

Carnegie Knowledge Network • What We Know Series on
Value-Added Methods and Applications

Webinar 15: What Do We Know About the Long-term Impacts of Teacher Value-Added?

Q&A with Stephen Raudenbush

April 9, 2014

Q: The studies you mention seem to be emphasizing long-term outcomes, such as earnings, over short-term ones, such as test scores. This may be difficult for schools because while they may care about the long-term outcomes they are not held accountable for most of them. How can school districts most productively balance the long-term benefits that you mentioned with their short-term accountabilities?

A: I think the rationale behind these studies is that we want to know whether it matters if a teacher boosts test scores. If so, then we first want to know if we are getting a valid estimate of the teacher's impact on test scores. Many studies focus on this, but if we assume that teachers are truly boosting test scores, does that really matter? Who cares? What are these tests measuring? I think that what these studies are showing is that yes, it does matter. These studies are saying that if you can boost test scores now, it will have an effect. It is something that kids can carry with them. It is giving evidence that teachers can have long-term effects and that value-added scores do carry information about that. So we should pay attention to value-added scores because to some extent it will stick with the kids. I do not think that anyone is suggesting that we should hold teachers accountable for whether their students ultimately get jobs. That does not make any sense. Rather, it is an exercise in indicating that boosting test scores is something that we should pay attention to.

Q: Do any of the studies you mentioned speak to value-added being used for teachers of students with disabilities? Do we have any data about the persistence of value-added for these types of teachers?

A: I am not sure if we have any information about the long-term effects for that particular group, which is a large, important group. That is something to think about looking at. We do have a brief by Dan McCaffrey on students with disabilities and value-added. I think that this is very informative about how reliable and valid value-added can be for students in special education. But I do not think that we are currently going to find any long-term impact data.

Q: You mentioned that value-added matters because it predicts some of these long-term and non-cognitive outcomes. Would it be more efficient to measure the non-cognitive factors directly? What are your insights on the emerging research on non-cognitive factors?

A: My own feeling is that we should make a big point of this. Certainly at the secondary level and in schools that serve more disadvantaged students who are at high-risk for dropping out, we should see whether effective teachers help students stay in school, come to class, pass courses, and accumulate credit. All of those things are very important. If you are in ninth grade and you fail a class or two, your chances of graduating are going to be very low. Are your teachers enabling you or encouraging you to do the work of schooling? A lot of that work is not what is measured on tests. We know that kids who get good grades have a number of other qualities. There is a very interesting book by James Heckman

Carnegie Knowledge Network • What We Know Series on Value-Added Methods and Applications

Webinar 15: What Do We Know About the Long-term Impacts of Teacher Value-Added?

and Tim Kautz that recently came out on what they call “character.” Clearly at the secondary level measuring some of these other outcomes, for example, whether students are getting in trouble, being suspended, and attending class, are all important. At the elementary level, it may be a little more difficult to capture these things. It may be that school attendance is actually more under the control of the parents, so I am not sure that we actually see the teacher effects with those kids. But we could be collecting portfolios of students. Ron Ferguson has his Tripod Survey of kids’ perceptions and their engagements, which is not strictly academic. He has gone all the way down to at least grade three and has validated that students as young as grade three can tell us very useful information about their engagement and how successful their teacher is. I think that we should really push very hard to diversify the range of outcomes that we look at for schools, teachers, and all educators.

Q: It seems that it would be wise for us to begin constructing a richer set of effectiveness indicators, but then what? What should schools and districts do with this richer set of metrics that include some of these non-cognitive measures?

