

NBER WORKING PAPER SERIES

HOW MUCH ARE PUBLIC SCHOOL TEACHERS WILLING TO PAY FOR THEIR  
RETIREMENT BENEFITS?

Maria Donovan Fitzpatrick

Working Paper 20582  
<http://www.nber.org/papers/w20582>

NATIONAL BUREAU OF ECONOMIC RESEARCH  
1050 Massachusetts Avenue  
Cambridge, MA 02138  
October 2014

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 Damon Clark, Liran Einav, Gopi Shah Goda, Caroline Hoxby, William Johnson, Damon Jones, Ed Olsen, Greg Rosston, John Shoven, Sarah Turner and many conference and seminar participants. Funding from the National Institute on Aging, through Grant Number T32-AG000186 to the National Bureau of Economic Research, is gratefully acknowledged. All errors are my own. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2014 by Maria Donovan Fitzpatrick. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

How Much Are Public School Teachers Willing to Pay for Their Retirement Benefits?

Maria Donovan Fitzpatrick

NBER Working Paper No. 20582

October 2014

JEL No. H55,H72,H75,I22,J26,J45

**ABSTRACT**

Public sector employees receive large fractions of their lifetime income in the form of deferred compensation. The introduction of the opportunity provided to Illinois public school employees to purchase additional pension benefits allows me to estimate employees' willingness-to-pay for benefits relative to the cost of providing them. The results show employees are willing to pay 20 cents on average for a dollar increase in the present value of expected retirement benefits. The findings suggest substantial inefficiency in compensation and cast doubt on the ability of deferred compensation schemes to attract employees.

Maria Donovan Fitzpatrick  
Department of Policy and Management  
Cornell University  
103 Martha Van Rensselaer Hall  
Ithaca, NY 14853  
and NBER  
[maria.d.fitzpatrick@cornell.edu](mailto:maria.d.fitzpatrick@cornell.edu)

## 1. Introduction

In 2012, relative to their counterparts in the private sector, U.S. public employers contributed nearly three times as much per hour worked to the pension benefits of their employees.<sup>1</sup> Moreover, this differential use of deferred compensation has been a key feature of the U.S. labor market for at least the past twenty five years as the private sector has trended toward the use of defined contribution rather than defined benefit plans to supplement Social Security while the public sector has engaged in making its defined benefit plans more and more generous.<sup>2</sup> Among other OECD countries, many do not have separate pension systems for public employees. Even in those that do, the differential generosity of the public pension relative to private sector pensions is rarely as large as it is in the U.S. where average replacement rates are around 45 percent and 80 percent in Social Security and public sector pensions, respectively.<sup>3</sup>

Why is so much of U.S. public sector compensation deferred? The answer to this question is central to our understanding of the operation of the labor market for public sector workers, particularly to questions involving the supply and

<sup>1</sup> Based on a comparison of compensation per hour worked of public school teachers versus professionals in the private sector using the Bureau of Labor Statistics' National Compensation Survey. <http://www.bls.gov/news.release/ecec.nr0.htm>

<sup>2</sup> Between 1980 and 2008, defined benefit coverage of private employees fell from 31 to 8 percent (Bureau of Labor Statistics 2008; Department of Labor 2002).

<sup>3</sup> Kings, Turkish and Manning (2007) discuss the generosity of public versus private pensions in OECD countries. The replacement rate measure for Social Security comes from the International Monetary Fund (2011) and the estimate for public employee pensions in the U.S. is the minimum replacement rate for career employees in a state defined benefit pension estimated by Beshears et al. (2011).

compensation of public sector workers. It also has important policy implications, since some estimates suggest that pension funds for public sector employees are \$3 trillion under-funded in present value terms (Novy-Marx and Rauh 2009). One explanation is that it reflects the preferences of public sector workers, who may have lower discount rates or may be more risk-averse than private sector workers (Gustman et al. 1994). The previous literature has tried to determine if this is the case, but the results are inconclusive (Ehrenberg 1980, Schiller and Weiss 1980, Smith 1981, Ehrenberg and Smith 1983, Montgomery and Shaw 1997). The main difficulty is that wages and pension benefits move together, making it difficult to disentangle public sector workers' valuations of each.

In this paper, I solve this problem and estimate the extent to which public sector employees are willing to pay for increased retirement benefits. To do so, I exploit a 1998 policy that allowed Illinois Public School (IPS) employees to purchase extra retirement benefits. The policy allowed teachers to choose between current dollars and increased pension benefits, thereby allowing me to estimate demand for additional pension benefits and compare it to the expected present discounted value of these additional retirement benefits. Since the price of the additional pension benefits depended on features of the system (experience and scheduled salary) that are arguably exogenous to unobserved characteristics of those choosing them, I can estimate unbiased valuations using instrumental variables strategies.

My main finding is that the majority of IPS employees value pension benefits at the margin much less than the cost of providing them. On average, these employees are willing to trade just 20 cents of current compensation for each expected dollar of future compensation. Of course, the valuation of additional pension benefits I estimate is unlikely to be the employees' average valuation of their entire pension, but this evidence suggests that, at the margin, public employees would prefer increases in current salary to increases in pension benefits of the same present discounted value. The employees' low valuation casts significant doubt that worker preferences can be used as the main justification for public sector pensions that are so much larger than those in the private sector. Instead, with some important caveats, it provides support for the political economy explanations for the use of relatively generous pension systems in the public sector (Glaeser and Ponzetto 2012, Inman 1982).

The next section describes a theoretical model of demand for the upgrade that underlies the empirical strategy in Section 3. Section 4 details the administrative data, which I obtained from the Illinois State Board of Education (ISBE) and the Illinois Teacher Retirement System (TRS). In Section 5, I present the results. Section 6 concludes.

## 2. A Theoretical Model of Demand for Additional Retirement Benefits

In 1998, the Illinois state legislature used an unanticipated surplus (relative to mandated levels) in pension fund assets to increase the rate at which pension benefits accrued for teachers' future years of service. As described in more detail below, benefits for service accrued prior to 1998 are subject to a less generous rate of accrual unless a teacher chooses to pay a fee to 'upgrade,' in which case benefits were subject to the new more generous rate of accrual. Put simply, employees were offered the opportunity to pay a one-time fee in order to receive an increased stream of payments in the future, much like a single-premium deferred annuity.

Since pensions can be considered long-life insurance, the theoretical model of demand for insurance introduced by Einav, Finkelstein and Cullen (2010) is useful for interpreting the upgrade purchasing behavior of public school employees in this setting. In the model, individuals choose between two insurance contracts. The first is a low-coverage insurance contract,  $l$ , available to the employee at no cost. Here,  $l$  is the employees' original pension. The choice employees make is of whether to pay a price,  $p$ , to purchase a high-coverage contract,  $h$ , a pension with the additional retirement benefits offered through the upgrade opportunity.

Individuals differ in their risk factors, preferences, financial resources, and expectations about retirement timing, health, future salaries, and the health of the

pension system. Define a vector,  $\zeta$ , to include these and any other characteristics of individuals that are related to demand for insurance. The distribution of these characteristics in the population is  $G(\zeta)$ . The utility to consumer  $i$  of purchasing the additional insurance is  $v^h(\zeta_i, p)$  and the utility associated with the low-coverage insurance is  $v^l(\zeta_i)$ .

Based on her characteristics, each individual chooses whether to purchase the high-coverage pension or not. An employee will purchase the extra retirement benefits if the increase in lifetime utility she gets from doing so is larger than the decreased utility from the loss in current income, i.e. if  $v^h(\zeta_i, p) > v^l(\zeta_i)$ . An individual's reservation price is defined as is the highest price she is willing to pay to purchase the additional retirement benefits,  $\pi(\zeta_i) \equiv \max\{p: v^h(\zeta_i, p) > v^l(\zeta_i)\}$ . Because consumers face a discrete choice of whether to purchase or not, the aggregate demand curve traces out the distribution of reservation prices in a population of employees,  $D(p) = \int 1(\pi(\zeta) \geq p) dG = \Pr(\pi(\zeta_i) \geq p)$ . In other words, employees who purchase at price  $p$  but do not purchase at price  $p+\varepsilon$  are those whose reservation price is  $p$ . Empirically determining willingness-to-pay for pensions is therefore a matter of estimating a demand as a function of price, controlling for the characteristics that also affect demand for long-life insurance, e.g.,  $D_i = \beta_0 + \beta_1 P_i + \gamma \zeta_i + \varepsilon_i$ .

Importantly, the model is not specific about the characteristics of employees that determine their willingness-to-pay for insurance. For example,

just as in a market for any other good or service, consumers will differ in their financial resources, or income, and these differences will be part of what determines their willingness-to-pay. Similarly, some employees will be more impatient than others, and, *ceteris paribus*, will be willing to pay less for the upgrade. Empirically, these sources of heterogeneity in willingness-to-pay will bias the estimates of demand if they are correlated with price and not adequately captured by other controls in the model,  $\zeta_i$ , i.e. if  $Cov(P_i, \varepsilon_i) \neq 0$ . The limited information available in administrative data makes it difficult to control for the full set of characteristics driving demand. Therefore, below I describe two instrumental variables strategies I employ to create measures of price that are uncorrelated with unobservable characteristics determining demand for long-life insurance.

