Student Outcomes based on Perceptions of Teacher Quality

Jack Murdoch

_____

A Thesis Submitted to The Honors College

In Partial Fulfillment of the Bachelors degree
With Honors in

Business Economics

THE UNIVERSITY OF ARIZONA

May 2023

Approved by:

Dr. Martin Gray Hunter
Department of Economics

Abstract

In this paper, we examine the effects of student perception of teachers on long term student outcomes. Using data from the National Longitudinal Survey of Youth with nearly 9,000 observations, we construct three multiple linear regression models with dependent variables: income, natural log of income, and highest grade completed regressed against student ratings of their teachers along with several relevant demographic variables. We find that students who rated their teachers highly were more likely to have a higher income later in life. Having a teacher perceived as bad had no statistically significant effect on long term income. Students who rate their teachers highly complete more years of education on average and those who rate their teachers poorly complete less years of education on average. These findings lead us to conclude that the quality of a school's teachers matters in the long run for a student's success both academically and financially.

Table of Contents

Introduction

The past few decades have seen an abysmal plateau in the skills of U.S. Students as measured by standardized tests. There is an alarming lack of progress according to the National Assessment of Educational Progress (NAEP) within the U.S., but if we compare the test scores of American students on the Programme for International Student Assessment (PISA) test to those of students in other developed countries, the results are distressing (Hanushek, Woessman, Peterson, & Summers, 2013). American students in the class of 2011 who were deemed proficient in math according to NAEP, came in 32nd on the PISA test compared to other developed countries (Hanushek, Woessman, Peterson, & Summers, 2013).

The purpose of tests is to measure a student's skill level in important areas, especially math and reading. American test scores point towards a severe deficiency in these skills compared to other economically advanced nations. American students are not developing the skills in reading and math that students in other economically developed countries are. Eric Hanushek concludes in "Endangering Prosperity that mathematical skills especially are important to a country's economic growth and that the American education system's inability to provide these mathematical skills will harm our GDP and overall economic competitiveness in the long run (Hanushek, Woessman, Peterson, & Summers, 2013).

If we accept the assumption that mathematical ability is accurately tested by standardized tests and that mathematical skills taught in the classroom help propel GDP, as Hanushek et al. propose, it becomes increasingly apparent that the current educational system is failing our students and our country (Hanushek, Woessman, Peterson, & Summers, 2013). There is a real problem and real consequences of a failure to address it.

In this paper, we examine the effects of student perception of teachers on long term student outcomes. Using data from the National Longitudinal Survey of Youth with nearly 9,000 observations, we construct three multiple linear regression models with dependent variables: income, natural log of income, and highest grade completed regressed against student ratings of their teachers along with several relevant demographic variables. We find that students who rated their teachers highly were more likely to have a higher income later in life. Having a teacher perceived as bad had no statistically significant effect on long term income. Students who rate their teachers highly complete more years of education on average and those who rate their teachers poorly complete less years of education on average. These findings lead us to conclude that the quality of a school's teachers matters in the long run for a student's success both academically and financially.

Motivation and Literature Review

Background

The field of study known as the economics of education seeks to answer empirical questions about the role of education in preparing students for the labor market. The Mincer

earnings function which posits that an individual's wage is a function of years of schooling and work experience is one of the most basic theorems in labor economics (Mincer, 1974). The basic assumption behind this function and others like it is that education builds human capital, which is valued by employers who are willing to pay for it. However, one weakness with this model is that it measures the effect of education in terms of years spent in school. This does not account for the fact that there is enormous variation between schools in terms of quality. Some schools are very good at instilling skills in students while others are less so. Measuring education merely by years equates a diploma from a high-quality school to that from a terrible school as this does not allow for the existence of quality differences. This may be acceptable if we assume that education is merely a signaling mechanism to employers, but if we actually believe that education builds human capital, then the variation between schools renders simple years in the classroom an inadequate measure of school quality.

The last few decades have seen an abundance of quality empirical research examining the various inputs that make a school good at educating its students. Dozens of economists, sociologists, and political scientists have studied what makes a good education and what leads to long term student success. The multitude of factors such as peer effects, teacher quality, class size, and family structure that affect the academic success of students makes the conversation surrounding how to create good educational systems complicated. However, the wealth of literature surrounding all of these topics has brought us, if not directly to the chalice of truth, at least to the tabernacle within which it lies. After decades of research there are still many questions left unanswered, but we have at least uncovered important trends in the literature.

The question that all researchers in this field are trying to answer is this: what are the inputs that make a good school. A good school is not simply a function of easily measurable inputs, but is instead an institution that improves the educational attainment of the students that attend it. A student who attends a good school should enter the workforce with more skills and thus earn more consistent with the Mincer earnings function (Mincer, 1974). Schools have a huge role in the life of a student and thus proper evaluation of individual schools is key to education policy.

Current literature in the economics of education indicates that human capital is one of the most important factors determining a successful educational institution, and the decisions administrators make regarding their staff have huge consequences for students in classrooms. Researchers have examined a variety of methods for evaluating teacher quality ranging from having a masters degree, to the actual effect a teacher has on test scores.

School Inputs in the Education Production Function

One of the most important concepts in the field of Economics of Education is the idea of an education production function in which student outcomes such as test scores, graduation rates, or years of schooling are a function of school, community, and family inputs. The idea of an education production function is usually attributed to sociologist James Coleman who produced the famous “Coleman Report” in 1966. This study, mandated by the Civil Rights Act of 1964 to examine the differences in educational opportunities between black and white students, created a model of student outcome in a particular grade as a function of peer effects, family inputs,

community inputs, teacher inputs, and school inputs (Coleman, 1968). Coleman found that the most important inputs were unrelated to schools and were instead related to family economic and educational background. He argued that factors like economic resources, the educational background, and parental interest in student achievement were far more important than school resources like funding, class size, and teachers. These school-related inputs were found to have little effect on student output (Coleman, 1968). However, he did find that black students benefited more from better school resources like better teachers than did white students where teacher quality is measured by a teacher's level of education and score on a verbal skills test (Coleman, 1968).

The effect of the Coleman report was enormous in the fields of sociology and economics at the time. However, many were skeptical of the conclusion that schools had little ability to influence the life of a student. Among these skeptics was Eric Hanushek, who claimed that Coleman had used flawed measures of school inputs, particularly teacher quality. Hanushek pointed out that measuring teacher quality is very difficult and that factors used by Coleman like teacher level of education and test score are not necessarily correlated with the actual effectiveness of a teacher. Because of the faulty measures in Coleman's education production function, teacher effects on student outcomes were biased.

One major problem in Coleman's study was that it was based upon the idea that measurable inputs of school resources, like school funding, teacher level of education, school size, curriculum, and lab facilities, were the appropriate inputs for predicting student achievement (Coleman, 1968). Coleman discovered that these inputs made relatively little difference in the educational gains in the children (Coleman, 1968). This led some to conclude that there was no way that education could make a meaningful difference in the life of a child and that the only way to improve academic achievement for disadvantaged students was by addressing the pre-existing conditions like poverty and family structure.

However, Coleman's analysis is overly simplistic in assuming that it can use simple quantitative measures of how good or bad a school is. Schools and their components are highly complex and there is no easy way to accurately quantify effective versus ineffective schools simply on the basis of easily definable inputs. This is especially true of teacher quality. Coleman used inputs like a teacher's level of education and score on a verbal skills test as a proxy for teacher quality (Coleman, 1968). While easily measurable, these inputs may be imperfect estimators of actual teacher quality. Teacher quality for the purposes of this paper is a teacher's human capital or the ability of a teacher to impart useful knowledge to students.

