

Pension Enhancements and the Retention of Public Employees

Cory Koedel
P. Brett Xiang

First Version: October 2014
This Version: October 2015

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. Our results indicate that the St. Louis enhancement – which was structurally similar to enhancements that were enacted in other public pension plans across the United States in the late 1990s and early 2000s – was not a cost-effective way to increase employee retention.

Acknowledgments

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, the Education Seminar Series at the Federal Reserve Bank of New York, and the CALDER conference 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.

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 state 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 (Gustman, Mitchell, and Steinmeier, 1994; 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).

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.

We contribute to the literature on how workers respond to pension incentives by examining the effect on retention of increasing pension backloading via benefit-formula enhancements. Improving our understanding of how workers respond to changes to plan rules within the DB pension framework 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), or over \$52,000 per teacher on average. This represents over one quarter of the entire operating budget for the district at the time.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, along with the change in the incentive to remain in covered employment, varied considerably across teachers.

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); for more on teacher plans specifically 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.

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 teachers within St. Louis to identify the enhancement's retention effects. We highlight two key findings. First, the largest behavioral response to the enhancement came in the form of a temporary delay in retirements during a "gap year" between the approval and enactment of the enhancement among retirement-eligible teachers. The delayed retirements are consistent with a temporary change to these teachers' retention incentives during the gap year. Second, 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 strong evidence of a meaningful behavioral response. Although our preferred estimates are imprecise, even at their upper bounds they indicate that the retention effects of the enhancement were too small to justify the cost of its implementation.

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 St. Louis enhancement 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.

We conclude by offering several explanations for why the large changes to teachers’ retention incentives created by the St. Louis pension enhancement generated only a limited behavioral response. We also consider alternative explanations for why the enhancement was enacted in the first place given that its significant cost cannot be justified in terms of workforce retention benefits. 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,

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.

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 workers. It was enacted for teachers retiring on or after June 30, 1999. Individuals who began collecting benefits after the 1998-1999 school year received their pensions based on the improved formula; individuals who began collecting their pensions prior to the conclusion of the 1998-1999 school year received 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:

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

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-

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

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 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, maximum 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

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.

enhancement until far into the future, teachers at or near retirement eligibility at the time when the enhancement was enacted received 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 teachers who were closest to benefit eligibility when it was enacted.

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 teachers' pension wealth profiles (again, see Appendix A for details). Descriptive statistics are reported in Table 1.⁸

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.⁹ 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 benefit collection for one additional

⁸ We end the data panel after the 1999-2000 school year because we are concerned about other factors outside of the enhancement influencing our estimates as we move further away from the policy event. Perhaps the biggest concern is that the economy entered a mild recession in 2001, which may have affected the teacher labor market. In an analysis omitted for brevity we examine the sensitivity of our results to extending the data panel beyond the year-2000 and our findings are qualitatively unaffected.

⁹ Our data are insufficient to determine whether the enhancement was offset by lower wages, or lower wage growth, by the district. A descriptive review of wage data in St. Louis over the course of our data panel suggests that real wages were fairly flat, but the counterfactual is unobserved. Although pensions are not collectively bargained with other dimensions of the compensation package, we cannot rule out an informal tradeoff, particularly because the St. Louis plan is municipal (as opposed to state-level). If wage growth was reduced in St. Louis as an informal tradeoff against the enhancement, it would likely lead to upward bias in our estimates of the relative effects of the enhancement on retention because younger teachers, whose pension incentives changed the least, would also likely be the ones who would be more responsive to a reduction in wage growth (Farber, 1999).

year. Note that a teacher who was planning to retire after the 1997-1998 school year could receive the improved benefit simply by delaying collection, regardless of whether she chose to work during the 1998-1999 school year. However, the opportunity cost of continued work during the 1998-1999 school year was greatly reduced for retirement-eligible teachers because the optimal decision in most circumstances would be to delay retirement collection until after the 1998-1999 school year regardless of the work choice. Put differently, DB-covered workers normally experience a sharp spike in the opportunity cost of continued work once they become eligible for benefit collection because pension payments are foregone while working (Koedel, Podgursky and Shi, 2013). However, because the stream of pension payments is so much more valuable under the enhanced formula, waiting until after the 1998-1999 school year to collect would be optimal for most teachers regardless of the work decision during that year, in which case there would not be foregone pension payments associated with continued work during the 1998-1999 school year.¹⁰

Although we were unable to find direct evidence to document the extent to which teachers knew about the approval of the enhancement prior to its enactment, 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 and that this information factored into their retirement decisions. Given that there is some uncertainty regarding the extent to which information about the enhancement was available to all teachers during the 1997-1998 school year, and the unique situation of retirement-eligible teachers at the conclusion of that 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

¹⁰ Related to this point is that retirement eligible teachers at the conclusion of the 1997-1998 school year who wanted to delay collection until after the following year may have faced liquidity constraints. This also would have pushed them to work during the 1998-1999 school year.

the enactment of the enhancement (1997-1998), and (3) after the enactment of the enhancement (1998-1999 and 1999-2000).

As a final note on the timing issue, the gap year was important for retirement-ineligible teachers at the conclusion of the 1997-1998 school year, who by definition would not be eligible to collect retirement benefits for at least one additional year regardless of their quit decision at the conclusion of the 1997-1998 school year, primarily because it may have affected general awareness of the enhancement during that year. Put differently, if we assume that all teachers were aware of the pending change, the 1997-1998 school year should be viewed no differently than any other post-policy year for teachers who were not eligible for collection until 1998-1999 or later. If some teachers were not aware of the pending change during the 1997-1998 school year, they would be expected to behave as in the pre-enhancement years.¹¹

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' pension wealth and 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-

¹¹ As an example, consider a 40-year old teacher with 10 years of experience who is deciding whether to continue teaching in St. Louis Public Schools or exit at the conclusion of the 1997-1998 school year. The fact that the teacher is vested ensures that she will be eligible for a pension, but in this case if she leaves she will not be eligible to file for retirement until age-65, which will occur well after the planned enactment in June of 1999. Whether she works during the 1998-1999 school year has no bearing on which formula will be used to determine her pension benefit.

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, maximum 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, maximum pension wealth (MPW) is the value of the pension at the top of the accrual curve. The enhancement's effect on MPW is an overstatement because all workers will not reach and retire at the maximum. Following Koedel, Ni and Podgursky (2014), we also calculate expected pension wealth (EPW) for each teacher with and without the enhancement, with these calculations serving as the foundation for our preferred estimate of the total cost. 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.

