Historically, many “pay-for-performance” plans have foundered because teachers didn't have confidence that the information used to gauge their performance was objective, accu-

rate and fair. In other cases, the dollars involved simply didn't represent enough extra money to make a difference. Political buy-in has often been contingent on merit pay programs being funded from “new money” of some kind, resources that are often the first to dry up when economic times turn bad and budgets are imperiled. Value added effectiveness data can be used to create teacher compensation plans that are objective, robust and long-lasting.

Barriers to Exit: Keeping Low-Performing and Mediocre Teachers in the Classroom

Our ability to create a functioning, efficient job market for teachers is also hampered by the tenure provisions that provide job security for teachers after a relatively short time on the job. If you look at the state laws that govern teacher tenure, you find that acceptable grounds for not granting what amounts to lifetime job security generally include things like “incompetence,” “immorality,” “conviction of a felony,” “alcoholism,” “moral turpitude,” “criminal sexual miscon-

Missed Opportunities

While significant regulatory barriers may be keeping many potentially effective teachers out of the classrooms that need them, a recent study from the New Teacher Project suggests that the problem may also be a matter of major flaws in the local hiring process itself.¹ A detailed study of human resources practices in a group of large urban school districts, the report found that because of complex work rules relating to internal teacher transfers and rigid job posting requirements negotiated in collective bargaining contracts, the districts often delayed offering jobs to qualified, willing teachers for months and months at a time. Frustrated, many of those teachers took more timely offers from surrounding suburban districts. Follow-up surveys found that the lost teachers would have preferred to work in the urban district. Tragically, for every month that the urban districts delayed in offering a job to willing candidates, the overall talent in the candidate pool shrank, with the best prospective teachers being the first to leave the process or be hired away by other districts.

This report suggests that urban schools may have more opportunities than they realize to hire strong candidates, eliminating whatever excuses that might have existed for falling back on the “any warm body” approach to filling teacher vacancies. Fixing this problem should be an immediate priority – finding enough people who are both willing and able to effectively work with low-income and minority students in urban settings is hard enough without stumbling at the finish line by simply failing to offer them jobs they want to accept.

¹ Jessica Levin and Meredith Quinn, *Missed Opportunities: How We Keep High-Quality Teachers Out of Urban Classrooms*, The New Teacher Project, 2003.

duct,” “mental derangement,” “participating in un-American or subversive doctrines,” or “knowing membership in the Communist Party.”²¹

You generally do *not* find words like “mediocrity,” “barely adequate,” or “just not good enough.” Rather than making the benefits of more-or-less-permanent employment a privilege granted to only the very best, it’s basically a right granted to all but the very worst. This system sends a clear message about our expectations for teachers: we don’t expect them to be excellent, we’ll just settle for no obvious signs of gross incompetence, addiction, perversion or sedition. These and other barriers to removing consistently low-performing teachers create a kind of arteriosclerosis in the job market for teachers, preventing the flow of ineffective teachers out of the system, and thus—unavoidably—the flow of effective teachers into the system.

Barriers to Entry: Keeping Potentially High-Performing Teachers Out of the Classroom

The teacher job market is also significantly regulated at the entry point. All states have laws spelling out minimum requirements for getting a license to teach, usually some combination of a Bachelor’s degree, specific training in teaching methods, and passage of a test of knowledge and teaching skills. These requirements make sense, up to a point. Since we can’t fairly assess how effective teachers are before they’ve actually taught anyone, we need to ensure a baseline level of quality for new teachers by using proxy measures like education, pedagogy, and knowledge in specific subject areas.

There are reasons to believe that current state certification standards can be vastly improved. On the one hand, the level of knowledge needed to pass teacher licensure tests is often depressingly low.²² In that sense, the licensure process does a bad job of guaranteeing professional quality.

On the other hand, a number of researchers and analysts have suggested that specific licensure

requirements for teacher preparation are unduly burdensome, expensive and time-consuming, and serve as an unnecessary barrier to entering the field.²³

The low licensure test standards send a bad signal to people choosing a career path. People who want to be excellent, who want a professional life with colleagues who are also striving for excellence, are going to be driven away by teaching standards that say, “All we expect from you is a bare minimum level of knowledge and skills.” In this way, the licensure process may do a bad job of bringing qualified people into the profession.

In recent years, a number of states have responded to some of these concerns by creating alternatives to the traditional certification process, often aimed at mid-career professionals who want to move to teaching but can’t go back to college for a new degree. At the national level, the Teach for America program recruits top recent college graduates from a variety of academic backgrounds—including those who didn’t major in education—to teach in high-poverty urban and rural schools. Similarly, the New Teacher Project is currently working to help school districts in large urban areas—including New York City, Los Angeles, Atlanta, and Baltimore—fill vacancies in hard-to-staff schools. They recruit both recent college graduates and mid-career professionals to teach in high-need schools, offering an alternative path to the profession that bypasses traditional teacher education in favor of more accelerated training. The project has brought thousands of motivated teachers into schools that need them.

There’s an ongoing back-and-forth debate about traditional vs. alternative certification, centered on whether one process produces better teachers than the other. Neither side can prove

These regulations exist because they are, in some respects, a natural reaction to the shortcomings of an information-starved market.

their case convincingly, because that would take better information about teacher effectiveness, which in most places we don't have.

If all of these regulations—price controls, barriers to exit, and barriers to entry—are so bad for public schools, why do they exist? Because they are, in some respects, a natural reaction to the shortcomings of an information-starved market. Until recently, few states have been able to get good information about teacher effectiveness *into* the market. As a result, they've spent their time passing regulations to keep bad information *out* of the market. For example, lacking hard data about teacher performance, many states and schools have created elaborate procedural safeguards to prevent teachers from being fired for potentially arbitrary judgments of non-performance, or for being paid extra for potentially arbitrary measures of good performance. So we see that the information void in the market isn't just preventing us from doing things that are smart, it's actually having the perverse effect of causing things that are counterproductive.

The new ability to create value-added measures of teacher effectiveness changes all of this. If we bring that information into the market, and if we pare down the elaborate regulatory structure that was created to compensate for the lack of that information, we can create a better job market for teachers, one that encourages strong candidates to enter the field, pays them what they're worth if they succeed, and moves them out of the system if they don't.

3. Start interfering with the market by implementing new regulations that make things better.

So far we've talked about two things: injecting vital teacher effectiveness information into the job market, and getting rid of regulations that limit the ability of schools to hire the teachers they need. These are important, necessary changes.

Indeed, it's hard to seriously think of a way to close the achievement gap without them.

But they're not enough. In fact, implementing these policy changes and *only* these changes could actually have the effect of making things *worse* for low-income and minority students. These changes are all about making the teacher job market more dynamic, more efficient—more free. Freedom has great virtues, but it also has its own set of consequences. In the case of the teacher job market, greater freedom and information for buyers and sellers means greater ability to act in self-interest. Unfortunately, the self-interest of individual schools and teachers doesn't always line up with our broad education priorities the way we'd like them to.

A recent study of teacher job market decisions helps illustrate this problem. Using a comprehensive database of employment records in Texas from 1993-1996, researchers examined tens of thousands of instances where teachers left a job in one school district for a job in another.²⁴ By looking at teacher pay and experience, as well as examining the characteristics of the districts the teachers left and those they entered, the researchers were able to draw conclusions about what factors drive teachers to change jobs.

The authors found that the biggest factor was student achievement: teachers tended to move away from districts with many low-performing students and toward districts where performance was higher. They also moved away from high-poverty schools toward lower-poverty schools. Most disturbingly, race was also a significant factor. White teachers, who made up the large majority of teachers studied, tended to move away from high-minority schools—even after factoring out the influence of student achievement.

A study from Georgia found similar results, finding that teachers who moved to new schools within a district went to schools with higher levels of student achievement and with lower propor-

tions of African American students and students in poverty. The authors noted, “The most striking finding is the relationship between the percentage of black students and teacher turnover. White teachers, who compose 80 percent of the Georgia teaching force, are much more likely to leave schools that serve higher proportions of black students, and these turnover rates increased dramatically over the late 1990s.”²⁵ Similar patterns of teachers transferring away from poor and minority students have been found in New York.²⁶

These studies suggest that teachers exercising free will in the job market tend to gravitate away from the low-achieving, low-income, and minority students who depend most on good teachers to support their learning. Not all teachers of course; many dedicated teachers act on a sincere desire to help such students. But the overall trend is disturbing, and it’s consistent with a lot of anecdotal

evidence about teacher hiring. Many high-poverty urban and rural schools are constantly grappling with high teacher turnover rates, struggling every fall to fill teacher slots with even minimally qualified candidates. Unfortunately, many schools with low-performing students also have poor working conditions for teachers—larger class sizes, outdated facilities, and other problems.

Wealthy, high-achieving suburban schools, by contrast, can have 50 applicants for every open position and often retain good teachers for years or even decades. As one teacher said in explaining why she didn’t want to work in a low-achieving school:

“You have to be a combination of a social worker and Mother Teresa to work in those schools...I worked so hard to get my license, I did all this schooling, and the last thing I heard, America was a country of free choice.”²⁷

What Tennessee educators say

Diana Green was a new principal at a middle school that was deep on Tennessee’s “needs improvement list,” ranked among the lowest in the state.¹ When the state began providing TVAAS data detailing student gain scores organized by subject and by teacher, it was “like an epiphany.” The data showed which teachers were most effective and which were not, as well as which subject areas and grades needed the most improvement. In response, she dedicated more instructional time to the lagging subject areas, and “started placing teachers strategically within the school, assigning the most effective teachers to the students and subject areas that needed the most help.” TVAAS data also changed her approach to hiring and personnel management, giving her more objective information to gauge which prospective job candidates from other schools in the system would be most effective, and which current teachers were the best candidates to be re-hired and retained in future years. A year after these TVAAS-driven reforms were in place, the school was off the “needs improvement list.” Three years later, it ranked among the top 20% in the state in terms of annual student gains in achievement.