The human capital possessed by teachers is not as simple as measuring the teacher's educational attainment or intelligence. We simply do not know yet what sort of factors make an individual good at instilling useful knowledge into young minds. The individual determinants of a skillful educator may ultimately be a question for psychologists, but for our purely economic analysis the relevant points established by the literature are that teaching ability is a skill that some individuals possess in a greater degree than others, and that a teacher's ability to impart knowledge makes a huge difference to a student's achievement both in school and in the labor

market. While it may be intuitive to think of a teacher's level of education as being a good measure of skill in the classrooms, the wealth of literature produced so far on this topic gives us reason to doubt this. A better and more direct measure for understanding the effectiveness of a given teacher is to actually observe how much knowledge the student has accumulated in the classroom. Researchers can do this by measuring the difference between standardized test scores before and after having a particular teacher. Using student achievement negates the need for arbitrary proxies of teacher skill and goes straight to measuring student outcomes. When this score is added up and averaged over several years of teaching, researchers construct what is called a value-added model of teacher effectiveness. The value-added model of teacher effectiveness is generally preferred by researchers today over the measures used by Coleman, as it more directly assesses teacher quality.

Teacher Education

There has been a wealth of literature examining the relationship between student achievement gains on standardized tests and teacher level of education. In 1976 David Armor et al. studied reading scores among disadvantaged students in Los Angeles Unified School District (Armor, et al., 1976). These researchers collected data on teachers' educational history including characteristics: college attended, undergraduate major, graduate school experience, and college education relating to teaching skills. Armor et al. found that these characteristics had no effect on student achievement in reading (Armor, et al., 1976). This result had strong implications for how schools evaluate teacher skill and it seemed to contradict Coleman's assumption that teacher quality is related to teacher educational background. As the literature surrounding this topic grew, more and more studies found little correlation between observable teacher characteristics and teacher quality. In 1981, a study by Murnane and Phillips in midwestern schools found no effect of teacher educational background on student achievement in their study of schools (Murnane & Phillips, What do Effective Teachers of Inner City Children Have in Common?, 1981).

As the field of economics of education grew, the number of studies examining student outcomes in relation to teacher characteristics increased. One often cited characteristic for teacher quality is a teacher's level of education. Coleman used this as an input in his education production function (Coleman, 1968). However, it is not clear that the mere attainment of a degree makes one a better teacher, and there could be many teachers who are very skilled at instilling knowledge but lack additional years of education. It may be the case that years of education is a biased measure of teacher quality. When examining the literature overall, our skepticism about the strength of teacher education as a measure of quality seems warranted. Eric Hanushek conducted a meta-analysis of over 400 such studies and found that in only a small minority of these studies was there a statistically significant effect on student outcomes of teacher education level, and an equally small number of studies found statistically significant negative effect of teacher level of education (Hanushek & Rivkin, Handbook of Economics of Education. Volume 2, Chapter 18, 2006). The majority of these studies yield statistically

insignificant estimates of the effect of a teacher's level of education on student outcomes (Hanushek & Rivkin, Handbook of Economics of Education. Volume 2, Chapter 18, 2006).

Teacher Compensation

Teacher salary is another input in the education production function that receives much attention in the debate around education policy and school funding. Teachers are paid little considering their level of education and their salary compared to other workers with a similar level of education has dropped significantly since the 1940s (Hanushek & Rivkin, Handbook of Economics of Education. Volume 2, Chapter 18, 2006). This has led many to speculate that low salaries have driven talented potential teachers away from the marketplace into other more profitable endeavors. From an economic perspective, this makes a lot of sense as the opportunity cost of a career in education is too much for many educated people. However, Michel Podgursky points out that there are many non-pecuniary "fringe benefits" to the teaching profession which make it particularly appealing to certain types of workers (Podgursky, Fringe Benefits: There is More to Compensation than a Teacher's Salary, 2003). Teachers have their summers off, and they work less hours per week than other similarly educated professionals (Podgursky, Fringe Benefits: There is More to Compensation than a Teacher's Salary, 2003). It is also the case that teaching allows one to share a schedule with one's own children, which is a very attractive aspect of this line of work for those with families, especially women. This, combined with a very attractive pension plan and high levels of job security indicate that there are many non-pecuniary benefits to the teaching profession (Podgursky, Fringe Benefits: There is More to Compensation than a Teacher's Salary, 2003).

Given that salary is sometimes considered a measure of an employee's value to the firm, it is often used as a proxy for teacher quality in the education production function. However, this ignores the fact that public education wages are very different from those of the private sector. Teacher wages in traditional public schools are often determined by a flat salary schedule meaning that salary increases are awarded based on time with the school rather than on any measure of teacher quality. It may be the case that it attracts certain types of individuals who value a good work life balance very highly and place a great value on time spent with family. Given this, it seems likely that teacher salary is an inaccurate estimator of a teacher's skill, as there are non-pecuniary benefits to the profession, and the salary schedule is based of seniority rather than on value added. Eric Hanushek confirms that salary is unlikely to be a useful proxy for teacher skill with another meta-analyses of the relevant research on this topic which finds that only a minority of studies find a statistically significant effect on student outcome of teacher salary (Hanushek & Rivkin, Handbook of Economics of Education. Volume 2, Chapter 18, 2006). Like the other easily measurable inputs, salary shows little promise as an accurate measure for teacher quality on the education production function.

Teacher Experience

Experience in the teaching profession is another commonly cited input for determining teacher quality. It makes sense that teachers who have been teaching longer might be better than those who have spent less overall time in the classroom. However, Rivkin, Hanushek, and Kane found that the effect of experience on teacher quality is nonlinear (Hanushek, Rivkin, & Kain, Teachers Schools and Academic Achievement, 2005). Teachers who are in the first 3 years of their career tend to be significantly worse at educating than more experienced workers (Hanushek, Rivkin, & Kain, Teachers Schools and Academic Achievement, 2005). However, after the first two to three years, additional years of education make little difference in student outcomes (Hanushek, Rivkin, & Kain, Teachers Schools and Academic Achievement, 2005). The largest gains in skill for a teacher seem to come in their first three years of teaching which indicates that total years of experience is an imperfect measure of teacher quality (Hanushek, Rivkin, & Kain, Teachers Schools and Academic Achievement, 2005). In another meta-analysis conducted by Eric Hanushek only 41 percent of high-quality studies reveal a statistically significant impact of teacher experience on student outcomes (Hanushek & Rivkin, Handbook of Economics of Education. Volume 2, Chapter 18, 2006). This is likely due to a nonlinear effect of teacher experience with large amounts of skill developed early in one's career and a somewhat flat effect of additional years of experience after that. The early concentration of experience effects are significant in considering the problems of high teacher turnover which is much more prevalent among new teachers (Ingersoll & Smith, 2003). However, it seems that overall years of experience is a flawed measure of teacher quality, and a teacher with 30 years in the classroom may be little better at educating than one with 5. Despite the ease of using experience to measure quality, we should be cautious considering that the marginal effects of another year of teaching are not constant.

