12-1-2017

Are Teacher Pensions "Hazardous" for Schools?

Patten Priestley Mahler
Centre College

Upjohn Institute working paper ; 18-281

Follow this and additional works at: https://research.upjohn.org/up_workingpapers

Part of the [Education Economics Commons](#), and the [Labor Economics Commons](#)

Citation

Mahler, Patten Priestley. 2018. "Are Teacher Pensions 'Hazardous' for Schools?" Upjohn Institute Working Paper 18-281. Kalamazoo, MI: W.E. Upjohn Institute for Employment Research. <https://doi.org/10.17848/wp18-281>

This title is brought to you by the Upjohn Institute. For more information, please contact repository@upjohn.org.

Are Teacher Pensions “Hazardous” for Schools?

Upjohn Institute Working Paper 18-281

Patten Priestley Mahler
Centre College
e-mail: patten.mahler@centre.edu

December 2017

ABSTRACT

I use a detailed panel of data and a unique modeling specification to explore how public schoolteachers respond to the incentives embedded in North Carolina’s retirement system. Like most public-sector retirement plans, North Carolina’s teacher pension implicitly encourages teachers to continue working until they are eligible for their pension benefits, and then leave soon afterward. I find that teachers with higher levels of quality, as measured by a teacher’s value-added to her students’ achievement test scores, are more responsive to the “pull” of teacher pensions. Younger teachers, those with higher salaries, and nonwhite teachers are also more likely to stay during the pension “pull.” All teachers show a strong response to the pension “push,” with about a quarter of teachers leaving every year once they become eligible for their pension. I depart from other models of teacher retirement by using a Cox proportional hazard model. Given that salaries are generally fixed by the state, I find that the number of years a teacher must work before she is eligible for her full pension benefit is the major driver of variation in pension wealth. This specification has the benefit of a flexible baseline hazard that can easily capture the sharp incentives driving a teacher’s retirement decision that are dependent on her proximity to retirement eligibility, and can flexibly account for differences driven by local labor market conditions. These analyses highlight important unintended effects that inform education policies going forward to ensure the retention of high-quality teachers in all types of schools.

Key Words: Teacher Retirement, Teacher Pensions, Public Expenditure, Public Pensions, State Finance, Nonwage Benefits

JEL Classification Codes: H55, H72, H75, J26, J32, I21

Acknowledgments: This research uses data accessed through the North Carolina Research and Data Center and was funded by the Center for Retirement Research at Boston College; the Institute of Education Sciences, U.S. Department of Education, through Grant R305B090002 to the University of Virginia; the Bankard Fund for Political Economy at the University of Virginia; and the W.E. Upjohn Institute for Employment Research. All opinions and errors are solely attributed to the author.

Upjohn Institute working papers are meant to stimulate discussion and criticism among the policy research community. Content and opinions are the sole responsibility of the author.

Are Teacher Pensions “Hazardous” for Schools?

Retirement of experienced teachers could have negative consequences on student learning and further exacerbate the disproportionate demand for high-quality teachers in disadvantaged schools. Many public pension systems provide incentives for teachers as young as age 50 to retire. States are considering, or have recently enacted, changes to these pension systems for purely fiscal reasons, ignoring how these changes might affect the retention of effective teachers who have important impacts on student learning or staffing in hard-to-staff schools.

Teacher pension funds fall short of their liabilities by an estimated \$1 trillion, causing many states to consider cost-saving changes, including restructuring their plans (Pew Center on the States 2010). Traditionally, public schoolteachers in almost every state have been covered under a defined benefit pension plan, which is markedly different from the defined contribution plans (e.g., a 401[k]) offered in most private-sector jobs. Defined benefit plans provide teachers with a fixed annuity paid regularly over their retirement after they have reached certain age-experience thresholds. In contrast, a defined contribution plan puts money aside in an account that can be accessed by the individual after retirement but does not provide any incentive to retire in one particular year or the next. The present structure of defined benefit plans provides incentives to “pull” midcareer individuals toward continuing teaching, while “pushing” later-career teachers out of the profession. These incentives may affect the retirement behavior of teachers differently depending on their alternative career options or enjoyment of teaching. This in turn could increase the challenges in already hard-to-staff schools or degrade the overall quality of education. Therefore, one must ask, “How responsive are teachers to the pull and push of pension incentives? Do teachers with different characteristics (qualifications, effectiveness,

demographics) have different responses to pension incentives? Do teachers in different school environments (in terms of student racial or ethnic composition, student poverty, or grade level) have different responses to pension incentives?”

I find that teachers are very responsive to the pull of teacher pensions, with fewer than 6 percent of teachers leaving each year. Teachers with higher levels of quality, measured by a teacher’s value-added (the amount she contributes to her students’ achievement test scores), are more responsive to the pull of teacher pensions, perhaps because they find teaching more satisfying. Teachers who are younger, have higher salaries, or are nonwhite are less likely to retire than their counterparts. Teachers with advanced degrees or those who attended either more competitive or less competitive colleges are more likely to retire than those with bachelor’s degrees or those who attended colleges rated as being of average competitiveness. Teachers in schools with a higher proportion of students who are economically disadvantaged are more likely to leave. Middle-school teachers are more likely to retire than elementary or high-school teachers and seem to be particularly sensitive to the racial composition of their school, as more of them exit during the pension pull at schools that have a larger proportion of black students.

All teachers show a strong response to the pension push, with about a fourth of teachers leaving every year once they become eligible for their pension. High- and low-value-added teachers exit at similar rates. As with the pension pull, teachers who are young, have higher salaries, or are nonwhite are more likely to stay. Teachers who attended less competitive colleges are more likely to stay during the pension push, perhaps because they have fewer alternative career opportunities. Middle-school teachers in schools with a large proportion of black students who are still teaching during the pension push are less likely to

leave (as many of those have already left prior to pension eligibility), while middle- and high-school teachers in schools with a large proportion of economically disadvantaged students are more likely to leave.

My theoretical model differs from the canonical models that have been used to study retirement because I include nonpecuniary benefits from work and utilize a different estimation strategy. Teachers are paid according to a rigid salary schedule, so these nonmonetary factors are an important extension to explain differences in retirement behavior that cannot be explained by differences in compensation. I estimate a teacher's propensity to retire using a Cox proportional hazard model, which changes the effective unit of analysis from a teacher's year-by-year retirement decision to a teacher's entire spell of employment. For example, a logit or probit identifies what factors are associated with retirement in a particular year, while a proportional hazard model identifies what factors are associated with shorter careers. A Cox proportional hazard model has the benefit of a flexible baseline hazard that can easily capture the sharp incentives driving a teacher's retirement decision that are dependent on her proximity to retirement eligibility, and can flexibly account for differences in this hazard rate driven by local labor market conditions that vary by year and region.

Given that salaries are generally fixed by the state, I find that the number of years a teacher must work before she is eligible for her full pension benefit is the major driver of variation in pension wealth. In North Carolina, eligibility for full pension benefits occurs at set age-experience combinations. I define the unit of time as the number of years until full pension receipt eligibility, which means all teachers who have the same number of years to pension receipt have a common baseline probability of retirement. By choosing the unit of time to be years until pension eligibility, my baseline hazard answers my first research

question about how teachers respond to the “pull” and “push” of pensions—and it does so in a fully flexible, nonparametric way. Even though the baseline hazard is not estimated, it can be calculated after estimation by setting all covariates equal to zero.

