A joint initiative of Ludwig-Maximilians University's Center for Economic Studies and the Ifo Institute for Economic Research

CESifo Area Conference on

Economics of Education

03 - 04 September 2010 CESifo Conference Centre, Munich

Do Public School Teachers Value Their Retirement Benefits?

Maria D. Fitzpatrick

CESifo GmbH
Poschingerstr. 5
81679 Munich
Germany

Phone: +49 (0) 89 9224-1410
Fax: +49 (0) 89 9224-1409
E-mail: office@cesifo.de
Web: www.cesifo.de

Do Public School Teachers Value Their Retirement Benefits?

Maria D. Fitzpatrick

Stanford Institute for Economic Policy Research

Stanford University

This draft: August 18, 2010

Stanford Institute for Economic Policy Research, Stanford University. Mail: 366 Galvez St., Stanford, CA 94305. Email: mfitzpat@stanford.edu Phone: (650) 724-6807. Webpage: www.stanford.edu/~mfitzpat

Acknowledgements:
The author is the Searle Freedom Trust Postdoctoral Scholar at the Stanford Institute for Economic Policy Research at Stanford University. Special thanks to Cynthia Fain, Ed Maybe, Sally Sherman and Sandra Cochran for their help with the Teacher Retirement Service data. This paper has greatly benefited from the comments of Gopi Shah Goda, Damon Jones, Greg Rosston, John Shoven and Sarah E. Turner. All errors are my own.

Abstract

Teachers, and public sector employees more generally, tend to receive relatively large portions of their employment compensation in pensions and other non-salary benefits. For example, teachers in the Illinois Teacher Retirement System received $3.5 billion in pension benefits payments in 2009. At 40 percent of the year's payroll for member teachers, these benefits represent large costs to taxpayers in the form of deferred compensation. The introduction of an option for public school employees in Illinois to purchase additional retirement benefits provides an opportunity to examine how teachers value such benefits through the measurement of their demand and cost schedules. The price of the annuity was based on the teacher's salary and accumulated experience as of 1998. Because a teacher's salary is likely correlated with her annuity demand and subsequent collection for reasons other than just the salary's determination of the price she pays, I use a simulated instrumental variables framework to estimate employee demand for the extra retirement benefits relative to the costs of providing them. On average, the willingness to pay for the upgrade is almost $110,000 less per employee than the cost of the increased benefits it provides, thus ruling out a worker-preference theory as the main justification for large pensions. Relative to other forms of compensation such as current wages, other benefits and working conditions, pensions are an extremely costly and unwieldy tool for attracting high quality teachers. What is more, currently there is great concern about the inability of state and local governments to pay their promised pension benefits (Leiber 2010) and to balance their budgets more generally (Powell 2010). The finding that pension benefits at the margin are worth much less to public employees than they cost to provide suggests reducing benefits may be a more efficient policy than increasing taxes for balancing pension budgets.

1. Introduction

At a time when most private firms have switched to defined contribution retirement plans, the defined benefit pensions paid to state and municipal employees in most states have swelled to enormous magnitudes. For example, teachers in the Illinois Teacher Retirement System received $3.5 billion in pension benefits payments in 2009. At 40 percent of the year's payroll for teachers eligible for the system, these benefits represent large costs to taxpayers in the form of deferred compensation. Although economists have long been interested in these forms of compensation, neither the broad analysis of public sector employment nor the specific analysis of teacher labor markets provides theoretical and empirical justification for these pension structures.[1]

The most straightforward explanation rests on there being a group of employees with preferences for insurance against outliving one's savings. These employees are willing to pay a premium (lower wages) in exchange for a guaranteed stream of income upon retirement (defined pension benefits). If these preferences are correlated with the qualities that make public sector employees more productive (e.g. higher quality teachers), offering a defined benefit pension plan will help the employer recruit the best employees. If employees are risk averse, the reduction in current wages can be larger than the present value of the future stream of payments. A common operating assumption in this type of model is that the employees are willing to bear the incidence of the pension benefits through lower wages. This is a difficult assumption to test, because sorting of workers and other factors make it impossible to view the wages that would occur in the counterfactual with no retirement benefits.

---

[1] To date, most of the literature on teacher pensions has focused on their effects on teacher labor supply decisions (e.g. Costrell and Podursky 2009, Costrell and Podursky Forthcoming, Brown 2010).

In this study, I offer an important test of employee preferences for retirement benefits by examining an opportunity for employees of Illinois Public Schools (IPS) to upgrade their pension benefits. In 1998, the Illinois state legislature implemented a new formula for calculating pension benefits in which future years of service increased the promised annual pension benefit at a faster rate than the previous formula. Meanwhile, the legislature offered existing employees of IPS the chance to increase their accumulated pension accrual to the new, more generous level by paying a one-time fee based on their current years of service and salary.

The offer of this annuity product is the natural experiment I use to estimate teachers' valuation of pension compensation. The price of the annuity-like product was based on the teacher's salary and accumulated experience in 1998. Because a teacher's salary is likely correlated with her education, her health and other factors that are correlated with annuity demand, I use a simulated instrumental variables framework to estimate employee demand for the extra retirement benefits relative to the costs of providing them. The results clearly show that IPS employees value pension benefits at the margin much less than the cost of providing them. On average, the willingness to pay for the upgrade is almost $110,000 less per employee than the cost of the increased benefits it provides. This rules out worker preferences as the main justification for large pensions.

Other prominent theories for justifying the use of generous defined benefit pension plans in the public sector rely on union voice models where the preferences of union management differ from those of the marginal worker (Freeman 1986). If management has stronger preferences for deferred compensation, we should see differences across teachers in the valuation of this product. Using new techniques adapted from the literature on adverse selection in insurance markets, I show that the willingness to pay for increased retirement benefits varies systematically with workers' future collection of benefits. As the price of the retirement benefit increases, ceteris paribus, the employees who purchase it experience larger increases in lifetime retirement benefits on average. As in the insurance literature, the conclusion is that "adverse selection" exists -- those who gain the most from the increased retirement benefits have the highest valuations. However, since the demand curve analysis shows that few of the employees of IPS value the extra benefits at their dollars worth, it is unlikely that union-voice models, like the median voter model, fully explain current levels of public pension funding.

This research provides new information critical to public policy decision making. In the last couple of decades, 46 states have maintained defined benefit pension plans as the main provider of retirement benefits to their state employees (GAO 2007). In 2010, Illinois was on track to spend $9.6 billion from 17 of its pension funds for state and local employees.[2] In the 1990s, as high stock market returns improved the balance sheets of pension funds, many states, including California, New York and Florida, increased the generosity of retirement benefits for their employees.[3] If the public school employees in Illinois are representative of teachers and other public employees in the rest of the U.S., there was a large deadweight loss associated with these benefit increases.

What is more, current public policy debates focus on the inability of state and local governments to pay their promised pension benefits (Leiber 2010) and to balance their budgets more generally (Powell 2010). One recent Pew Center on the States study estimated projected pension obligations outweigh available pension funds by $1 trillion (2010). Policies to close this gap fall into two categories: decrease benefits or increase taxes. In 2006, over half of the defined benefit pension plans for state and local employees had benefit levels between the original levels

[2] http://www.suntimes.com/news/maxedout/2122300,CST-NWS-pension25.article (Accessed July 7, 2010)
[3] Based on search of news articles between 1990 and 2000 using LexisNexis Academic.

offered to Illinois teachers and the higher amount provided by the upgrade, while another 20 percent had benefit rates above the 2.2 percent offered after the upgrade increase.[4] The results of this study imply that benefits this generous are worth much less to public employees than they cost to provide, leaving open the question of why to use these forms of deferred compensation (or at the very least leaving open the question of why the benefits are so generous) and suggesting that the more efficient policy may be to reduce benefits. Relative to other forms of compensation such as current wages, other benefits and working conditions, pensions are an extremely costly and unwieldy tool for attracting high quality teachers.

The next section explains the pension system for employees of IPS, including the benefit formula change in 1998 and the associated upgrade opportunity. Section 3 then details the administrative data I use, which I obtained from the Illinois State Board of Education (ISBE) and the Illinois Teacher Retirement System (TRS). The simulated instruments framework and results are described in Section 4. Then Section 5 concludes with a discussion of alternative models and their importance for future research and policy.

2. Illinois Teacher Retirement Benefits and the Annuity

Legislative History

I study an opportunity presented to IPS teachers to purchase additional retirement benefits. In 1998, years of political debate by legislators in Illinois culminated in state legislation granting teachers more generous pensions. The option to purchase additional retirement benefits was the byproduct of this pension 'sweetener'. The pension sweetener itself was the result of political one-upmanship by a few legislators representing Springfield, an area

[4] As calculated from data collected by the Center for Retirement Research at Boston College accessed at http://crr.bc.edu/frequently_requested_data/state_and_local_pension_data_4.html (August 15, 2010).

where a large fraction of residents are public employees. After passing bills proposing more and more generous changes back and forth between the House and Senate, and mounting political pressure from the unions, the first pension increase bill finally passed in July of 1997, enhancing pensions for the relatively small number of employees of the State Employees Retirement System (one of the almost 20 pensions funds for public employees in the state).[5] While these pension improvements were partially paid for by the state, much of the cost was borne by the employees themselves, who forwent cost-of-living increases and generous pay for accrued sick leave. Public school employees were left out of this bill, in part because increased benefits for the larger group would have been significantly more expensive and in part because of differences in the pension systems (e.g., different formulas).

Soon after the bill passed, however, the unions representing public school employees began to agitate for similar enhancements of their pensions. Despite quickly passing in the Senate, legislation extending the benefits to teachers and public school employees stalled in the House, where representatives were reluctant to support such an unfunded mandate, which required school districts to cover approximately 40 percent of the cost of the increase. Eventually, in July 1998, just before election season, unions and legislators reached a compromise whereby the teachers would bear a larger fraction of the costs of increasing pensions than originally proposed.

The new legislation increased the rate at which pension benefits accrued for teachers' future years of service. As described in more detail below, benefits for service prior to 1998 were subject to the old rate of accrual unless a teacher chooses to pay a one-time fee to 'upgrade'. This offer to upgrade retirement benefits is the natural experiment I use to estimate teacher's valuation of pension compensation. Put simply, teachers pay a one-time fee in order to receive

[5] The State Journal-Register, Springfield IL. May 22, 1997 and July 8, 1997.

an increased stream of payments in the future. The one-time fee depends on the teacher's service at the initial offering of the upgrade, future service and salary-at-time-of-purchase; the annual increase in retirement benefits depends on the teacher's service at the time of the initial offering, future service and her end-of-career salary.

Upgrade Payment and Pricing

The following statutory formula calculates the annual retirement amounts for years of service accrued by teachers and other public school employees after 1998:

$$Benefit = \min\,(0.022 \times Experience \times Salary, 0.75 \times Salary) \qquad (1)$$

The 'end-of-career salary', *Salary*, used in the calculation is the average of creditable earnings in the highest four consecutive years of the last 10 years of creditable service. *Experience* is equal to the number of years of service credit accumulated at the time of retirement.

As equation (1) shows, the annual benefit paid to a retired employee is equal to 2.2 percent of her *Salary* for each year of service, up to a maximum of 75 percent of her S*alary*. These annual benefits are available to members of the TRS when they terminate active service with IPS and meet the following age and service requirements: age 55 with 35 years of service, age 60 with 10 years of service, or age 62 with 5 years of service.[6] After retirement, benefits are increased by three percent annually to allow for increases in the cost-of-living. The benefit formula in (1), with a constant rate of 2.2 percent, applies to all creditable service accumulated after 1998.