The fact that these measurable inputs have little effect on student achievement may lead us to conclude, like Coleman, that ultimately teachers have little ability to influence student achievement, and that family and community inputs are the only ones which matter (Coleman, 1968). However, the fact that easily measurable inputs like years of education do not have a statistically significant coefficient in the education production function doesn't necessarily mean that teacher quality is irrelevant to student outcomes. It may be the case that we simply lack a set of obvious indicators of good vs bad educators, but that these are nonetheless still valid categories. Upon examining the literature, it seems that teachers are an extremely important variable in the education production function, but what matters is teacher skill rather than any arbitrary proxy for teacher quality like education or experience. A teacher's ability to give skills to a student is not determined by education, salary, experience, or any other of the easy to measure inputs used by Coleman.

Shortcomings and the Value-Added Model

Instead of these, we may consider a value-added measure of teacher quality in which teacher quality is measured by the difference a teacher makes in the knowledge of a student. A

value-added measure of teacher quality considers student scores on standardized tests before entering a teacher's class, and then the test scores after having been in the teacher's class. This value-added measure allows researchers to determine how much students learn from a teacher conditional on the other inputs like family, peer effects, and other relevant inputs. The advantage of the value-added model is that it shows the effect a teacher has on an individual class, or on several classes. Finding out which teachers consistently add to the knowledge of their students based on tests of mathematical skills helps to identify high quality teachers without having to resort to easily measurable, but inaccurate inputs like education level.

When researchers use the value-added model to measure teacher impacts on student achievement, they get much more significant results than those yielded by directly measurable proxies for quality. In 1992, Hanushek studied schools in Gary Indiana and found a dramatic difference in teacher effect on students as measured by a value-added model (Hanushek E. A., 1992). He classified teachers who consistently generated large gains in test scores as good and those who generated little or no improvement in test scores as bad. The results of this study showed that a good teacher can yield the equivalent of an extra year of learning when compared to a bad teacher (Hanushek E. A., 1992). Easily measurable inputs of teacher quality like education level had little predictive power over how much learning a teacher contributed to students (Hanushek E. A., 1992). Importantly, this study showed that these teacher differences existed within schools which further strengthens the argument for value added models. The fact that these teacher differences exist within schools serving the same population instead of teacher differences across different populations means that it is highly likely that teacher skill, rather than variance in population, is responsible for gains in student achievement.

One of the largest and most prominent studies of teacher value-added was conducted by Raj Chetty of Harvard in 2014. Chetty used data from school district records and federal income tax records to examine the outcomes, both academic and financial, of 2.5 million students based on teacher value-added measures. This study is divided into two papers which each examine a different aspect of the value-added model. The first paper examined the question: to what extent to value added measures of teacher quality indicate actual differences in teaching ability, and to what extent is this difference simply due to sorting among student populations with higher achieving students selecting into classes with better teachers (Chetty, Friedman, & Rockoff, Measuring the Impact of Teachers I: Evaluating Bias in Teacher Value-Added Estimates., 2014). The second paper addresses the question of whether teachers who have a high value added score are actually improving the long term prospects of their students, or if they are merely teaching to the test without really making a meaningful difference in student ability (Chetty, Friedman, & Rockoff, Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood., 2014).

The first paper lays out Chetty's value added measure which assigns teachers scores based on mean improvement in their students test scores over the course of several years (Chetty, Friedman, & Rockoff, Measuring the Impact of Teachers I: Evaluating Bias in Teacher Value-Added Estimates., 2014). This is the standard method by which value added is generally

measured. However, due to the fact that teacher quality actually changes over time, Chetty accounts for "drift" in teacher quality by weighing more recent classes at a higher value than classes that are further in the past. This method is more accurate than standard value-added models because it allows for change in teacher quality over time. Another reason that Chetty's study is innovative is that it uses teacher turnover and switches between school districts to demonstrate the effectiveness of the value-added model. This quasi-experimental method assumes that teacher turnover within the school is uncorrelated with characteristics of the student population. He justifies this assumption by pointing out that teacher turnover within schools is not correlated with any observable parent characteristics and that when teachers with high VA scores move across schools, scores change dramatically for the very year that they move. Chetty looks at the effect on students when teachers with high value-added scores transfer to a different school, and he finds that classes tend to do worse immediately after a high value-added teacher leaves, while students in the classroom that the teacher transfers to do much better on standardized tests. This study strengthens the literature around value-added measures and their ability to accurately predict which teachers raise student test scores.

In his second paper, Chetty goes on to argue that these VA models predict more than the ability to teach to the test, but instead, they actually predict long term effects, both academic and otherwise, of good teachers on students (Chetty, Friedman, & Rockoff, Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood., 2014). Using the same quasi-experimental model that he used in the first paper, Chetty finds that a one standard deviation increase in the value-added score of a teacher results in a 0.86% higher probability that a student will be in college at age 20, and raises the quality of colleges attended as well (Chetty, Friedman, & Rockoff, Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood., 2014). The effects on college attendance by themselves should indicate that good teachers are dramatically improving the educational outcomes of students, but there are also several impressive non-academic outcomes he examines as well. An individual's earnings increase by 1.34% given a teacher with a one standard deviation higher value-added score. An increase in one standard deviation of a teacher's value-added score decreases the probability that a female student will have a teenage pregnancy by 4.6%, which is a powerful result given the negative effects on earnings that teenage pregnancy causes. It is clear that the benefits that a high value-added teacher has on students extend far beyond higher standardized test scores. The results of this paper indicate that good teachers help students develop human capital which helps propel them to more success than they would have enjoyed under a worse teacher.

The effectiveness of the value-added model in predicting teacher quality has been supported in many different econometric analyses. However, as part of the Measures of Effective Teaching project, economists Kane, McCaffery, Miller, and Staiger have conducted experiments involving random assignment of teachers to students thus reducing the probability that unobservable student characteristics are influencing the effects of VA measures (Kane, McCaffery, Miller, & Staiger, 2013). Kane et al. used the value-added score of a teacher combined with student surveys, and a measure of teaching ability according to Charlotte

Danielson's framework for teaching as measured by four videos of classroom instruction by the teacher. These extensive measures were used to predict student outcomes, and these predictions were then measured against the actual achievement gains in the students randomly assigned to the teachers. These researchers found a clear, linear relationship between the predicted quality of the teacher according to their measures and the academic growth of the students.

Kane et al. used several different combinations of test score gains, student surveys, and classroom observations to find which model would best predict student achievement gains (Kane, McCaffery, Miller, & Staiger, 2013). In model one, gains on test scores were weighted the most, followed by student evaluations, and then classroom observations. Other models included one in which all three predictors were measured equally, and one in which classroom observations were weighted at 50% with test score gains, and student surveys were weighted at 25% each. They found that the best predictor of student achievement was model one in which value-added was weighted the most, followed by student surveys, with classroom observations carrying very little weight. This model also seems to make the most sense from a policy perspective considering the cost of the predictors used. The value-added model costs almost nothing as students take state tests already. In addition, student surveys are cheap and easy measures of teacher quality. Classroom observations, however, cost much more as it requires workers with expertise in rating teachers according to Danielson's framework for effective teaching to watch hours of classroom footage.