In the next section, I describe the literature on modeling pension incentives and provide evidence that, despite receiving similar salaries, teachers with different levels of quality (the value-added they bring to the job, or the competitiveness of the college they attended) or working conditions vary in terms of their persistence in the workforce. This evidence suggests that teachers’ responses to pension incentives may also vary along these dimensions. I use a unique set of detailed data on every North Carolina teacher and student over 14 years, as described in the third section. I can follow teachers over time and explore their responses to the North Carolina retirement plan as they approach and pass pension receipt eligibility. These data have several critical features that are often absent in other studies of retirement. First, I observe characteristics of their workplace necessary to decipher whether teacher retirement behavior differs by school characteristics, such as student demographics. Second, I link fourth- and fifth-grade teachers with their students who took end-of-grade achievement tests. This allows me to evaluate a teacher’s quality in terms of her prior teaching performance, measured using teacher value-added, to see how retirement behavior varies with teaching effectiveness.

I explain the North Carolina pension plan in the fourth section and describe how the pension plan affects teachers’ retirement incentives. In the fifth section, I present my estimation strategy and results, including the test of the proportionality assumption critical to the Cox model. I conclude in the sixth section with some policy implications of this research.

My Contributions and Review of the Literature

Modeling retirement incentives is complicated because the value of continuing to teach includes the value of the having the option to continue one year later, two years later, and so on. Rust (1989) and Berkovec and Stern (1991) model these incentives using a dynamic structural model. Stock and Wise (1990) make some simplifying assumptions and approximate a dynamic structural model using the “option value” approach where individuals continue to work because they value the option of being able to retire later. Coile and Gruber (2007) further simplify the dynamic nature of pensions by calculating just one number, the “peak value,” to symbolize the option value of continuing to work. The peak value is how much an individual’s pension wealth would grow should she continue working until she reached the peak of her pension wealth.

Economists have also used hazard models to study retirement (e.g., Hausman and Wise 1985). In an appendix to an earlier version of their 1990 published paper, Stock and Wise (1988) show that a hazard model can be thought of as a simplified version of their option value model, which is derived from utility maximization. I argue that there are advantages to using a hazard model for measuring teacher responses under one pension system. The main sources of variation in peak value across individuals of the same age and tenure are due to differences in pension plan rules or differences in earnings. Nationally representative data have variation across these two margins; state-level teacher data likely do not. Pension plans are often administered at the state level; salaries are also usually determined at the state level, with small differences at the district level, where everyone generally receives the same bump in salary. The only meaningful variable that is driving peak value differences is the number of years a teacher has until she is eligible for full pension benefits. I look at how retirement behavior is affected simply by the number of years until a teacher is eligible for full pension benefits, bypassing the assumptions about how individuals discount future wealth and the relevance of peak value to a teacher’s decision.

Teachers are arguably very cognizant of the number of years they have left, making it a simple, realistic variable that would naturally affect retirement behavior and is driven by the pension rules.

A few studies examine teacher pension incentives (Costrell and McGee 2010; Costrell and Podgursky 2009; Friedberg and Turner 2010) and describe how pension wealth accrues over a teacher's career, sometimes creating incentives to pull them to stay in teaching, and sometimes pushing them out. The pension pull is caused by an increase in the annuity due to additional experience and salary, and by an increase in the present discounted value of a teacher's pension because she is eligible to begin receiving the annuity sooner. After eligibility for her pension, the pension "push" is caused by the high cost of forgone pension benefits if she continues to teach. Although these studies generally concur that teachers behave predictably in the face of pension incentives, Fitzpatrick (2015) and Koedel and Xiang (2017) show that teachers likely value them at much less than they cost.

Three papers build on the peak value approach in Coile and Gruber (2007) by measuring differential responses among teachers: Costrell and McGee (2010), Friedberg and Turner (2011), and Koedel, Podgursky, and Shi (2013). Costrell and McGee (2010) estimate peak value models for Arkansas teachers and find that teachers are indeed responsive to peak value. Friedberg and Turner (2011) use the peak value approach to model retirement in nationally representative data in the Schools and Staffing Survey. They find evidence that teachers who are less satisfied with their jobs are likely to delay retirement during the pension pull and then retire abruptly during the pension push. My results suggest that the demographic composition of the school influences teacher retirement, perhaps because it is correlated with teacher satisfaction. I also look at the differences in retirement behavior by teacher value-added, and I find that high-value-added

teachers are the least likely to leave during the pension pull but behave similarly to low-value-added teachers during the pension push.

Koedel, Podgursky, and Shi (2013) examine the relationship between value-added and retirement, but through a different lens. They do not leverage variation in pension wealth or peak value, but instead use peak value to determine teacher type. Teacher type refers to whether one ends up retiring very early in one's career, around the time of pension receipt eligibility, or several years after pension receipt eligibility. Koedel, Podgursky, and Shi estimate the value-added of different types of teachers in Missouri using a student-level regression with indicators for teacher type. They find no evidence of teacher type being associated with teacher quality. Teachers who are still teaching at the end of their panel are included in the omitted category because, at the time of analysis, it is unknown what teacher type they will become; thus, the comparison group is a combination of teachers from the other teacher types. My hazard model specification allows teachers with censored retirement to be included without making the assumption that they are similar to one another. Instead, these teachers are simply counted as still teaching and I make no supposition about when they might retire.

There are a growing number of papers examining the effects of changes in retirement benefits (e.g., Fitzpatrick [2015]; Fitzpatrick and Lovenheim [2014]; Furgeson, Strauss, and Vogt [2006]; Koedel and Xiang [2017]). The most closely related study is Fitzpatrick and Lovenheim (2014), which finds that early retirement incentives have little effect on student test scores and may increase scores in schools with a large proportion of economically disadvantaged students. My results complement theirs in that I find that low-value-added teachers are the least attached to their jobs during the pension pull.

My study adds to a large body of literature on teacher retention, which most often has studied teachers who are early in their careers. Many papers study the effect of current compensation on teacher turnover using variation in pay scales across districts, which may be endogenous, or using salaries in alternative occupations, which are difficult to accurately measure (e.g., Hanushek, Kain, and Rivkin [2004]; Podgursky, Monroe, and Watson [2004]). I use variation in benefit eligibility thresholds, which are primarily driven by exogenous factors such as a teacher's age and years of experience.

Another branch of literature focuses specifically on the quality of retained teachers. These studies show that teacher retention varies by teachers' general-knowledge test scores and the competitiveness of the college attended, finding that those with better credentials are more likely to exit (Boyd et al. 2005; Lankford, Loeb, and Wyckoff 2002). Researchers have also developed the value-added measure of quality, which is the average growth in achievement that a teacher's students experience during the school year (Rivkin, Hanushek, and Kain 2005). Many studies compare the value-added measures of exiting early-career teachers to the measures of those who stay. These studies find that those with higher value-added are more likely to stay (e.g., Boyd et al. [2007]; Goldhaber, Gross, and Player [2011]). I study the behavior of mid- and later-career teachers with respect to their qualifications, and I measure their value-added over multiple years to get a more informed measurement of a teacher's effectiveness.¹

Researchers have additionally studied the sorting of different types of teachers across schools, concluding that schools with more disadvantaged students are likely to have high teacher turnover and generally less effective teachers (Boyd et al. 2007; Clotfelter, Ladd, and

¹ Rothstein (2009) notes the bias introduced in value-added measures due to the nonrandom assignment of students and teachers. Koedel and Betts (2011) show that measuring value-added with many years of data, as I do, reduces this bias.