[6] If a teacher is at least 55 years of age and has at least 20 years of experience, she may retire, but her benefits will be discounted by 6 percent for each year below age 60 she is. In addition, an Early Retirement Option (ERO) exists, whereby members who are at least 55 years old and who have at least 20 years of service can receive their full benefit (without the actuarial discounting of 6 percent per year). During the period studied, the ERO required a one-time member contribution of 7 percent for each year under the age of 60 or each year under 35 years of service, whichever is less. It also required a one-time employer contribution of 20 percent for each year the member is under the age of 60.

For service accrued prior to 1998, the benefit formula is nonlinear. The contribution proportions are 1.67, 1.9, and 2.1 percent of *Salary* per year for the first, second and third decades of service, respectively, and 2.3 percent per year for any service beyond 30 years. Figure 1 depicts the formulas in place for service accrual before and after the formula change. The flat rate formula, shown by the solid line, applies to service after July 1, 1998. Benefits for service prior to that point are based on the formula shown by the dashed line, unless members chose to "upgrade". With the flat rate for benefits, a teacher receives the maximum benefit with 35 years of service.[7] Under the preexisting formula, the maximum benefit was reached when a teacher accumulates 38 years of experience.

To receive the upgrade for service years prior to 1998, a teacher had to pay a one-time fee equal to one percent of her salary for each year of service accrued by 1998. The relevant salary for the upgrade fee, what I will call the salary-at-time-of-purchase, is one's highest annual salary in the four years prior to the decision to upgrade. The maximum upgrade fee is 20 percent of the salary-at-time-of-purchase. For every three years of service the teacher accrues after July 1998, the number of years used in the upgrade calculation is reduced by 1 year.

For example, consider a teacher with 20 years (or more) of experience in 1998. Panel B of Figure 1 depicts the present discounted value (PDV) of the cost of the annuity to the teacher, which depends in part on when she purchases it and whether she continues employment in IPS.[8] The horizontal axis measures the number of years after 1998 that the teacher continues working. First consider the price paid by the teacher if she upgrades immediately, depicted by the dashed line in the figure. She pays 20 percent of her highest salary between 1995 and 1998, as shown

[7] Over the period varying amounts of sick leave left at the end of a teacher's career could be counted for service credit. In 1984, the amount was increased to one year. Finally, two years became the limit in 2003.

[8] For the sake of illustration, in the figure I have assumed 5 percent annual growth in wages (the 25 percentile in my sample), and a discount rate and interest rate of 3 percent.

by point A. If this teacher continues working at least 3 more years for IPS, upon retirement she receives a refund (with interest) for the post-1998 service. If she teaches for three, four or five more years before retiring, she receives refund of 1 percent of her 1998 salary, plus interest. This is shown by the drop in the present value of the price when the horizontal axis reaches 2001 (though if she did not teach for three years straight it could happen in a later year).

The solid line in Panel B of Figure 1 depicts the PDV of the price paid by the teacher if she continues working and waits to purchase until the end of her career. For every three years after 1998 this teacher continues working and waits to purchase, the purchase price goes down by one percent of her 'salary-at-time-of-purchase'. However, the salary-at-time-of-purchase used to price the upgrade is based on the salary in the four years before purchase. Wage growth, either because a teacher is moving along the salary schedule or because salary schedules increase, will cause the price of the upgrade to rise. To return to the example teacher with 20 years of experience, if she continues teaching for 6 years after July 1998 before retiring she pays an upgrade fee of 18 percent of her highest salary between 2001 and 2004, the four years before she activates the upgrade. The optimal time to purchase therefore varies across employees, depending on wage growth, interest rates and one's discount rate. Under reasonable assumptions about each of these factors and regardless of future work experience, the PDV of the price was lowest if employees with more than 20 years of experience purchased in 1998, as indicated by the figure.[9]

Finally, once the teacher works the maximum years of service (between 34 and 38 years, depending on her experience in 1998), the annuity no longer increases her pension benefit

---

[9] Teachers with less experience in 1998 who face 5 percent wage growth experience prices whose PDV is lower if they wait to purchase. This is because the salary growth does not compound as quickly to put upward pressure on the price. This is another reason I focus on teachers with 25 to 32 years of experience in 1998, since I am most likely to observe their purchase decisions in my censored data (see Section 3).

amount. To compensate teachers who work so many years the upgrade will have no value, once a teacher is actively employed in a creditable position for more than 35 years, the purchase amount of the upgrade is refunded at a rate of 25 percent per year of service. This is shown by the drop off in the present value of the price after the example teacher (with 20 years of experience in 1998) continues teaching for 15 more years.

To get a better understanding of the price of the benefit relative to its payout it helps to examine using actual data. The median salary-at-time-of-purchase of an employee with 20 years of experience in 1998 is $49,548.[10] The purchase price of the upgrade to this teacher in 1998 was $9,909 (0.01*20*49,548).[11] If her ending salary is $81,729, the median maximum salary observed for employees with 20 years of service in 1998, the initial increase in her annual retirement benefit because of the upgrade is $6,784 ([(0.022-.0167)*10+(0.022-0.019)*10]*81,729). In future years of retirement, the nominal increase in her retirement benefit is even larger because of the three percent cost of living increases. Of course, if she waits to retire, discounting should be used to compare the value of the price to the extra payment from the upgrade.

Value of the Investment

In the effort to understand the purchase decisions of employees, it is helpful to have a sense of how valuable the upgrade is or what the return on it is relative to other assets. First, I examine a back-of-the-envelope calculation, as this may be how many people evaluate

[10] As described in the next section, all dollar amounts have been converted to $2010. The salary-at-time of purchase is a teacher's highest salary in the four most recent years of employment prior to purchase. This could be the years 1995 to 1998 if a teacher was employed in 1998, later years if the teacher waited to purchase or earlier years if the teacher was not employed by IPS in 1998.
[11] I assume 1998 was her highest paid year. Recall that some of the purchase price will be refunded if the teacher continues to teach for at least three years after 1998.

investment opportunities like this annuity. Consider the teacher in the example above, who had 20 years of experience in 1998. The price she pays for the annuity is 1 percent of her salary-at-the-time-of-purchase per year of experience, or 20 percent of her salary-at-time-of-purchase. The annual benefit to her will be 8.3 percent of her end-of-career salary in each year of retirement. If she retires right away, she will earn her money back on the investment (and then some) within 3 years (even with a reasonable discount rate to take account of the fact that the retirement benefits are not paid until the future). Of course, some teachers will not retire for many years after 1998 and each annual payment will be discounted at a higher rate to account for the greater delay in payment. However, it is also the case that most teachers can expect to collect many years of payments. The average employee in my data collects payments for 14 years, making it likely that, under reasonable assumptions, the upgrade is a good investment.

More generally, one can calculate an expected rate of return on the product and compare that to other assets the teachers could have purchased (Friedman and Warshawsky 1990, Mitchell et al. 1999). For each person in my data I calculate the expected internal rate of return, or the discount rate at which the price of the upgrade equals the expected present discounted value of the extra retirement benefits. More formally, I calculate the discount rate, *r*, in the following equation for each individual *i*:

$$P_i = \sum_{j=Age1998}^{119} \frac{A * M_j * R_j}{(1+r)^j}.$$

I define the price of the annuity, P, to be its price in 1998, not accounting for purchase timing decisions or future work experience. The expected payout when the teacher is age *j* is just the increase in her annual benefit, *A*, times the probability that she survives to that age, M_j, times the probability that she has retired by that age, R_j. I use projected mortality probabilities calculated from the cohort life tables underlying the 2007 Social Security Administration 2007 Trustees

Report to create a conservative estimate of the expected payout.[12] I use retirement probabilities estimated from the sample of teachers in Illinois who retired prior to the benefit change. In this exercise, the annual payment, *A*, is based on a teacher's highest earned salary as of 1998.

In expectation, the annuity was a rather good investment for most employees. Figure 2 is a histogram showing the distribution of the expected internal rates of return (IRR) on the upgrade purchase for eligible teachers. For 99 percent of eligible teachers, the expected IRR is above 7 percent. More than half of eligible teachers had expected IRR above the 15 percent Real S&P Index Return in the 1990s (Fama and French 2001).

3. Data Description

To conduct the analysis, I use data on employees of IPS amalgamated from a variety of sources. Because the data are an administrative record of all employed service in IPS over the entire period from 1987 to 2009, I can completely characterize the employment, retirement contribution and benefit receipt experiences of every employee of IPS over the period. In what follows, I describe each data set in turn and then the process by which I combine them. The section concludes with a description of procedures for creating the information used in the analysis, including the prices and costs of the upgrade.

Illinois Teacher Service Record

The first resource for my data on IPS teachers is the Teacher Service Record (TSR) covering the years from 1987 to 2009. The TSR is a database compiled by the Illinois State Board of Education (ISBE) from school district administrators to track employment and salaries

---

[12] Thanks to Gopi Shah Goda for sharing this data with me. The SSA provides mortality probabilities under three scenarios ranging from more to less conservative assumptions. I use the intermediate set of probabilities for men here, though the results are not sensitive to the cohort information used.

of teachers and administrators in public schools throughout the state. Each observation in the TSR is an employee record for a given school year. The TSR includes the following information about employees in IPS: name of the teacher or administrative employee, a unique identifier for the employee[13], the school and district in which the teacher is employed, total compensation (as reported to the relevant retirement system), number of months employed at the position, full-time equivalent percentage of the position and the percent of time that is administrative. The data also contain information on the number of years of experience (within the district, within Illinois and out-of-state) the employee has and the highest degree held by the employee.

The compensation reported includes, among other things, extra-duty pay (coaching, clubs, etc.), vacation and sick day buyouts, bonuses, school-board-paid retirement contributions, and other compensation that the Teachers Retirement System includes in total creditable earnings. This measure of compensation data does not include the cost of employer-paid health insurance or other benefits provided by the school-board to the employee. Importantly for the current work, the compensation measure recorded in the TSR is a precise measure of creditable earnings toward the retirement system. I will use the terms salary and compensation interchangeably to refer to this measure of creditable earnings.

Teacher Retirement System

I also use data collected by the Teacher Retirement System (TRS), the pension fund for teachers and administrators in IPS, but not Chicago Public Schools (CPS).[14] The TRS has provided data on the retirement benefits paid to its members. This includes information about

[13] This identifier has been provided for the years prior to 2005. For later years, I have used employees' names and employment information to match records to the years prior to 2006.
[14] Employees of Chicago Public Schools are covered by the Chicago Teachers Pension Fund. For more information, see the Data Appendix.

the name of the benefit recipient, the size of the annual benefit payment, the timing (beginning and ending) of benefit receipt, the creditable years of service and age of the member (employee). Additionally, the data contain information about the optional 'credits' purchased by members, including retirees and the employees who have yet to begin collecting benefits. Such credits include the 2.2 upgrade purchase that I study, as well as other service credits, e.g. for time spent on maternity leave (pre-1983), time spent teaching in private schools and time spent employed at a qualifying alternative public institution.[15]

Combining the Data

The two administrative sources (the ISBE and the TRS) do not use a common identifier for teachers. Therefore, I am forced to use fuzzy matching techniques to combine the data. Both sources provided me with the names of teachers as recorded in their systems. I also have information about recorded service accrued in both systems and employer information. I use this information to match teachers across the two data sources. (See the appendix for the exact algorithm.)