Two patterns in the table merit attention. First, the gaps between the enhancement gain measured in terms of MPW 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 at much higher risk of leaving the system before reaching the maximum. Second, and more importantly for the present analysis, the MPW and CPW gains in pension wealth owing to the enhancement vary considerably across bins. Among the retirement-ineligible workforce, teachers in lower-numbered bins experienced larger increases in both MPW and CPW. The marginal pension gain associated with remaining in the profession until full retirement induced by the enhancement for individual teachers – what is referred to by Coile and Gruber (2007) as “peak value” – is the difference between the MPW and CPW gains. For example, the gains in MPW and CPW on average for bin-2 teachers as reported in Table 2 were \$109,905 and \$77,312, respectively, which results in a peak-value increase of roughly \$33,000. Dividing this value by the average distance to full retirement for bin-2 teachers of 3.1 years, the average annualized increase in peak-value is \$10,500 per year until full retirement. For teachers in bins 3, 4, 5 and 6 the average annualized increases in the pension incentive are \$7,100, \$3,900, \$2,800 and \$2,100, respectively.

These cross-bin differences in the retention incentive change are substantial. For example, the annualized peak-value incentive increase is five times higher in bin-2 than bin-6, and more than twice as large for bin-2 teachers relative to bin-4 teachers (to put these numbers in context, note that the average annual salary for teachers in our analytic sample is \$48,916 in 2013 dollars). It is also notable that these annualized gains are conditional on survival until full retirement. The annual gains would be much smaller in the case of early exit, and younger workers are at a higher risk of not surviving in the profession until full retirement as indicated by the large gaps in EPW across bins as shown in the table.¹²

¹² The effect of the differential survival risk across bins on the retention incentive is difficult to pin down precisely because the survival risk itself can be a function of the retention incentive. Although we do not formally attempt to quantify the role of differential survival risks across bins in affecting the change in teachers' retention incentives, we note

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. Our estimates will be inclusive of any offsetting wealth effects of the enhancement owing to the increases in current pension wealth. Speaking empirically, every pension enhancement to a subnational plan that occurred during the enhancement boom in the late 1990s and early 2000s of which we are aware included some form of retroactive implementation. Therefore, from the perspective of a holistic evaluation of how pension enhancements affect workers behaviorally, total effect estimates that are inclusive of any wealth effects are most relevant.¹³

Finally, 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, which facilitates identification in the models we introduce in the next section.

4. Empirical Strategy

4.1 Primary Models

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

$$Y_{it} = \beta_0 + X_{it}\beta_1 + BIN_{it}\beta_2 + 1998_{it}\beta_3 + POST_{it}\beta_4 + (1998_{it} * BIN_{it})\theta_1 + (POST_{it} * BIN_{it})\theta_2 + \varepsilon_{it} \quad (3)$$

In (3), Y_{it} is an indicator variable equal to one if teacher i was retained and zero if she 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

that survival risk differences will exacerbate the incentive gaps by raising the value of the enhancement for teachers closer to retirement because of their lower risk of exit before reaching full retirement eligibility.

¹³ As for why retroactive implementation is so prevalent we can only speculate, but one possibility is that political constraints are such that it is necessary to induce large wealth effects to feasibly enact a pension enhancement. The political forces that drive pension changes merit additional attention in future research.

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).¹⁴ BIN_{it} is a vector of bin indicator variables based on the bin classifications established in the previous section (Table 2). The coefficient vector β_2 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) and the “1998” indicator being for the 1997-1998 school year, during which the enhancement was approved but prior to its enactment. 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.¹⁵

An identifying assumption in equation (3) is that pre-enhancement trends in retention rates are the same across bins. In Section 4.2 we provide evidence inconsistent with this assumption, which prompts the following expansion and modification of the model following Jacob (2005):

$$Y_{it} = \gamma_0 + X_{it}\gamma_1 + BIN_{it}\gamma_2 + 1998_{it}\gamma_3 + (1999_{it}\gamma_4 + 2000_{it}\gamma_5) + T_{it}\gamma_6 + (T_{it} * BIN_{it})\gamma_7 + (1998_{it} * BIN_{it})\pi_1 + (1999_{it} * BIN_{it})\pi_2 + (2000_{it} * BIN_{it})\pi_3 + u_{it} \quad (4)$$

¹⁴ 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).

¹⁵ 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.

Equation (4) can be characterized as a difference-in-difference of a short interrupted time series. It is of the same structure as equation (3) but makes two adjustments. First, we introduce bin-specific linear time trends via T_{it} , which is a linear time variable, and its interaction with the bin indicators. Second, to ensure that the time-trend parameters γ_6 and γ_7 are identified using variation from the pre-enhancement years only, we separate out the post-period years – 1999 and 2000 – and estimate different parameters for each. We make the latter adjustment purely for mechanical purposes as it ensures that there is no identifying variation in T_{it} in the post-enhancement period (because the value of T_{it} does not vary within a single year). Therefore, the only variation used to identify γ_6 and γ_7 comes from the pre-period years 1995, 1996, and 1997 (this modification to the model follows Jacob, 2005, who faces a similar identification issue). For consistency of reporting across the models in equations (3) and (4), and to improve our power in estimating the post-period effect in equation (4), when we show our results below we report a single post-period parameter estimate and standard error from equation (4) for each bin. The post-period parameters are linear combinations of our estimates of π_2 and π_3 (with the standard errors properly adjusted for the covariance).

When we report our results from equations (3) and (4), we focus on comparing teachers in bins 1 through 4 to teachers in bins 5 and 6. We combine bins 5 and 6 into a common control group for two reasons: (1) to improve statistical power, and (2) because the effect of the enhancement on teachers' annualized incentives is similar for teachers in bins 5-6, per the preceding section. We have also estimated all of our models using bin-6 as the only holdout group and we obtain qualitatively similar (albeit noisier) results.

In equation (4), which is our preferred specification for reasons that will become clear below, the identifying assumption is that deviations from the linear time trends across bins 1-4, relative to bins 5-6, 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 –teachers in bins 5-6 in our case – as zero, but teachers in bins 5-6 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 the above discussion.¹⁶

4.2 Retention Trends in St. Louis

Figure 2 shows annual retention rates for St. Louis teachers by bin in each year of our data panel. Teachers in bins 5-6 are combined to maintain consistency with the models described in the previous section. As would be expected based on the extant literature on teacher attrition (Boyd et al., 2011; Ingersoll, 2001; Loeb, Darling-Hammond and Luczak, 2009), the figure 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 figure 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

¹⁶ In an omitted analysis we also constructed models analogous to the models in equations (3) and (4) that compare 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 models that compare St. Louis teachers to outside teachers do 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 larger, limiting inference.

conditions in St. Louis, because they are less occupationally attached than their more senior counterparts (e.g., see Kambourov and Manovskii, 2008).