To be clear, this theoretical model, which maps directly into the empirical framework, is not explicit about the underlying preferences, budget constraints and private information that drive demand. However, even without specific assumptions about the nature of preferences or consumer characteristics driving demand, individuals revealed preferences for the increased retirement benefits tell us about the tradeoff between current income and deferred compensation that employees are willing to make. This, in turn, provides information about the efficiency of current compensation contracts.



### 3. Estimation of Employee Demand

In this model, estimating the demand curve is straightforward using standard techniques and information on prices and coverage choices. In order to credibly identify willingness-to-pay, the variation in price must be exogenous to other factors driving demand for additional pension benefits. For example, consider a scenario where subsets of the population are randomly assigned different prices for the same upgrade opportunity. Estimating take-up rates as a function of this randomly assigned price would trace out the aggregate demand curve in the population. In the current setting, prices are not randomly assigned, but are determined based on employees' accumulated experience and salary, which are likely related to factors driving demand for additional long-life insurance. I use two different instrumental variables frameworks to isolate variation in price that is exogenous to unobservable determinants of demand.

#### *Price of the Upgrade*

To receive the upgrade for service years prior to 1998, an employee had to pay a one-time fee equal to one percent of her salary for each year of service accrued by 1998, *Exp1998*. An employee can decide to upgrade at any point between 1998 and when she retires. The relevant salary for the upgrade fee 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. The price

of the upgrade to teacher  $i$  who purchases in year  $t$  is therefore  $Price_{it} = \min\left(\frac{Exp_{1998}}{100}, \frac{20}{100}\right) \times Salary_{it}$ . Since, as I describe in Section 4, I concentrate the analysis on the employees with the most experience whose recorded data is least likely to be plagued by censoring, i.e. those with 22 to 28 years of service in 1998,  $Price_{it} = 0.2 \times Salary_{it}$ .

The optimal time to purchase varies across employees depending on wage growth, interest rates, discount rates and future employment expectations, all of which may be related to demand for additional retirement benefits. If so, estimates of demand as a function of the price at the time of purchase may be biased. What is more, defining a salary at time of purchase is difficult for employees who do not purchase.

To sidestep these issues, I define a single price faced by each employee as the price of the upgrade calculated as if she purchased in 1998,  $Price_i = 0.2 \times Salary_{i1998}$ . I do this because, under reasonable assumptions about wage growth, interest rates and discount rates and regardless of future years of work, the present discounted value (PDV) of the price was lowest for most employees if they purchased in 1998.<sup>4</sup> This is true even though employees who continue working receive a discount of one percent of end-of-career salary for every three years of

<sup>4</sup> Teachers with little experience in 1998 who face greater than 5 percent wage growth experience lower PDV prices if they wait to purchase. This is because the salary growth does not compound as quickly to put upward pressure on the price. However, these low-experience employees are not in the sample.

continued employment. As an example, consider an experienced employee who faces a price of \$14,016, the mean price in the sample, in 1998. If she continues working for three more years, the price will be 19 percent of her end-of-career salary, rather than 20 percent. However, earnings of experienced workers increase by an average of 5 percent per year. Allowing for 3 percent inflation, the PDV of the real price in 2001 as of 1998 is \$14,549. The price continues to rise the longer one delays purchase.

That the price is lowest immediately upon the upgrade's introduction is likely why 73 percent of purchasers (52 percent of the sample) purchases within the first year it is offered. Another 10 percent of purchasers buy the upgrade in the second and third years of the program. The fraction purchasing in each year continues to decline each year, reaching just one percent in 2006 and every subsequent year. For those that purchase, the actual price paid is highly correlated with the price at the time the program is introduced. (The raw correlation coefficient is 0.78; the correlation conditional on other employee characteristics is 0.70.) Given this, the use of the price offered at the time the program was introduced seems a reasonable proxy for the actual price charged.<sup>5</sup>

<sup>5</sup> I experimented with various definitions of the price faced by an employee. These included using the actual price at the time of purchase for purchasers and the lowest price faced for non-purchasers, the actual price at the time of purchase for purchasers and the average price paid of similar purchasing employees for non-purchasers, and many others. The results were consistent with those presented here. If anything, they resulted in lower measures of willingness-to-pay. Therefore I use the measure of the price an employee faced in 1998 in part because it seems like

Even the price of the upgrade upon introduction of the program depends directly on an employee's salary (in 1998). A concern is that this salary is not randomly assigned to teachers, but that a teacher's salary is determined by factors that also affect demand, i.e.  $Cov(Salary_{i1998}, \varepsilon_i) \neq 0$ . There are two dimensions along which teachers' choices determine their compensation: effort and choice of employer. Each may be related to unobservable teacher characteristics driving demand for deferred compensation.

Consider the case of teachers who take on extra duties to earn extra pay, i.e. those who exert more 'effort.' Those who do so may be those without spouses or other family members to rely on for current or pension income, making them both want to take on extra duties for increased current income and to purchase the upgrade for its extra income upon retirement. If so, ordinary least squares estimates of the relationship between price and quantity demanded will be positively biased (in part because I am unable to control for marital status), possibly producing an upward sloping demand curve. To remove this bias, I use two distinct instrumental variables techniques to estimate demand. Each has different underlying assumptions about how teachers' salaries and their demand for additional retirement benefits are related to one another, but, as I detail in Section 5, they produce quite similar results.

---

the most straightforward definition and in part because it results in the most conservative estimates.

### *Cross-District Instrument for Price*

First, I create an instrument based on the prevailing salary schedule in a teacher's district of employment before the upgrade is offered. Specifically, the instrument for the price of the upgrade to teacher  $i$  in 1998 is the beginning salary paid to a teacher with a Bachelor's Degree according to the salary schedule in the employee's district of employment in 1998.<sup>6</sup> Because the instrument is the same for all teachers within a district, this instrument is invariant to employee effort. To instill confidence that the instrument is uncorrelated with the error term in the estimation of demand, I also include measures of employee characteristics as of 1998,  $X_i$ , such as Master's Degree attainment, year of experience fixed effects, age fixed effects (and in some specifications, year of experience by year of age fixed effects), fixed effects for position in the school (e.g. teacher, administrator, staff, etc.), and subject matter taught.

If employees systematically sort across employers in ways that are related to their demand for additional retirement benefits, estimates of demand that are based on district level salary schedules will be biased. For this reason, in some specifications, I include district-level characteristics,  $X_d$ .<sup>7</sup> When district-level characteristics are not available, I use county-level or Public-Use Microdata Area

<sup>6</sup> As will be seen in the discussion of the results, I also use alternative measures based on the scheduled salary for employees with other educational attainment levels.

<sup>7</sup> These district-level characteristics come from tabulations done by the U.S. Census and provided for the National Center for Education Statistics.

(PUMA)-level data (Ruggles et al. 2010).  $X_d$  includes measures of the distribution of household income, poverty rates, house prices, unemployment rates, marriage rates, educational attainment of residents, income distribution by educational attainment, and age and demographic make-up of households.<sup>8</sup> I also include county fixed effects,  $\varphi_c$ . As such, the assumption in these specifications is that teachers with similar characteristics and family circumstances are likely to be assigned at random to different, but *observationally similar*, districts within a county and therefore the differences in prices they face are exogenous.

#### *Within-District Instrument for Price*

In the second instrumental variables strategy, I rely on within-district variation in salaries paid to teachers with different amounts of experience at the time of the upgrade to instrument for price. Specifically, rather than using county-level fixed effects, I include district-level fixed effects,  $\varphi_d$ . Since the original instrument for price, the beginning salary paid to a Bachelor's Degree (BA) teacher in one's district of employment in 1998, is constant across all employees in the same district, in this instrumental variables framework, I use an alternate order statistic from the salary schedule: the scheduled salary paid to a teacher with

<sup>8</sup> The set of market-level variables is broad and defined by what is available from the U.S. Decennial Census, Common Core of Data and the Bureau of Labor Statistics. Here, when using county or PUMA level information from the Census, it is from the Decennial 2000 Census. Similar conclusions are drawn if more aggregated data from the 1990 Census are used instead.

a BA and the same amount of experience as teacher  $i$  in her district of employment in 1998,  $BAExp1998_{id}$ .