Of teachers that retire and begin collecting benefits, 96 percent are missing. For those who accrue some service in the retirement system, I match 90 percent of the individuals to the TSR. The lower match rate among non-retirees is driven largely by employees with just a couple of years or less of creditable service with the TRS, likely meaning that less time was available

[15] Teachers can pay a fee to have some years of service at qualifying 'out-of-system' institutions count as service in the IPS. Qualifying institutions include private schools designated by the state and public schools outside of Illinois. These service upgrade options can also be considered annuities and may be of interest in future work. Further, at no cost, teachers can apply for years of service spent employed in Chicago Public Schools count towards years of service for calculating the TRS benefit amount. See next section.

for correcting clerical errors that would reduce alignment with my fuzzy match techniques. I include all teachers that match between the two datasets in the sample.[16]

Defining the Prices and Costs

The terms of the annuity contract depend on the amount of service a teacher had in 1998, when the upgrade option was first offered. Recall that the additional retirement benefits paid to an upgrade purchaser are directly proportional to her service at the time of the initial offering and her end-of-career salary. Similarly, the price of the extra retirement benefits depends on the employee's service as of 1998, salary at time of purchase and, in the end, her service accrued after 1998. It is therefore the teacher's service accrued in 1998 that defines the terms of the upgrade contract. Regardless of when the teacher purchases the upgrade, how long the teacher works or what her end-of-career salary is, both the price and the payout of the upgrade will be directly proportional to her years of service in 1998.

As an example, consider the difference between two identical teachers, one has 10 years of experience in IPS in 1998 and the other has 11. Assume their salaries in 1998 were such that, initially, the price of the annuity offered to them is identical. If we further assume that the two teachers have the same end-of-career salary, amount of experience at retirement and length of life, the teacher with 11 years of experience would receive more total payout from the upgrade offered to her than the teacher with 10 years of experience would from the upgrade offered her.[17]

---

[16] Alternatively, I have estimated the demand and cost curves as described below using a sample that also includes any teachers in the TSR (the school data), but not in the TRS (the non-CPS retirement data) if they were employed by any school prior to 1998. With this sample, if the employee was not recorded in the TRS as having purchased the upgrade, either because she was in the TRS and did not purchase or because her name was not found in the TRS, I assume that she did not purchase the service upgrade. The results using this larger sample suggest slightly stronger price sensitivity, but are qualitatively similar. More information about the exact matching techniques and results with different samples is available from the author upon request.

[17] In this simple example, I have abstracted from the differences in present discounted value that likely occur because the two retirees begin collecting their benefits in different calendar years (though with the same experience).

The two teachers were offered different products. Each set of teachers with the same experience is, in essence, its own market. In what follows, therefore, I present descriptive statistics and estimates separately for each level of experience in 1998.[18]

Two crucial measures for the analysis are the price of the upgrade and how much it costs to provide the extra retirement benefits purchased. The price variable is easily defined as 1 percent of the highest salary earned in the last four consecutive years of earnings prior to purchase times the number of years of service the employee had accrued by 1998, up to a maximum of 20 percent. The cost to the pension fund of a person purchasing the upgrade is equal to the PDV of the extra amount that is paid out annually for each year that a person receives the retirement benefits. Conditional on the amount of experience accrued in 1998, therefore, the total cost of an upgrade purchaser to the pension fund depends on three things: the teacher's retirement age, age at death and end-of-career salary. Defining this cost in practice, however, is less straightforward than defining price for two reasons.

First, my data are censored. Even though my sample extends 10 years after the upgrade was introduced, I do not see complete pension collection history for the large majority of employees. For example, I only see the retirement age of people who retire by 2009. Similarly, I only see an employee's end-of-career salary if she begins collecting benefits by 2009. In fact, the only employees whose total cost I actually observe are employees that began receiving benefits and died before 2009.[19]

---

[18] To be precise, I conduct analyses separately by years of experience in 1998 for teachers with between 1 and 35 years of experience. There are relatively few teachers with more than 35 years of experience. Because, as described later, I am using two stage least squares, this low number of observations leads to a weak instruments problem.

[19] I assume that any teacher who begins benefit receipt and subsequently ends benefit receipt prior to 2008 is deceased. It seems implausible that most people would end benefit receipt for any other reason. This assumption was confirmed in discussions with the administrators of the TRS.

I deal with this censoring in a couple of ways. First, in the main set of results I use the information recorded on benefit collection and salary prior to the censoring point coupled with conservative projections about mortality and retirement behavior to create a conservatively realistic projection of the present discounted value of the upgrade for each teacher. Second, in later specification checks, I assume that there are no payments to any retirees after 2009; this is equivalent to assuming that everyone in the sample dies in 2009. Next, I describe the method of calculating costs for the main sets of results. The alternative strategy is defined in a later section.

First, the nature of the annuity product is such that the price is paid up front and the benefits are paid in future years. Because these extra benefit costs accrue over time, it's necessary to discount them to their present value to facilitate the comparison between costs and prices. I use the average interest rate for U.S. Treasury bills from 1994 to 1998, 4.77 percent, as the discount factor for measuring the costs of the increased retirement benefits in 1998 terms.[20] Additionally, I inflate monetary variables to $2010, as described in the Data Appendix.

Second, for teachers who have begun collecting retirement benefits, the annual cost of the upgrade in years in which benefits were collected is the present value of the additional retirement benefits they receive as a result of having purchased the upgrade. The annual cost for these retirees after 2009, when the data are censored, is defined as the present value of the additional retirement benefits observed prior to censoring adjusted for mortality probabilities and cost of living increases. The total cost for retired teacher *i* is therefore equal to the following:

$$Cost_i = \sum_{j=Age1998}^{Age2009} \frac{A_i}{(1.0477)^j} + \sum_{j=Age2009}^{119} \frac{A_i * M_j}{(1.0177)^j}.$$

[20] I use the Treasury bill as a measure of the risk free rate of return on investments. In the same way that public pensions must use historical returns when measuring their ability to cover promised pensions, I use the rate on Treasury bills from 1994 to 1998. The results are the same when I use other measures of the risk free rate for discounting.

The first term in the equation represents the observed payments to retiree *i*, while the second term represents the expectation of future payments to retiree *i*.

Third, since I do not see the total retirement benefit for teachers who have not yet retired, I use a teacher's highest observed salary to calculate her additional retirement benefit from the upgrade purchase. I perform an adjustment similar to the one for retirees, except that I must account for the unobserved retirement decision. The present value of the total cost for someone who has not retired by 2009 is

$$Cost_i = \sum_{j=Age2010}^{119} \frac{A_i * M_j * R_j}{(1.0177)^j}.$$

The annual amount is essentially weighted by the probability one collects a payment in that particular year, which in turn depends on the probability of surviving, M, and the probability of retiring, R.

Because of space constraints, in what follows, I present the descriptive statistics and results for the teachers with 25 to 32 years of experience in 1998. I choose the employees with a least 25 years experience in 1998 because they are those for whom the censoring problem is least severe. (The "normal" retirement age for these employees is 55 with 35 years of experience.) Similarly, because employees are allowed to purchase the pension upgrade at any point prior to retirement, I may not have the full purchase behavior of the cohorts with less experience in 1998. I therefore present the results for the most senior teachers, though results for the younger teachers are qualitatively equivalent and available upon request.

Table 1 presents information on IPS employees' take-up, price, cost and retirement decisions. Each row represents employees with the amount of creditable service accrued in 1998 indicated by the row header.[21] The second column reports the fraction of retirees who have

[21] Years of experience in 1998 are rounded the next highest year.

retired by 2009, the endpoint of my data. Here, retirement is defined as the collection of benefits, though it may not be synonymous with leaving the classroom. Teachers with less experience are less likely to have retired by 2009 than teachers with more experience. For example, 84 percent of teachers with 25 years of experience had begun collecting pension benefits by 2009, while 96 percent of the teachers who had 32 years of experience had retired by 2009. This difference illustrates the censoring problem with the data.

The third column of the table reports the fraction of teachers who have purchased the annuity by 2009. The take-up rate ranges from 74 percent for teachers with 25 years of experience in 1998 to 49 percent for teachers with 32 years of experience in 1998. Recall that the upgrade is a good investment for most teachers. It is therefore not obvious why less than three-fourths of them take it up. Also, it may seem curious that the teachers with 25 years of experience are more likely to have purchased the annuity by 2009 than those with 32 years of experience. Recall that if a teacher accumulates at least 38 years of experience, there is no reason to purchase the upgrade. If teachers who had accrued more than 30 years of experience are more likely to be those that intended to stay long enough in teaching that they reach the maximum benefit, they may be less likely to purchase the upgrade. Alternatively, these older teachers may be more credit constrained or have assets that are less liquid than their younger counterparts.

Columns 4 and 5 of Table 1 illustrate the prices offered to and paid by employees, respectively, while columns 6 and 7 present information about the costs to the pension fund of providing the upgrade to the whole set of employees and to the purchasers, respectively. These columns present the first suggestive evidence that the costs of providing the upgrade exceed the revenue generated by its purchasers. For example, compare the average price of the upgrade paid by purchasing employees with 25 years of experience in 1998 to the costs to date of providing them with the extra retirement benefits associated with it. The average price was $15,680 while the average amount collected is $106,547.

4. Simulated Instruments Estimates of Demand for and Costs of Upgraded Pension Benefits

Estimation Strategy

Estimating the demand function for the upgraded retirement benefits involves specifying the underlying relationships between take-up and the prices that consumers pay for the upgrade. I assume demand, D_i (a dummy variable equal to one if a teacher purchases the upgrade), is a function of the price of the upgrade, P_i:

$$D_i = \beta_0 + \beta_1 P_i + \varepsilon_i. \tag{2}$$

Similarly, I estimate the average cost curve by using information about the relationship between the costs to the pension fund of providing the extra benefits to a purchasing employee, C_i, and the purchase price, P_i:

$$C_i = \gamma_0 + \gamma_1 P_i + \upsilon_i. \tag{3}$$

Note that while I estimate the demand curve using the entire sample of employees, I use only the employees who purchase the upgrade to estimate the average cost curve. This is because the relevant cost measure is the cost of providing benefits to people who take-up the contract at a given price, not the cost of providing benefits to everyone in the market at that price. In estimating these equations, I make no assumptions about the shapes of individuals' utility functions or other structural parameters of behavior. As long as employees' choices reveal their preferences we have sufficient information to estimate the demand and cost curves.

The price of the upgrade contract offered to a teacher is proportional to her highest annual compensation by 1998, or at the time of purchase should she choose purchase after 1998. There are two pieces of this compensation measure: the basic salary paid by the district in which a teacher is employed and the extra compensation paid to a teacher for her extra-duty activities.[22] The basic pay component of compensation is largely determined by a set schedule based on the number of years of experience and the highest level of educational attainment a teacher has reached. Salary schedules vary quite a bit from district to district. The extra-duty pay component of compensation is mostly defined by the effort of the individual teacher, e.g. whether she is willing to stay after school hours to coach a sports team.

