NATIONAL CENTER *for* ANALYSIS *of* LONGITUDINAL DATA *in* EDUCATION RESEARCH

TRACKING EVERY STUDENT'S LEARNING EVERY YEAR

A program of research by the American Institutes for Research with Duke University, Northwestern University, Stanford University, University of Missouri-Columbia, University of Texas at Dallas, and University of Washington

CALDER

AIR

Pension Enhancements and the Retention of Public Employees: Evidence from Teaching

CORY KOEDEL

P. BRETT XIANG

WORKING PAPER 123 • FEBRUARY 2015

Pension Enhancements and the Retention of Public Employees: Evidence from Teaching

Cory Koedel
University of Missouri

P. Brett Xiang
University of Missouri

Contents

Acknowledgements

Koedel is in the department of economics and Truman School of Public Affairs, and Xiang is in the department of economics, at the University of Missouri, Columbia. The authors thank Shawn Ni, Eric Parsons, Michael Podgursky and participants in the Education Policy Seminar Series at the University of Virginia for useful comments, and the Missouri Department of Elementary and Secondary Education for access to data. They gratefully acknowledge research support from the Laura and John Arnold Foundation and CALDER. The views expressed here are those of the authors and should not be attributed to their institutions, data providers, or the funders. Any and all errors are attributable to the authors.

This research was supported by the National Center for the Analysis of Longitudinal Data in Education Research (CALDER) funded through Grant R305C120008 to the American Institutes for Research from the Institute of Education Sciences, U.S. Department of Education. CALDER working papers have not gone through final formal review and should be cited as working papers. They are intended to encourage discussion and suggestions for revision before final publication. The views expressed are those of the authors and should not be attributed to the American Institutes for Research, its trustees, or any of the funders or supporting organizations mentioned herein.

CALDER • American Institutes for Research
1000 Thomas Jefferson Street N.W., Washington, D.C. 20007
202-403-5796 • http://www.caldercenter.org/

Pension Enhancements and the Retention of Public Employees: Evidence from Teaching

Cory Koedel and P. Brett Xiang

CALDER Working Paper No. 123
February 2015

Abstract

We use data from workers in the largest public-sector occupation in the United States – teaching – to examine the effect of pension enhancements on employee retention. Specifically, we study a 1999 enhancement to the benefit formula for public school teachers in St. Louis that resulted in an immediate and dramatic increase in their incentives to remain in covered employment. To identify the effect of the enhancement on teacher retention, we leverage the fact that the strength of the incentive increase varied across the workforce depending on how far teachers were from retirement eligibility when it was enacted. Although we document substantial differences across teachers in how their retention incentives were affected by the enhancement, our results indicate that these differences did not translate into meaningful changes in behavior.

1. Introduction

Defined benefit (DB) pension plans have been in decline in the private sector for decades but are still prevalent in the public sector (Hansen, 2010; Wiatrowski, 2012). A distinguishing feature of public DB plans is that they backload retirement compensation. The degree of backloading was heightened in many states and municipal pension plans in the late 1990s and early 2000s when the benefit formulas in plans across the United States were enhanced (Koedel, Ni and Podgursky, 2014; Munnell, 2012; National Conference of State Legislatures, 1999, 2000, 2001).[1]

An economic rationale for the backloading built into public DB plans, and for the increases in backloading that occurred via the widespread pension enhancements around the turn of the century, is that deferred retirement compensation promotes employee retention (Lazear, 1990; Lazear and Moore, 1988). Particularly among teachers, the potential for the DB pension structure to improve retention is appealing given the well-documented attrition problems in public schools (Boyd et al., 2011; Ingersoll, 2001; Loeb, Darling-Hammond and Luczak, 2009). However, the literature on how workers, and teachers in particular, are affected by their incentives to remain in pension-covered employment has produced mixed results. Studies that examine temporary policies that modify workers' retention incentives suggest fairly large responses (e.g., Fitzpatrick and Lovenheim, 2014; Furgeson, Strauss and Vogt, 2006), while studies of permanent changes suggest much smaller responses (Brown, 2013; Smith and West, 2014). Comparisons between DB plans and alternative plans without backloading, like defined-contribution (DC) plans, find little evidence to suggest that workers' exit decisions are meaningfully affected by their DB retention incentives (Gustman and Steinmeier, 1993; Harris and Adams, 2007).

We contribute to the literature by examining the effect on retention of increasing pension backloading via benefit-formula enhancements. Improving our understanding of how workers respond

[1] There is nothing inherent to the structure of DB pension plans requiring that they backload compensation. However, as a practical matter the vast majority of public DB pension plans in the United States are significantly backloaded.

to changes to pension plan rules, without changing the underlying DB pension structure, is critical to informing contemporary pension policy. Current pension reform debates, and most reforms that have been enacted in recent years to lower the long-term obligations of pension funds, have focused on benefit modifications without changing the DB pension structure (e.g., changes in the rules governing retirement eligibility, replacement rates, cost-of-living adjustments, etc.).[2] Similarly, the sweeping reforms to public pension plans across the United States during the late 1990s and early 2000s – which unlike current reforms typically *improved* pension benefits – also took the form of changes to benefit formulas within the structure of pre-existing plans.[3]

We perform our analysis using data from public school teachers covered by the St. Louis Public School Retirement System. In 1999, the St. Louis plan enacted a generous benefit formula change that resulted in an immediate 60-percent increase in pension wealth for all workers. We estimate that the direct cost to the school district of providing the enhancement for the single cohort of teachers working at the time of its enactment was approximately $166 million (in 2013 dollars). This amounts to an average per-teacher cost of just over $52,000.[4] Although all teachers received an immediate 60-percent increase in pension wealth due to the enhancement, the backloading of retirement compensation in the system is such that the dollar value of the increase varied considerably across teachers. Those who were close to retirement eligibility had the most to gain from remaining in covered employment after the enhancement was enacted. Those far from retirement eligibility had much less to gain.

After documenting large differences across teachers in the change to their retention incentives owing to the enhancement, we estimate difference-in-difference models that compare behavioral responses across groups of teachers within St. Louis to identify the enhancement's retention effects. We

[2] For example, several state teacher plans have introduced new, less-generous pension "tiers" for incoming employees in recent years (e.g., Alabama, California, Connecticut, Illinois, Louisiana).
[3] See National Conference of State Legislatures (1999, 2000, 2001), and for more on teacher plans in particular see Koedel, Ni and Podgursky (2014).
[4] The formula improvement also obligated the school district (the employer) to provide richer pensions for future cohorts of teachers, which are not incorporated into the cost figures reported in the text.

use early-career teachers (more than 20 years away from retirement eligibility) as the baseline "control" group because the enhancement affected their incentives the least. Compared to this baseline group, we show that teachers who were working but already retirement-eligible at the time when the enhancement was approved clearly responded. These teachers delayed retirement during a "gap year" between the approval and enactment of the enhancement, which allowed them to take advantage of the retroactively improved benefit formula upon retirement. Among teachers not yet eligible for retirement, who make up most of the workforce and for whom retention outcomes are more policy relevant, we do not find any evidence that the substantial differences in the degree to which their retention incentives were affected translated into differences in retention behavior.

Conceptually, our study is most closely related to Brown (2013), who also evaluates the labor supply response to a pension enhancement. The two most notable differences between the St. Louis enhancement and the enhancement that Brown studies in California are (1) the St. Louis enhancement was significantly more generous, and (2) the California enhancement altered teachers' retirement-timing incentives (which Brown leverages to identify the labor-supply response), while the enhancement that we study did not change optimal retirement timing but rather the returns to surviving in covered employment until meeting fixed retirement-eligibility rules. Our study also differs from Brown (2013) in that we examine retention effects throughout the workforce whereas she focuses on senior teachers very close to retirement. Brown's focus on senior teachers is useful in that it allows her to cleanly estimate their labor supply elasticity with respect to pension-benefit changes, but from a policy perspective the question of retention effects for the larger workforce is also of interest. Our study complements Brown's work in this way, and our more broadly applicable findings are consistent with the small labor supply response that she estimates for older workers.

In the concluding section we discuss a number of potential explanations for why the St. Louis enhancement failed to generate a meaningful, differentiated response among teachers who experienced very different changes in their retention incentives. Given that the significant cost of the St. Louis enhancement is difficult to justify in terms of workforce retention benefits, we also discuss alternative explanations for why it was enacted in the first place. Finally, we consider the implications of our findings for current pension reform proposals, which as noted above are structurally similar to the enhancement that we study but aim to pare back rather than improve benefits.

2. Background

The St. Louis School District pension plan is a municipal plan and is structured similarly to other subnational public pension plans across the United States. The following formula is used to determine the annual benefit at retirement:

$$B = F * YOS * FAS \tag{1}$$

In (1), *B* represents the annual benefit, *F* is the formula factor, *YOS* indicates years of service in the system, and *FAS* is the teacher's final average salary, calculated as the average of the highest three years of earnings. *F*YOS* is commonly referred to as the "replacement rate." For example, in a system where the formula factor is 0.02, a teacher with 30 years of service will receive an annual pension that replaces 60 percent of her final average salary. Pension benefits are often adjusted for cost-of-living increases for retirees. In St. Louis, cost of living adjustments are *ad hoc* (as opposed to being mandated by statute). St. Louis teachers are also enrolled in Social Security.[5]

The pension enhancement that we study increased the formula factor in the St. Louis plan from 0.0125 to 0.0200. The improved formula factor was implemented retroactively – that is, individuals who retired under the enhanced rules had the higher rate applied to all service years. Thus, the enhancement resulted in an immediate, across-the-board 60 percent increase in pension wealth for all

[5] State and local workers were originally excluded from Social Security, but Congress passed legislation in the early 1950s that permitted states and municipalities to include their employees. The fact that St. Louis teachers are enrolled in Social Security is of limited practical importance for our analysis and all of the incentive changes that we study are driven by a municipal-plan rule change. Social Security benefits are much less lucrative than the benefits that pensioners can earn in most state and municipal plans, as will become clear below for St. Louis teachers.

workers. It was enacted for teachers retiring on or after June 30, 1999. Individuals who retired and began collecting benefits after the 1998-1999 school year received their pensions based on the improved formula; prior retirees collected a less remunerative stream of pension payments based on the original formula.

