Journal of Public Economics 98 (2013) 1–14

Contents lists available at SciVerse ScienceDirect

Journal of Public Economics

journal homepage: www.elsevier.com/locate/jpube

The link between pensions and retirement timing: Lessons from California teachers

Kristine M. Brown

University of Illinois at Urbana-Champaign, 504 E. Armory Ave., Champaign, IL 61820, United States

ARTICLE INFO

Article history:
Received 14 June 2011
Received in revised form 24 October 2012
Accepted 26 October 2012
Available online 17 November 2012

Keywords:
Pension reform
Retirement
Teacher labor supply

ABSTRACT

I exploit a major, unanticipated reform of the California teachers' pension to provide quasi-experimental evidence on the link between pension features and retirement timing. Using two large administrative data sets, I conduct a reduced-form analysis that leverages the nonlinearities in the return to work generated by the pension features and the reform-induced shifts of these nonlinearities for identification. The implied estimates of the elasticity of lifetime labor supply with respect to the return to work are centered around 0.04 in the medium-run and are less than 0.1 in the long-run.

1. Introduction

With the baby boomers reaching retirement age, public officials and private pension managers are scrambling to design policy that will reduce the burden of pension obligations on younger workers and shareholders, while still fulfilling the promises made to those nearing retirement. The proposed reforms will inevitably alter key pension financial incentives faced by members, such as the financial gain for an additional year of work, making the degree to which these incentives affect retirement timing central to the policy debate.

Although there is an extensive literature that addresses the relationship between pensions and retirement, there is no firm consensus on the magnitude of the behavioral response to pension incentives. Recent work has emphasized the importance of forward-looking pension financial incentives to individual retirement decisions and has utilized both structural and reduced-form approaches to estimate the behavioral response to these incentives.[1] These estimation techniques, which assume that retirees facing diverse pension incentives are otherwise identical after controlling for other observable characteristics, prove unsatisfactory as the potential for endogenous sorting makes it difficult to infer the true causal effects of the pension features.

In this paper, I address this concern by using a quasi-experimental approach to estimate the price elasticity of lifetime labor supply, a key parameter for predicting the response of individuals to pension reforms and for measuring the deadweight loss associated with retirement programs. I use two unique administrative data sets to exploit the exogenous variation in the return to work that is generated by the nonlinear features of the California teachers' pension benefits and by the reforms of these pension benefits. The distribution of retirements about the budget constraint nonlinearities reveals how much labor supply responds to changes in the return to work and is the basis for the estimates of the elasticity of lifetime labor supply. The results imply that California teachers' lifetime labor supply is relatively insensitive to the financial return to work.

This paper builds on both a growing literature that uses budget constraint nonlinearities to identify the causal effect of price changes on individual choice and a small literature that uses policy-reform based variation in pension financial incentives to address potential omitted variable bias. Saez (2010), which demonstrates that the compensated price elasticity is proportional to bunching at price schedule kink points in the context of income taxation, is a foundational work of this first literature.[2] More closely related to this paper, Manoli and

E-mail address: kmb@illinois.edu.

[1] Rust and Phelan (1997) present one of the most comprehensive models which incorporates the financial incentives of Social Security and the additional incentives generated by Medicare. Stock and Wise (1990), and Gustman and Steinmeier (1986) are also examples. Samwick (1998) develops a reduced-form variant of the option value model introduced by Stock and Wise (1990) to estimate the effect of Social Security and pensions on retirement behavior. Coile and Gruber (2007) introduce a modification of the reduced-form option value model to disentangle the effect of Social Security financial incentives on retirement. Their approach is also used extensively in a volume edited by Gruber and Wise (2004) and by Asch et al. (2005). Costrell and Podgursky (2009) and Costrell and McGee (2010) are examples of the teacher-pension specific retirement literature.

[2] Studies using related empirical methods examine such diverse topics as the impact of tax rates on earnings (Kleven and Waseem, 2012) and on labor supply (Chetty et al., 2011), the effect of EITC on non-labor income (Weber, 2012), the income and labor supply effects of CHIP (Pei, 2012), the effect of the Saver's Credit on income and savings (Ramnath, 2011), and the response of automakers to fuel economy policy (Sallee and Slemrod, 2012).

0047-2727/$ – see front matter
http://dx.doi.org/10.1016/j.jpubeco.2012.10.007

Weber (2011) find that the extensive margin labor supply is inelastic by exploiting discontinuities in the level of benefits in the Austrian pension system. I contribute to this literature by using discontinuities in the growth rate of pension benefits to estimate the lifetime labor supply elasticity. I also extend the estimation method used in this literature by using policy reforms that shift budget constraint nonlinearities to provide information about the counterfactual labor supply on a linear budget constraint.

The policy-reform retirement literature generally finds a smaller role for financial incentives than non-reform studies. For example, Burtless (1986) and Krueger and Pischke (1992) examine individuals' responses to changes in the level of Social Security benefits and find a small role for Social Security in the trend toward later retirement. More recently, Mastrobuoni (2009) examined the effect of the expected rise in the Social Security normal retirement age on retirement behavior and found somewhat larger effects.[3] This paper adds to this literature by focusing on distortions to the return to work rather than changes in the level of benefits.

The California public school system is an advantageous setting in which to estimate the impact of pension price incentives on retirement timing. California teachers are required to participate in a state pension system with a simple benefit formula, do not participate in Social Security, have tenure, and face a rigid collectively-bargained wage schedule, so there is little uncertainty in the financial return to work and it is both salient to the teachers and easily calculated with administrative data. Importantly, in contrast to the Social Security reforms addressed in the literature that primarily changed benefit levels and the focal retirement age, the California pension reform explicitly altered the financial return to an additional year of work. A further advantage of this study is that a large portion of the sample is women, a group which has arguably been understudied.

Given the minimal employment-related uncertainty faced by the teachers, I use a nonstochastic lifetime budget constraint framework to model their retirement decisions. One salient theoretical prediction of this model is that a bunching of retirements will be observed at budget constraint kinks and discontinuities. In the California teachers' case these nonlinearities are a product of the pension program. I first examine the response of individuals to their pension features and to the pension reform in a flexible way. I construct the prereform and postreform distributions of retirees over age and show that there is a spike in the distribution at the universal prereform budget constraint kink and that this spike shifts to the new kink following the reform. The distributions over service are also consistent with the discontinuity in the level of benefits that is introduced by the pension reform. The reform provides evidence that the distinct retirement pattern is shaped by pension financial incentives rather than other coincident factors.

Next, I incorporate the pension reform into the estimation method introduced by Saez (2010) to quantify the excess retirements at the budget constraint kinks and to estimate the elasticity of lifetime labor supply. I determine the excess retirements as the difference between the pre- and post- reform retirement distributions at points where the kinks are removed or introduced. The estimates of the labor supply elasticity with respect the financial return to work are relatively small with the preferred estimates centered at 0.04. The results imply that teachers are willing to adjust their retirement dates by less than two months in response to a 10% increase in compensation. I investigate the impact of potential extensive-margin frictions, specific to this setting, that may cause the elasticity estimates to be downward biased. These include a high implicit cost to retiring during the school year, a cost to adjusting retirement plans in response to the pension reform, and the cost of health insurance coverage. I find that these factors have little effect on the overall results.

Finally, I use an instrumental variable strategy to estimate an alternative measure of labor supply — the effect of the financial return to working on the probability of working an additional year. This alternative measure of labor supply allows me to compare the behavior of California teachers with findings in the literature. I find that California teachers behave similarly to the Social Security population in the U.S. and the estimated elasticity is similar to the findings of Manoli and Weber (2011) for the Austrian population.

The remainder of this paper is organized as follows. In Section 2, I provide an overview of the California teachers' defined benefit program, the reforms of the program, and the data used in this study. Section 3 introduces the empirical strategy which is based on a simple lifetime budget constraint model that captures the teacher retirement decision. Section 4 presents the main labor supply elasticity estimates based on retirement behavior at budget constraint nonlinearities. Section 5 includes robustness checks for the main results and Section 6 presents alternative labor supply estimates. Section 7 concludes.

2. Background and data

2.1. CalSTRS defined benefit program

California public school teachers are covered by a defined benefit retirement plan administered by the California State Teachers' Retirement System (CalSTRS) which ranks among the ten largest public retirement systems in the United States both in terms of assets and members. The main features of the defined benefit pension resemble those of most employer-sponsored defined benefit retirement programs and also of Social Security. Participation is mandatory for teachers employed full-time in California public schools and upon retirement each CalSTRS member receives a lifetime annuity with an annual value based on years of service, age and past salary. The retirement system is financed with contributions from active members (8% of salary), employing school districts, and the State General Fund, as well as with investment earnings.

The CalSTRS pension is likely to be a prominent component of California teachers' retirement portfolios and an important consideration in the retirement decision for several reasons. First, the CalSTRS pension is the only source of employment-based retirement income for career teachers because California teachers are not simultaneously covered by Social Security. Second, CalSTRS members' pensions are not disrupted as they move between employing public school districts within California. Third, CalSTRS is relatively generous; the average replacement rate for retired teachers is 59% of final annual salary, while the replacement rate for the average Social Security annuitant is only 41% of average annual lifetime earnings.[4]

Features of the pension benefit calculation and the reforms of these features are central to the empirical strategy employed in this paper. Each retired CalSTRS member receives a lifetime annuity with an annual value calculated according to the following formula:

$$B(R,S) = k(R,S) \times S \times w_S^f. \qquad (1)$$

[3] There are also several papers that identify the responsiveness of retirement to financial incentives using a shock created by a temporary retirement incentive program. These include Lumsdaine et al. (1992) and Pencavel (2001). Both find evidence that separation rates are responsive to the incentive but are unable to predict who will leave. The one most closely related to this paper is Furgeson et al. (2006) which uses administrative data to look at the retirement of Pennsylvania teachers and find that the substitution elasticity of retirement is strongly negative. It is not trivial to compare the findings from examining temporary retirement incentive programs to those from examining permanent pension reforms.

[4] Replacement rates are based on the author's calculations using data in the statistical section of the CalSTRS Comprehensive Annual Financial Report (CalSTRS, 2006a) and the Social Security Administration Performance and Accountability Report (Social Security Administration, 2006). The average replacement rate calculated for California teachers corresponds to an average retirement age that is just past 60, while the replacement rate reported for Social Security corresponds to retirement at age 65.

This "unmodified allowance" is a function of years of service *S*, retirement age *R*, final compensation w_S^f, and a benefit factor *k*. Final compensation is the average salary in the previous 1–3 years. The value of the benefit factor is increasing in retirement age *R* and years of service *S*. For individuals retiring prior to 1999 the benefit factor started at 1.4% at age 55, the early retirement age, and was increasing linearly in age at an annual rate of 0.12 percentage points up to the maximum of 2.0% at age 60, the normal retirement age. For example, an individual with 30 years of service retiring at age 60 would receive a benefit equal to 60% of her income.