Emily Baker is a principal at an urban elementary school in Chattanooga that serves many low-income and minority children. She uses TVAAS data as a tool for staff development and teacher improvement. “Every year I sit down with the data,” she said, “looking at which teachers are most effective in each subject, and where teachers have an opportunity to improve. Say Teacher A is getting 150% of the expected growth for 4th graders in reading, while Teacher B is only getting 85% of expected growth. I’ll have those two teachers sit down together and talk about their teaching practices and what material they’re covering, so Teacher B can learn from Teacher A and find out why she’s so successful.” TVAAS data provides the basis for a school-wide system of constantly learning from success, sharing that information, and improving. “Once the scores are looking you in the face,” she said, “you can’t deny what that the data tells you.”

¹ Telephone interview, July 25, 2003.

It's easy to imagine a situation where affluent suburban school districts, armed with better information about teacher effectiveness and more flexibility to pay teachers what they're worth, become even more skillful and effective at luring away the best and the brightest from high-poverty school districts, exacerbating a problem that's already at crisis levels across the country.

This is a tough problem, and there are some fundamental structural issues limiting our ability to fix it. One is the basic level of transaction in the marketplace between teachers and school districts. Certified teachers are free to sign on with any district they like, and districts are free to hire any teacher they like. So our ability to influence the hiring of teachers at the district level by direct regulation is very limited. That problem is better attacked from the financial side of things, which we'll get to later on.

The distribution of teachers *within* districts, however, is a different story. School districts potentially have significant capacity to determine where their teachers are deployed, much as businesses assign employees to different divisions within a company based on the organization's overall strategic goals. A recent study of the distribution of teachers in upstate New York found that two-thirds of the difference in teacher quality between schools was the result of differences among schools *within* districts, as opposed to only one-third *between* districts.²⁸ The issue of within-district assignment of teachers is particularly important for low-income children, since these students tend to be concentrated in large, unified, urban school districts with many schools. While there are over 14,000 active school districts nationwide, over 10% of all public school students living below the federal poverty line are educated in just five large urban school districts. The majority of all poor children are educated in 450 districts, which have an average of 57 schools apiece.²⁹

This means that, in theory, there are a lot of opportunities for school districts serving low-income and minority students to deploy their human resources in a manner that is rational, strategic, and fair, in a way that makes sense given their goals to raise achievement for all students and close the achievement gap.

Unfortunately, in practice such opportunities are severely limited. The employment contracts that many districts currently have with their teachers cede away the discretion of schools in choosing who to hire. These arrangements include strict seniority-based hiring protocols. If a teaching spot opens up in a school filled with lots of high-income, high-achieving students, then the district can't fill that position as it sees fit. It has to allow teachers with sufficient seniority to take the job, even if the district thinks that's not in the best interests of the students. Since more experienced teachers tend to gravitate away from low-achieving, low-income, and minority students, schools that serve those students are left to constantly struggle with high turnover and inexperienced staff—two factors that have a significant negative effect on student learning.

The internal financial systems used by most school districts support and exacerbate this sad state of affairs. Rather than giving each school the same amount of money, on a per-student basis, to hire teachers, most districts effectively give schools the amount of money they need to pay the teachers they have. A recent study of school district finances in three large urban districts—Cincinnati, Seattle, and Baltimore—found that this practice results in significant disparities in funding between schools within a given district.³⁰ High-poverty schools received significantly less per-student funding than low-poverty schools, because the high-poverty schools employed many more of the least experienced, least costly, teachers. One school in Cincinnati received almost *a million dollars less* per year under this system

than if it had simply been given the district average level of per-student funding. Meanwhile, the most affluent schools—the schools where teachers most want to teach—also get to pay their teachers the most money.

This study illustrates how seniority-based transfer provisions, single salary schedules, and within-district financial arrangements combine to create a system where experienced teachers get paid the most money to work with the students they most want to teach. What we need is a system where teachers get paid the most money to work where the *need* is greatest. In other words, we need a system of reasonable tradeoffs between teacher compensation, autonomy, and the needs of students. We can start by allocating funds to schools within school districts on *at least* an even per-student basis, so that schools serving low-income students start with the same amount of resources to hire effective teachers. We can continue by eliminating restrictive seniority-based hiring provisions that prevent school administrators from matching up teachers and students in ways that help close the achievement gap, and instead give principals authority to hire the best available teacher for the job.

Some school districts that are able to make strategic choices about teacher deployment are doing so. For example, the high-poverty urban schools in Chattanooga that have recently had success in using financial incentives to bring more high-performing teachers in have also worked to move less effective teachers out, using both voluntary and involuntary transfers to reconstitute their teaching force to better focus on the needs of the children.

The Wake County Public School District in North Carolina uses a different approach. The leadership there recently adopted a student assignment policy designed to ensure a diverse student body in each school in the district, including standards that no fewer than 25 percent of stu-

dents in each school are low-performing and no fewer than 40 percent are low-income. Since the policy was adopted in 2000, the achievement gap between poor and non-poor students in Wake County has narrowed by almost five percentage points.³¹ If economically diverse school districts create a balanced mix of students in each school, preventing broad disparities in terms of student achievement and income, they can blunt the tendency of teachers to make location decisions that reflect those disparities.

Charlotte-Mecklenburg provides another example, pursuing a policy of equity among schools in terms of teacher qualities. They try to ensure that teachers in schools designated as being at-risk of academic failure are similar to those in high-achieving schools in terms of years of experience, Master's degrees, and National Board Certification. To achieve this goal, the district recently instituted a policy of prohibiting 26 high-achieving schools with sufficient numbers of highly qualified teachers from hiring additional teachers away from other schools within the district.

In announcing this policy, the district noted that “in order to provide an excellent education to every student in the district, the talent of the district’s teaching staff must be spread throughout the system.”³² By limiting within-district teacher migration that would be harmful to vulnerable students, Charlotte-Mecklenburg is using information about teacher quality to manage its human resources strategically in order to ensure that low-income and minority students have their fair share of good teachers.

Reaction to this policy was not altogether positive. The president of the Charlotte Teacher’s Association said, “It shows they don’t have

We need a system of reasonable tradeoffs between teacher compensation, autonomy, and the needs of students.

enough faith in teachers to expend their once-a-year ability to improve [their own] working conditions.”³³ But really, this is exactly wrong. It’s because they have faith in teachers—faith in their importance to student learning, faith in their ability to help children who need the most help—that they decided to put some limits on the ability of teachers to exercise total free will within the district, to interfere with the teacher job market on behalf of the children who depend most on good teachers.

4. Level—and then un-level—the playing field for the buyers.

The changes in management and financing described above would go a long way toward improving the way we deploy teachers within districts. But what about *between* districts? While it’s reasonable to think that districts can adopt policies designed to optimize the distribution of teachers among schools, it’s probably unreasonable to think that states will be able to assert the right to assign teachers to districts. At the district level, teachers aren’t like dollars. Dollars go where we tell them to go, teachers go where they want.

The solution, then, is to put the school districts that serve high numbers of low-income students—many of which have been identified by the No Child Left Behind Act (NCLB) as needing improvement—in a position to accomplish with persuasion what we can’t accomplish with coercion. Those districts need enough money to go into the teacher job market and pay the market price, a price that reflects both the quality of the teacher and the difficulty of the assignment. They also need enough money to offer a working environment that will attract and support effective teachers by providing adequate facilities, reasonable class sizes, access to high-quality professional development, effective leaders, etc.

Unfortunately, these districts don’t have that much money right now. In fact, school districts

that educate the largest numbers of low-income and minority students receive, on average, hundreds of dollars *less* per student than school districts that educate the fewest numbers of low-income and minority students.³⁴ Those who need the most receive the least.

Fixing school funding disparities can be helliciously difficult. Funding schemes that short-change students living in low-wealth, high-poverty areas have resulted in contentious lawsuits in all but a handful of states. Some of these disputes took decades of litigation and hundreds of millions of dollars to resolve, and many continue to this day. But difficulty and expense don’t diminish the importance of funding education fairly. We can’t close the achievement gap until we identify the most effective teachers and match them up with the students who need their help. We can’t do that unless we pay those teachers the real market price for their services, and we can’t do that unless the school districts doing the paying actually have the money to write the checks.

Arguably, one reason state policymakers have been so reluctant to close these funding gaps is that they lacked information showing that additional funding could bring about a better education for students. Before value-added information came along, many people thought that low-performing students were essentially a lost cause, and that one teacher was as good as another. Now we know just how wrong those ideas were. The children in under-funded districts *can* learn, if they have effective teachers. And that knowledge has great implications for the pressing need to fix funding disparities, so school districts have enough money to hire the teachers they need, and address the working condition problems—facilities, safety, workload, availability of professional development, etc.—that drive away effective teachers.

That will cost money, and right now the financial playing field is uneven. It’s slanted in

favor of schools that serve middle and upper-income white children. But even if high-poverty schools finally get the same amount of money as their more affluent counterparts, they'll still be in the situation of offering teachers an equal salary to do an unequally difficult job. We can narrow the achievement gap by making sure that low-income and minority kids have access to the same level of financial and instructional resources as their peers, but we likely won't close the achievement gap until we give high-poverty schools the *extra* funding they need to really do their jobs.

In other words, once we level the playing field, we need to keep going and un-level it in the other direction, by providing extra funding to school districts based on the number of low-income children they have to educate. This isn't a novel idea; in fact it's already been implemented by both the federal government and most states. Title I, the single largest federal K-12 education program, directs more than \$12 billion a year to schools to improve education for poor students. Similarly, over three-fourths of the states have modified their school funding formulas to generate additional financial resources for schools that serve low-income children.³⁵

If that's the case, why hasn't it already worked? First, because too many high-poverty school districts continue to provide low-level, low-quality instruction to their students. If we start with very low expectations for children, we can almost guarantee the children won't exceed those expectations. Second, because the lack of value-added teacher effectiveness data limits the ability of high-poverty schools to identify the teachers they need to hire. Third, because artificial limitations on teacher hiring, assignment, and compensation limit schools' ability to hire those teachers, pay them what they're worth, and assign them to the students who need them. Fourth, because states still don't provide enough money through their current poverty-based

funding programs. While most states have some kind of funding stream directed to higher-poverty districts, many provide what amounts to only token funding levels.

If we use value-added data to give school districts educating low-income students the information they need to identify effective teachers, the freedom they need to hire and deploy them, and the resources they need to pay them the market rate and provide them with supportive working conditions, those districts will finally have a fair chance to close the achievement gap for their students.

