

NBER WORKING PAPER SERIES

THE EFFECT OF PENSION DESIGN ON EMPLOYER COSTS AND EMPLOYEE RETIREMENT CHOICES:
EVIDENCE FROM OREGON

John Chalmers
Woodrow T. Johnson
Jonathan Reuter

Working Paper 18517
<http://www.nber.org/papers/w18517>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2012

Previously circulated as "Pension Costs and Retirement Decisions in Plans that Combine DB and DC Elements: Evidence from Oregon.bv" Prepared for "Retirement Benefits for State and Local Employees: Designing Pension Plans for the Twenty-First Century," NBER Conference, August 17-18, 2012, Jackson Hole, WY. We thank John Shoven (discussant) and the other conference participants for encouraging us to extend our analysis in interesting new directions, and we thank seminar participants at the Boston College Center for Retirement Research for many helpful suggestions. We thank employees from Oregon's Public Employees Retirement System, who provided invaluable assistance by helping us to collect and interpret PERS data, Guy Tauer from the Oregon Employment Department, who helped us collect additional data, and Lenore Robbins. The authors acknowledge financial support from the Smith Richardson Foundation. Parts of this research were supported by the U.S. Social Security Administration through grant #10-P-98363-1-05 to the National Bureau of Economic Research as part of the SSA Retirement Research Consortium. The findings and conclusions expressed are solely those of the authors and do not represent the views of the SSA, any agency of the Federal Government, or the NBER. The Securities and Exchange Commission, as a matter of policy, disclaims responsibility for any private publication or statement by any of its employees. The views expressed herein are those of the authors and do not necessarily reflect the views of the Commission or of its staff.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w18517.ack>

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2012 by John Chalmers, Woodrow T. Johnson, and Jonathan Reuter. 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.

The Effect of Pension Design on Employer Costs and Employee Retirement Choices: Evidence from Oregon

John Chalmers, Woodrow T. Johnson, and Jonathan Reuter

NBER Working Paper No. 18517

November 2012, Revised March 2013

JEL No. D83,H55,J26

ABSTRACT

Oregon's Public Employees Retirement System (PERS) is a rich setting in which to study the effect of pension design on employer costs and employee retirement-timing decisions. PERS pays retirees the maximum benefit calculated using three formulas that can be characterized as defined benefit (DB), defined contribution (DC), and a combination of DB and DC. From the employer's perspective, we show that this "maximum benefit" calculation is costly. Average ex post retirement benefits are 54% higher than they if had been calculated using only the DB formula. Monte Carlo simulations verify that the higher cost could have been predicted at the start of our sample period. From the employee's perspective, we show that plan design distorts the retirement-timing decision: employees receiving DC benefits are significantly more likely to retire before the normal retirement age than employees receiving DB benefits. Exploiting two sources of exogenous variation in the level of the DC benefit, we show that employees respond to within-year variation in their retirement incentives and, consistent with peer effects, that they respond more strongly to these incentives when more of their coworkers face similar incentives. Finally, consistent with the emerging literature on financial mistakes by households, we show that a small but significant fraction of retirees would benefit from shifting their retirements by as little as one month.

John Chalmers
Charles H. Lundquist College of Business
University of Oregon
Eugene, OR 97403
jchalmer@lcbmail.uoregon.edu

Woodrow T. Johnson
Division of Risk, Strategy,
and Financial Innovation
U.S. Securities and Exchange Commission
100 F Street, NE
Washington, DC 20549
JohnsonW@sec.gov

Jonathan Reuter
Carroll School of Management
Boston College
224B Fulton Hall
140 Commonwealth Avenue
Chestnut Hill, MA 02467
and NBER
reuterj@bc.edu

1. Introduction

Employers must weigh the expected benefits of the pension plans they offer to employees against the expected costs. Among other benefits, offering a generous pension plan may allow an employer to attract and retain higher quality employees. Governor Tom McCall emphasized these potential benefits in 1967 when arguing to reform Oregon's Public Employees Retirement System (PERS):¹

“We are in a time of inflation and high employment. I have personal experience with the difficulty of recruiting top quality people at the available salaries and personal knowledge of the real sacrifices made by some who have accepted positions in my administration... At all levels our state employment has shown heavy turnover. This requires extensive recruiting and training programs and threatens a real loss of competency if not checked...”

The idea was that a more generous pension plan would improve the quality of the services provided by state and local employers while reducing the administrative and other costs associated with employee turnover. On the other hand, increasing expected retirement benefit payments imposes a direct cost on employers who must cover the larger pension payments. It may also impose indirect costs insofar as changes to plan generosity affect employee behavior.

PERS was created in 1946 and has been periodically modified over the years by the state legislature. By 1990, PERS had evolved into a complex pension plan with both defined benefit (DB) and defined contribution (DC) elements, serving essentially all non-federal public employees across the hundreds of non-federal public employers in Oregon.² When a member retires, PERS calculates retirement benefit using up to three different formulas that include DB and DC elements, and it automatically pays the largest benefit for which the member is eligible. The DB benefit depends upon the member's salary and years of service. The DC benefit depends upon the accumulation of assets in one or two DC-style retirement accounts. To be clear, the DC elements in PERS differ significantly from those in a traditional 401(k) plan: Oregon manages the investments, provides an annual return of at least 8% to certain plan members, and converts DC account balances into life annuity payments using annuity factors that Chalmers and Reuter (2012) show to be better than actuarially fair. The fact that members can expect to receive higher retirement benefits when equity market returns have been high makes the pension more

¹ The quote comes from page 12 of “The Oregon Public Employees Retirement System History, the First 60 Years,” published by PERS on July 6, 2010.

² See Snell (2012) and the following link for a listing of state and local plans that provide complex benefits plans that share some common characteristics with the Oregon plan. <http://www.nasra.org/resources/HybridBrief.pdf>.

generous to members—and more expensive to PERS employers—than if PERS used only its DB benefit formula. Similarly, the fact that members are insured against downside market risk makes the pension more generous to members—and more expensive to PERS employers—than if PERS used only its DC benefit formula.

In this paper, we study the effect of PERS' structure on both employers' pension costs and members' retirement-timing decisions. We begin by comparing the actual retirement benefits of PERS retirees to the hypothetical benefits they would have received if PERS used only its DB benefit formula. Using administrative data from 1990 to 2003, we find that the majority of retirees (87.5%) receive higher benefits than they would have if PERS used only the DB benefit formula. And, for the typical retiree whose career started at age 39 and ended after 21 years of service, actual benefits are 54% higher than those calculated using just the DB benefit formula. While some of the additional costs can be attributed to the effect of high equity market returns on DC account balances, more than half of them can be attributed to generous features of the plan, especially PERS' use of better-than-actuarially-fair annuity factors. When we replace PERS' annuity factors with those available each year from insurance companies, we find that actual benefits would have only been 24% higher than those calculate using the DB benefit formula.

To benchmark our *ex post* cost estimates, we simulate the ratio of DC benefits to DB benefits using the PERS plan features as of January 1990 and historical equity market return data from 1929 to 1989. Our simulation results show that, in expectation, most members earn larger DC benefits than DB benefits. For a simulated member whose PERS career starts at age 39 and ends when she reaches the early retirement age of 55, DC benefits are 18% higher than DB benefits at the 25th percentile of the *ex ante* distribution, 31% higher at the median, and 48% higher at the 75th percentile. Had the simulated member worked for an additional 5 years (matching the career-length of the typical retiree in our sample), her DC/DB ratio would have been even larger. In this case, the *ex post* ratio of the typical member falls below the 25th percentile of the *ex ante* distribution. These simulations show that PERS could have known in 1990 that its provision of the DC benefit formula would result in significantly larger retirement benefits than in a DB-benefit-only plan.

While our data do not allow us to assess the effectiveness of PERS in attracting and retaining high-quality employees, they do allow us to assess how plan design affects member retirement behavior. We find that as average retirement benefits increase above the levels they

would be in a DB-benefit-only plan, the probability of PERS members retiring before the normal retirement age increases. In part, this pattern reflects the fact the high equity market returns during our sample period allowed a subset of retirees to earn more in retirement benefits than they earned in salary. More generally, by allowing members to fund retirement after fewer years of service, PERS' structure increases employee turnover, which increases administrative costs associated with hiring and training more employees.³ It also makes member retirement-timing decisions more sensitive to equity market returns.

To provide more direct evidence on the link between pension design and retirement-timing decisions, we exploit two major sources of exogenous variation in the level of the DC benefit. The first arises from the fact that, until January 2000, PERS calculated returns earned in the DC-style accounts only once per year, in March. For members retiring in other months, their account balances were determined by extending the prior year's return forward, providing them with the opportunity to exploit "stale returns" (in the spirit of Stanton (2000)). Consider a member trying to decide, in February 1993, whether to retire in February or March. She earned an annual return of 15% in her member account in 1991. However, because equity market returns were significantly lower in 1992, she expects to earn an annual return of 8% in 1992. If the member retires in February 1993, before PERS finalizes the prior-year return, PERS will calculate the change in her DC account balance between January 1992 and February 1993 using the stale prior-year return of 15%. But, if she retires in March 1993, she expects PERS will calculate the change in her DC account balance between January 1992 and March 1993 using an updated prior-year return of 8%, resulting in significantly lower benefits. In this example, the member faces a strong incentive to retire in February instead of March.

We find that members' retirement-timing decisions respond to PERS' use of stale returns. This is true both in graphs showing the fraction of retirements occurring in January and February (when members are best able to estimate the retirement incentive due to stale returns) falls sharply after PERS eliminates the use of stale returns in January 2000, and in regressions using members' retirement incentives to predict their retirement dates. While our findings suggest that the typical member is able to determine whether she benefits from having her DC account balance calculated using stale returns, we also find that several hundred of the members who retired

³ Goda, Shoven, and Slavov (2009) discuss similar policy issues that arise from the retirement incentives built into the U.S. Social Security system.

in February, and therefore had DC benefits calculated using stale prior-year returns, would have been better off retiring in March, when the prior-year returns were finalized. In other words, just as Campbell (2006) finds that some households make costly financial mistakes, we find that some members make costly mistakes with respect to the retirement-timing decision.

The second major source of exogenous variation arises from PERS' adoption of updated annuity factors in July 2003. Because the old annuity factors were based on mortality tables from 1978, and because life expectancies have increased, the new annuity factors reduced DC retirement benefits between 1.4% and 17.8%. Consistent with members seeking to avoid this well-publicized reduction in benefits, we observe more retirements during the first six months of 2003 than during any other six-month period between 1990 and 2003. Furthermore, in our regressions, we find that members facing larger reductions in annuity factors are more likely to retire before July 2003. This highlights the challenge that employers face when seeking to reduce pension costs: attempts to lower pension costs by cutting future benefits are likely to trigger large numbers of retirements, which are likely to both attenuate the cost savings and impose administrative costs on employers, at least in the short run.

Finally, we test for peer effects in the retirement-timing decision. Our motivation is the fact that members are likely to learn about retirement incentives related to stale returns or upcoming changes in benefits from their coworkers. We use two different empirical strategies, both exploiting the fact that PERS covers hundreds of different employers. We find robust evidence that members respond more strongly to their own retirement incentives when more of their coworkers face the same retirement incentives. To the extent that peer effects amplify the reactions to retirement incentives, including those that might not have otherwise been salient to the typical member, they are likely to generate amplify retirement waves. More generally, the diffusion of information about retirement incentives across co-workers helps to reconcile survey evidence suggesting low levels of pension knowledge with empirical evidence that retirement time decisions respond to retirement incentives (Chan and Stevens (2008)).

Our paper is organized as follows. In Section 2, we compare and contrast the PERS pension plan with traditional DB and DC retirement plans, describing the different retirement incentives they create for members. In Section 3, we describe our sample of retirement-eligible members and retirees. In Section 4, we calculate the *ex post* cost of the PERS pension plan to PERS employers and then simulate *ex ante* costs under alternative assumptions about market returns.

In Section 5, we use individual-level data to study the effect of retirement incentives, member characteristics, and peer effects on the retirement-timing decision. In Section 6, we conclude.

2. Institutional details

PERS uses benefit formulas drawn from both DB and DC pension plans. Before describing how retirement benefits are calculated in PERS, we contrast PERS' DB and DC elements with those of traditional DB and DC pension plans.

2.1. Comparing PERS' DB elements to a traditional DB pension plan

In a traditional DB plan, the retirement benefit is determined by a formula that is based on the member's earnings history and years of service, rather than on the performance of the underlying investment portfolio. Given this formula, the only uncertainty that a member faces with respect to her retirement benefit after t years of service is the level of her future salary.

PERS offers a typical DB retirement benefit, which is the product of four inputs:

$$DB = \text{Payout Factor} \times \text{Final Salary} \times \text{Years of Service} \times \text{Early Retirement Factor}.$$

A member can affect the final three inputs through her choice of retirement date. The monthly benefit is increasing in both the member's monthly salary before retirement and the number of years of service. *Final Salary* is typically the member's average monthly salary over the past 36 months of covered employment, and *Years of Service* is the number of months that the member contributed into PERS divided by 12. For a member who chooses to retire before the plan's stipulated "normal" retirement age, the monthly retirement benefit is decreasing in the number of years until the member reaches the normal retirement age. *Early Retirement Factor* equals zero percent until the member reaches the early retirement age at which point it becomes positive. In each subsequent year, it increases by 8% until reaching 100% at the member's normal retirement age. For general service members, the early retirement is age 55, but the normal retirement age depends on when they were hired. It is 58 (or earlier with 30 years of service) for "Tier 1" members hired before January 1, 1996, and it is 60 (or earlier with 30 years of service) for "Tier 2" members hired between January 1, 1996 and August 28, 2003. (We describe "Tier 3" benefits, for members hired after August 28, 2003, in the appendix.) For police and fire officers, the early retirement age is age 50 and the normal retirement age is 55 (or earlier with 25 years of service). The *Payout Factor* is 1.67% for general service members and 2.00% for police and fire officers. After full careers, the general service members earn replacement rates of 50.1% (= 0.0167×30) while police and fire officers earn replacement rates of 50.0% (= 0.02×25).

2.2. Comparing PERS' DC elements to a traditional DC pension plan

In a traditional DC pension plan, such as a 401(k) plan, contributions are made into an individual account by the member and, in many cases, by the employer. The member's final retirement account balance depends on the sizes of these contributions and on the returns earned on investments within the individual account. In contrast to a typical DB plan, the member bears all financial market risk. Upon retirement, a member can use the account balance to purchase a life annuity from an insurance company, at a price that depends on both the retiree's expected life expectancy and the prevailing risk-free rate of return.