In August of 1998 the California State Legislature passed two bills, AB 1102 and AB 1150, that provided an unprecedented, and unanticipated increase in the generosity of the CalSTRS defined benefit pension for those that would retire on or after January 1, 1999.[5] Teachers could not plan for these changes as the bills were introduced in their final forms just days before the legislative vote and past attempts to change benefits were unsuccessful. However in Fall 1998, the regular CalSTRS newsletter (CalSTRS, 1998), which is mailed to all teachers, detailed the reforms and provided examples of how these reforms changed benefits.

The legislated reforms altered the pension program solely through changes to the benefit factor, *k* in Eq. (1), while the structure of the program, the general allowance formula, and the normal retirement age remained the same. The first reform, the Enhanced Age Factor (EAF), raised the maximum value for the benefit factor from 2.0% to 2.4%. The effect of this change can be seen in Fig. 1. The prereform schedule and the postreform EAF schedule are identical up to age 60. In the postreform EAF schedule the benefit factor continues to increase at an annual rate of 0.133 percentage points, reaching the new maximum of 2.4% at age 63. The second reform, the Career Bonus (CB), provides a onetime increase of 0.2 percentage points in *k* when 30 years of service is completed. However, the cap on the benefit factor is absolute, so individuals that reach 30 years of service between the ages of $61\frac{1}{2}$ and 63 gain less than 0.2 percentage points, with the bonus falling to zero at age 63. The postreform schedule with the EAF and the CB is represented by the black solid line in Fig. 1. In the postreform period, an individual moves from the "EAF Only" schedule to the "EAF + CB" schedule when she attains 30 years of service.

The reforms shifted the age location of the maximum benefit factor from age 60 to a later retirement age for all teachers. The interaction of the two reforms creates variation in the age location of the postreform maximum across the population. It occurs as early as age $61\frac{1}{2}$ if 30 years of service have been attained or as late as age 63 if 30 years are not worked before this age.

Despite the seemingly small nature of the reforms, their potential impact on retirement benefits was quite large. Most importantly for this paper, the shift of the benefit factor cap caused the present value of the financial return to working an additional year at age 60 to double on average. The effect on the financial return to work is discussed in greater detail in Section 3. Additionally, the unmodified allowance (B) increased by 20% for retirements at age 63 and by at least 10% for retirements after 30 years of service.

2.2. Data description

Two administrative data sets are used to estimate the sensitivity of lifetime labor supply to the financial return to work. The first data set includes counts of new teacher retirees in quarter-year age by half-year service bands for each calendar year 1994–2004 for all public schools in California. The second data set is individual-level administrative data for all retirement-eligible teachers that were employed by the Los Angeles Unified School District (LAUSD) during each academic year 1997–2004, with retirements indicated.[6] In addition to the number of retirements in each year, both data sets include the age and service of each retiree, which are used to calculate individuals' retirement benefits as a multiple of salary.[7] Details of the benefits calculation are in Appendix A.

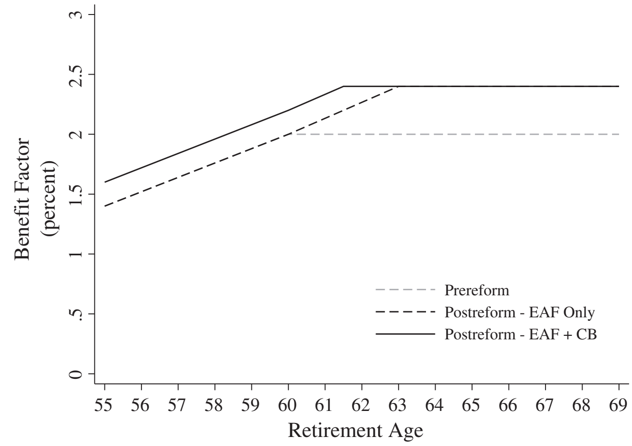

Fig. 1. Benefit factor schedule by retirement age.

The primary advantages of the systemwide retirement count data set are its size — there are over 84,000 retirement observations, and that it captures several years before and after the 1999 reforms. Summary statistics for this data are presented in columns (1) and (2) of Table 1. The annual number of retirements grew over time in proportion to the growth of the California teacher population over age 55. Following the reforms, the average age at retirement increased by less than a year and the average number of years worked under CalSTRS increased by about one year.

The LAUSD data is small by comparison and provides limited prereform coverage. Although there are over 44,000 individual-year observations, only 10% of retirement-eligible teachers retire each year leaving 4652 observed retirements. Also, years of service in LAUSD are recorded rather than total CalSTRS-covered service.[8] However, this data set is richer than the retirement count data providing an opportunity to examine heterogeneity by sex and to perform additional robustness checks. In addition, because active teachers and salary are included, I am able to estimate the effect a dollar increase in the return to work on the probability of working another year as an alternative labor supply measure, which can be compared with the findings of other studies. Summary statistics for observed retirees are shown in columns (3) and (4) of Table 1. Appendix B presents statistics from the Census 2000 to provide a better picture of the characteristics of California teachers than is possible with the administrative data.

3. Lifetime labor supply model and empirical strategy

A simple lifetime budget constraint model captures the major financial incentives of the CalSTRS pension and generates a number of unambiguous predictions for the retirement behavior of CalSTRS

[5] Contributions to the pension were not changed.

[6] The systemwide count data set was constructed with the assistance of the CalSTRS administrative office and begins in July 1994 and ends in June 2004. The LAUSD data set was compiled by Personnel Research and Assessment in LAUSD.

[7] Retirements as measured here require that teachers leave their current jobs and begin collecting pension benefits. However, it should be noted that teachers may return to work in the public schools and collect their pensions, but their benefits will be reduced dollar-for-dollar for earnings that exceed the plan limit. Fong and Makkonen (2012) estimate that 3 to 8% of retired teachers do return to work in this period, which would tend to bias the estimates of lifetime labor supply elasticity upward.

[8] Costs to transferring across districts limit most movement to the first four years of service (Reed et al., 2006).

Table 1
Summary statistics.

	Systemwide		LAUSD	
	Prereform (1994–98)	Postreform (1999–2004)	Prereform (1997–98)	Postreform (1999–2004)
Mean age	60.79	61.16	62.10	62.33
Mean service	26.61	27.72	25.69	27.32
Percent female			69.78	72.65
Mean salary (year 2000 $)			56,862	61,048
Observations	29,350	54,798	1029	3623

Note: The sample includes teachers age 55 or older at the time of retirement.

members in response to these incentives and to the reform of these incentives. This nonstochastic model, which abstracts from bequests and assumes retirement is an absorbing state, highlights the tradeoff between retirement leisure and consumption of market goods in the retirement timing decision, making it a good starting point for estimating lifetime labor supply elasticity with administrative data. In this model the sharp changes in the financial return to work are salient. The empirical strategy will exploit behavior at these budget constraint nonlinearities for identification and estimation.

Individual preferences are defined over two goods, lifetime consumption of market goods C and years of work S. An individual's utility in each period is assumed to be additively separable in consumption and leisure, so that $u_t(c_t,l)=v(c_t)-\phi_t\times l$, where ϕ_t is the disutility from working in period t and l takes the value one if the individual works in that period and is zero otherwise. For exposition, allow T to be the last period of life, known with certainty and let the interest and discount rates equal zero. These assumptions are relaxed in the empirical implementation (see Appendix A). Assuming $v(.)$ is concave with respect to c_t, the individual will maximize utility for any retirement date by perfectly smoothing consumption over the lifecycle so $c_t=\frac{C}{T}$ for all t. Then the lifetime utility can be written as $U(C,S)=T\times v(C/T)-\int_{t=0}^{S}\phi_t\,dt$ with $U_C=v'(C/T)=v'(c), U_S=-\phi_S$, and $U_{SS}=-\phi'_S$.

With a defined benefit retirement program, an individual's lifetime budget constraint is not the sum of wage earnings, as it is in the absence of the program, but rather the sum of lifetime wage earnings net of contributions to the program and pension wealth — the sum of pension income collected in retirement. The budget constraint for a CalSTRS member retiring at some retirement age R after S years of service can be written as $C=\int_{t=0}^{S}w_t(1-\tau_c)dt+\int_{t=R}^{T}B(R,S)dt$, where $B(R,S)$ is the annual retirement allowance as given by Eq. (1) and the last term is pension wealth. Because CalSTRS retirement benefits are a function of years of work and retirement age, I make the simplifying assumption that for teachers near retirement accruing an additional year of service is equivalent to retiring one year older. Specifically $R=S+a_0$, where a_0 is a constant that can be thought of as the effective starting age for an uninterrupted career.

The slope of the budget constraint is the total financial return to working. At any service level S, this "net wage" (w^{net}) is roughly a multiple of annual salary and can be written as

$$\begin{aligned}\frac{dC}{dS}&=w_S\times(1-\tau_c)+\frac{d}{dS}\int_{t=S+a_0}^{T}\left(k_S\times S\times w_S^f\right)dt\\&=w_S\times(1-\tau_c)-w_S^f\times S\times k_S\\&\quad+\left[w_S^f\times\left(\frac{dk_S}{dS}\times S+k_S\right)+\frac{dw_S^f}{dS}\times k_S\times S\right]\times(T-S-a_0).\end{aligned}\tag{2}$$

The second and third terms of Eq. (2) taken together are the pension wealth accrual, the total change to pension wealth for an additional year of work. The second term is the retirement allowance that could have been collected in the current year, but is forfeited to continue working. The third term is the change in annual allowance for delayed retirement accumulated over the shorter retirement period. A key determinant of the budget constraint slope at any point, and the one central to the empirical analysis, is the growth of the benefit factor for an additional year of work, $\frac{dk}{dS}$. Here, if $\frac{dk}{dS}>0$ the pension is generally acting as a subsidy to wage earnings, otherwise it is acting as a tax.

Faced with a convex, differentiable lifetime budget constraint, the labor supply S^* that solves the first order condition, $\frac{-U_S}{U_C}\big|S^*=w_{S^*}^{net}$, determines the utility maximizing retirement date provided $U_{SS}=-\phi'_S<0$. Heterogeneous preferences for working will generate a distribution of career lengths. If the preference parameter is smoothly distributed in the population, individual labor supplies will also be smoothly distributed according to some density function $h(S)$ when the population is faced with an approximately linear budget constraint.

In this framework, for a small change in the slope of the lifetime budget constraint income effects can be ignored and the substitution elasticity of lifetime labor supply at years of work S is

$$e=\frac{dS}{S}\times\frac{w^{net}}{dw^{net}}\tag{3}$$

(Saez, 2010). If the change in slope is not small the elasticity will include income effects. Estimates of the elasticity that rely on cross-sectional variation in the net wage will be biased if taste for work or other unobserved differences in the disutility of working are correlated with the net wages of individuals. The remainder of this section describes the empirical strategy that uses budget constraint nonlinearities and reform-induced shifts in these nonlinearities to address potential biases. The nonlinearities are a function of pension features that affect all teachers so it is unlikely that this source of variation in the return to work is correlated with individual preferences for work.