Kane et. al. make an important contribution to the literature around teacher quality in two important ways. Firstly, they use experimental data to show that the value-added model is actually measuring teacher skill, rather than simply measuring which teachers attract high achieving students (Kane, McCaffery, Miller, & Staiger, 2013). Secondly, this study introduces measures of teacher quality like student surveys and classroom observations that can be combined with value-added in order to create a more well-rounded predictive model. It is significant that student surveys proved to be useful in this model as this suggests that student perceptions of instructors may be somewhat useful in predicting teacher quality.

High Performing Schools Emphasizing Teacher Quality

Increasingly, researchers are finding particular school models which routinely generate large increases in student outcomes. The charter school model has been especially studied by many including Fryer and Dobbie who found in 2013 paper that students who attend high performing charter schools not only have higher test scores than their demographic peers in traditional public schools, but they also have many long-term non-academic improvements (Fryer & Dobbie, Medium-Term Impacts of High-Achieving Charter Schools on Non-Test Score Outcomes, 2013). The two found that those who attended high performing/no-excuses charter schools were far less likely to be incarcerated, less likely to become pregnant as a teenager, and more likely to attend college. These longer-term outcomes solidify the case that the no-excuses school model, which includes an emphasis on maximizing teacher's human capital and in turn has been found to positively impact student outcomes. This contributes to the long list of

rigorous empirical work which demonstrates the importance that education plays in the lives of students, and the fact that human capital is an integral part of the education production function. A good education does more than simply raise test scores, but it reduces risky behavior and improves long term outcomes for students.[1]

A broad view of the literature establishes that teachers have an enormous impact on the education production function and that we are coming closer to being able to measure teacher ability. The value-added model shows great promise, but alternative measures such as student perceptions of teacher quality may also be useful. In 1975 Richard Murnane asked principals to rate which teachers were most effective and which were least effective. These principles did not have access to the value-added data which had been collected, but their subjective assessments of good vs bad teachers were shown to be good predictors of which teachers raised the test scores of students (Murnane, The Impact of School Resources on the Learning of Inner City Children, 1975). A similar study of inner-city students by David Armor et al. in 1976 revealed the same ability of principles to predict teacher quality without any knowledge of the value-added model (Armor, et al., 1976). Both of these studies align with other findings that teachers have the ability to influence student achievement, but they also make the novel contribution that principle have good information about teacher quality independently of test score outcomes. This is important because it indicates that we have more than one way of assessing teacher ability. It is also significant because it substantiates the notion that subjective evaluations of teacher quality have some validity.

Data Overview and Measure of School Quality

Overview of the NLSY97

The National Longitudinal Survey of Youth 1997 (NSLY97) is an ongoing longitudinal study conducted by the Bureau of Labor Statistics of 8,984 individuals born between 1980 and 1984. The participants in this study were randomly selected and interviewed first in 1997 with follow up interviews on an annual basis until 2011 after which interviews were conducted biennially. The participants were between the ages of 12-18 when they were first interviewed in 1997. One of the remarkable aspects of the NLSY97 is its retention rate with 77.3% of the original sample responding to the latest survey in 2019. This, combined with the large sample size makes this dataset very useful for evaluating trends in the labor market and effects of education. In 1997, both the participants and their parents were asked about a wide range of topics about education, lifestyle, after-school programs, and work. Surveys given at later dates were more limited, only distributed to the participants, and asked questions about secondary education, lifestyle, employment, and income.

26% of the NLSY97 sample is black with a mean income of $45,025 and a mean highest grade completed of 13.58 as of 2019 when participants would have been between 34 and 40 years old. Hispanics account for 21% of the sample with a mean income of $52,791 and mean highest grade completed of 13.43. Non-black and non-hispanic respondents accounted for 52% of the sample with a mean income and highest grade completed of $64,530 and 14.67 respectively; these individuals were classified as “white” for the purposes of this study. Mixed

race respondents accounted for only 1% of the population with a mean income of $62,200 and a highest grade completed of 14.35. The sample is 52% male and 48% female. Mean male income is $66,503 as opposed to mean female income of $47,263. Mean highest grades completed for the male vs the female population were 13.75 and 14.46 respectively.

Table 1 – Participant Demographics

Demographics	n	%	Mean Income	Mean Highest grade completed
Male	4599	52%	$66503.34	13.75
Female	4385	48%	$47262.61	14.46
Non-Black/Non-Hispanic	4665	52%	$64530.30	14.67
Black	2335	26%	$45025.48	13.58
Hispanic	1901	21%	$52791.17	13.43
Mixed Race	83	1%	$62200.14	14.35

**Participants were between the ages of 34 and 40 when reporting their income and highest grade completed in 2019.*

Teacher "Quality" or Perception Measures

In the 1997 survey, students were asked several questions regarding their attitudes towards school and their teachers. Importantly for this research, the survey included statements about students' perceptions of their school. These statements in the survey were preceded by the question: "*Thinking about your (last) school in general, how much do you agree with each of the following statements about your school and teachers?*" Among the six statements that followed, three addressed the student's perception of the skill of their teachers. These were:

1. The teachers are good.
2. The teachers are interested in the students.
3. Students are graded fairly.

The students were asked to respond on a Likert scale that included the following options: agree, strongly agree, disagree, strongly disagree. These were the only questions in the survey given to students that directly pertained to teacher quality.

To explore the effects of subjective student perceptions of their teachers, we used the results of these questions to divide the ratings that students gave their instructors into four categories: excellent, good, neutral, and bad. We categorized a student as having rated their teachers "excellent" if they selected "strongly agree" in response to all three statements about teacher quality outlined above. We categorized a student as having rated their teachers "good" if they answered agree, or strongly agree to all three statements. Students were considered to have

"bad" teachers if they answered "disagree" or "strongly disagree" to all three statements. Finally, if a respondent didn't fall into any of the other categories into any of these three categories, they were classified as having "neutral" teachers. There were only 34 observations for which school attitudes were unavailable due to the respondent's refusal to answer. Considering the small proportion of non-response answers, we excluded these observations from our analysis.

Table 2 - Teacher Summary

	n	%	Mean income of students	Highest grade completed of students
Excellent Teachers	615	7%	$60,568.69	14.84
Good or Excellent Teachers	6,128	69%	$59,562.74	14.40
Neutral Teachers	2,512	28%	$51,931.84	13.54
Bad Teachers	310	3%	$47,057.66	12.86

**Note, data was not available due to non-responses for 34 students. They are thus excluded from this analysis.*

These measures of teacher quality at the school level are the closest we could come to being in line with previous research on factors that make a "good" teacher. The measures we've created reflect primarily the student's assessment of their teacher's "quality". In Armor's work, he established that teachers who were effective at increasing student achievement were teachers who "felt efficacious" maintained "orderly classrooms," engaged in high levels of contact with parents, and consulted with other teachers on reading programs (Armor, et al., 1976). Later work on the No-Excuses Charter school model rendered similar results about the importance of school cultures which emphasized high expectations and frequent communication with parents (Fryer & Dobbie, Getting Beneath the Veil of Effective Schools: Evidence from New York City, 2013) (Angrist, Pathak, & Walters, 2011). We concluded that the level of student belief that their teachers are good, interested in students, and fair graders, track very well with the current research on the high expectations culture. We assumed that teachers who have high expectations for their students are likely interested in student success and are probably more effective than those who have little interest in student success. It is also highly likely the case in our opinion that teachers who grade students fairly are more likely to be better at maintaining an orderly classroom which is another important aspect of the No-Excuses charter school model (Angrist, Pathak, & Walters, 2011) (Fryer & Dobbie, Getting Beneath the Veil of Effective Schools: Evidence from New York City, 2013). Given this and taking into consideration Thomas Kane's 2013 study establishing the efficacy of student ratings in predicting teacher quality, we decided to use these student ratings as a proxy for the overall quality of their teachers, but future research

is needed to examine if student perceptions of teacher quality correlate with traditional measures of teacher quality.