Vigdor 2006). I find that teachers in hard-to-staff schools are more likely to leave during both the pension pull and the pension push, and that less effective teachers are more likely than their effective counterparts to leave during the pull.

As stated above, many of my contributions are the result of an exceptionally detailed and encompassing data set and of rethinking the hazard model specification. I describe the specifics of the data in the next section and motivate the hazard model specification in the section after that.

Data

I use data from administrative records of all North Carolina public school students and teachers over the 1994–1995 through 2010–2011 school years, maintained by the North Carolina Education Research Data Center. These records include key individual and aggregate characteristics, follow individuals over time, and link students and teachers in school classrooms. I describe my sample selection process (Table 1), possible sources of measurement error, estimation strategy for teacher value-added (Table 2), and summary statistics of key variables (Table 3).²

I make two consequential exclusions in order to ensure that my model is properly specified. First, I am interested in teachers' exit behavior as it relates to pension incentives; therefore, I exclude teachers with less than five years of experience because they are not vested in the pension plan. Given that I see only the presence or absence of a teacher in the administrative data, I have to infer whether her absence is due to leaving teaching permanently

² I include those who have ever been full-time teachers with greater than five years of experience (making them eligible for the defined benefit pension plan) if their date of college graduation and other demographic variables are given, and if they have worked in a single school with nonzero salary during each year in the data.

(exited) or only temporarily. I define teachers as having exited if they are not observed as full-time employees in the North Carolina data for two consecutive years.³ I also exclude spells of employment for anyone under 40 years old, when the pension pull is likely very small, as they have at least 10 more years before they will be eligible for their pension. Exits from (and reentry to) teaching during a person's twenties and thirties may be correlated with child-rearing, a factor outside the scope of my model that could result in a temporary absence that is longer than two years, leading me to misclassify them as having exited when they have not.

Second, as explained in the section on "Estimation Strategy and Results," a critical assumption of a proportional Cox model is that all covariates must affect the baseline hazard proportionally. There are two teacher characteristics that did not satisfy basic tests of this proportionality assumption, likely because they are correlated with other factors that influence retirement, such as household assets or marital status. One is teacher gender. It appears that the baseline hazard for men is not proportional to the baseline hazard for women, implying that men have a different underlying attachment to teaching. Excluding men from the analysis leaves 81 percent of teachers. The other is the pension eligibility threshold that a teacher is approaching. In North Carolina, teachers are eligible for full pension benefits once they reach the earliest of three thresholds: 1) 30 years of service, 2) age 60 with 25 years of service, or 3) age 65 with 5 years of service. Some 78 percent of teachers reach the first threshold (30 years of service) before the others, meaning they started teaching in their twenties or early thirties. Those who reach the other two thresholds first did not start teaching until much later. While the inclusion of individuals across all three thresholds provides valuable variation in age, it also results in the age

³ A drawback of this definition of "exited" is that the last two years of data (2009–2010 and 2010–2011) cannot be included in the analysis because there are not two later years that can be used to distinguish between temporary and permanent leaves of absence.

covariate violating the proportionality assumption. Teachers who start teaching at a later age may be systematically different, resulting in a different hazard rate. Exploring the retirement behavior of men and those who started teaching at later ages could be a topic of future research but is less applicable to the behavior of teachers overall. After these exclusions, my sample size is 29,799 (58 percent) full-time female teachers.

I require that each individual in the sample have reasonable values for key covariates, including age, experience, and salary, and that each individual continually work as a teacher (as opposed to an administrator or staff member) during her tenure. I classify schools according to the lowest and highest grade levels of the school in the Common Core of Data (CCD) each year: elementary (PK–1 to 1–8), middle (4–7 to 5–9), and high (7–12 to 12). Thus, for instance, if the earliest grade falls between prekindergarten and 1 and the latest grade in the school is between 1 and 8, then it is classified as an elementary school. If the data on student characteristics are missing, I impute the value from the closest nonmissing year for that school. I lack data on the demographic composition of some schools, leading me to exclude less than 2 percent of teachers. I standardize measures of school demographics for each school level. For example, I use the mean and standard deviation of the proportion of black students in a school across all years. I standardize within school level because schools with lower grades tend to be smaller and more homogeneous than large high schools. I associate teachers with the competitiveness of the institution from which they received their undergraduate degree, where competitiveness is measured according to Barron’s rankings of “Most Competitive Schools” from the year closest to their graduation date (1984, 1986, 1988, 1990, or 1992). “Less competitive” is a Barron’s rating of either noncompetitive or less competitive; “competitive” is a rating of competitive; and “more competitive” is a rating of either more competitive or very competitive.

Although the data are quite rich, there are two limitations that affect the precision with which I can measure eligibility for full pension receipt. First, there is measurement error in my principal measure for pension eligibility, the number of years of service. A suitable proxy for years of service is years of experience, which I can identify given an individual's salary step; however, some teachers may have credit for additional years of service beyond their years of experience because of work in other state agencies or credit for unused personal or sick days. Second, I do not observe a teacher's age, but I do know the year she graduated from college. I impute age assuming she is 21 upon college graduation. Both measurement error in age and years of service imply that there are some individuals who are eligible for pension benefits whom I label as ineligible. The issue with age is less critical, given that I limit my sample to those who are approaching the 30-years-of-service threshold.

An integral part of my analysis is in identifying a teacher's quality using value-added measures. The sample of teachers that I am able to link with end-of-grade math and reading scores is the subset of 3,828 of these teachers who taught fourth or fifth grade during the 1998–1999 and 2007–2008 school years. I choose this subset because of two data limitations. First, student end-of-grade scores are associated with the teacher who proctored the exam. In elementary school grades (third through fifth, which are the elementary grades that take tests), the proctor is likely to be the student's instructor for math and reading, because students are generally with one teacher for most of the day. In sixth through eighth grades, the proctor may not be the student's instructor for math or reading; instead, it is likely that students are proctored by their homeroom teacher, who may instruct them in another subject. I follow Xu, Hannaway, and Taylor (2007) and compare the student composition (class size, number of white students, number of male students) of the tested class with that of the class that the proctor instructs (data

from a separate source). If the characteristics of tested classroom are similar to the instructor's classroom, then I deem the proctor to also be the instructor. Second, in order to calculate a teacher's value-added, I need her students' prior test scores and demographic characteristics. The prior test scores of third-grade students are unavailable because there is no test prior to third grade, so I exclude third-grade teachers from my value-added analysis. The available student demographic data varies from year to year, and only the 1998–1999 through 2007–2008 school years contain the student information that I include in my value-added specification. I exclude classes with fewer than 10 or more than 40 students. Across all years, teachers for whom I calculate value-added have between 10 and 252 students in all of their classes combined, with the average being 72.

