CROSS-SUBSIDIZATION OF TEACHER PENSION COSTS: THE CASE OF CALIFORNIA

Robert M. Costrell
(corresponding author)
Department of Education Reform
University of Arkansas
Fayetteville, AR 72704
costrell@uark.edu

Josh McGee
Department of Education Reform
University of Arkansas
Fayetteville, AR 72704
joshbmcgee@gmail.com

Abstract

The value of pension benefits varies widely, by a teacher's age of entry and exit. This variation is masked by the uniform rate of annual contributions, as a percent of pay, to fund benefits for all. For the first time, we unmask that variation by calculating annual costs at the individual level. In California, we find that the value of a teacher's benefits ranges from about 4 to 22 percent of pay, and exhibits some idiosyncratic patterns, as is endemic to traditional pension plans. The variation in individual cost rates generates an extensive but hidden array of cross-subsidies, as winners receive benefits worth more than the uniform contribution rate, and losers receive less. Almost two thirds of all entering teachers, past and present, are losers in California. By contrast, a prominently invoked study finds that nearly all active teachers are winners there. That result is shown to be highly skewed by excluding the losses of prior entrants who left early, thereby violating the funding fact that the gains and losses of winners and losers must offset each other. Our main policy conclusion is that cash balance plans can rationalize or eliminate the current system of cross-subsidies and provide the transparency lacking in traditional plans.

https://doi.org/10.1162/edfp_a_00253

1. INTRODUCTION AND SUMMARY

The funding plans for traditional teacher pension systems are built upon a highly uneven set of benefits, varying widely in value by age of entry and exit. These inequities are masked from individual teachers by a uniform fringe benefit rate for pensions. For example, the annual contribution to the pension fund (employer and employee contributions taken together) may be 15 percent of each teacher's salary. These contributions are designed to fund the future retirement benefits as they are earned,[1] for the system as a whole. However, the annual cost of benefits for individual teachers may deviate widely from this overall average—for example, early leavers may earn benefits worth 5 percent of salary per year whereas the benefits of those who stay until retirement are worth 25 percent. In effect, there is a large cross-subsidy from the former to the latter; this is a big part of how the plan funds benefits for career teachers. As we will show, there are also other patterns of cross-subsidies—for example, from younger to older entrants. In this paper, we present the methodology for calculating the annual cost (as a percent of pay) for individual benefits and the consequent array of cross-subsidies. We illustrate with the California State Teachers Retirement System (CalSTRS). Our goal is greater transparency and deeper understanding of the system of winners and losers embedded in the funding plans of traditional teacher pension systems. Without these cross-subsidies, the employer and/or employee would have to contribute much more to fund the benefits of career teachers. The existence of these cross-subsidies is well known to plan administrators, but not to typical participants.

Previous literature has established that teacher pension plans generate winners and losers (Costrell and McGee 2010; Costrell and Podgursky 2010a, b; McGee and Winters 2013; Aldeman and Rotherham 2014; Johnson and Southgate 2015; Lueken 2017). The issue has now reached the front page of the *New York Times* (Walsh 2017). This literature measures the value of individual benefits by "pension wealth," that is, the present value (PV) of the stream of pension payments. For winners, pension wealth exceeds the PV of the stream of contributions (unless specified otherwise herein, we refer to total contributions—employer plus employee), often by many thousands of dollars, and for losers the converse holds. This paper is the first to present a new metric for the value of individual benefits: the annual cost of funding them, as a percent of salary. This annual cost rate may be a more accessible metric than pension wealth (dollars of present value). In addition, the cost rate is firmly grounded in the standard method used to determine required funding as benefits accrue. Our contribution is to bring the individual cost rates out into the open. They are implicitly embedded within the actuarial calculation of the uniform cost rate, but they are not reported, so an individual teacher does not know the cost of his or her benefit. Without such information, she may well, understandably, perceive that the uniform contribution funds her pension benefit, even though, for many teachers, a good part of that contribution is funding other teachers'

1. In addition, the employer makes payments for the unfunded liability—benefits earned in the past, but not funded. This is a large problem, but is not the subject of this paper. The intergenerational cross-subsidies represented by these payments (Backes et. al. 2016) are a consequence of the failure to meet actuarial assumptions, particularly the return on investments (Costrell 2018a, b). In this paper, we take the assumptions as given to analyze the cross-subsidies within generations that are built into the system's funding plan, as distinct from the cross-subsidies between generations when the assumptions fail. In subsequent work, we incorporate cross-subsidies across generations that arise from the failure to meet assumed investment returns (Costrell and McGee 2017).

benefits.[2] Our purpose here is to make explicit the redistribution of contributions that pension systems build into their funding plan.

We find the cross-subsidies embedded in the CalSTRS plan to be widespread, substantial, and somewhat arbitrary. Individual annual cost rates vary from about 4 percent of pay (for those who enter and leave early) to 22 percent (for those who enter late and leave at the plan's normal retirement age). Almost two thirds of entrants are losers: Their benefits are worth less than their contributions and those seemingly made on their behalf. The cross-subsidies they provide to the winners are not small—they represent the vast majority of the employer contribution. Indeed, for most entrants, the entire employer contribution and some of the proceeds of their own contribution goes to fund others' benefits. As a result, the winners receive benefits of significantly greater value than the contributions by or for them.

A secondary purpose of this paper is to clarify that winners and losers under a particular pension plan can only be accurately described by looking at past and present entrants, not simply a snapshot (or cross section) of the currently active workforce. Such a snapshot excludes many losers who have already left, so it is skewed by survivorship bias to the winners. We offer this clarification because the picture of winners and losers previously established (and more fully explored here) has been challenged in a prominently invoked analysis of CalSTRS that finds nearly all active teachers are winners (Rhee and Fornia 2016). The main feature of their critique is a shift in focus from winners and losers among entrants to active teachers.[3] We show how dramatically the distribution of CalSTRS teachers differs between actives and entrants, past and present. Early leavers are almost completely absent from the distribution of actives. But without the losses of those early leavers, the actives' contributions simply do not add up to fully fund the winners' benefits. For a complete picture, we show the pattern of cross-subsidies between and among actives and prior leavers, and how they add up to zero, as they must.

The plan of the paper follows: We begin with a formal explanation of cost rates, both individual and uniform, and the associated cross-subsidies built into the funding plan. We show exactly how the cross-subsidies must add up to zero. We then illustrate with plan data from CalSTRS, followed by an analysis of the winners and losers, and an explanation of the very different results in Rhee and Fornia. Concluding remarks follow, including policy implications. We suggest that cash balance plans can rationalize or eliminate the current system of cross-subsidies and provide the transparency lacking in traditional plans.

2. NORMAL COST RATES AND CROSS-SUBSIDIZATION

Pension plans calculate a contribution rate that will fund future benefits as they are earned. This is called the *normal cost rate*. It is calculated at the aggregate level, but embedded within the calculation are the individual cost rates, based on age of entry

2. CalSTRS's *Member Handbook* (2016b, p. 24) states that "employers' contributions go toward paying all members' retirement benefits." The ambiguous term "all" may be easily, if incorrectly, construed as referring to "each."
3. "Winners and losers" are defined by Rhee and Fornia relative to a cash balance plan described later in this paper. They also (rightly) emphasize the full distribution of entry ages (rather than, say, age 25), in a state with many older entrants.

and exit, as shown in the Appendix. We describe here the basic concepts and provide the basic math.

Individual Normal Cost Rates

Consider an individual of type (e, s), where e is the age of entry and s (for separation) is the age of exit. For each type (e, s), we identify an individual normal cost rate, n_{es}, as a percent of salary. We calculate this rate to generate a stream of contributions sufficient to fund the individual's future benefits. That is, the PV of contributions must equal the PV of benefits.

Formally, for an individual of type (e, s), we must have $n_{es}W_{es} = B_{es}$, where W_{es} is the PV of earnings (so $n_{es}W_{es}$ is the PV of contributions, for a constant rate n_{es}) and B_{es} is the PV of benefits (both evaluated at entry). It then immediately follows that the individual cost rate is the ratio of the PV of benefits to that of earnings:

$$n_{es} = B_{es}/W_{es}. \tag{1}$$

This is the rate that, applied to the individual's annual earnings over her career, would prefund her benefits. It represents the value of her benefits earned annually, as a percent of earnings—an individual fringe benefit rate for pensions.