A: That is a huge question. I have a personal view on this, which is that it makes sense for the districts to get a lot of good information. But, in some sense, I think that to really help teachers become successful and to make schools function well requires that the school itself is an organization functions well. There is a lot of evidence about this. I am talking about principal leadership. Additionally, I am talking about some of the new work on what is being called distributed leaders, which is the engagement of teachers, master teachers or advanced teachers, as instructional leaders. My view is that this information that we are currently talking about can supplement all of the information that school leaders are able to get about how well students are doing and classrooms are functioning in order to get a much richer sense of not only how well the whole school is doing, but also how particular teachers are doing and what might be the needed areas for guidance or help. In my view, I imagine a two-level accountability system. Districts hold particular schools accountable for doing well by their students. Then all the teachers in a school have a stake in helping each other help the school to do well. The school would become an information-rich environment where there are a lot of opportunities for people to get guidance. For the small number of people who get guidance and learn who are still not doing well, they may learn that this may not be the right profession for them. We see this in the writing about how this works in effective schools. People have the opportunity to learn about how they are performing and where they could improve.

Q: To this point about the importance of the schools in student outcomes. Given the segregation in New York schools, the highest in the United States according to the person who asked the question, and if we cannot separate the teacher effect from the school effect in the in Chetty et al. study, why would we take his suggestions on the impact of the teacher on long-term outcomes seriously?

A: There is some evidence in the Chetty et al. study that individual teachers do really make a difference. They use a brilliant strategy to look at this. They looked at both high value-added and low value-added teachers and watched what happened when the teachers moved. When high value-added teachers left a school, the school value-added actually went down. There is another study by Kirabo Jackson that

Carnegie Knowledge Network • What We Know Series on Value-Added Methods and Applications

Webinar 15: What Do We Know About the Long-term Impacts of Teacher Value-Added?

shows when your peers have high value-added then you have higher value-added. I think that it is certainly hard to separate teachers from schools. What the Chetty et al. study shows is that if I am in an instructionally effective environment, it can have long-term effects on my outcomes. The study is not attributing how much of that came from the school as an organization versus how much came from the individual. I really resonate with this concern about segregation. I am in Chicago so maybe we can debate about who is more segregated. Chicago is certainly a very segregated city. New York maybe more segregated because of its high-end neighborhoods. In my other brief, I argue that comparing teachers who work in very, very different kinds of schools serving very different kinds of kids just does not make sense. All of the literature that we have showing that teacher value-added can be estimated validly and fairly is based on studies that compare teachers who are in the same school. To me that has a very important implication. We do not want to pit teachers from the same school against each other and have them compete for raises. That simply does not make sense. What we want to do is create a situation where teachers are motivated and are collaborating and helping each other. The visibility of differences in effectiveness helps people who are more expert determine who to help and in which particular domain. I think that there is a role for teacher value-added in schools, but it is mostly in determining how the school can give guidance.

Q: We have a question on teacher effects versus peer effects. Much of this analysis refers to a teacher's contribution to student achievement, but the long-term data at the classroom level shows larger effects. Might this be due in part to peer composition? How might we tease this out and what might it mean for district policy?

A: That is a wonderful question. It is one of the main things that I discuss in my second brief. In comparing teachers from different schools I discuss this very problem. It is very difficult with value-added data alone to separate them. I would say that it is nearly impossible. The Tennessee Class Size study and a couple of other experimental studies help us look at this because teachers within the same school are being randomly assigned and kids are being randomly assigned to teacher. This means that we should only see chance differences between classrooms in peer compositions, so they are likely to be small. There is much better evidence in that case that something is specifically happening that is not related to a student's peers. The other thing is that in the MET study we see the same teacher being randomly assigned and then we see the same teacher having different students. We see that experimental study value-added is predicting how much students learn the next year within the same school. The MET study is another experimental study that gives us some ability to separate out what you might think are the unique teacher effect. I think that there are very clever and excellent studies that enable us to see teacher effects. But, we are not able to separate teacher and school effects within conventionally commuted value-added, which is one of the motivations behind my second brief on comparing teachers from different schools.

Q: You mentioned that teachers likely vary in the extent to which the effect from their initial value-added persists. Can you say more about potential sources or causes of this variation in persistence?

Carnegie Knowledge Network • What We Know Series on Value-Added Methods and Applications

Webinar 15: What Do We Know About the Long-term Impacts of Teacher Value-Added?