Figure 1 shows pension-wealth accrual in the St. Louis system for a representative 24-year-old new entrant under the pre- and post-enhancement pension rules. Pension wealth is calculated as the present value of the stream of pension payments. Pension wealth at time *s*, with collection starting at time *j* where $j \geq s$, can be written as:

$$\sum_{t=j}^{T} Y_t * P_{t|s} * d^{t-s} \qquad (2)$$

In (2), Y_t is the annual pension payment in period *t*, $P_{t/s}$ is the probability that the individual is alive in period *t* conditional on being alive in period *s*, and *d* is the discount factor. Details about our pension-wealth calculations are provided in Appendix A.

Figure 1 separately shows St. Louis system wealth accrual, Social Security wealth accrual, and total wealth accrual (the latter combines the two). Similarly to other public DB plans, wealth accrual in the St. Louis plan is heavily backloaded. Wealth accrual in Social Security is relatively flat compared to wealth accrual in the system and does not decline.[6]

The backloading is the result of two features of the St. Louis plan. First, like other public pension plans nationwide (e.g., see Costrell and Podgursky, 2009; National Council on Teacher Quality, 2012), the St. Louis plan offers a generous retirement provision that depends on within-system experience. In particular, teachers in St. Louis can take advantage of the “rule of 85,” which allows for retirement with full benefits when age and experience sum to 85. For example, although the official retirement age in the system is 65, an individual who begins teaching at age 24 and works continuously can retire and

[6] Social Security wealth does not decline because unlike system pension payments, Social Security payments can be collected while working.

begin collecting a full pension immediately at age 54 with 31 years of service (54+31=85). The additional pension payments that can be collected via rule-of-85 are quite valuable. Note that in Figure 1, peak-value pension wealth is achieved when the representative teacher reaches the rule amount. Teachers who do not work long enough to take full advantage of rule-of-85, and therefore forgo pension payments while they wait to become eligible for pension collection, have much lower pension wealth.[7]

The second feature of the St. Louis system that causes backloading – and again, a feature common to public DB pension plans more generally – is that the final average salary (*FAS*) is frozen at the time of exit. It is not adjusted for inflation or life-cycle pay increases. An individual who exits the system mid-career will earn a pension that depends on a deflated *FAS* value relative to an individual who remains in the system until retirement.

As can be seen clearly in Figure 1, the pension-formula enhancement increased the degree of backloading in the pension system. Put differently, it exacerbated the uneven rate of pension-wealth accrual by implementing a fixed percentage increase across an uneven base. Also note that the gap between the wealth-accrual curves in the figure understates the unevenness in pension-wealth gains across the workforce because it does not account for differences in discounting over the career cycle (pension wealth in Figure 1 is discounted to the point of entry for a new teacher). For example, while the newly-entering teacher represented in Figure 1 would not see meaningful gains from the enhancement until far into the future, teachers at or near retirement eligibility at the time when the enhancement was enacted were set to receive their improved benefits with very little discounting. In summary, retention incentives were increased unevenly across the workforce by the enhancement, with the largest increases accruing to the teachers who were closest to benefit eligibility at the time when it was enacted.

[7] Like other public DB plans, the St. Louis plan allows teachers to collect benefits before reaching full retirement eligibility under some conditions and with a collection penalty. The early-collection options in the St. Louis plan are built into the accrual curves in Figure 1 and our calculations more generally – put differently, we allow teachers to collect their pensions under the most lucrative option at each potential exit point.

3. Data and Enhancement Details

3.1. Data

We use a six-year administrative data panel from the Missouri Department of Elementary and Secondary Education (DESE) covering the school years 1994-1995 through 1999-2000 for the empirical analysis. The data panel contains basic demographic information about teachers in St. Louis along with information about salary, age and experience, which we use to construct each teacher's pension wealth profile (again, see Appendix A for details). Descriptive statistics are reported in Table 1.

3.2. The Enhancement Legislation

The benefit-formula enhancement was enacted in June of 1999 and all teachers who filed for retirement after the 1998-1999 school year were eligible for the improved benefit. However, according to the 2009 actuarial report from the pension fund, which provides a legislative history of changes to the plan, the enhancement was approved by the board of education in the fall of 1997. Thus, teachers working during the 1997-1998 school year who were planning to retire at the conclusion of that year had a particularly strong incentive to delay retirement for one additional year. We were unable to find any evidence to document the extent to which teachers knew about the approval of the enhancement prior to its enactment. However, the results that we present below suggest that at the very least, retirement-eligible teachers at the conclusion of the 1997-1998 school year were aware of the benefit-formula improvement that was to come.

Given that there is some uncertainty regarding the extent to which information about the enhancement was available to teachers during the 1997-1998 school year, we construct the models below to compare teachers during three different time periods: (1) prior to the approval of the enhancement (1994-1995, 1995-1996 and 1996-1997), (2) after the approval but before the enactment of the enhancement (1997-1998), and (3) after the enactment of the enhancement (1998-1999 and 1999-2000).

3.3 Unevenness in the Effect of the Enhancement on Teachers' Retention Incentives

Figure 1 provides a graphical illustration of the unevenness with which teachers' retention incentives were strengthened across the workforce but as noted above, it understates differences across workers at different points in the career cycle because it discounts pension wealth to a fixed point in time. In reality, the unevenness is exacerbated by the fact that benefits for younger teachers are discounted further into the future.

Table 2 provides a more accurate depiction of the heterogeneous effects of the enhancement on teachers' retention incentives. To construct the table, we first identify the closest full-retirement option for each teacher in each year of our data panel (i.e., rule-of-85 or age-65). Then we group each teacher-year observation into one of six bins based on distance to full-retirement eligibility assuming continuous work. Bin-1 teachers are those who are already eligible for full retirement at the conclusion of year-*t* (bin-1 teachers may have been eligible for many years, or may be gaining eligibility for the first time after year-*t*). Bin-2 teachers are 1-5 years away from retirement eligibility. Teachers in bins 3, 4, 5 and 6 are 6-10, 11-15, 16-20 and 21+ years away from retirement eligibility, respectively.

Table 2 reports current, peak-value and expected pension wealth with and without the enhancement for teachers in each bin, excluding Social Security. Current pension wealth (CPW) measures the immediate value of the pension. Examining the effect of the enhancement on CPW is informative, but understates the total value of the enhancement because it does not incorporate the enhancement's effect on the option value of continued work (Coile and Gruber, 2007; Stock and Wise, 1990). At the other extreme, peak-value pension wealth (PV) is the value of the pension at the peak of the accrual curve. The enhancement's effect on PV is an overstatement because all workers will not reach and retire at the peak.[8] Following Koedel, Ni and Podgursky (2014), we also calculate expected pension wealth (EPW) for each teacher with and without the policy, with these calculations serving as

[8] For a more detailed discussion of the peak-value concept, see Coile and Gruber (2007).

the foundation for our preferred estimate of the total cost of the enhancement. Expected pension wealth is a weighted summation of pension wealth after each possible year of the career (forward looking), where the weight applied to each pension-wealth value is the conditional probability that the teacher exits the profession after that year. The exit probabilities that we use as weights are determined based on administrative attrition data for teachers with different age-experience profiles in St. Louis (see Appendix A). For ease of interpretation, the EPW numbers reported in Table 2 are based on a simple, static calculation where we hold exit probabilities for teachers fixed at their post-enhancement levels under the old and new rules.

Focusing on the PV and EPW calculations, two patterns in the table merit attention. First, the gaps between the enhancement gain measured in terms of PV and EPW, within bins and holding rules fixed, are much larger in the higher-numbered bins than in the lower-numbered bins. This reflects the fact that older teachers' career paths are more certain, or put differently, that younger teachers are much more at risk of leaving the system before reaching the peak. Second, and more importantly for the present analysis, the PV gains in pension wealth owing to the enhancement vary considerably across bins. The PV gains provide a measure of how much teachers' retention incentives were affected by the enhancement. In addition to the obvious pattern that PV pension wealth increased by substantially more for teachers in bins 1 and 2 (and to a lesser extent bin 3) relative to teachers in bins 4, 5 and 6, it is noteworthy that the differences are even larger when annualized. For example, not only do bin-2 teachers see an increase in their PV pension wealth more than twice as large as that of bin-6 teachers on average, they will also become eligible to collect benefits much sooner. Dividing the average PV gains in each bin by the average number of years until retirement eligibility (in column 1) reveals a more dramatic gap in the effect on retention incentives. The annualized value of the enhancement for the average bin-6 teacher conditional on reaching the peak is roughly $2,100 per year, whereas for the average bin-2 teacher it is approximately $35,000.

Our empirical strategy, which we describe in the next section, compares teachers across bins and over time to identify the effects on behavior of the differential increases in their retention incentives. Table 3 provides descriptive information about the distribution of age and experience across the bins. Although there are clear differences in the expected ways, there is also considerable overlap across bins along these dimensions.