If we compare individuals with different entry and exit ages, (e, s), we find their cost rates, n_{es}, vary widely. In general, for any given e, n_{es} rises with s, from the point of vesting up through a peak value retirement age. This is a manifestation of the well-known back-loading of benefits that favors long-termers under traditional pension formulas based on final average salary (FAS) (Costrell and Podgursky 2010a). The variation in n_{es} with e, for any given s, is less obvious. Later entrants with the same exit age have shorter service, so their pension and its PV, B_{es}, is lower, but so is that of their earnings, W_{es}. For CalSTRS, as discussed subsequently, the latter outweighs the former, so the normal cost rate, B_{es}/W_{es}, is generally higher for later entrants.

Uniform Normal Cost Rate for an Entering Cohort

Consider an entering cohort, whose entrants vary by age at entry, e, and projected age of exit, s. Denote the joint frequency of e and s, among entrants, as p_{es}. For example, entrants of age 25 years who are projected to stay to age 65 years constitute a fraction $p_{25,65}$ of the entering cohort. Let us now consider a uniform normal cost rate, call it n^*, applied to all members of the cohort (of varying entry ages) throughout their careers (of varying length). It can readily be shown that the uniform cost rate required to fund the cohort's projected benefits is:

$$n^* = \sum_e \sum_s n_{es}\,(p_{es}W_{es}) \Big/ \Big(\sum_e \sum_s p_{es}W_{es} \Big). \tag{2}$$

This is the ratio of the PV of the cohort's benefits[4] to the PV of the cohort's earnings; the same relationship we saw for the individual normal cost rate holds for the cohort as a whole. This expression also shows, importantly, that n^* is a weighted average of individual normal cost rates n_{es} across ages of entry and exit. The weights for

4. Substituting $n_{es} = B_{es}/W_{es}$ into the numerator gives $\Sigma_e \Sigma_s p_{es} B_{es}$.

n_{es} are $(p_{es}W_{es})/(\Sigma_e\Sigma_s p_{es}W_{es})$, representing the share of type (e, s) in the cohort's PV of earnings.[5]

Cross-Subsidies, Weights, and Adding Up

The deviations $(n_{es} - n^*)$ are positive and negative, corresponding to whether the cost of funding any individual's benefits exceeds or falls short of the cohort's uniform contribution rate, n^*. They constitute cross-subsidies. Moreover, by the nature of averages, these cross-subsidies must add up to zero, when properly weighted, by shares of the cohort's PV of earnings:

$$\sum_e \sum_s (n_{es} - n^*)\,(p_{es}W_{es}) \Big/ \left(\sum_e \sum_s p_{es}W_{es}\right) = 0. \tag{3}$$

Consider a simple two-group example—early leavers and career teachers, accounting for 40 percent and 60 percent, respectively, of the cohort's lifetime earnings. Suppose the individual normal cost rate for early leavers is 9 percent and that of career teachers is 19 percent, such that n^* is 15 percent. Thus, early leavers provide a cross-subsidy of −6 percent of their earnings, career teachers receive a cross-subsidy of +4 percent, and their weighted sum is zero.

We can also measure the cross-subsidies in present values, $(n_{es} - n^*)W_{es} = (B_{es} - n^*W_{es})$, that is, the difference between the PV of benefits for an individual of type (e, s) and the contributions by or for that individual. Equation 3 implies

$$0 = \sum_e \sum_s (n_{es} - n^*)\,(p_{es}W_{es}) = \sum_e \sum_s (B_{es} - n^*W_{es})\,p_{es}. \tag{4}$$

Here, the weights for adding up are simply the cohort's shares of entrants of type (e, s), p_{es}. Thus, to continue our simple example, suppose two thirds of the cohort leaves early and one third stays for a full career. Let the PV of earnings be \$400,000 for early leavers and \$1.2 million for career teachers (these numbers tie to the shares of earnings given earlier). Thus, early leavers receive benefits of 0.09 × \$400,000 = \$36,000 in PV, but the contributions by or for them are 0.15 × \$400,000 = \$60,000. For career teachers, the PV of benefits is 0.19 × \$1,200,000 = \$228,000, but contributions are 0.15 × \$1,200,000 = \$180,000. The cross-subsidies are, respectively, −\$24,000 and +\$48,000, and their sum is zero when weighted by shares of the cohort (two thirds and one third).

To summarize, the zero-sum result holds in two forms: (1) for cross-subsidies in PV terms, applying weights equal to the shares of the entering cohort; and (2) for cross-subsidies in normal cost rates, applying weights equal to the shares of the entering cohort's lifetime earnings. We identify both sets of weights because previous literature (including Rhee and Fornia, discussed later) measures gains and losses in PV terms, whereas the present paper measures them in normal cost rates. Both sets of weights are pertinent for constructing counts of winners and losers. A typical result will be that x percent of the cohort are losers, but they account for a smaller portion of the cohort's earnings because they tend to be early leavers.

5. These are not the exact weights used in actuarial practice, but are consistent with the approach (see the Appendix).

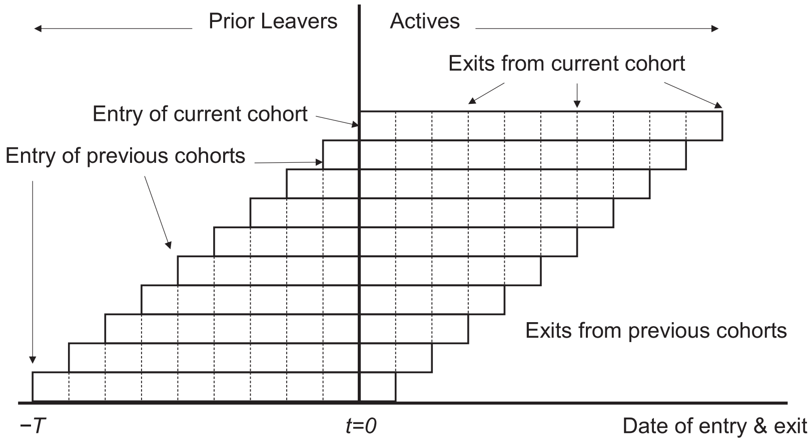

Figure 1. Layered Cohorts.

Cost Rates Across Cohorts

Thus far, we have considered individual and uniform cost rates for members of a given entering cohort. These cost rates (equations 1 and 2), and the zero-sum results (equations 3 and 4), pertain not only to the current cohort, but to any entering cohort, past or present. Figure 1 illustrates how the series of cohorts are layered among the current active workforce. The top bar represents the current cohort, entering at $t = 0$, with exits distributed among various dates thereafter. The bars below represent past cohorts, having entered at various dates $t < 0$. Among each of these cohorts, some remain active (those whose exits lie to the right of the vertical line at $t = 0$), and others (to the left of the line) have already exited. This partition will be formalized when we turn to the winners and losers among "actives vs. entrants," in a subsequent section.

In the meantime, we consider the unpartitioned cohorts. We may consider them singly, as we have thus far, or all cohorts taken together, past and present, currently working their way over time through the workforce. This aggregate is the set of layered cohorts depicted in figure 1, top to bottom and both sides of the vertical line. For any given benefit formula and set of actuarial assumptions, the cost rates n_{es} and n^* generalize across cohorts, independent of entry date. To see this, note first that even as entry wages vary across cohorts, the individual cost rates, $n_{es} = B_{es}/W_{es}$, do not vary, as both the numerator and denominator increase proportionally with the entering wage.[6] Hence, n_{es} is constant across cohorts. Under the additional assumption that the weights in equation 2, $(p_{es} W_{es})/(\Sigma_e \Sigma_s p_{es} W_{es})$, are constant across cohorts,[7] so is n^*. That is, we can consider n_{es} and n^* as referring to each or all cohorts, past and present, represented in the current workforce under the current benefit formula and assumptions.[8]

6. For given actuarial assumptions on the annual growth of wages, each year's wage over an individual's career increases proportionally with the entry wage. This means that the PV of earnings (the denominator of n_{es}) increases proportionally, and so does the FAS, which is proportional to the PV of benefits (the numerator).
7. This will be true if the relative frequencies and wages, by age of entry, are constant. If so, then the plan assumptions on exit rates and wage growth rates will generate constant p_{es} and constant relative values of W_{es}.
8. Actuarial practice applies the same normal cost rate to all members (in the same benefit tier), regardless of when they entered. The cost attributable to retroactive changes in benefits and changes in assumptions are allocated after the fact to the unfunded liability.