Statistical Analysis with Empirical Models

In our study, we regressed long term respondent outcomes on how they rated their teachers in 1997 along with controls for sex, race, parental education, and location. Assuming that students are accurate in assessing their teacher's skill level, this model should show the effects of teachers on students. The dependent variables that we examined were the highest grade ever completed, income, and the natural log of income of the respondent in 2019. We regressed these models against dummy variables for excellent, good, and bad teachers, which are the variables which interest us the most. In each model we included demographic variables of race, sex, highest grade completed of the respondent's father, and a dummy variable for whether the respondent lived in a rural or urban area.

Our study found that student perceptions of their teachers in 1997 had a statistically significant effect on long term outcomes both in terms of income and in terms of educational achievement. These results are in accordance with the conclusions of researchers like Chetty who claim that skilled teachers are able to influence the long-term outcomes of their students (Chetty, Friedman, & Rockoff, Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood., 2014). More importantly, however, they indicate that students' assessments of teacher quality matter. This confirms Kane's findings with the MET project that student surveys can be an accurate way to rate teachers. The findings also support past literature findings that teacher quality impacts student success.

Model 1 – Predicting Income

Equation 1 - Regression Model Predicting Income

$$Income = \beta_0 + \beta_1 Excellent\ Teacher + \beta_2 Good\ Teacher + \beta_3 Bad\ Teacher + \beta_4 Sex + \beta_5 Black + \beta_6 White + \beta_7 Hispanic + \beta_8 Father's\ Highest\ Grade + \beta_9 Urban + \varepsilon$$

Table 3 – Abbreviated Regression Output Predicting Income

	Dependent Variable
	Income
Excellent Teacher	173.51
	-2,709.51
Good Teacher	4248***
	-1607.19
Bad Teacher	-2522.75
	-4,215.89
Controls for Sex	Y
Controls for Race	Y
Controls for Parent Education	Y
Controls for Location	Y
Observations	5,396
R-Squared	0.073

*Note: *p<0.1, **p<0.05, ***p<0.01*
**Click here to see the full results*

From the first model that utilized income as the dependent variable, we found that having a teacher rated as "good or excellent" according to student perceptions was statistically significant at the 1% level. The estimate coefficient indicates that those who reported having "good or excellent" teachers had an income $4,248 higher than those who reported having neither excellent, good, or bad teachers. The coefficient of having a bad teacher was -$2522, and the coefficient of having an excellent teacher was $173, but neither were not found to be statistically significant. The specified model can be found in equation 1 above along with the simplified results presented in Table 3 with the full results included in the appendix.

Model 2 – Predicting ln(Income)

Equation 2 - Regression Model Predicting ln(Income)

$$ln(Income) = \alpha_0 + \alpha_1 Excellent\ Teacher + \alpha_2 Good\ Teacher + \alpha_3 Bad\ Teacher + \alpha_4 Sex + \alpha_5 Black + \alpha_6 White + \alpha_7 Hispanic + \alpha_8 Father's\ Highest\ Grade + \alpha_9 Urban + \varepsilon$$

Table 4 – Abbreviated Regression Output Predicting ln(Income)

	Dependent Variable
	ln(Income)
Excellent Teacher	0.026
	(0.048)
Good or Excellent Teacher	0.083***
	(0.028)
Bad Teacher	-0.07
	(0.075)
Controls for Sex	Y
Controls for Race	Y
Controls for Parent Education	Y
Controls for Location	Y
Observations	5,360
R-Squared	0.077

*Note: *p<0.1, **p<0.05, ***p<0.01*

**Click here to see the full results*

We also estimated a log level regression analysis on income motivated by the original Mincer equation (Mincer, 1974). From this model, we found that the effect of having a "good or excellent" teacher was an 8.3% increase in future income compared to those who reported having neither excellent, good, or bad teachers. Similarly, the coefficients of an excellent teacher and a bad teacher were not found to be statistically significant. The specified model can be found in equation 2 above along with the simplified results presented in Table 3 with the full results included in the appendix.

Model 3 – Predicting Highest Grade Completed

Equation 3 - Regression Model Predicting Highest Grade Completed (HGC)

$$HGC = \theta_0 + \theta_1 Excellent\ Teacher + \theta_2 Good\ Teacher + \theta_3 Bad\ Teacher + \theta_4 Sex + \theta_5 Black + \theta_6 White + \theta_7 Hispanic + \theta_8 Father's\ Highest\ Grade + \theta_9 Urban + \varepsilon$$

Table 5 – Abbreviated Regression Output Predicting Highest Grade Completed (HGC)

	Dependent Variable
	Highest Grade Completed
Excellent Teacher	0.474**
	(0.212)
Good Teacher	0.666***
	(0.121)
Bad Teacher	-0.630**
	(0.296)
Controls for Sex	Y
Controls for Race	Y
Controls for Parent Education	Y
Controls for Location	Y
Observations	6,850
R-Squared	0.065

*Note: *p<0.1, **p<0.05, ***p<0.01*

**Click here to see the full results*

In our model using highest grade completed as the dependent variable, we found that having a "good or excellent" teacher had a statistically significant positive coefficient of 0.67 on a scale of 0 to 20. Having an excellent teacher and having a bad teacher had statistically significant coefficients of 0.47 and -0.63 respectively. The specified model can be found in equation 3 above along with the simplified results presented in Table 5 with the full results included in the appendix.

Study Limitations

While the results outlined above take a unique approach, leverages publicly available longitudinal data on nearly 900 individuals, and findings are consistent with past literature, there are several areas in which it could be improved. One major problem is that we did not have access to any value-added measures. Value-added is the gold standard for measuring teacher quality, the NLSY97 provided no way of measuring this. The only direct measures of teacher quality in the data were the subjective student ratings of their teachers. In a dataset with value-

added measures, the correlation between student ratings and measurable teacher quality would be far clearer and more reliable. This raises some significant concerns about possible endogeneity in our model. Students who like their teachers may share several unobservable characteristics which make them more likely to succeed academically and financially. Similarly, students who are low achievers may be more likely to rate their teachers poorly.

Our study also lacks variables pertaining to the school in which the respondents attended both initially and later. While such data exists, we were unable to retrieve it from the Bureau of Labor statistics as it requires prior approval. School level data is very important and our lack of access to this data raises more concerns about endogeneity. Students who attend schools which have a higher concentration of quality teachers may share several relevant but unobservable characteristics. Studies done by Hanushek have measured teacher effects within schools which allowed him to control for and examine the impact of certain teacher and school specific characteristics than we were unable to more effectively (Hanushek E. A., 1992).

Another problem with our analysis is that it only includes student ratings from the year 1997. Unfortunately, the NLSY97 survey only asked questions concerning school attitudes in that year. A dataset with multiple years of student ratings would strengthen our case for the reliability of student ratings. Both Chetty and Hanushek have performed such analyses with value-added data, and it allowed them to measure the effects of multiple years under a good teacher on students (Hanushek E. A., 1992) (Chetty, Friedman, & Rockoff, Measuring the Impact of Teachers I: Evaluating Bias in Teacher Value-Added Estimates., 2014).