4. Empirical Strategy

4.1. Model

We estimate the following difference-in-difference model to compare teacher responses to the enhancement across bins, specified as a linear probability model:

$$Y_{it} = \beta_0 + X_{it}\beta_1 + BIN_{it}\lambda + 1998_{it}\gamma_1 + POST_{it}\gamma_2 + T_{it}\gamma_3 + (T_{it} * BIN_{it})\gamma_4 + (1998_{it} * BIN_{it})\theta_1 + (POST_{it} * BIN_{it})\theta_2 + \varepsilon_{it} \quad (3)$$

In (3), Y_{it} is an indicator variable equal to one if teacher *i* was retained and zero if teacher *i* exited covered employment at the conclusion of year *t*. Given that some teachers temporarily leave and return later, to ensure that we capture exits accurately we define an exit as occurring whenever a teacher leaves and does not return for five consecutive years (we use data from beyond the frame of the analytic data panel to code exits as necessary). X_{it} is a vector of observable characteristics about the teacher including race, gender, education level and age. We use unique indicators for each age in the model as in Coile and Gruber (2007).[9] BIN_{it} is a vector of bin indicator variables based on the bin classifications established in the previous section (Table 2), with bin-6 being the omitted group. The

[9] The age indicators pick up the same factors captured by the baseline in a Cox proportional hazard model (Coile and Gruber, 2007). We combine indicators for several sparsely populated age values in the data (i.e., for particularly young and old teachers).

coefficient vector λ captures the constant factors associated with the bin assignments that contribute to retention. The variables 1998_{it} and $POST_{it}$ divide the sample by time period, with the "post" indicator being for the years after the enhancement was enacted (1998-1999 and 1999-2000 cohorts) and the "1998" indicator being for the 1997-1998 school year, during which the enhancement was approved but prior to its enactment. T_{it} is a linear time control, which we interact with the bin indicators to allow for bin-specific linear time trends in retention attributable to non-enhancement factors – γ_3 and γ_4 are identified using variation over time within the pre- and post-enhancement periods. The coefficients θ_1 and θ_2 represent the difference-in-difference estimates of the enhancement effects and are of primary interest. Finally, ε_{it} is the error term. We cluster our standard errors at the individual level because our data panel includes repeat observations for individual teachers.[10]

We focus on comparing teachers in bins 1 through 5 to teachers in bin 6. The identifying assumption is that deviations from the linear time trends across bins 1-5, relative to bin-6, that coincide with the discrete approval/enactment of the enhancement can be attributed to the policy change. Within the standard difference-in-difference framework it is typical to think of the effect of the intervention on the control group – bin-6 teachers in our case – as zero, but bin-6 teachers were not exempt from the enhancement. Nonetheless, examining heterogeneity in the effect of the enhancement across bins within St. Louis is of interest given the large differences in the retention-incentive changes for teachers in different bins (per Table 2). Although we are unable to evaluate the effect of the enhancement on bin-6 teachers directly, in Appendix C we evaluate bin-6 teachers

[10] There is no time dimension to our clustering structure so our standard errors will not be artificially deflated by the serial correlation issue raised in Bertrand, Duflo and Mullainathan (2004). In an analysis omitted for brevity, we also estimated models clustered at the school level to allow for peer effects in retirement behavior (Brown and Laschever, 2012). The higher level of clustering increases our standard errors by 15-20 percent and thus further weakens our results, which per below already show no evidence of enhancement effects on retention.

indirectly, using a cost-benefit framework to estimate the size of the effect on bin-6 teachers that would be required for the enhancement to pass a cost-benefit test.[11]

4.2. Retention Trends in St. Louis

Table 4 shows annual retention rates for St. Louis teachers by bin in each year of our data panel. Consistent with the literature on teacher attrition (Boyd et al., 2011; Ingersoll, 2001; Loeb, Darling-Hammond and Luczak, 2009), the table shows that young and inexperienced teachers have the lowest retention rates. In contrast, retention rates are much higher for teachers within 10 years of retirement eligibility – averaged across years, the annual retention rate for teachers in bins 2 and 3 exceeds 96 percent. The table also shows that retention rates were declining for teachers in all bins over the course of our data panel, and that the declining trend clearly preceded the enactment of the enhancement in 1999. The declining retention trend may reflect a number of factors, ranging from worsening working conditions in St. Louis public schools to the availability of more and better non-teaching options brought on by a booming economy during the second half of the 1990s. A concern related to our identification strategy is that early-career teachers' retention outcomes may be more responsive to macroeconomic factors, and/or worsening working conditions in St. Louis, because they are less occupationally attached than their more senior counterparts (e.g., see Kambourov and Manovskii, 2008).

To explore this issue further we estimate the following model based on Fitzpatrick and Lovenheim (2014):

$$Y_{it} = \pi_0 + X_{it}\pi_1 + BIN_{it}\pi_2 + G_{it}\pi_3 + (G_{it} * BIN_{it})\pi_4 + u_{it} \qquad (4)$$

[11] In an omitted analysis, we also constructed a model analogous to the model in equation (3) that compares St. Louis teachers to teachers outside of St. Louis who are covered by a different pension plan. However, the differential retention trends for teachers outside of St. Louis, both for novices and more senior teachers, are quite large and the alternative model does not fit the data well. We interpret this as evidence that outside teachers do not make for a good comparison group for St. Louis teachers. That said, the model that compares St. Louis teachers to outside teachers does not lead to different substantive conclusions than what we show below (which are essentially null results), although our standard errors for the coefficients of interest from the alternative model are significantly larger, limiting inference.

In equation (4), G_{it} is a vector of year indicator variables. It replaces the timespan variables (*1998*, *POST*) and linear time trend controls from equation (3). The other variables are specified as above.

Based on the output from equation (4), in Figure 2 we construct time trends in retention rates for teachers in all six bins relative to bin-6 teachers in 1995 (the first year of our data panel). We note two important aspects of the trends shown in the figure. First, previewing our main findings, there is no visible evidence of a meaningful differential retention response to the 1999 enhancement for teachers outside of bin-1, whose response is during the "gap year" between approval and enactment. Second, consistent with the above discussion and of direct relevance for the modeling, a larger decline in retention rates for bin-6 teachers relative to other teachers is observed before the enhancement was enacted. We account for the trend differences illustrated in Figure 2 with the linear time trend controls in equation (3), which again are identified using variation within the pre- and post-policy time periods (i.e., holding policy regime fixed). We also show results from models without the time-trend controls. The results from these restricted models will overstate the effects of the pension enhancement on differential retention in St. Louis by failing to account for the pre-policy divergence in retention rates across bins.[12]

5. Results

Table 5 presents estimates from our primary model in equation (3). The table shows the difference-in-difference coefficients from several variants of the model, with the estimates in the fourth column coming from the full specification with the linear time trend controls. Coefficients for the other control variables not shown in the table are reported in Appendix B.

[12] A related issue is that these factors may have influenced the composition of bin-6 teachers over time. For example, if non-teaching opportunities were improving during the second half of the 1990s then teacher quality may have been declining among new entrants into St. Louis over the course of our data panel. Unfortunately our data are not sufficient to directly investigate this issue. However, if the quality of new entrants was indeed declining over time during our data panel, available evidence suggests that the effect on retention would be modest and that if anything, the compositional change would lead to further overstatement of the enhancement effect (Goldhaber, Gross and Player, 2011; Krieg, 2006; West and Chingos, 2009).

We begin by discussing our findings for retirement-ineligible teachers (bins 2, 3, 4 and 5; relative to bin-6 teachers). Focusing on our full specification in column (4) and the estimates from the post-enhancement period, we find no evidence to suggest that teachers in bins 2-5 were differentially affected by the pension enhancement relative to teachers in bin-6. This is a notable result in light of the substantial differences in how teachers' retention incentives were affected across bins (per Table 2).

A caveat to this interpretation is that there are non-negligible costs in terms of statistical power associated with identifying the enhancement effects conditional on the linear time trends, as can be seen by the increase in the size of our standard errors moving from Model 3 to Model 4. An alternative is to focus on Model 3, which does not include the linear time trend controls but produces estimates that are more precise. As illustrated by Figure 2, failing to account for the linear time trends will cause positive bias in our estimates of the relative effects of the enhancement on retention. Consistent with this expectation, the estimates from Model 3 for teachers in all bins relative to bin-6 are more positive, and some are statistically significant. One could interpret the findings from Model 3 as (very) upper-bound estimates of the enhancement's differential effects on retention. In results omitted for brevity, we show that even using these upper-bound estimates, and under reasonable assumptions about the baseline effect of the enhancement on bin-6 teachers, the enhancement fails a cost-benefit test.[13]

Next we turn to bin-1, retirement-eligible teachers. The estimate in row (1) and column (4) of Table 5 shows that retention for these teachers spiked by 7.23 percentage points during the 1997-1998 school year. This is as expected if retirement-eligible teachers in 1997-1998 were aware of the pending enhancement – the increase in pension benefits associated with waiting for the enhancement to take effect, even for teachers on the downward sloping portion of the pension-accrual curve, was dramatic (per Table 2).

[13] We use the general cost-benefit framework described in Appendix C for these calculations.

Although the behavior of retirement-eligible teachers in 1997-1998 indicates that at least some teachers were aware of the enhancement legislation during that year, we were unable to uncover evidence regarding the mechanism for information transmission. One possibility is that teachers who submitted paperwork to retire at the conclusion of the 1997-1998 school year were informed at that time of the value of delaying retirement. Given our uncertainty about how bin-1 teachers knew about the pending enhancement during the 1997-1998 school year, we refrain from drawing strong inference from our estimates for other teachers at the conclusion of that year, although as a practical matter this interpretation issue is of limited consequence given our results in Table 5 (or lack thereof).

As a final note on our findings, we return to the issue that our control group of novice teachers in St. Louis is not entirely untreated. Large retention effects on bin-6 teachers seem unlikely based on outside evidence (French and Jones, 2012; Smith and West, 2014), but we are unable to examine bin-6 teachers directly. While our results are still informative about the relative effects of the enhancement without knowing the baseline effect on novices, information about the effect on novices would help us to evaluate the enhancement policy holistically. As noted above, in Appendix C we provide some indirect insight on the issue of novices by performing a cost-benefit analysis to determine the size of the effect on bin-6 teachers that would be required in order for the enhancement to pass a cost-benefit test. Our findings indicate that the enhancement's effect on retention for bin-6 teachers would need to be implausibly large for it to be a break-even policy.

6. Discussion and Conclusion

We examine the effect on teacher retention of the 1999 benefit-formula enhancement in the St. Louis Public School Retirement System. The enhancement was implemented in such a way that teachers experienced very different changes to their retention incentives. However, we find that these differences did not translate into differences in retention behavior, with the exception of retirement-eligible teachers who strategically delayed retirement for one year to take advantage of the retroactively-improved benefit formula. This is a notable result in light of the substantial cost associated with providing the enhancement, which we estimate to be approximately $166 million for the single cohort of teachers working in the school district at the time when it was enacted.