3. INDIVIDUAL NORMAL COST RATES FOR CALSTRS

We estimate the individual normal cost rates for CalSTRS, $n_{es} = B_{es}/W_{es}$, for all entry and exit ages, $e, s = 20, \ldots, 75$. The data on which our calculations are based are the CalSTRS actuarial assumptions and benefit formula. The actuarial assumptions cover wage growth (merit salary increases by entry age and years of service, and general wage inflation), discount rate, exit rates for retirement (by age and years of service), exit rates prior to retirement (by years of service), probability of refund (by entry age and years of service), and mortality rates (for female actives and future retirees). These assumptions are provided in the annual valuation report (CalSTRS 2016a), supplemented with more granular data (annual rates vs. selected rates) from CalSTRS. We also use CalSTRS data on the distribution of entry ages. The benefit formula is delineated in the annual valuation report, as well as the member handbook (CalSTRS 2016b). This includes the retirement eligibility conditions, age-specific multipliers (described below), cost of living adjustments (COLA), employee contribution rate, and interest rate on refunds.

Because B_{es} and W_{es} are proportional to the entry wage, $n_{es} = B_{es}/W_{es}$ is independent of it, so we can normalize B_{es} and W_{es} per dollar of entry wage. For W_{es} we have:

$$W_{es} = \sum_{a=e}^{s} (1 + r)^{(e-a)} \prod_{\alpha=e}^{a} \left(1 + g^{w}_{e,(\alpha-e)}\right). \tag{5}$$

Here, $g^{w}_{e,(\alpha-e)}$ is the wage growth (merit plus general wage inflation) by entry age and service, as given by CalSTRS actuarial assumptions, and r is the CalSTRS discount rate of 7.5 percent.[9]

Benefits can be in the form of a pension or refund of employee contributions.[10] If a teacher takes the refund, she forgoes the pension (or possibility of future pension), and receives the cumulative value of the employee (but not employer) contributions, with accumulated interest at the rate set by CalSTRS. Teachers who leave before vesting, without the expectation of returning and becoming eligible for a pension, would certainly take the refund because it is the only benefit to which they are entitled. Teachers who leave after vesting, but too young to draw a pension, may either take the refund or leave the money in the fund to draw a pension in the future, upon reaching an eligible age. Finally, teachers who leave service and are eligible for an immediate pension may still choose the refund, although it is generally not financially prudent to do so. CalSTRS assigns probabilities of taking the refund in these various situations, by entry age e and service $(s - e)$, $(prob\ refund)_{e,s-e}$. If a teacher takes the refund, B_{es} is the present value of the cumulative employee contributions (9.205 percent of pay) with interest (4.5 percent in 2015[11]), discounted back to entry, $PV(Refund_{es})$, a straightforward calculation.

If a teacher takes the pension, B_{es} is the present value of the stream of pension payments, weighted by the survival probabilities, discounted to entry, $PV(Pension_{es})$. The payments begin with a starting pension determined by CalSTRS's formula, and augmented annually with a 2.0 percent simple cost of living adjustment. Specifically,

9. We use the actuarial assumptions that CalSTRS adopted through the 2015 valuation report, for comparison with Rhee and Fornia (2016). CalSTRS has since revised its assumptions. Among other changes, the discount rate has been cut to 7.25 percent for the 2016 valuation and 7.00 percent thereafter. For estimates of the individual normal cost rates using the revised assumptions, see Costrell (2018c).
10. We leave aside disability and death benefits, which constitute about 5 percent of normal cost, less than 1 point.
11. For the 2016 valuation report, interest on member accounts is reduced to 3.0 percent.

Table 1. Individual Normal Cost Rates (percent of pay), n_{es}, Selected Ages of Entry and Exit

Estimated Using 2015 CalSTRS Assumptions and Benefit Formula for New Hires									
		Age of Exit, *s*							
		30	35	40	45	50	55	60	65
Age of Entry, *e*	25	4.4%	4.9%	5.5%	6.6%	8.5%	11.2%	14.4%	16.4%
	30		5.1%	5.8%	7.0%	8.9%	11.9%	15.4%	17.6%
	35			6.0%	7.2%	9.2%	12.3%	16.2%	18.5%
	40				7.1%	9.3%	12.7%	18.1%	20.9%
	45					8.8%	12.1%	16.5%	21.7%

Notes: The individual normal cost rate is the percentage of salary to be collected, from entry to exit, that will fund that individual's benefits. Each entry represents the estimated contributions (employer plus employee) required to fund an individual entering at age *e* (row) and exiting at age *s* (column). For example, to cover benefits of an individual entering at age 25 years and retiring at age 65 years, annual contributions must be 16.4 percent of salary. CalSTRS = California State Teachers Retirement System.

the starting pension for an individual of type (*e*, *s*) is calculated under the "2% at 62" program for new hires (since 2013). This is an FAS formula (average of final three years) with age-specific multipliers ranging from 1.16 percent at age 55 years to 2.0 percent at age 62 years, and 2.4 percent at 65 years, after five-year vesting. Thus, a 25-year-old entrant who works to age 65 years will retire at $40 \times 2.4 = 96$ percent of FAS. Vested employees who withdraw before age 55 years but do not cash out are assumed by CalSTRS to defer the pension to age 60 years (when the multiplier is 1.76 percent). So a 25-year-old entrant exiting at age 50 years will receive at age 60 years a pension of $25 \times 1.76 = 44$ percent of her FAS (frozen ten years earlier).

Taking $PV(Pension_{es})$ weighted together with $PV(Refund_{es})$, we have:

$$B_{es} = (prob\ refund)_{e,s-e} PV(Refund_{es}) + [1 - (prob\ refund)_{e,s-e}] PV(Pension_{es}). \qquad (6)$$

Using W_{es} from equation 5, we have $n_{es} = B_{es}/W_{es}$. This give us the contribution rate required, over one's career, to fund the benefits of an individual entering at age *e* and exiting at age *s*.

Variation in Normal Cost Rates by Age of Entry and Exit

Table 1 depicts the normal cost rates, n_{es}, for selected ages of entry[12] and exit. The variation is wide, from 4.4 percent to 21.7 percent (the full range, for entry ages not shown, is 4.3 to 22.4 percent). The cost rates shown here rise with age of exit (looking along the rows). For example, annual contributions to fund benefits for a 25-year old entrant need only be 4.4 percent of pay if she exits upon vesting at age 30 years, but rise to 16.4 percent if she stays to 65 years. The cost rates also generally rise with age of entry (looking down the columns), for example, the contribution required to fund retirement at age 65 years rises to 21.7 percent for entry at age 45 years. Finally, if we look down and across any diagonal, where years of service is constant, we find a strong rise in the cost rate. For example, the benefits for a 25-year career cost 8.5 percent of pay for one

12. These are representative of the vast majority of entrants. The proportion between ages 23 and 47 years is 84 percent. The full distribution is shown in figure 3 of Rhee and Fornia (2016) and reproduced in figure A.1.

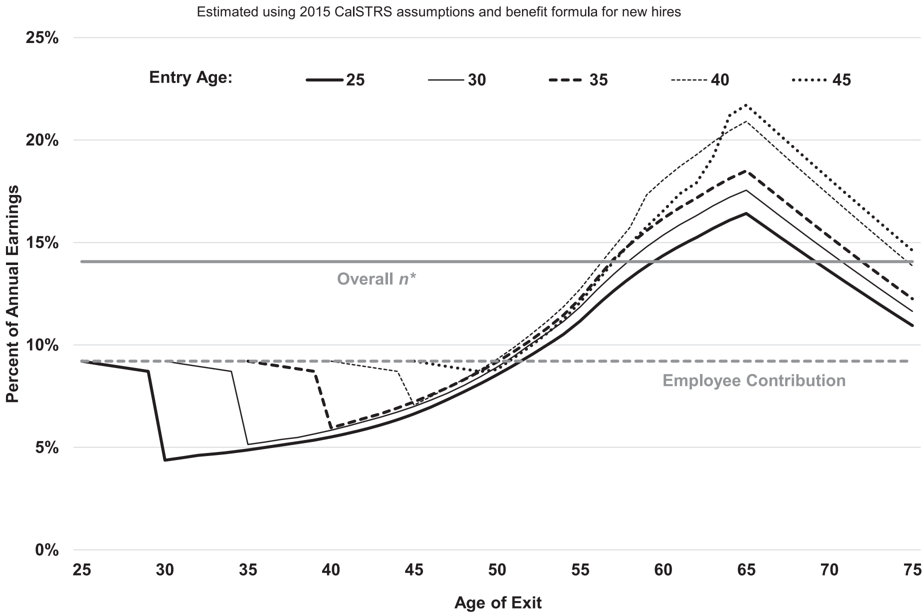

Note: The curves depict n_{es}, the annual contribution rate required to fund benefits of an individual entering at age *e* and exiting at age *s*. Variation in cost by age of exit is shown along each curve; variation by age of entry is shown across curves.