5. Increase the overall effectiveness of teachers.

Everything we've discussed thus far has been in a zero-sum context, as if there's a fixed population of teachers, some of whom are effective, some of whom are ineffective, and most of whom are in between. We know that the current distribution of teachers among different kinds of students is badly skewed against low-income and minority children, and steps must be taken to fix that problem if we're going to close the achievement gap.

Such steps may be initially difficult. But fortunately, teacher quality is not a fixed commodity. If we can significantly raise the overall level of teacher effectiveness, if we can increase the collective level of skill, talent and effectiveness possessed by America's education workforce, then solving our current distribution problem becomes less a matter of redistribution and more a matter of directing new resources to where they're needed most.

Once we level the financial playing field, we need to keep going and un-level it in the other direction, by providing extra funding to school districts based on the number of low-income children they have to educate.

For starters, value-added data can be used to better understand and focus on the specific characteristics and attributes of effective teachers. There is a lively ongoing dialogue in the academic community on this point, centered on the question of which factors—education, experience, certification, content knowledge, etc.—most contribute to good teaching. But there’s surprisingly little solid, convincing data to resolve the debate.³⁶

Value-added information can change this. School districts in Tennessee, Dallas, and Minneapolis have used their value-added data to study the qualities of effective teachers. When value-added measures become more widespread, vast new opportunities for research will open up, letting us eventually move beyond debating which factors are associated with effective teachers, and move toward devoting more of our time and energy to actually improving effectiveness itself among both our practicing teachers and those preparing for entry into the profession.

Improving effectiveness through professional development

Already, many education dollars are committed each year to the ongoing improvement of practicing teachers. Virtually every school sets aside a few days or afternoons a year for professional development. But too often these are temporary and scattershot. Rare is the case where a school follows up to see if the programs were truly effective—to see if the students being taught by the teachers who took the program are doing any better than they were before.

Schools in need of improvement are going to need a much smarter, more effective approach to in-service teacher training than this. NCLB means that schools will know which students are having problems in which grades and which subjects. In fact, NCLB requires schools that need improvement to spend at least 10% of their Title I funds on professional development for teachers.

If schools have good data about teacher effectiveness—if they can see how effective their teachers are *before* going through professional development compared to how effective they are *after* professional development—they can find out which programs are worthwhile and which are not, matching up teachers with the specific programs that best fit their particular needs.

Dallas Independent School District uses teacher effectiveness data to design individualized training programs for teachers who are struggling with student performance. The data system allows them to see which specific skills, subject areas, and kinds of students need improvement the most, providing information that leads to a customized plan for professional development. A number of school leaders in Tennessee report similar efforts. They use value-added TVAAS data to determine, for each teacher, which subjects and students most need to improve and what professional development strategies will best help the teacher accomplish those goals.

Improving higher education

Value-added data can also be used to improve the process of teacher training prior to entering the classroom. State policymakers are well-positioned to make this happen. In addition to being the primary buyer of teachers (nine out of ten schoolchildren are enrolled in public schools) state governments are also a major producer of teachers. Roughly 70% of all new teachers receive their postsecondary education in public, state-funded colleges and universities.³⁷ If states, through their management and regulation of colleges and universities, can improve the quality of the emerging teacher workforce, then effective teachers will be more abundant.

Once again, this is an issue where creating value-added teacher effectiveness information plays a vital role. To date, there has been very little in the way of good information about the

quality of the teachers that universities produce. Local schools that are really paying attention may have a general sense of which colleges and universities are supplying them with effective teachers, but they are unlikely to have any hard data to back that up. Too many schools of education, meanwhile, pretty much just let their graduates go out into the workplace to do as they will, with little interest in knowing whether these new teachers actually help students learn once they get there.

There has recently been some progress in trying to fix this problem. When Congress reauthorized the federal Higher Education Act (HEA) in 1998, it included provisions requiring states and schools of education to report, for the very first time, on the quality of the teachers they

produce. States were asked to hold education schools accountable for the results by identifying those teacher preparation programs that need to improve. Lacking data on the actual effectiveness of education school graduates, most states used a proxy measure, the success rates of graduates on teacher licensure exams. The general response from states to HEA was underwhelming, to say the least. In the first year, states reported that a grand total of one education school—out of over 1,300 nationwide—was “low-performing,” with another 13 “at-risk.”³⁸ Congress is currently working to tighten up the HEA provisions.

One very promising effort to link teacher education to teacher effectiveness is being led by the non-profit Carnegie Corporation of New York, which recently launched an initiative to

Complementary not Contradictory: How value-added works with NCLB

Value-added measures are, fundamentally, indicators of progress. They tell us how much additional learning has occurred within a fixed time period, and they're adjusted to isolate each individual teacher's contribution to student learning.

By contrast, under the federal Title I law (commonly referred to as “No Child Left Behind” or NCLB), schools and school districts are evaluated on the percentage of students who have met a common state-wide proficiency standard. The “AYP” system under NCLB is designed to give information about whether students have acquired the skills and knowledge they need given their current grade level.

Are these two systems—one based on progress in performance, one based on absolute performance—contradictory? They are not. They are, in fact, complementary.

Take, for example, a hypothetical elementary school for students in grades K-4 that did not meet the state's AYP target. The principal analyzed school data using the value-added model and found that the effect of instruction in Grade 3 was below par: the students only made 50% of the expected progress that year. By contrast, the effect of instruction in Grade 4 was very good: these students made 125% of their expected progress. The Grade 4 gains were enough for students make up *some* of the progress they lost in Grade 3, but, significantly, not all.

Value-added teacher effectiveness data is extremely useful in this situation, because it helps principals and teachers identify the strengths and weaknesses of instruction in the school. Value-added shows that the issue is in Grade 3, not Grade 4, so the principal knows where to focus efforts for professional development and instructional improvement. In schools that are making AYP, principals can use value-added data proactively to ensure that they are on target for continuous improvement.

Absolute measures of performance, such as those under NCLB, are essential for identifying not just progress, but whether students have the skills and knowledge they need to succeed. Value-added is just as vital for helping districts and schools make well-informed decisions about instruction to make sure all students meet state standards.

increase the overall level of teacher quality by developing state-of-the-art schools of education.³⁹ The first design principle for the education schools that receive funding under this program is “a respect for evidence, including attention to pupil learning gains accomplished under the tutelage of teachers who are graduates of the program.” This initiative recognizes the crucial fact that education schools can’t get better unless they

use value-added data to gain some sense of what they’re trying to improve.

We need what amounts to a whole new attitude in higher education, a sense of responsibility for recruiting and training quality teachers, and a genuine commitment to finding out whether or not those teachers are successful with students.

The Ohio Partnership for Accountability, which includes all 51 schools of education in the state, the Department of Education, and the Board of Regents, recently announced a project to use value-added teacher effectiveness data to better understand, study, and improve university teacher preparation programs. The dean of the University of Dayton school of education, the co-chair of the project, noted, “We’ve got tons of graduates, but no mechanism to assess the relative effectiveness of those teachers.”

Implemented with the participation of the state’s teachers unions, Ohio’s project is the first state-wide effort of its kind and should be a model for all other states to follow.⁴⁰

In another example, the Renaissance Partnership for Improving Teacher Quality involves eleven university schools of education in a number of states that work with partner K-12 schools to assess the effectiveness of new teachers by examining a variety of evidence, including student learning gains.

Our colleges and universities also have a lot of potential to improve the quality of the teaching workforce by recruiting a larger number of talented students into the profession. For example, the Texas A&M University System recently set ambitious goals to improve the amount, diversity, and quality of its teaching candidates. The results have been impressive—the total number of new teaching candidates that passed the state certification test increased by 20% from 2000 to 2002, while the number of African American teacher candidates increased by 116% and the number of bilingual/ESL candidates jumped by 84%. The number of new teacher candidates also increased by 64% in special education, 41% in math, and 34% in science. And despite an increase of hundreds of new teaching candidates, the overall pass rate on the state licensure exam did not decline.⁴¹ This shows that it’s possible to raise the quantity of teachers and enhance diversity without lowering standards of teacher quality.

We need what amounts to a whole new attitude in higher education, a sense of responsibility for recruiting and training quality teachers, and a genuine commitment to finding out whether or not those teachers are successful with students. Education schools that adopt this attitude will reap the benefit of increased demand at both ends of the pipeline—more students will want to enroll, more schools will want to hire their students. Those that don’t will be forced to adapt and change.

Helping Teachers and Managing Schools

Better professional development and pre-service education are important, but many important increases in teacher effectiveness don’t take place in the university lecture hall or the after-school professional development seminar, but in the classroom itself. Using value-added information rationally in the day-to-day management of public schools and classrooms will, in and of

itself, have the effect of improving instruction and increasing the quality of teachers. Here are three examples of how:

- *Teacher effectiveness data can help teachers help themselves.* Good teachers work and learn and get better. They “reflect on their practice,” to use a phrase from education circles. Value-added data can help this process immensely. It takes a lot of the guesswork out of self-evaluation. It shows teachers which students are making the most—or least—growth in which subject areas. In breaking down teacher effectiveness information by topic or concept, the data often shows that teachers are quite effective in some subjects areas, but less so in others. It also shows that teachers are often more effective in helping different kinds of students, in terms of prior achievement, than others.

- *Teacher effectiveness data can help teachers help each other.* One of the reasons new teachers struggle is that they’re often thrown into the classroom with little in the way of guidance or assistance. To fix this problem, a number of states and districts are putting in place mentoring and “induction” programs for new teachers. Value-added information can help schools to match up new or struggling teachers with their most effective colleagues, ensuring that they learn from the best.

In this way, teacher effectiveness data sets the stage for a work environment that is much more geared towards sharing knowledge and expertise within the school building. Teachers that are particularly effective and skilled in working with low-performing students, for example, can help their colleagues who struggle with those students, while other teachers may excel with higher-performing students, or in certain subjects, or even parts of subjects. The point being, almost all teachers have something important to share with their colleagues. Value-added data helps teachers and principals

discover what that “something” really is.

For example, the Public Education Foundation in Tennessee sponsored a study of teachers who were consistently rated as highly-effective using TVAAS. After videotaping the effective teachers in the classroom, the teachers themselves then provided commentary and critiques on their own work, pointing out their strengths and weaknesses, explaining particular techniques and approaches that had proven successful. These tapes are now being used by other teachers as part of ongoing professional development, allowing them to learn from the success of their peers.