One explanation for the limited differential effect of the enhancement on teacher retention is that teachers do not value their pension benefits, at least at the margin, at nearly the cost of providing them. This could be due to some combination of teachers being oversaturated with retirement compensation (Fitzpatrick, forthcoming) and their lack of knowledge about the full value of their pensions (Brown et al., 2013; Gustman and Steinmeier, 2005). Another possibility is that income effects may have worked to offset some potential retention benefits of the enhancement, although this seems less likely given that Brown (2013) estimates a small labor response to pension-benefit changes in her study despite taking an approach designed to minimize contamination from income effects. A third factor potentially contributing to our findings is that senior teachers, whose incentives were most affected by the pension enhancement, already have very high retention rates and thus a limited scope for response. It may be that the St. Louis enhancement, which like enhancements to other state and municipal plans was structured to strengthen retention incentives for senior workers the most, was poorly designed if an objective was to improve employee retention.

The enhancements in St. Louis and other subnational pension plans may also have been motivated by other factors. Koedel, Ni and Podgursky (2014) argue that pension enhancements during this time period represent a form of rent capture by senior workers, facilitated by an extended run of above-average stock market returns in the late 1990s and early 2000s that temporarily led to favorable pension-fund balance sheets. They draw a parallel between retroactive benefit-formula enhancements (rent capture by senior workers) and skipped pension payments by government agencies on behalf of employees (rent capture by governments), with both activities being common at the time. Glaeser and

Ponzetto (2014) take a different approach but their study is also consistent with pension enhancements being viewed as a form of rent capture. In short, they argue that because pensioners understand the value of their benefits better than other taxpayers, pension plans serve as a medium whereby politicians can transfer resources to pensioners without sacrificing votes. A consequence in their model is that public-sector compensation is disproportionately and inefficiently delivered in the form of backloaded pension payments.[14]

Although our data are not sufficient to pin down the precise mechanism by which the enhancement failed to meaningfully affect teacher retention, our finding that it did not have a large retention effect is important for several reasons. First, although the last period of widespread pension enhancements in subnational DB pension plans occurred 15 years ago, and new enhancements do not appear to be on the immediate horizon, there may come a time in the not-so-distant future when economic expansion again leads to calls to enhance public-sector pensions. Our study will be useful for informing public policy at that time. Second, to the extent that there is symmetry to our findings, we would also expect teachers' behavioral responses to pension-benefit reductions to be modest. This has implications for current pension reform debates given that it is now apparent that the costs associated with maintaining public plans are larger than commonly reported by actuaries (Biggs, 2011; Munnell, 2012; Novy-Marx and Rauh, 2009, 2011, 2014), which has put pressure on state and municipal governments across the United States to lessen benefits in an effort to reduce long-term obligations.[15] Third, as noted above, our findings are consistent with Fitzpatrick (forthcoming), who shows that

[14] Indirect evidence consistent with this result comes from Chingos and West (forthcoming) and Goldhaber and Grout (forthcoming). In different contexts, both studies show that a substantial fraction of teachers are willing to transfer retirement compensation from the backloaded defined-benefit structure to the more mobile defined-contribution structure. DeArmond and Goldhaber (2010) provide survey evidence that is also consistent with teachers preferring that marginal retirement compensation be delivered in the form of a more mobile benefit.

[15] Examples of recent reforms that have already been enacted include less-lucrative pension-plan "tiers" for newly-entering public-sector workers in some states and municipalities (e.g., Alabama, California, Connecticut, Illinois, Louisiana).

teachers value their pension benefits at much less than it costs to provide them, and more generally, with a number of studies suggesting that retention decisions for teachers and other workers are more strongly influenced by other factors (Boyd et al., 2005; Kersaint et al., 2007; Loeb, Darling-Hammond and Luczak, 2009; Manoli and Weber, 2011). This opens up the possibility for Pareto improving policies that pare back pension benefits in state and municipal plans. Reductions in benefits could be offset for public workers by less-costly increases in wages, and the residual could be used for other spending priorities and/or even higher salaries.

References

Bertrand, M., Duflo, E., & Mullainathan, S. (2004). How much should we trust differences-In-differences estimates? *Quarterly Journal of Economics, 119*(1), 249-275.

Biggs, A. (2011). An options pricing method for calculating the market price of public sector pension liabilities. *Public Budgeting and Finance, 31*(3), 94-118.

Boyd, D., Grossman, P., Ing, M., Lankford, H., Loeb, S., & Wyckoff, J. (2011). The influence of school administrators on teacher retention decisions. *American Educational Research Journal*, *48*(2), 303-333.

Boyd, D., Lankford, H., Loeb, S., & Wyckoff, J. (2005). Explaining the short careers of high-achieving teachers in schools with low performing students. *American Economic Review, 95*(2), 166-171.

Brown, K. (2013). The link between pensions and retirement timing: Lessons from California teachers. *Journal of Public Economics, 98*(1), 1-14.

Brown, K. M., & Laschever, R. A. (2012). When they're sixty-four: Peer effects and the timing of retirement. *American Economic Journal: Applied Economics*, *4*(3), 90-115.

Brown, J.R., Arie, K., Luttmer, E. F. P., & Mitchell, O. S. (2013). Do consumers know how to value annuities? Complexity as a barrier to annuitization. NBER Working Paper No. 19168.

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, 104*(9), 2633-79.

Chingos, M.W., & West, M. R. (Forthcoming). Which teachers choose a defined contribution pension plan? Evidence from the Florida retirement system. *Education Finance and Policy*.

Clotfelter, C. T., Ladd, H. F., & Vidgor, J.L. (2006). Teacher-student matching and the assessment of teacher effectiveness. *Journal of Human Resources, 41*(4), 778-820.

Coile, C., & Gruber, J. (2007). Future social security entitlements and the retirement decision. *The Review of Economics and Statistics, 89*(2), 234-246.

Costrell, R. M., & Podgursky, M. J. (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.

DeArmond, M., & Goldhaber, D. (2010). Scrambling the nest egg: How well do teachers understand their pensions and what do they think about alternative pension structures? *Education Finance and Policy, 5*(4), 558-586.

Fitzpatrick, M. (Forthcoming). How much do public school teachers value their pension benefits? *American Economic Journal: Economic Policy*.

Fitzpatrick, M., & Lovenheim, M. (2014). Early retirement incentives and student achievement. *American Economic Journal: Economic Policy, 6*(3), 120-154.

French, E., & Jones, J. (2012). Public pensions and labor supply over the life cycle. *International Tax and Public Finance, 19*(2), 268-287.

Furgeson, J., Strauss, R. P., & Vogt, W. B. (2006). The effects of defined benefit pension incentives and working conditions on teacher retirement decisions. *Education Finance and Policy, 1*(3), 316-348.

Glaeser, E.L., & Ponzetto, G.A.M. (2014). Shrouded costs of government: The political economy of state and local public pensions. *Journal of Public Economics, 116*, 89-105.

Goldhaber, D., Gross, B., & Player, D. (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.

Goldhaber, D., & Grout, C. (Forthcoming). Which plan to choose? The determinants of pension system choice for public school teachers. *Journal of Pension Economics and Finance*.

Gustman, A.L., & T. L. Steinmeier. (2005). Imperfect knowledge of social security and pensions. *Industrial Relations: A Journal of Economy and Society, 44*(2), 373-397.

Gustman, A. L., & T. L. Steinmeier. (1993). Pension portability and labor mobility: Evidence from the survey of income and program participation. *Journal of Public Economics, 50*(3), 299-323.

Hansen, M., & Janet, S. (2010). An introduction to teacher retirement benefits. *Education Finance and Policy, 5*(2), 402-437.

Hanushek, E.A. (2011). The economic value of higher teacher quality. *Economics of Education Review 30*(3), 266-479.

Hanushek, E. A., & Rivkin, S. G. (2010). Generalizations about using value-added measures of teacher quality. *American Economic Review, 100*(2), 267-271.

Harris, D.N., & Adams, S.J. (2007). Understanding the level and causes of teacher turnover: A comparison with other professions. *Economics of Education Review, 26*(4), 325-337.

Ingersoll, R.M. (2001). Teacher turnover and teacher shortage: An organizational analysis. *American Educational Research Journal, 38*(3), 499-534.

Kambourov, G., & Manovskii, I. (2008). Rising occupational and industry mobility in the United States: 1968-97. *International Economic Review, 49*(1), 41-79.

Kane, T.J., Rockoff, J.E., & Staiger, D.O. (2008). What does certification tell us about teacher effectiveness? Evidence from New York City. *Economics of Education Review, 27*(6), 615-631.

Kersaint, G., Lewis, J., Potter, R., & Meisels, G. (2007). Why teachers leave: Factors that influence retention and resignation. *Teaching and Teacher Education, 23*(6), 775-794.

Koedel, C., Ni, S., & Podgursky, M. (2014). Who benefits from pension enhancements? *Education Finance and Policy, 9*(2), 165-192.

Krieg, J.M. (2006). Teacher quality and attrition. *Economics of Education Review, 25*(1), 13-27.

Lazear, E.P. (1990). Pensions and deferred benefits as strategic compensation. *Industrial Relations*: *A Journal of Economy and Society, 29*(2), 263–281.

Lazear, E. P., & Moore, R. L. (1988). Pensions and turnover. In Bodie, Z., Shoven, J.B., & Wise, D.A (Eds.), *Pensions in the U.S. Economy* (pp. 163-188). Chicago, IL: University of Chicago Press.

Loeb, S., Darling-Hammond, L. & Luczak, J. (2009). How teaching conditions predict teacher turnover in California schools. *Peabody Journal of Education, 80*(3), 44-70.

Manoli, D.S., & Weber, A. (2011). Nonparametric evidence on the effects of financial incentives on retirement decisions. NBER Working Paper No. 17320.

American Psychiatric Association. (2000). *Practice guidelines for the treatment of patients with eating disorders* (2nd ed.). Washington, DC: Author.