3.1. Convex kink

In the CalSTRS case, the first type of budget constraint distortion generated by the pension is a convex kink that occurs when the benefit factor cap is reached. As S increases $\frac{dk}{dS}>0$ is relatively constant, until k reaches the legislated maximum (2% in the prereform period and 2.4% in the postreform period) causing $\frac{dk}{dS}$ to fall to zero. The sharp decrease in the slope at this point creates a convex kink in the lifetime budget constraint — at age 60 in prereform period and between ages $61\frac{1}{2}$ and 63, depending on years of service, in the postreform period. The discussion below is framed in terms of a kink at a particular service level, but the empirical implementation will be adjusted to account for the age, rather than service, locations of CalSTRS kinks.

To see how a kink affects the retirement pattern, consider a population that faces a linear budget constraint with a slope of w_H^{net}. Each individual retires at S^*, such that his ratio of marginal utility of labor to marginal utility of consumption is equal to the net wage. As above, assuming heterogeneous preferences smoothly distributed across the population and no discontinuities in the cost of working, S^* is also smoothly distributed according to some function $h(S)$. If a kink is introduced to the budget constraint, so that the budget constraint slope falls to $w_L^{net}<w_H^{net}$ for $S\in[S_K,T]$ individuals with $S^*>S_K$ may adjust their retirement dates.

Following Saez (2010), with the introduction of a small kink, there exists an individual that will adjust retirement from S_H^* to S_K, and whose indifference curve will be exactly tangent to the upper segment of the budget constraint at S_K, so that $\frac{-U_S}{U_C}\big|_{S_K}=w_L^{net}$, as shown by U_L in Fig. 2. The change in lifetime labor supply for this individual, $S_H^*-S_K$, is dS for $dw^{net}=w_H^{net}-w_L^{net}$. S_H^* can be estimated by noting that all individuals with a lifetime labor supply of $S^*\in[S_K,S_H^*]$ when faced with the linear budget constraint of slope w_H^{net} will also locate

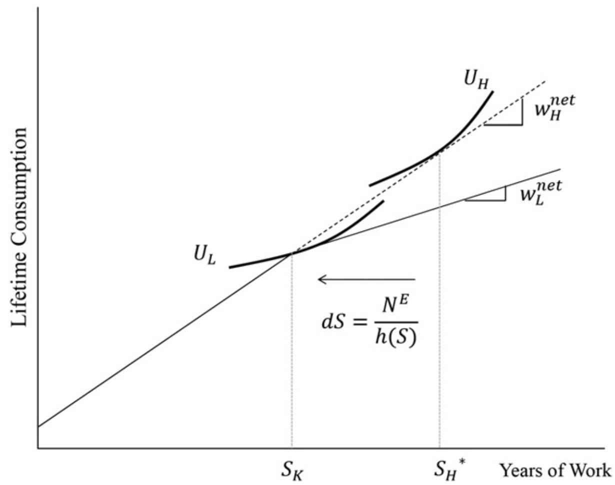

Fig. 2. Elasticity estimation with a convex kink. Notes: This figure depicts the retirement decision when a kink is introduced to the budget constraint at S_K. With the introduction of the kink all those with $S_K<S^*<S_H^*$, will move to S_K. The distance between S_H^* and S_K is the total movers, N^E, divided by the retirement density across these service levels when the budget constraint is linear.

at S_K when the kink is introduced causing bunching in the retirement distribution.

The total number of excess kink retirements is $N^E = \int_{S_K}^{S_H^*} h(S)dS$, where $h(S)$ is the density of retirees when the budget constraint is linear. The distance between S_K and S_H^* for a given change in net wage at the kink is increasing in the labor supply elasticity. Thus the bunching at the kink point is proportional to the substitution elasticity of lifetime labor supply. This relationship is used to estimate dS in Eq. (3). The elasticity can then be approximated as

$$e = \frac{N^E/h(S)}{S_K * ln\left(w_H^{net}/w_L^{net}\right)}, \tag{4}$$

where $N^E/h(S)$ is dS (Saez, 2010; Weber, 2012). The net wage and the change in net wage at the kink (as a multiple of annual salary) can be calculated using the pension formula. The details of the empirical implementation are described in Section 4.2.

Eq. (4) will only yield the lifetime labor supply elasticity of substitution in the case of a small net wage change, otherwise it will be contaminated with income effects. The slope change at the CalSTRS kinks is a relatively large 50% decrease. The assumption of a small change is relaxed in Section 4.2. The assumption that the underlying distribution $h(S)$ is smooth is examined in Section 5.

3.2. Discontinuity

The second distortion created by the pension is a discontinuity in the budget constraint. This discontinuity is the result of the career bonus legislation and occurs at 30 years of service for some participants in the postreform period only. At the service threshold, hereafter S_D, the change in the benefit factor is positive for an infinitesimal change in service, $\frac{dk}{dS}\rightarrow\infty$, creating the discontinuity in the lifetime budget constraint. The discontinuity only exists for those teachers that do not reach the binding benefit factor cap before completing 30 years of service. I define the High Service group as those that will receive the full bonus ($\Delta k = 0.2$ percentage points) at 30 years of service and the Low Service group as those that will receive no bonus at 30 years of service.

To see how this discontinuity affects the retirement distribution, again assume that preferences for work are smoothly distributed across the population and that when faced with a linear budget constraint each individual will retire such that $\frac{-U_S}{U_C}|_{S^*} = w^{net}$, resulting in a smooth distribution of retirements across along the service dimension. If a discontinuity is introduced to the linear budget constraint at S_D for all individuals, individuals with S^* to both sides of the discontinuity have an incentive to adjust their retirement dates toward S_D.

In this context, there exists an individual with $S_L^*<S_D$ that will be exactly indifferent between her old retirement date and S_D on the new budget constraint such that $U(S_L^*) = U(S_D)$ as shown in Fig. 3. This individual will move to S_D and the change in lifetime labor supply in response to the discontinuity is $dS = S_D - S_L^*$. Each individual with $S_L^* \le S^* \le S_D$ will also delay retirement to the discontinuity, while those with $S^*<S_L^*$ will not adjust their retirement dates at all. With a homogeneous elasticity across the population and no frictions, this creates a gap in the retirement distribution such that there will be zero retirements between S_L^* and S_D. The "missing" retirements will reappear exactly at S_D creating a spike in the distribution at this point.

These shifted retirements, $N^G = \int_{S_L^*}^{S_D} h(S)dS$, are proportional to the lifetime labor supply elasticity of substitution. Analogous to the kink case, $dS = \frac{N^G}{h(S)}$, where $h(S)$ is the density when the budget constraint is linear. Empirically the density does not go to zero preceding the gap, but in the absence of optimization frictions the estimated elasticity can be interpreted as the average elasticity of the population.

However, unlike the kink case, the discontinuity does not create exogenous variation in the marginal return to work, so calculating the elasticity of lifetime labor supply with the estimate of dS is not as straightforward. I approximate the marginal return to work with the average return between S_L^* and S_D. This is similar to the approach in Kleven and Waseem (2012), though in their context the nature of the discontinuity allowed them to estimate optimization frictions, an advantage not present here. In this strategy $dw^{net} = w_D^{net} - w^{net}$ where w^{net} is the return to work at S_L^* when there is no discontinuity and w_D^{net} is the average return to work between S_L^* and S_D when the discontinuity at S_D is present. This effectively imposes a nonconvex kink at S_L^* years of service as shown by the dashed line in Fig. 3. Of course, faced with such a budget constraint, the individual that was located at S_L^* on the linear budget constraint with slope w^{net} would be better off by locating somewhere between this point and S_D, so the estimated elasticity will be an upper bound for the true parameter. The alternative structural estimation strategy requires imposing a functional form on utility and using the constraint $U(S_L^*) = U(S_D)$ to directly estimate the elasticity.

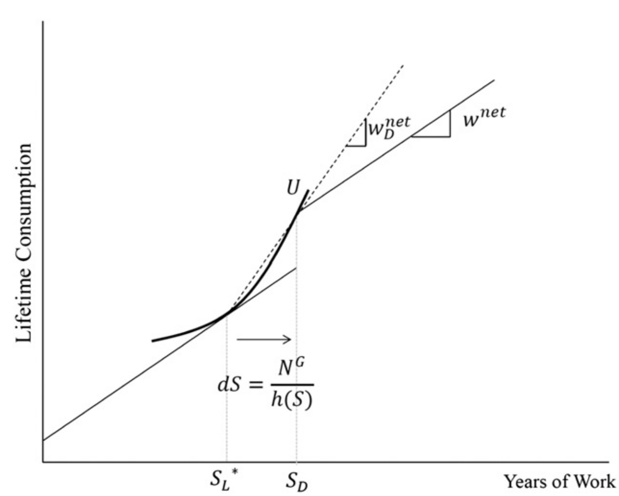

Fig. 3. Elasticity estimation with a discontinuity. Notes: This figure depicts the retirement decision when a discontinuity is introduced to the budget constraint at S_D. With the introduction of the discontinuity individuals with $S_L^*<S^*<S_D$ will move to S_D. The distance between S_L^* and S_D is the total movers from below, N^G, divided by the retirement density across these service levels when the budget constraint is linear. Individuals with $S^*>S_D$ will also reduce their labor supplies.

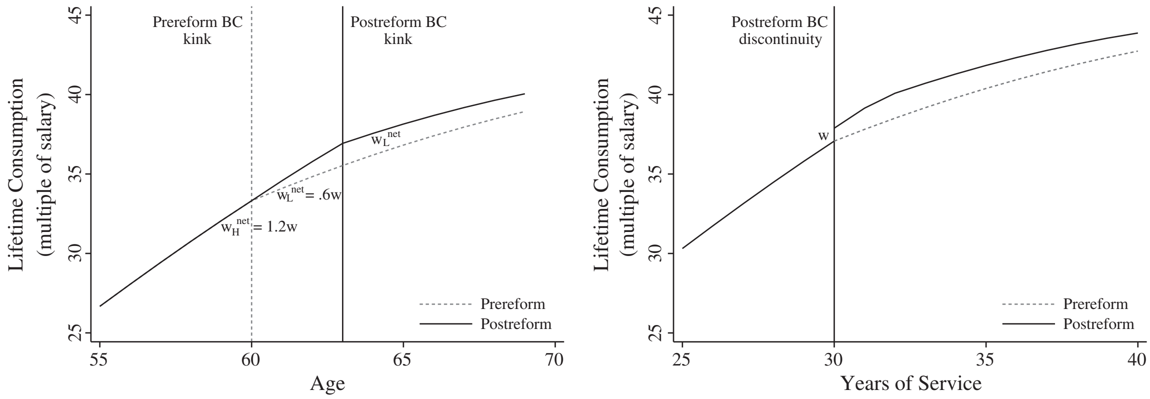

Fig. 4. Stylized budget constraints in age and service. Notes: The budget constraint in age (left) depicts the lifetime budget constraint for the median CalSTRS member that would have 27 years of service at age 60. The budget constraint in service (right) depicts the lifetime budget constraint for a CalSTRS member that would have 30 years of service at age 60. For both, total consumption is the present discounted value of salary and future pension payments (discounted to age 55) scaled by the median salary at age 60 ($55,000). The discount factor is assumed to be 0.97, salary is assumed to grow by $1000 per year and the individual is assumed to live to age 80 with certainty. Changes to the discount rate change the level of the budget constraint but the percentage change in the slope at the kink point remains the same.