- *Teacher effectiveness data can help create a professional culture that is more oriented toward achievement and continuous improvement.* We want our public schools to be full of people who want to be good and get better. Although you might not guess it from reading the newspaper, many teachers really like having their performance assessed. They see it as a way to identify their strengths and weaknesses, to advance and distinguish themselves from their peers. As one teacher recently asked after viewing teacher effectiveness data from another state:

“Why can’t we get information like that at our schools? I would really like to see a bar graph like that for my own teaching. I even think teachers could be fairly evaluated on such data.”⁴²

Teachers like this are as frustrated as anyone by our current information-starved system. They’re never completely satisfied with their students’ performance; they’re always looking for ways to get better, even though the system doesn’t recognize their improvement.

There are a lot of teachers like this in our public schools, but not enough. There are also many more who would demonstrate these qualities if

schools were managed in a way that encouraged that approach to their jobs. Unfortunately, too many schools operate in a culture that seems to take the statement “you can improve” as a harsh criticism of their performance in the present, rather than a vote of confidence in their potential performance in the future.

We can all agree that student learning is the overriding goal of education so we should want to know about teachers’ effect on that learning. But other kinds of information about teachers are certainly useful and important as well.

There are also, inevitably, going to be some teachers who don’t have the capacity or the inclination to meet high standards of effectiveness. Some people don’t work well in an environment that requires rigorous evaluation and high standards of performance. As a nation we just can’t afford to have teachers with these qualities in our public schools.

Using teacher quality information to identify, recognize, encourage, and reward effective teachers can change the professional culture of teaching in a way that places far more emphasis on high achievement and continuous improvement. This will:

(1) help transform teachers with the potential to be high-achieving; (2) bring more people into the profession that thrive in that environment naturally; and (3) move people out of the profession who can’t or won’t make the change.

If This Is Such an Obviously Good Idea, Why Isn’t Everybody Doing It?

Up until recently many states haven’t had the standards, testing, and computers they need to create real measures of teacher effectiveness. That excuse is now rapidly fading away. But other obstacles remain. Some states actually have laws on the books that make it *illegal* to use infor-

mation about student learning to find out how effective teachers are in helping students learn, essentially moving teacher effectiveness data into the realm of forbidden knowledge. For example, Indiana complies with NCLB by giving its students a test called ISTEP. It also has a law requiring school districts to have a plan for evaluating the performance of its teachers. That law couldn’t be more clear. It says:

“...the plan may not provide for an evaluation that is based in whole or in part on the ISTEP test scores of the students in the school corporation.”⁴³

Even if states don’t require schools to hide and conceal vital information about teacher effectiveness, some school districts agree to do so voluntarily, through collective bargaining. These districts have negotiated contracts with the local teachers union that either prohibit entirely the use of student test scores to evaluate teachers, or make such an evaluation dependent on the consent of the teacher in question.⁴⁴

Why construct these elaborate walls around what is obviously vital information?

There are generally four objections that are raised. We hear them a lot. But they ultimately don’t make the case.

OBJECTION #1. *It’s unreasonable to hold teachers accountable for tests of student learning and **only** tests of student learning, because other things about teachers are important to know as well.*

This isn’t an objection so much as a true statement masquerading as an objection. We can all agree that student learning is, in the end, the overriding goal of public education and so we should want to know about teachers’ effect on that learning. But other kinds of information about teachers are certainly useful and important as well. Schools can, for example, supplement value-added measurements of effectiveness with

observations and evaluations of classroom teaching and practices. Any kind of high-stakes personnel decisions ought to give weight to the judgment and observations of the persons responsible for managing teachers and running the school.

Some very detailed processes have been developed to examine teacher practices by looking at portfolios of lesson plans and assignments, videotapes of classroom teaching, in-class observations, and other means of assessing the teaching process. This kind of information can be quite valuable for helping teachers examine how they teach and become more effective. But while evaluations of pedagogy, classroom management, and other teacher practices can add to our understanding of teachers, we have to be careful that they never *replace* or even *supersede* measures of student learning. Effective teaching is often idiosyncratic, an art as well as a science. Different teachers may use very different approaches to achieve the same bottom-line results. In the end, we should be mindful that “effective” doesn’t mean “seeming effective” or “acting like an effective person.” It means being effective, getting actual results with students.

OBJECTION #2. *It’s unreasonable to hold a teacher accountable for student learning because student learning isn’t fully under a teacher’s control.*

There are many things that occur outside the classroom that influence how much students learn, so how can you hold teachers accountable for the result?

It’s a good question, but it’s a question with an answer. The value-added systems used in places like Tennessee and Dallas are specifically designed to adjust for outside factors. The TVAAS system controls for each student’s individual learning history. So if a student’s family background, aptitude, motivation, or any other

possible factor has resulted in low achievement and minimal learning growth in the past, all that is taken into account when the system calculates the teacher’s contribution to student growth in the present. A teacher who produces average growth with students who had previously grown very slowly will receive a high rating; a teacher who produces identical success with students who had previously excelled would receive a lower rating. The system recognizes that every student is different and adjusts accordingly. It also uses multiple years of data to smooth out the effects of year-to-year differences. Value-added methods isolate the teacher’s influence on student learning, producing a rating based on teachers and teachers alone.

These safeguards *protect* teachers from being unfairly evaluated on what isn’t under their control. And the need for these protections is steadily growing. With schools being held more and more accountable for the performance of their students on standardized tests, it’s no longer a question of whether or not we equate teacher effectiveness with student performance. That’s going to happen one way or another; the only question now is whether or not we do a good job of it. Teachers need and deserve a carefully designed value-added system to accurately, fairly gauge their effectiveness.

OBJECTION #3. *It’s unreasonable to hold teachers accountable for student learning because the measure of student learning is imperfect.*

The third big objection to creating and using teacher effectiveness data is that even if you can effectively isolate the teacher’s influence on student test scores, the tests themselves are flawed. People note that standardized tests are imperfect measures of student learning, and say that makes the tests inappropriate measures of teacher quality.

Information that isn't perfect can still be very useful. A teacher might be rated as 90% more effective than average—almost twice the norm—in one year, 80% in the next, 95% the year after that. These differences might be a function of changes in actual performance, imperfections in the test, or both. But in each year it is quite clear that she's a great teacher.

It's quite true that tests aren't perfect instruments of measuring student learning. Any test covers only a sample of the total amount of knowledge and skills a student has acquired.

Some states address the general problem of test imperfection by using multiple assessments to evaluate student learning, reasoning that using more than one testing instrument provides a fuller picture of how well students are doing. The results from these additional tests can, in turn, be used to generate additional information about teacher effectiveness. Many schools use periodic mid-year tests to gauge

progress. This information can also be fed into the value-added system. Tennessee uses multiple years of data in their TVAAS calculations to reduce the effect of anomalies that may crop up.

More data can make teacher effectiveness measures better, but not perfect. *But this imperfection is by no means an argument against using student test information to measure teacher effectiveness.* Information that isn't perfect can still be very useful. A mercury thermometer might only be accurate within plus or minus 5 degrees, but if it says minus 10 on a January morning, you're not leaving the house without a winter coat. A teacher might be rated as 90% more effective than average—almost twice the norm—in one year, 80% in the next, 95% the year after that. These differences might be a function of changes in actual performance, imperfections in the test, or both.

But in the end the specific reason doesn't matter, because in each year it is quite clear that she's a great teacher and needs to be paid, supported, and assigned accordingly.

There's an old saying that "perfection is the enemy of the good." This is just plain common sense—don't let the fruitless pursuit of an unattainable ideal keep you from something that's imperfect but really worthwhile. Opponents of getting new information about teacher quality have essentially flipped this idea on its head and used as a rhetorical strategy.

One reason we can reject these arguments and use teacher effectiveness information despite its imperfections is the fact that research indicates that teacher quality varies a lot. The difference between high-performing teachers and low-performing teachers just doesn't fall within the statistical margin of error—it doesn't even come close. Even after making every reasonable effort to adjust for non-teacher factors and give people the benefit of the doubt, we still find that some teachers are, year after year, much more effective than others.⁴⁵ As a school administrator in Dallas said,

"Use of the system for ten years has clearly demonstrated that effective and ineffective classrooms can be clearly, reliably, and fairly identified based on the achievement of their students... The District considers this to be a closed issue."

OBJECTION #4. *OK, fine—it's just flat-out a bad idea.*

Eventually you get beyond the technical arguments against measuring teacher effectiveness to what is likely the real reason—many people simply don't want to do it. They just think it's a bad idea.

Some of the objections come from administrators who really have no interest in upsetting the status quo and acknowledging how some of

their teachers are much more effective than others. If they did, they would have to act on that information, to do something about it, and that would involve making difficult and uncomfortable choices.

And nowhere is the objection to finding and using teacher effectiveness data more strongly articulated than in the statements of teachers unions—particularly when the issue involves how and how much teachers get paid. For example, this is what the president of the local teachers union had to say when leaders in Chattanooga first announced the plan to pay high-performing teachers more money to teach in schools that needed the most help:

“We think [the plan] is a potentially divisive issue for the system”—Bill Bowman, president of the Hamilton County Education Association.

-Chattanooga Times Free Press, March 14, 2002.

One union official chose to address the issue of differences in teacher effectiveness by denying that such differences exist:

“I hesitate to say one teacher is better than another teacher.”—Jeff Cloutier, Executive Director, United Teachers of Richmond, California, in response to questions about policies that require schools to lay off teachers based on seniority, rather than merit.

-Contra Costa Times, April 8, 2003.

And here’s the outgoing head of the California teacher’s union, explaining why his organization doesn’t support paying teachers more money to do more difficult work, despite the fact that this stance hurts low-income children:

“The state’s largest and most powerful teachers union, the California Teachers Association, eschews giving teachers financial or other incentives for a difficult assignment. ‘It would be psychologically bad,’ said Wayne Johnson, CTA’s

outgoing president, adding teachers would resent colleagues with the same experience making more money. ‘The teachers in the [more affluent areas] would say we’re doing the same job.’ Johnson acknowledged that ‘poor kids are getting jobbed,’ arguing that smaller class sizes and more individual control over curriculum would attract more teachers to those neighborhoods. But Johnson wants those reforms for every school.”