To formalize the relative patterns in retention rates shown in Figure 2, we estimate the following model based on Fitzpatrick and Lovenheim (2014):

$$Y_{it} = \alpha_0 + X_{it}\alpha_1 + BIN_{it}\alpha_2 + G_{it}\alpha_3 + (G_{it} * BIN_{it})\alpha_4 + u_{it} \quad (5)$$

In equation (5), G_{it} is a vector of year indicator variables. It replaces the timespan variables (*1998*, *POST*) from equation (3). The other variables are specified as above.

Based on the output from equation (5), in Figure 3 we construct time trends in retention rates conditional on teacher characteristics for teachers in all groups relative to teachers in bins 5-6 in 1995 (the first year of our data panel). We note two important aspects of the trends shown in the figure. First, consistent with the retention rates shown in Figure 2 and 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, of direct relevance for the modeling, the results from equation (5) confirm a larger decline in retention rates for bin 5-6 teachers relative to other teachers prior to the enhancement (i.e., from 1995 through 1997). We account for the trend differences illustrated in Figure 3 with the linear time trend controls in equation (4). Results from the restricted model as shown in equation (3) will overstate the effect of the pension enhancement on differential retention in St. Louis by failing to account for the pre-policy divergence in retention rates across bins.¹⁷

¹⁷ A related issue is that outside factors may have influenced the composition of the young teaching workforce over time. For example, if non-teaching opportunities were improving during the second half of the 1990s due to the booming economy, then teacher quality may have been declining among new entrants into St. Louis over the course of our data panel (e.g., see Nagler, Piopiunik and West, 2015). 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 in our models (Goldhaber, Gross and Player, 2011; Krieg, 2006; West and Chingos, 2009).

5. Results

Table 4 presents estimates from equations (3) and (4). The table shows the difference-in-difference coefficients from several variants of the model in equation (3), with the estimates in the third column coming from the full specification. The fourth column shows estimates from equation (4). Coefficients for the other control variables not shown in the table are reported in Appendix B.

We begin by discussing our findings for retirement-ineligible teachers (bins 2, 3, and 4 relative to teachers in bins 5-6). 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 2-4 were differentially affected by the pension enhancement relative to teachers in bins 5-6. This is a notable result in light of the substantial differences in how teachers' retention incentives were affected across bins as discussed above, and the overall cost of the enhancement. However, although our point estimates in column (4) are nominally negative and provide no indication of a differential retention effect, a caveat is that the estimates are imprecise. One reason 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 in the table.¹⁸

Given the statistical power issue, one option is to focus on the results from Model 3, which excludes the linear time trend controls but produces estimates that are more precise. As noted above, 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 bins 5-6 are more positive, and the estimates for teachers in bins 2

¹⁸ The lack of evidence of a differential behavioral response among retirement-ineligible teachers in our preferred specification is also consistent with findings reported by Chan and Stevens (2008). These authors show that while most workers understand the basic framework of their pension plans (e.g., the retirement age, the formula factor, etc.), they are much less knowledgeable about the more complex aspects of their plans, like the value of remaining in covered employment until full retirement eligibility. Note that the St. Louis enhancement did not alter the basic pension framework, but rather the return to meeting fixed benchmarks.

and 3 are statistically significant. The findings from Model 3 are best interpreted as upper-bound estimates of the enhancement's differential effects on retention, but taken at face value they imply some behavioral response to differential changes in teachers' retention incentives, at least for teachers closest to retirement relative to younger teachers. To properly contextualize these estimates, below we evaluate them within a cost-benefit framework to determine whether the size of the implied behavioral response under the favorable conditions of Model 3 can justify the cost of the enhancement legislation.

Next we turn to bin-1, retirement-eligible teachers. The estimate in row (1) and column (4) of Table 4 shows that retention for these teachers spiked by 7.0 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, in which case they would be better off delaying their retirement filings until the conclusion of the 1998-1999 school year regardless of their decision to continue teaching beyond 1997-1998. Given these circumstances, this particular cohort of retirement-eligible teachers did not face the high opportunity cost of continued work that is typical for retirement-eligible workers covered by DB pension plans. Their workforce retention behavior is consistent with this one-year incentive change.

Although the behavior of retirement-eligible teachers at the conclusion of the 1997-1998 school year 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, and chose to return to teaching for an additional year given that they would not begin collecting their pensions immediately. Given our uncertainty about how bin-1 teachers knew about the pending enhancement at the conclusion of 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 the results, or lack thereof, in Table 4.

As a final note on our findings, we return to the issue that our control group of younger/less-experienced teachers in bins 5-6 is not entirely untreated. Large retention effects for these teachers seem unlikely based on outside evidence (Fitzpatrick, forthcoming; French and Jones, 2012; Smith and West, 2014) but we are unable to examine them directly via equations (3) and (4). While our results are still informative about the relative effects of the enhancement without knowing the baseline effect on teachers in bins 5-6, we cannot directly evaluate the enhancement policy holistically with the estimates from Table 4 alone. We gain some indirect insight in the next section by using a cost-benefit framework to determine the size of the effect on teachers in bins 5-6 that would be required in order for the enhancement to pass a cost-benefit test, and then assessing the plausibility of the break-even effect size.

6. Cost-Benefit Analysis

6.1 Overview

In this section we perform a cost-benefit analysis of the pension enhancement. For simplicity, we focus 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.¹⁹

¹⁹ 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 preceding analysis informs our parameterization of the relative effects of the enhancement on teachers at different points in the career cycle – e.g., the marginal increase in retention for a bin-2 teacher relative to a teacher in bins 5-6. We use two different parameterizations of the relative retention effects, both of which offer a generous interpretation of our findings: (1) we use the upper bounds of the 95 percent confidence intervals of the estimates from Model 4 in Table 4, and (2) we use the point estimates from Model 3 in Table 4 (both parameterizations are made without regard to statistical significance of the individual coefficients; e.g., the parameterized annual effect for bin-4 teachers relative to teachers in bins 5-6 is 1.87 percentage points based on the results from Model 3).²⁰ In both scenarios we parameterize the one-year spike effect for bin-1 teachers in 1998 but do not parameterize a long-term effect for these teachers. Our analysis gives no indication that bin-1 teachers responded permanently to the enhancement, but the one-year retention spike for bin-1 teachers influences our benefit calculations.

With the relative retention effects across the workforce in hand per the above parameterizations, we can use a cost-benefit framework to recover the break-even retention effect of the enhancement for teachers in bins 5-6. We proceed in the following steps. 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. Next we specify a formula for the enhancement’s cost. The free parameter in both formulas is the level of retention caused by the enhancement. For any hypothetical retention effect that we specify for teachers in bins 5-6, and with the relative effects for teachers in lower-numbered bins in hand, we can use the cost and benefit formulas to evaluate whether the enhancement would pass a cost-benefit test.