While other work has focused on the spike in retirements at S_D (Manoli and Weber (2011) for example), I focus on the gap preceding S_D for two reasons. First, the missing retirements will not be equivalent to the bunching at S_D as individuals with $S^*>S_D$ will reduce their labor supplies in the presence of non-zero income effects. Estimates of the substitution elasticity of lifetime labor supply based on the excess retirements at S_D will then be overstated. Second, because the discontinuity was introduced by the reform, it will take time for those that delayed retirement to reach 30 years of service and retire, but the gap can be observed immediately.[9]

4. Empirical results

In this section, I examine the evolution of the aggregate retirement distributions over age and service, with particular attention to the action at budget constraint nonlinearities that were shifted by the pension reform. I then quantify the responsiveness of individuals to pension financial incentives by estimating the lifetime labor supply elasticity using retirement behavior around the budget constraint nonlinearities, as discussed in Section 3.

4.1. Graphical evidence

A comparison of the prereform and postreform retirement distributions across age and years of service, relative to the changes in the budget constraint faced, provides visual evidence of the response of CalSTRS members to the financial return to work and to the pension reform. The retirement distributions are constructed for the prereform period (1995–98) and the postreform period (1999–2003) using the systemwide CalSTRS administrative data. For each year, the number of retirements in each age or service band as a fraction of total retirements is calculated. The annual densities are averaged with equal weight to construct the distributions for the prereform and postreform periods.

In the prereform period all individuals have a convex kink in their budget constraints at age 60, where the benefit factor reaches the 2.0% cap and the net wage for continued work falls from approximately 1.2 times the annual salary to about 0.6 times the annual salary. The location of the kink in the postreform period is the age at which the benefit factor hits the cap of 2.4%. The earliest possible kink age is $61\frac{1}{2}$, for the High Service group, and the latest possible age is 63 for the Low Service group. These High Service and Low Service groups, those attaining 30 years of service before age $61\frac{1}{2}$ and after age 63 respectively, include over 95% of the retiring population in each year, and will be the focus of the graphs. A stylized budget constraint depicting this reform is in Fig. 4.

The average age distributions of annual retirements for the prereform and postreform periods are shown for the High Service group in Fig. 5a and for the Low Service group in Fig. 5b. For each retirement age R, the fraction includes all retirees that retire at ages $\in[R, R+.25)$. As expected there is a spike in the density of retirements at age 60 ± 3 months in the prereform period and the density at this age drops by over $\frac{1}{3}$ in the postreform period. However the persistent spike at age 60 indicates potential adjustment frictions that will be discussed in Section 5. For the High Service group, the fraction of retirees locating at the new kink of $61\frac{1}{2}\pm3$ months has doubled from 4% to 8%. For the Low Service group, the increase in retirements at the new kink of age 63, from 3.5% to 4.5%, is not as pronounced. The smaller effect at this later age may be partially attributable to delayed transition between the prereform and postreform equilibria, which is discussed in more detail in Section 5. The Low Service group provides a cleaner test of the predicted shift of the retirement distribution because its budget constraint is altered by only a shift in the kink point, whereas the High Service group also faces a new discontinuity at 30 years of service in the postreform period. It is also worth noting the spike in retirements at the early claiming age of 55, a feature observed in most social insurance systems.

The shift of the spike in the retirement distribution from age 60 to the age that coincides with the group-specific postreform kink (age $61\frac{1}{2}$ or 63) is the key evidence supporting the existence of a causal link between the financial return to work and retirement timing. The reform provides a test that the strong preference for retiring at age 60 is influenced by financial incentives rather than non-pension coincident incentives.

The reforms also introduce a discontinuity at 30 years of service for those that do not hit the benefit factor cap beforehand. The stylized budget constraint for the High Service group, who were eligible for the maximum bonus at exactly 30 years of service, with consumption as a function of service, is shown in Fig. 4. The service location of the postreform kink varies across individuals and so does not make aggregate predictions for the retirement distribution on the service dimension. The effect of the reform is reflected by the shift from the dashed line to the solid line. The magnitude of the discontinuity at 30 years of service is approximately equal to the annual salary for the median CalSTRS retiree. There is no change in the budget constraint in service for the Low Service group because they are unaffected by construction.

The prereform and postreform retirement distributions along the service dimension for the affected High Service group and for the

[9] This issue is compounded by a shorter observation window due to contamination by a pension reform in 2001 that also affected benefits at 30 years of service.

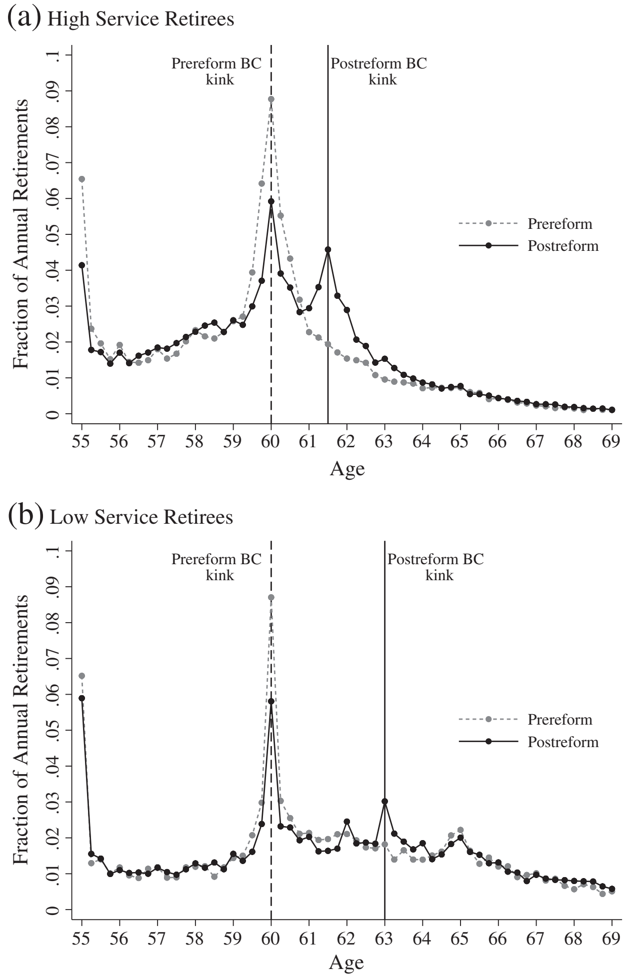

Fig. 5. Age distribution of retirees by service group. Notes: This figure depicts the density of retirements between ages 55 and 69. The annual age distributions were constructed using the CalSTRS systemwide administrative data. These were then averaged, with equal weight, to create the prereform (1995–98) and postreform (1999–2003) distributions. The High Service group is defined as having the ability to work 30 years by age $61\frac{1}{2}$ and the Low Service group is defined as being unable to work 30 years before age 63. The prereform and postreform budget constraint kinks are indicated by the dashed vertical line and the solid vertical line, respectively. The return to working an additional year falls by 50% at the kinks.

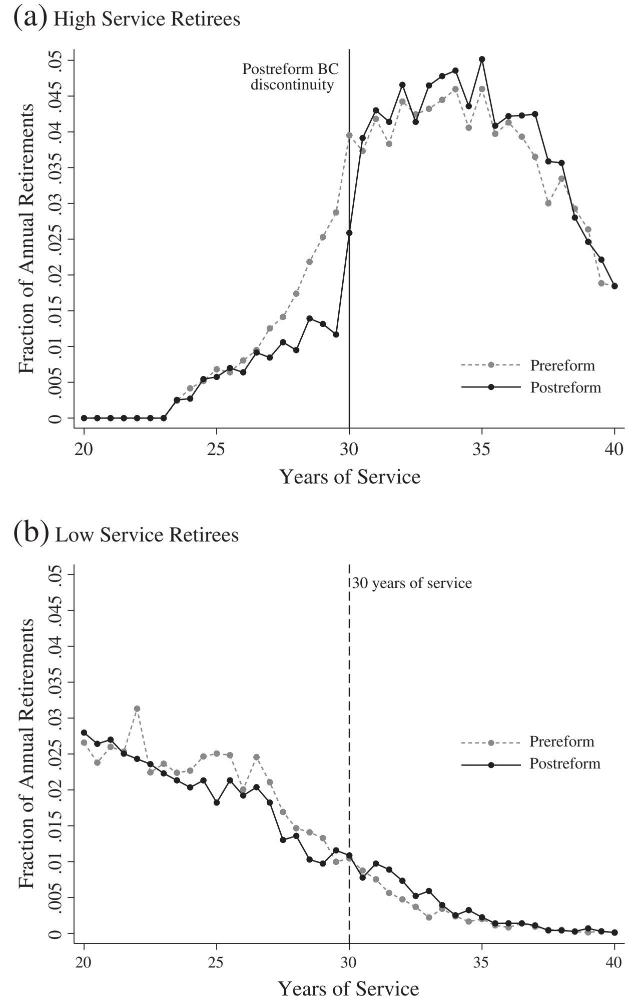

Fig. 6. Service distribution of retirees by service group. Notes: This figure depicts the density of retirements between 20 and 40 years of service. The annual service distributions were constructed using the CalSTRS systemwide administrative data. These were then averaged, with equal weight, to create the prereform (1995–98) and postreform (1999–2000) distributions. The High Service group is defined as having the ability to work 30 years by age $61\frac{1}{2}$ and the Low Service group is defined as being unable to work 30 years before age 63. The postreform budget constraint discontinuity for the High Service group is indicated by the solid vertical line. There is no discontinuity for the Low Service group. This sample includes only retirements through 2000 because a later reform, effective in 2001, also provided incentive to delay retirement to 30 years of service.

unaffected Low Service group are shown in Fig. 6a and b respectively. As expected, there is a gap between the pre- and post-reform distributions for the High Service group. The cumulative gap is over 7% of the population, with most of the shift occurring between 27 and 30 years of service. While the ratio of the fraction of retirements occurring at 30 years of service to the fraction at 29 years of service is larger in the postreform period, a spike at 30 years of service is not observed. This is in part because those that have delayed retirement must accumulate 30 years of service before their retirements will be observed. For the Low Service group, the prereform and postreform distributions are similar and reveal teachers have no preference for retirement at a particular service level; this is in line with their financial incentives.