Like the other instrument, the use of  $BAExp1998_{id}$  purges unwanted variation across employees in choices of effort, since it relies only on the scheduled salaries. With this specification, the underlying assumption is that there are no differences between teachers in the same district who have similar characteristics but had different levels of experience in 1998. In other words, there are no differences in the demand for pension benefits across cohorts of entering teachers or selection into who leaves IPS after accumulating 22 and 28 years of experience that are correlated with a district's salary schedule and not adequately captured by the included controls. Since I include statewide experience fixed effects, the underlying identification assumption is only violated if the steepness in the salary schedule within a district is directly related to differences across teachers with different levels of experience in unobservable characteristics that drive demand for additional retirement benefits. Although I cannot directly prove the assumptions underlying each of these IV strategies are not violated, I am reassured by the concordance of the results using these different identification strategies that rely on rather dissimilar identification assumptions.

### *Expected Benefit of Purchasing the Upgrade*

The fact that price depends on an employee's salary complicates the traditional demand estimation framework slightly because the value of the upgrade also depends on an employee's salary, specifically her salary at the end of her career. To help fix ideas about the value of the upgrade, Figure 1 depicts the formulas in place for service accrual before and after the formula change. For service accrued prior to 1998 (the dashed line), the statutory formula for calculating the annual retirement benefit is nonlinear and depends on the accumulation of creditable service in the system.<sup>9</sup> The contribution proportions are 1.67, 1.9, and 2.1 percent of end-of-career 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. After 1998, each year of service accrued by any teacher contributes 2.2 percent of a teacher's end-of-career salary to the annual benefit amount (the solid line).<sup>10</sup> The relevant end-of-career salary is the average of the teacher's creditable earnings in the four consecutive highest salary years of the previous 10 years of creditable service. The increase in annual retirement benefits

<sup>9</sup> Creditable service includes years employed in IPS as well as some years spent on medical or military leave and, for a fee, years of service spent in private schools. The full 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. An actuarially discounted benefit is available for retirement after age 55 with at least 20 years of experience, but there is no actuarial discounting for people who delay retirement.

<sup>10</sup> Should an employee with at least 30 years of experience choose not to upgrade, she will still receive the 2.3 rate of accrual for her remaining experience.



from purchasing the upgrade is equal to the difference between the two accrual formulas (in Figure 1, the vertical difference between the two lines at a given level of service) times a teacher's end-of-career salary. For teachers in my sample, i.e. those with 22 to 28 years of experience in 1998, this difference is between 8.5 and 9.1 percent of end-of-career salary.<sup>11</sup>

If an employee's salary in 1998 and her salary at the end of her career are correlated, estimating demand as a function of price alone will produce biased estimates of the relationship between price and take-up. I therefore estimate demand,  $D_i$  (a dummy variable equal to one if a teacher purchases the upgrade), as a function of the initial price at which she is offered the upgrade,  $P_i$ , and the expected value to her of purchasing the upgrade,  $B_i$ :

$$D_i = \beta_0 + \beta_1 P_i + \beta_2 B_i + \varepsilon_i. \quad (1)$$

<sup>11</sup> The maximum retirement benefit is 75 percent of one's end-of-career salary. In Figure 1, this is indicated by the flattening out of the two benefit formula lines at the benefit factor of 0.75. For employees of IPS who expect to continue working until they reach this maximum benefit, purchasing the upgrade would provide no additional benefit. Twenty-two percent of employees do not purchase the upgrade. One may be concerned that non-purchasing employees expected to work so long that the upgrade would be of no benefit to them. However, before the upgrade was introduced, employees reached this maximum benefit point with 38 years of experience so only a teacher who works at least 38 years would find the upgrade worthless. In order for the investment to have no expected return, one would have to expect to work 38 years with at least a probability of 0.84. Just 3 percent of IPS employees retiring before 1998 did so with 38 years of experience or more. Also, of employees who did not purchase the upgrade, just ½ of one percent reached 38 years of service. Therefore, the population reaching the maximum benefit is not large enough to explain the take-up rates (and their variation with prices) seen later in the paper.

<sup>12</sup> This is akin to controlling for the product's characteristics in a hedonic framework. Because I am estimating 2SLS models, I lay out the empirical model as a linear probability model where demand is a linear function of price. Other specifications give qualitatively similar results.

This benefit to an employee of purchasing the upgrade is equal to the present discounted value (PDV) of the extra amount that is paid out annually for each year that she receives retirement benefits.<sup>13</sup> IPS employees make their purchase decisions without observing their actual payout of benefits. I therefore define the benefit to the employee to be the average expected PDV of the extra retirement benefits for an employee of her same age and experience using information on highest expected salary and expected retirement and mortality behavior.<sup>14</sup>

These extra benefits accrue over time, so it is necessary to discount them to their present value in 1998. I use the true interest cost rate for Illinois state bonds sold in 1998, 5.10 percent per year, as the discount factor for measuring the value of the increased retirement benefits in 1998 terms.<sup>15</sup> TRS increases benefits by 3 percent per year, to keep pace with inflation, and I include these cost-of-living increases in my calculation of the expected benefit of purchasing the upgrade. The expected present value of the total benefit of the upgrade to an employee is therefore

$$B_i = \sum_{j=Age1998}^{119} \frac{(1.03)^j A_i * M_j * R_{ij}}{(1.051)^j}.$$

<sup>13</sup> In calculating  $B_i$ , I exclude any survivor benefits. About 10 percent of benefit recipients are spouses and survivors who collect benefits after an employee is deceased. This results in my understating the actual benefit of purchasing the upgrade.

<sup>14</sup> To calculate highest expected salary I use data on salaries from 1987 to 1998 to estimate wage growth from one year to the next and use these estimates to predict wage growth in each future year. Mortality probabilities are calculated from the Social Security Administration life tables for women (since 80 percent of teachers are women).

<sup>15</sup> Illinois Economic and Fiscal Commission (1998). The choice of discount rate does not affect the qualitative results presented.

The annual additional benefit to teacher  $i$ ,  $A_i$ , is essentially weighted by the probability the employee collects a payment in that particular year, which in turn depends on her cohort's probability of surviving until year  $j$ ,  $M_j$ , and retiring before year  $j$ ,  $R_{ij}$ .<sup>16</sup>

For reasons similar to those described for the endogeneity of prices, I also treat the expected benefit from purchasing,  $B_i$ , as endogenous. To remove potential endogeneity in end-of-career salary across employees, I create a simulated instrumental variable,  $BMAMax_d^{Exp, Age}$ . I use the maximum salary paid to a teacher with an MA degree in district  $d$  in 1998 to create an expected value of benefits,  $BMAMax$ . The instrument is the same value for all employees within a district that are the same age and have accrued the same amount of experience as of 1998. The maximum scheduled salary for MA teachers in the same district as a teacher is a good measure to use to calculate expected future benefits because most teachers i) have MA degrees and ii) reach the maximum scheduled salary before retiring. The resulting set of estimation equations is:

$$D_{idc} = \beta_0 + \beta_1 P_{idc} + \beta_2 B_{idc} + \theta X_{idc} + \pi X_{cd} + \varphi_c + \varepsilon_{idc}, \quad (2)$$

$$P_{idc} = a_0 + a_1 BABegSalary1998_d + \alpha_2 BMAMax_d^{Exp, Age} + \mu X_{idc} + \kappa X_{dc} + \vartheta_c + \omega_{idc}, \text{ and} \quad (3)$$

<sup>16</sup> The retirement probabilities in each future school-year are calculated conditional on age and experience using actual data on retirement behavior from before the upgrade was offered. These observed retirement probabilities take into account the likelihood of both early (and late) retirement behavior as it is observed in the data.

$$B_{idc} = \delta_0 + \delta_1 BABegSalary1998_d + \delta_2 BMAMax_d^{Exp, Age} + \tau X_{idc} + \omega X_{dc} + \rho_c + \epsilon_{idc}. \quad (4)$$

The set of estimation equations for the within-district IV strategy is:

$$D_{idc} = \beta_0 + \beta_1 P_{idc} + \beta_2 B_{idc} + \theta X_{idc} + \varphi_d + \epsilon_{idc}, \quad (5)$$

$$P_{idc} = a_0 + a_1 BAExp1998_{id} + a_2 BMAMax_d^{Exp, Age} + \mu X_{idc} + \vartheta_d + \omega_{idc}, \text{ and} \quad (6)$$

$$B_{idc} = \delta_0 + \delta_1 BAExp1998_{id} + \delta_2 BMAMax_d^{Exp, Age} + \tau X_{idc} + \rho_d + \epsilon_{idc}. \quad (7)$$

Note that even with district fixed effects, the coefficients on  $B_i$  and  $BMAMax_d^{Exp, Age}$  are identified because there is within district variation in these variables across employees with different combinations of age and experience in 1998.

### *Expected Costs of the Upgrade*

Estimating employees' willingness-to-pay for the benefits alone does not allow me determine whether employees' valuation of the upgrade exceeds its costs. To determine this, I also need information on the expected costs to the pension fund of providing the upgrade. If this were a market for a traditional good or service, costs would be determined by production technologies. Because this is an insurance product, the average cost curve is determined by the set of purchasers at a given price. (In this setting, the supplier is the government, who offers to sell the higher level of retirement benefits to any employee who wishes

to purchase at a set price.) I implicitly assume there are no administrative or other additional costs involved in providing the extra retirement benefits.