-Oakland Tribune, June 17, 2003.

This last quote is Exhibit A in the case that our education system’s priorities are dangerously skewed toward what’s good for adults to the detriment of what’s good for children. The head of the teachers union in the largest state in the nation admits, freely and for attribution, that union objections to giving teachers more money to do work that is more difficult are harmful to poor children. His justification for this position is that teachers would be resentful—perhaps even psychologically damaged—because they believe something that isn’t true. Which is worse? That he said it, or that nobody cares? Why would someone go on the record and endorse a policy that he acknowledges is bad for poor children?

The key to that question probably lies in the common words that appear in these quotes and others like them—words like “resentful” and “divisive.” Teachers unions know as well as anyone that there are vast differences among their membership in terms of effectiveness. But many—although certainly not all—see any attempt to identify those differences as a potential source of dissension in the ranks, as breeding resentment, as being inherently *dis-unifying*.

In a way it makes sense that anything that seems “divisive” would be so scary. After all, what is “division” if not the opposite of “union”? The history of unionization in America is rife with dishonorable attempts by management to break unions through divide-and-conquer strategies, to dilute union strength by turning members against

one another—often by offering higher wages or benefits to those who refused to join the union.

Given that history, the impulse to value unity above all else is in some ways understandable. But the response has been to create unity through an artificial uniformity, to create a system that pretends that teachers are what they're not—all the same.


Teacher quality information is divisive, in a sense. By definition, it divides teachers among those who are really good, those who are really bad, and those who are in between. So the question isn't whether or not it's divisive. The question is whether we accept or reject the underlying assumption of those who apply this label—that distinguishing among teachers based on effectiveness is necessarily wrong.

The response of some union officials to this issue masks significant differences among their members on exactly the same question. In fact, there are good reasons to think that many rank-and-file teachers actually support the idea of creating value-added measurements of teacher effectiveness. For example, a recent poll of teachers found a virtual dead heat on the issue. When asked about a system of “measuring teacher effectiveness by assessing students’ skills and knowledge when they first come to a teacher and measuring them again when students leave to see what progress was made,” 48% of teachers surveyed rated the idea “Excellent” or “Good,” compared to 18% who rated it “Poor.”⁴⁶ Despite the fact they would presumably rate lower using such a system, 50% of teachers with less than five years experience supported the idea, compared to 44% of those who had been teaching for more than 20 years.

Some local unions have taken positive steps. The Denver Classroom Teachers Association, for example, is an active partner in an ongoing pay-for-performance program that is based on student test scores.

chart 5

How teachers view measuring their effectiveness by their students' growth



Source: Public Agenda, “Stranded By We: What Teachers Really Think About Unions, How They’re Paid and Other Professional Matters,” 2008

And to their credit, union leaders in Tennessee are recognizing and supporting the improvements that have occurred in Chattanooga. Regarding the reform efforts there, which include paying teachers identified as high-performing by value-added data more money to teach in high-poverty schools, an official of the Tennessee Education Association recently said, “When you look at what is best for students, sometimes it does take some radical changes and requires us to experiment and try some options.”⁴⁷

But the fact is that many local and national union officials oppose creating and using data about teacher effectiveness in any meaningful kind of way. This is truly a shame, because in the long run information about teacher success in helping children learn has great potential to be unifying, not divisive. Every piece of value-added data we find that describes the relationship between teachers and student achievement reinforces the central importance of teachers in education. It makes the compelling case that the only way to close the achievement gap and raise performance for all children is to focus our energies and resources on getting low-income and minority children the best possible instruction, to make good teaching the first priority in our public schools.

Many teachers feel under-recognized and under-rewarded. They feel that society doesn't afford them the respect or the compensation they deserve. And they're right. But the reason isn't malice or hostility. It isn't lack of respect, it's lack of information. The inherent complexity of teaching is its downfall in this sense: because it's hard to conceptualize and quantify, because it's hard to measure objectively, people are unwilling to recognize and reward what they can't firmly understand.

It doesn't have to be this way. We now have the ability to find that missing information, in many states for the very first time. And if good teachers can unify around the idea of measuring effectiveness, then society can unify around the idea of good teachers—the need to recognize them, to reward them, and to elevate their professional status to where it really belongs.

An Ambitious Agenda

We've outlined an ambitious agenda for change, a series of reforms that need to be enacted in order to find out who our best teachers are and get more of them to the children who desperately need them. None of our proposals is completely new; many have been discussed and tried before, and are gaining new credibility. The Teaching Commission, for example, recently endorsed many similar reforms.⁴⁸

The idea of finding and using information about teacher effectiveness is so obviously sensible that it never goes away, but so very threatening to the status quo that it's never implemented deeply, broadly and consistently over time. Proponents and opponents alike settle for experimental studies and "demonstration projects." As a result, policies and programs that use teacher effectiveness data to improve student learning remain intermittent, infrequent, and scattershot, dependent on a particular school leader, legislative initiative, non-profit foundation, or special budget line item. They are the exception. They need to become the rule.

We can't rely on whims and pilot projects any more. We need to break through this entrenched stasis and move toward broad, deep, systemic reform. We need to make the creation and use of basic information about teacher effectiveness a standard, central practice in public education nationwide.

And still, there's one more thing. For this to really work, we have to start changing some of the basic values that underpin the way people think about teachers, students, and schools.

It is quite clear that our education system is very much designed to bring the best teachers together with the wealthiest, highest-achieving students. Our schools reflect the norms and values of our time, and teachers respond. Many of those who are really great and really successful aspire to and eventually achieve the goal of engaging our brightest young minds, the future leaders of tomorrow. They teach in brightly-lit suburban classrooms with the latest equipment, surrounded by a small, selective class of smart, motivated students. This, for teachers, is success.

On a very basic level we approve of this as a society. We look at these students and say "These children will lead us someday. They'll be our doctors and scientists, our business leaders and representatives. They're the future. They don't just deserve the very best, we're all better off if they have it."

Of course all of us have great respect for the talented, energetic, highly-effective teachers who go into the "inner city classroom," people who devote their professional lives to children who need the most help, who work in an education

The idea of finding and using information about teacher effectiveness is so obviously sensible that it never goes away, but so very threatening to the status quo that it's never implemented deeply, broadly and consistently over time.

system that doesn't recognize them or pay them what they're worth. But our respect for them is detached and abstract. They're like saints in a stained-glass window or nuns who spend their lives ministering to the diseased and desperately poor. They're objects of our admiration, but not our aspiration. We admire them, but we don't want to *be* them.

There is, after all, a certain comfort in knowing that there are selfless people in this world. Their virtue reflects well on all of us. But in the end we're kind of glad that there aren't too many of our most talented educators who want to do this work, that there remains a plentiful supply of talented, hard-working teachers ready to teach the best and the brightest, and that the parents of the best and the brightest can marshal the resources to pay them.

What can we say? This is a basic ethical failure in our society. All parents want what's best for their kids, that's right and proper and to be expected. But our current education system—both in the practical reality of its relentless sorting of low-achieving, low-income, minority children into the classrooms of our least effective teachers, and in the underlying values and expectations for teachers in general—only responds to the desires of some parents, not to others. All children have the right to good teachers, but only the rights of some are being respected. This problem is so pervasive and so ingrained that we've stopped seeing it as a problem at all, and instead adjusted our values and beliefs to accommodate it. If we're ever going to truly make progress, we have to be honest with ourselves and name injustice for what it is.

The end result of the status quo is a massive lost opportunity for our economy and society as a whole. There are millions of perfectly bright students in low-income households who have vast untapped potential to embody whatever utilitarian ideas of engaged, productive, worthwhile citizens we can dream up. It's not just that they deserve better teachers, though they do. It's not just that justice *demand*s better teachers, though it does. It's that we're shooting ourselves in the foot economically and socially by not giving these children opportunities that they desperately want to take advantage of, if only they were given the chance.

In the end, we have a choice to make.

We can either take action and choose to fix this problem, or through inaction choose not to fix it. We can't fool ourselves into thinking that one is any less of a decision than the other.

The idea of effective teachers helping needy students has tremendous power. It re-affirms the promise of public education and its ability to make all the difference in students' lives. It is a powerful solvent to the inertia and sense of helplessness that have infiltrated the ideas and culture of our public schools. It is a catalyst for radical improvement in almost every facet of education. Good teachers can close the achievement gap, if only we can find them and let them do their work.

The truth about good teachers is out there. All we have to do now is find it, and use it well.

Getting the Most Value Out Of Value-Added Data

Actions for states, higher education, districts, administrators and teachers

Research is clear that teachers have a tremendous impact on student learning. This report has shown how value-added data can tell us how great that effect is and identify which teachers are most effective with which students and in what subjects. When used wisely, this information provides a strong basis for actions that will help states, districts and schools improve teacher quality, raise overall student achievement, and close the achievement gap, as follows:

State Policymakers:

- **Develop and support data systems for the collection and analysis of value-added data.** A value-added data system requires common academic standards, yearly standardized assessments aligned to those standards, and data systems to store and analyze the information at the individual student and teacher level. Many states already have these pieces in place. Those that don't will have them soon because of new provisions of the No Child Left Behind Act. States need to invest in data systems that will pull this data together in a way that will produce valid and fair information about teacher effectiveness. This value-added data should be provided to district and school administrators, to the individual teachers and to the schools of education that prepared them.
- **Make sure state assessments are coherent, rich and accurate measures of learning and academic growth.** The importance of good assessments cannot be overstated. For one thing, good tests drive good instruction. But a strong assessment system is also an essential element of value-added data. Because the information is derived from test scores, the "value" is seriously deflated if the tests lack rigor, are overly reductive or don't show steady progress from year to year.
- **Examine state funding policies to make sure high-poverty and high-minority districts have the means to attract and retain highly effective teachers.** Across the country, schools and districts serving large numbers of poor and minority students are more likely to have underqualified and inexperienced teachers than other better-funded communities. Several states and districts have experimented with signing bonuses, pay-for-performance plans, equitable funding formulas, and better working conditions as ways to make sure our neediest students get teachers that will help them succeed. These programs need to be rigorously evaluated and successful ones need to be greatly expanded. Ensuring an equitable distribution of teacher talent cannot be relegated to pilot programs or demonstration projects. State policy must prioritize this critical issue.
- **Require evidence of student learning as part of teacher preparation and licensure.** There is currently a small, but growing movement to require teacher candidates to show some evidence of producing gains in student learning as a condition for graduation from education school. States should insist that these value-added measures be part of their accreditation process for schools of education and be incorporated into state initial licensing requirements. Likewise, alternate routes into teaching should be evaluated according to the effectiveness of teachers taking that path. In addition, states should include value-added data in their definition of a "highly qualified teacher" for NCLB reporting purposes.