I measure teacher value-added, under the variable $ValueAdded_i$, as a time-invariant trait related to the average growth of student standardized achievement scores for a particular teacher relative to the growth of other teachers' student scores after netting out the average effects of other observable factors (Table 2). To address concerns about bias in value-added measures, I follow Koedel and Betts (2011) and estimate a teacher's value-added using the test scores of her students over multiple years:

$$(1) \quad Achievement_{jit} = ValueAdded_i + \mathbf{Y}_{jt}\boldsymbol{\rho} + v_{jit}.$$

The dependent variable is the normalized test score $Achievement_{jit}$ of student j of teacher i in year t . I regress this on a vector of student, class, and school attributes $\mathbf{Y}$, which includes student j 's test score from the previous year, demographic and achievement measures for other students in j 's classroom and school, and grade and year fixed effects. I do not include teacher experience in this specification. It is well established by the literature (Boyd et al. 2008; Rockoff 2004) that teachers' value-added increases as they gain experience, especially in the first few years of

teaching. Teachers included in the value-added specification are in their fifth or more year of teaching, so it is likely that these teachers have little variation in quality because of increases in experience. Additionally, the most policy-relevant comparison is to compare the retirement behavior of all teachers to one another. If I controlled for experience, the measure would only compare a teacher to her same-experienced peers. In the case of a teacher with 34 years of experience, this would only be those who also stayed four years past pension eligibility, which is a select group of teachers. Because my value-added measure compares all non-novice teachers to one another, I am able to describe the effect of value-added on the selection that caused some teachers to stay until they had 34 years of experience. I exclude school fixed effects because I want to compare the retirement behavior of individuals with different value-added measures across, not within, schools. I use empirical Bayes shrinkage (Value-Added Research Center 2010) to account for measurement error in the value-added estimates, although using an estimated coefficient as a covariate will admittedly be tenuous even after shrinkage.

I have value-added measures for 3,730 teachers, 12.5 percent of my analytic sample. Because there are separate math and reading tests, I compute a teacher's average value-added, the average of reading and math. The shrunken standard deviation of these measures is in line with other estimates from other value-added research—0.19 for math and 0.10 for reading—as are the magnitude and significance of covariates in the value-added regression (Table 2) (e.g., Clotfelter, Ladd, and Vigdor [2007]; Hanushek and Rivkin [2010]).⁴ In order to ease interpretation of value-added in my hazard model specification, I use a standardized measure so that teachers whose value-added is one standard deviation above the mean have a value-added equal to one.

⁴ The standard deviation of value-added may differ from other estimates because novice teachers are excluded, teacher experience is not accounted for, and there are no school or school-by-year fixed effects.

Table 3 shows descriptive statistics for teachers in the beginning and end of the time period I study for the full sample as well as the sample for whom I have value-added. As is consistent with the large cohort of baby boomers moving toward retirement, the average teacher age, experience, and teaching salary are increasing over time. Over half of teachers in the full sample are in elementary schools, with the remaining half split between middle and high school. Teachers in the value-added sample are primarily in elementary schools because these are individuals who have taught fourth or fifth grade. Most teachers attended colleges that fell into the “competitive” or “less competitive” categories, and around 40 percent have degrees beyond a bachelor’s. The vast majority are white.

Modeling Retirement Incentives

Teachers in North Carolina and most other states are eligible for a defined benefit pension plan that provides an annuity paid regularly after they retire from teaching. This annuity equals 1.82 percent of the average salary of their last four years, multiplied by the years of service in North Carolina public schools. These retirement benefits make up a large portion of a teacher’s total compensation; for instance, a 30-year teacher receives over half of her highest salary each year following retirement.⁵ Teachers are entitled to full pension benefits when they reach certain age-experience thresholds: 30 years of service, age 60 with 25 years of service, or age 65 with 5 years of service.⁶ As described in the third section, I limit my sample to teachers who will first become eligible for pension benefits at the 30-years-of-service threshold.

⁵ Teachers in North Carolina are also eligible for Social Security.

⁶ Teachers can retire short of these thresholds and receive a reduced benefit. Reduced benefit eligibility: age 50 with 25 years of teaching, or age 60 with 5 years of teaching. Benefits are reduced by 3–5 percent per year short of full retirement threshold, depending on age at retirement. Note: a teacher may retire before reaching an age threshold but not receive benefits until her age and experience are past the threshold—e.g., a 58-year-old could retire after her

Any promise of wealth or income after retirement is likely to affect teacher exit behavior. It is important to isolate the “push” and “pull” effects of defined benefit pensions from the overall income effects. The income effect can be summarized by the pension’s present discounted value, called “pension wealth.” Figure 1A shows pension wealth as a function of her experience at the time of exit for a hypothetical North Carolina teacher who starts teaching at age 21. Figure 1B shows how much pension wealth accrues at each level of experience. This defined benefit plan provides incentives to continue teaching up until one reaches 30 years of experience (52 years old) and can begin to receive her annuity immediately upon exiting teaching. Before a teacher reaches 30 years of experience, her pension wealth increases for several reasons. First, as is true throughout her career, an additional year of work, along with any associated increase in salary, significantly raises the annuity amount. Second, during this period, pension wealth increases because she is closer to receiving her benefits—one more year of work implies (at least) one less year she must wait to get her annuity. The accrual rate is at its highest when she has 27 years of experience, at which point she would receive an additional \$56,000 worth of pension wealth for teaching just one more year. This additional pension wealth may “pull” her to continue teaching. Once she reaches 30 years of experience (age 52) she can receive her annuity immediately upon exiting teaching, at which point the accrual drops from \$38,000 to \$15,000. If she continues to work, the annuity amount would still increase, but her pension accrual nonetheless declines because each additional year of teaching imposes an implicit cost—it is one less year that she could be receiving benefits. After 42 years of experience (age 65), this cost exceeds the additional annuity amount, making the accrual negative. These decreased accruals

twenty-fifth year and receive reduced benefits immediately and thereafter, or wait two years and get full benefits upon turning 60.

after one has attained 30 years of experience may “push” her out of teaching because they are either not worth waiting for or they effectively decrease the value of her retirement income.

In contrast, the accruals in a defined contribution plan are relatively flat and (theoretically) never decline. Defined contribution plans do have set rules for when a person can have access to her retirement account (generally after age 59½). Early access is possible by paying an early withdrawal penalty, exercising a withdrawal due to hardship, or borrowing against the retirement funds. Thus, these strict pension eligibility thresholds and sharp, changing accrual patterns are characteristics of defined benefit pensions alone.

In a defined benefit plan, the relative value of continuing to teach depends on how long one intends to teach. For example, someone who does not intend to stay until pension eligibility should not be prevented by the utility gains from staying until then. This attachment to teaching likely varies with teachers’ attributes, causing individuals with different characteristics to respond differently to retirement incentives. For example, teachers may have different predilections for their work because they enjoy teaching, get satisfaction from being an effective teacher, or have a fondness for a particular school environment. Teachers with high nonpecuniary benefits would be less influenced by the value of pension accruals as they relate to the teachers’ retirement decisions, making them less sensitive to both the pull and the push of pensions. On the other hand, teachers with talents that are highly valued in alternative professions may be less sensitive to the pension pull but more sensitive to the pension push because the opportunity cost of teaching is high. To capture differences in responsiveness to pension incentives, I see how hazard rates vary with characteristics like teacher quality, having an advanced degree, teaching different grades, and student demographics.

While the value and growth in pension wealth can describe the magnitude of pension incentives in great detail, pension eligibility thresholds effectively capture the same incentives, are more salient to individual teachers, and require fewer assumptions about individuals' discount rates or expectations about future salaries. Figure 2 shows the pension wealth for three hypothetical teachers that represent the spectrum of teachers in my sample. Around three-quarters of my sample begin teaching prior to age 25, while the remaining quarter begin by the time they are 30. Because of their different start ages, the teachers have different pension wealth distributions and vary in the year in which they are eligible for pension receipt as well as the accrual rates. Figure 2A shows these pension wealth distributions indexed by age. Figure 2B shows them indexed by pension eligibility: the number of years an individual has been eligible to receive full benefits immediately upon exiting teaching. While there are still differences across teachers with different start ages, there are many similarities, including the timing of large versus small pension accruals.