²⁰ As a practical matter these two scenarios end up being fairly similar because the upper bound of the 95 percent confidence intervals from Model 4 are fairly close to the point estimates from Model 3.

We identify the retention effect on teachers in bins 5-6 that equates the benefit and cost formulas as the “break-even” or “cost-neutral” policy effect. This value answers the question “How large of an effect would the enhancement need to have on teachers in bins 5-6 for it to be a cost-neutral policy?”

6.2 Monetizing Retention Benefits

The research literature on teacher quality consistently identifies more experienced teachers as more effective (e.g., see Clotfelter, Ladd and Vigdor, 2006; Kane, Rockoff and Staiger, 2008; Sass et al., 2012), and teacher effectiveness as 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.²¹ 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:

$$EB(R) = (0.11 / 0.18) * (R * CS) * (b * Y) \tag{6}$$

The first term in parenthesis in equation (6) 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.²² 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

²¹ One reason that this parameterization is generous is that it captures the 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 5-6 retention for the enhancement to be cost neutral.

²² 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).

to 14.08 for our calculations.²³ 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.²⁴

6.3 *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 Appendix A for details about our EPW calculations).²⁵ The change in EPW depends on the effect of the enhancement on retention because increased workforce persistence generates higher pension payments. The function mapping retention 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:

²³ We do not have linked student-teacher data from St. Louis during this time. We approximate class size by the student-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).

²⁴ 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.

²⁵ As noted above, 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.

$$EC(R) = DC(R) + IC(R) \tag{7}$$

In equation (7), $DC(R)$ is the direct cost in terms of the change in EPW and $IC(R)$ is the indirect salary cost.²⁶

Under an initial scenario where the retention effect on all retirement ineligible teachers is set 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.²⁷ The high direct cost estimate reflects the fact that even with a very small retention effect, the enhancement is still costly because all teachers receive the improved benefit formula regardless of their behavioral response.

6.4 *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. Starting with the scenario where the effect during bin 5-6 work years is set to be zero, and using the above-described parameterized effects for teachers during work years in bins 2, 3 and 4, the enhancement dramatically fails the cost-benefit test. From this point, we can increase the assumed retention effect during bin 5-6 work years, which reverberates throughout the workforce and raises R , until we reach the break-even value of the policy where $EB=EC$.

The variable 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

²⁶ Per the discussion preceding equation (7), 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.

²⁷ The former number is what we report above as the direct cost of the enhancement based on a “static” calculation that does not allow for a retention effect.

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

$$R = \sum_{t_1}^{T_1} I_{t_1}^{ae} + \sum_{t_2}^{T_2} I_{t_2}^{ae} + \sum_{t_3}^{T_3} I_{t_3}^{ae} + \sum_{t_4}^{T_4} I_{t_4}^{ae} + \sum_{t_{5,6}}^{T_{5,6}} I_{t_{5,6}}^{ae} \quad (8)$$

where the subscripts in equation (8) 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

$$I^{ae} = \sum_{j=0}^{Z-a} \left(\prod_{i=0}^j A_{a+i,e+i}^p \right) - \sum_{j=0}^{Z-a} \left(\prod_{i=0}^j A_{a+i,e+i}^{np} \right) \quad (9)$$

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 75.²⁸

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 (9) 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 zero throughout the career, the vectors A^p and A^{np} will be identical for each age-experience combination.

We numerically solve for R by equating equations (6) and (7), while imposing the constraint that heterogeneity in the enhancement's effect on annual retention throughout the career is

²⁸ 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.

represented by the two parameterizations discussed above. For example, denoting the specified annual retention effect on teachers in bins 5-6 as X , the effect on annual retention for bin-2 teachers in the post-policy period is set to $(X + 0.0317)$ in the parameterization based on Model 3 in Table 4, where 0.0317 is the differential effect of the enhancement during bin-2 work years in the post-policy period. Using the parameterization based on the 95-percent-confidence-interval upper-bound estimates from Model 4, the break-even effect on retention for teachers in bins 5-6 is 3.6 percentage points annually; using the point estimates from Model 3 the break-even effect size is 4.5 percentage points.

How large are these break-even values? We illustrate using the smaller estimate of 3.6 percentage points by considering 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 (9), 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 3.6 percentage-point increase in the annual retention rate during work years in bins 5-6, 5.47 (3.6+1.87) percentage points during work years in bin-4, etc. This calculation indicates that the break-even effect corresponds to more than a *six-fold increase* in the likelihood of survival until full retirement for a representative age-24 entrant.²⁹

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 estimates the effect of a change in pension benefits on the retention of new workforce entrants. Smith and West (2014)

²⁹ 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 A^p in equation (9) 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 about the 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 (9) – and recalculate the probability of survival until full retirement, it falls to just 0.7 percent.

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 effect for teachers in bins 5-6 suggests a level of responsiveness for teachers far out of line with what Smith and West (2014) find for military personnel. However, 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 directly estimate the effect of the enhancement on teachers in bins 5-6, we conclude that this effect would need to be implausibly large for the enhancement to pass a cost-benefit test.

7. 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 do not find strong evidence to suggest that these differences translated into differences in retention behavior, with the exception of retirement-eligible teachers who were more likely to delay their retirements for one year in response to a short-term change to their incentives. Although estimates from our preferred specification are too imprecise to rule out a moderate behavioral response among retirement-ineligible

teachers, our cost-benefit analysis shows that the enhancement was not a cost effective way to promote employee retention.

One explanation for our findings 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; Chan and Stevens, 2008; Gustman and Steinmeier, 2005). Another possibility is that income effects may have worked to offset some potential retention benefits of the enhancement. This is a potentially important concern given the retroactive implementation of the benefit improvement, which greatly increased the immediate value of teachers' pensions, although it is notable 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.³⁰ A third factor that may contribute 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. Again, the retroactive implementation of the benefit formula is an issue because it resulted in the enhancement increasing the retention incentives of late career teachers the most, and late career workers are more inertial in their careers than their early-career counterparts (Farber, 1999; Harris and Adams, 2007). 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.

³⁰ Per above, it is also an empirical point that every pension enhancement that occurred during the enhancement boom in the late 1990s and early 2000s of which we are aware included some form of retroactive implementation. Therefore, total effect estimates that are inclusive of any wealth effects are most relevant for informing policy decisions with regard to pension enhancements.