Manhattan Institute (2013). *Civic report: Better pay, fairer pensions: Reforming teacher compensation*. New York City, NY: McGee, J. & Winters, M.A.

The Brookings Institution (2012). *State and local pensions: What now?* Washington, DC: Munnell, A.H.

National Conference of State Legislators. (1999). *Pensions and retirement plan issues in 1999 state legislatures.* Washington, DC.

National Conference of State Legislators. (2000). *Pensions and retirement plan issues in 2000 state legislatures.* Washington, DC.

National Conference of State Legislators. (2001). *Pensions and retirement plan issues in 2001 state legislatures.* Washington, DC.

National Council on Teacher Quality. (2012). *No one benefits: How teacher pension systems are failing both teachers and taxpayers*. Washington, DC.

Novy-Marx, R., & Rauh, J. (2014). Revenue demands of public employee pension promises. *American Economic Journal: Economic Policy, 6*(1), 193-229.

Novy-Marx, R., & Rauh, J. (2011). Public pension promises: How big are they and what are they worth? *Journal of Finance, 66*(4), 1207-1245.

Novy-Marx, R. & Rauh, J. (2009). The liabilities and risks of state-sponsored pension plans. *Journal of Economic Perspectives, 23*(4), 191-210.

Sass, Tim R., Hannaway, J., Xu, Z., Figlio, D.N., & Feng, L. (2012). Value added of teachers in high-poverty schools and lower poverty schools. *Journal of Urban Economics, 72*, 104-122.

Smith, J. S., & West, J.E. (2014). *Reduced retirement benefits: Should I stay or should I go?* Unpublished manuscript. Baylor University.

National Center for Education Statistics, Institute of Education Sciences, U.S. Department of Education (2012). *Digest of education statistics 2011* (NCES 2012-001). Washington, DC: Snyder, T. D., & Dillow, S. A.

Stock, J. H., & Wise, D.A. (1990). Pensions, the option value of work, and retirement. *Econometrica, 58*(5), 1151-1180.

West, M.R., & Chingos, M.M. (2009). Teacher effectiveness, mobility and attrition in Florida. In Matthew G. Springer (Ed.), *Performance incentives: Their growing impact on American K-12 Education* (pp. 251-271). Washington, DC: Brookings Institution Press.

Wiatrowski, W. J. (2012). The last private industry pension plans: A visual essay. *Monthly Lab. Rev.*, *135*, 3.

Figures

Figure 1. Pension Wealth Accrual for a Representative 24-year-old Entrant into Teaching in St. Louis, Discounted to the Point of Entry, Using the Original (Left Panel) and Enhanced (Right Panel) Pension Benefit Formulas.

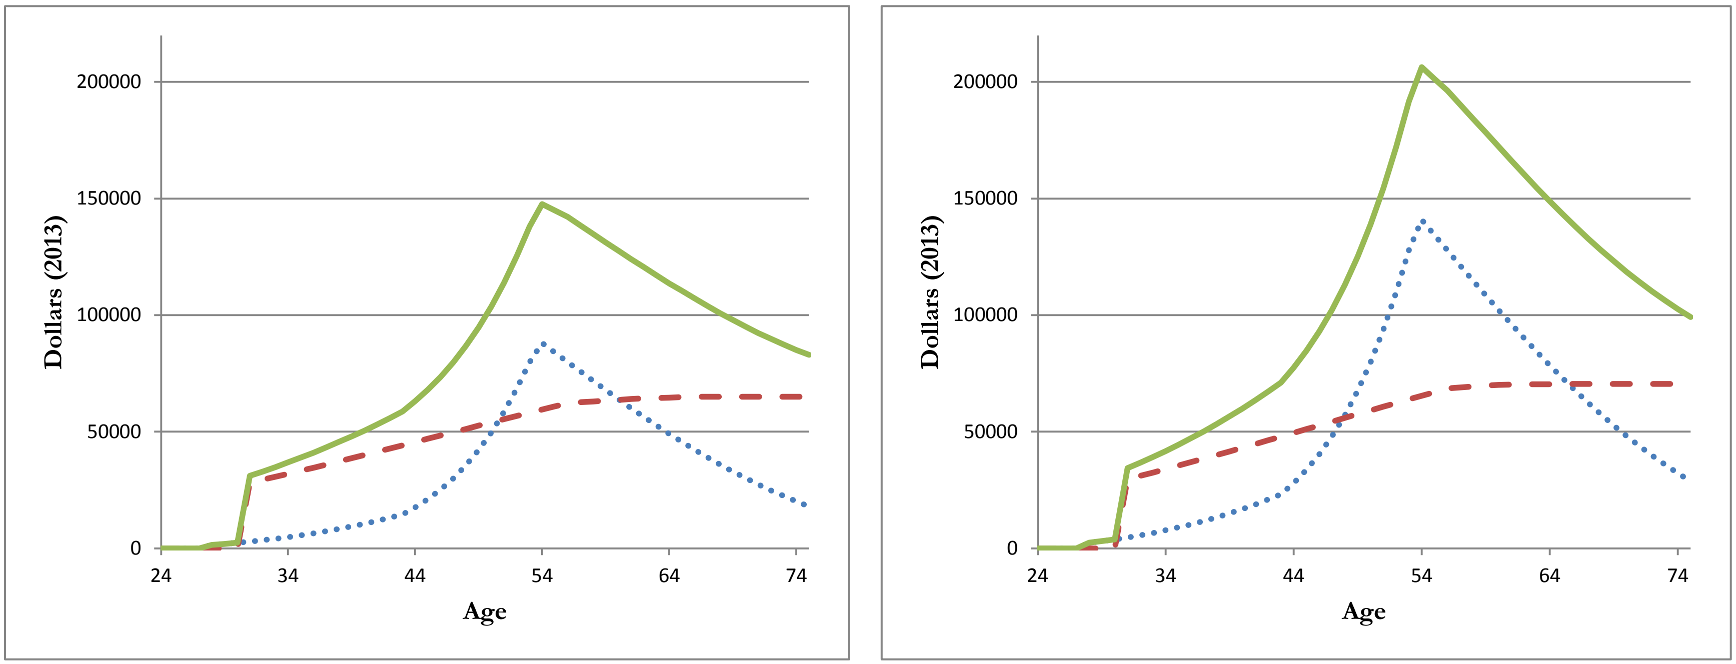

Notes: The left panel of the figure shows pre-enhancement wealth accrual and the right panel shows post-enhancement accrual. Both panels compute pension wealth holding all else equal (e.g., discount rate, salary growth, etc.) The dashed red line shows Social Security wealth accrual, which does not change due to the enhancement (we assume the teacher began contributing to Social Security at age-22 prior to entry into teaching, although this assumption is of no practical consequence). The dotted blue line shows system wealth accrual. The solid green line shows total wealth accrual, combining system and Social Security pension wealth.

Figure 2. Retention Trends from 1994-1995 through 1999-2000, by Bin, Relative to Bin-6 Teachers in 1994-1995.