Figure 2. Normal Cost Rate, by Age of Entry and Exit, n_{es}

who enters at 25, but 20.9 percent for entry at 40. This is due to the plan's eligibility conditions, as the pension for the first individual is deferred to age 60 years, after exit at age 50 years, whereas the second individual's pension is immediate upon exit at age 65 years. Thus, the benefits awarded for the same years of service, but different entry age, are of widely varying value and cost.

For a more complete look, consider figure 2, which depicts the normal cost rates for the same selected entry ages, but for all exit ages *s*, including those before vesting and after age 65 years. The pattern for any given entry age (e.g., age 25 years) is depicted along each curve, as the exit age varies.

Prior to vesting, the benefit is the refund of employee contributions. The normal cost rate, therefore, starts at the employee contribution of 9.205 percent—each curve begins at the dashed gray horizontal line representing that rate. The cost rate then gently declines, falling a bit below the employee contribution rate. That is because the interest credit of 4.5 percent is below the fund's assumed return, 7.5 percent, with the fund retaining the difference. The plan's cost rate to cover the refund falls as this difference accumulates, but not by much over these five years.

Then, upon vesting at five years, the normal cost rate drops precipitously, as depicted for each curve in figure 2. The reason is that one is now eligible for a deferred pension with a very low normal cost, and CalSTRS assumes a sizeable probability of eschewing the refund, which is worth much more. Figure 3 illustrates for age-25 entrants. The dotted curve represents the cost rate for refunds, extended beyond the vesting period. The dashed curve is the normal cost rate for the pension; this is quite low immediately upon vesting—far below the employee contribution rate—since FAS is frozen and the

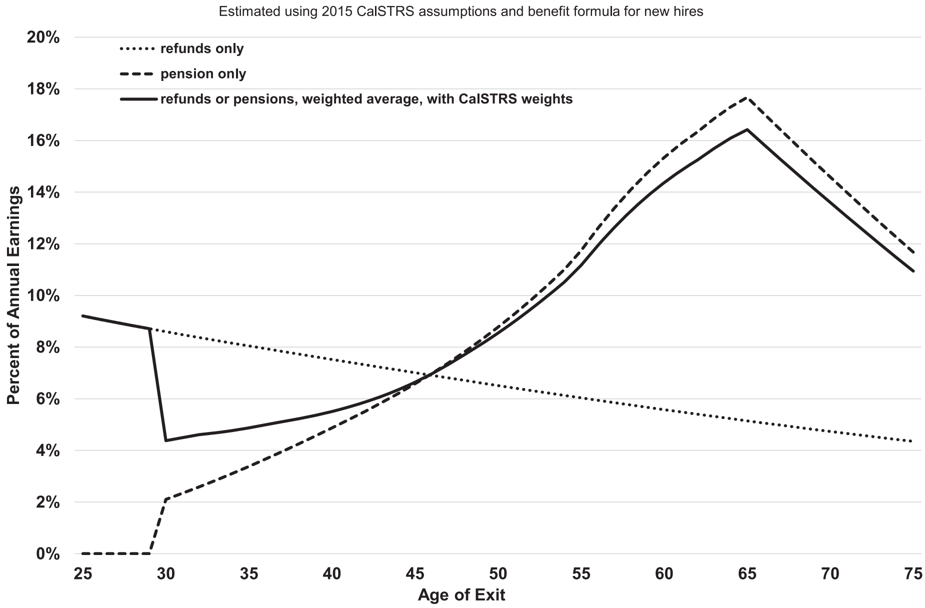

Note: The curves depict the annual contribution rate required to fund an individual's benefits—refund or pension. The normal cost rate, n_{es}, for an individual entering at age *e* and exiting at age *s*, is a weighted average of the two.

Figure 3. Normal Cost Rate, Entry Age 25 Years, Pension and Refunds

pension is deferred until age 60 years. The solid curve, which is the normal cost rate, n_{es}, is an average of the pension and refund, weighted by the probability of taking the refund. Prior to vesting, the refund rate is 100 percent, so the solid and dotted curves coincide. Immediately upon vesting, CalSTRS assumes only 40 percent probability of refund, even though the refund is of higher present value. This accounts for the sharp drop in normal cost 5 years after entry.[13] For age-25 entrants, the normal cost at this point is 4.4 percent, the minimum value referred to in table 1, and depicted in both figures 2 and 3.

From vesting to age 55 years, the normal cost rate for the pension rises as the deferral to age 60 becomes shorter. Beyond age 55 years, there is no deferral, but n_{es} continues to rise as the age-specific multiplier grows.[14] Each year of delayed retirement beyond age 55 years is a year of forgone pension payments, but prior to age 65 years, the growth in the multiplier outweighs this effect. After age 65 years the multiplier stops growing, and the normal cost declines. This pattern, before and after age 65 years, is reflected in figure 2 along each curve, corresponding to any given entry age.

13. As figure 3 shows, the normal cost rate for 25-year-old entrants is reduced by those who decline refunds up to age 46 years, and by those who continue to take refunds after age 46 years. (These calculations do not reflect the probability of re-entry.) See Lueken (2014) for the first analysis of suboptimal cash-out decisions by Illinois teachers. In our analysis, optimal cash-out decisions would raise n^* by 1 percent and the range for n_{es} would be 6.5 to 24.9 percent.
14. Interestingly, there is no discontinuity at age 55 years, when the pension draw becomes immediate, instead of the six-year deferral assumed for exits at age 54 years. That is because the age-specific multiplier is higher for age 60 than for 55 years, a feature that smoothes out the accrual pattern for CalSTRS, as noted previously by Costrell and Podgursky (2009).

In addition to the variation within entry-age cohorts, figure 2 also depicts the (vertical) variation across entry ages for the same exit age. As stated earlier, this can go either way, and we see this in figure 2. The predominant pattern is for normal cost to rise with entry age (the cost curves are generally higher for later entrants), but not always.[15] Thus, for exit at age 65 years, the normal cost rate is highest, at 21.7 percent, for age-45 entrants, as referred to in table 1 and shown in figure 2. This general pattern represents a reward (in effect, earning benefits at a higher rate) for those with shorter years of service, when that is due to later entry. This seems hard to square with the usual rationale for traditional FAS plans of rewarding longer years of service.[16]

Cross-Subsidy Rates

The wide variation among individual cost rates contrasts with the uniform contribution rate. That is the weighted average, n^*, as given in equation 2. This is the normal cost rate that will fund the benefits of each or all cohorts, past and present, represented in the current workforce, under the current benefit formula and assumptions. We calculate n^* to be 14.1 percent of pay, depicted in figure 2 as the solid gray horizontal line.[17] The deviations of individual cost rates from n^* represent the cross-subsidy rates, $(n_{es} - n^*)$. Those above the line receive cross-subsidies from those below the line. For example, the extreme points depicted, of 21.7 and 4.4 percent, represent cross-subsidies of +7.6 and −9.7 percent of pay.

The dashed gray line, representing the employee contribution rate of 9.2 percent, allows us to identify the employer contribution, 4.9 percent of pay, as the distance between that line and the total contribution line, n^*. This provides further perspective on the magnitude of the cross-subsidies. Consider the extreme cases. The individual with a normal cost rate of 21.7 percent receives employer-funded benefits of 12.5 percent of pay (21.7 − 9.2), well over double the employer contribution seemingly made on her behalf. Conversely, the individual with a cost rate of 4.4 percent receives no employer-funded benefit, and loses the value of over half of her own contributions (4.8 out of 9.2 percent) to fund the benefits of others.