All of that said, if the decision makers who approved the enhancement had been asked at the time to identify its likely benefits to students in St. Louis Public Schools, it is hard to imagine a more compelling policy rationale than improving teacher retention. Perhaps a competing rationale would be to improve recruitment, but if improving recruitment were truly the objective then the funds devoted to support the enhancement could have been targeted much more effectively toward new entrants. Of course, the enhancement in St. Louis and other similar enhancements in subnational pension plans across the United States may 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.³¹

We conclude by noting several policy implications of our results. First, although the last period of widespread pension enhancements in subnational DB pension plans occurred 15 years ago, and

³¹ 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.

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, our analysis suggests that any adverse behavioral response by teachers to pension-benefit reductions will likely be less costly than the associated savings. 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.³² Third, and related to the previous point, our findings are consistent with Fitzpatrick (forthcoming), who shows that 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.

³² 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).

References

Bertrand, Marianne, Esther Duflo and Sendhil Mullainathan. 2004. How Much Should We Trust Differences-In-Differences Estimates? *Quarterly Journal of Economics* 119(1), 249-275.

Biggs, Andrew. 2011. An Options Pricing Method for Calculating the Market Price of Public Sector Pension Liabilities. *Public Budgeting and Finance* 31(3), 94-118.

Boyd, Donald, Pam Grossman, Marsha Ing, Hamilton Lankford, Susanna Loeb, and James Wyckoff. 2011. The Influence of School Administrators on Teacher Retention Decisions. *American Educational Research Journal* 48(2): 303-333.

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.

Brown, Kristine. 2013. The Link Between Pensions and Retirement Timing: Lessons from California Teachers. *Journal of Public Economics* 98(1), 1-14.

Brown, Kristine M. and Ron A. Laschever. 2012. When They're Sixty-Four: Peer Effects and the Timing of Retirement. *American Economic Journal: Applied Economics* 4(3), 90-115.

Brown, Jeffrey R., Arie Kapteyn, Erzo F. P. Luttmer, and Olivia S. Mitchell. 2013. Do Consumers Know How to Value Annuities? Complexity as a Barrier to Annuitization. NBER Working Paper No. 19168.

Chan, Sewin and Ann Huff Stevens. 2008. What you Don't Know Can't Help You: Pension Knowledge and Retirement Decision Making. *Review of Economics and Statistics* 90(2), 253-266.

Chetty, Raj, John N. Friedman and Jonah E. Rockoff. 2014. Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood. *American Economic Review* 104(9), 2633-79.

Chingos, Matthew W. and Martin R. West. 2015. Which Teachers Choose a Defined Contribution Pension Plan? Evidence from the Florida Retirement System. *Education Finance and Policy* 10(2), 193-222.

Clotfelter, Charles T., Helen F. Ladd and Jacob L. Vidgor. 2006. Teacher-Student Matching and the Assessment of Teacher Effectiveness. *Journal of Human Resources* 41(4): 778-820.

Coile, Courtney and Jonathan Gruber. 2007. Future Social Security Entitlements and the Retirement Decision. *The Review of Economics and Statistics* 89(2), 234-246.

Costrell, Robert M. and Michael J. 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.

DeArmond, Michael and Dan Goldhaber. 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.

Farber, Henry S. 1999. The Dynamics of Job Change in Labor Markets in *Handbook of Labor Economics* (Eds. Orley Ashenfelter and Richard Layard), Volume 3B, 2439-2483.

Fitzpatrick, Maria (forthcoming). How Much Do Public School Teachers Value Their Pension Benefits? *American Economic Journal: Economic Policy*.

Fitzpatrick, Maria and Michael Lovenheim. 2014. Early Retirement Incentives and Student Achievement. *American Economic Journal: Economic Policy* 6(3), 120-154.

French, Eric and John Jones. 2012. Public Pensions and Labor Supply Over the Life Cycle. *International Tax and Public Finance* 19(2), 268-287.

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.

Glaeser, Edward L. and Giacomo A.M. Ponzetto. 2014. Shrouded Costs of Government: The Political Economy of State and Local Public Pensions. *Journal of Public Economics* 116, 89-105.

Goldhaber, Dan, Bethany 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.

Goldhaber, Dan and Cyrus Grout (forthcoming). Which Plan to Choose? The Determinants of Pension System Choice for Public School Teachers. *Journal of Pension Economics and Finance*.

Gustman, Alan L, Olivia S Mitchell, and Thomas L Steinmeier. 1994. The Role of Pensions in the Labor Market: A Survey of the Literature. *Industrial and Labor Relations Review* 47(3), 417-38.

Gustman, Alan L. and Thomas L. Steinmeier. 2005. Imperfect Knowledge of Social Security and Pensions. *Industrial Relations* 44(2), 373-397.

Gustman, Alan L. and Thomas 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, Janet S. 2010. An Introduction to Teacher Retirement Benefits. *Education Finance and Policy* 5(2): 402-437.

Hanushek, Eric A. 2011. The Economic Value of Higher Teacher Quality. *Economics of Education Review* 30(3), 266-479.

Hanushek, Eric A. and Steven G. Rivkin. 2010. Generalizations about Using Value-Added Measures of Teacher Quality. *American Economic Review* 100(2), 267-271.

- Harris, Douglas N. and Scott J. Adams. 2007. Understanding the Level and Causes of Teacher Turnover: A Comparison with Other Professions. *Economics of Education Review* 26(4), 325-337.
- Ingersoll, Richard M. 2001. Teacher turnover and teacher shortage: An organizational analysis. *American Educational Research Journal* 38(3): 499-534.
- Jacob, Brian. 2005. Accountability, Incentives and Behavior. The Impact of High Stakes Testing in Chicago Public Schools. *Journal of Public Economics* 89, 761-796.
- Kambourov, George and Iourii Manovskii. 2008. Rising Occupational and Industry Mobility in the United States: 1968-97. *International Economic Review* 49(1), 41-79.
- Kane, Tom J., Jonah E. Rockoff, and Douglas O. Staiger. 2008. What Does Certification Tell us about Teacher Effectiveness? Evidence from New York City. *Economics of Education Review* 27(6), 615-631.
- Kersaint, Gladis, Jennifer Lewis, Robert Potter and Gerry Meisels. 2007. Why Teachers Leave: Factors that Influence Retention and Resignation. *Teaching and Teacher Education* 23(6), 775-794.
- Koedel, Cory, Shawn Ni and Michael Podgursky. 2014. Who Benefits from Pension Enhancements? *Education Finance and Policy* 9(2), 165-192.
- Koedel, Cory, Michael Podgursky and Shishan Shi. Teacher Pension Systems, the Composition of the Teaching Workforce, and Teacher Quality. *Journal of Policy Analysis and Management* 32(3), 574-596.
- Krieg, John M. 2006. Teacher Quality and Attrition. *Economics of Education Review* 25(1), 13-27.
- Lazear, Edward P. 1990. Pensions and Deferred Benefits as Strategic Compensation. *Industrial Relations* 29(2): 263-281.
- Lazear, Edward P. and Robert L. Moore. 1988. Pensions and Turnover, in *Pensions in the U.S. Economy* eds. Zvi Bodie, John B. Shoven and David A. Wise. Chicago: University of Chicago Press, pp. 163-188.
- Loeb, Susanna, Linda Darling-Hammond and John Luczak. 2009. How Teaching Conditions Predict Teacher Turnover in California Schools. *Peabody Journal of Education* 80(3), 44-70.
- Manoli, Dayanard S. and Andrew Weber. 2011. Nonparametric Evidence on the Effects of Financial Incentives on Retirement Decisions. NBER Working Paper No. 17320.
- McGee, Josh and Marcus A. Winters. 2013. Better Pay, Fairer Pensions: Reforming Teacher Compensation. Civic Report: Manhattan Institute.
- Munnell, Alicia H. 2012. *State and Local Pensions: What Now?* Washington, DC: The Brookings Institution.