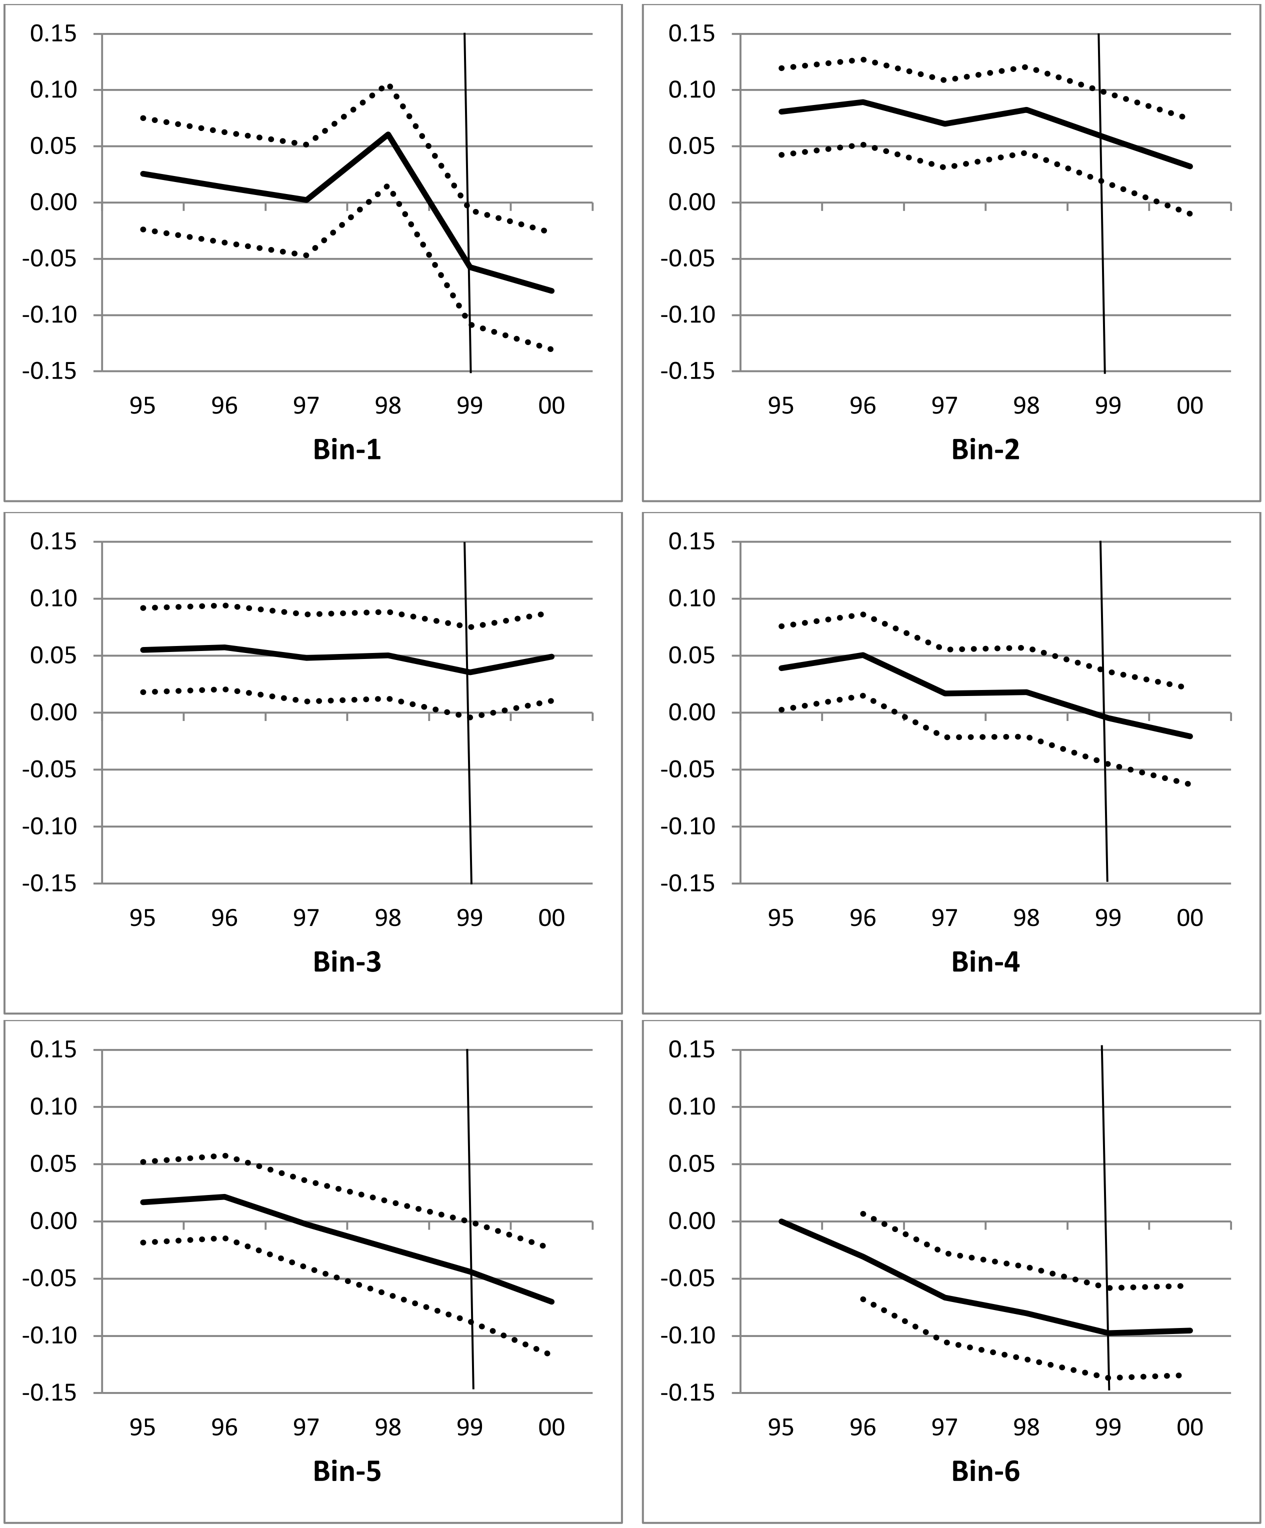

Notes: Each graph shows the retention trend relative to the omitted group (bin-6 teachers in 1995), with the dotted lines showing the 95-percent confidence interval. Note that because bin-6 teachers in 1995 are the comparison group, their retention rate is normalized to zero and presented without a confidence interval.

Tables

Table 1. Descriptive Statistics for Analytic Sample of St. Louis Teachers.

Teacher Characteristics	
Teacher-Year Observations	18828
Unique Teachers	4610
Average Age	47.7
Average Experience	15.5
Share with Master's or PhD	0.468
Share Female	0.783
Share White	0.367
Share African American	0.623
Share Asian	0.004
Share American Indian	0.003
Share Other	0.003
Distance to Full Retirement Eligibility	
Share Retirement Eligible (Bin 1)	0.153
Share Retirement Eligible in 1-5 years (Bin 2)	0.173
Share Retirement Eligible in 6-10 years (Bin 3)	0.167
Share Retirement Eligible in 11-15 years (Bin 4)	0.175
Share Retirement Eligible in 16-20 years (Bin 5)	0.156
Share Retirement Eligible in 21+ years (Bin 6)	0.177
School Characteristics	
Share of Students Eligible for Free/Reduced Price Lunch	0.848
Share of Students Who Are Disadvantaged Minority	0.829

Notes: All shares reported in the table are simple averages across teacher-year observations (teacher-years are the units of analysis).

Table 2. Pension Wealth Under New and Old Rules Using Various Measures, for Teachers by Distance from Full Retirement Eligibility.

	Avg. Years Until Retirement Eligible	Current Pension Wealth			Peak-Value Pension Wealth			Expected Pension Wealth		
		Old Rules	New Rules	Difference	Old Rules	New Rules	Difference	Old Rules	New Rules	Difference
Retirement Eligible (Bin 1)	0	$209,013	334,421	125,408	209,219	334,774	125,555	174,166	278,666	104,500
Eligible in 1-5 Years (Bin 2)	3.1	128,854	206,166	77,312	183,084	292,989	109,905	148,650	237,840	89,190
Eligible in 6-10 Years (Bin 3)	8.0	43,887	70,219	26,332	137,722	220,532	82,810	105,394	168,631	63,236
Eligible in 11-15 Years (Bin 4)	13.1	14,979	23,966	8,987	99,324	159,540	60,216	62,090	99,344	37,254
Eligible in 16-20 Years (Bin 5)	17.9	5,366	8,586	3,220	86,814	139,873	53,059	35,485	56,775	21,291
Eligible in 21+ Years (Bin 6)	25.2	678	1,084	407	89,111	143,026	53,916	13,589	21,742	8,153

Notes: Hypothetical pension wealth under the new and old rules is computed for all teachers in all years, in 2013 dollars. The value of 0 for "years until retirement eligibility" for bin-1 teachers indicates that these teachers can retire at the conclusion of the current year and do not need to remain in the system for any additional years to be eligible for full retirement. Expected pension wealth in the table is based on a simplified calculation holding exit probabilities fixed at their post-enhancement levels for all teachers. Also note that expected pension wealth for bin-1 teachers is lower than either current or peak-value wealth because the expected-pension-wealth calculations include positive probabilities of continued work past the peak, as is observed in the real data.

Table 3. Descriptive Statistics for Teachers by Distance from Full Retirement Eligibility.

	Age 25th, 50th, 75th percentiles			Experience 25th, 50th, 75th percentiles		
Retirement Eligible (Bin 1)	57	60	64	28	32	35
Eligible in 1-5 Years (Bin 2)	51	53	58	22	26	28
Eligible in 6-10 Years (Bin 3)	47	49	54	14	20	23
Eligible in 11-15 Years (Bin 4)	44	47	51	8	12	16
Eligible in 16-20 Years (Bin 5)	40	43	45	4	7	10
Eligible in 21+ Years (Bin 6)	28	32	36	1	2	4

Table 4. Annual Teacher Retention Rates in St. Louis, by Bin and Year.

	1995	1996	1997	1998	1999	2000
Retirement Eligible (Bin 1)	85.6	85.1	84.5	90.6	78.6	76.3
Eligible in 1-5 Years (Bin 2)	97.5	98.1	96.4	97.6	95.3	92.1
Eligible in 6-10 Years (Bin 3)	97.8	97.9	96.8	97.0	95.3	96.6
Eligible in 11-15 Years (Bin 4)	96.7	97.9	94.5	94.5	92.0	90.4
Eligible in 16-20 Years (Bin 5)	95.1	95.3	93.0	90.7	88.6	86.1
Eligible in 21+ Years (Bin 6)	91.8	88.6	84.8	83.2	81.0	82.1

Notes: The columns indicate school years by the spring year (e.g., 1996 indicates the 1995-1996 school year).

Table 5. The Effects of the Pension Enhancement on Teacher Retention.

	Model 1	Model 2	Model 3	Model 4
Bin-1*1998 (retirement eligible)	0.1058	0.1035	0.0931	0.0723
	(0.0244)**	(0.0243)**	(0.0238)**	(0.0372)*
Bin-2*1998	0.0531	0.0521	0.0485	0.0192
	(0.0197)**	(0.0196)**	(0.0195)**	(0.0291)
Bin-3*1998	0.0447	0.0438	0.0428	-0.0045
	(0.0201)**	(0.0200)**	(0.0200)**	(0.0293)
Bin-4*1998	0.0315	0.0298	0.0286	0.0049
	(0.0211)	(0.0210)	(0.0210)	(0.0311)
Bin-5*1998 (16-20 years from eligibility)	0.0119	0.0110	0.0104	-0.0133
	(0.0232)	(0.0231)	(0.0231)	(0.0333)
Bin-1*POST (retirement eligible)	-0.0097	-0.0084	-0.0192	-0.0569
	(0.0210)	(0.0209)	(0.0205)	(0.0537)
Bin-2*POST	0.0304	0.0299	0.0267	-0.0259
	(0.0159)*	(0.0158)*	(0.0159)*	(0.0406)
Bin-3*POST	0.0508	0.0513	0.0506	-0.0335
	(0.0154)**	(0.0154)**	(0.0155)**	(0.0405)
Bin-4*POST	0.0151	0.0170	0.0142	-0.0288
	(0.0166)	(0.0165)	(0.0166)	(0.0429)
Bin-5*POST (16-20 years from eligibility)	-0.0052	-0.0043	-0.0072	-0.0507
	(0.0188)	(0.0187)	(0.0188)	(0.0454)
Teacher Characteristics (Race, Gender, Education Level)		X	X	X
Age Indicators			X	X
Linear Time Trends				X
R-Squared	0.0464	0.0533	0.0719	0.0729
N	18828	18825	18825	18825

*Indicates statistical significance at the 10 percent level
**Indicates statistical significance at the 5 percent level
Notes: The omitted group is bin-6 teachers, who are more than 20 years away from full retirement eligibility. Standard errors are clustered at the teacher level. Three observations are dropped in columns (2), (3) and (4) because teacher characteristics are missing. Coefficients for the other control variables are reported in Appendix Table B.1 for Model 4.