These cross-subsidies are built into the funding plan (which includes the assumed rate of return). For those individuals below the solid gray line, the plan is counting on using some or all of the employer contributions—plus, for many, part of the employee contributions—along with the assumed returns on those contributions to help finance the benefits of those above the line.

4. THE DISTRIBUTION OF NORMAL COST RATES: HOW MANY WINNERS AND LOSERS?

We have identified the normal cost rates for all ages of entry and exit, as displayed for selected entry ages in figure 2. We now examine the frequency distribution of entrants with those normal cost rates. To do so, we calculate the joint frequency of entrants by

15. For any given exit age, the starting pension for older entrants, with shorter service, is a lower percent of FAS, which reduces normal cost. But their FAS is higher relative to cumulative earnings (since it is a shorter stream), which raises normal cost. The latter generally outweighs the former, as in most of figure 2, although we see the reverse case for entry ages 40 and 45 years up to around exit age 63 years.
16. Appendix figure A.3 depicts the relationship between n_{es} and years of service, within and across entry ages.
17. This is within a point of CalSTRS's estimate of the normal cost rate for new entrants (net of death and disability).

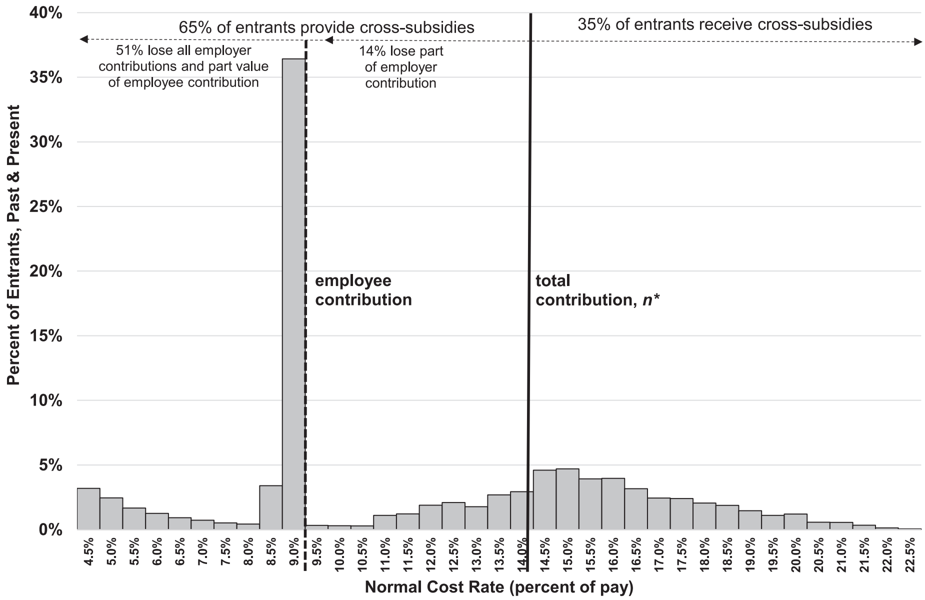

Note: Each bar represents the fraction of entrants with given normal cost rate. For example, 36.4% of entrants have normal cost rate centered on 9.0% of pay (8.75%–9.25%).

Figure 4. Frequency Distribution of Individual Normal Cost Rates

ages of entry and exit,[18] p_{es}, and link them to the corresponding normal cost rates n_{es}. Figure 4 depicts the result. Each bar represents the share of entrants with the normal cost rate indicated (±0.25%).

We can immediately see the top-line results for winners and losers. The winners, who receive benefits worth more than the contribution rate, are those to the right of the vertical line for n^*.[19] These make up 35 percent of all entrants. They receive cross-subsidies from the other 65 percent, to the left of the n^* line. Among the latter, 14 percent lose some, but not all, of the employer contribution (those between the vertical lines for employee and total contributions), and 51 percent lose the employer contribution plus part of the value of their own contribution.

At a more fine-grained level, we see a large concentration of entrants with normal cost around 8.5 to 9 percent. These are almost entirely those who leave before vesting. Those with lower normal costs are those who leave shortly after vesting, especially entrants in their 20s, with the lowest normal costs. Those with normal costs around n^* are predominantly individuals who enter in their 20s and stay until age 55 years or later. Finally, those toward the right of figure 4, with normal costs from 16 percent and beyond, generally enter in their 30s or 40s and stay to retirement. These are the individuals who receive the largest cross-subsidies.

18. For each entry age e, the conditional frequency of exit age s, $p_{s|e}$, is found by constructing the "survival" rate (i.e., probability of still working), from CalSTRS's assumed exit rates (retirement or withdrawal), and differencing that series. The joint frequency, p_{es}, is given by $p_{s|e} \cdot p_e$, where p_e is the distribution of entry ages.
19. These are the number of entrants corresponding to the normal cost rates above the n^* line in figure 2.

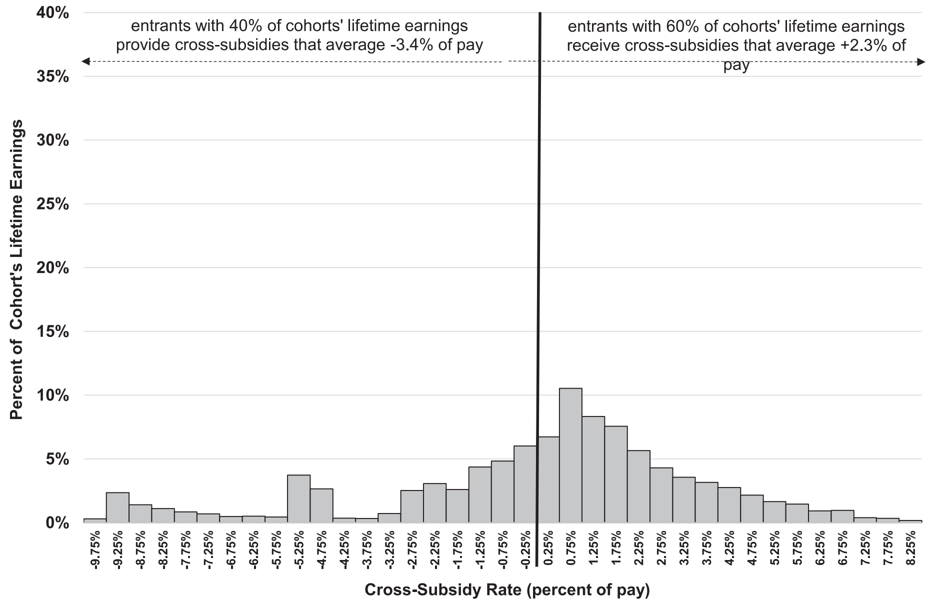

Note: Each bar represents entrants' fraction of cohort's lifetime earnings with given cross-subsidy rate. For example, entrants with 10.5% of cohort's lifetime earnings receive cross-subsidy centered on +0.75% of pay.

Figure 5. Distribution of Cross-Subsidies: Their Weighted Sum Is Zero

Cross-Subsidies Sum to Zero

As discussed earlier, the weights for n_{es} used to construct n^* (the uniform contribution rate that funds each cohort) are the shares of each cohort's PV of earnings, $(p_{es}W_{es})/(\Sigma_e\Sigma_s p_{es}W_{es})$. Thus, as shown in equation 3, the weighted sum of the cross-subsidy rates, $(n_{es} - n^*)$, is zero. Figure 5 shows the distribution of cross-subsidy rates, with these weights as frequencies. It is worth reiterating that the cross-subsidies and their distribution pertain not only to the entering cohort (the top bar in figure 1), but the aggregate of all past and present cohorts (the entire outlined set of cohorts in figure 1), under the assumptions stated previously.

The cross-subsidy rates in figure 5 are simply the normal cost rates in figure 4 minus n^* (14.1 percent). However, using shares of PV of earnings instead of the proportion of entrants, p_{es}, shifts the weight of the distribution to the right. The reason is simple: Those with lower normal cost rates (negative cross-subsidies) tend to be early leavers with shorter earnings streams, so smaller shares of the cohort's PV of earnings. Thus, although the losers constitute 65 percent of all entrants (as shown in figure 4), they account for only 40 percent of their lifetime earnings (as shown in figure 5).