Future Research

The NLSY97 includes surveys sent to the participants' schools in 1997. These surveys cover many different school characteristics, and questions about staff and student body composition. This data would allow us to directly examine objective school characteristics rather than having to use subjective student attitudes about their school/teachers, however due to confidentiality reasons, it is highly restricted. Future research should obtain this sensitive data and examine the relationship between value-added measures of teacher quality and subjective student ratings to help us assess with more confidence how good students are at rating their teachers and potentially identify a more accurate relationship between teacher quality and downstream student success.

Unlike Chetty and Hanushek, our study looked only at aggregate teacher quality across the school. The survey asked students to respond to questions about their "School in general" which doesn't allow us to draw information about an individual teacher, but rather about the quality of teachers across the entire school. A future study could improve upon our results if they are able to examine student perceptions of individual teachers within the school.

Conclusion

This study utilized the nearly 9000 participants in National Longitudinal Survey of Youth 1997 (NSLY97) to examine the relationship between student perceptions of their teachers while between the ages of 12-18 and their relationship with future income and educational attainment. It found that, even controlling for a number of relevant demographic variables, students who think highly of their teachers were more likely to have higher incomes and higher educational attainment when reported between the ages of 34 to 40. This is an important finding because it

points to student ratings of their teacher's quality as a significant predictor of future success, further suggesting that teachers play an important role in the lives of students both in the short and long-run. Current literature in the economics of education indicates that teachers have a profound effect on students that is typically measured through value added models. Our study builds on the work of Kane in indicating that student surveys may be an important tool in teacher evaluation.

Appendix

Table 5 – Participant Demographics

Demographics	n	%	Mean Income	Mean Highest grade completed
Male	4599	52%	$66503.34	13.75
Female	4385	48%	$47262.61	14.46
Non-Black/Non-Hispanic	4665	52%	$64530.30	14.67
Black	2335	26%	$45025.48	13.58
Hispanic	1901	21%	$52791.17	13.43
Mixed Race	83	1%	$62200.14	14.35

**Participants were between the ages of 34 and 40 when reporting their income and highest grade completed in 2019.*

Table 6 - Teacher Summary

	n	%	Mean income of students	Highest grade completed of students
Excellent Teachers	615	7%	$60,568.69	14.84
Good or Excellent Teachers	6,128	69%	$59,562.74	14.40
Neutral Teachers	2,512	28%	$51,931.84	13.54
Bad Teachers	310	3%	$47,057.66	12.86

**Note, data was not available due to non-responses for 34 students. They are thus excluded from this analysis.*

Figure 1 - Mean Student Income Bar Graph

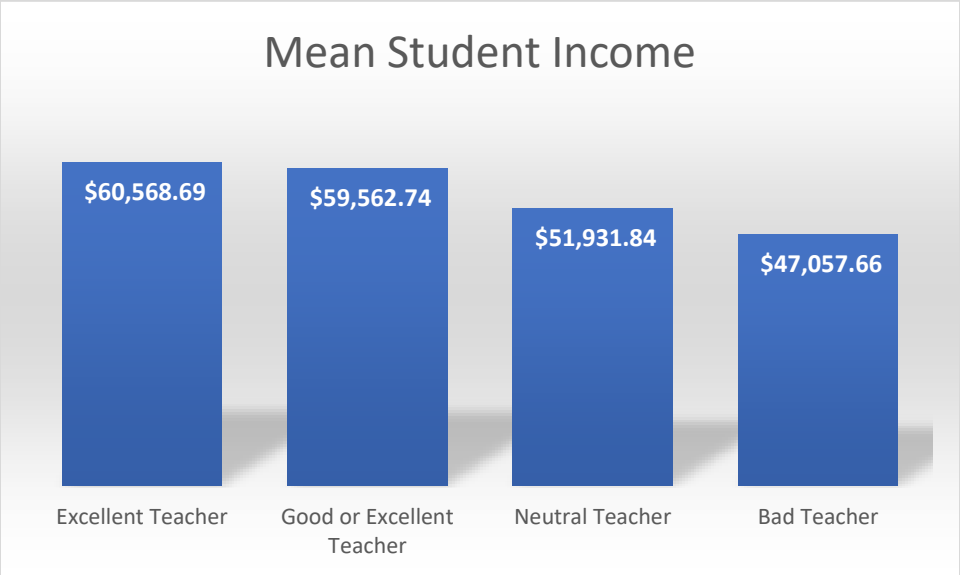

Figure 2 - Mean Highest Grade Completed

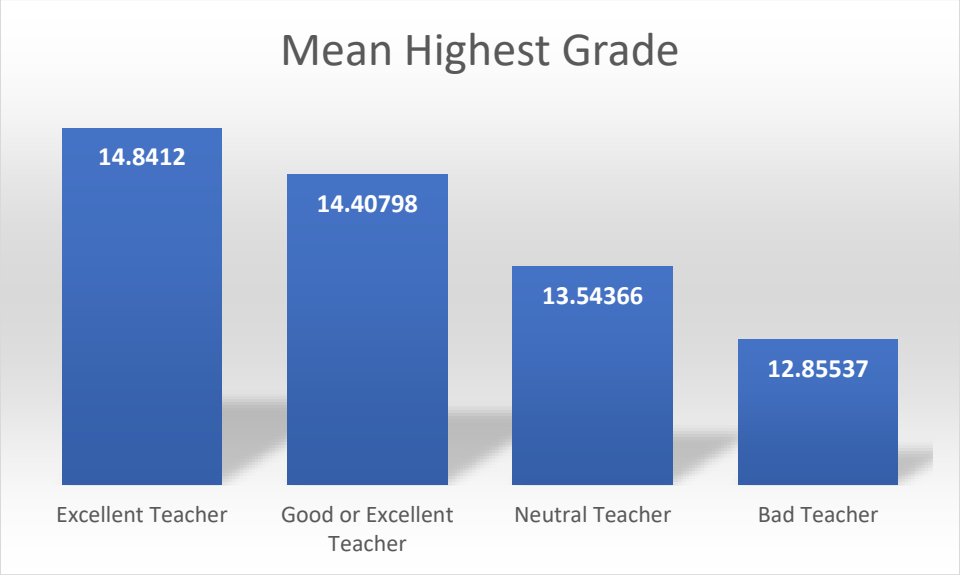

Table 7 - Full Regression Results Predicting Income

	Dependent variable:
	income
good_teacher	4,247.997***
	(1,607.189)
bad_teacher	−2,522.752
	(4,215.892)
excellent_teacher	173.508
	(2,709.511)
sex	−18,015.240***
	(1,380.735)
black	−14,033.560*
	(7,356.383)
hispanic	−7,851.795
	(7,377.824)
white	1,185.277
	(7,299.774)
fathers_highest_grade	839.395***
	(101.443)
urban	6,486.857***
	(1,580.589)
Constant	72,503.900***
	(7,754.428)
Observations	5,396
R^2	0.073
Adjusted R^2	0.071
Residual Std. Error	50,533.420 (df = 5386)
F Statistic	46.983*** (df = 9; 5386)
Note:	*p<0.1; **p<0.05; ***p<0.01

*Return to income regression model discussion

Table 8 - Full Regression Results Predicting ln(Income)