I estimate the aggregate average cost curve by using information about the relationship between the expected costs to the pension fund of providing the upgrade to a purchasing employee,  $C_i$ , and the price at which the employee is offered the upgrade,  $P_i$ :

$$C_i = \gamma_0 + \gamma_1 P_i + u_i. \quad (8)$$

This expected cost to the pension fund,  $C_i$ , is equal to the extra benefit collected by the employee who purchases the upgrade,  $B_i$ . In estimating equation (8), I include employee and district characteristics as controls and use information on beginning scheduled salaries to instrument for the price of the upgrade. Because  $C_i$  and  $P_i$  are both functions of employees' salaries, there may be a mechanical relationship between the two. That is, employees in districts with more generous salary schedules will face higher prices and benefits from purchasing than their counterparts in districts with lower salaries. To isolate the true behavioral relationship between  $C_i$  and  $P_i$ , I also include a measure of maximum salaries paid to MA degree teachers in the district when estimating equation (8).<sup>17</sup>

<sup>17</sup> The results are qualitatively similar when I include do not include this measure. They are also similar when I include alternative measures to control for the mechanical relationship between price and costs, e.g. *BMAMax*.

#### **4. Data Description**

To conduct the analysis, I use data on employees of IPS amalgamated from two sources, the Teacher Service Record (TSR) and the TRS. Because the data are from administrative records of all employed service in IPS over a period of more than 20 years, I can completely characterize the employment, retirement contribution and benefit receipt experiences of every employee of IPS over the period. Next, I describe each data set in turn and then the process by which I combine them.

The first resource for my data on IPS teachers is the 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 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, the school and district of employment, total compensation, number of months employed, full-time equivalent percentage of the position.<sup>18</sup> The data also contain information on the

<sup>18</sup> The reported compensation 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 TRS 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

number of years of experience (within the district, within Illinois and out-of-state) and the highest degree held by the employee.

I also use data collected by the TRS, the pension fund for teachers and administrators in IPS, but not Chicago Public Schools (CPS).<sup>19</sup> The TRS data contains information on the retirement benefits paid to its members. This includes information about the name of the benefit recipient, the size of the annual benefit payment, the timing of benefit receipt, the creditable years of service and age of the employee, and upgrade purchase timing and price.

The two administrative sources (the ISBE and the TRS) do not use a common identifier for teachers. Therefore, I 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 in both systems about recorded service accrued and employer. Using this, I find matches for 97 percent of

---

system. I will use the terms salary and compensation interchangeably to refer to this measure of creditable earnings.

<sup>19</sup> Employees of CPS are covered by the Chicago Teachers 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. Similar information to that provided by TRS is not available from the CTPF. 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.

teachers in the TRS data who were eligible for the upgrade. I include all employees that match between the two datasets in the sample.

As mentioned earlier, I limit the sample to employees with 22 to 28 years of experience in 1998. I exclude more senior cohorts because members of those cohorts had already begun retiring when the upgrade was introduced and the remaining employees are therefore a selected sample. On the other hand, I exclude less senior cohorts because teachers have until retirement to purchase an upgrade. Therefore, because my data are censored in 2009, I may not observe the true purchase decisions of cohorts with less experience.

Table 1 presents information on upgrade take-up, prices and expected benefits for employees in the sample. (Note that all prices and benefits have been appropriately inflated to 2010 dollars.) Each row contains statistics on these measures for employees with the amount of creditable service accrued in 1998 indicated by the row header. The second column reports the fraction of retirees who have retired by 2009, the endpoint of my sample. Here, retirement is defined as the collection of benefits, though it may not be synonymous with leaving the classroom because retirees can work part-time. As expected, employees with less experience are less likely to have retired by 2009 than teachers with more experience.

The third column of the table reports the fraction of teachers who have purchased the annuity by 2009. The take-up rate ranges from 70 percent for



teachers with 22 years of experience in 1998 to 78 percent for teachers with 25 years of experience in 1998. In columns (4) and (5) of Table 2, I present summary statistics on the prices of and additional benefits provided by purchasing the upgrade, respectively. The average price of the upgrade offered to employees with 25 years of experience in 1998 was \$15,264 while the average expected benefit of purchasing was \$96,943, which means the average ratio of expected benefits to price was 6.37 (column (6)). The average prices, benefits and ratio of benefits to price are similar for employees with other levels of experience in 1998. Importantly for identification purposes, the standard deviations in columns (4) and (5) indicate there is variation in both the price of and expected benefit from purchasing the upgrade.

## 5. Results

Standard economic theory predicts demand should be negatively related to price. Table 2 presents both ordinary least squares (OLS) and two-stage least squares (2SLS) estimates of the effect of price on take-up. The OLS results in the first column of the table show a positive relationship between price and demand, suggesting an upward sloping demand curve.<sup>20</sup> However, because the prices used

<sup>20</sup> The OLS results suggest that employees are 0.9 percentage points *more* likely to purchase the upgrade when the price is \$1,000 higher. This positive relationship exists even though all of the available controls for individual characteristics like age and employee position are in the regression, which suggests the positive relationship is not driven by older employees with higher salaries or the administrative employees who also have higher salaries. Also, all of the market- and

in the OLS estimation are not exogenous, they do not reveal employees' underlying demand for the upgrade.

To identify employees' true willingness-to-pay for extra retirement benefits, I conduct the two 2SLS estimation strategies described in the previous section. In the first three columns of Panel A of Table 2, the instruments are as specified in equations (3) and (4) – the beginning salary paid to BA degree holding teachers in 1998 and the expected benefit to purchasing calculated using the maximum scheduled salary for a MA degree teacher in employee's district of employment. Column (2) presents the estimated coefficients on  $P_i$  and  $B_i$  using a parsimonious specification that includes controls for individual characteristics and county fixed effects. In columns (3) and (4), I add in district level controls and year-of-age-by-year-of-experience fixed effects, respectively. Inclusion of the latter captures nonlinearities in the value of purchasing the upgrade across people of different ages and with different accumulated experience, two characteristics connected to the value of the upgrade. To be sure the estimated effect is not an artifact of my choice to use the BA teachers' salary schedule in estimation, in Panels B and C of Table 2, I repeat the estimation procedures just described using

---

district-level controls are included in Column (1), including county fixed effects. Their inclusion changes the OLS estimated relationship between price and take-up hardly at all, suggesting public school employees are not sorting across markets based on their propensity for take-up.

the beginning salaries paid to MA degree and PhD degree holding teachers, respectively.<sup>21</sup>

The 2SLS estimates accord with economic theory: demand decreases as price increases. All of the estimates of the effect of price on take-up in Columns (2), (3) and (4) are statistically significant at the one percent level. The estimates suggest that a \$1,000 increase in price leads to a decrease in take-up of between 6.5 and 9.2 percentage points. The estimates change little with the inclusion of district-level controls or age-by-experience fixed effects and are statistically indistinguishable from one another. The estimated relationship between price and take-up varies little using the preferred specification with the full set of individual, district and market characteristics.

All else equal, an increase in the expected payout from the upgrade should increase take-up. The coefficient estimates in Table 2 suggest that a \$1,000 increase in expected benefit from purchasing the upgrade has the effect of increasing take-up by between 1.1 and 1.6 percentage points. In other words, employee demand is less sensitive to a \$1,000 increase in the payout from purchasing than in the same sized decrease in price. Dividing the coefficient on price by that on benefits give a sense of how relatively sensitive these employees are to benefits versus price. This ratio of coefficients is between 0.17 and 0.18,

<sup>21</sup> Some districts do not report having salary schedules specific to teachers with PhDs. For these districts I use the salary schedule for MA degree holders as the schedule for PhD holders.

suggesting employees are willing to trade just 17 or 18 cents of current compensation for one dollar of future compensation.

Estimates of the aggregate demand curve using the instrumental variable strategy that relies on the within-district variation in the scheduled salary across teachers with different amounts of experience are in columns (5) and (6). The latter reports coefficients when I include experience-by-age fixed effects, which are not in the former. Again the three panels present estimates using salaries for BA, MA and PhD teachers separately. The estimated effect of a thousand dollar increase in the price of the upgrade is to decrease take-up by 11.5 to 17.1 percentage points. These point estimates are statistically similar to those in the previous columns, but the inclusion of district fixed effects causes renders some of the estimates statistically insignificant.

The coefficients on  $B_i$  in columns (5) and (6) are essentially zero. Since the salary used to calculate  $B_i$  is the same for all employees of a district and these specifications include district fixed effects, the remaining variation in  $B_i$  comes from differences in take-up of people of different age and experience levels within a district. But, I include statewide age, experience and, sometimes, age-by-experience fixed effects in these specifications. These fixed effects capture average differences in behavior of employees at different points in the life and career cycles. Therefore, the fact that  $\beta_2$  equals zero in columns (5) and (6) suggests that there is little variation in take-up behavior related to the size of the

benefit that remains after removing average differences in behavior across districts and for employees of different ages and experience levels.