4.2. Elasticity of lifetime labor supply

Supported by the visual evidence in Section 4.1 that the teachers responded to the change in pension financial incentives created by the reform, I quantify their sensitivity to these incentives in this section. As described in Section 3, I use a reduced-form approach to estimate the elasticity of lifetime labor supply using retirement behavior around budget constraint nonlinearities. The main results focus on the kink point bunching at age 60 and the change in the return to work at that age, but I also estimate the elasticity using the postreform discontinuity at 30 years of service.

The bunching at age 60 is estimated as the difference between the number of retirees at the kink in the prereform period and the number of retirees at age 60 in the postreform period when the budget constraint is linear at this age. It is likely that there are some impediments to retiring exactly at age 60. For example most teachers retire in the summer, between academic years. To address this a retirement is treated as a kink retirement if the retirement age is within some bandwidth δ of age 60, $R \in [R_K - \delta, R_K + \delta)$. The specific summer-retirement constraint is also addressed directly in Section 5. Because the kink occurs in the age dimension rather than the service dimension, $h(S)$ and S_K in Eq. (4) will be replaced with $h(R)$ and the mean service level at R_K respectively in the empirical estimation. In the modified Eq. (4) the excess retirements at the kink are estimated as $\hat{N}^E = H_K(R) - H(R)$ where $H_K(R)$ and $H(R)$ are the raw cumulative distributions within the specified bandwidth of age 60 in the pre- and post-reform periods respectively and $\hat{h}(R) = H(R)/2\delta$. S_K is estimated as the average service of

Table 2
Elasticity of lifetime labor supply estimates.

	Kink: age 60	Kink: age 60	Kink: age 60	Kink: age 60	Kink: age 60	Discontinuity: service 30
	(1)	(2)	(3)	(4)	(5)	(6)
	Kink = 60 years +/− 3 months					*Gap = 27–30 years*
Elasticity	0.0174***	0.0114***	0.0431***	0.0310***	0.0396***	0.0586***
	(0.0000)	(0.0000)	(0.0002)	(0.0003)	(0.0006)	(0.0011)
	Kink = 60 years +/− 6 months					*Gap = 25–30 years*
Elasticity	0.0256***	0.0175***	0.0646***	0.0543***	0.0668***	0.0932***
	(0.0001)	(0.0001)	(0.0003)	(0.0005)	(0.0009)	(0.0023)
	Kink = 60 years +/− 9 months					*Gap = <30 years*
Elasticity	0.0301***	0.0212***	0.0751***	0.0580***	0.0653***	0.1648***
	(0.0001)	(0.0001)	(0.0003)	(0.0005)	(0.0009)	(0.0042)
Sample	Systemwide, all	Systemwide, High Service	Systemwide, Low Service	LAUSD, all	LAUSD, men only	Systemwide, High Service

Notes: Each coefficient is an estimate of the elasticity of lifetime labor supply. Columns (1)–(5) are based on the bunching of retirees at the kink retirement age of 60 in the prereform period and column (6) is based on the absence of retirements preceding the discontinuity at 30 years of service in the postreform period. Columns (1)–(3) use retirement observations from the years 1994–97 and 2000–04, columns (4) and (5) use 1997 and 2000–04. Column (6) only includes observations through 2000 (including 1999) due to an additional pension reform in 2001 that changed benefits for retirement at 30 years of service. Bootstrapped standard errors (500 replications) are in parentheses.
***Significant at the 1% level, **significant at the 5% level, and *significant at the 10% level.

retirees within the specified bandwidth of the kink, and $\hat{w}_H^{net}$ and $\hat{w}_L^{net}$ are the average net wages to either side of the kink within the specified bandwidth in the prereform period. I exclude years 1998–99 from the estimation to remove reform-anticipation effects and any response to wealth shocks. All results are shown for δ values of 0.25, 0.50, and 0.75 years with nonparametric bootstrapped standard errors.

The results for the full population using the systemwide count data are in column (1) of Table 2 and the estimates using the individual-level data from the Los Angeles school district are in column (4). Each column of Table 2 includes three elasticity estimates with each column corresponding to the indicated data set and sample and each row corresponding to the indicated bandwidth. For both data sets the estimated elasticity increases as the bandwidth increases, as expected. The results range from 0.017–0.030 with the systemwide data to 0.031–0.058 with the LAUSD data. The elasticities are precisely estimated and significantly different from zero at the 1% level in all cases, but the magnitudes are relatively small. Taken together, these estimates imply that teachers were willing to extend their working careers by about one year in response to the pension reform which increased the financial return to working by 100%. Columns (2) and (3) provide the same estimates for service subgroups of the teaching population. The estimated elasticities for the Low Service group are approximately three and a half times those for the High Service group but are still relatively small, implying less than a two month increase in lifetime labor supply in response to a 10% increase in the return to working. Although the movement in the retirement distributions is large, the elasticity estimates are relatively small because the change in the return to work at the kink is also large.

This estimation method assumes that the change in the return to work at the kink is small. This assumption can be relaxed by imposing a specific utility function. I use the quasilinear function $U = C - \frac{S^{1+1/e}}{1+1/e} \times \alpha$ with elasticity e and taste shifter α. The results are shown in column (5) of Table 3.

There is little demographic information available in the administrative data, but the individual-level data set from LAUSD does include sex. Column (5) of Table 2 replicates the estimates in column (4) using only the subsample of men, approximately 30% of the population. In contrast to the usual findings, the results suggest that men's labor supply is mildly more elastic than that of the population as a whole. In this context, the result is consistent with men having better outside options for

Table 3
Elasticity of lifetime labor supply estimates — robustness checks.

	Kink: age 60	Kink: age 60	Kink: age 61.5	Kink: age 63	Kink: age 60	Kink: age 60, 63
	Summer birthday	Counterfactual	Counterfactual	Counterfactual	Semi-parametric	MLE
	(1)	(2)	(3)	(4)	(5)	(6)
	Kink = age +/− 3 months					
Elasticity	0.0236***	0.0257***	0.0042***	0.0113***	0.0163***	0.0641***
	(0.0004)	(0.0001)	(0.0000)	(0.0002)	(0.0000)	(0.0079)
	Kink = age +/− 6 months					
Elasticity	0.0437***	0.0563***	0.0092***	0.0194***	0.0241***	0.1282***
	(0.0007)	(0.0001)	(0.000)	(0.0003)	(0.0001)	(0.0115)
	Kink = age +/− 9 months					
Elasticity	0.0610***	0.0964***	0.0076***	0.0358***	0.0280***	0.1730***
	(0.0009)	(0.0002)	(0.0000)	(0.0003)	(0.0001)	(0.0140)
Data set	LAUSD	Systemwide	Systemwide	Systemwide	Systemwide	LAUSD
Sample	All born in summer	All	High service	Low service	All	Low service

Notes: Each coefficient is an estimate of the elasticity of lifetime labor supply. Column (1) is the reform-based elasticity estimate at age 60 for individuals with summer birthdays in the years 1997 and 2000–2004. Columns (2)–(4) are non-reform-based elasticity estimates. Column (2) shows the elasticity estimate based on the prereform bunching at age 60 relative to the estimated counterfactual in the years 1994–97. Columns (3) and (4) show the same estimate at each of the postreform kinks for the years 2000–04 and 2002–04 respectively. The postreform periods are adjusted for this estimation to allow those that have delayed retirement in response to the reforms to reach the new kink. Column (5) is a semiparametric reform-based estimation using the same sample as column (1) of Table 2 (years 1994–97, 2000–04). Column (6) presents the maximum likelihood results, which used pooled data for the years 1997–2004. Bootstrapped standard errors (500 replications) are in parentheses.
***Significant at the 1% level, **significant at the 5% level, and *significant at the 10% level.

work, as individuals can collect their pensions and work outside the California public school system without penalty.

In the final column of Table 2 are the elasticity estimates using the discontinuity at 30 years of service, for the affected the High Service group only. The gap in the retirement distribution just before the discontinuity is $\hat{N}^G = H(S) - H_D(S)$ where $H(S)$ and $H_D(S)$ are the cumulative distributions of retirements for $30-\delta$ years of service in the pre- and post-reform periods respectively. The labor supply response to the discontinuity in Eq. (3) is estimated as $\hat{d}S = N^G/h(S)$, where $\hat{h}(S) = H(S)/\delta$. The net wage absent the discontinuity is estimated as the average net wage at $30 - dS$ and $\hat{w}_D^{net}$ is the average return to work between $30 - \hat{d}S$ and S_D, including the discontinuity. The elasticity estimates for the systemwide data High Service group are shown for δ values of 3 years, 5 years and 30 years. The discontinuity estimates are larger than the kink estimates and imply that a teacher will work about 4 additional months for a 10% increase in the financial return to working. It is especially high considering that only the High Service group is included in the sample. As discussed in Section 3 this estimate should be treated as an upper bound, but other factors may contribute to its higher value. First, the institutional feature of summer retirements should not affect individuals' ability to delay retirement to exactly 30 years of service. Second, unlike the kink estimate, it does not capture the response to income effects.

5. Robustness to frictions and unobserved costs

The reform-based lifetime labor supply elasticity estimates in the previous section were small. This section explores potential frictions and other unobserved costs associated with retirement in the California teachers' context that would cause the elasticity to be underestimated.

5.1. Frictions

The first type of friction is an optimization friction that may be created by the institutional features of employment as a public school teacher. Traditionally K-12 schools are in session September–June, with students advancing to the next grade the following September. Empirically, over 80% of all teachers retire in the summer months between sessions. An implicit cost to retiring mid-year is a potential explanation for the breadth of the spike at the kink retirement age of 60 as teachers with birthdays distributed across all months retire in this short window. While the bandwidth of $\delta = 0.75$ should capture all of the kink bunching in the presence of this friction, I address this issue directly with the individual-level LAUSD data. I estimate the elasticity of lifetime labor supply at the kink of age 60 as in Section 4.2 but the sample is limited to those teachers whose birthdays are during the summer months and thus should be able to frictionlessly retire at the exact kink age. The results are in column (1) of Table 3 for bandwidths of 0.25, 0.50, and 0.75 years around age 60, though the larger δ values are not needed here to address the potential summer retirement friction. The results are on average smaller than those that included the full sample in column (4) of Table 2.

The second type of optimization friction is a dynamic friction. Retirement-age teachers have likely already made some plans regarding their impending retirements and adjusting these plans may be costly. The persistence of the spike in the retirement distribution at age 60 following the reform is consistent with such an adjustment cost. A dynamic friction would mitigate the response to the pension financial reforms in the short run and cause the excess retirements at the kink and thus the labor supply elasticity would be understated.