It is likely that there are common unobserved factors driving both a teacher's compensation in 1998 and her propensity to purchase insurance. For example, research has found that individuals with more education, more financial assets and larger family incomes are more likely to purchase annuities as are healthier people, older people, women and single people (Brown 2001; Banks and Emerson 1999; Inkmann et al. forthcoming). Some of these characteristics have also been tied to individuals being more highly paid. As an example, consider educational attainment. For teachers, having a Master's Degree (MA) offers a bump up in annual salary of a few thousand dollars, on average. If having an MA is also correlated with better health and longer lifespan, then teachers with higher salaries, who face higher prices for the upgrade, may be more likely to purchase the upgrade, which provides extra insurance against living longer. In this case, the estimate of β_1 will be biased upward. A similar potential bias problem exists with the estimation of the cost equation.

[22] There is actually another component comprised of other compensation forms paid by the district, but these are generally common across all teachers in a district, so can be subsumed into the basic pay category.

This endogeneity issue leads to the employment of an instrumental variables strategy to estimate the demand and cost curves. To do so, it is necessary to find an instrument that is correlated with a teacher's salary-at-time-of-purchase, but uncorrelated with her propensity to purchase insurance or the cost of insuring her. To create the instrument, I establish in which district the teacher taught in 1998 or the last year in which she was teaching before the upgrade option was introduced. I then determine the maximum salary paid to teachers with a Bachelor's Degree in that district according to the district's salary schedule. My instrument is the average of these maximum salaries across all districts in the same region other than the district in which the teacher was teaching. A region in Illinois is a construct of the state, used to determine geographically similar areas for distribution and coordination of resources from the state. In 1998, there were 45 regions Illinois (Figure 3) each consisting of at least 8 districts. The first stage equation is therefore $P_i = \alpha_0 + \alpha_1 RegionAverageBAMaxSalary_i + \upsilon_i$.

As mentioned, a good instrument must be strongly correlated with an employee's salary, but uncorrelated with either her demand for the upgrade or the amount of extra retirement benefits she collects from it. As I show in the next section, my instrument is highly correlated with a teacher's salary. Of course, satisfaction of the exclusion restriction can never be proven. Because my instrument is measured at the region level, the underlying assumption is that teachers do not sort into regions of Illinois based on things correlated with the propensity to purchase insurance or the cost of providing insurance. According to the 2000 Decennial Census, the majority of teachers in Illinois report being the spouse of the household head, meaning residential location decisions are probably made based on their spouse's employment.[23] Combined with the fact that less than 20 percent of teachers in Illinois commute more than 30

[23] Based on the author's calculations using the IPUMS USA from Ruggles et al. (2005).

minutes to their place of employment (also from the 2000 Decennial), it seems unlikely that teachers are sorting across regions in order to manipulate their compensation.[24]

As discussed in Section 2, employees were not required to purchase the annuity in 1998. They can purchase it at any point between 1998 and when they retire. 52 percent of teachers who purchase the annuity do so at some point after the first year it was offered. When defining a measure of the annuity price a teacher faces, one option would be to use her salary at the time of purchase (or more precisely her highest in the four years before purchase). Such a measure may be problematic. A teacher's salary at her chosen time of purchase is even more likely to be correlated with unobservable characteristics driving the teacher's demand for insurance and the cost of insuring her than her salary at the time of program announcement. For this reason, I use the teacher's salary in 1998 rather than her salary at time of purchase.

Estimates of the Demand and Cost Schedules

The coefficient estimates from my two-stage least squares (2SLS) estimation in Table 2 show the demand for the upgrade decreases with the price charged. For example, teachers with 30 years of experience are 2.7 percentage points less likely to purchase the extra retirement benefits when the cost increases by $1,000. I combine the estimates in Table 2 about demand's sensitivity to price (dQ/dP) with the average take-up (Q) and price information (P) for each experience level group from Table 1 to obtain an estimate of the elasticity of demand. These estimates of the price elasticity of demand for the upgrade at the average take-up and price level are also shown in Table 2. In general, the estimates range from -0.5 and -0.75. Demand for the upgrade is fairly inelastic and does not vary systematically with experience.

[24] I have also used an "area" level measure of average maximum salary according the salary schedules as the instrument for a teacher's compensation. The area measure is a more aggregate level of geography and is even less likely to be plagued by issues of sorting. Doing so makes little difference for the estimates.

Table 3 presents the results of estimating the effects of price on the cost function shown in equation (3) with 2SLS. The coefficient estimates on the relationship between price and costs are all positive. This means the cost of providing the extra retirement benefits increases as the price paid for the benefits increases. For example, the extra retirement benefits paid to purchasers with 30 years of experience in 1998 increase by $9,422 when their purchase price increases by $1,000. The estimates range from cost increases of $7,003 to $9,648 with a $1,000 price increase, depending on teachers' experience in 1998.[25]

The coefficient estimates from Tables 2 and 3 can be used to draw estimated demand. average cost and marginal cost curves. Comparing these curves can give a sense of the cost of providing the upgrade to IPS employees relative to the value the employees place on the benefits. To be more precise, equation (2) defined the demand curve. Equation (3) combines with equation (2) in order to draw the average cost curve. In order to determine the marginal cost curve of providing the upgrade to IPS employees, I use the coefficient estimates from the estimation of (2) and (3) in the following formula:

$$MC(p) = \frac{\partial TC(p)}{\partial D(p)} = \frac{\partial (AC(p) \cdot D(p))}{\partial D(p)} = \left(\frac{\partial D(p)}{\partial p}\right)^{-1} \frac{\partial (AC(p) \cdot D(p))}{\partial p}.$$

Figure 4 presents examples of these estimated curves (and the underlying data) for IPS employees by level of experience in 1998. In the figure, the circles represent the average take-up rates for employees with a given predicted price (predicted from the first stage of the 2SLS process and rounded to the nearest $1000). The square boxes represent the average total costs for the set of purchasers indicated on the horizontal axis who purchase at a particular price. Both

[25] Looking across the coefficient estimates for employees with different amounts of experience in 1998 suggests that the relationship between costs and price is smaller for teachers with less experience in 1998. This may be the result of the censoring problem, since fewer of the teachers with less experience in 1998 have retired by the time of data collection. It could also be because the teachers with less experience in 1998 retire later, making the PDV of their payments worth less due to discounting.

the circles and squares present information on actual data and the size of a circle or square is proportional to the number of employees in the data that the object represents. The relatively flat solid lines are the demand curves as predicted using the estimates from Table 2. The steeper dashed and dotted lines represent the average and marginal cost curves, respectively, as calculated from the estimates in Tables 2 and 3. Since the data on teachers with 30 years of experience is the least affected by the censoring problem, I again focus on the estimated demand and cost curves for these teachers. Figure 5 replicates the estimated demand, average cost and marginal cost curves for the teachers with 30 years of experience in 1998.

The first and foremost noteworthy conclusion from the figure is that the estimated demand curve is almost everywhere below the estimated marginal cost curve. (The points, representing data tell the same story.) This suggests that the willingness of most teachers to pay for the upgrade is much less than the cost of providing them with the extra retirement benefits. For teachers with between 25 and 32 years of experience in 1998, the average difference between employees' willingness to pay for and the future costs of the upgrade is $109,674. In a perfectly competitive equilibrium, the price charged would be that which equates the revenue generated with the costs of provision. In other words, the perfectly competitive price would be set where the demand and average cost curves intersect. In this market for extra retirement benefits, a perfectly competitive equilibrium does not exist. What is more, since average cost is always larger than the employees' willingness to pay (as represented by the demand curve), there would be no private provision of this annuity product.

Welfare Implications of the Results

Of course this is not a private market, but one in which the government intervenes to provide the product. In order to understand the welfare implications of the government's provision of the extra retirement benefits, it is helpful to consider an example. Imagine the government decides to offer the upgrade at the price P_G. Anyone who chooses to can purchase the upgrade at this price, but only Q_G fraction of the people do.

The total revenue generated by these purchasers is $TR=P_G*Q_G$. Since the average cost per purchaser when Q_G people buy the product is C_G, the total cost of providing the insurance to these consumers is $TC=C_G*Q_G$. The difference between the total cost and total revenue is $TC-TR = (C_G-P_G)*Q_G$. Some of this extra cost is consumer surplus to the purchasing employees, but the rest of it is just dead-weight-loss. In Figure 5, the total amount of consumer surplus generated by the upgrade offered at price P_G is represented by the area of the triangle below the demand curve and above the price charged, $.5*(A-P_G)*Q_G$. The remainder of the extra cost, represented by the area of trapezoid $ABEC_G$, is dead-weight-loss.

Who is responsible for paying the cost of the providing the upgrade that exceeds the revenue generated? There are two possibilities, the employees themselves or the taxpayers. To understand how the employees themselves may bear the costs, consider a perfectly competitive labor market where firms pay workers their marginal product. In such an environment, the increase in lifetime compensation offered by the upgrade would come at the price of lower forms of other compensation (wages, benefits) earlier in the lifecycle. In this way, even though the sticker price of the extra retirement benefits is less than its full cost, the employees themselves still pay the full price. But the demand curve in Figure 5 tells us that, when given the choice, these employees are unwilling to trade-off enough current dollars to cover the cost of the retirement benefits. Either the employees are unaware of the tradeoff between current and future compensation or something blocks the link between lifetime compensation and a worker's marginal product.

The public nature of public employee pension increases makes it unlikely that public employees are unaware of the potential tradeoffs between retirement benefit increases and current wages. In particular in Illinois, a search for articles relevant to the late-1990s round of pension increases shows that newspapers often reported the need for employees to help defray the costs of the increases. Furthermore, every other change to the public pension system in Illinois in the 1990s included some payment from teachers, e.g. annual contribution increases, salary freezes, cuts in other benefits, etc. The high rate of unionization in the public sector renders it even more likely that workers understand wage trade-offs, since one role of unions is information dissemination. Since it is likely that public sector workers are aware that at least some of the incidence of pension benefit increases may fall on them yet they do not fully value the pension benefit increases, it is worth considering whether the full incidence of the pension increase falls on their shoulders.

Why might the tradeoff between current income and future income be less than one for public employees? The reason lies in the ability of public workers, particularly unionized public workers, to use political pressure to increase the demand for public goods and, hence, the demand for their own labor. The increase in labor demand can lead to increased employment, higher wages or both, depending on the elasticity of labor supply and the elasticity of residential location decisions to tax rates (Courant et al. 1979). The ability of public sector workers to increase wages in this way is amplified when public budgets are allowed to be unbalanced, as is the case with underfunded pensions (Inmann 1982, Inmann and Albright 1990).

Heterogeneity in Costs Implies Adverse Selection

Despite the finding that, on average, teachers do not value the increased pension benefits, it could still be the case that large defined pension benefit plans exist because of the preferences of employees. Consider, for example, a model in which the government must negotiate with union representation about the compensation package for public employees. If the union management has a strong preference for pension benefits, it may negotiate for larger pension benefits than is optimal from the perspective of the marginal worker. Positing that this type of difference between union management and members exists has been popular in the literature (Freeman 1986), but does not yet have much empirical support.

The nature of the upgrade allows me to test whether preference for the annuity varies systematically across teachers. A finding that teachers have different preferences over retirement benefits leaves open the possibility that preference differences between union management and members could be a driving force behind large public employee pension packages. The test is only suggestive, however, because I cannot rule out that other mechanisms (such as the political economy of the legislative budgeting process) also play a role.