Appendix A: Pension-Wealth Calculation Details

A.1. Basic Calculations

We use the following information from the administrative data file for each teacher to calculate pension wealth: (1) age, (2) system experience and (3) earnings. We determine teacher's survival probabilities over the life cycle using the Cohort Life Tables provided by the Social Security Administration (by gender and birth year). We project out future wages based on current earnings and a wage-growth function that depends on teaching experience. The parameters for the growth function come from a regression based on a 16-year data panel from Missouri where we regress teacher wages on a cubic function of experience. The function captures real wage growth, and wages are also adjusted for inflation. The representative teacher in Figure 1 starts with the base wage for a typical 24-year-old entering teacher in St. Louis, and the growth function adjusts the wage profile moving forward so that *FAS* can be calculated after each possible exit date. For the real-data calculations, teachers' reported wages in each year are used to project wages forward (and backward when necessary).

Our calculations require that we specify a real discount rate. We use a real rate of four percent, which allows for a positive real interest rate and a time preference in earnings.[16] For each teacher, after each year of work, we identify the optimal collection age assuming that the teacher exits after that year, then calculate the present discounted value of the stream of pension payments over the life cycle per equation (2) in the text.

A.2 Expected Pension Wealth

Expected pension wealth (EPW) for each teacher in each year of the data panel is calculated as a weighted average of pension wealth accrued after each potential year of work from the current year forward. The baseline weights are age-and-experience conditional exit probabilities estimated using

[16] Our choice of a four-percent real discount rate falls somewhere in between what others have used in the literature. For example, Coile and Gruber (2007) use 6 percent and Costrell and Podgursky (2009) use 2.5 percent.

data from St. Louis teachers during the last two years of the data panel (post-enhancement). They are estimated from the regression of a binary indicator for retention on interactions of age-and-experience groups. The age and experience groups are reported in the row and column headers in Table A.1, which also shows the estimated conditional retention probabilities. All teachers are assumed to exit with certainty at age-75 regardless of experience. Our approach to calculating EPW is based on Koedel, Ni and Podgursky (2014).[17]

[17] With more data, interactions for each unique age-and-experience combination could be used as in Koedel, Ni and Podgursky (2014), but given the size of the St. Louis sample this was not possible. The grouping of age-and-experience combinations into bins as shown in Table A.1 results in some smoothing of the estimated retention probabilities over the career.

Appendix Table A.1. Estimated Age-and-Experience Conditional Retention Probabilities for St. Louis Teachers.

Age/Experience	1	2	3	4	5	6-9	10-13	14-17	18-21	22-25	26-29	30-33	34-37	38+
<=24	0.78	0.60	0.75	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
25-29	0.81	0.79	0.78	0.83	0.93	0.73	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
30-34	0.81	0.79	0.78	0.79	0.76	0.88	0.88	N/A	N/A	N/A	N/A	N/A	N/A	N/A
35-39	0.90	0.81	0.88	0.92	0.85	0.81	0.89	0.87	0.92	N/A	N/A	N/A	N/A	N/A
40-44	0.86	0.80	0.90	0.84	0.84	0.88	0.89	0.92	0.93	0.95	N/A	N/A	N/A	N/A
45-49	0.78	0.94	0.86	0.84	0.85	0.89	0.95	0.91	0.95	1.00	0.96	N/A	N/A	N/A
50-54	0.84	0.83	0.90	0.92	0.91	0.91	0.96	1.00	0.97	0.98	0.98	0.92	N/A	N/A
55-59	0.78	0.95	0.89	1.00	0.95	0.88	0.91	0.90	0.97	0.93	0.96	0.81	0.87	0.79
60-64	0.54	1.00	1.00	1.00	0.63	0.88	0.90	0.80	0.82	0.65	0.81	0.75	0.84	0.66
65-69	0.55	0.77	0.87	0.84	0.81	0.73	0.83	0.45	0.68	0.63	0.57	0.76	0.50	0.76
70+	0.67	0.73	0.81	0.83	0.80	0.81	0.88	0.60	0.59	0.60	0.80	0.67	0.67	0.82

Notes: Due to the availability of large sample sizes along with quickly-changing retention rates, we use separate bins for low levels of teaching experience. Retention rates for several cells in the table are estimated using very few observations and in some cases we imputed values using the simple average of the values in adjacent cells (e.g., age=70+, experience=5). However, these cases are of little practical consequence because their weight in the EPW calculations is very small. All teachers' retention probabilities are set to zero at age-75.

Appendix B: Supplementary Tables

Appendix Table B.1 reports coefficients for the other variables from Model 4 in Table 5. We omit the coefficients on the age indicators and liner time trends for brevity.

Appendix Table B.1. Coefficients for Control Variables from Model 4 in Table 5.

	Model 4
Female	0.0097
	(0.0052)*
Asian	0.0543
	(0.0325)*
Black	0.0441
	(0.0045)**
Hispanic	0.0266
	(0.0416)
American Indian	-0.027
	(0.0584)
Master Degree	-0.004
	(0.0042)
Doctoral Degree	-0.0343
	(0.0265)
Other Non-Bachelor Degree	0.0150
	(0.0301)
Bin-1 Indicator	0.0360
	(0.0231)
Bin-2 Indicator	0.0977
	(0.0187)**
Bin-3 Indicator	0.0619
	(0.0178)**
Bin-4 Indicator	0.0556
	(0.0176)**
Bin-5 Indicator	0.0318
	(0.0168)*
Year-1998 Indicator	0.002
	(0.0262)
Post-Policy Indicator	0.0234
	(0.0360)
Age Indicators	X
Linear Time Trends	X
R-Squared	0.0729
N	18825

*Indicates statistical significance at the 10 percent level
**Indicates statistical significance at the 5 percent level
Notes: See notes to Table 5.

Appendix C: Using a Cost-Benefit Framework to Gain Insight about Novice Teachers

C.1 Overview

Table 5 shows that the enhancement generated a one-year retirement delay for bin-1, retirement-eligible teachers at the conclusion of the 1997-1998 school year. For retirement-ineligible teachers we do not find any evidence of a differential retention response across bins. If bin-6 teachers were unaffected then this would imply a null effect throughout the retirement-ineligible workforce. Although previous research suggests that it is not unreasonable to expect young teachers to be unaffected, or at least only marginally affected (French and Jones, 2012; Smith and West, 2014), as we note in the text we cannot evaluate novices directly because we use them as the baseline comparison group. Thus, our findings do not rule out the possibility that teachers in all retirement-ineligible bins (bins 2-6) were equally and positively affected by the enhancement. The purpose of this appendix is to examine the issue of novice teachers indirectly by estimating how large of an effect the enhancement would need to have on novices in order for it to be a cost neutral policy. For simplicity, we focus our cost-benefit analysis on the single cohort of retirement-ineligible teachers working during the 1998-1999 school year, plus retirement-eligible teachers in 1997-1998, because of a number of complications that arise in attempting to project costs and benefits into future years.[18]

[18] The most notable confounding issue that arises in attempting to calculate the long-term costs and benefits of the enhancement is that the long-term incidence of the costs is not clear. For example, if new, post-enhancement entrants in St. Louis went on to bear most of the enhancement's cost in the form of lower salaries over the course of their careers, then the cost burden on the district could be small. However, based on evidence from Fitzpatrick (2014), who shows that teachers do not value their pension benefits at the cost of providing them, in such a scenario there might also be a reduction in workforce quality, which could offset any retention benefits. For the 1998-1999 teaching cohort, many of whom had already been paid wages for large fractions of their careers, this is less of an issue. Our focus on the 1998-1999 cohort allows for a relatively clean cost-benefit analysis, and as will become clear below, it is sufficient to show that the enhancement is far from passing a cost-benefit test. At the very least, due to the presence of political and legal barriers to pension reform, one longer-term cost of the enhancement to the district (and future teachers) is that it imposed a constraint on the district's future expenditure choice set.

The procedure that we use to estimate the break-even effect of the enhancement on bin-6 teacher retention is straightforward. First, based on previous research we construct a general formula that can be used to calculate the monetary value of retaining experienced teachers over presumed novice replacements. Using this formula, we can specify a retention effect of the enhancement of any size and determine the dollar value of that effect. Second, we specify a formula for the cost of the St. Louis pension enhancement, which has direct and indirect components. The free parameter in both formulas is the level of retention caused by the enhancement. For any hypothetical retention effect that we specify, we can use the cost and benefit formulas to evaluate whether the enhancement would pass a cost-benefit test. We identify the retention effect on bin-6 teachers that equates the benefit and cost formulas as the "break-even" or "cost-neutral" policy effect.

C.2 Monetizing Retention Benefits

The research literature on teacher quality consistently identifies more experienced teachers as being more effective (e.g., see Clotfelter, Ladd and Vigdor, 2006; Kane, Rockoff and Staiger, 2008; Sass et al., 2012), and teacher effectiveness as being valuable (Chetty, Friedman and Rockoff, 2014; Hanushek, 2011). In our calculations we parameterize the effect of each additional year of retained experienced teaching over an assumed novice replacement at 0.11 standard deviations of student achievement based on Clotfelter, Ladd and Vigdor (2006). This is a per-student gain, and is a generous parameterization given our application.[19] We draw on Chetty, Friedman and Rockoff (2014) to estimate the dollar value of retaining *R* additional years of experienced teaching, realized through higher lifetime student earnings. To do this we use the following formula:

$$EB(R) = (0.11/0.18)*(R*CS)*(b*Y) \tag{C.1}$$

[19] One reason that this parameterization is generous is that it captures the relative value of teachers with more than 12 years of experience over novices (from the math models in Clotfelter, Ladd and Vigdor, 2006), but we apply it to retained teachers at any experience level. Also, Sass et al. (2012) estimate returns to teaching experience that are much lower than Clotfelter, Ladd and Vigdor (2006). Using an estimate based on their study in place of the estimate from Clotfelter, Ladd and Vigdor (2006) would result in our calculation of the enhancement's benefit falling by roughly half and require an even larger effect on bin-6 retention for the enhancement to be cost neutral.