	Dependent variable:
	logincome
good_teacher	0.083***
	(0.028)
bad_teacher	−0.070
	(0.075)
excellent_teacher	0.026
	(0.048)
sex	−0.352***
	(0.025)
black	−0.274**
	(0.131)
hispanic	−0.135
	(0.132)
white	−0.014
	(0.130)
fathers_highest_grade	0.014***
	(0.002)
urban	0.104***
	(0.028)
Constant	10.972***
	(0.138)
Observations	5,360
R^2	0.077
Adjusted R^2	0.076
Residual Std. Error	0.894 (df = 5350)
F Statistic	49.634*** (df = 9; 5350)
Note:	*p<0.1; **p<0.05; ***p<0.01

*Return to ln(income) regression model discussion

Table 9 - Full Regression Results Predicting Highest Grade Completed

	Dependent variable:
	highest_grade
good_teacher	0.666***
	(0.121)
bad_teacher	−0.630**
	(0.296)
excellent_teacher	0.474**
	(0.212)
sex	0.816***
	(0.105)
black	−0.603
	(0.560)
hispanic	−0.842
	(0.563)
white	−0.068
	(0.557)
fathers_highest_grade	0.114***
	(0.008)
urban	0.239**
	(0.121)
Constant	11.549***
	(0.591)
Observations	6,850
R^2	0.065
Adjusted R^2	0.064
Residual Std. Error	4.335 (df = 6840)
F Statistic	53.048*** (df = 9; 6840)
Note:	*p<0.1; **p<0.05; ***p<0.01

*Return to HGC regression model discussion

Charter School Effectiveness

The struggles of traditional public schools to raise educational attainment has led to a search for innovative alternatives. One such alternative that has become increasingly popular since its inception in the 1990s is the Charter School. Charter Schools are publicly funded schools, but which have a large degree of independence from institutions like school boards and

teacher's unions. In theory, this autonomy allows charter schools to make innovative decisions unconstrained by the bureaucratic apparatus of traditional public education. Proponents of charter schools argue that these schools can make difficult but beneficial choices about personnel, discipline, curriculum, and allocation of resources because they do not have to worry about teachers unions or school boards.

The empirical research surrounding the effectiveness of these schools has yielded results indicating that charter schools have a very significant effect for certain targeted populations, but often yields little, or negative effects for other populations. The Center for Research on Education Outcomes (CREDO) has published dozens of high-quality studies using standardized test scores and a value-added model to evaluate the effectiveness of charter schools on a state by state and a national level. CREDO's 2013 survey of charter schools in 26 states found that charter schools on average delivered an additional eight days of learning compared to neighboring traditional public schools (Raymond, et al., 2013). Large scale studies like these continue to bolster the idea that charter schools may be a promising model for improving student outcomes in the U.S. This may have much to do with the way that charter schools are able to manage human capital.

Recent literature has examined differences in school structure and performance and has found some important differences in the factors which predict an effective school. Researchers Angrist, Pathak, and Walters looked at high achieving charter schools in Massachusetts and speculated on the factors that made them so successful. In particular, they looked at the "No-Excuses" charter school model in which school days are longer, academic rigor is high, discipline is strictly enforced, and students are held to high standards. Angrist et al. found significant differences between urban charter schools and traditional public schools in Massachusetts (Angrist, Pathak, & Walters, 2011). Urban charter schools were far more likely to identify with the No-excuses model, and measurable characteristics like school time, parent school contracts, and extended tutoring confirm this (Angrist, Pathak, & Walters, 2011). Non-Urban charter schools in Massachusetts were far less likely to identify with the No-Excuses model (Angrist, Pathak, & Walters, 2011). Urban charters were found to be far better at raising student test scores conditional on demographic variables than traditional charter schools (Angrist, Pathak, & Walters, 2011). Urban charters also proved to be far more effective than non-urban charter schools which seems to indicate that something about the No-Excuses model is particularly effective at influencing student achievement (Angrist, Pathak, & Walters, 2011).

Some of the most important studies regarding charter school effectiveness have come from Harvard economist Roland Fryer. Fryer and Will Dobbie studied New York city charter schools and found, like Angrist et al. that the No-Excuses school model was particularly effective at raising the academic achievement of disadvantaged students (Fryer & Dobbie, Getting Beneath the Veil of Effective Schools: Evidence from New York City, 2013). Specifically, they found that a culture of high expectations, extended instructional time, high dosage tutoring, data driven instruction, and frequent teacher feedback were important in determining school effectiveness. These are what Fryer and Dobbie term the five tenets of effective schools, which tend to be prevalent in schools that identify with the No-Excuses model. Fryer and Dobbie found that the five tenets of effective schools had a significant impact on the test scores of students attending those schools even after controlling for student demographics.

Specifically, they assigned much of the gain in test scores to the human capital practices of these schools which included more teacher feedback from administrators, and more data driven instruction.

Fryer and Dobbie found in another (Fryer & Dobbie, Medium-Term Impacts of High-Achieving Charter Schools on Non-Test Score Outcomes, 2013) 2013 paper that students who attend high performing charter schools not only have higher test scores than their demographic peers in traditional public schools, but they also have many long term non-academic improvements. The two found that those who attended high performing/no-excuses charter schools were far less likely to be incarcerated, less likely to become pregnant as a teenager, and more likely to attend college. These longer-term outcomes solidify the case that the no-excuses school model, which includes an emphasis on maximizing teacher's human capital. This contributes to the long list of rigorous empirical work which demonstrates the importance that education plays in the lives of students, and the fact that human capital is an integral part of the education production function. A good education does more than simply raise test scores, but it actually reduces risky behavior and improves long term outcomes for students.

One important difference between charter schools and traditional public schools is in how they utilize and manage human capital. One essential aspect of a charter school is its level of autonomy in making personnel decisions. Traditional public schools are often bound by contracts with teachers unions which limit their ability to fire teachers and frequently institute strict seniority requirements. Charter schools, because they are so often union-free, can fire teachers much more easily, and seniority has much less influence in any personnel decisions made by the school. The differences in personnel policy between charters and traditional public schools have been extensively documented by researchers Podgursky and Ballou (Podgursky & Ballou, Personnel Policy in Charter Schools, 2001). They find that charter schools generally demand much more from teachers which is likely due to the paucity of collective bargaining agreements in charters. Teachers at charter schools generally work longer hours and spend more total days in school per year than those at public schools (Podgursky & Ballou, Personnel Policy in Charter Schools, 2001). Charter schools often do not require teachers to be certified which allows them a much wider range of options on the labor market compared to traditional public schools. The pay scales are also radically different between school types (Podgursky & Ballou, Personnel Policy in Charter Schools, 2001). Charter schools are more likely to use merit-based pay than public schools, and they also tend to pay teachers more in hard to staff areas like science and mathematics (Fryer & Dobbie, Medium-Term Impacts of High-Achieving Charter Schools on Non-Test Score Outcomes, 2013).

Podgursky and Ballou posit that the relative freedom afforded to charter schools allows them to innovate and maximize the productivity of their human capital (Podgursky & Ballou, Personnel Policy in Charter Schools, 2001). Charter schools use a variety of tools to measure teacher quality when determining merit-based pay, but among these are student surveys, parent surveys, peer evaluation, and student value-added measures. Traditional public schools on the other hand tend to have a seniority-based schedule in which salaries are raised by a set amount at set intervals. This gives teachers more wage security, but critics of this system maintain that it fails to promote teacher achievement.