How large are the cross-subsidies? Taken together, the losers provide cross-subsidies that total −3.4 percent of their lifetime earnings. That is the weighted average for the left side of figure 5. By contrast, the winners receive cross-subsidies that average +2.3 percent. One can readily verify the zero-sum result: $0.60 \times 2.3\% - 0.40 \times 3.4\% = 0.0\%$. Stated differently, the cross-subsidies reduce the effective employer contribution to an average of 1.5 percent of pay for the losers, and raise it to 7.1 percent for the winners—a sizeable gap in employer-funded benefits.

Comparison with Fiscally Neutral Cash Balance Plan

We have defined winners and losers within the plan by how the annual cost rate to fund an individual's benefit compares with the uniform contribution rate charged for all, n_{es} vs. n^*. The division between winners and losers is the same if we compare present values instead, B_{es} vs. $n^* W_{es}$. Using either measure, we can equivalently identify the plan's winners and losers in comparison with an alternative, fiscally neutral cash balance (CB) plan. In a CB plan, each individual's benefit is given by a (notional) retirement account balance (to be annuitized or drawn down) equal to the cumulative value of contribution credits (employee contributions and employer credits), plus accumulated interest credits. For a fiscally neutral CB plan (i.e., the plan's total cost is the same as the FAS plan), the contribution credits equal the FAS plan's contribution rate, n^*, and interest is credited at the same assumed rate of return, r.[20] In such a CB plan, the normal cost rate is identical for all, regardless of age of entry or exit—it is simply the uniform contribution rate, n^*—so there are no cross-subsidies.

The winners and losers in a traditional FAS plan versus this CB plan are identified by comparing the present value or annual cost of individual benefits in the two plans. This division is the same as comparing the individual cost and uniform contributions within the FAS plan, n_{es} vs. n^*. Thus, in comparing the CalSTRS plan with such a CB plan (14.1 percent contribution credit and 7.5 percent interest credit), we have the same results as given earlier: 65 percent of all entrants (with 40 percent of their cohort's lifetime earnings) would do better under the CB plan.

Actives vs. Entrants

The picture of winners and losers under traditional FAS plans presented here and previously has been challenged by Rhee and Fornia (2016; hereafter RF). We find 65 percent of all CalSTRS entrants are losers, whereas RF (p. 31) find only 21 percent of "active teachers" are worse off, compared with a CB plan. Because the RF results have been prominently invoked in the policy debate over teacher pension plans, it is important to understand why there is such a difference.[21] The key difference is their focus on "active teachers." Entrants who exit early are largely excluded. For example, RF find that although "40% of new hires leave before vesting, these leavers represent just 6% of teaching positions" (p. 5). Rhee (2017a, b) claims their approach is the comprehensive one, adjuring readers in the title of a recent commentary to "Look at the Whole Workforce, Not Just New Entrants." However, the claim is misleading, as can be visualized in figure 1. New entrants (the top bar) are representative of all cohorts, past and present. For the complete picture one must think of the full tilted layer cake, as depicted, moving

20. In practice, CB plans can differ in two ways: (1) contribution rates can vary with years of service, as in Kansas' Tier 3 plan for teachers (Schmitz 2016); and (2) the interest credit is typically below the assumed rate of return.
21. CalSTRS (2016c) highlights the study in their Comprehensive Annual Financial Report (p. 8, "UC Berkeley Study Validates Teacher Defined Benefit Pensions") and features it on their Web site (CalSTRS 2016d), along with three promotional videos, "Key Findings," a "Fact Sheet," and a press release. The advocacy group "Californians for Retirement Security" criticized (via Twitter) *New York Times* reporter Mary Williams Walsh (2017) for not drawing on RF in her front-page story on teacher pensions. Most recently, Randi Weingarten (2017) cites RF to show "how well defined benefit pension plans serve the teaching profession," in an American Federation of Teachers release critiquing Aldeman and Robson (2017). (See also Aldeman 2017 for his reply.) It should be noted that in contrasting FAS plans with CB plans, the term "defined benefit" is inaccurately deployed, since both FAS and CB plans are defined benefit.

through the workforce from entry to exit. This is the full cycle of work lives for which the contributions are designed to fund the benefits, with the full set of winners and losers. By focusing on "active teachers" (those to the right of the vertical bar) RF slice off the portion of the layer cake to the left, excluding all previous entrants who have left by the date of their snapshot. By our estimate, this exclusion amounts to 70 percent of all relevant entrants, past and present. Because those excluded are disproportionately losers by virtue of their shorter tenure, more of those who remain are winners—an example of "survivorship bias."

The bias here arises from the difference between the joint frequencies p_{es} among entrants and actives. The number of entrants, past and present, who enter and exit at ages e and s can be partitioned at time 0 into those who have already left ("inactives"[22]) and those who are projected to leave later ("actives"). Specifically, the actives among type (e, s) are those who entered since time $e - s$ and inactives entered earlier. Let $-T$ represent the earliest cohort with any remaining actives at time 0 (the bottom bar's entry date in figure 1) and let N_t be the size of cohort t. Then we can partition all entrants of type (e, s) between inactives and actives as follows:

$$\sum_{t=-T}^{0} p_{es} N_t = \sum_{t=-T}^{e-s-1} p_{es} N_t + \sum_{t=e-s}^{0} p_{es} N_t. \tag{7}$$

These three sums are, respectively, those of type (e, s) on both sides of the vertical bar in figure 1, those to the left, and those to the right. The joint frequency of type (e, s) for each of these three groups, is the corresponding sum in the term above, divided by its sum over all types (e, s). For all entrants, this simplifies to p_{es}, as previously stated, but not for actives:[23]

$$p_{es} = \frac{\sum_{t=-T}^{0} p_{es} N_t}{\sum_e \sum_s \sum_{t=-T}^{0} p_{es} N_t} \neq \frac{\sum_{t=e-s}^{0} p_{es} N_t}{\sum_e \sum_s \sum_{t=e-s}^{0} p_{es} N_t} = p_{es}^{active}. \tag{8}$$

Figures 6 and 7 depict the joint frequencies p_{es} and p_{es}^{active} for selected entry ages.[24] The contrast is striking. Among all entrants, in figure 6, the frequency of early exits is high and is also sizeable for exits at ages 55 to 65 years. Among actives, in figure 7, those who exit early are almost all gone, and those who remain are far more likely to stay beyond age 55 years. This skews the calculation of winners and losers, since the early leavers are losers and many of those who stay beyond age 55 years are winners. This impact can be readily verified by applying these frequencies to the PV of benefits as calculated by RF (their Appendix B) for CalSTRS's FAS plan and RF's CB plan.[25] Applying the frequencies we have calculated for actives (figure 7) to their PVs, we closely replicate their estimated percent of losers. Replacing those frequencies with the ones for entrants (figure 6), raises the estimate by 37 percentage points, the vast majority of

22. At the risk of confusion with the actuarial meaning of "inactives," we use the term as shorthand for prior leavers.
23. Factoring out $\sum_{t=-T}^{0} N_t$ from the first ratio, and using $\sum_e \sum_s p_{es} = 1$ gives the first equality. We cannot factor out $\sum_{t=e-s}^{0} N_t$ from the second denominator, since the range of t varies with e and s, hence the inequality.
24. Unlike the distribution of entrants, the distribution of actives varies with the growth of cohort size, N_t. We take N_t as constant (as does CalSTRS), but the results are insensitive to reasonable variation in its growth.
25. They used the same selected entry ages as above to estimate the winners/losers under CalSTRS's FAS plan vs. CB.