To address this I again estimate the elasticity using the bunching at the prereform budget constraint kink point, but estimate the excess retirements as the difference between the observed retirement distribution in the presence of the kink and the estimated counterfactual distribution for a linear budget constraint. The counterfactual distribution is estimated from the observed distribution in the presence of the kinked budget constraint. Following Saez (2010), the excess kink retirements are estimated as $\hat{N}^E = H_K(R) - H_-(R) - H_+(R)$ where $H_K(R)$ is the raw cumulative distribution between $R_K - \delta$ and $R_K + \delta$, $H_-(R)$ is the raw cumulative distribution between $R_K - 2\delta$ and $R_K - \delta$, and $H_+(R)$ is the raw cumulative distribution between $R_K + \delta$ and $R_K + 2\delta$. The counterfactual density in the modified Eq. (4) is estimated as $\hat{h}(R) = (H_-(R) + H_+(R))/2\delta$. Looking at the empirical retirement distributions, it is clear that this estimate will attribute virtually all of the mass at age 60 in the prereform period to a behavioral response to the sharp change in the financial return to working.

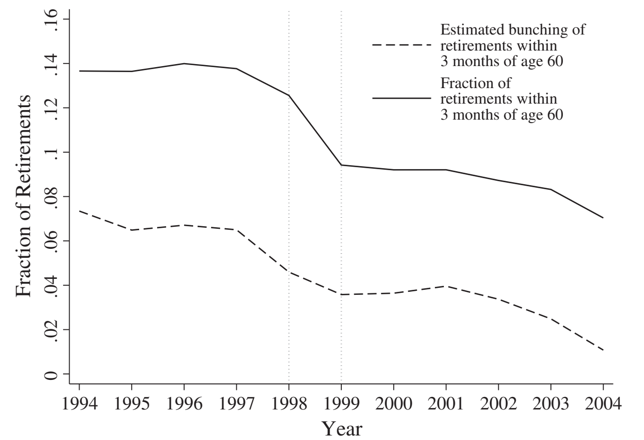

Fig. 7. Time trend in retirement density at the prereform budget constraint kink. Notes: The trends reflect the density of retirements within three months of age 60. The solid line is the raw fraction of teachers at the kink and the dashed line is the estimated bunching. There is one observation per year. The first vertical line indicates the year in which individuals were first notified of the pension reforms and the second vertical line indicates the year the reforms were effective.

As expected, the elasticity estimates in column (2) of Table 3 are larger (than column (1) of Table 2) but still predict that teachers will increase their labor supply by less than half a year in the long-run in response to a 10% increase in the financial return to work. In addition, it appears that it will be some time before the (assumed) equilibrium postreform distribution of retirements is observed. Fig. 7 shows the time trend of the raw fraction of CalSTRS retirees that retire within 3 months of age 60 and the estimated excess retirements at this age. There is a clear decrease in retirements at age 60 when the reforms are announced and become effective, but the continued decline is much slower. The counterfactual elasticity estimates assume bunching is zero.

5.2. Health insurance and unobserved costs of work

Potential unobserved costs or financial returns to work may affect the elasticity estimates. Health insurance is perhaps the most notable form of compensation excluded from the preceding modeling and analysis of the retirement decision. Employer-sponsored health insurance may be valued by older workers, because it increases access to and reduces the cost of coverage.[10] California teachers' retiree health insurance is not centralized and although school districts are required to offer continued coverage there is substantial variation in employer premium support.

Omitted health insurance may cause the elasticity estimate to be downward biased if district premium support differs for active and

[10] The evidence of the importance of health insurance coverage in the retirement decision is mixed. Several studies (Blau and Gilleskie, 2001; French and Baily Jones, 2011; Gruber and Madrian, 1995; Johnson et al., 2003; Rust and Phelan, 1997) find that health insurance availability increases retirement a moderate to significant amount, while others (such as Gustman and Steinmeier, 1994) find a much smaller role for health insurance in explaining retirement behavior. See Currie and Madrian (1999) for a review.

retired teachers or if premium support increases discontinuously at the budget constraint kink at age 60. It is not likely that such factors will substantively change the results in the CalSTRS case. First, on average health insurance premiums are valued at up to 20% of salary, so even if a district pays 100% of premiums for active teachers and 0% for retirees the small elasticity estimates would increase by about 25%.[11] However, it is most common for employers to contribute to premiums until Medicare eligibility at age 65 while only 19% offer no contributions for retirees (CalSTRS, 2006b), mitigating this effect in the aggregate. Second, evidence from large school districts suggests that premium support is more commonly tied to years of service in the district rather than retirement age. This is the case in LAUSD, which employs 10% of all teachers. There is no evidence that the LAUSD service requirements coincide with age 60 and as shown in Table 2, the elasticity estimates for LAUSD are larger but not out of line with the estimates from the full state.

Another possible cause of the persistent spike at age 60 after the reform is that there is an increase in the cost of working at this age. A retirement age norm of age 60 is perhaps the most plausible example. If this is the case, the difference in bunching between the prereform and postreform periods does accurately represent the response to the pension reform, but does not map to the elasticity in a straightforward way. To address this, I estimate the lifetime labor supply elasticity at the postreform retirement age kinks (ages $61\frac{1}{2}$ and 63). At these ages there is no excess density before the reform, suggesting that there is no unobserved change in the cost (or return) to work at these ages. A reform-based estimation is not possible at these ages because the budget constraint preceding the kink has also changed, so I estimate the elasticity using the same counterfactual method that was used for the estimates in column (2) of Table 3. The sample period is limited to allow individuals that delayed retirement in response to the reform to reach the new kink ages. The estimates are in columns (3) and (4) of Table 3. As before, the Low Service group is more responsive than the High Service group, but the elasticities are on average smaller than the reform-based estimates.

6. Alternative specifications

The above analysis focuses on estimating the elasticity of lifetime labor supply using retirement behavior around discontinuities in the lifetime budget constraint. Here I briefly consider two alternative estimations. The first takes a more structural approach to estimate the elasticity of lifetime labor supply and the second estimates instead the effect of pension financial features on the probability of working an additional year.

6.1. Structural estimation

The simple structural estimation explicitly models the full budget constraint faced by each teacher, including the nonlinearities, to remove the endogeneity of the net wage that is due to its simultaneous determination with the retirement age.[12] I estimate this model by maximum likelihood with the pooled pre- and post-reform LAUSD administrative data set of all teachers that retired in 1997–2004. I restrict the sample to Low Service teachers, those that will not be affected by the discontinuity in pension accrual at 30 years of service, in order to simplify the structural modeling and to minimize the effects of measurement error in service. They face a two segment piecewise linear budget constraint with a kink at age 60 before the reform and at age 63 after the reform. With a labor supply function of the form $lnS_i = elnw_i^{net} + \varepsilon_i$ where ε_i is the individual taste for work distributed normally with mean μ_ε and variance σ_ε^2 and the elasticity of lifetime labor supply with respect to return to work is denoted by e, the log likelihood is

$$logL = \sum_i s_i \times log\left\{\phi\left(\frac{lnS_i - elnw_i^{net} - \mu_\varepsilon}{\sigma_\varepsilon}\right)\right\}$$
$$+\sum_i K_i \times log\left\{\Phi\left(\frac{lnS_{K,i} - elnw_{L,i}^{net} - \mu_\varepsilon}{\sigma_\varepsilon}\right) - \Phi\left(\frac{lnS_{K,i} - elnw_{H,i}^{net} - \mu_\varepsilon}{\sigma_\varepsilon}\right)\right\},$$

where s_i is an indicator for retirement on a budget constraint segment and K_i is an indicator for retirement on a kink.[13] The estimated parameters are the price elasticity e, μ_ε and σ_ε.

With the pooled pre- and post-reform data it is expected that the elasticity is effectively estimated from reform-based variation in the net wage, similar to Friedberg (2000). However, it is important to note that the reform is not otherwise explicitly incorporated in the estimation strategy, and identifying variation is also coming from the cross section. The estimates are shown in column (6) of Table 3 for each of the three kink bandwidths used in the reduced form analysis. The results range from 0.064 for the 3 month bandwidth to 0.173 for the 9 month bandwidth. These are 1.5 to 2 times larger than the reform-based kink estimates for this population (not included in table). This suggests that even after leveraging the substantial reform-induced variation, these estimates may be upward biased. The direction of the bias is consistent with a positive correlation between the taste for working and wages.

6.2. Probability of working

As a complement to the elasticity estimates in the previous section, I estimate an additional measure of labor supply — the probability that a teacher does not retire each period subject to his or her financial return to working. This allows me to compare the findings of this study to the recent literature that uses survey data and studies more diverse populations, a context in which budget constraint nonlinearity methods are often infeasible.

The empirical specification takes the general form $Pr(y_{i,t} = 1|DBP_{i,t}, X_{i,t}) = \beta_0 + DBP'_{i,t}\beta_1 + X'_{i,t}\beta_2 + \varepsilon_i$ where $y_{i,t}$ takes the value of one if the individual does not retire in year t and is zero if the individual retires in year t (which is the last year she is included in the panel data), $DBP_{i,t}$ are the pension financial variables which include the level of benefits and the financial return to continued work, and $X_{i,t}$ includes other controls, such as age dummies. As seen in Section 3, the 1-year net wage will only be sufficient to capture the forward-looking financial return to delaying retirement in the absence of budget constraint discontinuities or non-convexities. For this reason, I capture the forward-looking financial incentive to delay retirement in a flexible way by estimating specifications that include pension wealth accrual over several time periods and "peak value," a measure introduced by Coile and Gruber (2007). The peak value is the difference between the pension wealth associated with retirement in the current year and the pension wealth associated with retirement at the future age with the highest expected pension wealth.

The model is estimated with OLS using the LAUSD data, which includes active teachers, for academic years 1997–2000 (academic year t is September t–August $t+1$). Potential omitted variable bias is addressed with an instrumental variable approach that exploits the 1999 pension reform and an additional pension reform, effective in January 2001, that granted an unexpected increase in pension wealth to those with 30 or more years of service. For each IV specification the pension variables are instrumented using the unexpected, reform-induced shock to that variable (Brown and Laschever, 2012). For

[11] The estimate is a back-of-the-envelope calculation using the average premiums in the *2006 Health Benefits Survey of Employers* (CalSTRS, 2006b).

[12] As demonstrated by Burtless and Hausman (1978), Hausman (1985), and Moffitt (1990), realized labor supply and the financial return to work are simultaneously determined on a kinked budget constraint. Individuals that are observed retiring before the kink, at a relatively high net wage, must have a low taste for work, while those that are observed retiring after the kink, when the net wage is relatively low, have a high taste for work. The unobserved taste parameter is then negatively correlated with the net wage and will bias the elasticity estimate.

[13] This is a product of the same utility function as was assumed to estimate the results in column (5) of Table 3, $U = C - \frac{S^{1+1/e}}{1+1/e} \times \alpha$, where $\alpha = exp(-\eta)$ and $e \times \eta = \varepsilon$.

Table 4
The probability of continued work.