PERS' DC elements differ in several ways from a traditional DC plan. PERS offers only two investment choices: the "regular" account and the "variable" account. Both accounts are managed by Oregon's Office of the State Treasurer. Members are permitted to allocate 0%, 25%, 50%, or 75% of their member contribution to the variable account, which invests in equity. The remaining member contribution, and all of the employer contribution, is allocated to the regular account, which invests in a mixture of debt and equity. Tier 1 members receive market returns in the variable account, but a minimum annual return of 8% in the regular account. Tier 2 members receive market returns in both accounts.

PERS calculates DC benefits using the following formula:

$$DC = 2 \times \text{Member's Account Balance} \times \text{Actuarial Equivalency Factor}.$$

The member's account balance depends on how her employee contributions are allocated across the regular and variable accounts and on the annual returns credited to each account. To capture the accumulated value of employer contributions, which are not reflected in the member's account balance, PERS doubles the member's account balance at retirement. During the period when PERS calculates member account balances using stale returns allows, the choice of retirement month can affect the level of the DC benefit.

The actuarial equivalency factor (AEF) is an age-based, gender-neutral annuity factor that Chalmers and Reuter (2002) show to be better than actuarially fair. There are two sources of this generosity. First, during most of our sample period, PERS bases its AEFs on life expectancy estimates from 1978. In contrast, the AEFs available from life insurance companies are periodically updated to reflect prevailing mortality risks. Second, PERS' AEFs are based on the assumption that PERS can earn a risk-free rate of return of 8 percent every year whereas the AEFs available from insurance companies are based on market rates, such as the prevailing yield on

10-year U.S. Treasury notes. Because the yield on 10-year U.S. Treasury notes trends down from 7.94% on January 2, 1990 to 4.26% on December 31, 2003 and mortality risk declines over the same period, the life annuity payments PERS pays are significantly higher than the life annuity payments that could be purchased from insurance companies. Retiring before the normal retirement age results in benefits being calculated using a lower AEF, but there is no explicit penalty for early retirement.

2.3. Calculating retirement benefits in PERS

We describe PERS as a combination pension plan because it offers both DB-style and DC-style retirement benefits. PERS uses the following three benefit formulas, and it pays each retiring member the largest benefit for which she is eligible:⁴

- (1) **DB** = $1.67\% \times \text{Final Salary} \times \text{Years of Service} \times \text{Early Retirement Factor}$,
- (2) **DC** = $2 \times \text{Member's Account Balance} \times \text{Actuarial Equivalency Factor}$,
- (3) **DCDB** = $1.00\% \times \text{Final Salary} \times \text{Years of Service} \times \text{Early Retirement Factor}$
 $+ 1 \times \text{Member's Account Balance} \times \text{Actuarial Equivalency Factor}$
 $= 0.6 \times \text{DB} + 0.5 \times \text{DC}$.

Members who contributed to PERS before August 21, 1981 are eligible to receive the maximum benefit calculated using all three formulas. Other members are eligible to receive the maximum benefit calculated using the DB and DC formulas. Although many public employers allow employees the *ex ante* choice of participating in either a DB plan or a DC plan, PERS is unusual in giving members the highest *ex post* benefit.⁵ PERS is portable between PERS employers, and members vest after either five years of service or the attainment of age 50.

The benefit formulas stated above assume that members choose to receive all of their retirement benefits in the form of life annuity payments.⁶ Members can instead choose to receive a lump sum payment based on the value of their account balance and lower life annuity payments. Chalmers and Reuter (2012) study this choice and find that 15% of PERS retirees choose the par-

⁴ These formulas apply to general service employees. Payout factors in the *DB* and *DCDB* formulas are higher for police and fire officers (2.00% versus 1.67% in equation (1) and 1.35% versus 1.00% in equation (2)).

⁵ Brown and Weisbenner (2012) study the choice between DC and DB plans within the State Universities Retirement System of Illinois. Based on a survey of other plans, they write “our best estimate is that approximately half of all states offer at least a subset of higher education employees a choice between a DB and a DC system.”

⁶ The benefit formulas also assume that the member chooses a single life annuity, which ends when the retiree dies. PERS adjusts monthly payments downward when the member instead chooses a joint life annuity.

tial lump sum option.

2.4. Employer pension costs

For a traditional DC plan, the employer's pension costs are limited to its retirement contribution and any administrative costs associated with offering the plan. For a traditional DB plan, the employer is liable for the expected present value of promised retirement benefits. By offering multiple benefit formulas, some of which depend on realized equity market returns, PERS is offering members a guaranteed DB benefit and the chance to receive instead a higher DC benefit. Therefore, PERS is liable for the expected present value of the promised DB retirement benefits *plus* the expected cost of providing higher benefits when equity market returns have been high.

PERS is funded by employers. Employers contribute six percent of gross wages to the regular and variable accounts, as directed by the employee.⁷ Additional employer contributions are assessed by PERS to meet the plan's funding needs, and they are invested exclusively in the regular account. Irrespective of the benefit formula used to calculate a member's benefit, benefit payments come from the regular and variable accounts. When the combined value of these accounts is below the expected present value of the retirement benefits owed to current and future retirees, the plan is underfunded and employer contribution rates are increased.

The way that employee and employer contributions are invested is too risky to match PERS' obligations under the DB benefit formula but, in some cases, not risky enough to match its obligations under the DC benefit formula. Consider a member retiring under the DB benefit formula. Because this situation is most likely to arise when realized equity market returns have been low, the accumulated value of employee and employer contributions is likely to be lower than if they had been invested in risk-free assets, resulting in fewer plan assets with which to pay the DB benefits. Now consider a member retiring under the DC benefit formula. PERS doubles the member's account balance before applying the annuity factor. The implicit assumption is that, over the member's career, employer contributions have grown at the same rate as employee contributions. However, when employee contributions are directed to the variable account this assumption is violated. The larger the difference between the member's account balance and the accumulated value of the employer contributions, the larger the resulting underfunding.

⁷ Although we refer to this amount as the "employee contribution", PERS employers agreed to make the ongoing contributions on behalf of employees in lieu of increasing nominal wages in 1979.

There are three other potential sources of underfunding built into PERS. First, the fact that Tier 1 members earn a minimum annual return of 8% in the regular account can result in members' account balances growing faster than PERS' regular account. Second, the use of stale returns to calculate account balances until January 2000 provides members with the option to retire when stale returns increase member account balances above current their market value and delay retirement otherwise. (Stanton (2000) shows that the option to exploit stale returns within 401(k) plans is valuable.) Finally, PERS' annuity factors generate higher benefit payments than can be purchase from insurance companies.

2.5. Retirement incentives differ across DB and DC benefit formulas

In their seminal paper, Stock and Wise (1990) contrast the retirement incentives embedded in traditional DB and DC retirement plans. They demonstrate that the DB retirement plan gives a member a stronger incentive to work until she is eligible for normal retirement benefits, but a weaker incentive to continue working thereafter. The intuition for their finding is that once a member with a DB retirement plan is eligible for normal retirement benefits, increases in the expected present value of retirement benefits from an additional year of service are typically dominated by the prospect of receiving one less year of retirement benefits. This is because the DB benefit does not adjust for the member's life expectancy. In contrast, for a member with a DC retirement plan, there is no explicit early retirement penalty to avoid through an additional year of labor, but the factors used to convert the DC account balance into life annuity income increases monotonically with the member's age. For these same reasons, we expect PERS retirements under the DB benefit formula will be more closely related to the early retirement penalty and normal retirement age than PERS retirements under the DC benefit formula.

Consider a member who will retire under the DB benefit formula and who is eligible for early retirement benefits in month t (at age 55, with 27 years of service) and normal retirement benefits in month $t+36$ (at age 58, with 30 years of service). To determine the effect of each additional year of employment on the member's retirement benefits, we can compare the initial retirement benefits in months $t-1$, t , $t+12$, $t+24$, $t+36$, and $t+48$. We state the initial monthly retirement benefit as a replacement rate, which measures the monthly benefit as a fraction of the member's final average monthly salary. The replacement rate jumps from 0% to 34.3% in month t when the member becomes eligible to receive early retirement benefits, and it rises rapidly thereafter as the member moves from early retirement to normal retirement.

Month	Retirement Eligibility	Years of Service		Early Retirement Factor		DB Payout Factor		Replacement Rate (RR)	Percentage Increase in RR
t-1	Ineligible	26.9	×	0%	×	1.67%	=	0.0%	0%
t	Early	27.0	×	76%	×	1.67%	=	34.3%	∞
t+12	Early	28.0	×	84%	×	1.67%	=	39.3%	14.6%
t+24	Early	29.0	×	92%	×	1.67%	=	44.6%	13.4%
t+36	Normal	30.0	×	100%	×	1.67%	=	50.1%	12.4%
t+48	Normal	31.0	×	100%	×	1.67%	=	51.8%	3.3%

Once the member is eligible for normal retirement benefits, however, the replacement rate only increases by 1.67 percentage points per additional year of service. The percentage change of the replacement rate (last column) highlights the differential retirement incentive. The replacement rate increases by around 13% each year during the early retirement period, but it only increases by around 3% each year after reaching the normal retirement age. Layered on top of this is the effect of age: members receive the benefit for fewer years when they retire at older ages, which reduces their effective benefit by about eight percent per year. Consider, for example, two members who retire under DB benefits after 30 years of service with identical earning profiles. The younger member will get the same initial benefit as the older worker, but she will get it for more years by virtue of her younger age.

A member who retires under the DC benefit formula faces different incentives. On the one hand, the DC formula does not have an early retirement penalty, which might encourage comparatively early retirements. On the other hand, the AEF increases with age, which might encourage comparatively delayed retirements. Potentially overshadowing these incentives, especially for Tier 2 members who do not enjoy the minimum annual return of 8% in their regular accounts, is the possibility that market returns will lower the member's account balance and, thereby, her DC benefits.

2.6. Retirement incentives embedded in PERS pension plan

In the prior subsection, we focus on variation in retirement incentives across the DB and DC benefit formulas. In this subsection, we focus on one source of exogenous variation in the level of the DB benefits and three sources of exogenous variation in the level of the DC benefits. We use these sources of variation to study the average sensitivity of the retirement-timing decision to retirement incentives. We also use them to document that while the majority of members appear to successfully time their retirements to exploit these features, a minority of members

would have received higher benefits had they made small changes to their retirement date.

2.6.1. Early retirement penalties and DB benefits

The first source of exogenous variation arises from a change on January 1, 1997 in how often the early retirement penalty is updated. Before 1997, the penalty is updated once per year, in the member's birth month. Beginning in 1997, the penalty is updated monthly. Consequently, between 1990 and 1996, a member who retires one month before her birth month receives DB benefits that are 92.0% of the DB benefits that she would receive if she waited to retire in her birth month. Between 1997 and 2003, the 8% percent penalty is spread evenly over 12 months, so that the corresponding number is 99.3%. Primarily, this change eliminates the possibility that a member retiring under the DB benefit formula would earn a significantly higher benefit by delaying retirement one month. However, by reducing the disincentive to retire in the months immediately before the normal retirement age, this change has the potential to increase early retirements under the DB benefit formula between 1997 and 2003.

2.6.2. Stale returns and DC benefits

The second source of exogenous variation in retirement benefits arises from PERS' use of stale returns between 1990 and 1999 to calculate the member's account balance. Every April, PERS provides members with a statement that reports the contributions and returns credited to the member's account over the prior calendar year, as well as the account balance at the end of the prior calendar year. Prior to January 1, 2000, the timing of this report reflected the fact that PERS did not finalize annual returns for the regular and variable accounts in year $t-1$ until March in year t . For example, when a member retires in March of year t , PERS uses the newly finalized annual returns for year $t-1$ to determine the returns between January and March of year t . In this case, there is considerable uncertainty about whether these stale returns are higher or lower than the returns that PERS will finalize in March of year $t+1$. On the other hand, when a member retires in February of year t , PERS uses the finalized annual returns for year $t-2$ to determine the returns between January of year $t-1$ and February of year t . In this case, the stale returns are over one year old, and the member's retirement incentive in February depends on how the finalized returns for year $t-2$ compare to the member's forecast of the not-yet-finalized returns for year $t-1$.

To capture the retirement incentives due to stale returns, we define DC_delta as the monthly return that the member receives from retiring in month t instead of retiring in the month that PERS next finalizes returns for the regular and variable accounts (i.e., March 1990, March

1991, ..., and March 1999). For example, in February 1998, this is the percentage change in the member's account balance from retiring under the stale returns available in February versus retiring under the finalized returns in March. In Figure 1, between 1990 and 1999, we plot the median, minimum, and maximum fluctuations in retirement benefits due to stale returns (*DC_delta*) available to members who would retire under the DC benefit formula in January or February. The fact that *DC_delta* ranges from -4.6% to 3.7% in February 1998 reflects the fact that stale returns affected the regular and variable accounts differently. In the regular account, the finalized 1996 return (21.00%) was higher than the finalized 1997 return (18.70%), providing members invested primarily in the regular account with an incentive to retire in February. However, in the variable account, the finalized 1996 return (21.06%) was lower than the finalized 1997 return (28.87%), providing members primarily invested in the variable account with an incentive to delay retirement until March.

2.6.3. Actuarial equivalency factors and DC benefits

The final two sources of variation in retirement incentives come from changes to the AEF tables used to determine DC benefits. Before January 1, 1997, PERS updated the AEF used to determine a member's retirement benefits only once a year, in her birth month. During this regime, DC benefits could be as much as 4.10% lower (1.75% lower at the median) if the member chose to retire in the month immediately before her birth month rather than in her birth month. On January 1, 1997, PERS adopted AEF tables that increased with each month of age. As with the change from annual to monthly early retirement penalties, this change eliminated the incentive for a member to delay retirement until her birth month. Nevertheless, the new monthly factors were still based on mortality tables from 1978.

The more significant source of variation comes from PERS' adoption in 2003 of AEF tables based on then-current forecasts of retiree life expectancy. The new AEFs were between 1.4% and 17.8% lower than the old AEFs, with the largest decreases for older retirees. For members between the ages of 58 and 65, AEFs decreased between 5.8% and 10.2%. Although Oregon House Bill 2004 requiring the use of updated AEFs was not signed into law until May 9, 2003, the plan to adopt new AEFs was well publicized beginning in September 2010, when

PERS adopted the mortality assumptions on which the new AEFs were based.⁸ This change created a strong incentive for members who expected to receive DC benefits to retire before the new AEFs took effect on July 1, 2003. It is worth noting, however, that because the new AEFs continued to assume a constant annual risk-free rate of return of 8%, they remained significantly more generous than the AEFs available each year from life insurance companies.