In order to conduct this test, I adapt methods from the literature on adverse selection in insurance markets to determine whether teachers who gain the most from the upgrade are the most likely to purchase it (Einav, Finkelstein and Cullen 2010). The insurance nature of the upgrade combined with the fact that its price is determined by only a couple of employee characteristics implies that the upgrade's value will vary across teachers who are offered it at the same price. This is because, as discussed in Section 2, the value of the upgrade, *ceteris paribus*, will depend on a teacher's retirement date, length of life and salary trajectory between purchase and retirement. Workers vary along these three dimensions in actuality and in their expectations of their value at any point in time. Since the variation along these dimensions is not captured in the pricing of the upgrade, it may be the case that two people offered the upgrade at the same price will have different purchase propensities. A shown by Einav, Finkelstein and Cullen (2010), if the (exogenous) purchase price and costs of providing an insurance product are positively correlated, adverse selection exists.

In other words, if the people willing to pay a high price for insurance are also those who accrue the most in insurance costs, there is evidence of adverse selection. Intuitively, moving from a higher price to a lower one entices 'less risky' consumers to purchase the upgrade and the marginal cost of insuring these marginal consumers will be less than the cost of insuring the inframarginal consumers. My instrumental variables strategy implies that the price an employee is charged for the upgrade is essentially random. This implies a straightforward test: if the marginal cost curve is downward sloping, adverse selection exists or, put more broadly, there are differences in valuation of the insurance product among employees.

Referring back to the panels of Figure 4, it is clear that the marginal cost curves for the upgrade are downward sloping. As the price of the retirement benefit increases, ceteris paribus, the employees who purchase it experience larger increases in retirement benefits. As in the insurance literature, the conclusion is that adverse selection exists -- those who gain the most from the increased retirement benefits have the highest valuations. However, since the demand curve analysis shows that few of the employees of Illinois Public Schools value the extra benefits at their dollars worth, it is unlikely that seniority bias models of unions like the one previously described fully explain current levels of pension funding.

Teachers Versus Other School Employees

One may also be interested in whether the results just presented hold true for all employees of IPS or whether there are differences in employee valuation for certain groups of employees. For example, it is of interest whether teachers and administrators value retirement benefits similarly. If they do, there may be room to provide different retirement benefits to teachers and administrators.

Columns 1 and 2 of Table 4 contain the results for people who were teachers in 1998, while columns 3 and 4 contain employees who were not teachers in 1998. Employees who were not teachers are generally administrators of some type, although the category also includes school staff such as nurses and psychologists. Even though in some instances, there is not enough statistical power to render the results statistically significant, the results for both teachers and other types of public school employees follow the same general patterns already seen for the combined sample. Demand for the upgrade is negatively associated with price while the average cost of providing the extra retirement benefit is positively related to its price. Notably, there is no statistically significant difference between the price sensitivity of people who were teachers or other types of employees, nor is there a difference in the relationship between price and average costs. Finally, it is also the case that the costs of providing the benefits to purchasers who were teachers or who were other types of employees greatly dominated the valuation the employees placed on the extra benefits.

Teachers of Different Subjects

Another dimension along which we might see variation in teacher behavior is by the subject matter taught. From their training, math and science teachers may have an easier time calculating the financial rewards associated with annuity purchase. To investigate whether there are differences in behavior, I compare two subsets of my overall sample: teachers whose modal teaching assignment over the period in which they taught was i) some form of math or science class or ii) some type of English or language arts class.[26] These two groups have pretty similar propensities to purchase the upgrade. 42 percent of math and science versus 44 percent of English teachers purchase the upgrade, and the proportions who purchase show no differential patterns across the teachers' subject matter specialty.

Columns 5 and 6 of Table 4 show the instrumental variables estimates of the relationships between the price of the upgrade and it's demand and cost, respectively, for Math and Science teachers, while columns 7 and 8 do the same for English teachers. The estimates for the latter group generally follow the same pattern as those for the overall sample of all public school employees. The math and science teachers, however, appear to be somewhat less sensitive to the price of the upgrade when purchasing it than other teachers. The estimated coefficients on price using the sample of just math and science teachers are more likely to be positive and close to zero and are sometimes, though not always, statistically different from the coefficients for English teachers. Finally, the estimates of the relationship between costs and price for teachers of different subjects are statistically equivalent. Taken together, this may be considered suggestive evidence that, math and science teachers are more likely to understand the value of the upgrade, though it is still unclear why more teachers do not purchase the upgrade.

[26] The TSR data report information on a teacher's main teaching assignment and, in more recent years, her other teaching assignments as well. For consistency in the definition over time, I use only a teacher's reported main teaching assignment to classify her.

Results with Alternative Data Specifications

As discussed, the data on retirement and career decisions are censored in 2009. This means I cannot observed the full collection of benefits for the majority of teachers who had the opportunity to purchase the upgrade. In the results presented thus far, I made conservative assumptions about employee mortality and retirement behavior to impute the extra retirement benefits collected by IPS employees after 2009. In Table 5, I make an even more conservative assumption: there are no payments to any retirees after 2009. This is equivalent to assuming that everyone in the sample dies in 2009. Of course this decreases the magnitude of the estimated costs, which can be seen by comparing the estimates in Table 5 to those in Table 3.

One other potential bias stems from the fact that I use administrative data from multiple sources and the populations of the samples do not line up perfectly. The inability to match some teachers in the employment data to the retirement data may mean the results presented are biased. For example, if for some reason the people who are not matched across samples are those that value the annuity the least (at any price), my estimates of valuation would be upward biased. I have also estimated the demand and cost schedules for three different samples. First, I include all employed teachers (not just those for whom I can find retirement information) and assume the 'unmatched' teachers *did not* purchase the upgrade. Second, I include the same set of teachers just described but assume the 'unmatched' teachers *did* purchase the upgrade. Finally, I have used datasets that include the employees in Chicago Public Schools (using the above assumptions). None of the results using these various samples are qualitatively different than the ones presented thus far.

5. Conclusion

In this paper, I empirically test whether teachers and other public school employees value pension benefits at the margin as much as it costs to provide them. Almost all states use a form of defined benefit program similar to the one studied as their main pension system for public employees. By guaranteeing a large fraction of pre-retirement salary until they death, these programs offer generous deferred compensation to teachers and other public employees. Although economists have long been interested in these forms of compensation, neither the broad analysis of public sector employment nor the specific analysis of teacher labor markets provides theoretical and empirical justification for these pension structures.

The most straightforward argument for the large deferred compensation packages offered through defined benefit pension programs is that public employees prefer the guaranteed stream of income they provide to an equivalent increase in current wages. An opportunity offered to public school employees allows me to estimate the willingness-to-pay for the benefits and the cost of supplying them. I show that teachers valuation of the increased pension benefits was much less than their cost. On average, the willingness to pay for the upgrade is almost $110,000 less per employee than the cost of the increased benefits it provides. By showing that workers value the increased benefits at only a fraction of the cost of providing them, the results of this paper offer clear evidence against a worker preference rationale for generous defined benefit pension packages. If teachers were able to choose between higher pension benefits and equivalent increases in current salary, as is often assumed, they would choose higher current salaries. This evidence suggest that workers are not paying for the generous pension benefits with decreases in current wages (as is often assumed).

Another potential explanation for the use of defined benefit pension plans is that there is a subset of employees who prefer such benefits. If these employees become union management, they may negotiate more generous pension benefits than the marginal worker would prefer. Using techniques adapted from the insurance literature, I show that employees who are willing to pay the highest prices are those that collect the most in extra benefits, on average. This is akin to a finding of adverse selection and suggests there is some heterogeneity in teacher valuation. However, it does not appear the differences across teachers are enough to explain the enormous gaps between the costs and valuation of the increased pension benefits.

Currently, many researchers, policymakers and taxpayers are concerned about the ability of state and local governments to pay their promised pension benefits (Leiber 2010) and to balance their budgets more generally (Powell 2010). Policies to close this gap fall into two categories, decrease benefits or increase taxes. The results of this study imply that these benefits are worth much less to public employees than they cost to provide, leaving open the question of why to use these forms of deferred compensation (or at the very least leaving open the question of why the benefits are so generous) and suggesting that the more efficient policy may be to reduce benefits.

References

Banks, J. and C. Emmerson. 1999. UK Annuitants, Briefing Note no. 5, London: Institute for. Fiscal Studies (http://www.ifs.org.uk/pensions/annuities.pdf)

Brown, Jeffrey. 2001. "Private Pensions, Mortality Risk, and the Decision to Annuitize." *Journal of Public Economics*. Issue 82: 29-62.

Brown, Kristine M. 2010. "The Link Between Pensions and Retirement Timing: Lessons From California Teachers." Manuscript. University of Illinois.

Costrell, Robert M. and Michael Podursky. 2009. "Peaks, Cliffs, and Valleys: The Peculiar Incentives in Teacher Retirement Systems and Their Consequences for School Staffing."

Costrell, Robert M. and Michael Podursky. Forthcoming. "Distribution of Benefits in Teacher Retirement Systems and Their Implications for Mobility." *Education Finance and Policy*.

Courant, Paul, Edward M. Gramlich and Daniel L. Rubinfeld. 1979. "Public Employee Market Power and the Level of Government Spending." *The American Economic Review*. 69(5), December 1979, p 806-817.

Einva, Liron, A Finkelstein, MR Cullen. Forthcoming. "Estimating Welfare in Insurance Markets Using Variation in Prices" *Quarterly Journal of Economics.*

Fama, Eugene F. and Kenneth R. French. “The Equity Premium.” Center for Research in Security Prices, Working Paper No. 522; April 2001: 1-22, Table 2.

Friedman and Warshawsky. 1990. "The cost of annuities: Implications for saving behavior and bequests." *Quarterly Journal of Economics* 105, no. 1 (February): 135-54.

Freeman, Richard B. 1986. "Unionism Comes to the Public Sector." Journal of Economic Literature, Vol. 24, March 1986, pp. 41-86

General Accounting Office. 2007. *State and Local Government Retiree Benefits: Current Status of Benefit Structures, Protections and Fiscal Outlook for Funding Future Costs*. GAO-07-1156. September 2007.

Inkmann, Joachim, Paula Lopes and Alexander Michaelides. Forthcoming. "How Deep is the Annuity Market Participation Puzzle?", *Review of Financial Studies*.

Inman, Robert P. 1982. "Public Employee Pensions and the Local Labor Budget," Journal of Public Economics, 19(1): 49-71.

Inman, Robet P. and Albright, David, "Local Debt in a Federalist Public Economy: Teacher Pensions as a Case Study," mimeo., University of Pennsylvania, 1990.

Lieber, Ron. 2010. "Battle Looms Over Huge Costs of Public Pensions." *The New York Times*. August 6, 2010.

Mitchell, Olivia S., James Poterba, Mark Warshawsky, & Jeffrey Brown. 1999. “New Evidence on the Money’s Worth of Individual Annuities.” *American Economic Review*. December 1999: 1299-1318.

Pew Center on the States. 2010. *The Trillion Dollar Gap*. Report, February 2010, Washington D.C.

Powell, Michael. 2010 "Illinois Stops Paying Its Bills, but Can't Stop Digging Hole." *The New York Times*. July 2, 2010.