- **Evaluate professional development programs using value-added data.** States allocate millions of dollars each year to the continuing professional development of the teaching workforce. Yet up until recently, few programs could offer data on their ultimate effect on student learning. Value-added information should be a primary criterion in identifying specific professional development needs and measuring their effectiveness. States should further target professional development resources to districts and schools with the greatest need.
- **Review state laws and regulations that affect districts' ability to use value-added measures as part of tenure and teacher compensation policies.** Value-added information is most useful when administrators and teachers can act on it. Whether or not teachers are effective at growing student learning should be part of teacher evaluations, including the decision to award tenure, and should inform teachers' professional development plans. At minimum, this information should be available to administrators for making smart hiring decisions and assigning teachers to schools and classrooms that represent the best match for their individual talents.

Higher Education:

- **Track the effectiveness of graduates from teacher preparation programs and report this data as an indicator of program quality.** Value-added data about teachers should be analyzed and used to improve and strengthen the university programs that prepared them. Institutions with highly effective programs can use this information to attract talented candidates, while weaker programs can identify models to emulate. Such information will also ensure that local schools have better information about where to recruit new teachers.
- **Support more research into what makes teachers effective.** Value-added data identifies which teachers are effective. It doesn't tell us why. While research has revealed some things about the characteristics of effective teachers—they know their subjects well, for example—there remains much more to learn. Universities should encourage more research into this area and use the findings to improve teacher preparation.

District policymakers and administrators:

- **Use value-added data as a central part of recruitment, hiring, compensation and performance evaluations of teachers.** To make sure districts have the faculty with the skills they need, they need to consider teachers' effectiveness alongside other, more traditional measures of teacher quality when making employment decisions. If needed, school boards should negotiate teachers' contracts that allow for such measures to be used.
- **Make sure that teacher quality is distributed equitably throughout the district, and make doubly sure that low-performing students get effective teachers.** Districts should foster a professional environment that confers the most prestige to teachers who take on the most challenging cases. Some districts are doing this by offering incentives, professional recognition, differential pay or working conditions commensurate with the job, for example, by lessening the student load for teachers who work with struggling students. School boards should further make sure that all schools have the funds they need to compete for the best and most experienced teachers in the district.

- **Improve the effectiveness of the current teacher workforce by providing time and resources for ongoing professional development informed by value-added data and teacher effectiveness research.** Districts need to invest in their faculties so that all teachers can become more effective. We all have individual strengths and weaknesses that we bring to our jobs, and teachers are no exceptions. Consequently, even the best teachers typically show more effectiveness in some areas—whether subject areas, particular topics, or different students—than in others. Value-added data is most useful when it can be parsed to show teachers where they are strongest and where they can improve, and give principals specific information for planning professional development for their staffs. School principals should also use this information when making teacher assignments in order to place teachers where they can have the greatest impact.
- **Use value-added data as a component in plans to help ineffective teachers to become better and as part of the process for fairly, but absolutely, removing persistently ineffective teachers from the classroom.** While most teachers will improve given the right support, there are some who will still fail to measure up in the end. We do no one any favors—certainly not the students, but not even the teachers themselves—by allowing these individuals to remain in the classroom.

Teachers:

- **Use value-added data to reflect on your own practices and plan your own professional development.** Even an informal analysis of classroom assessments can yield insights into instruction that works or that needs to be strengthened.
- **For teacher unions: allow teacher effectiveness to outweigh seniority in staffing decisions.** Negotiate for fair pay-for-performance plans based in large part on teacher effectiveness, particularly in shortage areas and underserved schools. Teachers unions that have long agonized about the subjectivity of existing merit pay systems should welcome value-added as a more objective basis for performance pay.

Federal policymakers:

- **Make “value-added” an element of evaluating federally supported activities to improve teacher preparation and professional development.**
- **Support more value-added research.** The federal government can significantly influence the national research agenda by convening researchers and design teams, and by awarding research grants. Not only does much remain to be learned about what makes effective teachers effective, there is more work to be done to advance the design and use of value-added systems. The federal government should also create a clearinghouse for dissemination of best practices in generating and using value-added data.
- **Appropriate funds to help states upgrade their data collection and analysis systems.** Creating these systems will take dollars that pinched state policymakers are hard-pressed to come by, but an infusion of federal funds would jumpstart the process. Federal support could ensure that such data systems meet common standards and use comparable indicators to allow for more robust and reliable research.

The Opportunity Gap

No Matter How You Look At It, Low-Income and Minority Students Get Fewer Good Teachers

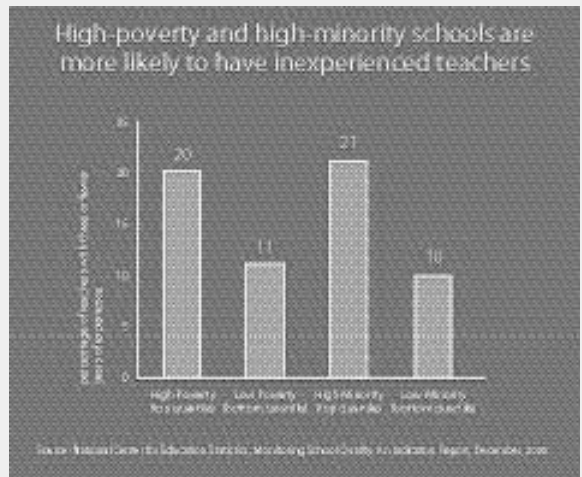
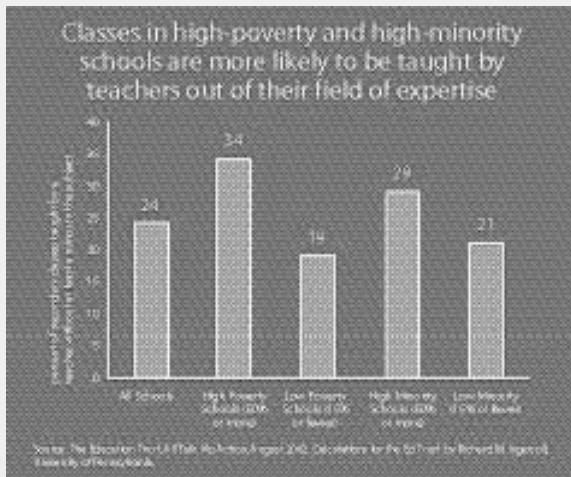
The ability to collect and analyze value-added data is new. Because of this, we have only a limited amount of current information on teacher effectiveness. As a consequence, we have only a limited amount of information with respect to the distribution of effective teachers to different kinds of students. What we do know suggests there's a big problem. In Tennessee, for example, one study using TVAAS data found that African American students were significantly more likely than white students to be assigned to an ineffective teacher.¹

Below we see five different ways of looking at the distribution of teachers with different qualities to different kinds of students. While none of these qualities is a perfect proxy for teacher effectiveness, research shows that those teachers who demonstrate some or all of these qualities in combination are, on average, more effective than those who do not. No matter how you define it, low-income and minority students are systematically assigned to the least qualified teachers.

Knowledge: Teachers need to know the subject they're teaching. A person with a B.A. in math is likely to be a better math teacher than someone with a B.A. in art history.² Yet one out of four high school courses in the core academic subjects is being taught by teachers without a college major, or a even a college minor, in that field. In high-poverty schools the ratio climbs to more than one in three. Students in high-poverty classrooms are

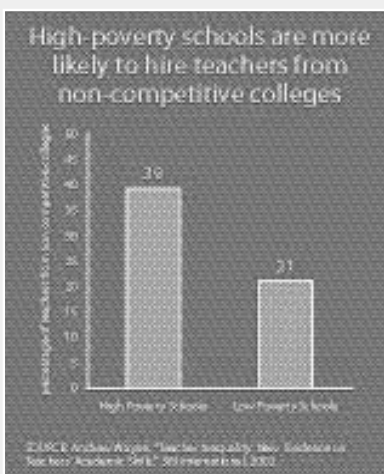
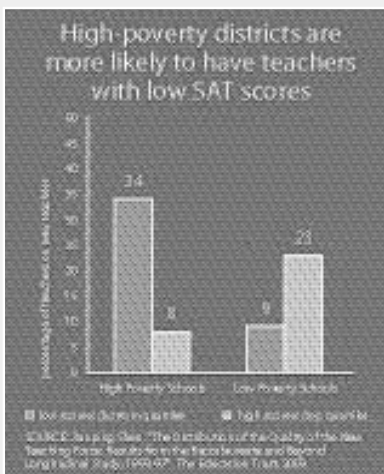
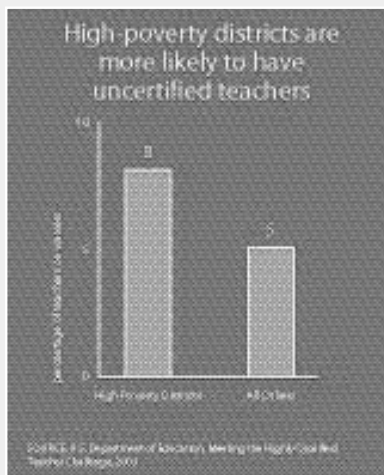
77% more likely than students in low poverty classrooms to be assigned to an out-of-field teacher.³ The pattern repeats itself for minority students: 21% of the courses in low-minority high schools are taught by teachers without a major or minor in field, compared to 29% of the courses in high-minority schools—and a whopping 35% in schools that are more than 90% African-American.