As mentioned, the coefficients on price using the two instrumental variables strategies are remarkably similar to one another. There is no reason to have expected this to be the case. The first instrumental variables strategy would produce biased estimates if teachers sorted across districts according to their preferences for additional retirement benefits. The second would do so if cohorts of teachers within the same district were different from one another or if pre-1998 retirement behavior led to differences in preferences across teachers within a district who had different amount of experience in 1998.

Meanwhile, the cost to the pension fund of providing the extra benefits is positively related to price. Table 3 presents the results of estimating the relationship between price and the cost shown in equation (8). In the table, I use the beginning salary for a teacher with MA degree as the instrument for price and include increasingly more explanatory variables for market and district characteristics. The estimates in the table suggest that the average cost of providing the extra retirement benefits to an employee increases as the price of the upgrade increases. For example, controlling for individual and market-level characteristics including county fixed effects, the extra retirement benefits paid to purchasers is \$4,905 higher for every \$1,000 increase in price. The estimates range from \$4,905 to \$5,614, depending on the set of controls included in the

regression, however none of the estimates are statistically different from one another.

*Employee Valuation of Current Income Relative to Future Income*

The coefficient estimates from equation (1) and equation (8) can be used to trace out the upgrade's demand, average cost, and marginal cost curves. Comparing these curves gives 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 (1) defines the demand curve, while equation (8) can be used to draw the average cost curve.<sup>22</sup> In order to determine the marginal cost curve of providing the upgrade to purchasing IPS employees, I use the coefficient estimates from the estimation of (1) and (8) in the following formula:

$$MC(D) = \frac{\partial TC(D)}{\partial D} = \frac{\partial (AC(D) \cdot D)}{\partial D} = \frac{\partial AC(D)}{\partial D} D + AC(D). \quad (9)$$

Since the demand curve traces the marginal willingness-to-pay a particular price and the marginal cost curve presents the cost of providing the upgrade to those who purchase at a particular price, the ratio between the vertical heights of these

<sup>22</sup> Note that equation (1) is slightly different than the actual equations underlying the estimates in Table 3, which were equations (2) and (5). The difference is that equation (1) only includes price and the expected benefits proxy as the explanatory variables, but no other covariates. I use equation (1) because tracing out the demand curve involves a choice of vertical intercept. In equations (2) and (5) any choice of the intercept would be conditional on specified values of the characteristics. With equation (1), the only characteristic to fix when drawing a demand curve is the value of  $B_i$ , or the expected benefits proxy. The results vary little based on this choice. I use the average level of  $B_i$  (calculated using the maximum MA Degree salary) across districts to assign the intercept of the demand curve.

lines is the cost of providing the upgrade to IPS employees relative to the value the employees place on the benefits.

Figure 2 presents these curves for IPS employees with 22 to 28 years of experience in 1998. The solid line is the demand curve. The dashed and dotted lines represent the average and marginal cost curves, respectively. The data underlying the demand curve and average cost curves are represented by circles and squares, respectively, the size of which are weighted by the number of observations each represents.

The first and most noteworthy conclusion from the figure is that the demand curve is everywhere below the marginal cost curve. This implies 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. Taking the estimated demand curve as the true curve, the highest marginal willingness-to-pay for the upgrade of these employees – as indicated by the point of intercept between the demand curve and the vertical axis – is just \$25,996. Yet, this high valuation employee can expect to collect over \$121,000 in benefits suggesting that even the highest value employee is willing to trade just 21 cents of current compensation for a future dollar of benefits.<sup>23</sup> Averaging along the entire aggregate demand curve,

<sup>23</sup> The estimated valuation of the highest valuation person depends on the value of the vertical intercept of the demand curve. In this setting, the demand curve was drawn conditional on the expected benefit from upgrade being equal to the average of expected benefits in the sample. Using a higher expected benefit value results in a demand curve that lies vertically above the one

which represents the willingness-to-pay in the population of workers, employees in IPS are willing to trade just 20 cents for a dollar's worth of future benefits.

Another way to calculate the ratio of willingness-to-pay relative to cost is to estimate demand for the upgrade as a function of the ratio of expected benefits to price. To assuage concerns about endogeneity of this measure, the actual ratio for an individual employee can be instrumented with the ratio of expected benefits to price calculated using the moments of the salary schedule in an analogous manner to the instrumental variables already used. For example, the ratio of *BMAMax* to the price calculated using the *BABegSalary* in a district can be used to instrument for the ratio of benefits to price of the employees in that district of the same age and with the same experience. When I conduct such analyses using different moments from the salary schedule for different educational attainment levels, I find that the coefficient on this ratio is between 0.18 and 0.29 and statistically significant at the one percent level.<sup>24</sup> This suggests that employees are willing-to-trade between 18 and 29 cents of current income for an additional dollar of future income, which concords with the estimates of equations (1) and (8) that allow for different coefficients on role of a dollar increase in price and benefits.

---

in Figure 2. If the expected benefit is equal to the highest in the sample, the highest valuation person is willing to pay just 33 cents on the dollar and the average valuation in the population is estimated to be 30 cents on the dollar. Alternatively, if I use the value of the lowest expected benefits to draw the demand curve, the highest valuation is 17 cents per dollar and the population average is 13 cents per dollar of expected benefits.

<sup>24</sup> Results are not reported in the tables, but are available from the author upon request.



The second notable aspect of the figure is that, although I have quite a bit of data, little of its mass sits in the range of demand proportions below 40 percent of the population.<sup>25</sup> This makes it difficult to be certain what the levels of valuation are among the employees with the highest willingness-to-pay because I am forced to go out-of-sample to make predictions. However, even just comparing the demand and marginal cost curves over the region of observed data suggests that over half of employees are unwilling to trade just 20 cents of current income for a dollar of expected future compensation. The fact that the majority of employees' willingness-to-pay for future compensation lies so far below the costs of providing it implies that they are also not willing to pay for the costs of these pensions benefits with decreased wages. Thus, the theories that rely on worker preferences cannot explain why we see generous defined benefits pensions in the public sector.

#### *Placing the Estimates in Context*

The estimates presented indicate public school employees would prefer \$2 in current wages to \$10 in PDV of annuitized wealth in their retirement package. This estimate places the discount rates of teachers at the high end of those for employees in other settings. For example, in a setting of public school

<sup>25</sup> Six percent of the data fall in the range where take-up is between zero and forty percent. Note that estimating the demand and cost curves without the outlying data leads to a similar demand curve and a slightly steeper cost curve.

employees in Oregon, Chalmers, Johnson and Reuter (2012) found only 15 percent of employees chose a lump sum option that offered about 70 cents, on average, for every dollar of annuity it replaced. Warner and Pleeter (2001) found that, when given the opportunity to choose, the vast majority of enlisted military took a lump sum payment instead of a retirement annuity worth twice as much. Brown, Ivkovich and Weisbenner (2013) find 70 percent of Croatians chose an immediate payment option, even when the deferred payment stream offered a 26 percent nominal internal rate-of-return.<sup>26</sup>

As these studies suggest, the preference for current funds relative to annuitized payments is not uncommon in the literature. Indeed the annuity puzzle has existed for years and has been explained using people's desire for liquidity and/or diversification of assets, bequest motives, or reliance on family members for insurance (Kotlikoff and Spivak 1981). What is striking about the present study is the magnitude of money that some employees in this setting are willing to leave on the table. There are two main reasons that employees in this and other settings may reveal such low willingness-to-pay for annuities relative to lump sums.

<sup>26</sup> Most of the previous literature involved tradeoffs where the lump sum was lower in PDV than the annuity. Clark, Morrill and Vanderweide (2013) find that less than a third of public sector workers in North Carolina choose a lump sum option even though the lump sum option was likely more valuable in PDV terms than the alternative annuity option.

One possibility is that the public employees in this setting strongly favored the lump sum payments, perhaps in part because they are oversaturated with this particular annuitized illiquid asset. While private sector employees earn \$1 in compensation through employer provided retirement benefits for every hour they work, teachers earn \$3. Taking spousal benefits and Social Security payments into account makes it even more likely that teachers would prefer current compensation to "over-annuitization" through additional retirement benefits. Recall that the discount rate used to calculate the cost of the pensions here was the financial discount rate (the Illinois bond rate), which is determined by the value of bonds to a diversified investor. If the public school employees are not diversified, say, because the majority of their savings is in the form of their pensions, they may discount pension increases at a lower rate than the one used to calculate costs, which would lead to the result presented. This is consistent with recent survey evidence in which teachers were more likely to report wanting additional pension dollars invested in a defined contribution type plan, rather than a defined benefit style plan (DeArmond and Goldhaber 2010). And, the Illinois constitution forbids borrowing against public sector pensions. Therefore, if employees have preferences either for liquidity or diversification, not purchasing the upgrade may be optimal even if the rate-of-return is high.