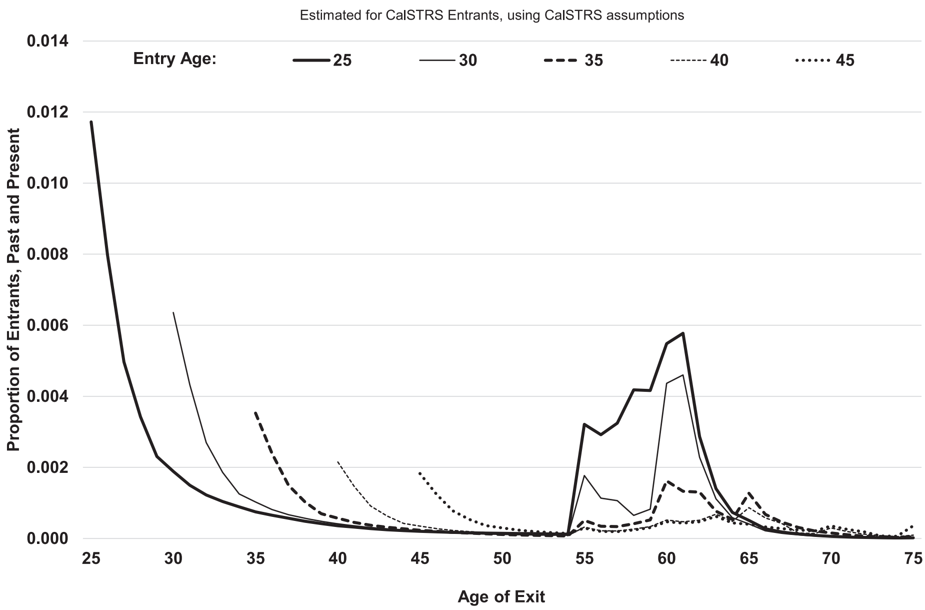

Note: The curves depict p_{es}, the joint frequency of individuals entering at age *e* and exiting at age *s*, among all entrants. Variation by age of exit is shown along each curve; the height of each curve reflects the share of entrants of age *e*.

Figure 6. Joint Frequency of Entry and Exit Ages, all Entrants, p_{es}

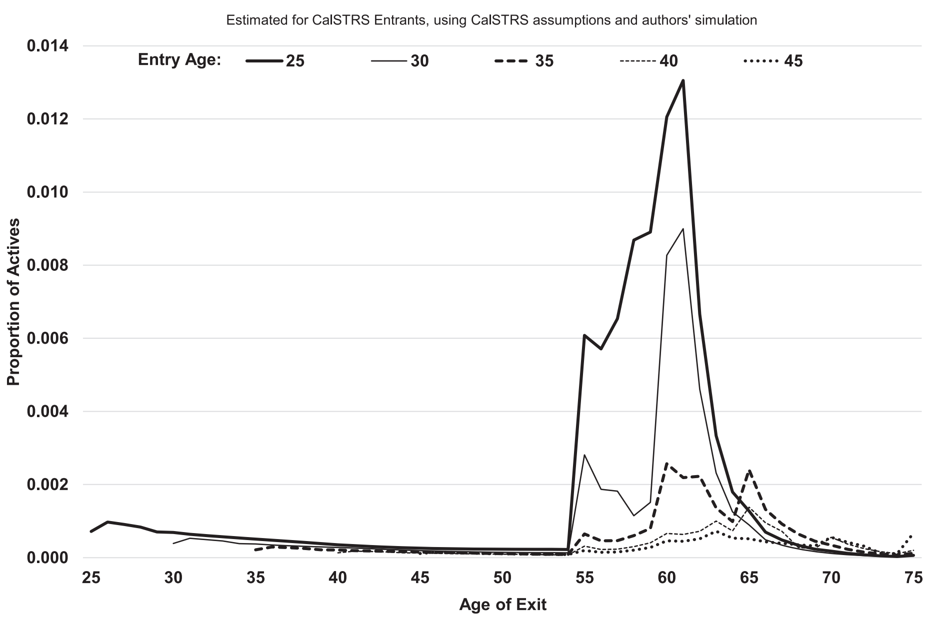

Note: The curves depict p_{es}^{active}, the joint frequency of active individuals who entered at age *e* and will exit at age *s*. Variation by age of exit is shown along each curve; the height of each curve reflects the share of actives who entered at age *e*.

Figure 7. Joint Frequency of Entry and Exit Ages, Actives, p_{es}^{active}.

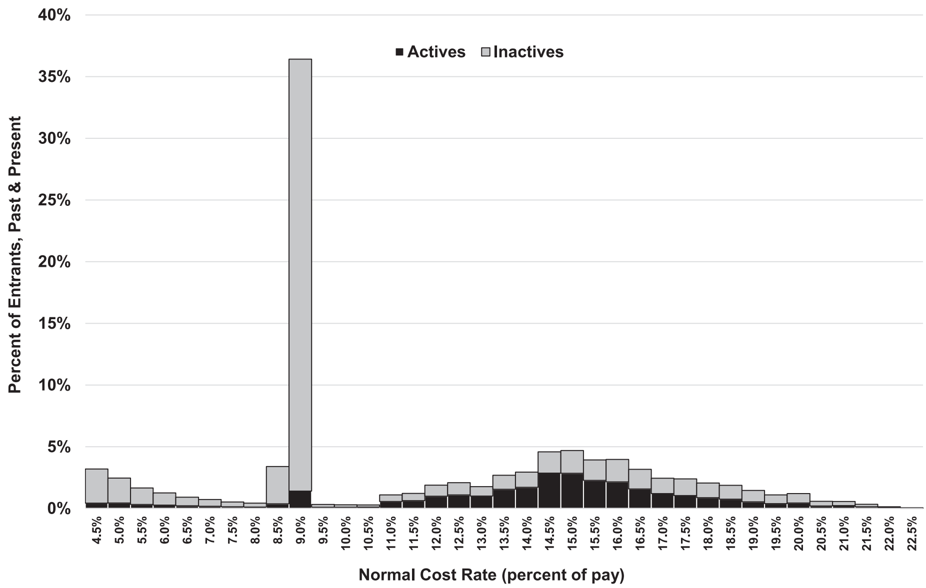

Note: Each bar represents the fraction of all entrants with given normal cost rate, broken down between actives and inactives. For example, 36.4% of entrants have normal cost rate centered on 9.0%, of which 1.4% are active and 35.0% are inactive.

Figure 8. Frequency Distribution of Normal Cost, Actives and Inactives

the 44-point difference between our studies (21 percent vs. 65 percent).[26] Clearly, RF's shift in focus from entrants, past and present, to actives dramatically reduces the count of losers.

RF argue that "focusing on new-hire attrition is misleading" (p. 5) because early leavers are unimportant "from a public education policy perspective" (p. 6). But even leaving aside issues of equity and retirement security for early leavers, the focus on winners and losers among actives (instead of all entrants, active and inactive) provides an incomplete financial picture. Focusing on actives omits the cross-subsidy left behind by early leavers to help fund benefits of those who stay. Consequently, whereas the gains and losses among entrants must add up to zero (as shown in equations 3 and 4), the gains outweigh the losses among actives. Thus, among actives, we find (using RF's PV estimates) that the average benefit under CalSTRS's FAS plan is about 25 percent higher than under RF's CB plan. This is not surprising, as it would be hard to imagine the losses of the 21 percent losers offsetting the gains of the 79 percent winners. Using RF's PV estimates again, but taking the average for all entrants, active and inactive, brings the average much closer to parity.[27]

Finally, we return to our metric of normal costs. Because early leavers are those with the lowest normal costs, their absence leaves the distribution of actives skewed toward the middle and higher range of costs. Figure 8 illustrates. The full bars represent the

26. The remainder is due to the fact that RF do not compare CalSTRS's FAS benefit to a fiscally neutral CB plan. Instead they compare to a less generous and thus less costly CB plan by using an interest credit and annuitization rate (7 percent) that is lower than the CalSTRS earnings assumption in effect at the time of their study (7.5 percent).

27. Our estimate of the remaining gap is 6 percent, attributable to the fiscal non-neutrality of RF's CB plan.

Table 2. Cross-Subsidy Rates (percent of pay): Winners/Losers and Actives/Inactives

	Inactives	Actives	All
Losers	−4.3% (0.23)	−2.2% (0.17)	−3.4% (0.40)
Winners	+2.6% (0.26)	+2.0% (0.34)	+2.3% (0.60)
All	−0.6% (0.49)	+0.6% (0.51)	0.0% (1.00)

Notes: Each cell gives the average cross-subsidy rate (percent of pay) for the group indicated by column and row, along with that cell's share of all entrants' PV of earnings (in parentheses). For example, inactive losers, who account for 0.23 of the PV of earnings, provide a cross-subsidy of −4.3 percent of their earnings.

distribution of normal cost rates among all entrants, active and inactive (including all entry ages, not just selected ones), reproducing figure 4. The RF approach excludes the inactives (the light top segment of each bar), leaving only the actives (the dark bottom segments). Clearly, the vast preponderance of those with low normal cost rates are excluded, far more so than for those with average and higher normal costs. In other words, RF are correct to state that a majority of actives are winners under CalSTRS—we calculate that majority at 60 percent (the share of the short dark bars to the right of $n^* = 14.1$ percent). But they constitute only 18 percent of the full entering cohorts from which the actives originate (their share of the full bars in figure 8). Their gains, together with those of the winners among the inactives (17 percent of the full cohorts), are funded by the losses of the much larger group of losers—65 percent of all entrants, past and present.