Nagler, Markus, Marc Piopiunik and Martin R. West. 2015. Weak Markets, Strong Teachers: Recession at Career Start and Teacher Effectiveness. NBER Working Paper No. 21393.

National Conference of State Legislators. 1999. Pensions and Retirement Plan Issues in 1999 State Legislatures. Policy Report.

--. 2000. Pensions and Retirement Plan Issues in 2000 State Legislatures. Policy Report.

--. 2001. Pensions and Retirement Plan Issues in 2001 State Legislatures: Policy Report.

National Council on Teacher Quality. 2012. No One Benefits: How Teacher Pension Systems are Failing Both Teachers and Taxpayers. Policy Report.

Novy-Marx, Robert and Joshua Rauh. 2014. Revenue Demands of Public Employee Pension Promises. *American Economic Journal: Economic Policy* 6(1), 193-229.

Novy-Marx, Robert and Joshua Rauh. 2011. Public Pension Promises: How Big Are They and What Are They Worth? *Journal of Finance* 66(4), 1207-1245.

Novy-Marx, Robert and Joshua Rauh. 2009. The Liabilities and Risks of State-Sponsored Pension Plans. *Journal of Economic Perspectives* 23(4), 191-210.

Sass, Tim R., Jane Hannaway, Zeyu Xu, David N. Figlio and Li Feng. 2012. Value Added of Teachers in High-Poverty Schools and Lower Poverty Schools. *Journal of Urban Economics* 72, 104-122.

Smith, Jeffrey S. and James E. West. 2014. Reduced Retirement Benefits: Should I Stay or Should I Go? Unpublished manuscript, Baylor University.

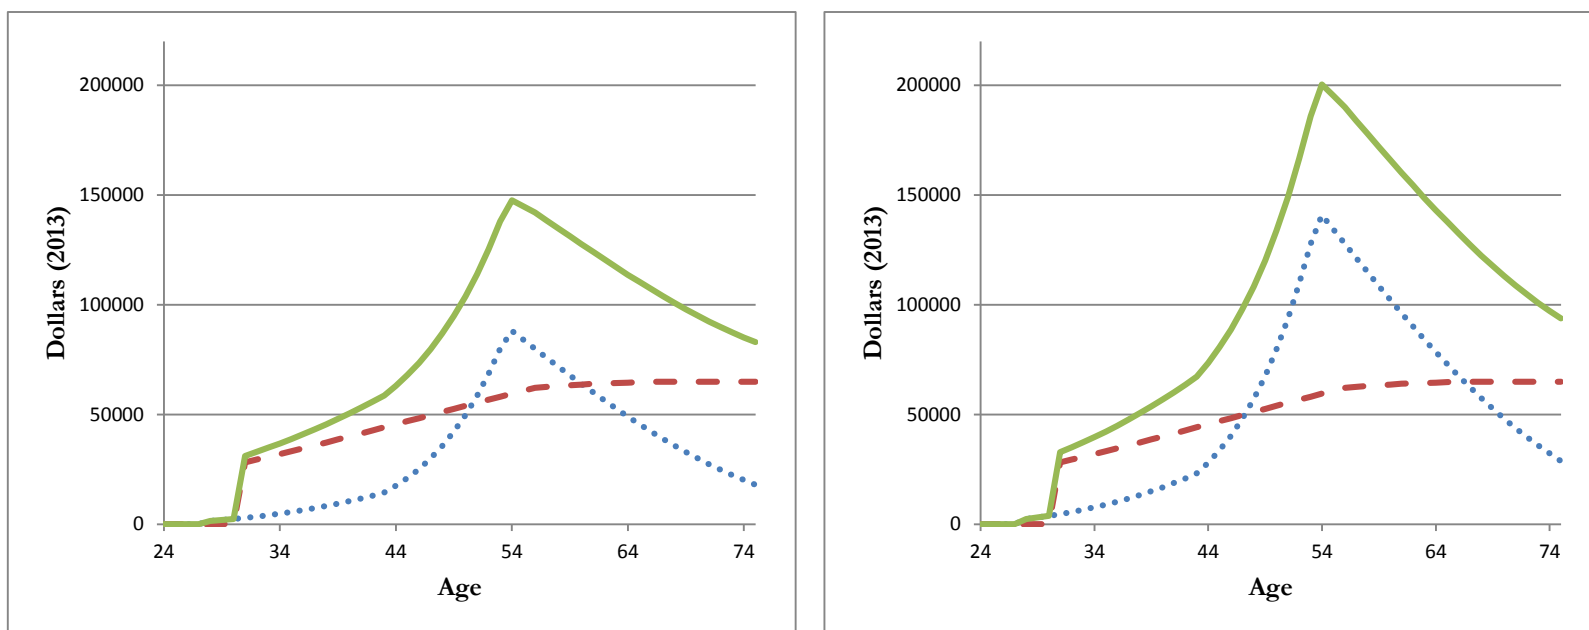
Snyder, Thomas D. and Sally A. Dillow. 2012. *Digest of Education Statistics 2011* (NCES 2012-001). National Center for Education Statistics, Institute of Education Sciences, U.S. Department of Education. Washington, DC.

Stock, James H. and David A. Wise. 1990. Pensions, the Option Value of Work, and Retirement. *Econometrica* 58(5), 1151-1180.

West, Martin R. and Matthew M. Chingos. 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.