Using pension eligibility instead of pension wealth could disregard some potentially important details. First, if teachers were not paid according to a rigid, statewide salary scale, then their pension wealth distributions could vary substantially. Second, if teachers were subject to different pension systems across the state, then the pension growth patterns could be noticeably distinct from one another. These situations, which are easily found when comparing individual behavior across states or employers, could create meaningful differences within a given level of pension eligibility that should not be ignored. However, these factors are stable across many public pension systems, including those for emergency and safety personnel as well as government employees. In addition, I include controls for age and salary per month to account for these small differences in pension wealth conditional on pension eligibility, as well as the

effects of these variables on retirement through other channels (such as an increasing value of leisure).

Figure 3A shows teachers' hazard rate by years of experience and pension eligibility. This is the marginal proportion of teachers who exit at each experience level, conditional on working up to that point. Figure 3B shows the survival function, the proportion of teachers with 20 years of experience who are still working at each level of experience—in other words, the cumulative effect of the hazard rate. The large jump in the hazard rate starts at 29 years of experience, one year short of the 30-year threshold. Theoretically, this is when the incentive to exit teaching is very low, as waiting just one more year entitles you to full pension benefits upon leaving. Leaving one year shy of this threshold implies you must wait to receive pension benefits until you age into the next threshold (age 60 and 25 years of service), which could be as many as nine more years. I suspect that there is measurement error, in that years of experience is not the perfect proxy for years of service. Although I do not have the data to test this hypothesis, the likely culprit for this error is that teachers can take their unused sick days and turn them in as years of service on top of their existing years of experience. Teachers can have as many unused sick days as would equal a full school year, leading to this discrepancy. I cannot identify how many unused sick days a teacher has, making it impossible to correct this issue. Instead, I analyze exit behavior for those with 20 to 28 years of experience (pension eligibility –10 to –2, or between 10 and 2 years until eligibility), and those with 30 to 35 years of experience (pension eligibility 0 to 5, or having been eligible for zero to five years), omitting the suspect year.

Estimation Strategy and Results

As discussed in the previous section, given that salaries are generally fixed by the state and the pension system is statewide, I find that the number of years a teacher must work before she is eligible for her full pension benefit is the major driver of variation in pension wealth. I depart from other models of teacher retirement by using a Cox proportional hazard model, which changes the effective unit of analysis from a teacher's year-by-year retirement decision to a teacher's entire spell of employment. Binary models, such as a logit or probit, ignore the effect of time, treating the retirement decision every year as an independent decision related to different values of covariates. I define the unit of time as the number of years until full pension receipt eligibility, which means all teachers who have the same number of years to pension receipt have a common baseline probability of retirement, $\lambda_0(t)$. The functional form of the baseline hazard is not specified, meaning it can have sharp spikes or a smooth curve, depending on what fits the data. The model is shown in Equation (2):

$$(2) \quad \lambda(t|X_{it}) = \lambda_0(t) \cdot \exp(X_{it}\beta) .$$

A strength of the Cox proportional hazard model is that it explicitly handles censored data, meaning that individuals who had not yet retired at the end of my panel are properly included in estimating the effects of the covariates. Through a process known as stratification, these models can flexibly account for local labor market conditions by allowing individuals in different school districts (i.e., local education agencies) to have distinct baseline hazards. For example, a teacher in a rural district may be more attached to teaching regardless of pension eligibility simply because there are not a lot of other jobs available. Similarly, I stratify by school year to account for statewide variation in school

policies (e.g., the No Child Left Behind Act) or sweeping changes in job availability (e.g., the Great Recession).⁷ Stratification is akin to fixed effects, but because the baseline hazard is not parametrically identified, I cannot test the significance of differences across years or districts. I cluster standard errors at the teacher level.

A limitation of the proportional hazard model is the assumption that all covariates have a proportional effect on the baseline hazard. This can be tested by whether the slope of Schoenfeld residuals with respect to time is zero. To satisfy the proportional hazard assumption, the individuals in the sample need to be reasonably similar. In the literature review section, I describe why this implies that I focus solely on women who first reach pension eligibility when they have 30 years of service. A consequence of the latter provision is that I cannot separately control for experience. A second issue is that in years prior to 2007–2008, the last salary step is 30 years of experience and over, meaning that teachers will still likely get paid more if they teach the following year, but only because of an increase in the overall salary schedule, not a salary step.

I analyze the pension “pull” and “push” separately, allowing covariates to affect the baseline hazard in different ways depending on whether the teacher is not yet eligible for her pension or past eligibility. I exclude the year just prior to pension eligibility because of the measurement error described in the fourth section.

The results in Tables 4 and 5 are presented as hazard ratios. A hazard ratio of 1 implies that individuals with different values of this covariate have the same hazard of retirement. Note that even a large proportional effect on the baseline hazard is ultimately a

⁷ There was a policy in place from 1999 to 2009 in which North Carolina allowed retired teachers to return to work. My preliminary results suggest that, because the policy was constantly on the verge of being ended, it did not affect selection into retirement in a meaningful way. Nonetheless, flexibly controlling for year effects should negate the effects of this policy.

small effect if the baseline hazard is small. As shown in Figure 4, the baseline hazard prior to pension eligibility is below 0.06. This implies that fewer than 6 percent of teachers leave each year prior to pension eligibility. Figure 5 shows that the baseline hazard after eligibility is consistently over 0.22 for the full sample, yet consistently below 0.22 for the value-added sample—sometimes as low as 0.07. Note that because I drop the year prior to pension eligibility, I cannot calculate the survival function posteligibility conditional on teaching 10 years prior to eligibility (as in Figure 4B). Instead, I show the survival function as a function of those who were teaching at pension eligibility.

First, I discuss my results during the years leading up to pension eligibility (Table 4). The three columns have the results for value-added teachers only (1) and for two specifications with all teachers (2 and 3). Across all specifications, a one-year increase in age is associated with a 9 percent higher likelihood of retirement. An increase in salary of \$100 per month decreases the likelihood of retirement by at least 5 percent. Having a master's degree or higher could represent additional general human capital, as those with advanced degrees are 27 to 67 percent more likely to leave than their bachelor's degree counterparts. Teachers who attended competitive colleges are less likely to leave than their colleagues who went to either less competitive or more competitive colleges. Nonwhite teachers are over 20 percent less likely to retire than their white counterparts.

Teachers with value-added one standard deviation above the mean are 20 percent less likely to retire than average teachers. Perhaps the pension “pull” is not strong enough to keep those who might be dissatisfied with their jobs because they are not as successful (as measured with this metric). As shown in Figure 6A, the hazard rate for teachers whose value-added is -1 , 0 , or 1 standard deviations above the mean does not look very different;

however, these small differences in hazard rates result in 8 percent more high-value-added teachers than low-value-added teachers (86 versus 78) continuing in the classroom until almost eligible for their pension.

With respect to school characteristics, middle-school teachers are 14 percent more likely to leave than high-school or elementary-school teachers. A one-standard-deviation increase in the proportion of students in the school who are economically disadvantaged leads to a 9-to-10 percent increase in retirement. A one-standard-deviation increase in the proportion of black students leads to an 11 percent increase in retirement for middle-school teachers but no significant difference in elementary or high school exits.