Experience: Studies show that inexperienced teachers are significantly less effective than their peers.⁴ Nationwide, children in high-poverty or high-minority schools are almost twice as likely as other children (20% vs. 11%) to have novice teachers.⁵ State data confirm this. While 23% of New York City teachers had fewer than three years of experience in 1999, for example, only 14% were similarly inexperienced in the neighboring Lower Hudson and Long Island districts.⁶ In Texas, low-income elementary school students are 20% more likely to be taught by teachers with one year of experience or less.⁷ California is much worse. Students in high-poverty, high-minority schools are almost *twice as likely* as students attending predominantly white schools to be taught by a teacher in the first year or two of teaching.⁸ We're using the classrooms of our least advantaged students as a training ground where new teachers can make mistakes, learn from them, and then take that knowledge elsewhere.



Certification: While every state has a teacher certification process, some schools still end up hiring teachers who aren't fully certified, and the proportion of teachers on "waivers" from certification is 61% higher in high-poverty school districts nationwide than in all other districts.⁹ In California 18% of the teachers in high-poverty schools are uncertified, compared with 10% of the teachers in other schools. Over 28% of African American students in high-poverty California schools have an uncertified teacher.¹⁰ Of 114,638 teachers in high-poverty school districts in New York, 13,357, or 12%, are uncertified. Of the 103,875 teachers in the rest of the state, 143, or 0.1%, are uncertified, meaning that 99% of all uncertified teachers in New York are teaching in high-poverty school districts.¹¹ In Maryland uncertified teachers comprise 20% of teachers in high-poverty schools and 12% in other schools; New Mexico shows a disparity of 16% to 7%. Most state certification standards for teachers are not particularly high, and yet we let people teach who don't even meet those standards, and they disproportionately end up teaching low-income and minority children.

Test Performance: Those who teach low-income and minority students are less likely to score well on standardized tests, such as teacher licensing tests, assessments of basic skills, and college admissions exams.¹² Illinois children in high-poverty schools are five times as likely to be taught by teachers who failed the state teacher licensure exam



at least once, and twenty-three times as likely to be taught by teachers who failed it at least five times.¹³ Of those teaching minority students in New York, 21% failed one of the state's certification exams, compared to 7% of those who teach white students. A study found that 34% of new teachers in high-poverty schools were in the bottom quartile of performance on the SAT, while only 8% were from the top quartile. By contrast, only 9% of those in low-poverty schools were from the bottom SAT quartile, compared to 23% from the top quartile.¹⁴ We consistently assign teachers whose performance lags on important tests to low-income and minority students, whose performance in turn lags on important tests.

Quality of Undergraduate Education: National data indicates that 21% of the teachers in the lowest-poverty schools attended "non-competitive" colleges, compared to 39% of those in high-poverty schools.¹⁵ Similarly, state-specific data from New York suggest big differences along these lines for minority students. Ranking all schools by the institutions their teachers attended, the top ten percent of schools in New York employ almost no teachers who attended "least-competitive" institutions, while the bottom ten percent average more than 30%. Minority students in New York are more than twice as likely as white students to be taught by teachers from the least competitive institutions.¹⁶

How states, districts and schools are using value-added data

Tennessee has the most comprehensive value-added system in the country. It's the only state so far to put in place a statewide process of gathering the information it needs about the effect of teaching on students' academic growth. Other states and districts, however, are starting to catch on. Because Texas has annual testing in all the elementary and middle school grades, the Dallas school system has been able to calculate value-added measures of teacher quality for a number of years. Other states with school districts using value-added analysis of schools and/or teachers include Colorado, Minnesota, North Carolina, and Wisconsin.

What's amazing when you talk to the educators in these schools and districts is the many different ways they've been able to use information about teacher effectiveness to help improve student achievement. Here are just a few examples:

Chattanooga, Tennessee

Chattanooga, Tennessee, is a leader in analyzing and using value-added data to make the best possible matches between effective teachers, subjects and students. With the help of the community and local businesses, the district offers substantial incentives to highly effective teachers to teach in low-performing, high-poverty elementary schools where eligible teachers earn as much as \$5,000 per year more. Other perks include housing benefits and free graduate education.

The net result of these policies and other efforts in Chattanooga to improve funding, management, leadership, and community support has been a significant increase in student achievement. Each of the nine schools targeted by this initiative was previously among the 20 lowest-scoring in the state on 3rd grade reading scores. All nine schools have shown

statistically significant gains in all five tested subject areas in each of the last two years, growing faster than the other schools in the district. During that time, the proportion of students reading at grade level in grade three increased from 22.6% to 36.9%.

Dallas Independent Schools

The Dallas Independent School District used value-added data to study highly effective teachers in order to learn what makes them so successful. They found that the most effective teachers taught more higher-order skills and challenging subject matter than other teachers. These teachers were knowledgeable in their subjects, de-emphasized "drill and practice" approaches and consistently taught the entire curriculum to their classes.

Community Advocates in Tennessee

The non-profit Chattanooga Public Education Foundation (PEF) in Tennessee used TVAAS scores to conduct a study of 92 teachers who were particularly effective in helping students learn. They found that high-performing teachers—those who were in the top 25% on TVAAS measures and also nominated as high-performing by their principals—shared certain distinctive qualities and practices.¹ These teachers combined consistently high expectations and standards for themselves and their students with a high level of flexibility and student engagement. This research is now being used for professional development activities and gives principals a more concrete basis for choosing future teaching candidates. PEF is also working with higher education institutions to use TVAAS data to improve the effectiveness of new teacher candidates by analyzing the relationships between curriculum and coursetaking in higher education with subsequent success in the classroom.

Milken's Teacher Advancement Program

Through the Teacher Advancement Program (TAP) sponsored by the Milken Family Foundation, a group of schools in Arizona recently began a program of school improvement centered on reforms including market-based teacher compensation, multiple career paths, improved professional development, and teacher accountability for their students' value-added achievement. The results have been encouraging. All participating schools made achievement gains in both years they implemented the reform, out-gaining demographically similar comparison schools by a total of 13% over two years. The schools that most rigorously implemented the program had student gains that were 51% greater than in comparison schools.² Similar TAP programs are currently underway in Arkansas, Colorado, Indiana, Louisiana, and South Carolina.

Minneapolis Public Schools

Minneapolis Public Schools recently used value-added measures of teacher effectiveness to identify a group of highly effective elementary schoolteachers, those who produced achievement gains that were significantly above expected levels for three consecutive cohorts of 2nd grade students. The district now plans to follow up with research and classroom observations of the highly effective teachers, learning from their success and using that information to increase teacher effectiveness districtwide.

¹ Chattanooga Public Education Foundation, 2002.

² John Schacter et al, *The Impact of the Teacher Advancement Program on Student Achievement, Teacher Attitudes, and Job Satisfaction*, Milken Family Foundation, 2002.

³ Interview, June 25, 2002.

Guilford County North Carolina

A district administrator in Guilford County, North Carolina, told us how value-added data helps them with professional development and classroom assignment: "Until adopting the value-added system, we had no process for identifying teacher effectiveness in terms of whether their students were making a full year's growth. This has been a major void in the process used to identify and correct weaknesses of individual teachers or groups of teachers... We can tell that some teachers are producing growth in math while others are getting better results in reading, enabling principals to make scheduling decisions and classroom assignments based on hard data."³

Ohio Partnership for Accountability

The Ohio Partnership for Accountability is a unique collaboration of the state department of education, board of regents, and all 51 schools of education. Early this year, the partnership announced their plan to use value-added teacher effectiveness data to better understand, study and improve the preparation of teachers in Ohio colleges and universities. Over the next five years, Ohio researchers will study the math and reading scores of the students of both new and veteran teachers as a means to evaluate the quality of teacher preparation in the state and to identify the practices and policies that have the most effect.

ENDNOTES

¹ The U.S. Department of Education, FY 2004 Education Budget Summary and Background Information, Appendix 1, shows that federal, state, and local expenditures for elementary and secondary education for the 2002-2003 school year totaled to \$429 billion. By comparison, the Office of Management and Budget, FY 2004 Budget, Historical Tables, shows that total outlays for national defense in FY 2003 were \$376 billion.

² The most prominent and enduring example of this phenomenon can be found in the 1966 federally-sponsored analysis usually referred to as the “Coleman Report.” In analyzing available data, the report was much more successful in finding correlations between academic success and the qualities of students, such as race, income, socio-economic status, etc., than it was in finding links between academic success and the qualities of schools. What we now know is that this was essentially a function of the fact that specific, meaningful data about student characteristics was, and is, relatively easy to come by, while data of similar quality about schools is not. Some things are easier to measure than others, but the degree of measurability and degree of importance are no way related. Thus, the researchers couldn’t find school effects in their analysis not because it wasn’t there, but because they were looking for it in the statistical equivalent of a dark room using weak flashlights. The room has gotten brighter and the flashlights stronger since then, but despite this some commentators continue, unaccountably, to rely on the Coleman Report as an important source of information for contemporary education commentary (See for example, George Will, *Shame: School Reform’s Weak Weapon*, The Washington Post, March 2, 2003).

³ Kati Haycock, *Good Teaching Matters*, The Education Trust, 1998.

⁴ W.L. Sanders and J.C. Rivers, *Cumulative and Residual Effects of Teachers on Future Student Academic Achievement*, University of Tennessee Value-Added Research and Assessment Center, 1996.

⁵ W.L. Sanders, quoted by David Hill, “He’s Got Your Number,” *Teacher Magazine*, May 2000.

⁶ Steven G. Rivkin, Eric A. Hanushek, and John F. Kain, *Teachers, Schools and Academic Achievement*, University of Texas-Dallas Texas Schools Project, 2002.

⁷ Samuel Bratton, Sandra Horn, S. Paul Wright, *Using and Interpreting Tennessee’s Value-Added Assessment System: A Primer for Teachers and Principals*, 1996. For a detailed explanation of the statistical methods used to create TVAAS measures, see Sanders, Saxton, Horn, “The Tennessee Value Added Assessment System: A Quantitative Outcomes-Based Approach to Educational Assessment,” *Grading Teachers, Grading Schools: Is Student Achievement a Valued Education Measure?*, Jason Millman ed., 1997.

⁸ The TVAAS system adjusts its rating for the amount of student achievement information available for each teacher. If insufficient data is available to provide a strong rating—as with newly-hired teachers, for example—the system gives

them the benefit of the doubt and assumes that their performance is equal to the system average.