We define AEF_delta as the change in DC retirement benefits that a member (eligible for DC benefits) would receive if she retired now rather than waiting for the next known change to her AEF. It is measured as a monthly return, from the date of the possible retirement to the date of the change. Between January 1990 and December 1996, the next known change occurs in the member's birth month or in January 1997, whichever comes first. During this period, AEF_delta measures the cost to members eligible for DC benefits of retiring in the months leading up to her birth month. In contrast, the large positive returns between January 2003 and June 2003 measure the growing incentive for members retiring under DC to retire before the change to the new AEFs on July 1, 2003. For the median member eligible to retire under DC, average monthly life annuity payments are 5.3% higher if she retires in June 2003 instead of July 2003. However, the incentive to retire in June 2003 ranges from 2.7% to 21.1%, with the strongest incentives for the oldest members. In Figure 1, we plot the average, minimum, and maximum fluctuations in retirement benefits in 2003 due to the predictable changes in AEFs (AEF_delta).

3. Data

In 2006, PERS held nearly \$56 billion in assets, making it the 22nd largest public or private pension fund in the country. PERS covers approximately 95% of all non-federal public employees in Oregon. Participating employers include all state agencies, universities, and school districts; and almost all cities, counties, and other local government units. Administrative data obtained from PERS allow us to calculate PERS member i 's retirement benefits under the DB,

⁸ What was not well publicized was the "look-back" provision created for members retiring on or after July 1, 2003. This provision compares benefits calculated using the member's current account balance and new AEFs to benefits calculated using the member's account balance as of June 30, 2003 and old AEFs, and then pays out the higher of the two benefits. According to State Representative Tim Knopp, "We don't want to force employees into retiring early because they are afraid their retirement benefits will shrink." However, the only reference that we could find to the look-back provision in the main state newspaper, *The Oregonian*, was in the article "Committee Endorses PERS Changes" published on February 21, 2003, which contained Mr. Knopp's quote. Furthermore, under this provision the larger the decline in AEFs, the longer that a member's account balance is effectively frozen at June 2003 levels, reducing the member's ability to increase future pension benefits through continued employment.

DCDB, and DC benefit formulas if she chooses to retire in month t . These data also allow us to determine when member i becomes eligible to receive PERS retirement benefits and, when member i is currently employed, the PERS employer code.

Our sample includes 62,953 unique members who work for a PERS-covered employer and are eligible to retire at some point between January 1990 and December 2003.⁹ Members enter our sample when they become eligible to retire or, if they are already eligible to retire, in January 1990. They exit our sample when they stop working for a PERS-covered employer. Table 1 Panel A provides annual summary statistics for all retirement-eligible members. The average replacement rate, calculated as the monthly benefit that the member would receive upon retirement divided by the member's salary over the prior 12 months, increases from 27% in 1990 to 39% in 1998, and then decreases to 33% in 2003. The (unreported) unconditional probability of retirement in any given month among the individuals represented in Panel A is 1.46%.

Table 1 Panel B provides annual summary statistics for the 35,129 members who choose to stop working for a PERS-covered employer and immediately begin collecting their PERS retirement benefits. Comparing Panels A and B, we see that retirees have replacement rates that are 24-68% higher, and three to seven more years of service than their non-retiring peers. The average retirement age falls from 60.6 years at retirement in 1990 to 58.5 years old in 2003 while, over the same period, average years of service increase from 18.9 to 21.2 years. The time-series correlation between the average replacement rate and the average retirement age is -0.95, suggesting that higher retirement benefits allow for earlier retirements.

In Figure 2, we graph the fraction of retirement eligible members who retire each year. We distinguish retirements by members who are eligible to receive normal DB benefits (i.e., for whom there is no early retirement penalty in the DB benefit calculation) from early retirements. We find that annual retirements are lumpy. Among those eligible for normal benefits, the fraction who choose to retire ranges from 11.6% in 2000 to 28.8% in 2003. Among those eligible

⁹ Our sample includes all PERS participants except for legislators and judges because PERS declined to provide us with the data needed to calculate their potential retirement benefits. Because our peer effects analysis requires knowing the geographic location of employees within firms (i.e., the city) and PERS does not have that information, we exclude members working for employers with multiple locations (for example, PERS data does not allow us to determine whether an employee of the Oregon University System works at Eastern Oregon University, the University of Oregon, or another location). This filter allows us to maintain a consistent sample across all tables, but it reduces the number of retirement-eligible-employee-year observations from 303,689 to 245,808 (a loss of 19.2%), and the number of retirees from 42,862 to 35,129 (a loss of 18.1%). Nevertheless, the summary statistics for the full sample of employees are virtually identical to those in Table 1.

for early benefits, the fraction who choose to retire ranges from 4.3% in 1996 to 14.9% in 2003. In Figures 3a and 3b, we graph the fraction of retirements that occur in each month, for different sample periods. Figure 3a focuses on retirees whose benefits are calculated using the DC benefit formula, while Figure 3b focuses on retirees whose benefits are calculated using the DB benefit formula. Retirements by teachers at the end of the school year help to explain the retirement spikes in June. The fact that retirement spikes in February are limited to retirees receiving DC benefits, and to the period 1990 to 1999, imply that some members timed their retirements to benefit from stale returns.

4. Estimating the effects of PERS' design on employer costs

In this section, we estimate the incremental costs that Oregon public employers incur through PERS' provision of three benefit formulas instead of just the DB benefit formula. Our first set of estimates is based on *ex post* realizations of the data during our sample period. Our second set of estimates is based on *ex ante* Monte Carlo simulations that use only what was known in 1990.

4.1. Realized employer costs

In Table 2, we compare the actual benefits received by retiring PERS members to two counterfactual benefits. First, we re-calculate each retiree's replacement rate using only PERS' DB benefit formula. Second, we re-calculate each retiree's replacement rate using all three benefit formulas, but replacing PERS' better-than-actuarially-fair annuity factors with those available from insurance companies. In both cases, we are benchmarking PERS against mechanically less generous alternative pension plans while holding members' retirement-timing decisions constant. Our goal is to determine how much of the *ex post* benefit to members—and the associated *ex post* cost to employers—was due to the use of multiple benefit formulas, and how much was due to the use of better-than-actuarially-fair annuity factors when calculating DC benefits. Note that the sample is slightly smaller than in Table 1 Panel B. Our data on the annuity factors available each year from insurance companies are limited to ages 50 through 70, and 632 of the retirees in Table 1 Panel B are either younger than 50 or older than 70.¹⁰

The average retiree in Table 2 receives 51.2% of her final monthly salary as a retirement benefit. The average counterfactual replacement rate based on the DB benefit formula is 33.4%.

¹⁰ Our data on annuity factors come from TIAA. They were previously used in Chalmers and Reuter (2012).

Therefore, actual retirement benefits are 53.5% higher, on average, than the counterfactual “DB only” benefits. Furthermore, these higher benefits are widespread. Comparing each retiree’s “DB only” and actual benefit, we find that 87.5% of retirees earn higher benefits because of the use of multiple benefit formulas.

The average replacement rate calculated using counterfactual annuity factors is 41.3%, which falls approximately halfway between that calculated using retiree’s “DB only” and actual benefits. The fact that 70.6% of retirees would have earned higher than “DB only” benefits if PERS had used annuity factors available from insurance companies highlights the significant *ex post* costs associated with offering multiple benefit formulas. Nevertheless, we find that more than half of the incremental pension costs are due to PERS’ use of better-than-actuarially-fair annuity factors.

Providing retirees with the maximum benefits for which they are eligible also increases dispersion in realized benefits (holding inputs like salary and years of service constant). Average replacement rates under the DB formula only range from 30.8% in 2000 to 35.2% in 2002. Yet, the ratio of actual benefits to “DB only” benefits ranges from 113.9% in 1990 to 180.9% in 2000. Because PERS does not know the benefit formula that will ultimately be used to calculate each member’s retirement benefits, it cannot easily manage plan assets to meet its *ex post* pension obligations. Investing in risky assets makes it easier to cover DC benefits when equity market returns have been high, but harder to cover DB benefits when equity market returns have been low.

In the appendix, we estimate PERS’ pension liabilities for these 34,497 retirees to be \$11.9 billion based on the “DB only” benefits, \$14.8 billion based on the “Counterfactual AEF” benefits, and \$18.3 billion based on their actual benefits. The difference between the actual and “DB only” benefits is \$6.3 billion, or \$183,737 per retiree. By way of comparison, PERS estimates the difference between pension liabilities and pension assets to be \$17 billion in 2003 and \$15 billion in 2009, and Novy-Marx and Rauh (2011), using a discount rate similar to ours, estimate the difference to be \$38 billion in 2009. Hence, our estimate of \$6.3 billion in incremental pension costs is economically significant, especially given that it applies to a subset of current retirees and ignores the retirement benefits owed to current employees.

4.2. Expected employer costs

In this subsection, we simulate PERS’ expected incremental costs of offering the DC benefit formula. Our Monte Carlo simulations follow the PERS policies and practices that were

in effect for new members at the start of our sample period, including Tier 1 benefit rules, stale member-account returns, and no DCDB benefit formula. Our measure of PERS' expected incremental costs is the ratio of simulated DC benefits to simulated DB benefits. For each age between 20 and 57, we generate 50,000 trials of the DC and DB benefits for a sample member who joins PERS at that age and works through retirement. We report statistics describing the *ex ante* distribution of the DC/DB ratio for each initial-age cohort.

We make a number of assumptions in our simulations. First, we assume that no member leaves PERS employment until retirement. Second, we assume that wages grow by 3.59% per year.¹¹ Third, we assume that the contribution to the member account is 6% and that it happens at the start of each year (the doubling of the member account at retirement in equation (2) effectively makes the contribution rate 12% while working). Fourth, we incorporate both the guaranteed annual return of 8% and the fact that members and employers share the gains when market returns are greater than 8%. Specifically, we randomly draw from the historical distribution of S&P 500 annual returns between 1926 and 1989 for each member-year, and we set the portfolio return equal to $\text{MAX}(8\%, 8\% + 0.5 \cdot (\text{SP500} - 8\%))$. Fifth, to capture the effect of stale returns, we calculate DC benefits using the higher of this year's return and last year's return. Finally, we assume that members use the following rule to determine when to retire. For each year between the early retirement age (55) and the normal retirement age (58), the member compares the annuitized DC benefit at the current age with the DB benefit at the normal retirement age. If the DC benefit is larger, the member retires immediately. Otherwise, the member waits one year and re-evaluates the two alternatives until mandatory retirement at the normal retirement age. This rule implies that the DC/DB ratio, for initial ages 26 and older, might have fewer working years embedded in the numerator than in the denominator. Allowing early retirements under DC benefits but not DB benefits is consistent with the realized retirements in our database, but we do not know whether PERS anticipated this in 1990. In unreported simulations, we confirm that the DC/DB ratio is larger when we remove the possibility of early DC retirements.

We present our simulation results in Figure 4. The horizontal axis depicts the initial-age

¹¹ The simulated DC/DB ratio is sensitive to the choice of wage growth. A higher growth rate increases the DC benefit due to higher account contributions throughout the career. However, because the DB benefit is based on the final average salary, it increases the DB benefit by even more. Because we do not know PERS' 1990 model of wage growth, we use the annualized growth rate of retirement-eligible members within our sample period. Future research could model the earnings growth of career public-sector employees.

cohort, and the primary vertical axis depicts the DC/DB ratio at the simulated retirement date. The figure plots the mean, 25th percentile, 50th percentile, and 75th percentile of the DC/DB ratio distribution. For example, the mean DC/DB ratio for members who join PERS when they are 39 years old is 1.35. This implies that members who first join PERS at age 39 could expect DC benefits at age 55 that are 35% larger than the DB benefits at age 58.¹² The DC/DB ratio is 1.18 at the 25th percentile; it is 1.31 at the 50th percentile; and it is 1.48 at the 75th percentile. The secondary vertical axis (the right hand side) plots the fraction of simulated members whose DC benefits exceed their DB benefits at retirement.

The DC/DB ratio increases for each year between age 20 and 25 because the DC benefit is age adjusted through the AEF while the DB benefit remains fixed for these members who all retire after 30 years of service. The DC/DB ratio falls for older workers, all of whom retire between age 55 and age 58, because they have fewer years to earn the more generous DC benefits. All simulated members retire with DC benefits through initial age 34. For initial-age cohorts 35 through 49, at least 97% of the simulated members retire with DC benefits. Thereafter, a sharply increasing fraction of the simulated members take the DB benefits. The intuition for these results is that for members with long careers ahead of them, 12% effective retirement contributions, guaranteed annual returns of at least 8%, and better-than-actuarially-fair annuity factors combine to generate large expected DC benefits.

We conclude from our simulations that the *ex post* incremental retirement benefits that we document above are consistent with what could have been modeled using *ex ante* data in 1990. A model that included more details such as the DCDB benefit formula and the possibility to have longer careers would increase the estimated *ex ante* costs.

5. The effects of PERS' design on members' retirement-timing decisions

Our analysis of the member retirement-timing decision proceeds in five steps. First, motivated by the predictions in Stock and Wise (1990), we present evidence on the retirement ages of members receiving retirement benefits under the three different benefit formulas. Second, we use variation in the level of retirement benefits within the DB and DC benefit formulas from one month to the next to identify the numbers of members who appear to have successfully and un-

¹² The 39 year old new member could expect to retire under the DB benefit at age 58 with 19 years of service, earning a replacement rate of $31.7\% = 1.67\% \times 19$. The simulation shows that the same member could expect to retire under the DC benefit at age 55 with 16 years of service, earning a replacement rate of $42.8\% = 1.35 \times 31.7\%$.

successfully exploited this variation. Third, we estimate a baseline model to predict whether a retirement-eligible individual will choose to retire in month t . We use individual-specific information such as age, gender, job type, projected retirement benefit, and *ex post* mortality measures, as well as the exogenous variation in retirement incentives described above. The baseline model allows us to test whether members respond to the different retirement incentives generated by the combination structure, and it allows us to quantify the effects of those incentives.

Fourth, to test for peer effects in the retirement-timing decision, we add the actual retirement decisions of an individual's coworkers to the baseline model. To help distinguish peer effects from alternative explanations such as unobserved heterogeneity among employers, we include controls that vary at the employer-date level, such as the fraction of non-retirement eligible members leaving the employer in month t . In addition, we instrument coworker retirements with several sources of exogenous variation in coworker retirement incentives. Finally, to determine whether peer effects reflect the diffusion of information about retirement incentives, we test whether members are disproportionately more likely to respond to retirement incentives when more of their coworkers face the same incentives.

5.1. Retirement ages and retirement benefit formulas

Following Stock and Wise (1990), we predict that retirements under the DB benefit formula will be more sensitive to the size of the early retirement penalty and attainment of the normal retirement age than retirements under the DC benefit formula. In Table 3, we find strong support for this prediction. In Panel A, we report the distribution of retirement ages for members who receive benefits under the DB, DCDB, and DC formulas. To facilitate comparisons across benefit formulas, we focus on 29,554 retirees for whom the early retirement age is 55 and the normal retirement age is 58. (We begin with the sample of retirees described in Panel B of Table 1, but then exclude 2,385 police and fire officers, 632 retirees for whom the normal retirement age is 60, and 2,568 retirees for whom the normal retirement age is 58 but who qualify for normal retirement benefits based on their years of service rather than their age.)