A: I know that Susanna Loeb has done some very interesting work on this and one of her Carnegie briefs is related to the variability in the persistence of teachers. It is true that some teachers are more persistent than others. I think that I, and maybe any of us, could only speculate on why one person may be more persistently good and stable in how effective they are. Who knows what is going on in the lives of these people or whether it is an interaction between the environment they are in and their own characteristics. These could make people more stable in their effectiveness, while others are bouncing around. I would look hard at Susanna's work. I think that she has some of the best work on it and has taken a very careful look at it.

Q: Someone noticed the variation in the long-term effects of value-added. It looks like the biggest long-term effect is on earnings, which is the most distal in terms of years from schools. Why do you think that the effects on earnings are larger than on college attendance, college quality, or teen pregnancy?

A: I could only speculate about that. I do not think that the studies themselves give us insight on this. These studies are a bit like a black box. You have students coming in and then later on they just pop up with the results. There is not much in the studies themselves. One thing that I would say though is whether you attend college or not may be very dependent on your academic record. It certainly would also be based on other things, like how much information you had about college, but your academic record plays a central role. I think an intriguing hypothesis is the idea that schools are like the work place. Showing up, getting the job done, getting good grades, and persisting are analogous to showing up to a work place, doing work, getting paid, and being promoted. This is something that sociologists have been talking about for sixty or seventy years: what is the function of the school in the development of the child? It is not just learning subject matter. It is about learning how to be in society and, in particular, learning how to be in the work place that is the classroom, which has many similarities to the work place that supports the economy. So there is reason to believe that kids who do well, are given good instruction, learn how to be good students, and are rewarded for that may in fact become better employees because they are more productive on the job and may be paid more.

Q: Have any studies looked at the cumulative effect of a sequence of high value-added teachers for individual students and whether fade out is still an issue in that case?

A: There is some work on cumulative effects, but I think it has more to do with where we think a child's achievement would be with a series of good teachers. Bill Sanders has done some work on this in Tennessee. I think that you can find his work in our brief. But, I do not think that we have any evidence of the cumulative effects of value-added on long-term outcomes. I was wondering why we did not get that out of the Chetty et al. study. They had great long-term data, but it seems that we did not get more on the cumulative effects. If you think about what a school or district should be, it should put together a sequence of experiences that cumulatively help shape a good trajectory. That is really what it is all about. It is not just about randomly having good teachers who are going to affect you in one particular year. I think that looking at teachers in isolation as individuals, not together, is a limitation of our work.

Carnegie Knowledge Network • What We Know Series on Value-Added Methods and Applications

Webinar 15: What Do We Know About the Long-term Impacts of Teacher Value-Added?

Q: Clearly long-term studies are slow and expensive, what can districts and schools do now in light of the research that you brought to us today?

A: That is a huge question. What schools and districts should do cannot be derived from any discussion we have on value-added because the task of thinking how to create a powerful, coherent educational system for kids, how to make school function well, how to help schools support teachers, how to ensure students can be safe and develop well cannot be addressed simply by value-added. That task is much bigger. The question might be how does value-added help us get to that vision of a school. I think that it gives us a piece of useful information about average learning data for kids. It is not terribly precise, but it is worth paying attention to and using in combination with other information, such as classroom observations (conducted multiple times a year, ideally by more than one person), looking at student perception, and examining some of the other kinds of data that I was mentioning. I think the real story is about how schools get to be better at being creators and users of information. I have not mentioned formative assessments, but I think that there is an exciting number of schools using them. It seems that they are being used more frequently to see if kids are making progress and devising instructional strategies in response. Then teachers can check later assessments to see if the earlier instructional strategies are paying off. It is very exciting. If those formative assessments are linked to a good summative assessment, which can be used for a value-added score, we could use the summative assessments to hold the school accountable and the formative assessments to help teachers devise good instructional plans. You might end up with a system that could capitalize on the new ability to access massive amounts of information. That is really the goal. Whether we can figure out how to do that intelligently is probably one of the things that determines whether we can dramatically improve the school system.