The first term in parenthesis in equation (C.1) is the per-student effect on achievement of retained teaching experience, which is the experience effect divided by 0.18 to convert it into standard deviations of the distribution of teacher quality.[20] Per above, *R* is the total number of retained "experienced years" of teaching attributable to the enhancement, and the average class size is denoted by *CS*, which we set to 14.08 for our calculations.[21] Thus, the total size of any specified retention treatment in terms of the number of student-years affected is *14.08*R*. The impact of a one-standard-deviation improvement in teacher quality on earnings is denoted by *b*, which is assumed to be constant over the work life and parameterized at 1.34 percent based on Chetty, Friedman and Rockoff (2014). *Y* is expected average lifetime earnings for each student, which is parameterized at $416,026. This figure is also taken from Chetty, Friedman and Rockoff (2014), but adjusted to align the discount rate with our pension-wealth calculations and converted to 2013 dollars.[22]

C.3 Enhancement Costs

The cost of the enhancement, *EC*, has two components. First is the direct cost, which we estimate using the increase in expected pension wealth (EPW) caused by the enhancement for the 1998-1999 teacher cohort (see the main text and Appendix A for details about our EPW calculations).[23] The change in EPW caused by the enhancement is a function of the specified retention effect because increased workforce persistence generates higher pension payments. The function mapping retention

[20] Our parameterized value of the standard deviation of teacher quality, 0.18, is the simple average of the 10 estimates reported in Hanushek and Rivkin (2010) for math (see Table 1 in their paper).
[21] We do not have linked student-teacher data from St. Louis during this time. We approximate class size by the student-to-teacher ratio in the district, which over the time period of our data panel was 14.08. In the year-2000, the student-teacher ratio in Missouri public schools was 14.1; across the United States it was 16.0 (Snyder and Dillow, 2012, Table 71).
[22] We use a 4-percent real discount rate for our pension-wealth calculations (see Appendix A). We thank John Friedman for assistance in making the discount-rate adjustment to the Chetty, Friedman and Rockoff (2014) figure for our application. Note that our cost-benefit findings are not qualitatively sensitive to alternative, reasonable discount rates as long as the same discount rate is applied on both the cost and benefit sides.
[23] We include the gains in teaching experience from the one-year spike for bin-1 teachers in our benefit calculations. The costs associated with generating this spike are captured by our EPW calculations for the 1998-1999 cohort because retirement-eligible teachers from 1997-1998 who did not retire after that year are included in the 1998-1999 cohort.

effects to changes in total EPW is complicated due to the nonlinearities in pension-wealth accrual for individual teachers and is captured by our EPW calculations. A second, indirect cost component is the salary difference between retained teachers and their potential novice replacements, which also depends on the retention effects of the enhancement. We parameterize the salary cost using the 2013 teacher salary schedule in St. Louis. The formula for the enhancement cost can be written as:

$$EC(R) = DC(R) + IC(R) \tag{C.2}$$

In equation (C.2), *DC(R)* is the direct cost in terms of the change in EPW and *IC(R)* is the indirect salary cost.[24] When we set the retention effect on teachers during work years in bins 2-6 to zero, in which case the only retention effect is for bin-1 teachers in 1997-1998, we estimate the direct cost of the enhancement to be $166 million and the indirect cost to be $1 million.[25] The high direct cost estimate reflects the fact that even with a very small retention effect (the bin-1 retention spike alone), the enhancement is still costly because all teachers receive the improved benefit formula regardless of the behavioral response.

C.4 Finding the Break-Even Retention Effect

The enhancement's benefits and costs are both increasing in *R*, but because of the substantial value of teacher quality as established by Chetty, Friedman and Rockoff (2014), the benefits increase with *R* faster than the costs. Thus, starting with the scenario where the effect during bin-6 work years is assumed to be zero and the only retention effect is on bin-1 teachers – in which case the enhancement dramatically fails the cost-benefit test – we can increase the assumed retention effect during bin-6 work years, which reverberates throughout the workforce and raises *R*, until we reach the break-even value of the policy where *EB=EC*. The value of *R* that equates *EB* and *EC* answers the question: "How many years

[24] Per the discussion preceding equation (C.2), it is a simplification to specify the cost as a simple function of *R*, where *R* is defined above as the total years of retained experienced teaching. In fact, the cost depends on which years are retained. Our cost calculations account for this nuance, although for presentational convenience we do not expound on it further here.

[25] The former number is what we report in the main body of the paper as the direct cost of the enhancement.

of retained experienced teaching must be caused by the enhancement in order for it to be a cost neutral policy." *R* is operationalized in our calculations in terms of enhancement effects on annual retention over the course of teachers' careers. The annual retention effects compound over years to generate the increase in total *R*.

The formula for total *R* in terms of compounded, individual annual retention effects can be written as follows:

$$R = \sum_{t_1}^{T_1} I_{t_1}^{ae} + \sum_{t_2}^{T_2} I_{t_2}^{ae} + \sum_{t_3}^{T_3} I_{t_3}^{ae} + \sum_{t_4}^{T_4} I_{t_4}^{ae} + \sum_{t_5}^{T_5} I_{t_5}^{ae} + \sum_{t_6}^{T_6} I_{t_6}^{ae} \quad \text{(C.3)}$$

where the subscripts in equation (C.3) indicate bin numbers and T_j is the number of teachers in bin *j*. I^{ae} for a teacher with age *a* and experience *e* is expressed as follows

$$I^{ae} = \sum_{j=0}^{Z-a} \left(\prod_{i=0}^{j} A_{a+i,e+i}^{p}\right) - \sum_{j=0}^{Z-a} \left(\prod_{i=0}^{j} A_{a+i,e+i}^{np}\right) \quad \text{(C.4)}$$

and represents the cumulative increase in expected work years attributable to the enhancement's annual retention effects. $A_{a+i,e+i}^{p}$ and $A_{a+i,e+i}^{np}$ are estimated age- and experience-conditional annual retention rates over the course of the potential career starting with the current year (*j*=0), with and without the enhancement in place, respectively. *Z* is the terminal exit age, which we set to 75.[26]

Post-policy annual retention rates conditional on age and experience in St. Louis are reported in Appendix Table A.1. The annual retention rates that we use in equation (C.4) with the policy, A^{p}, are taken directly from the table. We can specify any hypothetical policy effect of the enhancement within this framework by changing (lowering) the values of A^{np} relative to A^{p}, where A^{np} represents counterfactual post-policy annual retention rates. In the case where the policy effect on retention is

[26] Put differently, at age-75 all teachers are assumed to exit with probability one. Forcing teachers to exit at age-75 is of no practical consequence for our calculations because the probability of surviving to age-75 is very low.

zero throughout the career, the vectors A^p and A^{np} will be identical for any given age-experience combination.

We numerically solve for *R* by equating equations (C.1) and (C.2), while imposing the constraint that the enhancement's effect on annual retention for bin-6 teachers is maintained throughout bins 2-5 as well (i.e., no effect heterogeneity per Table 5). Denoting the annual retention effect on teachers in bins 2-6 as *X*, the effect on bin-1 teachers in 1997-1998 is set to (*X* + 0.0723), where 0.0723 is the differential effect of the enhancement on bin-1 teachers in 1997-1998.

We estimate that the break-even value for the annual retention effect on teachers in work years covered by bins 2-6 is 6.5 percentage points annually. Correspondingly, the 1997-1998 spike for bin-1 teachers at the break-even effect size is 13.7 percentage points (6.5 + 7.2). How large is this break-even value? To illustrate, consider the implied effect on retention until full retirement for a representative age-24 entrant into St. Louis. To estimate this effect we perform a calculation based on equation (C.4), modified to estimate the effect of the enhancement on the likelihood of the representative new teacher surviving until full retirement. We use the break-even scenario where the enhancement effect is a 6.5 percentage-point increase in the annual retention rate throughout the retirement-ineligible career. This calculation indicates that the break-even effect corresponds to more than a *ten-fold increase* in the likelihood of survival until full retirement for a representative age-24 entrant.[27]

We provide some context for this number by comparing it to findings from Smith and West (2014), which to the best of our knowledge is the only study that directly investigates the effect of a

[27] More concretely, we use age- and experience-conditional annual retention rates in the post-policy years in St. Louis (from Appendix Table A.1) to construct the equivalent of A^p in equation (C.4) for the representative age-24 entrant. These annual retention rates can be used to predict the likelihood of survival until full retirement, which is 4.2 percent (two-thirds of new entrants in St. Louis leave within the first five years alone; see McGee and Winters (2013) for more on the general urban teacher retention problem in the pension context). If we subtract out the break-even enhancement effects to construct counterfactual annual retention rates under the break-even scenario – the equivalent of A^{np} from equation (C.4) – and recalculate the probability of survival until full retirement, it falls to just 0.3 percent.

change in pension benefits on the retention of new workforce entrants. Smith and West (2014) examine a change to the pension formula for military personnel that effectively reduced the payoff to surviving until retirement eligibility by 20 percent. Although there are obvious issues in comparing their results to ours, ranging from differences in working conditions between urban teachers and military personnel, differences in how long it takes to reach retirement eligibility across sectors (military personnel are eligible for retirement at an earlier age and are much more likely to reach eligibility), and differences in the pension-benefit change across studies (a 20 percent reduction versus 60 percent increase), it is notable that Smith and West (2014) estimate a very small labor-supply elasticity to the pension-benefit change for new entrants into the military. In particular, they find that the 20-percent reduction in benefits caused just a 2-3 percent decline in the probability of survival to retirement eligibility for new entrants. Obviously, our comparable estimate using the break-even bin-6 effect suggests a level of responsiveness for teachers far out of line with what Smith and West (2014) find for military personnel, but this is the level of responsiveness that would be required for the pension enhancement to have been a cost-neutral policy. Although our study is not designed to examine the effect of the enhancement on bin-6 teachers directly, we conclude that this effect would need to be implausibly large for the enhancement to pass a cost-benefit test.