We find that retirees receiving DB and DCDB benefits are five to seven times more likely to retire at age 58 than they are at age 55. In contrast, retirees receiving DC benefits are more likely to retire at age 55, when they first become eligible to collect retirement benefits, than at age 58. At age 55, we observe 17.2% of the retirements under DC versus only 3.1% under DB. By age 58, we observe 50.6% of the retirements under DC versus only 29.1% under DB.

These patterns are consistent with the higher benefits made possible by the DC benefit formula allowing members to retire at earlier than normal ages. Another interesting pattern is that retirements under DB are almost twice as likely to happen at age 62 as retirements under DC. One interpretation is that receiving a higher-than-DB benefit via the DC benefit formula makes it easier for a member to finance her retirement before she is eligible to collect benefits from Social Security. Overall, the patterns in Panel A suggest that the combination pension plan increases early retirements, both by increasing expected retirement benefits and by eliminating the penalty associated with early retirements.

In Panel B of Table 3, we explicitly classify each of the retirees described in Panel B of Table 1 as “retiring early” if she would have been subject to an early retirement penalty under the DB benefit formula. We then calculate the fraction of early retirements for each of the three benefit formulas. As predicted, early retirements are much less common under the DB benefit formula than under the DC benefit formula. This is true both overall (14.0% versus 33.8%) and within the 1990-1996 and 1997-2003 subsamples. As expected, between 1996 and 1997, when early retirement penalty calculations change from annual to monthly updates, there is a modest increase in early retirements under the DB benefit formula, from 8.1% to 13.7%, a difference which is statistically significant at the 10-percent level.

5.2 How many retirees respond to exogenous variation in their retirement incentives?

In this section, we use discrete, predictable changes in the level of retirement benefits between months t and $t+1$ to ask what fraction of retirees appear to be responding to this high-frequency variation in their retirement incentives. Using a narrow but operational definition of what constitutes a mistake, we also ask how many members would have been better off if they had retired one month earlier or one month later than their chosen retirement month.

5.2.1. Early retirement penalty update frequency changed from annual to monthly

Between 1990 and 1996, the early retirement penalty is updated for each member in her birth month. This creates a disincentive for a member whose retirement benefit will be reduced by the early retirement penalty to retire in the month immediately before her birth month. We observe 197 retirements during this period where the value of the DB benefit is reduced by an early retirement penalty. Surprisingly, we find that members are unwilling to postpone their retirements until their birth months: 19 of these retirements occur in the month immediately before

the member's birth month, which is higher than the 16 that would occur if the retirements were spread evenly across the year.

5.2.2. Actuarial equivalency factor update frequency changed from annual to monthly

Between 1990 and 1996, the actuarial equivalency factor is updated once per year in the member's birth month. This creates a disincentive for a member whose retirement benefit depends on the AEF to retire in the month before her birth month. We observe 5,323 retirements during this period in which benefits are determined using the DC benefit formula, 398 of which occur in the month immediately before the birth month (which is slightly less than the 444 retirements that would occur were they spread evenly across the year). The DC benefits of these 398 retirees are 2.2% lower, on average, than they would have been if the members had retired one month later. Additionally, we find that the fraction of members retiring in their birth month falls from 18.4% when the AEF is updated annually, to 14.0% when it is updated monthly. Together, these findings suggest that, between 1990 and 1996, some members delayed retirement until their birth month to receive a higher AEF.

5.2.3. Using stale returns to determine members' account balances

Through December 1999, PERS used stale returns to determine the value of each member's account balance. In Figures 3a and 3b, we find strong evidence that retirement-timing decisions respond to the use of stale returns. Between 1990 and 1999, the fraction of members retiring in February is 31.3% among those receiving DC benefits but only 7.5% among those receiving DB benefits. And, between 2000 and 2002, when PERS no longer uses stale returns, the fraction of members receiving DC benefits who retire in January falls to 6.1%.

In Table 4 Panel A, we divide the 4,100 members who retired in February into two groups. The first group benefited from retiring in February rather than March. Among these 3,552 retirees, the average value of *DC_delta* is 3.0% and the maximum is 28.6%. (Recall that *DC_delta* measures how much higher the member's account balance will be in February than in March because of the stale returns.) The other 548 retirees did not benefit from the stale returns. For this second group, the average value of *DC_delta* is -2.6% and the minimum is -20.5%. The fact that 13.4% of the members who retired in February between 1990 and 1999 were made worse off by stale returns helps to explain the relatively low coefficient on *DC_delta* that we estimate below in Table 5.

We also consider the 232 members who were eligible to retire in February but chose instead to retire in March. Within this sample, 142 earned higher average DC benefits because of this one month delay. Because we continue to measure *DC_delta* in February, the average value of -5.7% implies that these 142 retirees' DC benefits would have been, on average, 5.7% lower if they had retired in February. On the other hand, the 90 members who retired in March would have earned DC benefits that were 6.4% higher if they had retired one month earlier. This is a costly mistake; albeit one that affects few retirees.

In Panel B, we compare the characteristics of the 3,694 members who benefited from stale returns and the 638 who did not. Based on existing evidence that levels of financial literacy are lower among women (e.g., Lusardi and Mitchell (2007) and Lusardi and Tufano (2008)) and those with lower incomes (e.g., Campbell (2006) and Levy and Seefeldt (2008)), we expect the likelihood of a mistake to be higher for women and lower for those with higher incomes. Instead, we find fewer mistakes by women and no differences in final average salary. We also find significantly fewer mistakes by school district employees and slightly more mistakes by employees who die within 48 months of retirement (a rough proxy for health). The only other meaningful difference we find is that retirees who invested some of their employee contributions in the variable account are more likely to make mistakes. Because members are required to direct at least 25% of each retirement contribution to the regular account, members who invest in the variable account have two sets of stale returns to consider, making the net effect of stale returns on the member's account balance more difficult to determine.

5.2.4. Updating actuarial equivalency factors in July 2003

In July 2003, PERS updated its actuarial equivalency factors to reflect prevailing life expectancies. As Figure 2 shows, a significant fraction of retirement-eligible members retired each month between January 2003 and June 2003. If these retirements were in response to the changes in AEFs, then they should have been concentrated among members facing the DC and DCDB benefit formulas. Indeed, we find that 26.6% of the 13,864 members eligible for DC benefits and 33.2% of the 642 members eligible for DCDB benefits retire during this six month period, compared to only 6.7% of the 7,982 members eligible for DB benefits. Between July 2003 and December 2003, when the new AEF tables are being used, all three fractions fall, to 6.4%, 7.7%, and 1.8%.

Focusing on retirements under the DC benefit formula in the two months surrounding the

change in AEFs, we observe 728 in June but only 30 in July. We find that those retiring in June earned DC benefits between 2.8% and 21.1% higher than they would have been under the new AEFs; the average increase was 6.6%. The 30 members retiring in July would have earned DC benefits between 2.9% and 14.0% higher if they had retired under the old AEFs.

5.3. Baseline retirement-timing model

In Table 5, we use linear probability models to explain the retirement-timing decisions of retirement-eligible members. Because PERS retirement incentives can vary significantly from coworker to coworker and from month to month, the dependent variable equals one if member i retires from employer j in month t , and zero otherwise. (This contrasts with the existing literature, which focuses on predicting the year or age of retirement.) In column (1), we focus on the full sample of retirement-eligible members. In columns (2) through (4), we restrict the sample to the following three subsets: female members, active police and fire officers, and members whose birth month is month t . We multiply coefficient estimates by 100, so that 1.00 represents 1 percentage point. To allow for correlated behavior within employers, which we test for below, standard errors are clustered on employers.

In addition to the member characteristics and retirement incentives variables described below, we include fixed effects for each of the 34 ages (measured in years) between 46 and 79.¹³ We also include a separate fixed effect for each of the 168 months in our sample period (January 1990 through December 2003). However, because our sample combines school districts that operate on a nine-month schedule with employers that operate on a year-round schedule, we interact each date fixed effect with a dummy variable that indicates whether employer j operates on a nine-month schedule. These date-by-employer-type fixed effects allow us to control for the fact that school district members are more likely to retire in June at the end of the school year. More generally, by including date-by-employer-type fixed effects, we “remove” the average retirement effects due to PERS plan changes and any other time-specific event within our sample period. In other words, we use within-period, within-employer-type, within-age variation to estimate the coefficients in Table 5.

Because we predict that member i will be more likely to retire when her expected retirement benefits are more generous, we include two measures of generosity. The first is the frac-

¹³ Although we limit our sample to ages between 46 and 79, doing so only causes us to drop 27 members.

tion of member i 's current monthly income that she would receive each month from PERS in retirement.¹⁴ Consistent with our prediction, the coefficient on the replacement rate is positive and statistically significant at the 1-percent level. The estimated coefficient of 3.456 implies that a one-standard deviation increase in the replacement rate (0.245) increases the probability of retirement by 0.85 percentage points. This effect is economically large—the unconditional probability of retiring in a given month is only 1.46%.

Our second measure of generosity is the “option value” introduced by Stock and Wise (1990), which is a forward-looking measure that estimates the utility gain from deferring retirement until the optimal retirement time. The more that a worker gains from delaying retirement the less likely she is to retire today. We implement the Stock and Wise (1990) model by calculating the present value of a member's dollar wealth when retiring on the optimal date (including both labor and pension income) and subtracting the present value of a member's dollar wealth when retiring today.¹⁵ When the optimal retirement date is today, the difference between these numbers is zero. When the optimal retirement date is in the future, the difference between these numbers is strictly positive, and it measures the present value of the benefit of deferring retirement.¹⁶ The measure that we include in our regressions is divided by member i 's average annual salary over the past 12 months. The predicted sign is negative. The estimated coefficient is statistically significant in none of the four specifications.

Of more interest to us are the three variables that isolate the short-run retirement incentives (or disincentives) generated by the use of stale returns in the calculation of the member's account balance (DC_delta) and by changes in annuity factors (AEF_delta). Each variable measures the change in retirement benefits (as a monthly return) from retiring in month t relative to waiting for the updated annual returns or annuity factors to take effect. Therefore, the pre-

¹⁴ This is defined as the expected monthly retirement income that employee i would receive if she retired in month t scaled by her average monthly salary over the past 12 months

¹⁵ Variations of the Stock and Wise measure have been used by Samwick (1998), Chan and Stevens (2004), Chan and Stevens (2008), Coile and Gruber (2007), and others.

¹⁶ Our estimation requires several assumptions. We assume that annual wage growth is 2% and that the annual discount rate is 3%. PERS makes COLA adjustments to the benefit each August that is set at the smaller of Portland's CPI and 2%. Since Portland's CPI was rarely under 2%, we assume the annual adjustments would always be 2%. Consistent with prior research, we assume that members are risk averse and that members value retirement income more than labor income (i.e., members would rather not work). We pick the same parameter values as Samwick (1998). Specifically, we set $\gamma=0.75$ for risk aversion and $k=1.5$ for the preference for retiring. When $k=1.5$, members are indifferent between working to earn \$3 and retiring to collect \$2. Last, we forced members to retire by age 80 because PERS does not calculate AEFs beyond age 80. Given the very small number of members who actually chose to retire beyond age 80, this last assumption does not seem unreasonable.

dicted sign on each variable is positive.

Our stale return variable takes on non-zero values in January and February, between 1990 and 1999. Because prior equity market returns have been fully realized, these are the months when the incentive (or disincentive) associated with having DC benefits calculating using stale returns should be the clearest. The coefficient on the variable measuring retirement incentives in January and February is statistically significant at the 1-percent level, but its economic significance is modest. A one standard deviation increase is associated with a 0.36 percentage point increase in the probability of retirement.

To measure the effect of PERS only updating a member's actuarial equivalency factor in her birth month, we interact *AEF_delta* with a dummy variable indicating whether month *t* is between January 1990 and December 1996. The coefficient on this variable is statistically significant at the 5-percent level, but economically insignificant. A one-standard deviation increase in *AEF_delta* during this period increases the probability of retirement by 0.05 percentage points. To measure the effect of the new AEF tables in July 2003, we interact *AEF_delta* with a dummy variable that indicates whether month *t* is between January 2003 and June 2003. The coefficient on this variable is both statistically significant at the 1-percent level and economically significant. Here, a one standard deviation increase is associated with a 1.85 percentage point increase the probability of retirement. Note that since 1-percentage point increases in *DC_delta* and *AEF_delta* have the same effect on DC benefits, the fact that coefficients differ across these two measures is consistent with the retirement incentives due to stale returns being less well known than the retirement incentives due to the changing actuarial equivalency factors. The fact that 14.5% of the members who retire in February have negative values of *DC_delta* (i.e., appear to have made retirement-timing mistakes) also helps to explain the difference in magnitudes.

To study whether member retirement decisions are constrained by retirement eligibility rules, we introduce dummy variables to indicate whether member *i* became eligible for early retirement benefits in month *t*, in months *t-1* through *t-11*, or prior to month *t-11*, and to indicate whether member *i* became eligible for normal retirement benefits in months *t* or in months *t-1* through *t-11*. (The omitted category is being eligible for normal retirement for twelve or more months.) Similarly, to control for the possibility that members are more likely to retire in their birth month, we introduce a dummy variable that indicates whether month *t* is member *i*'s birth month. We find that individuals are much more likely to retire in a birth month (0.88 percentage

points) and in the first month that they are eligible for normal PERS retirement benefits (3.03 percentage points).

As *ex post* measures of health, we include dummy variables that indicate whether the member dies between months 1 and 12 and between months 13 and 48. Since we possess information on member deaths through the end of 2007, we are able to define these dummy variables for every retirement-eligible member in every year of our sample. To the extent that these future deaths are good proxies for relatively poor health today, the predicted signs on both coefficients are positive. Consistent with this prediction, both *ex post* mortality measures are economically significant predictors of retirement. A retiree who dies within the next 48 months is 1.023 percentage points (0.815 plus 0.208) more likely to retire today.

Other continuous variables include years of service, which is positively correlated with the retirement decision, and the unemployment rate within the county in month t , which is negatively correlated with the retirement decision in some specifications. For completeness, we also include dummy variables indicating whether member i is female; actively employed as a police or fire officer, eligible for Tier 2 pension benefits; or would receive benefits calculated under DC, DB, or DCDB (the omitted category).

When we restrict our sample to the subset of members who are female (column (2)) or active police and fire (column (3)), the estimated coefficients on the variables of interest are qualitatively similar to those found in the earlier specifications. Perhaps the most interesting difference is that police and fire officers are more likely than the other participants to retire in the first month in which they are eligible for normal PERS retirement benefits (10.12 percentage points versus an unconditional probability of 1.25 percent). When we restrict the sample to members who have a birthday in month t , we find greater sensitivity to the replacement rate and to variation in the level of the DC benefit due to the changing AEF tables.