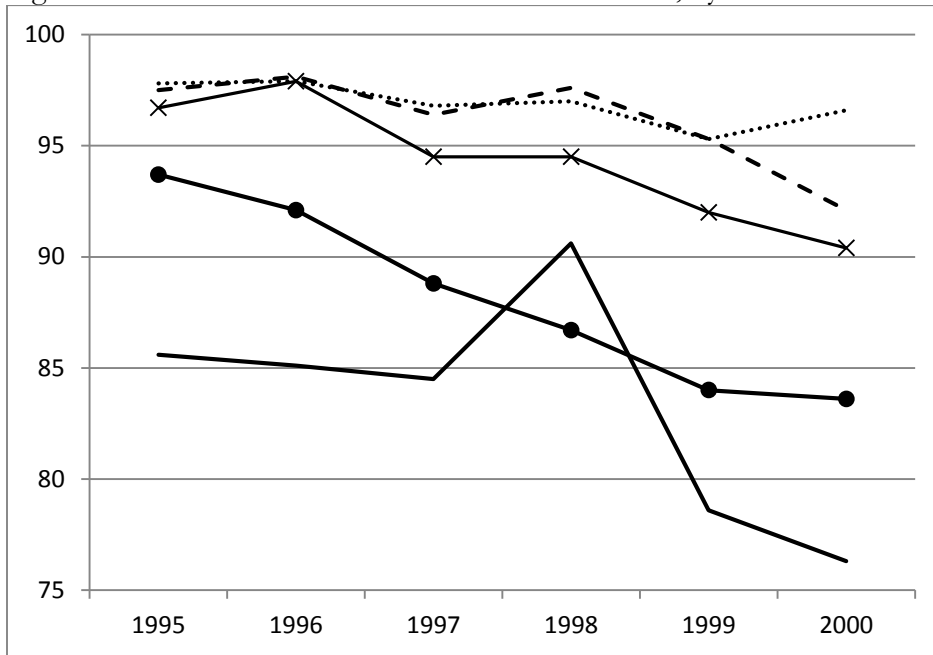
Wiatrowski, William J. 2012. The Last Private Industry Pension Plans: A Visual Essay. *Monthly Labor Review* (December).

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.



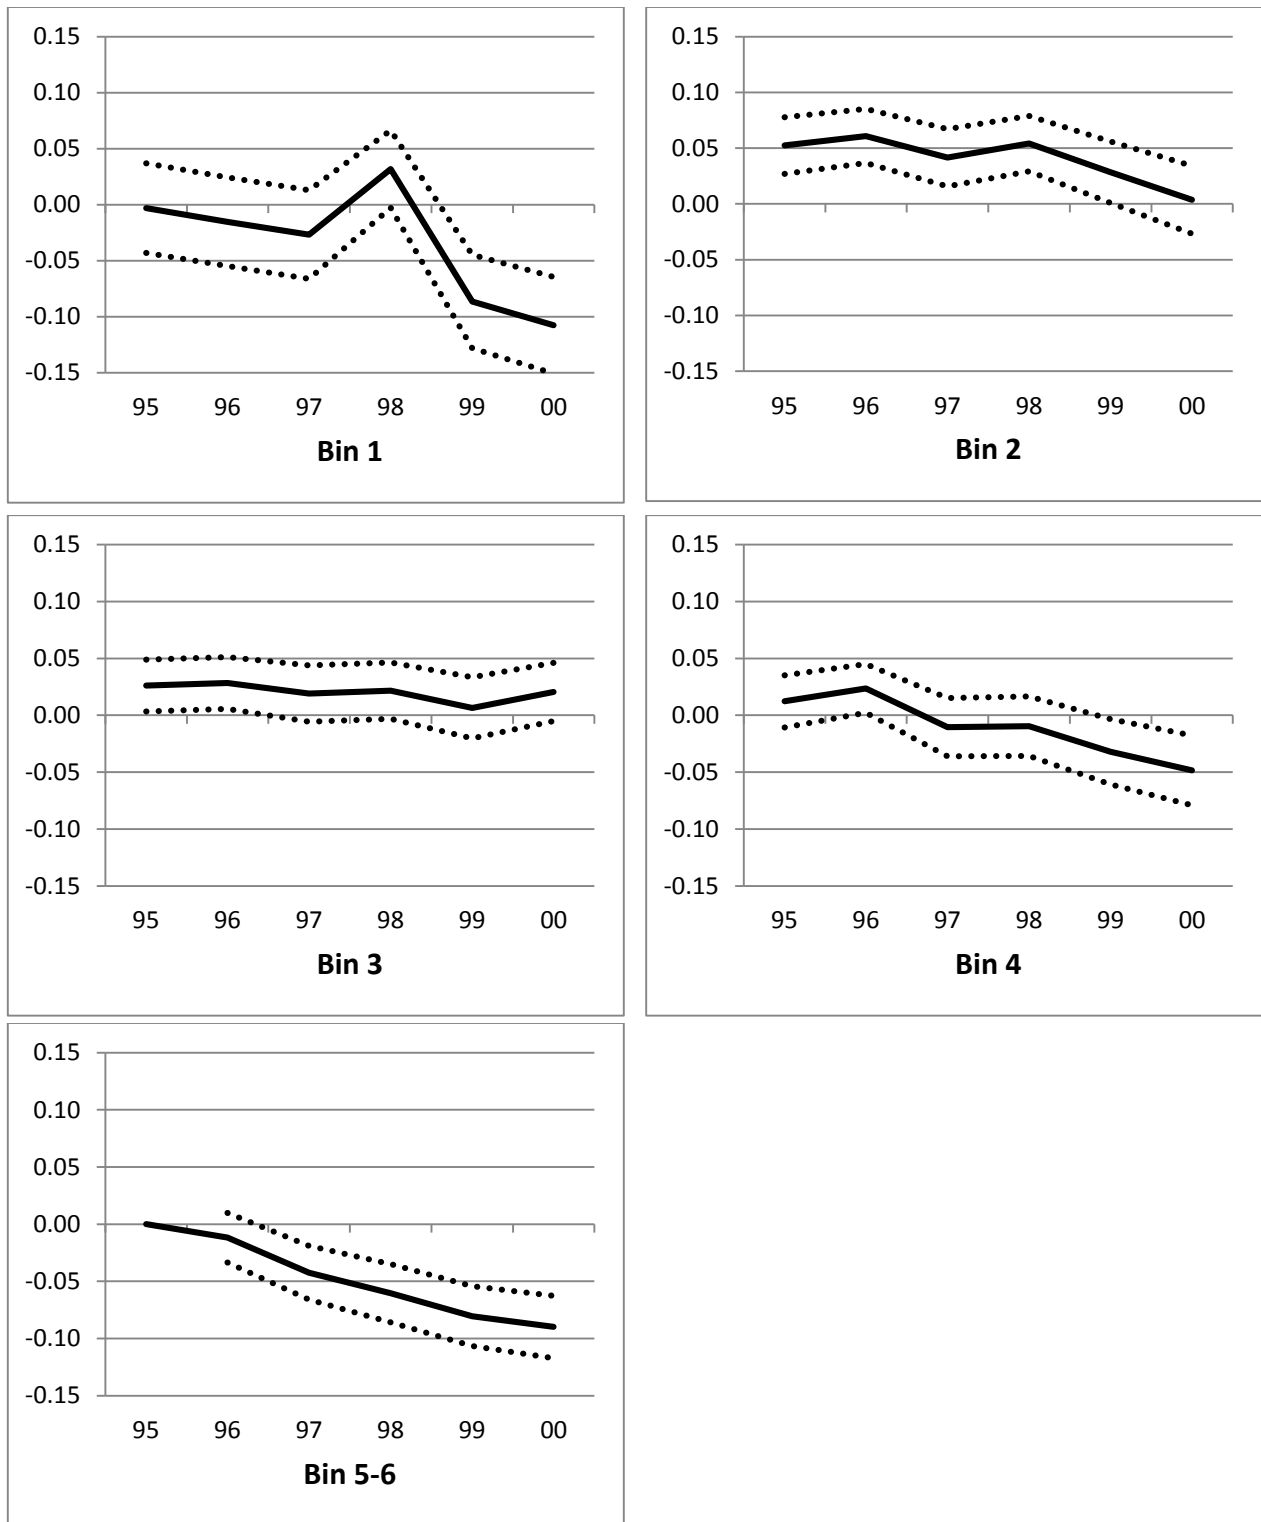
Notes: The left panel of the figure shows pre-enhancement wealth accrual and the right panel shows post-enhancement accrual. Both panels compute pension wealth holding all else equal (e.g., discount rate, salary growth, etc.) The dashed red line shows Social Security wealth accrual, which does not change 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. Annual Teacher Retention Rates in St. Louis, by Bin and Year.