The second possibility is that employees do not accurately make decisions about their purchase based on valuations relative to price. This could occur for

several reasons, including uninformed consumers, default behavior, credit constraints or a lack of financial literacy. There are several features of this setting that make each of these behavioral explanations less likely than in other previous research settings. First, the employees in IPS were repeatedly presented with information about the upgrade. The 2.2 formula change and the upgrade were widely publicized; both were the subject of many news reports and newspaper articles in 1998 (e.g. Finke 1998, Erikson 1998, Thomson and Mask 1998). Also, after the formula change, detailed mailings were sent out to all IPS employees documenting how the price and benefit of the upgrade could be calculated.<sup>27</sup> In addition to the initial worksheets for employees to calculate their own price and benefits of the upgrade, the TRS also sends members annual personalized statements of the upgrade cost and the increase in annual retirement benefits from purchasing. Moreover, teachers and other public school employees are unionized and one major role for unionization is information provision. Existing survey evidence has shown that teachers in traditional defined benefit pension plans are quite knowledgeable about their pension plans, more so than workers in the private sector more generally (DeArmond and Goldhaber 2010). In fact, the teachers unions in Illinois also sent information to their members about the

<sup>27</sup> Examples of worksheets mailed to employees in the weeks following the formula change are in Appendix A.

upgrade and its value to them. It is therefore unlikely that a lack of information prevents take-up for many, if any, employees.

Second, default behavior likely plays a smaller role in this setting than others. The features of the upgrade purchase decision are much simpler than other investment decisions in which the literature finds people behaving in ways at odds with the rational expectations framework (e.g. Madrian and Shea 2001). For example, the price of the upgrade and the potential benefit from purchasing are more straightforward than the calculations necessary to make decisions about defined contribution pensions, which include complicated decisions about the portfolio choice.<sup>28</sup> Recent literature suggests that people are more likely to opt-out of the default investment option when the returns to doing so are large (Beshears et al. 2010), as is the case here.

Third, the fact that employees could spread payments over as many as five years lessens concerns about credit constraints. The ability to have the payment for the upgrade deducted from one's paycheck, deducted from one's retirement benefit check (if retiring immediately after the new policy came into place) or even transferred from other retirement accounts also diminishes the likelihood of binding credit constraints. Moreover, employees without sufficient assets could

<sup>28</sup> Consider, for example, the relative simplicity of a defined benefit pension program as compared to a defined contribution pension program. While the former has essentially two parameters - the contribution rate and the benefit accrual rate - that are constant across employees, the latter has many pieces - including contribution rates, matching rates, asset distribution decisions, interest rate-risk tradeoffs, etc. - that may vary across employees.

delay purchase to allow themselves time to save. The ability to smooth payments over time and have them deducted from one's paycheck lessens the role for credit constraints in this setting.

Finally, in a well-designed survey, Brown, Kapteyn, Luttmer and Mitchell (2013) find that financial literacy is strongly correlated with people's consistent valuation of hypothetical lump sum versus Social Security benefit trade-offs across a variety of framing set-ups. However, financial literacy likely plays a smaller role in this setting than in Brown et al. (2013) for a few reasons. First, the employees with an opportunity to upgrade were directly and repeatedly provided with a great deal of personalized information about the two most important pieces of financial information regarding the upgrade: its price and the annual additional retirement benefits associated with purchase. Second, even relatively naïve assessments of the price and additional retirement benefits associated with the upgrade provide the employee with enough information to "make a rough guess" of the upgrade's value. This is particularly true because I focus on employees who are close to retirement, which eliminates much of the need for a precise understanding of discounting, a potential financial literacy pitfall. Finally, the decision to upgrade is an actual choice, rather than hypothetical, making it more likely that even people with low financial literacy will invest in effort to make a decision regarding the upgrade.

Of course, in this study, I am unable to determine whether the willingness-to-pay I estimate represents employees' true valuation of their pension benefits or whether there is some alternative behavioral explanation underlying their decisions. Regardless, the parameter I estimate – public school employees' willingness-to-pay for increased retirement benefits relative to the cost of those benefit increases – is an important one for determining efficient public sector compensation because the employees' low valuation in this setting makes clear that employees would have been unwilling to forgo salary in exchange for these benefits.

## **6. Conclusion**

In this paper, I test empirically whether teachers and other public school employees are willing to pay as much for additional pension benefits as it costs to provide them with the benefits. Almost all states use a form of defined benefit program similar to the one studied here as their main pension system for public employees. By guaranteeing a large fraction of pre-retirement salary from retirement until death, these programs offer generous deferred compensation to public employees.

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 increases in benefits and the cost of supplying them. I show that teachers' valuation of the increased pension benefits was much less than their cost. Based on the estimates, IPS employees on average prefer an increase in current wages of just over \$2 to a \$10 increase in the PDV of annuitized wealth in their retirement package.

This research provides new information critical to public policy decision making. Forty-six states have maintained defined benefit pension plans as the main provider of retirement benefits to their state employees (GAO 2007). These plans are becoming quite costly. In 2012, the taxpayers of Illinois contributed \$5 billion to the pension funds for state and local employees and those funds were still underfunded by \$96.8 billion.<sup>29</sup> Meanwhile, Illinois is no more generous than many other states; the 1998 pension 'sweetener' brought pension benefit levels to the median of what they were in other states where teachers do not receive Social Security. Nor was Illinois the only state that decided to provide more generous benefits. 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. Of course, the valuation of additional pension benefits I estimate is unlikely to be the employees'

<sup>29</sup><http://www.illinoispolicy.org/simplereport/illinois-taxpayers-bear-the-brunt-of-rising-pension-costs/> (Accessed October 2, 2014)



average valuation of their entire pension. However, the margin at which I estimate pension benefit valuation is important for policymakers because it is quite representative of the margin on which many other states are currently making decisions about increasing or decreasing their pension benefits.

In this context, the main finding of this paper, that the majority of IPS employees value their pension benefits at 20 cents on the dollar, has two important implications. First, it suggests a possible Pareto-improving and politically feasible solution to the current inability of states to pay their promised pension benefits to public employees. Governments could offer to buy back pension benefits from teachers and other public sector employees. If the results here generalize, governments may be able to buy back promised employee pension benefits, or at least some of these promised benefits, for as little as twenty cents on the dollar. Doing so would draw down the pension obligations of governments both significantly and immediately, rather than waiting for a reduction in benefits to take effect years in the future.

Of course, just because employees were only willing to pay about 20 cents, on average, for an increase in pension benefits of \$1, does not mean they will be willing to sell the benefits back at such a low price. A difference between willingness-to-pay for and willingness-to-sell could lie behind the low take-up of a generous lump sum offer for the public school employees in Chalmers and Reuter (2012). In a hypothetical setting, Brown et al. (2013) showed large

discrepancies between willingness-to-pay for and willingness-to-sell Social Security benefits, though there were no differences for those classified as financially literate.<sup>30</sup> However, because there may be large potential to reduce the burden on taxpayers while also making public sector employees better off such a buyback policy is worth more attention.

Second, 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 such generous defined benefit pension packages. If the choice in the structure of teacher compensation is between higher pension benefits and equivalent increases in current salary, as is often assumed, employees would be better off with increases in current salary. In fact, even small increases in current salary, or one-time bonuses, would make the public school employees better off. This evidence suggests that workers are not paying for the generous pension benefits with decreases in current wages, as is also often assumed. This should therefore turn our attention to other explanations for the existence of these generous benefits.

One possible alternative explanation for the use of defined benefit pension plans is that there is a subset of employees who prefer such benefits to current

<sup>30</sup> More generally, there is a long literature documenting differences between willingness-to-pay and willingness-to-accept (for a summary of recent evidence, see DellaVigna 2009). Although debate exists about what drives these differences across settings, it is probably safe to assume that public employees will require more from a buy-back program than they are willing-to-pay when offered the opportunity to purchase additional benefits.

compensation. If these employees become union leadership, they may negotiate more generous pension benefits than the marginal worker would prefer. Positing that this type of difference between union leadership and members exists has been popular in the literature (Freeman 1986), but does not yet have much empirical support. At the crux of this set of theories is the idea that some employees value future compensation more than current compensation.<sup>31</sup> Although it is an out-of-sample prediction, my estimates suggest that even the highest willingness-to-pay employee values the increased benefits at just 29 cents on the dollar. This is suggestive evidence that the differences in valuation across teachers are able to explain the existence of generous pension benefits in the public sector.