5.4. Testing for peer effects in retirement-timing decision

Because we are interested in testing whether members learn about their retirement incentives from co-workers, we define peers as those people who work for the same employer and are eligible for retirement in the same month. In many cases, this gives relatively fine peer groups. For example, employers include individual school districts (e.g., Jackson County School District #1 and Jackson County School District #10), city employers (e.g., City of Madras and City of Klamath Falls), and fire districts (e.g., Rainier Fire Department and Keizer Fire Department).

Many of our employers are quite small and have only a few members (e.g., the Oregon Hazelnut Commission) while a few are quite large and have thousands of members (the largest is the Portland School District). In our empirical work, we exclude employers in months where the employer has fewer than two retirement-eligible members because we cannot test for peer effects when the PERS member has no retirement-eligible coworkers.

In Table 6, we extend our empirical specification to test for peer effects. With the notable exception of Brown and Laschever (2012), the existing retirement literature does not allow for peer effects in the retirement-timing decision. Our measure of peer retirements, *frac_retire*, is the fraction of a member's retirement-eligible coworkers (excluding herself) that retire from employer j in month t . Our test for peer effects is whether the probability that member i retires in month t is increasing in *frac_retire*. The decision to focus on retirements in month t (instead of in a particular year y or at a particular age) is driven by the within-year, time-varying retirement incentives in the PERS system.

In column (1), we add *frac_retire* to an extended version of the specification in column (1) of Table 5. The estimated coefficient is 27.024, which is both statistically significant at the one-percent level and economically significant. Interpreted as a peer effect, a one-standard deviation increase in the fraction of peers retiring (3.36 percent) increases the probability of retirement by 0.91 percentage points, which is large given that the unconditional probability of retirement in month t is 1.46 percent. Therefore, within our sample, there is a strong correlation between individual retirement decisions and average retirements within the same employer and month, even controlling for individual-level predictors of retirements, age fixed effects, and month-by-employer type fixed effects (which allow the likelihood that school employees retire in month t differs from the likelihood that non-school employees retire in month t). In fact, the estimated coefficients on the other variables—including member i 's short-run retirement incentives based on stale returns and changing actuarial equivalency factors—are almost identical to those estimated in Table 5, suggesting that *frac_retire* is essentially uncorrelated with our set of individual-level determinants.

5.4.1. Controls for correlated and exogenous effects

A key question is whether the error term in column (1) is correlated with the peer effects variable due to unobserved member characteristics or employer shocks. Since *frac_retire* varies at the employer-month level, to help rule out correlated and exogenous effects, column (1) in-

cludes three control variables that also vary at the employer-month level.

First, to control for time-series variation in the quality of the employee's workplace (for example, whether the new boss is overbearing), we include turnover of non-retirement eligible employees within the same employer and month. Second, we control for the retirement behavior of PERS members who work for other employers located in the same county. We conjecture that these individuals might retire together because of common economic factors in their county, or because they are responding to common information in local media outlets. Third, under the assumption that the former employees of employer j are a good control group for the current employees of employer j , we control for the fraction of former employees that retire in month t . The fact that the estimated coefficient on *frac_retire* is positive and statistically significant with these controls in the regression increases our confidence that we are identifying a peer effect.

5.4.2. Instrumental variables

To provide further evidence that we are identifying a peer effect, in the remaining columns of Table 6, we switch our estimation from OLS to instrumental variables. Our goal is to isolate variation in the fraction of coworker retirements that is being driven by exogenous variation in coworker's retirement incentives—rather than variation due to selection, firm-specific shocks, or other unobserved commonality in individual characteristics—and ask whether this variation helps to predict the retirement of member i in month t . In each column between (2) and (5), we estimate a different instrumental variables regression using a different instrument. Each instrument is calculated using all retirement-eligible members who work at employer j in month t , excluding member i . In column (6), we use estimate a single regression using all four instruments. As discussed below, the estimated peer effect is positive in each of the five new specifications, but it is statistically significant in only three of them.

The first instrument is the average retirement incentive due to stale pricing; the second instrument is the average retirement incentive due to the change in actuarial equivalency factors in July 2003; and the third instrument is the average retirement incentive due to actuarial equivalency factors being updated annually, in the member's birth month. The larger each of these instruments, the stronger the short-term retirement incentives faced by an individual's retirement-eligible coworkers. When member i is eligible for the DC or DCDB retirement benefit calculations, the first and second instruments will be positively correlated with member i 's own retirement incentives, which we control for directly in the regression. In contrast, the third instrument

captures variation in coworker retirement incentives driven by the distribution of coworker birth months over the calendar year, which should be uncorrelated with member i 's own retirement incentives. In other words, whereas the first and second instruments correspond to situations in which coworker retirements are informative about general retirement incentives, the third instrument does not.

When we use coworkers' average retirement incentives due to stale returns in January and February as our instrument in column (2), the estimated peer-effect coefficient increases to 45.40, and remains statistically significant at the 1-percent level despite an 8-fold increase in its standard error. In contrast, when we use coworkers' average retirement incentives due to changes in the actuarial equivalency factors in 2003 as our instrument in column (3), the estimated coefficient falls to 16.02 and loses statistical significance (with a p -value of 0.234). One possible explanation for the different results in columns (2) and (3) is a difference in saliency. Whereas PERS repeatedly told members about changes to the actuarial equivalency factors in July 2003, allowing members to determine their own retirement incentives, PERS did not tell members about the retirement incentive due to the use of stale returns, forcing coworkers to learn about this incentive from coworkers.

When we use coworker retirement incentives based on the number of months to their birthday as our instrument in column (4), the estimated peer-effect coefficient is large and negative although the standard error is even larger. This suggests that peer effects only arise from the diffusion of information about retirement incentives that generalize to other coworkers. In column (5), we use the fraction of member i 's coworkers that have a birthday in month t as our instrument to explain variation in the fraction of member i 's coworkers who retire in month t . Our original thinking was that members who retire in their birth month will be less sensitive to retirement incentives, and that this instrument will allow us to measure peer effects driven by non-financial retirements. However, we found in Table 5 that members retiring in their birth month are at least as sensitive to expected retirement benefits, local labor market conditions, and (some of) their own short-run retirement incentives as other members. Therefore, the estimated coefficient of 22.85 in column (5) may also reflect the diffusion of information driven by time-series variation in the fraction of recently-informed coworkers.

In column (6), when we use all four instruments at the same time, the estimated peer-effect coefficient is 17.58, and it is statistically significant at the 5-percent level (p -value of

0.044). According to this estimate, a one standard deviation increase in *frac_retire* increases the probability that member *i* retires in month *t* by 0.69 percent, which is slightly less than half of the unconditional probability of 1.46 percent. The overall evidence in Table 6 suggests both that we are identifying true peer effects, and that these peer effects reflect the diffusion of information about retirement incentives that generalize to coworkers.

5.4.3. Do peer effects reflect shared retirement incentives?

To test more directly whether peer effects reflect the diffusion of information about retirement incentives, we adopt the identification strategy of Bertrand, Luttmer, and Mullainathan (2000).¹⁷ Because PERS retirement benefits are calculated using three different benefit formulas (DC, DCDB, and DB), different coworkers can face different retirement incentives within the same month. For example, while members facing the DC benefit formula can time their retirement to exploit stale returns, members facing the DB benefit formula cannot. We use this fact to test whether individuals are *disproportionately* more likely to respond to their own retirement incentive when more of their coworkers face the same incentive.

In Table 7, we replace the fraction of member *i*'s coworkers retiring in month *t* with variables that measure the quantity and expected behavior of coworkers facing the same retirement benefit calculation as member *i*. First, for each member, we calculate the fraction of her retirement-eligible coworkers who would retire under her same retirement benefit calculation (i.e., DC, DCDB, or DB) that she does in month *t*. The larger this fraction, the larger the number of peers with whom member *i* can discuss her own retirement incentives. Second, within the full sample of employers, we calculate the fraction of retirement-eligible coworkers facing each retirement benefit that retire in month *t*. This variable measures the average strength of the retirement incentives that members expecting DC, DCDB, and DB benefits face in month *t*. For example, by controlling for the fraction of members who retire under the DC benefit in February 1995, we capture the average retirement incentive due to the use of stale returns in the calculation of members' account balances. Finally, we interact the fraction of coworkers facing the same retirement benefit calculation as member *i* in month *t* with the average fraction of retirement-eligible workers facing this retirement benefit calculation who retire in month *t*. This in-

¹⁷ Bertrand, Luttmer, and Mullainathan (2000) study the decision by individuals to participate in welfare programs. To test for peer effects, they interact the quantity of people who live in the same area and speak the same language as employee *i* with the average welfare participation rate for people who speak that language in the full cross section. They find that the interaction term is positive and statistically significant.

teraction term is our new variable of interest.

In the first column of Table 7, we report coefficients for the linear probability model:

$$\Pr(\text{retire}_{ijkt}) = \left(\text{frac}_{\text{same}_{-ijkt}} \times \overline{\text{retire}_{kt}} \right) \alpha + (\text{frac}_{\text{same}_{-ijkt}}) \gamma + X_{ijkt} \beta + \eta_{kt} + \delta_{jt} + \varepsilon_{ijkt}$$

where $\text{frac}_{\text{same}_{-ijkt}}$ is the fraction of employee i 's retirement-eligible coworkers at employer j facing retirement benefit k in month t , $\overline{\text{retire}_{kt}}$ is the fraction of retirement-eligible employees facing retirement benefit k that retire in month t (measured across all employers), and X_{ijkt} contains many of the control variables from Table 6, including all of employee i 's individual retirement incentives. Including a separate fixed effect for each retirement benefit calculation-month combination (η_{kt}) allows us to control for the average effect of benefit-specific retirement incentives on retirements in month t (and causes $\overline{\text{retire}_{kt}}$ to drop from the regression). Because we are focused on the interaction term, we are also able to include a separate fixed effect for each employer-month combination (δ_{jt}). These fixed effects allow us to control for the average characteristics and incentives faced by employee i 's coworkers in month t and to control for employer-month specific shocks—something that we could not control for in Table 6. (Of course, we can no longer include the fraction of employee i 's coworkers that retire in month t , or any other variable that varies solely at the employer-month level.) Standard errors are clustered on employer.

If employees are disproportionately more likely to respond to their retirement incentives when more coworkers face the similar incentives, α will be positive. Indeed, the coefficient on the interaction term in column (1) is positive and statistically significant (p-value of 0.000). It is also economically significant. Following Bertrand, Luttmer, and Mullainathan (2000), we estimate that PERS employees are 90.9% more likely to respond to their aggregate retirement incentives than they would be in the absence of any peer effects.¹⁸ The fact that employees are significantly more likely to respond to incentives when more of their coworkers face them strongly suggests that peer effects reflect the diffusion of information about retirement incentives.

The test in column (1) assumes that employee i 's peer group is best defined by her benefit

¹⁸ We calculate the multiplier in column (1) of Table 7 as $\sum_{k=\{DB, DCDB, DC\}} \text{frac}_k \left(1 / (1 - \alpha(\text{frac}_{\text{same}_k}) - 1) \right)$ which depends on the average value of $\text{frac}_{\text{same}}$ for each of the three retirement benefit calculations ($\text{frac}_{\text{same}_{DB}}$, $\text{frac}_{\text{same}_{DCDB}}$, and $\text{frac}_{\text{same}_{DC}}$), and the fraction of retirees whose retirement benefits are determined by DB, DCDB, and DC (frac_{DB} , frac_{DCDB} , and frac_{DC}). When we interact retirement benefit calculation type with job type, for example, the number of categories doubles from $\{DB, DCDB, DC\}$ to $\{DB, DCDB, DC\} \times \{\text{PF}, \text{not PF}\}$.

formula or, alternatively, that any employee is equally likely to talk about retirement with any other employee. In the remaining columns of Table 7, we include interaction terms based on alternative definitions of employee i 's coworkers.¹⁹ In column (2), we include the fraction of coworkers who are the same gender as employee i in employer j in month t , the average fraction of coworkers who are the same gender as employee i that retire (from any employer) in month t , and the interaction between these variables. The coefficient estimate on the interaction term is negative and statistically indistinguishable from zero, while the coefficient estimate on the original retirement benefit calculation interaction term is almost identical to the one in column (1). In column (3), we include an interaction term based on the fraction of coworkers who are the same gender and face the same retirement benefit calculation as employee i . Relative to column (1), this specification allows for the possibility that employees are more likely to discuss retirement incentives with coworkers of their own gender. While the estimated coefficient on this interaction term is positive, it is statistically indistinguishable from zero. Moreover, it has little effect on the size or significance of the original interaction term.

Police and fire officers have their own, more generous versions of the DCDB and DB benefit calculations, and they may be more likely to interact with other police and fire officers than with general service employees. (Some employers employ both general service employees and police or fire officers.) Therefore, in the remaining columns of Table 7, we distinguish police and fire officers from other employees. In column (4), we include the fraction of coworkers who have the same PERS classification (i.e., general service or not) as employee i in employer j in month t , the average fraction of coworkers who have the same PERS classification as employee i that retire (from any employer) in month t , and the interaction between these variables. The coefficient estimate on the new interaction term is negative, but statistically indistinguishable from zero, while the coefficient on the retirement benefit calculation interaction term remains both economically and statistically significant.

In column (5), we include an interaction term based on the fraction of coworkers who

¹⁹ When testing for peer effects in the decision by university employees to participate in a supplemental tax-deferred retirement savings account, Duflo and Saez (2002) argue that a priori restrictions on which coworkers are peers can be used to help identify peer effects. They construct subgroups based on gender, years of service, age, faculty versus staff, and academic department. Because we are focused on the retirement-timing decision, we do not attempt to construct subgroups based on years of service or age, but we do construct subgroups based on gender, police and fire versus normal, and employer (which is our analog to department).

have the same job type and face the same retirement benefit calculation as employee i . This specification allows for the possibility that police and fire officers are more likely to respond to their own retirement incentives when more of their police and fire officer coworkers face the same retirement incentives, and that the same is true for general service employees. Indeed, the coefficient on this final interaction term is positive and statistically significant, while the coefficient on the original interaction term declines. In other words, whereas the findings in the earlier columns suggest that peers groups can be defined as those coworkers facing the same retirement benefit calculation in month t , the findings in column (5) suggest that employees are slightly more likely to discuss retirement incentives with coworkers with the same job type.

6. Conclusion

Oregon's Public Employees Retirement System (PERS) offers employees a pension plan that is both generous and complex. We find strong evidence that plan members see through the plan's complexity and time their retirements to maximize their monthly benefits. This behavior imposes direct costs on employers through higher benefits as well as potentially large administrative costs arising from shortened careers and lumpy retirements.. Exploiting exogenous variation in the level of benefits, we also find strong evidence that employees learn about their retirement incentives from their coworkers. At the same time, we find that a minority of members make mistakes. For example, we identify retirees who could have retired one month earlier and received benefits that were as much as 20% larger.