All LAUSD teachers age 55–75	Linear probability model; dependent variable — work an additional year					
	(1)	(2)	(3)	(4)	(5)	(6)
Specification	OLS	2SLS	OLS	2SLS	2SLS	2SLS
Panel A.						
Pension wealth	−0.021***	−0.031***	−0.026***	−0.030***	−0.025***	−0.020***
	(0.004)	(0.006)	(0.003)	(0.006)	(0.006)	(0.006)
Peak value	0.042***	0.009				
	(0.007)	(0.015)				
Net wage			0.102***	0.043***	0.034	0.037*
			(0.011)	(0.021)	(0.022)	(0.021)
2nd-year pension wealth accrual					0.082***	0.083***
					(0.027)	(0.026)
3rd-year pension wealth accrual						0.028
						(0.028)
4th-year pension wealth accrual						0.018
						(0.042)
5th-year pension wealth accrual						0.024
						(0.035)
R^2	0.072	0.071	0.074	0.073	0.075	0.075
Sample size	21,529	21,529	21,529	21,529	21,529	21,529
Panel B. First stage for select pension financial incentives						
Unexpected change in pension wealth		1.756***		1.392***	1.436***	1.635***
		(0.016)		(0.015)	(0.015)	(0.018)
Unexpected change in peak value		0.765***				
		(0.017)				
Unexpected change in net wage				0.973***	0.964***	1.050***
				(0.017)	(0.016)	(0.016)

Notes: The sample of the LAUSD administrative data includes all teachers that are age 55–75 in each academic year (Sept.–Aug.) 1997–2000. All pension financial measures and their instruments are measured in one-hundred thousand dollars. In addition to the variables shown, all specifications include integer age indicators, years of service, salary and year fixed effects. The standard errors, clustered at the individual level, are in parentheses. The F-statistic, adjusted for multiple endogenous variables using Angrist and Pischke (2008), is greater than 1000 for all endogenous variables in all specifications and the p-value for each test is 0.000. In each specification all pension-related financial variables (pension wealth, peak value, net wage, and accruals) are instrumented with the corresponding set of financial shocks. Select first-stage coefficients of the effect of the pension financial shock on the corresponding financial measure are in columns (2) and (4)–(6) of panel B. The first-stage coefficient of the effect of the unexpected change in the second-year pension wealth accrual on second-year pension wealth accrual is 0.763*** (0.016) in column (5) and 0.838*** (0.015) in column (6). In column (6), the first stage coefficients of the effect of the unexpected change in the third-, fourth-, and fifth-year pension wealth accrual on the corresponding pension wealth accrual are 0.941*** (0.026), 0.597*** (0.019), and 0.718*** (0.015) respectively. Additional first-stage results are available in Appendix C. The probit estimates (for the non-IV specifications) are similar.
***Significant at the 1% level, **significant at the 5% level, and *significant at the 10% level.

example, the instrument for pension wealth is calculated as the pension wealth of individual i at time t calculated under the postreform benefit formula minus the pension wealth for individual i at time t calculated under the prereform benefit formula: $PWIV_{i,t} = PW_{i,t}^{post} - PW_{i,t}^{pre}$. Individual characteristics that enter the benefit calculation are held constant, so that the IV is capturing only the change due to the reform of the pension formula. The instruments take a non-zero value in the reform years (academic years 1998 and 2000) only. This instrument is unlikely to be correlated with individuals' unobserved tastes for work or other retirement-related factors (after controlling for age and service) because it is determined only by an unexpected change in pension rules that applied to teachers in all schools across the state.

The results are presented in Table 4. All specifications include age and year dummies, years of service, and salary (as part of net wage in columns (3)–(6)). In all specifications, as expected, a larger pension wealth is negatively correlated with continued work, while a greater return to working in future years is positively correlated with working an additional year. The IV specifications suggest this is a causal relationship. The magnitudes of the OLS estimates in the peak value specification (column (1)) are in line with studies using this estimation method to examine the effect of international Social Security programs on retirement (see Gruber and Wise, 2004). However, the coefficient on peak value is significantly reduced once instrumented.[14] The remaining results in Table 4 suggest that the financial return to work is important in the decision to work an additional year, but that the return in the next two years is most important. A back-of-the-envelope calculation of the elasticity of the probability of retirement with respect to the net wage is 0.29, which falls within Manoli and Weber's (2011) estimates using Austrian data. That this labor supply decision is relatively frictionless implies that the associated elasticity estimate may well-represent the long-run response (Chetty et al., 2011), though it is larger than the long-run kink estimates of the elasticity of lifetime labor supply reported in Section 5.

7. Discussion and conclusion

I use the nonlinearities of the California teachers' pension coupled with a pension reform to estimate the elasticity of lifetime labor supply with respect to the financial return to work. The results suggest that retirement-eligible individuals will work less than an additional month in the short-run and less than an additional half year in the long-run in response to a 10% increase in the financial return to work. Due to the large size of the budget constraint kink for California teachers it is possible that the lifetime labor supply elasticity estimate includes both income and substitution effects. However focusing on behavior at the kink minimizes the contamination from income effects, while the discontinuity estimate further supports a small substitution elasticity. The small elasticity estimates imply that defined benefit retirement programs do not greatly distort retirement timing and that the deadweight burden of the programs is smaller than suspected.

This paper also highlights that identification and estimation strategies based on policy reforms and reduced-form nonlinear budget constraint methods are complementary. Attention to changes in the nonlinearities of a pricing schedule provides more precise predictions for individual

[14] I am not aware of another study that has provided a reform-based instrumental variable estimation with this specification of the pension financial incentives. However in a similar spirit, Liebman et al. (2009) estimate the labor supply response to the effective Social Security tax rate by using only the discontinuities in tax rates that are the product of the Social Security benefit rules.

behavior following a reform, while reforms provide more information regarding individual behavior in the absence of price nonlinearities. In the California teacher context, the precise predictions for changes in the retirement distribution also provided an opportunity to verify that the pension financial incentives were salient and well-understood. This is beneficial for addressing concerns about financial literacy, tax salience, or norms, that may confound results in such contexts.

The findings of this research have timely and direct policy implications. Many teachers' retirement programs are facing large expected shortfalls. The management of these deficits will not only affect teachers' retirement security but also education security. Understanding how teachers will respond to the inevitable reforms of their retirement programs is essential to safeguard the financial and human resources needed to produce a high quality education environment. Further, although California teachers are a select group, the features of their pension are similar to those of Social Security, a program that covers a broader population. Social Security serves as the only defined benefit plan for a growing number of workers in the United States and it is both important and salient to this population as coverage is retained when individuals change jobs. Also like the California teachers' pension, Social Security has key program ages at which covered workers exhibit a propensity to retire. Given these strong parallels, the finding that California teachers' retirement behavior was little affected by a large pension reform raises concerns about how much of an impact the recent increase in the Social Security "full retirement age" will have on the labor supply of older workers.

Acknowledgments

I thank Emmanuel Saez for his generous advice and support. I also thank Alan Auerbach, David Card, Tom Davidoff, and Stefano DellaVigna for many helpful comments and guidance. This paper also benefited from comments by Courtney Coile, Elizabeth Weber Handwerker, Kostis Hatzitaskos, Jane Leber Herr, Vivian Hwa, Salar Jahedi, Ron Laschever, three anonymous referees, and seminar participants at University of California, Berkeley, University of Illinois at Urbana-Champaign, University of Missouri, RAND, Public Policy Institute of California, Urban Institute and at the 2010 ASSA and Joint Statistical meetings. I also thank Ed Derman and Rhonda Webb at CalSTRS and Steven Cantrell and Marianne Bellacommo at LAUSD for their assistance in securing the data for this project. Financial support from the UC Labor and Employment Research Fund and from the Center for Retirement Research at Boston College is gratefully acknowledged. All errors are my own.

Appendix A. Calculation of financial variables

Pension wealth for an individual retiring in year t is based on the service years, age, and salary of the teacher in year t and is calculated according to the formula

$$PW_{i,t} = \sum_{a=t}^{T} \pi_{a|t} \left(\frac{1}{1+r}\right)^{a-t} B_a\left(k_t, S_t, w^f_{S,t}\right),$$

where $\pi_{a|t}$ is the probability of living to each future year a given having lived to year t, and was computed from the CDC life tables by age and sex,[15] T is the maximum possible age that can be attained and is assumed to be 100, r is the real interest rate and is assumed to be 0.03. Pension benefits also increase by 2% of the initial benefit in each retirement year and this is incorporated into the calculation. Salary is not available for the systemwide data set so $w^f_{S,t}$ is set to one and pension wealth is a multiple of annual salary. Current-year salary is available in the LAUSD data and is used as the final compensation, $w^f_{S,t}$ in the calculations. Salary is assumed to grow by 2% per year, but this is only relevant for calculating the future accruals and the peak value in Section 6.2. The service used for S is years of CalSTRS-covered service in the systemwide data set and years of teaching in LAUSD for the individual-level data set. This measure will understate a teacher's true pension-relevant service if he or she has worked in another district in California. For the kink elasticity estimation pension wealth is calculated at each quarter-age and for the discontinuity estimation pension wealth is calculated at each half-year of service, both corresponding to finest bins available in the count data. For the estimations of probability of working an additional year in Section 6.2, pension wealth is calculated as of the last day of the academic year. In this case, all financial variables are adjusted to year 2000 dollars.

The annual net wage is calculated as the sum of salary net of pension contributions plus pension wealth accrual. It is $w^{net} = .92w_t + (PW_{t+\delta} - PW_t)/\delta$ where δ is 0.25 for the kink-age estimates, 0.5 for the discontinuity-service estimates, and 1 for the probability estimates. The remaining pension variables used in Section 6.2 is calculated as follows. Pension wealth accrual in any year t is $(PW_{t+1} - PW_t)$ and peak value is $(PW_{t^{Max}} - PW_t)$, where t^{Max} is the future retirement date with the highest expected pension wealth.

Unfortunately, due to the limited nature of the administrative data, it is not possible to include taxes in the calculation of the return work.

[15] United States Life Tables, 2000. National Vital Statistics Report Volume 51, No. 3.

Appendix B. Census tabulations

Table B.1
Comparison of California teachers to the California population.

	California teachers	California workers	All Californians
Wage income			
$0	0.09	0.23	0.43
$1–50 k	0.54	0.55	0.41
$50 k–100 k	0.36	0.17	0.13
$100 k–150 k	0.00	0.03	0.02
Greater than $150 k	0.00	0.02	0.02
Homeownership rates	0.87	0.76	0.74
House value			
Less than $500 k	0.76	0.65	0.64
Greater than $500 k	0.25	0.35	0.36
Difficulty working	0.08	0.13	0.15
Fraction female	0.72	0.47	0.52
Marital status			
Married	0.69	0.67	0.60
Separated	0.02	0.03	0.03
Divorced	0.18	0.18	0.19
Widowed	0.05	0.05	0.12
Never married/single	0.07	0.07	0.06
Race			
White	0.85	0.72	0.70
Black	0.06	0.06	0.06
American Indian or Alaska native	0.00	0.01	0.01
Chinese	0.02	0.03	0.03
Japanese	0.02	0.01	0.01
Other Asian or Pacific Islander	0.02	0.07	0.07
Other or multiple	0.04	0.10	0.11
Weighted observations	151,306	4,094,681	5,516,039
Actual observations	7783	210,899	284,422

Notes: Sample restricted to those age 50–69 (Ruggles et al., 2004).