⁹ An independent evaluation of TVAAS performed by prominent statisticians found that “Both the analysis of TVAAS estimates of teacher gain effects and our own analysis of a large sample of gain scores in teacher classrooms showed that, although these estimates were...variable from year to year, the results were stable enough to permit identification of teachers with notably meritorious or problematical instructional effectiveness, as measured by test-score gain.” See R. Darrell Bock and Richard Wolfe, *A Review and Analysis of the Tennessee Value-Added Assessment System*, Tennessee Comptroller of the Treasury, 1996.

¹⁰ Not her real name.

¹¹ A study of teachers in New York found a high degree of correlation between various teacher qualities in schools. Schools with more inexperienced teachers, for example, also tended to have more teachers who were uncertified, attended a less competitive university, etc. See Lankford, Loeb, and Wyckoff, 2002.

¹² Sitha Babu and Robert Mendro, *Teacher Accountability: HLM-Based Teacher Effectiveness Indices in the Investigation of Teacher Effects on Student Achievement in a State Assessment Program*, presented at the American Education Research Association annual meeting, April 2003.

¹³ In this analysis, “effective teachers” are defined as those who rank in the top two quintiles of all teachers in terms of effectiveness; “ineffective teachers” are those who rank in the bottom two quintiles. Low-, medium-, and high-performing students are identified by sorting all students into equal thirds based on prior performance.

¹⁴ Babu, Mendro, 2003.

¹⁵ Kevin Carey, *The Funding Gap*, The Education Trust, 2003.

¹⁶ Ibid.

¹⁷ School funding lawsuits have been filed in 44 states over the last 35 years, in many cases forcing major changes to state funding schemes that unconstitutionally discriminate against low-income students. For a recent example of a successful lawsuit, see Campaign for Fiscal Equity v. New York, www.cfequity.org.

¹⁸ Approximately 95% of school districts use the single salary schedule, see Allan Odden and Carolyn Kelley, *Paying Teachers for What They Know and Do: New and Smarter Compensation Strategies to Improve Schools*, 1997.

¹⁹ Despite the fact that numerous studies indicate that the link between increased experience and increased effectiveness drops off substantially after at most 5—10 years of experience, teachers continue to get paid more for every year of service—above and beyond any inflationary increases—all the way until retirement. Many studies of teacher effectiveness find no relationship at all between effectiveness and attaining a Master’s Degree (see, for example, Rivkin, Hanushek, and Kain, 2002).

²⁰ This phenomenon has been studied with respect to other high-demand teachers, such as those who teach in subject areas like science, math, and special education. One study of teacher salary data found that instructors who teach in these hard-to-fill subject areas in private schools, which generally aren't subject to rigid salary schedules and price controls, are paid more than their private-school peers. No such salary differences existed between similar sets of teachers in public schools. (Dale Ballou and Michael Podgursky, "Let the Market Decide," *Education Next*, 2001.)

²¹ Andrea Holland Larue, *The Changing Face of Teacher Tenure*, The University of Austin, 1996. See also various state tenure statutes.

²² Ruth Mitchell and Patte Barth, *Not Good Enough: A Content Analysis of Teacher Licensing Examinations*, The Education Trust, 1999.

²³ For a thoughtful discussion of the potential benefits of making the teacher certification process significantly less burdensome, see Frederick M. Hess, *Tear Down This Wall: The Case for a Radical Overhaul of Teacher Certification*, Progressive Policy Institute 21st Century Schools Project, November 2001.

²⁴ Eric Hanushek, John Kain, and Steven Rivkin, *Why Public Schools Lose Teachers*, November 2002.

²⁵ Catherine Freeman, Benjamin Scafidi, and David L. Sjoquist, *Racial Segregation in Georgia Public Schools 1994—2001: Trends, Causes, and Impact on Teacher Quality*, Andrew Young School of Policy Studies, Georgia State University, 2002.

²⁶ Hamp Lankford, Jim Wyckoff, Frank Papa, "The Labor Market for Public School Teachers: A Descriptive Analysis of New York State's Teacher Workforce," University at Albany, 2000.

²⁷ *The New York Times*, September 1, 2000.

²⁸ This encompassed all schools and school districts in New York other than New York City. When New York City was included, within-district differences accounted for 35% of the variation in teacher quality, reflecting the substantial differences in teacher quality between New York City and the rest of the state. See Lankford, Loeb, and Wyckoff, 2002.

²⁹ Common Core of Data, National Center for Education Statistics.

³⁰ Marguerite Roza and Paul Hill, *How Within-District Spending Inequities Help Some Schools to Fail*, Center on Reinventing Public Education, University of Washington, May 2003.

³¹ *Making Choices: Diversity, Student Assignment and Quality in Wake's Public Schools*, Wake Education Partnership, 2003.

³² Press release, Charlotte-Mecklenburg Schools, March 12, 2003.

³³ "CMS Teachers Told To Toe the Line in Class," *Rhinotimes.com*, April 3, 2003.

³⁴ Carey, 2003.

³⁵ Kevin Carey, *State Poverty-Based Education Funding: A Survey of Current Programs and Options for Improvement*, Center on Budget and Policy Priorities, November 2002.

³⁶ Andrew Wayne and Peter Youngs, *Teacher Characteristics and Student Achievement Gains: A Review*, Review of Educational Research, 2003.

³⁷ U.S. Department of Education, National Center for Education Statistics, *1993 Baccalaureate and Beyond Longitudinal Study*.

³⁸ Sandra Huang, Yun Yi, and Kati Haycock, *Interpret With Caution: The First State Title II Reports on the Quality of Teacher Preparation*, The Education Trust, 2002.

³⁹ <http://www.carnegie.org/sub/program/teachers.html>

⁴⁰ Scott Stephens, "Ohio Looks for the key to perfect teachers," *Cleveland Plain Dealer*, January 11, 2004.

⁴¹ Texas A&M University System, Regents Initiative for Excellence in Education, 2002 Annual Report.

⁴² Email correspondence, October 21, 2003.

⁴³ Indiana Code 20-6.1-9-3

⁴⁴ Pamela Riley et al., *Contract for Failure: The Impact of Teacher Union Contracts on the Quality of California Schools*, Pacific Research Institute for Public Policy, 2002.

⁴⁵ One study of value-added measures of teacher effectiveness grouped teachers into quintiles based on effectiveness ratings and studied how those ratings changed over time. Among teachers who taught the same grade and subject for three consecutive years, it found that the majority of teachers saw their effectiveness rankings change by 0 or 1 quintiles, while close to 90% changed by 2 or fewer quintiles. The teachers whose ratings were least likely to change were those ranked in the bottom (least effective) quintile and the top quintile. (Dash Weerasinghe and Mark Anderson, *Validation Studies and Post-Hoc Analyses of Classroom Effectiveness Indices*, Dallas Independent School District, 2001.)

⁴⁶ *Stand By Me: What Teachers Really Think about Unions, Merit Pay and Other Professional Matters*, Public Agenda, 2003.

⁴⁷ Jay Mathews, "A Move to Invest More in Effective Teaching," *The Washington Post*, February 10, 2004.

⁴⁸ The Teaching Commission, "Teaching At Risk: A Call to Action", 2004.

Notes to “The Opportunity Gap” page 36

- ¹ In a study of third grade teacher assignments in two large metropolitan school systems, the study found that while black students made up 38% of all students, they comprised 51% of students assigned to teachers in the bottom quintile of effectiveness, but only 29% of those assigned to teachers in the top quintile of effectiveness – Sanders and Rivers, 1996.
- ² D.D. Goldhaber and D.J. Brewer, “Evaluating the Effect of Teacher Degree Level on Educational Performance,” *Developments in School Finance 1996*, NCES, 1997
- ³ Craig Jerald and Richard Ingersoll. *All Talk, No Action: Putting an End to Out-of-Field Teaching*. Education Trust, 2002.
- ⁴ Rivkin, Hanushek, and Kain, 2002.
- ⁵ National Center for Education Statistics, “Monitoring Quality: An Indicators Report,” December 2000.
- ⁶ Hamilton Lankford, James Wyckoff, and Papa. *The Labor Market for Public School Teachers: A Descriptive Analysis of New York State’s Teacher Workforce*; 2000.
- ⁷ Rivkin, Hanushek and Kain, 2002.
- ⁸ Christopher Jepsen and Steven Rivkin, *Class Size Reduction, Teacher Quality, and Academic Achievement in California Public Elementary Schools*, Public Policy Institute of California, 2002.
- ⁹ *Meeting the Highly Qualified Teacher Challenge: The Secretary’s Second Annual Report on Teacher Quality*, U.S. Department of Education, Office of Policy Planning and Innovation, 2003.
- ¹⁰ Jepsen and Rivkin, 2002.
- ¹¹ U.S. DOE, 2003.
- ¹² Kain and Singleton, “Equality of Educational Opportunity Revisited,” *New England Economic Review*, May/June 1996.
- ¹³ Rossi et al, *Chicago Sun-Times*, September, 2001.
- ¹⁴ Jianping Shen, “The Distribution of the Quality of the New Teaching Force: Results from the Baccalaureate and Beyond Longitudinal Study 1993-97” The Education Trust, Unpublished Study, 2003, p.8.
- ¹⁵ Wayne, Andrew, “Teacher Inequality: New Evidence on Teachers’ Academic Skills.” SRI International, 2002.
- ¹⁶ Hamilton Lankford, Susanna Loeb, James Wyckoff, “Teacher Sorting and the Plight of Urban Schools: A Descriptive Analysis,” *Educational Evaluation and Policy Analysis*, Spring 2002.

Thinking K-16 is published by

The Education Trust, Inc.

1725 K Street, N.W., Suite 200 • Washington D.C., 20006

Phone: 202-293-1217 • Fax: 202-293-2605

www.edtrust.org.

Editor: Patte Barth

Editorial Assistant: Christa Jackson

Graphics Designer: Autumn Richardson

The Education Trust was created to promote high academic achievement for all students at all levels, kindergarten through college. While we know that all institutions could better serve their students, our work focuses on the schools and colleges most often left behind in efforts to improve education: those serving Latino, African American, Native American and low-income students.

Thinking K-16 is published with the intent to share lessons learned in these communities with policy-makers as well as with educators and members of the public concerned with the quality of education provided our neediest young people.