One message from our paper is that there is a distinction between the expected level of retirement benefits and the form in which those benefits are delivered to employees. For example, if PERS had intended merely to increase average pension generosity, it could have done so by offering instead a one-formula DB pension plan with a larger payout factor. Simplifying the form of the benefits from multiple complex formulas to a single simple formula would result in two broad changes. First, it would make the plan more predictable: the underlying PERS investment portfolio would be easier to manage; the employers' annual pension costs would be less volatile; employees would have more certainty about the level of future retirement benefits; and both employees and employers would have more certainty about future retirement dates. Second, simplified plan benefits would increase transparency: the true value of the pension benefit would be better understood by employees, employers, and tax payers. Future research could investigate why employers build complexity into their pension plans.

References

- Bertrand, Marianne, Luttmer, Erzo F. P., and Mullainathan, Sendhil, 2000. Network effects and welfare cultures. *Quarterly Journal of Economics* 115 (3), 1019–1055.
- Brown, Jeffrey, and Scott Weisbenner, 2012. Why do individuals choose defined contribution plans? Evidence from participants in a large public plan. Unpublished working paper.
- Brown, Kristine, and Ron Laschever, 2012. When they're sixty-four: Peer effects and the timing of retirement. *American Economic Journal: Applied Economics* 4 (3): 90–115.
- Campbell, John, 2006. Household finance. *Journal of Finance* 61, 1553–1604.
- Chan, Sewin, and Ann Huff Stevens, 2004. Do changes in pension incentives affect retirement? A longitudinal study of subjective retirement expectations. *Journal of Public Economics* 88, 1307–1333.
- Chan, Sewin and Ann Huff Stevens, 2008. What you don't know can't help you: Pension knowledge and retirement decision-making. *Review of Economics and Statistics* 90, 253–266.
- Chalmers, John, and Jonathan Reuter, 2012. How do retirees value life annuities? Evidence from public employees. *Review of Financial Studies* 25, 2601–2634.
- Choi, James, David Laibson, and Brigitte Madrian, 2011. \$100 bills on the sidewalk: Suboptimal investment in 401(k) plans. *Review of Economics and Statistics* 93, 748–763.
- Choi, James, David Laibson, and Brigitte Madrian, 2005b. Are empowerment and education enough? Under-diversification in 401(k) plans. *Brookings Papers on Economic Activity* 2005(2), 151–198.
- Coile, Courtney and Jonathan Gruber, 2007. Future Social Security entitlements and the retirement decision. *Review of Economics and Statistics* 89, 234–246.
- Duflo, Esther, and Emmanuel Saez, 2002. Participation and investment decisions in a retirement plan: The influence of colleagues' choices. *Journal of Public Economics* 85, 121–148.
- Duflo, Esther, and Emmanuel Saez, 2003. The role of information and social interactions in retirement plan decisions: Evidence from a randomized experiment. *Quarterly Journal of Economics* 118 (3), 815–842.
- Goda, Gopi Shah, Shoven, John B. and Slavov, Sita N., 2009. Social Security policy in a changing environment: Removing the disincentives in Social Security for long careers.
- Goda, Gopi Shah, Shoven, John B. and Slavov, Sita N., 2010. Does stock market performance influence retirement intentions? NBER Working Paper #16211.


- Gustman, Alan and Thomas Steinmeier, 2001. Imperfect knowledge, retirement and saving. Dartmouth College Working Paper 01-11.
- Levy, Helen and Kristin Seefeldt, 2008. How do lower-income families think about retirement? University of Michigan working paper.
- Lusardi, Annamaria and Peter Tufano, 2009. Debt literacy, financial experiences, and overindebtedness. Harvard Business School working paper.
- Madrian, Brigitte and Dennis Shea, 2001. The power of suggestion: Inertia in 401(k) participation and savings behavior. *Quarterly Journal of Economics* 116 (4), 1149–1187.
- Manski, Charles, 1993. Identification of exogenous social effects: The reflection problem. *Review of Economic Studies* 60, 531–542.
- Snell, Ronald, 2012. State cash balance, defined contribution and hybrid retirement plans,” National Conference of State Legislatures,
<http://www.ncsl.org/Portals/1/Documents/employ/State-DC-and-Hybrid-Plans-July2012.pdf>.
- Novy-Marx, Robert, and Joshua Rauh, 2011. Public pension promises: How big are they and what are they worth? *Journal of Finance* 66, 1207–1245.
- Novy-Marx, Robert, and Joshua Rauh, 2013. Linking benefits to investment performance in US public pension systems. NBER Working Paper #18491.
- Sacerdote, Bruce, 2001. Peer effects with random assignment: Results for Dartmouth roommates. *Quarterly Journal of Economics* 116 (2), 681–704.
- Samwick, Andrew, 1998. New evidence on pensions, social security, and the timing of retirement. *Journal of Public Economics* 70, 207–236.
- Stanton, Richard, 2000. From cradle to grave: How to loot a 401(k) plan. *Journal of Financial Economics* 56, 485–516.
- Stock, James H and David A. Wise, 1990. Pensions, the option value of work, and retirement. *Econometrica* 58, 1151–1180.
- Sundaresan, M. Suresh, and Fernando Zapatero, 1997. Valuation, optimal asset allocation, and retirement incentives of pension plans,” *Review of Financial Studies* 10, 631–660.

Figure 1. Fluctuations in DC Benefits
 Due to use of stale returns in 1990-1999 and updated AEFs in 2003



Note: For 1990-1999, we plot the median value of *DC_delta* (blue), as well as the range between the minimum and maximum values (red). For 2003, we plot median value of *AEF_delta* (blue), as well as the range between the minimum and maximum values (red). The sample is limited to members who would receive DC benefits.

Figure 2. Fraction of Retirement-Eligible Members Who Retire each Calendar Year



Note: We graph two retirement ratios. The first is based on the number of DB, DC, and DCDB retirements within the set of member eligible to receive full DB benefits. The second is based on the number of DB, DC, and DCDB retirements within the set of members eligible to receive reduced DB retirement benefits (i.e., within the set of members whose DB and DCDB benefits would be reduced by the early retirement penalty).

Figure 3a. Retirement-Timing Decisions of DC Retirees within the Calendar Year






Figure 3b. Retirement-Timing Decisions of DB Retirees within the Calendar Year



Note: We graph the fraction of retirements that occur each month in three periods: 1990-1999, 2000-2002, and 2003. Figure 3a is restricted to retirees receiving DC benefits. Figure 3b is restricted to retirees receiving DB benefits.

Figure 4. *Ex Ante* Estimate of DC/DB Ratio at Retirement



Note: Using PERS characteristics as of 1990, we plot percentiles from the *ex ante* distribution of the ratio of DC benefits to DB benefits at retirement for members who join PERS at each age between 20 and 57. We also plot the fraction of the 50,000 simulated pension paths that result in DC benefits. See Section 4.2. for a discussion of the assumptions underlying these Monte Carlo simulations.

Table 1. Individual-level summary statistics

	EMPLOYED & ELIGIBLE TO RETIRE (#)	REPLACE- MENT RATE (%)	YEARS OF SERVICE (# years)	AGE (# years)	RETIRE WITH DCDB BENEFIT? (%)	RETIRE WITH DC BENEFIT? (%)
<i>Panel A. Eligible to Retiree</i>						
1990	12,872	27.2%	15.1	59.1	49.6%	21.8%
1991	13,540	32.7%	15.3	59.1	24.7%	57.7%
1992	13,869	28.6%	15.4	59.0	41.8%	27.6%
1993	14,641	30.3%	15.5	59.0	31.6%	39.0%
1994	14,807	29.7%	15.6	58.9	32.1%	34.0%
1995	14,563	36.1%	15.2	58.7	8.9%	73.0%
1996	15,968	33.2%	15.6	58.7	14.0%	58.8%
1997	17,573	37.9%	15.7	58.5	7.3%	73.5%
1998	18,734	39.4%	15.6	58.4	4.8%	79.5%
1999	18,757	38.7%	14.9	58.3	3.4%	83.0%
2000	18,930	37.3%	14.5	58.3	2.7%	78.3%
2001	21,809	36.9%	14.8	58.3	3.0%	70.4%
2002	24,344	35.9%	15.0	58.3	3.0%	63.9%
2003	25,001	33.3%	14.8	58.3	3.0%	58.9%
1990-2003	245,408	34.5%	15.2	58.6	13.8%	60.7%
<i>Panel B. Choose to Retire</i>						
1990	1,506	35.4%	18.9	60.6	59.7%	19.7%
1991	1,707	40.8%	19.3	60.5	41.1%	45.7%
1992	1,678	38.1%	19.5	60.6	52.6%	29.0%
1993	1,955	40.6%	19.7	60.5	42.1%	41.6%
1994	2,775	41.0%	20.6	60.0	46.0%	34.6%
1995	1,454	44.6%	18.8	60.4	17.7%	67.2%
1996	1,785	44.0%	19.9	60.0	25.4%	56.6%
1997	2,238	50.4%	20.4	59.7	15.6%	72.7%
1998	3,607	57.6%	21.3	59.2	7.6%	86.1%
1999	3,462	58.8%	20.9	58.6	6.1%	88.2%
2000	1,669	55.5%	19.2	58.7	4.3%	88.0%
2001	2,348	57.7%	20.6	58.8	5.2%	84.7%
2002	3,628	60.4%	21.7	58.7	4.9%	84.3%
2003	5,317	56.0%	21.2	58.5	4.8%	82.2%
1990-2003	35,129	51.0%	20.5	59.4	19.2%	68.3%

Note: The unit of observation is retirement-eligible employee i in year t . For employees who do not retire in year t , variables are measured in December. For employees who do retire in year t , variables are measured in the last month of employment.

Table 2. Comparing actual and counterfactual retirement benefits owed to new retirees

	Number of Retirements	Replacement Rate		Replacement Rate as % of DB only		% Retirees Receiving Benefits Above Those Based on DB		
		Counterfactual		Counterfactual		Counterfactual		
		DB only	Actual	AEF	Actual	AEF	Actual	AEF
1990	1,482	31.1%	35.4%	34.6%	113.9%	111.2%	79.4%	74.6%
1991	1,669	32.1%	40.9%	39.5%	127.8%	123.3%	86.8%	83.5%
1992	1,634	32.4%	38.3%	36.7%	118.3%	113.3%	81.5%	75.6%
1993	1,917	32.5%	40.7%	36.6%	125.4%	112.6%	83.7%	69.7%
1994	2,734	34.0%	41.0%	36.2%	120.7%	106.3%	80.5%	51.6%
1995	1,425	30.8%	44.8%	38.9%	145.7%	126.4%	84.9%	74.7%
1996	1,758	32.7%	44.0%	37.0%	134.9%	113.1%	82.1%	64.6%
1997	2,190	33.9%	50.5%	42.0%	149.5%	124.1%	88.3%	72.9%
1998	3,539	34.6%	57.7%	45.2%	167.3%	130.6%	93.7%	80.3%
1999	3,408	33.8%	59.1%	43.9%	175.0%	129.8%	94.3%	75.4%
2000	1,641	30.8%	55.6%	45.7%	180.9%	148.4%	92.3%	77.5%
2001	2,306	33.5%	57.9%	45.8%	172.9%	136.5%	89.8%	73.3%
2002	3,563	35.2%	60.7%	46.0%	172.2%	130.5%	89.1%	72.8%
2003	5,231	33.9%	56.2%	40.4%	165.7%	119.0%	86.9%	59.1%
Total	34,497	33.4%	51.2%	41.3%	153.5%	123.6%	87.5%	70.6%

Note:

This table compares PERS retirement benefits calculated in three different ways. "DB only" is the counterfactual replacement rate that each retiree would have received from PERS under the DB benefit formula. "Actual" is the replacement rate that each retiree actually receives, determined by the maximum of the DB, DC, and (when eligible) DCDB benefit formulas. "Counterfactual AEF" is the counterfactual replacement rate that each retiree would have received if PERS converted each retiree's DC account balance into life annuity payments using the actuarial equivalency factors (AEFs) available in the private market. Because our data on private market AEFs are limited to ages 50 to 70, the sample of retirees is slightly lower than in Table 1 Panel B. All three benefit calculations hold the members' retirement date choices constant, and are based on the assumption that members choose to receive all of their retirement benefits in the form of life annuity payments.

Table 3. Probability of early retirement and distribution of retirement ages by benefit formula

Panel A. Sample of retirees for whom early retirement age is 55 and normal retirement age is 58

Age at Retirement	Retires with DB Benefits		Retires with DCDB Benefits		Retires with DC Benefits	
	#	%	#	%	#	%
55	105	3.1%	141	2.9%	3,660	17.2%
56	135	4.0%	191	3.9%	2,039	9.6%
57	206	6.1%	383	7.8%	2,112	9.9%
58	542	15.9%	1,101	22.4%	2,941	13.8%
59	313	9.2%	560	11.4%	1,630	7.7%
60	289	8.5%	480	9.8%	1,437	6.8%
61	316	9.3%	506	10.3%	1,587	7.5%
62	566	16.7%	654	13.3%	1,841	8.7%
63	230	6.8%	244	5.0%	807	3.8%
64	203	6.0%	201	4.1%	768	3.6%
65	238	7.0%	256	5.2%	1,019	4.8%
66	92	2.7%	75	1.5%	337	1.6%
67	58	1.7%	41	0.8%	227	1.1%
After 67	106	3.0%	74	1.4%	833	3.6%
All	3,399		4,907		21,238	

Panel B. Full sample of retirees

	Retires with DB Benefits		Retires with DCDB Benefits		Retires with DC Benefits	
	#	% Retiring Early	#	% Retiring Early	#	% Retiring Early
1990-2003	4,388	14.0%	6,750	12.2%	23,991	33.8%
1990-1996	2,245	8.8%	5,292	12.0%	5,323	26.9%
1997-2003	2,143	19.6%	1,458	12.8%	18,668	35.7%
Change		10.8% ***		0.8%		8.8% ***
1996 only	320	8.1%	454	9.5%	1,011	23.8%
1997 only	262	13.7%	350	11.7%	1,626	24.2%
Change		5.6% *		2.2%		0.3%

Note:

In Panel A, we report the distribution of retirement ages for employees retiring under each of the three benefit formulas. We begin with the full sample of retirees described in Table 1 Panel B. However, to facilitate comparisons across benefit formulas, we exclude 2,385 retirees who are police and fire officers, 632 retirees whose normal retirement age is 60, and 2,568 retirees whose normal retirement age is 58 but who qualify for normal retirement benefits based on years of service rather than age. This leaves us with a sample of 29,544 retirees for whom the early retirement age is 55 and the normal retirement age is 58. In Panel B, we focus on the full sample of retirees described in Table 1 Panel B. For employees retiring under each of the three benefit formulas, we report the total number of retirements and the fraction of retirements by employees who face (or would face) an early retirement penalty under the DB and DCDB benefit formulas. Because the frequency with which the early retirement penalty is updated changes from annual to monthly on January 1, 1997, we also report separate statistics for these two regimes. To test whether the likelihood of an early retirement changes across these regimes, we regress a dummy variable indicating whether employee i faces (or would face) an early retirement penalty under the DB and DCDB formulas on a dummy variable indicating whether the early retirement penalty is updated monthly. To allow for correlated behavior within employers, we cluster the standard errors on employee i 's employer identification number.