The inability of the first two classes of theories to explain the use of generous defined benefit pension plans for public employees points to a last set of models. These models rely on the political nature of public pensions to drive a wedge between the actual costs of pension benefits and the perceived costs of those benefits used to make decisions about compensation for public employees (e.g. Glaeser and Ponzetto 2012). Since in these models politicians care mostly about being elected (or reelected), their discount rates for future budgetary expenses may be low. As such, they are concerned more with voters' perceived costs of pension benefits relative to current salaries, rather than the actual relative

<sup>31</sup> Otherwise, the union leadership would be maximizing an objective function orthogonal to its members' preferences. While possible, this does not seem rational or to be the basis for the popular argument of union leadership's involvement in generous public sector pensions.

costs of these two compensation mechanisms. If for some reason the costs of pension benefits are less salient to voters (perhaps because they are less publicized), the perceived costs of pensions may be much lower than the actual costs of pensions. The true costs of pension benefits also may not be relevant for politicians if their constituents are likely to move out of the jurisdiction before the payment comes due (Inman 1982). The difference between real and perceived costs is exacerbated in a system that allows underfunding of public pensions, as most states do today. Understanding the role of salience, myopia and/or residential transitions in the determination of the lifetime compensation mix for public employees is an important avenue for future research.

## REFERENCES

- Beshears, John, James J. Choi, David Laibson, and Brigitte Madrian. 2010. "The Limitations of Defaults." NB-10. Working Paper Prepared for the 12th Annual Joint Conference of the Retirement Research Consortium.
- Brown, Jeffrey, Arie Kapteyn, Erzo Luttmer, and Olivia Mitchell. 2013. "Decision Complexity as a Barrier to Annuitization." NBER Working Paper No. 19168.
- Brown, Jeffrey, Zoran Ivkovich and Scott Weisbenner. 2013. "Empirical Determinants of Intertemporal Choice." NBER Working Paper No. 18517.
- Bureau of Labor Statistics. 2008. National compensation survey: Employee benefits in the United States, March 2008. Bulletin 2715. Washington, DC: Bureau of Labor Statistics. Available at <http://www.bls.gov/ncs/ebs/benefits/2008/ownership/private/table02a.pdf>
- Chalmers, John, Woodrow Johnson and Jonathan Reuter. 2012. "How Do Retirees Value Life Annuities? Evidence from Public Employees." *Review of Financial Studies*. 25(8):2601-2634.
- 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.
- DellaVigna, Stefano. 2009. "Psychology and Economics: Evidence from the Field." *Journal of Economic Literature*. 47(2): 315-372.
- Ehrenberg, Ronald G. 1980. "Retirement System Characteristics and Compensating Differentials in the Public Sector." *Industrial and Labor Relations Review*. 33(4):470-483.
- Ehrenberg, Ronald and Robert Smith. 1983. "Estimating Wage-Fringe Trade-offs: Some Data Problems." in *The Measurement of Labor Cost*, ed. Jack Triplett. University of Chicago Press. pp. 347-370.

- Einav, Liran, Amy Finkelstein, Mark Cullen. 2010. "Estimating Welfare in Insurance Markets Using Variation in Prices" *Quarterly Journal of Economics*. 123(3):877-921.
- Erikson, Kurt. 1998. "Deal Made for Teacher Pensions." *The Pantagraph* (Bloomington, IL). May 22, 1998. Page A3.
- Finke, Doug. "Agreement on Teachers' Pensions Reached/Lawmakers Set to Wrap Up Today". *The State Journal Register*. May 22, 1998. Page 1.
- 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.
- Glaeser, Edward L. and Giacomo A. M. Ponzetto. 2012. "Shrouded Costs of Government: The Political Economy of State and Local Pensions." *Universitat Pompeu Fabra Working Paper No 1336*.
- Gustman, Alan, Olivia Mitchell and Thomas Steinmeier. 1994. "The Role of Pensions in the Labor Market: A Survey of the Literature." *Industrial and Labor Relations Review*. 47(3): 417-438.
- Illinois Economic and Fiscal Commission. 1998. *Illinois Bond Reference*
- Inman, Robert P. 1982. "Public Employee Pensions and the Local Labor Budget," *Journal of Public Economics*, 19(1): 49-71.
- Kings, J., E. Turkisch and N. Manning (2007), "Public Sector Pensions and the Challenge of an Ageing Public Service", *OECD Working Papers on Public Governance*, No. 2, OECD Publishing. <http://dx.doi.org/10.1787/285530706017>
- Kotlikoff, Laurence and Avia Spivak. 1981. "The Family As an Incomplete Annuities Market." *Journal of Political Economy*. 89(2): 372-91.
- Madrian, Bridget and Dennis Shea. 2001. "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior." *Quarterly Journal of Economics*.

116(4): 1149-1187.

- Montgomery, Edward and Shaw, Kathryn, 1997. "Pensions and Wage Premia," *Economic Inquiry*, Oxford University Press, vol. 35(3), pages 510-22, July.
- Novy-Marx, Robert and Joshua D. Rauh. 2009. "The Liabilities and Risks of State-Sponsored Pension Plans" *Journal of Economic Perspectives*, 23(4): 191-210.
- Schiller, Bradley R. and Randall D. Weiss. 1980. "Pensions and Wages: A Test for Equalizing Differences." *The Review of Economics and Statistics*. 62(4): 529-538.
- Smith, Robert S. 1981. "Compensating Differentials for Pensions and Underfunding in the Public Sector." *The Review of Economics and Statistics*. 63(3):463-468.
- 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.
- Thomson, Dona and Teresa Mask. 1998. "Educators Give High Marks to Teacher Pension Plan." *Springfield Daily Herald*. May 23, 1998. Page 11.
- Warner, John and Saul Pleeter. 2001. "The Personal Discount Rate: Evidence from Military Downsizing Programs." *American Economic Review*. 19(1): 33-53.

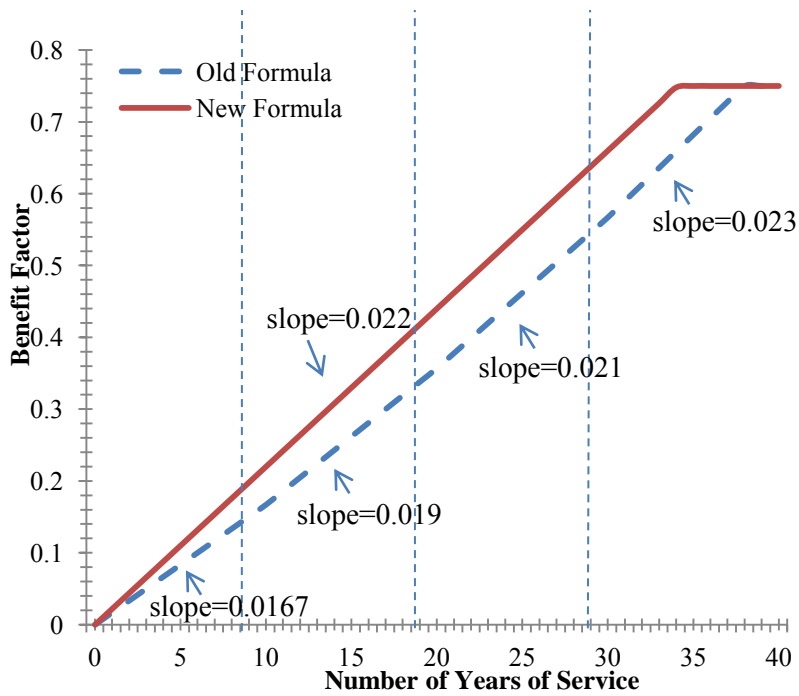


FIGURE 1. TEACHER RETIREMENT BENEFIT FORMULAS IN ILLINOIS, BEFORE AND AFTER 1998

Note: 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.

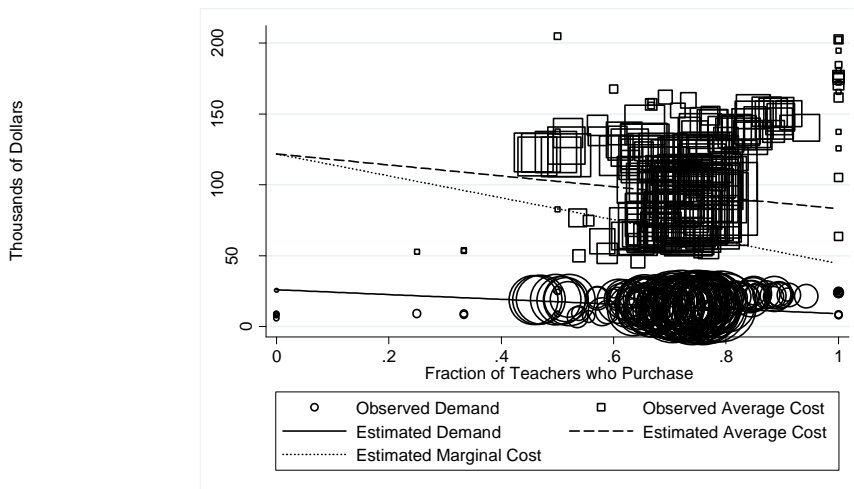


FIGURE 2. ESTIMATED DEMAND AND COST CURVES FOR EMPLOYEES WITH 22 TO 28 YEARS OF EXPERIENCE IN 1998

Note: Based on the author's calculations using the TRS and TSR.