A greater fraction of California teachers are women than the general working population. California teachers are also more likely to own a house and are less likely to report health-related difficulty working. These last two factors are likely important in retirement decisions. However, they will work in opposite directions on retirement age, with real estate wealth enabling an earlier retirement and good health enabling a longer career.

Appendix C. First stage for the probability of continued work

Table C.1
First stage estimates for Table 4.

All LAUSD teachers ages 55–75; first stage of linear probability model							
Corresponding specification in Table 4	Column 2		Column 4		Column 5		
Variable being instrumented	Pension wealth	Peak value	Pension wealth	Net wage	Pension wealth	Net wage	2nd-year accrual
Unexpected change in pension wealth	1.756***	−0.002	1.393***	−0.020***	1.436***	−0.014***	−0.008***
	(0.016)	(0.007)	(0.015)	(0.003)	(0.015)	(0.003)	(0.003)
Unexpected change in peak value	0.861***	0.765***					
	(0.026)	(0.017)					
Unexpected change in net wage			0.193***	0.973***	0.123***	0.964***	−0.016**
			(0.042)	(0.017)	(0.042)	(0.016)	(0.008)
Unexpected change in 2nd-year accrual					0.461***	0.064***	0.763***
					(0.044)	(0.012)	(0.016)
Sample size	21,529	21,529	21,529	21,529	21,529	21,529	21,529

Notes: Primary column titles indicate the location of the corresponding second stage results in Table 4 of Section 6.2. For a given specification, each column shows the variable being instrumented. The coefficients are for the instrument indicated by the row label. The standard errors, clustered at the individual level, are in parentheses.
***Significant at the 1% level, **significant at the 5% level, and *significant at the 10% level.

Table C.2
First stage estimates for Table 4 continued.

All LAUSD teachers ages 55–75; first stage of linear probability model						
Corresponding specification in Table 4	Column 6					
Variable being instrumented	Pension wealth	Net wage	2nd-year accrual	3rd-year accrual	4th-year accrual	5th-year accrual
Unexpected change in pension wealth	1.635***	0.070***	0.053***	0.054***	0.002	−0.004
	(0.018)	(0.004)	(0.003)	(0.004)	(0.003)	(0.003)
Unexpected change in net wage	0.320***	1.050***	0.043***	0.099***	0.057***	0.045***
	(0.044)	(0.016)	(0.007)	(0.008)	(0.007)	(0.006)
Unexpected change in 2nd-year accrual	0.680***	0.094***	0.838***	−0.036***	0.017**	0.040***
	(0.048)	(0.011)	(0.015)	(0.009)	(0.008)	(0.007)
Unexpected change in 3rd-year accrual	0.188***	0.272***	0.031***	0.941***	0.002	0.036***
	(0.059)	(0.013)	(0.010)	(0.026)	(0.011)	(0.009)
Unexpected change in 4th-year accrual	0.599***	0.123***	0.200***	0.011	0.597***	−0.085***
	(0.062)	(0.012)	(0.011)	(0.012)	(0.019)	(0.010)
Unexpected change in 5th-year accrual	0.654***	0.259***	0.200***	0.316***	0.119***	0.718***
	(0.056)	(0.011)	(0.010)	(0.011)	(0.019)	(0.015)
Sample size	21,529	21,529	21,529	21,529	21,529	21,529

***Significant at the 1% level, **significant at the 5% level, and *significant at the 10% level.

References

Angrist, Joshua D., Pischke, Jorn-Steffen, 2008. Mostly Harmless Econometrics: An Empiricist's Companion. Princeton University Press, Princeton.

Asch, Beth, Haider, Steven J., Zissimopoulos, Julie, 2005. Financial incentives and retirement: evidence from federal civil service workers. Journal of Public Economics 89 (2–3), 427–440.

Blau, David M., Gilleskie, Donna B., 2001. Retiree health insurance and the labor force behavior of older men in the 1990s. The Review of Economics and Statistics 83 (1), 64–80.

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

Burtless, Gary, 1986. Social Security, unanticipated benefit increases, and the timing of retirement. The Review of Economic Studies 53 (5), 781–805.

Burtless, Gary, Hausman, Jerry A., 1978. The effect of taxation on labor supply: evaluating the Gary negative income tax experiment. Journal of Political Economy 86 (6), 1103–1130.

California State Teachers' Retirement System, 1998. Bulletin, vol. 13(1). California State Teachers' Retirement System, Sacramento, CA.

California State Teachers' Retirement System, 2006a. CalSTRS Comprehensive Annual Financial Report of the California State Teachers' Retirement System for the Fiscal Year Ended June 30, 2006. California State Teachers' Retirement System, Sacramento, CA.

California State Teachers' Retirement System, 2006b. 2006 Health Benefits Survey of Employers. California State Teachers' Retirement System, Sacramento, CA.

Chetty, Raj, Friedman, John N., Olsen, Tore, Pistaferri, Luigi, 2011. Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: evidence from Danish tax records. Quarterly Journal of Economics 126 (2), 749–804.

Coile, Courtney C., Gruber, Jonathan, 2007. Future Social Security entitlements and the retirement decision. The Review of Economics and Statistics 89 (2), 234–246.

Costrell, Robert M., McGee, Joshua, 2010. Teacher pension incentives, retirement behavior, and potential for reform in Arkansas. Education Finance and Policy 5 (4), 492–518.

Costrell, Robert M., Podgursky, Michael, 2009. Peaks, cliffs, and valleys: the peculiar incentives in teacher retirement systems and their consequences for school staffing. Education Finance and Policy 4 (2), 175–211.

Currie, Janet, Madrian, Brigitte C., 1999. In: Ashenfelter, Orley C., Card, David (Eds.), Health, health insurance, and the labor market. : Handbook of Labor Economics, vol. 3C. Elsevier, Amsterdam, pp. 3309–3415.

Fong, Anthony B., Makkonen, Reino, 2012. Retirement patterns of California prekindergarten–grade 12 educators. Issues and Answers Report, REL 2012-130. U.S. Department of Education, Institute of Education Sciences, National Center for Education Evaluation and Regional Assistance, Regional Education Laboratory West, Washington, D.C.

French, Eric, Jones, John Bailey, 2011. The effects of health insurance and self-insurance retirement behavior. Econometrica 79 (3), 693–732.

Friedberg, Leora, 2000. The labor supply effects of the social security earnings test. The Review of Economics and Statistics 82 (1), 48–63.

Furgeson, Joshua, Strauss, Robert P., Vogt, William B., 2006. The effects of defined benefit pension incentives and working conditions on teacher retirement decisions. Education Finance and Policy 1 (3), 316–348.

Gruber, Jonathan, Madrian, Brigitte C., 1995. Health-insurance availability and the retirement decision. American Economic Review 85 (4), 938–948.

Gruber, Jonathan, Wise, David A. (Eds.), 2004. Social Security Programs and Retirement around the World: Micro Estimation, vol. 2. University of Chicago Press, Chicago.

Gustman, Alan L., Steinmeier, Thomas L., 1986. A structural retirement model. Econometrica 54 (3), 555–584.

Gustman, Alan L., Steinmeier, Thomas L., 1994. Employer provided health insurance and retirement behavior. Industrial & Labor Relations Review 48 (1), 124–140.

Hausman, Jerry A., 1985. The econometrics of nonlinear budget sets. Econometrica 53 (6), 1255–1282.

Johnson, Richard W., Davidoff, Amy J., Perese, Kevin, 2003. Health insurance costs and early retirement decisions. Industrial & Labor Relations Review 56 (4), 716–729.

Kleven, Henrik Jacobsen, Waseem, Mazhar, 2012. Behavioral responses to notches: evidence from Pakistani tax records. Unpublished.

Krueger, Alan B., Pischke, Jorn-Steffen, 1992. The effects of social security on labor supply: a cohort analysis of the notch generation. Journal of Labor Economics 10 (4), 412–437.
Liebman, Jeffrey B., Luttmer, Erzo F.P., Seif, David G., 2009. Labor supply responses to marginal social security benefits: evidence from discontinuities. Journal of Public Economics 93 (11–12), 1119–1284.
Lumsdaine, Robin L., Stock, James H., Wise, David A., 1992. Three models of retirement: computational complexity versus predictive validity. NBER Working Paper 3558.
Manoli, Day, Weber, Andrea, 2011. Nonparametric evidence on the effects of financial incentives on retirement decisions. NBER Working Paper 17320.
Mastrobuoni, Giovanni, 2009. Labor supply effects of the recent social security benefit cuts: empirical estimates using cohort discontinuities. Journal of Public Economics 93 (11–12), 1224–1233.
Moffitt, Robert, 1990. The econometrics of kinked budget constraints. Journal of Economic Perspectives 4 (2), 119–139.
Pei, Zhuan. 2012. Dynamic opting-in incentives in income-tested social programs: evidence from Medicaid/CHIP. Unpublished.
Pencavel, John, 2001. The response of employees to severance incentives: the University of California's Faculty, 1991–94. The Journal of Human Resources 36 (1), 58–84.
Ramnath, Shanthi P. 2011. Taxpayers' response to notches: evidence from the saver's credit. Unpublished.
Reed, Deborah, Rueben, Kim S., Barbour, Elisa, 2006. Retention of New Teachers in California. Public Policy Institute of California, San Francisco.
Ruggles, Steven, Sobek, Matthew, Alexander, Trent, Fitch, Catherine A., Goeken, Ronald, Kelly Hall, Patricia, King, Miriam, Ronnander, Chad, 2004. Integrated Public Use Microdata Series: Version 3.0 [Machine-readable Database]. Minnesota Population Center [producer and distributor], Minneapolis, MN (http://usa.ipums.org/usa/).
Rust, John, Phelan, Christopher, 1997. How Social Security and Medicare affect retirement behavior in a world of incomplete markets. Econometrica 65 (4), 781–831.
Saez, Emmanuel, 2010. Do taxpayers bunch at kink points. American Economic Journal: Economic Policy 2 (3), 180–212.
Sallee, James M., Slemrod, Joel, 2012. Car notches: strategic automaker responses to fuel economy policy. Journal of Public Economics. http://dx.doi.org/10.1016/j.jpubeco.2012.06.005.
Samwick, Andrew A., 1998. New evidence on pensions, social security, and the timing of retirement. Journal of Public Economics 70 (2), 207–236.
Social Security Administration, 2006. Social Security Administration Performance and Accountability Report, Fiscal Year 2006. Social Security Administration, Washington, DC.
Stock, James H., Wise, David A., 1990. Pensions, the option value of work, and retirement. Econometrica 58 (5), 1151–1180.
Weber, Caroline E. 2012. Does the earned income tax credit reduce saving by low-income households?. Unpublished.