Table 4. Retirement Timing Decisions and Stale Returns

Panel A. Impact of Stale Returns on DC Benefits in February

Benefit from Stale Returns?	February Retirees				March Retirees					
	#	Mean	Min.	Med.	Max.	#	Mean	Min.	Med.	Max.
Yes	3,552	3.0%	0.0%	1.9%	28.6%	142	-5.7%	-22.5%	-4.1%	-0.2%
No	548	-2.6%	-20.5%	-2.3%	0.0%	90	6.4%	0.4%	6.0%	19.9%
All Retirees	4,100	2.3%	-20.5%	1.9%	28.6%	232	-1.0%	-22.5%	-3.1%	19.9%

Panel B. Retiree Characteristics

Benefit from Stale Returns?	#	Monthly Salary	Monthly Replace. Rate	Years of Service	Age	Female?	Police or Fire?	School?	Tier 2?	Die in months 1-48?	Invest in Variable Account?
No	638	3,872	65.7%	21.7	59.7	48.0%	2.5%	57.8%	1.1%	2.8%	42.9%
All Retirees	4,332	3,857	62.2%	21.9	59.4	55.7%	2.1%	67.1%	0.3%	1.8%	19.9%

Note:

Through December 1999, PERS calculates member's PERS account balances using stale returns. For those retiring in February of year t , PERS uses the annual returns for year $t-2$ to grow the member's account balance between January of year $t-1$ and February of year t . For those retiring in March of year t , PERS instead uses the newly-determined annual returns for year $t-1$. When the stale returns for year $t-2$ are higher than the returns that PERS is expected to set for year $t-1$, members receive higher DC benefits from retiring in February. When the stale return for year $t-2$ is lower than this expected return, members receive higher DC benefits from waiting to retire in March. In Panel A, we focus on the subset of members who retire in February or March under the DC benefit formula. We find that 3,694 members receive higher DC benefits through their choice of retirement month, while 638 members received lower DC benefits. We also summarize the impact that stale returns have on the member's DC benefit. For example, the 140 retirees who benefited from stale returns by retiring in March would have received DC benefits that were, on average, 8.0% lower had they instead retired in February. In Panel B, we report the average values of the final average monthly salary, replacement rate, years of service, and age at retirement. We also report the fraction of retirees who are female, police or fire officers, employed by a school, eligible for Tier 2 benefits, die within 48 months of retirement, and allocate a positive fraction of their PERS contributions to the variable account

Table 5. Linear Probability Model Predicting Choice of Retirement Date

Dependent Variable: Sample: Estimation:	<i>1 if employee i retires on date t, 0 otherwise</i>			
	All members	Female = 1	PF = 1	Birth Month = 1
	OLS	OLS	OLS	OLS
	(1)	(2)	(3)	(4)
Replacement rate	3.456 *** [0.187]	3.099 *** [0.218]	2.806 *** [0.718]	6.767 *** [0.524]
EPV benefit of waiting until t*	0.003 [0.002]	0.001 [0.002]	-0.009 [0.006]	-0.001 [0.007]
Incentive in Jan. and Feb. due to stale returns	7.333 *** [0.703]	9.381 *** [1.005]	-2.546 [3.456]	6.616 [4.055]
Incentive due to use of annual AEFs	9.027 *** [3.276]	5.604 * [3.198]	38.735 ** [16.803]	
Incentive due to change in AEFs in July 2003	87.877 *** [5.839]	101.797 *** [6.899]	85.205 *** [25.475]	104.965 *** [15.825]
Eligible for early retirement -- month 1	0.006 [0.224]	0.194 [0.264]	-1.171 ** [0.502]	-2.004 *** [0.711]
Eligible for early retirement -- months 2-12	-1.181 *** [0.131]	-0.954 *** [0.124]	-1.977 *** [0.355]	0.280 [0.275]
Eligible for early retirement -- month 13+	-0.496 *** [0.081]	-0.398 *** [0.079]	-1.309 *** [0.212]	1.273 *** [0.220]
Eligible for normal retirement -- month 1	3.030 *** [0.330]	2.309 *** [0.311]	10.123 *** [1.895]	3.340 *** [0.341]
Eligible for normal retirement -- months 2-12	0.328 *** [0.078]	0.093 [0.080]	1.100 *** [0.361]	0.934 *** [0.238]
Birth month?	0.876 *** [0.092]	0.699 *** [0.105]	0.566 *** [0.182]	
Member dies in months 1-12?	0.815 *** [0.191]	0.769 *** [0.258]	0.477 [0.564]	1.207 [0.805]
Member dies in months 13-48?	0.208 *** [0.055]	0.123 [0.086]	0.342 [0.313]	0.743 ** [0.299]
Years of Service	0.036 *** [0.005]	0.043 *** [0.006]	0.016 [0.021]	0.055 *** [0.015]
Unemployment rate within county	-2.653 *** [0.941]	-2.525 *** [0.912]	-0.368 [2.832]	-5.410 ** [2.486]
Female?	-0.049 * [0.025]		0.039 [0.063]	-0.323 *** [0.092]
Police or Fire Fighter?	-0.032 [0.092]	-0.088 [0.114]		-0.332 [0.268]
PERS Tier Two?	-0.014 [0.039]	0.025 [0.049]	0.194 [0.143]	-0.479 *** [0.135]
Full life annuity calculated under DB	-0.089 ** [0.045]	-0.032 [0.035]	-0.239 [0.174]	-0.512 *** [0.154]
Full life annuity calculated under DC	-0.259 *** [0.037]	-0.150 *** [0.037]	-0.468 *** [0.142]	-0.742 *** [0.153]
Separate fixed effect for each age (in years)?	Yes	Yes	Yes	Yes
Date-by-employer type FE?	Yes	Yes	Yes	Yes
N	2,407,980	1,386,758	174,114	221,770
R-Squared	0.0629	0.0659	0.0487	0.0692

Note: Dependent variable equals 1 if employee i retires in month t and 0 otherwise. Employees who retire in month t are dropped from the sample in month t+1. For consistency with tests for peer effects in later tables, the sample is restricted to employers with two or more employees eligible to retire in month t. Independent variables are described in Section 5.3. Standard errors are clustered on employer. Statistical significance at the 10-percent, 5-percent, and 1-percent levels is denoted by *, **, and ***. Coefficients are multiplied by 100, so that 1.000 represents 1 percentage point.

Table 6. Testing for Peer Effects in the Choice of Retirement Date

	<i>I if employee i retires on date t, 0 otherwise</i>					
	All members OLS	All members IV	All members IV	All members IV	All members IV	All members IV
	(1)	(2)	(3)	(4)	(5)	(6)
Fraction of current employees retiring	27.042 *** [2.110]	45.400 *** [16.023]	16.017 [13.452]	-122.505 [339.821]	22.853 *** [8.545]	17.576 ** [8.720]
Fraction of current employees retiring from other employers in same county in month t	0.088 [2.984]	0.040 [2.210]	0.117 [3.446]	0.479 [9.203]	0.099 [3.158]	0.119 [3.376]
Fraction of non-retirement eligible employees leaving employer j in month t	2.793 *** [0.780]	2.234 *** [0.765]	3.129 *** [0.959]	7.351 [10.798]	2.921 *** [0.850]	3.081 *** [0.898]
Fraction of former employees retiring in month t	0.083 [0.272]	0.085 [0.207]	0.081 [0.312]	0.062 [0.825]	0.082 [0.287]	0.080 [0.306]
Incentive in Jan. and Feb. due to stale returns	7.109 *** [0.673]	6.956 *** [0.693]	7.200 *** [0.684]	8.349 *** [3.021]	7.143 *** [0.679]	11.790 *** [0.881]
Incentive due to use of annual AEFs	8.770 *** [3.263]	8.583 *** [3.265]	8.883 *** [3.259]	10.303 * [5.769]	8.814 *** [3.259]	8.354 ** [3.265]
Incentive due to change in AEFs in July 2003	87.133 *** [5.736]	86.646 *** [5.701]	87.425 *** [5.741]	91.092 *** [11.613]	87.243 *** [5.765]	87.335 *** [5.754]
Birth month?	0.874 *** [0.092]	0.872 *** [0.092]	0.875 *** [0.092]	0.888 *** [0.094]	0.874 *** [0.092]	0.876 *** [0.092]
Other controls and FEs from Table 5 Column (1)?	Yes	Yes	Yes	Yes	Yes	Yes
Date-by-employer type FE?	Yes	Yes	Yes	Yes	Yes	Yes

Note: Estimation is via OLS in column (1) and 2SLS in the remaining columns. The dependent variable and sample restrictions are the same as in Table 5. The instrument in column (1) is the average value of DC_delta for employee i 's coworkers at employer j in month t . The instruments in columns (2) and (3) are the average values of AEF_delta for employee i 's coworkers at employer j in month t , in the periods 2003 and 1990-1996, respectively. The instrument in column (4) is the fraction of coworkers with a birthday in month t . We use all four instruments in column (5). Additional control variables include the fraction of current employees retiring from other employers in same county in month t , the fraction of non-retirement eligible employees leaving employer j in month t , and the fraction of employer j 's former employees retiring in month t . Standard errors are clustered on employer. Statistically significance at the 10-percent, 5-percent, and 1-percent levels is denoted by *, **, and ***. Coefficients are multiplied by 100, so that 1.000 represents 1 percentage point.

Table 7. Interaction based tests for peer effects

	<i>1 if employee i retires on date t, 0 otherwise</i>				
Dependent Variable:	(1)	(2)	(3)	(4)	(5)
Fraction of coworkers with same benefit type (t) *	86.501 ***	86.550 ***	82.047 ***	88.421 ***	57.419 ***
Fraction of all employees with same benefit type who retire (t)	[15.779]	[15.745]	[19.023]	[15.379]	[18.772]
Fraction of coworkers with same benefit type (t)	-0.818 ***	-0.817 ***	-0.624 ***	-0.843 ***	-0.566 ***
	[0.162]	[0.162]	[0.208]	[0.158]	[0.214]
Fraction of coworkers with same gender (t) *		-2.797			
Fraction of all employees with same gender who retire (t)		[5.829]			
Fraction of coworkers with same gender (t)		-0.179 **			
		[0.091]			
Fraction of coworkers with same benefit type and gender (t) *			7.540		
Fraction of all employees with same benefit type and gender who retire (t)			[9.621]		
Fraction of coworkers with same benefit type and gender (t)			-0.216 *		
			[0.116]		
Fraction of coworkers with same PF status (t) *				-10.739	
Fraction of all employees with same PF status who retire (t)				[9.962]	
Fraction of coworkers with same PF status (t)				0.122	
				[0.257]	
Fraction of coworkers with same benefit type and PF status (t) *					26.815 **
Fraction of all employees with same benefit type and PF status who retire (t)					[11.035]
Fraction of coworkers with same benefit type and PF status (t)					-0.281
					[0.179]
Controls and FE from Table 5 Column (1)?	Yes	Yes	Yes	Yes	Yes
Date-by-benefit calculation-by-gender FE?	Yes	Yes	Yes	---	---
Date-by-benefit calculation-by-police/fire FE?	---	---	---	Yes	Yes
Date-by-employer FE?	Yes	Yes	Yes	Yes	Yes
N	2,407,959	2,407,959	2,391,709	2,407,959	2,402,768
Adj. R-Squared (demeaned)	0.0224	0.0224	0.0224	0.0222	0.0225
Multiplier for interaction term introduced in column (1)	90.9%	91.0%	81.5%	95.3%	44.2%
Multiplier for additional interaction term	---	-1.5%	4.1%	-8.8%	16.6%

Note: We describe our empirical strategy, which is based on Bertrand, Luttmer, and Mullainathan (2000), and the calculation of the multiplier effect due to peer effects in Section 5.4.3. We include all of the independent variables and fixed effects from Table 6 that vary within employer-month, but do not report the estimated coefficients. Standard errors are clustered on employer. Statistically significance at the 10-percent, 5-percent, and 1-percent levels is denoted by *, **, and ***. Coefficients are multiplied by 100, so that 1.000 represents 1 percentage point.

Web Appendix for “The Effect of Pension Design on Employer Costs and Employee Retirement Choices: Evidence from Oregon”

1. Brief History of Changes to PERS

The history of the PERS plan reflects the countervailing tensions between providing competitive pension benefits and managing the costs of providing and managing those benefits.¹ PERS came into existence on July 1, 1946. At the time, it was argued that an orderly pension system would help Oregon state and local employers compete more effectively for employees. It initially resembled a DC plan, with employees contributing into an account that earned interest, but with benefits capped at \$125 per month. In 1953, PERS employees began contributing into Social Security. In 1967, PERS became a hybrid system, combining DC-style and DB-style benefits into a single benefit formula. Specifically, PERS began to calculate the retirement benefits of new and existing employees using the DCDB (“Formula Plus Annuity”) benefit formula. In the same year, PERS began investing up to 10% of its portfolio in equities.

As quoted in the introduction, Governor Tom McCall argued that benefits for Oregon state employees were hampering the State’s ability to attract and retain the talented employees that it needed to prosper. Between the late 1960s and the early 1980s, as a result of perceived labor market pressures and periods of high inflation, PERS increased employees’ expected retirement benefits. In 1969, employees were given the choice between two investment options. The “regular” option guaranteed a minimum return of 5.5% per year, whereas the “variable annuity” option invested more heavily in U.S. and international equity. PERS also added an annual cost of living adjustment (COLA) to retirement benefits calculated under the existing DCDB benefit formula. In 1973, PERS increased the maximum annual COLA from 1.5% to 2.0%. In 1979, in lieu of increasing nominal wages, employers began to “pick up” the 6% employee contribution on behalf of their employees. Because these changes increased the generosity of the PERS pension plan, between 1978 and 1981, PERS’ actuary issued warnings about a potential underfunding problem. In 1981, the DB benefit formula (“Full Formula”) was introduced, and the DCDB formula was discontinued for new employees.