Table 5 includes the hazard ratios for teachers once they are eligible for their pension; these hazard ratios show more similarities than differences across teachers. Note that even though the effects on the covariates are generally smaller than during the pension pull, the overall effects may be larger because the baseline hazard is approximately five times larger. This specification is slightly different from the one in Table 4 because age² is an additional covariate. I found that when age² was excluded from the covariates, the specification failed the proportional hazards test. As with the results before pension eligibility, increased age increases retirement probability; however, after one becomes eligible for a pension, an increase in age increases retirement probability at a higher rate. Specifically, depending on the specification, the hazard rate increases by 3 to 4 percent more for each additional year. An increase of \$100 per month in salary only decreases retirement by 3 percent (instead of 5). As with the pension pull, nonwhite teachers are less likely to leave than white teachers. Teachers who attended less competitive colleges are 7 percent less likely to leave after pension eligibility, perhaps because they have lower human capital

and fewer alternative job prospects. There are no differences in retirement probability by value-added, meaning pensions are “pushing” out good teachers and bad teachers at the same rates.

Retirement behavior across racial, ethnic, and socioeconomic student characteristics does not reveal any pronounced pattern. Again, middle-school teachers have some different responses to pensions. Although teachers from all levels of schools have a similar likelihood of retiring, middle-school teachers in schools with a one-standard-deviation increase in black students are less likely to leave than teachers in the average school. This may be due to the selection out of these schools prior to pension eligibility. Teachers in middle and high schools with a one-standard-deviation higher proportion of economically disadvantaged students are around 7 to 8 percent more likely to leave.

Summary and Policy Implications

I argue that a Cox proportional hazard model is a suitable, simple, and realistic way to estimate retirement behavior when all employees are subject to a similar salary schedule and pension rules. I find that North Carolina teachers are very responsive to pension incentives: few leave in the 10 years prior to pension eligibility, and many leave upon or after attaining eligibility. As expected, younger and more highly paid teachers are more attached to the workforce throughout the pension pull and push. The current pension pull is strongest for teachers who have high value-added and those who are not white. I provide evidence that teachers with higher levels of general human capital (as measured by having an advanced degree) are more likely to leave during the pension “pull,” and teachers with lower levels of general human capital (as measured by attending less competitive colleges) are least likely to

leave during the pension “push.” Teachers’ retirement behavior is affected by student demographic composition, as teachers are more likely to exit when teaching at a school with a higher proportion of economically disadvantaged students; however, my model cannot show what the counterfactual scenario would be if the pension “pull” were not in place—it could be that the retirement rate in hard-to-staff schools would be even higher than it is now. Overall, pension eligibility is a strong, salient driver of teacher retirement behavior. Given that teachers have different responses to the current defined benefit plan incentives, they will likely have different responses to any proposed changes to defined contribution plans as well.

References

- Berkovec, James, and Steven Stern. 1991. "Job Exit Behavior of Older Men." *Econometrica* 59(1): 189–210.
- Boyd, Donald, Pamela Grossman, Hamilton Lankford, Susanna Loeb, and James Wyckoff. 2007. "Who Leaves? Teacher Attrition and Student Achievement." Working paper. Stanford, CA: Center for Education Policy Analysis.
- Boyd, Donald, Hamilton Lankford, Susanna Loeb, Jonah Rockoff, and James Wyckoff. 2008. "The Narrowing Gap in New York City Teacher Qualifications and Its Implications for Student Achievement in High-Poverty Schools." *Journal of Public Policy Analysis and Management* 27(4): 793–818.
- Boyd, Donald, Hamilton Lankford, Susanna Loeb, and James Wyckoff. 2005. "Explaining the Short Careers of High-Achieving Teachers in Schools with Low-Performing Students." *American Economic Review* 95(2): 166–171.
- Clotfelter, Charles T., Helen F. Ladd, and Jacob L. Vigdor. 2006. "Teacher-Student Matching and the Assessment of Teacher Effectiveness." *Journal of Human Resources* 41(4): 778–820.
- . 2007. "How and Why Do Teacher Credentials Matter for Student Achievement?" CALDER Working Paper No. 2. Washington, DC: Urban Institute, Center for Analysis of Longitudinal Data in Education Research (CALDER).
- Coile, Courtney, and Jonathan Gruber. 2007. "Future Social Security Entitlements and the Retirement Decision." *Review of Economics and Statistics* 89(2): 234–246.
- Costrell, Robert M., and Josh B. McGee. 2010. "Teacher Pension Incentives, Retirement Behavior, and Potential for Reform in Arkansas." *Education Finance and Policy* 5(4): 492–518.
- Costrell, Robert M., and Michael Podgursky. 2009. "Peaks, Cliffs, and Valleys: The Peculiar Incentives in Teacher Retirement Systems and Their Consequences for School Staffing." *Education Finance and Policy* 4(2): 175–211.
- Fitzpatrick, Maria D. 2015. "How Much Are Public School Teachers Willing to Pay for Their Retirement Benefits?" *American Economic Journal: Economic Policy* 7(4): 165–188.
- Fitzpatrick, Maria D., and Michael F. Lovenheim. 2014. "Early Retirement Incentives and Student Achievement." *American Economic Journal: Economic Policy* 6(3): 120–154.
- Friedberg, Leora, and Sarah Turner. 2010. "Labor Market Effects of Pensions and Implications for Teachers." *Education Finance and Policy* 5(4): 463–491.

- . 2011. *Pensions and Public School Teacher Retirement: An Analysis Using National Teacher Data*. Research Dialogue No. 99. New York: TIAA-CREF Institute.
- Furgeson, Joshua, Robert P. Strauss, and William B. Vogt. 2006. “The Effects of Defined Benefit Pension Incentives and Working Conditions on Teacher Retirement Decisions.” *Education Finance and Policy* 1(3): 316–348.
- Goldhaber, Dan, Betheny Gross, and Daniel Player. 2011. “Teacher Career Paths, Teacher Quality, and Persistence in the Classroom: Are Public Schools Keeping Their Best?” *Journal of Public Policy and Management* 30(1): 57–87.
- Hanushek, Eric A., John F. Kain, and Steven G. Rivkin. 2004. “Why Public Schools Lose Teachers.” *Journal of Human Resources* 39(2): 326–354.
- Hanushek, Eric A., and Steven G. Rivkin. 2010. *Using Value-Added Measures of Teacher Quality*. Calder Brief No. 9. Washington, DC: Urban Institute, Center for Analysis of Longitudinal Data in Education Research (CALDER).
- Hausman, Jerry A., and David A. Wise. 1985. “Social Security, Health Status, and Retirement.” In *Pensions, Labor, and Individual Choice*, David A. Wise, ed.. Chicago: University of Chicago Press, pp. 159–192.
- Koedel, Cory, and Julian R. Betts. 2011. “Does Student Sorting Invalidate Value-Added Models of Teacher Effectiveness? An Extended Analysis of the Critique.” *Education Finance and Policy* 6(1): 18–42.
- Koedel, Cory, Michael Podgursky, and Shishan Shi. 2013. “Teacher Pension Systems, the Composition of the Teaching Workforce, and Teacher Quality.” *Journal of Policy Analysis and Management* 32(3): 574–596.
- Koedel, Cory, and P. Brett Xiang. 2017. “Pension Enhancements and the Retention of Public Employees.” *Industrial and Labor Relations Review* 70(2): 519–551.
- Lankford, Hamilton, Susanna Loeb, and James Wyckoff. 2002. “Teacher Sorting and the Plight of Urban Schools: A Descriptive Analysis.” *Educational Evaluation and Policy Analysis* 24(1): 37–62.
- Pew Center on the States. 2010. *The Trillion Dollar Gap: Underfunded State Retirement Systems and the Roads to Reform*. Washington, DC: Pew Center on the States.
- Podgursky, Michael, Ryan Monroe, and Donald Watson. 2004. “The Academic Quality of Public School Teachers: An Analysis of Entry and Exit Behavior.” *Economics of Education Review* 23(5): 507–518.