Steven Ruggles, J. Trent Alexander, Katie Genadek, Ronald Goeken, Matthew B. Schroeder, and Matthew Sobek. *Integrated Public Use Microdata Series: Version 5.0* [Machine-readable database]. Minneapolis: University of Minnesota, 2010.

Figure 1. Details of the Upgrade Opportunity

Panel A. The Teacher Retirement Benefit Formulas in Illinois, Before and After 1998

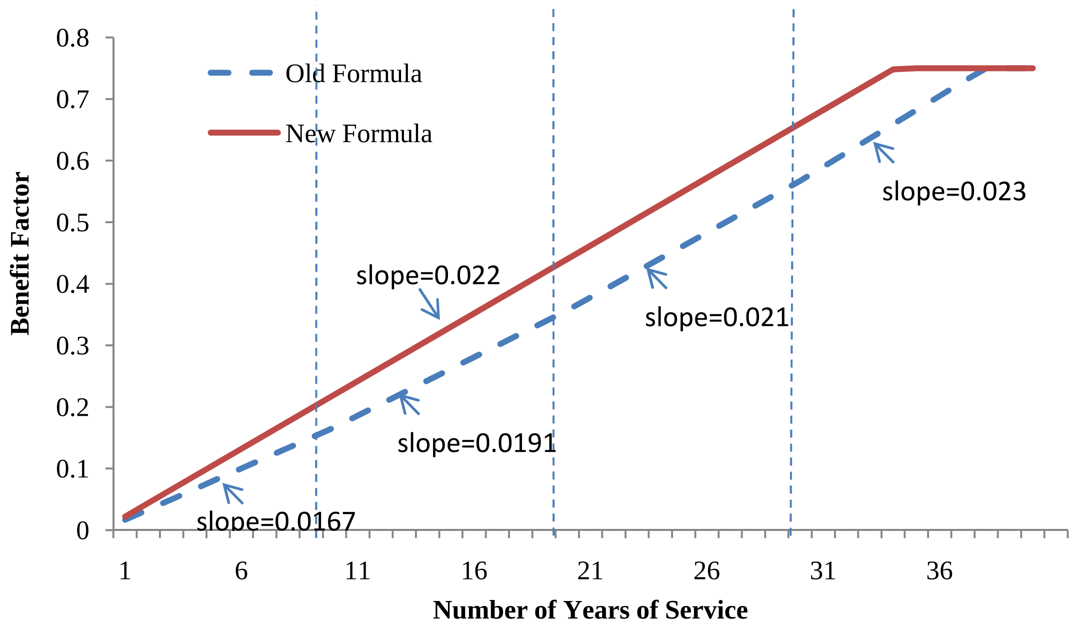

Panel B. Changes in Upgrade Price With Timing of Purchase for Teachers with 20 Years of Experience or More in 1998 Who Continue Working

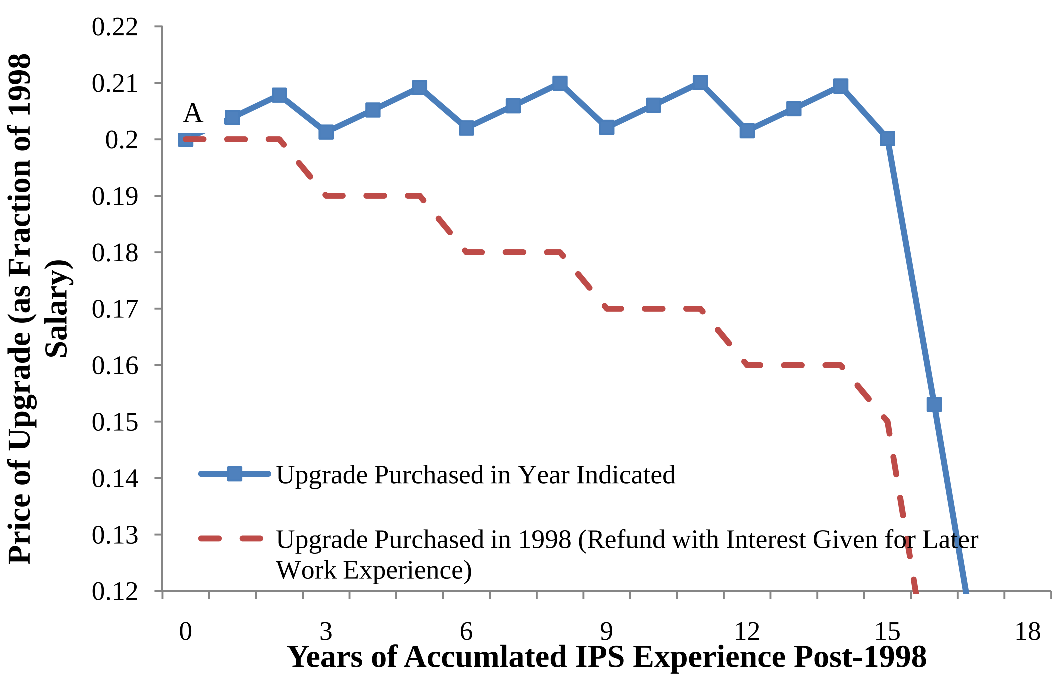

Note: In Panel A, the solid line depicts the flat rate formula used for calculating pension benefits for years of service accrued after 1998. The dashed line illustrates the graduated pension formula used before 1998. In Panel B, I assume 5 percent growth in annual salary, a discount rate of 3 percent and paid interest rate of 3 percent.

Figure 2. Distribution of the Expected Rate of Return on the Upgrade Purchase

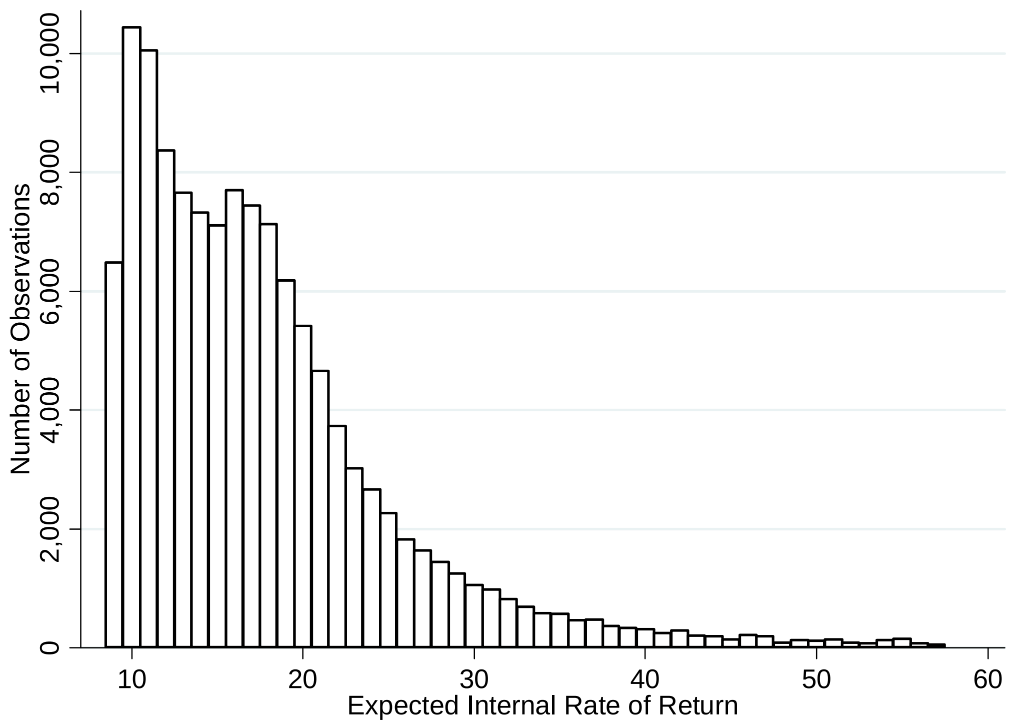

Note: Based on the author's calculations using the TRS and TSR. The expected internal rate of return (IRR) is calculated based on expected mortality and retirement behavior as of 1998. Exact methods for calculating the rate of return are described in the text. For readability, only the expected rates of return between the 1st and 99th percentiles are included in the graph. The lowest expected IRR is 3 percent and the highest is 239 percent. The 1st percentile falls at 8 percent, the median at 16 percent and the 99th percentile at 58.

Figure 3. Map Defining "Regions" and "Areas" of Illinois

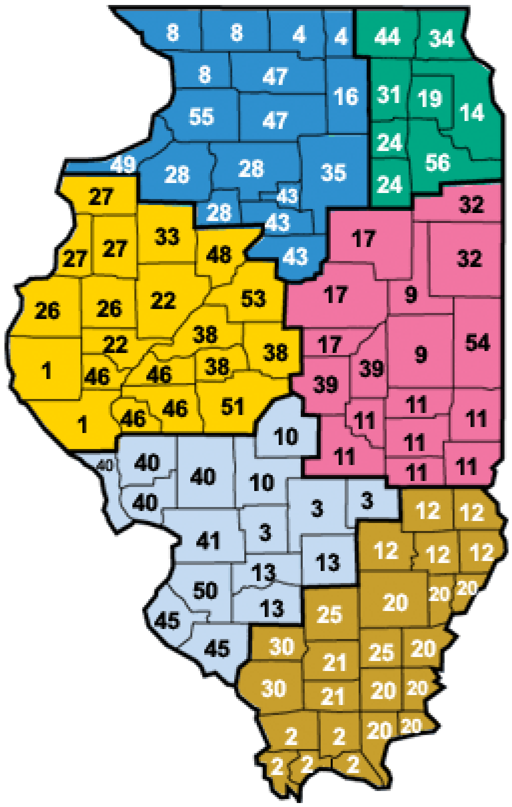

Note: Regions and Areas as defined by the Illinois State Board of Education, http://www.iarss.org/directory/state-map.html. Lightweight lines delineate county boundaries, regions are identified by number and areas are identified by heavyset lines and different colors. Chicago Public School district is not shown on the map and consists of its own region.