Another important difference in human capital management between charters and traditional public schools is that charter schools have significantly higher turnover rates than public schools (Podgursky & Ballou, Personnel Policy in Charter Schools, 2001) (Stuit & Smith, 2010). A 2010 study by Stuit and Smith estimated that the attrition rate at charter schools is 130 percent higher than at traditional public schools (Stuit & Smith, 2010). There are probably several reasons for this. Charter school teachers tend to have very different characteristics than those of traditional public schools (Podgursky & Ballou, Personnel Policy in Charter Schools, 2001). They are younger, less educated, less likely to have studied education, and less likely to have a teaching certificate (Podgursky & Ballou, Personnel Policy in Charter Schools, 2001). This population may be much less committed to a career in teaching than a population which has gone to the trouble of getting a teaching certificate or studied education in college. Teaching is a daunting profession, and it may be the case that charter schools are attracting workers who are not prepared for the demands of the classroom, resulting in high turnover.

In explaining teacher turnover, it cannot be ignored that working conditions are often much worse in charter schools (Stuit & Smith, 2012). Charter school teachers are asked to work longer hours, have less time off, and are paid less than traditional public school teachers which may reduce incentives to remain in a demanding profession (Podgursky & Ballou, Personnel Policy in Charter Schools, 2001) (Stuit & Smith, 2012). Stuit and Smith 2012 point to data from the 2004 NCES teacher follow up survey that indicates that charter school teachers are three times more likely than teachers in traditional public schools to cite poor working conditions as the reason for leaving a school (Stuit & Smith, 2012). Stuit and Smith's analysis proposes that the most important contributing factor to the charter TPS turnover gap is the lack of teacher union membership in charter schools (Stuit & Smith, 2012). This is significant because the inability of teachers unions to get a foothold in charter schools has been a topic much praised by charter advocates. However, it seems likely that the benefits which union membership confers on teachers may be a large reason why turnover rates are much lower in traditional public schools. Ultimately more data is needed to understand the tradeoffs between the restraints that a teacher's union puts on administrator's ability to fire bad teachers, and the lower turnover rate that may accompany a unionized workplace.

It is important to consider charter schools in any conversation regarding how to improve student outcomes. Charter schools are a promising alternative to traditional public schools largely due to their relative autonomy over human capital decisions. Charter schools do not always have to pay for inputs like a masters degree or an educational certificate for teachers (Podgursky & Ballou, Personnel Policy in Charter Schools, 2001). They are more free to pay STEM teachers more, and most importantly, they have more freedom to pay teachers based on merit, rather than on a flat salary schedule (Podgursky & Ballou, Personnel Policy in Charter Schools, 2001). This last consideration may be one of the most important for charter schools. The ability to reward good teachers more allows administrators to incentivize better performance, and to retain good human capital. When considering solutions to America's education crisis, we cannot ignore school models which maximize the utility of their human capital.

Charter Management Organizations

It is important to note that there are significant differences in human capital policies between types of charter schools as well. While most charter schools are "freestanding" meaning that the charter is held locally, about 44% of charter schools are managed by a "management organization which manages several schools (NAPCS). These management organizations are classified as either a non-profit "charter management organization" or a for-profit "educational management organization." While management organizations allow for economies of scale, teachers who work under their banner seem to have much less freedom, and much higher turnover rates (Roch & Sai, 2018).

Bibliography

Angrist, J. D., Pathak, P. A., & Walters, C. R. (2011). Explaining Charter School Effectiveness. *National Bureau of Economic Research.*

Armor, D., Conry-Oseguera, P., Cox, M., King, N., McDonnell, L., Pascal, A., . . . Zellman, G. (1976). *Analysis of the School Perferred Reading Program in Selected Los Angelas Minority Schools.* Santa Monica: RAND Corporation.

Chetty, R., Friedman, J. N., & Rockoff, J. E. (2014). Measuring the Impact of Teachers I: Evaluating Bias in Teacher Value-Added Estimates. *American Economic Review*, 2593-2632.

Chetty, R., Friedman, J. N., & Rockoff, J. E. (2014). Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood. *American Economic Review. Vol. 104. No. 9.* , 2633-79.

Coleman, J. S. (1968). Equality in Educational Opportunity. *Equity and Excellence in Education*, 19-28.

Fryer, R. G., & Dobbie, W. (2013). Getting Beneath the Veil of Effective Schools: Evidence from New York City. *American Economic Association. Vol 5, No. 4*, 28-60.

Fryer, R. G., & Dobbie, W. (2013). Medium-Term Impacts of High-Achieving Charter Schools on Non-Test Score Outcomes. *National Bureau of Economic Research*, 985-1037.

Hanushek, E. A. (1992). The Trade-off between Child Quantity and Quality. *The Journal of Political Economy. Vol 100. Issue 1.*, 84-117.

Hanushek, E., & Rivkin, S. G. (2006). *Handbook of Economics of Education. Volume 2, Chapter 18.* Elsevier.

Hanushek, E., Rivkin, S., & Kain, J. F. (2005). Teachers Schools and Academic Achievement. *Econometrica: Journal of the Econometric Society, Vol. 73. Issue 2*, 447-448.

Hanushek, E., Woessman, L., Peterson, P. E., & Summers, L. H. (2013). *Endangering Propsperity: A Global View of the American School.* Brookings Institution Press.

Ingersoll, R. M., & Smith, T. M. (2003). The Wrong Solution to the Teacher Shortage. *Educational Leadership. Vol. 60. Number 8.*, 30-33.

Kane, T. J., McCaffery, D. F., Miller, T., & Staiger, D. O. (2013). *Have We Identified Effective Teachers? Validating Measures of Effective Teaching Using Random Assignment.* Measures of Effective Teaching Project.

Mincer, J. A. (1974). Schooling and Earnings. *National Bureau of Economic Research, Schooling, Experience, and Earnings*, 41-63.

Murnane, R. J. (1975). *The Impact of School Resources on the Learning of Inner City Children.* Cambridge, MA. : Balinger Publishing Company.

Murnane, R. J., & Phillips, B. R. (1981). What do Effective Teachers of Inner City Children Have in Common? *Social Science Research*, 98.

Podgursky, M. (2003). Fringe Benefits: There is More to Compensation than a Teacher's Salary. *Education Next, Vol 3. No. 3*, 71-76.

Podgursky, M., & Ballou, D. (2001). *Personnel Policy in Charter Schools.* Washington D.C.: Thomas B. Fordham Foundation.

Raymond, M. E., Woodworth, J. L., Negassi, Y., Lawyer, K., Dickey, K., Davis, D., & Cremata, E. (2013). *National Charter School Study 2013.* Stanford, CA: Center for Research on Educational Outcomes, Stanford University.

Roch, C. H., & Sai, N. (2018). Stay or Go? Turnover in CMO, EMO, and Regular Charter schools. *The Social Science Journal. Vol 55. Issue 3*, 232-244.

Stuit, D. A., & Smith, T. M. (2012). Explaining the Gap in Charter and Traditional Public School Teacher Turnover Rates. *Economics of Education Review. Vol. 31. Issue 2*, 268-279.

Stuit, D., & Smith, T. M. (2010). *Teacher Turnover in Charter Schools.* Nashville, TN: National Center on School Choice, Vanderbilt University.