TABLE 1. TAKE-UP RATES, RETIREMENT RATES, PRICES, AND COSTS BY LEVEL OF EXPERIENCE IN 1998

| Years of Experience in 1998 | Fraction Who Retire by 2009 | Fraction Who Purchase Upgrade by 2009 | Mean Price [\$1,000s] (Std. Dev.) | Mean Benefit [\$1,000s] (Std. Dev.) | Mean Ratio of Benefits to Price (Std. Dev.) | Number of Obs. |
|-----------------------------|-----------------------------|---------------------------------------|-----------------------------------|-------------------------------------|---|----------------|
| 22                          | 0.6                         | 0.7                                   | 14.02<br>(4.08)                   | 86.06<br>(28.41)                    | 6.29<br>(2.87)                              | 2,409          |
| 23                          | 0.67                        | 0.71                                  | 14.61<br>(4.35)                   | 90.41<br>(30.62)                    | 6.27<br>(1.99)                              | 2,589          |
| 24                          | 0.78                        | 0.76                                  | 14.90<br>(4.40)                   | 93.01<br>(31.55)                    | 6.27<br>(1.20)                              | 2,991          |
| 25                          | 0.83                        | 0.78                                  | 15.26<br>(4.51)                   | 96.94<br>(33.62)                    | 6.37<br>(1.26)                              | 3,121          |
| 26                          | 0.88                        | 0.77                                  | 15.40<br>(4.47)                   | 100.40<br>(35.39)                   | 6.54<br>(1.43)                              | 2,909          |
| 27                          | 0.89                        | 0.74                                  | 15.75<br>(4.73)                   | 102.73<br>(37.56)                   | 6.53<br>(1.37)                              | 2,674          |
| 28                          | 0.91                        | 0.73                                  | 15.90<br>(4.63)                   | 102.17<br>(36.42)                   | 6.42<br>(1.27)                              | 2,736          |
| Whole Sample                |                             |                                       |                                   |                                     |   | 19,429         |

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. ESTIMATES OF THE RELATIONSHIP BETWEEN DEMAND (PURCHASE OF THE UPGRADE) AND PRICE

|   | OLS      | IV Cross-District |           | IV Within-District |         |         |
|---|----------|-------------------|-----------|--------------------|---------|---------|
|   | 1        | 2                 | 3         | 4                  | 5       | 6       |
| <b>Panel A. Salary for BA Degree Holder Used in Creating the Instruments</b>    |          |                   |           |                    |         |         |
| Price (\$1,000)   | 0.009*** | -0.082***         | -0.071*** | -0.070***          | -0.167* | -0.171* |
|   | 0.001    | (0.007)           | (0.008)   | (0.008)            | (0.089) | (0.094) |
| Benefit (MA Maximum Salary)   | 0.001*** | 0.014***          | 0.012***  | 0.012***           | 0.002   | 0.003   |
|   | 0.0001   | (0.001)           | (0.001)   | (0.001)            | (0.002) | (0.002) |
| <b>Panel B. Salary for MA Degree Holder Used in Creating the Instruments</b>    |          |                   |           |                    |         |         |
| Price (\$1,000)   |          | -0.074***         | -0.066*** | -0.065***          | -0.115  | -0.121  |
|   |          | (0.008)           | (0.009)   | (0.009)            | (0.104) | (0.116) |
| Benefit (MA Maximum Salary)   |          | 0.013***          | 0.012***  | 0.011***           | 0.002   | 0.002   |
|   |          | (0.001)           | (0.001)   | (0.001)            | (0.002) | (0.002) |
| <b>Panel C. Salary for Ph.D. Degree Holder Used in Creating the Instruments</b> |          |                   |           |                    |         |         |
| Price (\$1,000)   |          | -0.092***         | -0.090*** | -0.090***          | -0.129  | -0.130  |
|   |          | (0.012)           | (0.015)   | (0.015)            | (0.106) | (0.116) |
| Benefit (MA Maximum Salary)   |          | 0.016***          | 0.016***  | 0.015***           | 0.001   | 0.001   |
|   |          | (0.002)           | (0.002)   | (0.002)            | (0.002) | (0.003) |
| County Fixed Effects  | X        | X                 | X         | X                  |         |         |
| District Characteristics  | X        |                   | X         | X                  |         |         |
| Experience X Age Fixed Effects  |          |                   |           | X                  | X       | X       |
| Experience X Age Fixed Effects<br>X District Characteristics                    |          |                   |           | X                  |         |         |
| District Fixed Effects  |          |                   |           |                    | X       | X       |

Note: Based on the author's calculations using data from the TSR and TRS. Sample includes all teachers eligible for the upgrade with between 22 and 28 years of experience in 1998. Each column-panel set represents estimates from a separate regression. The dependent variable is a dummy variable equal to one if the employee purchased the upgrade. As indicated, estimates are from OLS or 2SLS estimation. In the columns labeled IV Cross-District, the actual price paid for the annuity is instrumented with the beginning scheduled salary paid to teachers with a BA in one's district of employment in 1998 (measured in thousands of \$2010). In the columns labeled IV Within District, the instrument used is the scheduled salary paid to a teacher in the district with the same amount of experience as the teacher. Each panel presents estimates using the scheduled salary for a different group of teachers. All specifications include individual characteristics. Market-level characteristics are added in groups as described in the text. Standard errors (in parentheses) are clustered at the district level; \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.

TABLE 3. INSTRUMENTAL VARIABLES ESTIMATES OF THE RELATIONSHIP BETWEEN UPGRADE TOTAL COST AND PRICE

|                          | (1)                | (2)                 | (3)                 |
|--------------------------|--------------------|---------------------|---------------------|
| Price (\$1,000)          | 5.614**<br>(2.352) | 5.027***<br>(0.209) | 4.905***<br>(0.276) |
| PUMA Characteristics     | X                  | X                   |                     |
| County Fixed Effects     |                    | X                   | X                   |
| County Characteristics   | X                  |                     |                     |
| District Characteristics |                    |                     | X                   |

Note: Based on the author's calculations using data from the TSR and TRS. Sample includes teachers with 22 to 28 years of experience in 1998. The dependent variable is the expected cost (at time of purchase) to the TRS if the employee purchased the upgrade, measured in \$1,000. All specifications include individual characteristics. Market-level characteristics are added in groups as described in the text. Standard errors (in parentheses) are clustered at the district level; \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.

## For Online Publication

### Appendix A. Worksheet Examples of Information Provided by TRS to IPS Employees for the Purpose of Assessing Upgrade Purchase Decisions

#### ***Calculating Cost to Upgrade***

|   |   |  |  |     |  |             |  |        |  |      |  |    |  |
|---|---|--|--|-----|--|-------------|--|--------|--|------|--|----|--|
| Enter total years under old four-step formula or <b>20</b> , whichever is less<br>Minus years worked under new formula divided by 3 (drop any decimals)<br>Equals number of years to upgrade<br>Multiply by 1% (.01)<br>Multiply by highest salary earned during last four years (do not include year of application)<br><b>COST TO UPGRADE</b> | <table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 50%; height: 20px;"></td> <td style="width: 50%;"></td> </tr> <tr> <td style="text-align: center;">(-)</td> <td></td> </tr> <tr> <td style="text-align: right;">_____ years</td> <td></td> </tr> <tr> <td style="text-align: right;">x    %</td> <td></td> </tr> <tr> <td style="text-align: right;">x \$</td> <td></td> </tr> <tr> <td style="text-align: right;">\$</td> <td></td> </tr> </table> |  |  | (-) |  | _____ years |  | x    % |  | x \$ |  | \$ |  |
|   |   |  |  |     |  |             |  |        |  |      |  |    |  |
| (-)   |   |  |  |     |  |             |  |        |  |      |  |    |  |
| _____ years   |   |  |  |     |  |             |  |        |  |      |  |    |  |
| x    %  |   |  |  |     |  |             |  |        |  |      |  |    |  |
| x \$  |   |  |  |     |  |             |  |        |  |      |  |    |  |
| \$  |   |  |  |     |  |             |  |        |  |      |  |    |  |

***Option: You do not upgrade to 2.2 and have less than 30 years of service credit as of July 1, 1998.***

From the retirement formula table, find the total accrual percentage for years of service prior to 7/1/98 (Example: 20 years = .3570) \_\_\_\_\_ (1)

Years of service after 7/1/98 x .022 \_\_\_\_\_ (2)

Total years of service factor (add lines 1 and 2) = \_\_\_\_\_

Multiply by average salary (from page 5) x \_\_\_\_\_

Annual annuity = \_\_\_\_\_

Divide by 12 (months) /12

Monthly annuity = \_\_\_\_\_

***Option: You upgrade all service credit to 2.2.***

Average salary (page 5) \$ \_\_\_\_\_

Multiply total years of service by .022 x \_\_\_\_\_

Annual annuity = \_\_\_\_\_

Divide by 12 (months) /12

Monthly annuity \$ \_\_\_\_\_