Figure 4. Demand and Cost Curves

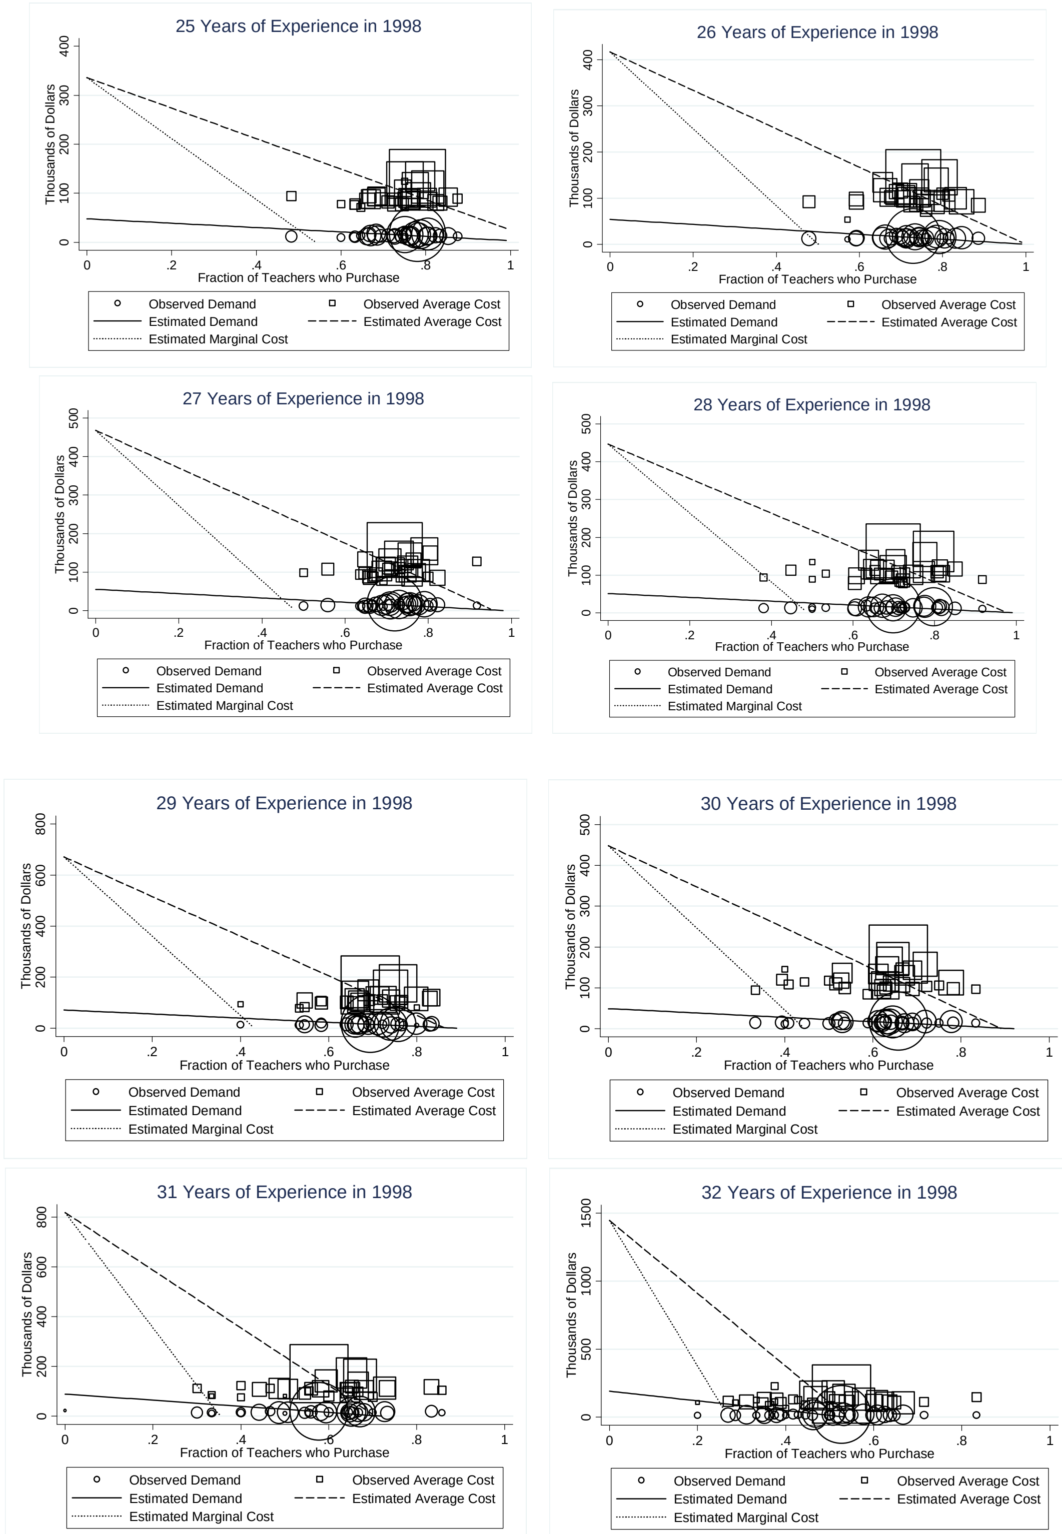

Note:

Figure 5. Estimated Demand and Cost Curves for Employees with 30 Years of Experience in 1998.

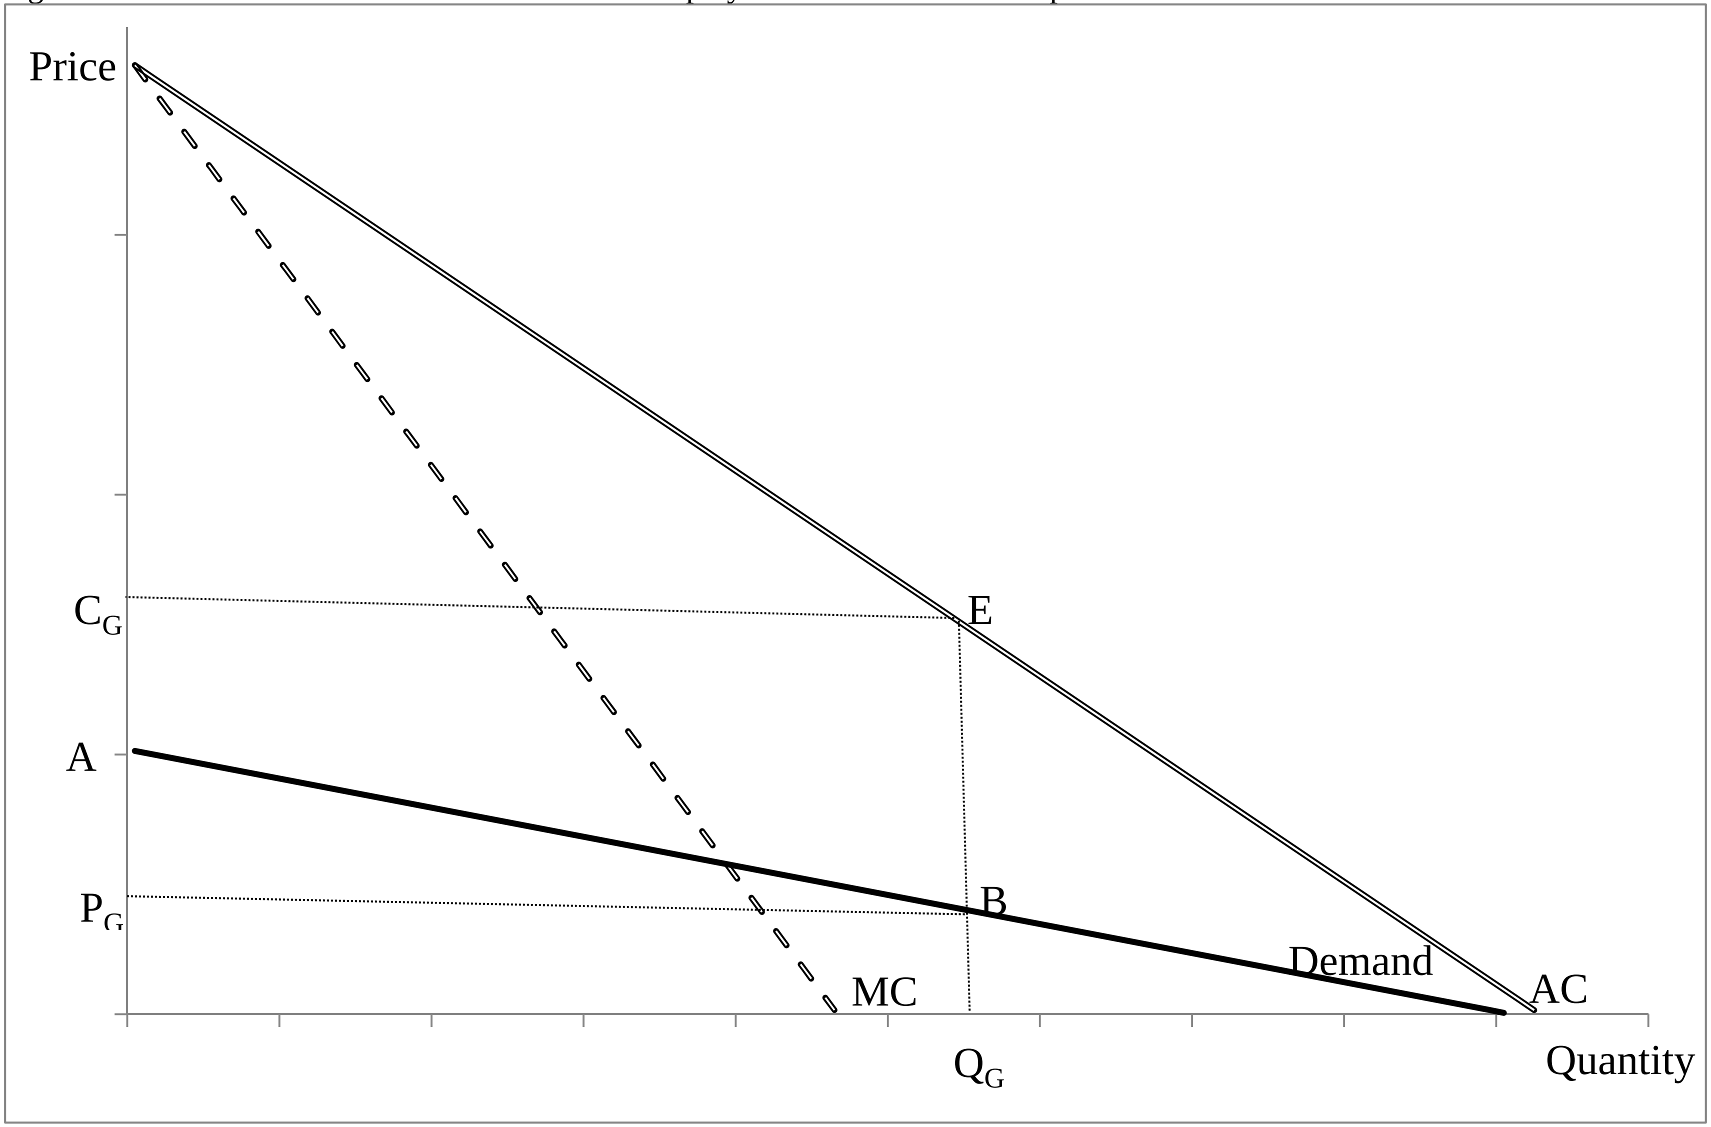

Table 1. Take-up Rates, Prices, Costs and Retirement Rates by Level of Experience in 1998

Years of Service in 1998	Fraction Who Retire by 2009	Fraction Who Purchase Upgrade by 2009	Average Price of Upgrade	Average Price of Upgrade to Purchasers	Average Cost of Providing Upgrade	Average Cost of Providing Upgrade to Purchasers	Number of Observations
25	0.84	0.74	$15,160	$15,680	$106,547	$112,433	3,272
26	0.89	0.72	$15,280	$15,670	$116,682	$121,995	3,045
27	0.91	0.70	$15,650	$16,060	$122,879	$128,614	2,790
28	0.91	0.69	$15,790	$16,290	$127,647	$135,532	2,854
29	0.93	0.69	$16,240	$16,710	$134,249	$142,073	2,835
30	0.93	0.61	$16,330	$17,010	$136,225	$147,405	2,345
31	0.94	0.60	$16,450	$16,890	$135,362	$142,477	1,741
32	0.96	0.49	$16,630	$17,440	$136,199	$145,819	1,425

Note: Based on the author's calculations using data from Illinois TRS and TSR. Years of service in 1998 is the number of creditable years of service the teacher has accrued by 1998. The fraction who retire is the fraction of the teachers with the indicated number of years of experience in 1998 who have begun collecting retirement benefits as of 2009. The fraction who purchased the upgrade is the fraction of teachers with the recorded amount of service who have purchased the upgrade by 2009. The average price of the upgrade is based on the teacher's salary and experience at the time of purchase (and is in thousands of $2010). The cost of the upgrade (in thousands of $2010) is the present value in 1998 of the extra retirement benefits paid out as of 2009 as explained in the text.