In 1994, Ballot Measure 8 eliminated the employer’s pickup of the employee’s 6% contribution and the guaranteed minimum return of 8% per year offered in the regular investment account. However, in 1996, the Oregon Supreme Court overturned Ballot Measure 8 for violat-

¹ The sources for this discussion are web-based documents available from the PERS website. They include “The Oregon Public Employees Retirement System History: The first 60 Years,” published by PERS on July 6, 2010 and “Public Employee Retirement in Oregon: Where does the system stand and where could Oregon go from Here?,” prepared by ECONorthwest for The Chalkboard Project and The Oregon Business Council on August 31, 2007.

ing the contract clause of the U.S. Constitution. The basic argument was that PERS could not change the expected retirement benefits of existing employees. In response, “Tier 2” retirement benefits were established for employees hired after December 31, 1995. The normal retirement age was increased from 58 to 60, and the guaranteed return of 8% per year in the regular account was eliminated. In January 1997, PERS began updating early retirement penalties and actuarial equivalency factors monthly instead of once per year, in the employee’s birth month.² This removed the disincentive for employees expecting to retire with DC or DCDB benefits to retire immediately before their birth month.

PERS employees retired in waves during the late 1990s. One likely explanation was the political uncertainty created by ongoing court cases and proposed legislation. Another explanation, for which we find strong empirical support, is that the use of stale returns to calculate member account balances generated lumpy retirements. There is anecdotal evidence that these lumpy retirements had real costs. Beginning in 1997, Oregon school districts began reporting teacher shortfalls because they were unable to replace all of the retiring teachers. The use of the “last known rate” to calculate retirement account balances was eliminated in January 2000.

In the spring of 2003, the difference between PERS’ assets and liabilities was \$17B. In 2003, PERS took several steps to close this funding gap. First, PERS established a third tier of pension benefits, called the Oregon Public Service Retirement Program (OPSRP), for employees hired after August 29, 2003. OPSRP lowered the payout factor used in the DB benefit formula from 1.67% to 1.50%, and increased the normal retirement age from 60 to 65. These changes reduced the generosity of the DB benefit formula. More significantly, PERS eliminated the DC benefit formula. Instead, employee retirement contributions were directed into individual retirement (“IAP”) accounts in which employees receive market returns and, therefore, bear market risk. In other words, for new employees, PERS became a system in which employer contributions fund a traditional DB retirement plan and employee contributions are invested just as in a traditional DC retirement plan. It is an open empirical question whether this reduction in expected retirement benefits, which reduced the expected pension costs of employers, lowered the average quality of new state and local employees.³

² The decision to use gender-neutral AEF grew out of a 1974 lawsuit filed by four female PERS members, which alleged that the use of gender-based mortality tables violated Title VII of the Civil Rights Act of 1964. The resulting consent decree also formed the basis for subsequent lawsuits that challenged the legality of updating the AEF to reflect the subsequent secular change in mortality risk.

³ We explored the idea of testing whether Oregon public schools were less able to attract high-quality teachers after 2003, but were unable to obtain any proxies for teacher quality before or after 2003. Even basic measures that might

Second, PERS began offering a full-lump-sum payout option at retirement (equal to twice the member's account balance) starting in January 2003. Because PERS life annuity payments were better than actuarially fair when compared to the existing partial lump sum payout option (Chalmers and Reuter (2012)), this option had the potential to reduce underfunding. We estimate that demand for the full lump sum reduced underfunding by approximately \$0.1B in 2003.

Third, the legislature made changes to PERS that reduced the expected retirement benefits of existing Tier 1 and Tier 2 employees. The retirement account underlying the DC benefit formula was closed to new contributions. While existing investments in the regular and variable options continued to earn returns, the absence of new contribution reduced the expected member's account balance at retirement. Instead, employee contributions were directed into the same IAP retirement account as employees hired into OPSRP, exposing the existing employees to market risk. New regulations also required PERS to update every two years the actuarial equivalency factors used to convert retirement account balances into initial monthly retirement benefits, with the first update occurring in July 2003. Because this change reduced life annuity payments under the DC and DCDB benefit formulas, it created strong incentives for certain employees to retire before July 2003. We find strong empirical support for the hypothesis that employees recognized this incentive. Finally, as a result of court decisions, PERS changed the DC benefit formula in July 2004 so that only the counterfactual returns that the employee would have earned by investing 100% in the regular account (and not the variable account) are doubled.⁴ As a result of these many changes, the likelihood that the DC benefit formula provides the maximum retirement benefit has declined substantially. While over 85% of retirees retired under DC in 2003, just over 50% did so in 2010.⁵ Furthermore, the 2003 changes significantly reduced the level of underfunding. According to a 2011 actuarial report by Mercer on PERS, the fraction of funded liabilities exceeded 86% on December 31, 2010.⁶

allow us to measure teacher shortages, such as teacher-to-student ratios, are unavailable from the Department of Education before 2004.

⁴ This change, which did not apply to retirements during our sample period, was the consequence of litigation between various public entities and the Public Employees Retirement Board. See section 1.1 of the settlement agreement: http://www.oregon.gov/pers/docs/board_information/board_meeting_2005/settlement_agreement.pdf.

⁵ Despite the widespread belief that the effect of the 2008 financial crisis on retirement assets forced employees to delay retirements, Goda, Shoven, and Slavov (2010) find only modest evidence that market returns delay retirement. The effect of market returns on the decision to retire is a potentially interesting research question to explore within PERS, as the effect of market returns on retirement account balances varies from Tier 1 to Tier 2 to OPSRP.

⁶ See the Mercer's report to PERS, "December 31, 2010 Actuarial Valuation Oregon Public Employees Retirement System," September 30, 2011.

2. Additional Summary Statistics

In Table A1, we report additional summary statistics for the PERS members in our sample. Panel A focuses on retirement-eligible members while Panel B focuses on retirees. We distinguish between members who are eligible to receive normal retirement benefits (based on age or years of service) and members who would receive lower DB or DCDB benefits due to an early retirement penalty. We also report average nominal salaries, the fraction of members who are female, the fraction of members who qualify for the more generous benefit formulas available to police and fire officers, the fraction of members who are classified as “Tier 2” (i.e., those hired after December 31, 1995), and the fraction of members who choose to allocate a portion of the employee retirement contribution to the variable account.⁷ Among retirement-eligible members, the average nominal monthly final average salary ranges from \$2,319 in 1990 to \$3,667 in 2003, an annualized increase of 3.59%. The fraction of female members increases from 54.7% in 1990 to 60.5% in 2003, while the fraction working as police and fire fighters ranges from 5.7% to 8.1%. Comparing Panel B with Panel A, we see that retirees are significantly more likely to be eligible for normal retirement benefits, and significantly less likely to be Tier 2 members (who were first hired by a PERS employer after December 31, 1995).

3. Additional Measures of Ex Post Pension Costs

In Table 2, we compare the actual benefits that PERS retirees received to two counterfactual benefits. In Table 2A, we calculate the present value costs of the actual and counterfactual retirement benefits from the employer’s perspective. “DB only” is the counterfactual benefit that each retiree would have received from PERS under the DB benefit formula. “Actual” is the benefit based on the retirement benefit that each retiree actually receives, which is the maximum of the DB benefit formula, the DC benefit formula, and (when eligible) the DCDB benefit formula. “Counterfactual AEFs” is the counterfactual benefit that each retiree would have received if PERS converted each retiree’s DC account balance into life annuity payments using the actuarial equivalency factors (AEFs) available in the private market. “PV Retirement Benefits Owed to New Retirees” sums the expected present value of the retirement benefits owed to the members retiring during each calendar year, under each of the three benefit calculation scenarios. We use retiree i’s gender and age at retirement and life tables from the Social Security Administration for 2004 to determine the probability that she receives each future monthly payment, we as-

⁷ Despite the doubling of the member’s account balance (see equation (2) in the main paper), all employer contributions go into the regular account. Members that allocate a portion of their contributions to the variable introduce funding risk for their employers. Namely, there is a funding shortfall whenever the returns in the variable account exceed the returns in the regular account.

sume an constant annual cost of living adjustment of 2.00%, and we use the yield on 10-year U.S. Treasury notes at the time of retirement as our discount rate. To calculate “Total” dollar costs, we use the yield of the 10-year U.S. Treasury notes at the time of retirement to convert the present value of the benefits owed to retiree i in year t into a future value measured on December 30, 2003; we then sum across retirees. To allow meaningful comparisons between dollar totals, we limit the sample to the 34,497 retirees between the ages of 50 and 70, for whom we possess data on private market AEFs from TIAA. We estimate PERS’ *ex post* pension liabilities for these 34,497 retirees to be \$18.3 billion. Consistent with Table 2, we find the *ex post* costs to be 153.2% higher than the \$11.9 billion estimated using the DB benefit formula and 124.0% higher than \$14.8 billion estimated using the counterfactual annuity factors.

There are three things to note about our dollar cost estimates. First, unlike the liability estimates in Novy-Marx and Rauh (2011), which are based on the benefits owed to all current employees and retirees, our liability estimates only apply to the subset of members who choose to retire during our sample period. Second, the relatively high pension costs in the last several years of the sample are driven by the combination of a risk-free rate near four percent and the cost-of-living adjustment of two percent. Novy-Marx and Rauh (2013) show that cost-of-living adjustments are an important determinant of pension liabilities. Third, we are measuring the net effect of the particular way that the PERS pension plan combines elements from DB and DC plans, including the net effect of the changes they made during our sample period. As a result, our specific costs estimates are unlikely to generalize to other pension plans.

Appendix Table 1. Additional Annual Individual-Level Summary Statistics

	EMPLOYED & ELIGIBLE TO RETIRE NORMAL (#)	EMPLOYED & ELIGIBLE TO RETIRE EARLY (#)	AVERAGE MONTHLY SALARY PAST 12 MONTHS (\$)	FEMALE (%)	POLICE OR FIRE (%)	PERS TIER 2 (%)	POSITIVE ALLOCATION TO VARIABLE ACCOUNT (%)
<i>Panel A. Eligible to Retiree</i>							
1990	7,724	5,148	\$2,319	54.7%	5.7%	0.0%	21.1%
1991	8,061	5,479	\$2,454	54.4%	6.0%	0.0%	22.0%
1992	8,266	5,603	\$2,587	54.9%	6.4%	0.0%	21.8%
1993	8,554	6,087	\$2,694	55.7%	6.4%	0.0%	24.7%
1994	8,516	6,291	\$2,809	56.1%	6.8%	0.0%	27.3%
1995	7,937	6,626	\$2,794	56.5%	6.9%	0.0%	29.1%
1996	8,839	7,129	\$2,928	56.6%	7.5%	1.1%	29.9%
1997	9,381	8,192	\$3,055	56.7%	8.0%	3.8%	33.1%
1998	9,495	9,239	\$3,177	57.1%	7.9%	6.9%	34.3%
1999	9,114	9,643	\$3,239	57.9%	8.1%	11.4%	34.5%
2000	9,123	9,807	\$3,299	58.8%	8.1%	15.5%	37.1%
2001	10,734	11,075	\$3,442	59.4%	7.7%	18.5%	38.3%
2002	11,634	12,710	\$3,593	60.0%	7.0%	21.0%	33.7%
2003	11,439	13,562	\$3,667	60.5%	6.7%	24.1%	24.5%
1990-2003	128,817	116,591	\$3,089	57.5%	7.2%	9.1%	30.1%
<i>Panel B. Choose to Retire</i>							
1990	1,266	240	\$2,801	52.4%	4.7%	0.0%	23.4%
1991	1,424	283	\$2,994	49.8%	5.7%	0.0%	25.4%
1992	1,402	276	\$3,183	49.2%	6.8%	0.0%	23.4%
1993	1,602	353	\$3,351	51.0%	6.1%	0.0%	25.4%
1994	2,282	493	\$3,506	53.9%	7.5%	0.0%	28.2%
1995	1,141	313	\$3,260	54.1%	4.3%	0.0%	29.6%
1996	1,475	310	\$3,581	53.0%	6.5%	0.2%	31.3%
1997	1,768	470	\$3,735	55.0%	8.6%	0.4%	35.1%
1998	2,469	1,138	\$3,866	55.0%	6.5%	0.8%	34.1%
1999	2,247	1,215	\$4,019	55.0%	7.5%	1.4%	34.9%
2000	1,055	614	\$4,093	54.5%	7.3%	4.5%	40.1%
2001	1,653	695	\$4,303	57.0%	7.9%	4.1%	38.7%
2002	2,503	1,125	\$4,534	57.7%	6.8%	4.2%	30.0%
2003	3,295	2,022	\$4,490	59.0%	6.6%	4.5%	21.6%
1990-2003	25,582	9,547	\$3,859	54.9%	6.8%	1.9%	30.0%

Note: The unit of observation is retirement-eligible employee i in year t . For employees who do not retire in year t , variables are measured in December. For employees who do retire in year t , variables are measured in the last month of employment.

Appendix Table 2. Present Value of Retirement Benefits Owed to New Retirees

	Number of Retirements	PV Retirement Benefits Owed to New Retirees (\$ millions)			PV Retirement Benefits Owed to New Retirees as % of DB only	
		DB only	Actual	Counterfactual AEF	Actual	Counterfactual AEF
1990	1,482	178.4	203.1	199.3	113.8%	111.7%
1991	1,669	228.7	294.7	285.8	128.9%	124.9%
1992	1,634	262.0	308.3	296.9	117.6%	113.3%
1993	1,917	365.7	460.1	414.8	125.8%	113.4%
1994	2,734	499.0	603.2	531.9	120.9%	106.6%
1995	1,425	249.6	371.4	322.7	148.8%	129.3%
1996	1,758	348.9	471.2	396.7	135.1%	113.7%
1997	2,190	483.5	726.2	606.5	150.2%	125.4%
1998	3,539	902.9	1,513.0	1,188.0	167.6%	131.6%
1999	3,408	887.0	1,559.4	1,160.0	175.8%	130.8%
2000	1,641	391.3	712.4	588.1	182.0%	150.3%
2001	2,306	676.8	1,169.4	929.2	172.8%	137.3%
2002	3,563	1,241.2	2,123.8	1,617.0	171.1%	130.3%
2003	5,231	1,926.2	3,184.8	2,291.1	165.3%	118.9%
Total	34,497	11,911.9	18,250.3	14,765.5	153.2%	124.0%

Note: This table extends the analysis of ex post pension costs in Table 2. "PV Retirement Benefits Owed to New Retirees" sums the expected present value of the retirement benefits owed to the members retiring during each calendar year, under each of the three benefit calculation scenarios. We use retiree i's gender and age at retirement and life tables from the Social Security Administration for 2004 to determine the probability that she receives each future monthly payment, we assume an constant annual cost of living adjustment of 2.00%, and we use the yield on 10-year U.S. Treasury notes at the time of retirement as our discount rate. To calculate "Total" dollar costs, we use the yield of the 10-year U.S. Treasury notes at the time of retirement to convert the present value of the benefits owed to retiree i in year t into a future value measure on December 30, 2003, and then sum across retirees.