Notes: The labels on the horizontal axis indicate school years by the spring year (e.g., 1996 indicates the 1995-1996 school year). Bin 1 teachers are represented by the solid line, bin 2 teachers by the dashed line, bin 3 teachers by the dotted line, bin 4 teacher by the line with x's, and bin 5-6 teachers by the line with circles.

Figure 3. Retention Trends from 1994-1995 through 1999-2000, by Bin, Relative to Teachers in Bins 5-6 in 1994-1995.



Notes: Each graph shows the retention trend relative to the omitted group (bin 5-6 teachers in 1995), with the dotted lines showing the 95-percent confidence interval. Because bin 5-6 teachers in 1995 are the comparison group, their retention rate is normalized to zero and presented without a confidence interval.

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

<u>Teacher Characteristics</u>	
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
 <i>Distance to Full Retirement Eligibility</i>	
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
 <u>School Characteristics</u>	
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			Maximum Pension Wealth			Expected Pension Wealth		
		Old	New	Difference	Old	New	Difference	Old	New	Difference
		Rules	Rules		Rules	Rules		Rules		
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 maximum wealth because the expected-pension-wealth calculations include positive probabilities of continued work past the maximum, as is observed in the real data.

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

	Age			Experience		
	25 th , 50 th , 75 th percentiles			25 th , 50 th , 75 th 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. Estimated Effects of the Pension Enhancement on Teacher Retention.

	Model 1	Model 2	Model 3	Model 4
Bin-1*1998 (retirement eligible)	0.1045 (0.0200)**	0.1018 (0.0200)**	0.0893 (0.0194)**	0.0700 (0.0329)**
Bin-2*1998	0.0517 (0.0140)**	0.0505 (0.0139)**	0.0448 (0.0139)**	0.0134 (0.0217)
Bin-3*1998	0.0434 (0.0146)**	0.0423 (0.0145)**	0.0392 (0.0145)**	0.0036 (0.0220)
Bin-4*1998 (11-15 years from eligibility)	0.0301 (0.0160)*	0.0283 (0.0159)*	0.0243 (0.0159)	0.0046 (0.0244)
Bin-1*POST (retirement eligible)	0.0013 (0.0185)	0.0008 (0.0184)	-0.0143 (0.0181)	-0.0480 (0.0496)
Bin-2*POST	0.0414 (0.0126)**	0.0394 (0.0125)**	0.0317 (0.0126)**	-0.0232 (0.0314)
Bin-3*POST	0.0619 (0.0119)**	0.0608 (0.0119)**	0.0557 (0.0120)**	-0.0064 (0.0307)
Bin-4*POST (11-15 years from eligibility)	0.0262 (0.0134)**	0.0268 (0.0133)**	0.0187 (0.0134)	-0.0161 (0.0343)
Teacher Characteristics (Race, Gender, Education Level)		X	X	X
Age Indicators			X	X
Linear Time Trends				X
R-Squared	0.0422	0.0499	0.0711	0.0722
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 5-6 teachers, who are more than 15 years away from full retirement eligibility. Standard errors are clustered at the teacher level. For Model 4, the "POST" interaction coefficient estimates are linear combinations of the individual parameter estimates for the 1999 and 2000 school years, as described in the text. 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 B 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) and 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.³³ 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

³³ 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.

year forward. The baseline weights are age-and-experience conditional exit probabilities estimated using 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).³⁴

³⁴ 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 4. We omit the coefficients on the age indicators and linear time trends for brevity.

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

	Model 4
Female	0.0101 (0.0052)*
Asian	0.0514 (0.0321)
Black	0.0452 (0.0045)**
Hispanic	0.0253 (0.0419)
American Indian	-0.0285 (0.0591)
Master Degree	-0.0028 (0.0042)
Doctoral Degree	-0.0335 (0.0266)
Other Non-Bachelor Degree	0.0156 (0.0302)
Bin-1 Indicator	-0.0064 (0.0190)
Bin-2 Indicator	0.0539 (0.0124)**
Bin-3 Indicator	0.0248 (0.0110)**
Bin-4 Indicator	0.0165 (0.0109)
Year-1998 Indicator	0.002 (0.0174)
Post-Policy Indicator	0.0072 (0.0240)
Age Indicators	X
Linear Time Trends	X
R-Squared	0.0722
N	18825

* Indicates statistical significance at the 10 percent level

** Indicates statistical significance at the 5 percent level

Notes: See notes to Table 4. The “Post-Policy” indicator reported for Model 4 is a linear combination of the year-1999 and year-2000 indicator variables as described in the text.