Table 2. Instrumental Variables Estimates of the Relationship Between Annuity Demand (Purchase of the Upgrade) and Price

Experience in 1998	25 years	26 years	27 years	28 years
Price of Annuity	-2.280***	-2.501**	-2.180***	-2.116***
	(0.675)	(1.040)	(0.843)	(0.684)
Price of Annuity Squared	0.072***	0.079**	0.067***	0.066***
	(0.021)	(0.033)	(0.026)	(0.021)
Constant	17.224***	18.899**	16.789***	16.213***
	(4.890)	(7.612)	(6.223)	(5.038)
Observations	3,272	3,045	2,790	2,854
First Stage F Statistic	657	643	534	552
Elasticity of Demand	-0.771	-0.501	-0.687	-0.645

Experience in 1998	29 Years	30 Years	31 Years	32 Years
Price of Annuity	-1.652***	-4.348	-1.794	-0.997**
	(0.538)	(3.165)	(1.302)	(0.491)
Price of Annuity Squared	0.050***	0.128	0.055	0.029**
	(0.016)	(0.093)	(0.040)	(0.014)
Constant	13.188***	34.416	14.086	8.212**
	(4.105)	(24.716)	(9.840)	(3.857)
Observations	2,835	2,345	1,741	1,425
First Stage F Statistic	488	356	315	240
Elasticity of Demand	-0.469	-0.761	-0.287	-0.646

Note: Based on the author's calculations using data from the TSR and TRS. Sample includes all teachers eligible for the upgrade. Each column row set represents estimates from a separate regression limited to teachers with the number of years of experience in 1998 indicated. The dependent variable is a dummy variable equal to one if the employee purchased the upgrade. Reported estimates are from 2SLS estimation where the actual price paid for the annuity is instrumented with the average maximum scheduled salary paid to teachers in one's region of employment in 1998 (both are measured in thousands of $2010). Standard errors are clustered at the district level. Standard errors are in parentheses; *** p<0.01, ** p<0.05, * p<0.1.

Table 3. Instrumental Variables Estimates of the Relationship Between Annuity Total Cost and Price

Years of Experience in 1998	25 years	26 years	27 years	28 years
Price of Annuity	7.003***	7.756***	8.689***	8.915***
	(0.366)	(0.433)	(0.549)	(0.546)
Constant	0.251	-1.789	-12.157	-11.164
	(5.548)	(6.580)	(8.407)	(8.655)
Observations	2,320	2,118	1,877	1,898
First Stage F Statistic	682	673	502	488
Years of Experience in 1998	**29 Years**	**30 Years**	**31 Years**	**32 Years**
Price of Annuity	9.648***	9.422***	9.424***	7.535***
	(0.516)	(0.589)	(0.714)	(1.149)
Constant	-21.099**	-14.970	-18.353*	12.309
	(8.245)	(9.581)	(11.062)	(20.245)
Observations	1,874	1,379	984	660
First Stage F Statistic	466	348	252	140

Note: Based on the author's calculations using data from the TSR and TRS. Sample includes all teachers eligible for the upgrade. Each column row set represents estimates from a separate regression limited to teachers with the number of years of experience in 1998 indicated. The dependent variable is a variable measuring the present value in 1998 of the extra retirement benefits paid out as of 2009 as explained in the text. Reported estimates are from 2SLS estimation where the actual price paid for the annuity is instrumented with the average maximum scheduled salary paid to teachers in one's region of employment in 1998 (both are measured in thousands of $2010). Standard errors are clustered at the district level. Standard errors are in parentheses; *** p<0.01, ** p<0.05, * p<0.1.

Table 4. Instrumental Variable Estimates of the Effects of Price on the Demand for and Average Cost of the Annuity by Employee Characteristics

Experience in 1998	Teachers		Other Employees		Math & Science Teachers		English Teachers	
	Demand	Cost	Demand	Cost	Demand	Cost	Demand	Cost
25 Years	-0.021***	6.955***	-0.031**	7.211***	0.007	8.160***	-0.053***	5.863***
	(0.005)	(0.397)	(0.014)	(0.985)	(0.011)	(1.286)	(0.013)	(0.628)
	2,665	1,887	607	433	481	339	621	433
26 Years	-0.019***	7.297***	-0.018	10.070***	0.015*	8.637***	-0.032**	6.730***
	(0.005)	(0.409)	(0.015)	(1.711)	(0.009)	(1.187)	(0.013)	(1.004)
	2,441	1,696	604	422	432	312	557	388
27 Years	-0.012**	8.899***	-0.045***	7.963***	-0.019	8.807***	-0.004	8.772***
	(0.006)	(0.533)	(0.015)	(1.612)	(0.012)	(1.253)	(0.013)	(0.805)
	2,201	1,494	589	383	446	301	541	365
28 Years	-0.014***	8.430***	-0.050***	11.185***	0.004	8.090***	-0.054***	6.675***
	(0.006)	(0.552)	(0.017)	(1.690)	(0.010)	(1.255)	(0.015)	(1.392)
	2,294	1,519	560	379	421	279	519	324
29 Years	-0.013**	9.079***	-0.007	12.519***	-0.005	9.779***	-0.007	9.047***
	(0.005)	(0.486)	(0.015)	(1.971)	(0.012)	(1.067)	(0.011)	(1.161)
	2,269	1,505	566	369	490	339	474	299
30 Years	-0.016**	9.518***	-0.028*	9.118***	-0.006	10.078***	-0.019	9.845***
	(0.007)	(0.634)	(0.016)	(1.335)	(0.013)	(1.360)	(0.016)	(1.567)
	1,871	1,096	474	283	429	263	401	221
31 Years	-0.002	9.483***	-0.048	9.127***	0.014	11.062***	-0.014	9.624***
	(0.007)	(0.785)	(0.037)	(2.345)	(0.013)	(1.777)	(0.021)	(1.565)
	1,391	805	350	179	333	194	293	176
32 Years	0.001	8.157***	-0.019	4.789	0.006	5.500**	0.004	8.093**
	(0.008)	(1.066)	(0.019)	(4.016)	(0.013)	(2.258)	(0.014)	(3.142)
	1,114	541	311	146	272	152	240	102

Note: Based on the author's calculations using data from the TSR and TRS. Sample includes all teachers indicated by the column heading. Each column row set represents estimates from a separate regression limited to teachers with the number of years of experience in 1998 indicated by the row header. The dependent variable in columns labeled *Demand* is a dummy variable equal to 1 if the person purchased the upgrade and 0 otherwise. The dependent variable in columns labeled *Cost* is a variable measuring the present value in 1998 of the extra retirement benefits paid out as of 2009 as explained in the text. Reported estimates are from 2SLS estimation where the actual price paid for the annuity is instrumented with the average maximum scheduled salary paid to teachers in one's region of employment in 1998 (both are measured in thousands of $2010). Standard errors are clustered at the district level. Standard errors are in parentheses; *** p<0.01, ** p<0.05, * p<0.1.

Table 5. Instrumental Variables Estimates of the Effects of Price on the Average Cost of the Annuity, Sample Censored in 2009

Years of Experience in 1998	25 years	26 years	27 years	28 years
Price of Annuity	2.543***	2.815***	2.869***	2.896***
	(0.199)	(0.178)	(0.260)	(0.363)
Constant	-2.240	-1.310	2.540	6.366
	(2.963)	(2.682)	(4.245)	(6.120)
Observations	2,406	2,173	1,953	1,976
First Stage F Statistic	682	673	502	488
Years of Experience in 1998	29 Years	30 Years	31 Years	32 Years
Price of Annuity	3.148***	3.191***	3.697***	2.473***
	(0.349)	(0.414)	(0.422)	(0.684)
Constant	5.575	9.507	3.190	25.62**
	(5.792)	(7.290)	(6.513)	(12.05)
Observations	1,943	1,434	1,019	696
First Stage F Statistic	466	348	252	140

Note: Based on the author's calculations using data from the TSR and TRS. Sample includes all teachers eligible for the upgrade. Each column row set represents estimates from a separate regression limited to teachers with the number of years of experience in 1998 indicated. The dependent variable is a variable measuring the present value in 1998 of the extra retirement benefits paid out as of 2009 as explained in the text. Reported estimates are from 2SLS estimation where the actual price paid for the annuity is instrumented with the average maximum salary according to the salary schedule paid to teachers in one's region of employment in 1998 (both are measured in thousands of $2010). Standard errors are clustered at the district level. Standard errors are in parentheses; *** p<0.01, ** p<0.05, * p<0.1.

Appendix A. Data Appendix

A.1. Chicago Teachers Pension Fund

Public school teachers in Chicago are covered by a separate pension system from the TRS, called the Chicago Teacher's Pension Fund (CTPF). The state legislature sets most of the rules concerning retirement contributions and benefits, so the CTPF and TRS offer identical options to employees and employees in both systems were offered the 2.2 upgrade. I have requested information from the CTPF that is identical in form to what I have received from the TRS and am assured it will be provided (by law they have to give it to me and they have been quite kind and cooperative). In the meantime, however, I proceed without this information. The implication of this is that I may be misclassifying the upgrade purchase decision for teachers who, because they retire with CTPF, purchased the upgrade with the CTPF. The TRS and CTPF suggest that a teacher collect pension benefits from the system in which they recorded the most service and apply for reciprocal service benefits from the other system. The more service a teacher had in Chicago, therefore, the more likely she is to purchase the upgrade from and retire with the CTPF. Because of this, I include only teachers who either never taught in Chicago or who have more than two years of service recorded in the TRS in my estimation sample. This excludes the most likely candidates for retirement in CTPF, for whom I may be missing information. In order to be sure the misclassification of the purchase decision is not driving the results, I repeat the estimation excluding teachers who spent any time teaching in CPS and the results are qualitatively similar.

A.2. A Note on Inflating Monetary Variables

I inflate the monetary variables, e.g. compensation, in the datasets to 2010 dollars using the Consumer Price Index for the appropriate year, e.g. for the year in which the salary was earned. Since my salary information is reported on a school-year basis, I use the inflation measure corresponding to the spring of the school year in which a salary was earned. The price of the upgrade is based on the highest salary earned in the last four consecutive years of earnings prior to purchase. For some teachers, this will be her salary in 1998, so that the price of the upgrade is inflated to 2010 dollars using the relevant CPI measure for 1998. Other teachers, however, had their last year of service much earlier. Because they had not yet begun collecting benefits in 1998, they were still eligible for the upgrade. The price charged to them was based on the highest nominal salary earned in the last four years of service before purchase. For these teachers I calculate the price of the upgrade based on this nominal salary measure and inflate it using the relevant inflation measure for 1998. For example, if a teacher's last stint of teaching before retirement in 2000 was in 1980, I use her nominal 1980 salary to price the upgrade, but inflate the price using the CPI for 1998. Finally, approximately 38 percent of teachers who purchased the upgrade did not do so immediately. The price to these teachers is inflated based on the relevant CPI measure for the year in which they purchase.