- Rivkin, Steven G., Eric A. Hanushek, and John F. Kain. 2005. "Teachers, Schools, and Academic Achievement." *Econometrica* 73(2): 417–458.
- Rockoff, Jonah E. 2004. "The Impact of Individual Teachers on Student Achievement: Evidence from Panel Data." *American Economic Review* 94(2): 247–252.
- Rothstein, Jesse. 2009. "Student Sorting and Bias in Value Added Estimation: Selection on Observables and Unobservables." NBER Working Paper No. 14666. Cambridge, MA: National Bureau of Economic Research.
- Rust, John. 1989. "A Dynamic Programming Model of Retirement Behavior." In *The Economics of Aging*, David A. Wise, ed. Chicago: University of Chicago Press, pp. 359–398.
- Stock, James H., and David A. Wise. 1988. "Pensions, the Option Value of Work, and Retirement." NBER Working Paper No. 2686. Cambridge, MA: National Bureau of Economic Research.
- . 1990. "Pensions, the Option Value of Work, and Retirement." *Econometrica* 58(5): 1151–1180.
- Value-Added Research Center. 2010. *NYC Teacher Data Initiative: Technical Report on the NYC Value-Added Model*. Madison, WI: University of Wisconsin, Wisconsin Center for Education Research.
- Xu, Zeyu, Jane Hannaway, and Colin Taylor. 2007. "Making a Difference? The Effects of Teach for America in High School." CALDER Working Paper No. 17. Washington, DC: Urban Institute, Center for Analysis of Longitudinal Data in Education Research (CALDER).

Table 1 Sample Selection

	<i>N</i>	%
Full-time, vested teachers over the age of 40 with 10 years or less until eligible for full pension benefit	51,470	100
Not in sample because:		
Unknown/unreasonable salary or hours worked	765	1.5
Unreasonable/inconsistent values of experience	1,281	2.5
Unreasonable/unknown value of age	6,061	11.8
Unknown sex or race/ethnicity	2,852	5.5
Unknown school characteristics	982	1.9
Unknown college competitiveness	1,103	2.1
Able to retire with less than 30 years of experience	8,627	16.8
In analytic sample:	29,799	57.9

NOTE: All North Carolina public schoolteachers employed between 1994–1995 and 2008–2009 school years. Teachers are vested after five years of service. I have value-added measures for 3,730 teachers in the analytic sample (12.5 percent of the analytic sample).

Table 2 Teacher Value-Added Specification

Dependent variable = standardized (mean 0, std. dev. 1 in grade and year) end-of-grade test score

	Math	Reading
Previous score (standardized by grade and year)	0.740 *** (0.001)	0.695 *** (0.001)
Female	-0.010 *** (0.001)	0.015 *** (0.001)
Black	-0.099 *** (0.002)	-0.131 *** (0.002)
Hispanic	0.014 *** (0.003)	-0.027 *** (0.003)
Other race	0.010 *** (0.003)	-0.023 *** (0.003)
Limited English proficiency status	-0.020 *** (0.005)	-0.063 *** (0.005)
Economically disadvantaged	-0.071 *** (0.001)	-0.091 *** (0.002)
Student variables (switching schools, repeating a grade, age in third grade)	X	X
Year indicators	X	X
Grade 4 indicator	X	X
Student exceptionality status (gifted, speech or language disability, physical disability, emotional disability, mental disability, learning disability, or other disability indicators)	X	X
Class-level variables (membership, lagged achievement, % nonwhite, % female, % LEP, % economically disadvantaged)	X	X
School-level variables (% black, % Hispanic, % economically disadvantaged, school undergoing structural change)	X	X
<i>N</i> (student test scores)	778,734	783,572
<i>R</i> -squared	0.75	0.69

NOTE: Sample is made up of students of all North Carolina public schoolteachers teaching between 10 and 40 students in a given year who took the fourth or fifth grade end-of-grade test between 1999–2000 and 2008–2009 (the years in which all independent variables are available). Value-added is measured at the teacher level. Standard errors shown in parentheses. * significant at the 0.10 level (two-tailed test); ** significant at the 0.05 level (two-tailed test); *** significant at the 0.01 level (two-tailed test). The standard deviation of value-added is 0.22 (math), 0.15 (reading), and 0.17 (average) prior to Bayes shrinkage. After shrinkage, the standard deviations are 0.19 (math), 0.10 (reading), and 0.13 (average). “LEP” = “limited English proficiency.”

Table 3 Descriptive Statistics for Full Sample and Value-Added Sample

		Full sample		Value-added sample	
Year		1995–1996	2008–2009	1995–1996	2008–2009
<i>N</i>		14,359	12,034	1,567	1,736
Experience (in fall)	mean	23.9	25.5	22.9	25.4
	std. dev.	3.8	4.2	3.1	4.2
	min.	19	19	19	19
	max.	34	34	34	34
Age (in fall)	mean	47.5	49.8	46.3	49.6
	std. dev.	4.4	4.6	3.6	4.7
	min.	40	40	40	40
	max.	63	64	60	63
Salary per month/\$100	mean	34.5	50.7	33.9	50.7
	std. dev.	2.5	4.8	2.2	4.6
	min.	21.0	29.8	29.4	43.5
	max.	42.2	66.7	40.9	66.7
% attended college rated as					
More competitive		5.2	8.0	3.8	8.2
Competitive		50.7	51.0	51.9	44.7
Less competitive		44.1	41.0	44.4	47.1
B.A. degree		58.5	57.5	57.9	56.7
M.A. degree and above		41.5	42.5	42.1	43.3
White		79.3	83.3	78.6	79.9
Nonwhite		20.7	16.7	21.4	20.1
Teaches at a school with					
% black students	mean	30.0	26.9	29.1	26.6
	std. dev.	22.3	22.5	23.6	23.1
	min.	0	0	0	0
	max.	99.9	100.0	99.7	98.1
% economically disadvantaged students	mean	30.0	36.8	34.9	39.7
	std. dev.	17.5	19.1	17.7	20.6
	min.	0	0.2	0	0.2
	max.	94.4	96.5	90.2	96.5
Elementary		57.1	56.6	94.1	90.8
Middle		18.8	20.7	5.6	8.7
High		24.1	22.6	0.3	0.5

NOTE: Descriptive statistics in the first and last years for the 29,799 teachers in the full sample and the 3,730 of those for whom I can calculate value-added.

Table 4 Hazard Ratios Pre-Pension Eligibility

	Value-added	All teachers	All teachers with interactions
Age	1.0939 *** (0.0254)	1.0962 *** (0.0069)	1.0963 *** (0.0069)
Salary per month/\$100	0.9277 * (0.0362)	0.9483 *** (0.0096)	0.9486 *** (0.0096)
M.A. degree and above	1.665 ** (0.3307)	1.269 *** (0.0615)	1.2681 *** (0.0615)
Nonwhite	0.7356 * (0.1312)	0.7887 *** (0.0388)	0.7854 *** (0.0388)
% black students in school (standardized)	1.0179 (0.1271)	1.0863 ** (0.0384)	1.0648 (0.0424)
% economically disadvantaged students in school (standardized)	1.0937 (0.1223)	1.086 *** (0.0347)	1.0939 ** (0.0388)
Teacher value-added (standardized)	0.7933 *** (0.0467)		
Attended more competitive college	1.1779 (0.1577)	1.1161 *** (0.0402)	1.1175 *** (0.0402)
Attended less competitive college	1.6715 ** (0.3442)	1.1325 ** (0.0675)	1.1339 ** (0.0675)
Middle school	1.0764 (0.2400)	1.1465 *** (0.0444)	1.1405 *** (0.0452)
High school	0.4911 (0.3300)	1.0207 (0.0400)	1.0148 (0.0423)
Middle school $\times$ % black			1.1134 ** (0.0595)
High school $\times$ % black			1.0021 (0.0518)
Middle school $\times$ % economically disadvantaged			0.9694 (0.0537)
High school $\times$ % economically disadvantaged			0.976 (0.0584)
<i>N</i> of observations	20,402	146,782	146,782
<i>N</i> of subjects	3,482	26,437	26,437
<i>N</i> of failures	413	4,332	4,332
Log pseudolikelihood	-441.2	-12246	-12243
Schoenfeld residual test of proportional hazards assumption			
Chi ²	3.39	14.17	16.12
Degrees of freedom	11	10	14
Prob > chi ²	0.98	0.17	0.31

NOTE: Standard errors shown in parentheses. * significant at the 0.10 level (two-tailed test); ** significant at the 0.05 level (two-tailed test); *** significant at the 0.01 level (two-tailed test). Hazard ratio > 1 implies higher likelihood of retirement; < 1 implies lower likelihood of retirement. Sample is composed of all teachers with between 10 and 2 years to go until pension eligibility.

Table 5 Hazard Ratios Post-Pension Eligibility

	Value-added	All teachers	All teachers with interactions
Age	0.1086 *** (0.0551)	0.1825 *** (0.0223)	0.1821 *** (0.0222)
Age ²	1.0201 *** (0.0047)	1.0153 *** (0.0011)	1.0153 *** (0.0011)
Salary per month/\$100	0.9696 (0.0212)	0.9663 *** (0.0056)	0.9663 *** (0.0056)
M.A. degree and above	1.0031 (0.1205)	0.9708 (0.0302)	0.9722 (0.0302)
Nonwhite	0.9508 (0.0873)	0.8859 *** (0.0230)	0.8871 *** (0.0231)
% black students in school (standardized)	0.9434 (0.0623)	0.9641 * (0.0190)	0.9771 (0.0216)
% economically disadvantaged students in school (standardized)	0.983 (0.0577)	1.0324 * (0.0187)	1.0128 (0.0201)
Teacher value-added (standardized)	0.9989 (0.0331)		
Attended more competitive college	0.9478 (0.0712)	0.9904 (0.0201)	0.99 (0.0201)
Attended less competitive college	1.2376 (0.1891)	0.9264 * (0.0385)	0.9279 * (0.0385)
Middle school	0.8129 (0.1127)	0.9707 (0.0216)	0.9753 (0.0221)
High school	0.4243 (0.3350)	0.974 (0.0202)	0.9831 (0.0208)
Middle school × % black			0.9334 ** (0.0289)
High school × % black			0.9791 (0.0275)
Middle school × % economically disadvantaged			1.0697 ** (0.0333)
High school × % economically disadvantaged			1.0594 * (0.0321)
N of observations	5,588	41,913	41,913
N of subjects	1,883	15,200	15,200
N of failures	1,244	10,847	10,847
Log pseudolikelihood	-1,021.3	-23,942	-23,939
Schoenfeld residual test of proportional hazards assumption			
Chi ²	3.45	11.57	13.47
Degrees of freedom	12	11	15
Prob > chi ²	0.99	0.40	0.57

NOTE: Standard errors shown in parentheses. * significant at the 0.10 level (two-tailed test); ** significant at the 0.05 level (two-tailed test); *** significant at the 0.01 level (two-tailed test). Hazard ratio > 1 implies higher likelihood of retirement; < 1 implies lower likelihood of retirement. Sample is made up of all teachers who have been eligible for pension benefits for up to five years.

Figure 1 Pension Wealth and Change in Pension Wealth by Experience

NOTE: Pension wealth for a hypothetical teacher who starts teaching at age 21 and faces the 2000–2001 North Carolina salary schedule during her entire career. $PensionWealth_{it} = \sum_{s=t}^T \beta^{s-t} \pi_{s|t} annuity_{it}$, the pension wealth of individual i who exits teaching in year t and receives annuity $annuity_{it}$ multiplied by the probability $\pi_{s|t}$ that one is alive in the later period s to receive that payment (conditional on being alive in period t). I assume β to be 0.95. I calculate $\pi_{s|t}$ using life tables by gender and race (white, black, Hispanic, and other) from the National Center for Health Statistics for the year 2006 (Arias 2010a,b).

Figure 2 Pension Wealth for Hypothetical Teachers by Age and Pension Eligibility

NOTE: Pension wealth for hypothetical white females who started teaching at age 21, 25, and 29. This range includes the starting age of approximately 70 percent of all teachers in the sample. I assume they are always paid according to the 2000–2001 salary scale. The magnitude of pension eligibility corresponds to how many years an individual has until (if < 0) or since (if ≥ 0) pension receipt eligibility.

Figure 3 Hazard Rates and Survival Functions by Years of Experience and Pension Eligibility

NOTE: Hazard rate is the marginal proportion of those leaving North Carolina public school teaching. Survival function is the proportion of teachers working in North Carolina public schools at 10 years until eligibility who are still present at a given level of pension eligibility. Graphs above include all individuals in analytic sample and do not control for any covariates.

Figure 4 Estimated Hazard Rate and Survival Function Pre-Pension Eligibility

Panel A. Hazard Rate by Pension Eligibility

Panel B. Survival function by pension eligibility

NOTE: Hazard rate is the marginal proportion of those leaving North Carolina public school teaching. Survival function is the proportion of teachers working in North Carolina public schools at 10 years until eligibility who are still present at a given level of pension eligibility. Graphs above are estimated using the average values for covariates for each level of pension eligibility and coefficients from specifications (1) and (2) in Table 4.

Figure 5 Estimated Hazard Rate and Survival Function Post-Pension Eligibility

Panel A. Hazard Rate by Pension Eligibility

Panel B. Survival Function by Pension Eligibility

NOTE: Hazard rate is the marginal proportion of those leaving North Carolina public school teaching. Survival function is a ratio where the denominator is the number of teachers working when they became eligible to receive their pension and the numerator is the number who are still working at a given level of pension eligibility. Graphs above are estimated using the average values for covariates for each level of pension eligibility and coefficients from specifications (1) and (2) in Table 5.

Figure 6 Estimated Hazard Rate and Survival Function Pre-Pension Eligibility by Teacher Value-Added

Panel A. Hazard Rate by Pension Eligibility

Panel B. Survival Function by Pension Eligibility

NOTE: Hazard rate is the marginal proportion of those leaving North Carolina public school teaching. Survival function is the proportion of teachers working in North Carolina public schools at 10 years until eligibility who are still present at a given level of pension eligibility. Graphs above are estimated using the average values for covariates for each level of pension eligibility and coefficients from specification (1) in Table 4, with value-added set to -1, 0, or 1.