Specifically, we find the net cross-subsidy received by actives from inactives is about 0.6 percent of pay. This is not huge, but not trivial either. It represents a 13 percent boost to employer-funded benefits, raising the effective employer contribution rate to 5.5 percent of pay for actives and reducing it to 4.2 percent for inactives.[28] Of course, there remain substantial cross-subsidies among actives, as there are among inactives, and entrants as a whole. Table 2 provides a cross-tab of the cross-subsidies from losers to winners, among inactives, actives, and all entrants taken together. Each cell in the table gives the average cross-subsidy rate (as a percent of pay) and, below it, that cell's share of all entrants' PV of earnings. Inactive losers provide an average cross-subsidy of −4.3 percent of pay (reducing their employer contribution below 1 percent). Active losers provide an additional cross-subsidy of −2.2 percent, such that active and inactive winners gain +2.0 and +2.6 percent of pay, respectively. The weighted sums across rows give the cross-subsidies from all losers to all winners, −3.4 and +2.3 percent of pay, as stated previously. Similarly, the weighted sums down the columns give the

28. There are several reasons the net cross-subsidy received by actives is smaller than one might expect, given the large number of winners among them. First, the number of winners among the actives in our calculations is not as large as in RF, because our CB is fiscally neutral. Second, because the inactives have shorter tenures than the actives, their share of the entrants' PV of earnings is lower, as is, accordingly, the magnitude of the cross-subsidy they provide. Finally, the actives have younger entry ages, which reduces normal cost for any given exit age.

cross-subsidy rate from all inactives to all actives, −0.6 to + 0.6 percent of pay, also as stated above. For the table as a whole, the weighted sum is zero.

5. CONCLUSION

The distinguishing characteristic of traditional FAS pension plans is that the benefit is delinked from contributions, unlike CB and other account-based plans. Some individuals receive benefits that cost more than the contributions made by or for them, and some receive less. This creates a system of hidden cross-subsidies, varying by age of entry and exit. Previous work on teacher pensions has analyzed such redistribution with the metric of pension wealth. The present paper adds to previous literature by measuring the value of individual benefits as the annual contributions required to fund them, as a percent of pay. The wide variation in these individual cost rates contrasts with the uniformity in the contribution rate actually levied, clearly revealing the system of cross-subsidies. In effect, many entrants help fund the benefits of others through some or all of their employer contribution and even some of their own contribution. Without these cross-subsidies, the employer and/or employee would have to contribute much more to fund the benefits of career teachers.

Our analysis covers the full range of entry and exit ages, illuminating additional patterns of cross-subsidies to those previously identified between short-termers and career teachers. Benefits vary not only by age of exit, but also, as we show, by age of entry. Older entrants are clearly, and strongly, subsidized by younger entrants leaving with the same years of service. Indeed, older entrants are generally subsidized by younger entrants leaving at the same age, even though the older entrants have fewer years of service. These patterns are difficult to reconcile with the usual claim that traditional FAS systems rationally reward longer years of service.

We have illustrated the analysis with CalSTRS plan data, portraying the full pattern of normal cost rates by age of entry and exit, ranging from about 4 to 22 percent of pay. Compared with the uniform contribution rate, these represent cross-subsidies of about −10 percent to +8 percent of pay. Overall, we find that 65 percent of all entrants are losers, receiving benefits that cost less than the contributions made by or for them. On average, the losers provide a cross-subsidy of −3.4 percent of their pay to enhance the benefits of the winners by +2.3 percent of their pay.[29] Stated differently, about two thirds of all entrants receive employer-funded benefits that average only 1.5 percent of their earnings, compared with 7.1 percent for the other third.

We have also addressed a dispute in the policy debate over the number of winners and losers between a traditional FAS plan and a fiscally neutral CB plan (i.e., where the total cost is the same). Rhee and Fornia (2016) have argued that the educationally relevant shares of winners and losers in a teacher pension plan (CalSTRS, specifically) are those of the active teaching force, rather than new entrants. However, the current entry cohort is representative of all past entry cohorts, of whom only a fraction remains active. The snapshot of actives excludes all the prior leavers among previous cohorts,

29. We should reiterate the caveat that these are our estimates of the cross-subsidies embedded in the funding plan, which rests on the actuarial assumptions, including the investment return. If the assumptions are not met (as has often been the case with the investment returns), unfunded liabilities will ensue, and the actual cross-subsidies will include another dimension, from future generations of teachers and taxpayers.

including a disproportionate number of losers. This approach skews the comparison, leading to the conclusion that the vast majority of active teachers are winners under the traditional FAS plan, compared with a fiscally neutral CB plan (and even more so in comparison to non-neutral plans, as in RF). Focusing on the snapshot of active teachers results in funding math that simply does not add up. For a fiscally neutral comparison, the gains of winners must add up to the losses of losers and that will not hold if we omit the losses of prior leavers. To accurately understand how an FAS plan distributes contributions among its members, we must consider all entrants, past and present, active and inactive. As we have shown, this includes cross-subsidies between (and among) prior leavers and actives.

What are the policy implications of this analysis? At the very least, any good policy should be transparent. Where traditional FAS plans are used, the system of hidden cross-subsidies should be laid bare. The uniform contribution rate, designed for funding purposes, masks the wide variation in individual cost rates. These rates can be readily calculated, by age of entry and exit, as a byproduct of the annual actuarial valuations, and should be made publicly available, so that members can better understand how their plan may affect them.[30]

There is reason to go further, by reducing the actual variation in cost rates. One of the reasons employers offer retirement plans is to help workers save enough across their careers to reach a secure retirement. The low savings rates effectively offered to workers who leave early (in order to help fund those with high normal costs) have the potential to endanger this goal, placing them at a big retirement savings deficit. There are steps that can be taken within traditional FAS plans to somewhat mitigate the system of cross-subsidies (Costrell and Podgursky 2011), such as indexing FAS to inflation or general wage growth (as under Social Security and the state plan of South Dakota). However, the complex system of hidden cross-subsidies will never be eliminated under FAS plans. This will require a CB or other account-based system, where the employer contribution is the employer-funded benefit, transparent to all. If the contribution is uniform, so are the rewards.

In addition to the issues of transparency and equitable provision of retirement security, it is important to consider the implications of pension reform for recruitment and retention of teachers, with potential consequences for teacher quality. Advocates of FAS plans claim their back-loading of benefits provides incentives to reduce turnover (e.g., Rhee and Fornia 2016; Weingarten 2017). If so, this would reduce the proportion of novice teachers, and, other things equal, raise average teacher quality. On the other hand, as shown by Costrell and Podgursky (2009), FAS systems generate not only strong “pull” incentives to stay on until retirement eligibility, but also “push” incentives to leave shortly thereafter—often in one’s mid-50s—so the impact on average experience is theoretically ambiguous. Moreover, if the incentive to stay on for one’s

30. Of course, individuals do not know in advance when they will exit, but foreknowledge of the ex post normal cost rates at various exit ages would be informative. More informative yet would be the marginal cost rates, for each extra year of service, as shown in Costrell and Podgursky (2009). Such information would also inform the policy debate over plan design. As McGee and Winters (2014) show, rational risk-averse teachers, who do not know at entry when they will exit, should, ex ante, consistently prefer a CB plan, with uniform rewards, to the uncertain individual rewards of FAS plans, and the preference would be strong under plausible degrees of risk aversion.

pension disproportionately alters the decisions of those who are burned out and otherwise ready to leave, FAS systems may have perverse effects on teacher quality, through self-selective retention.

The empirical evidence suggests that the policy impact of a move from FAS to CB or other account-based plans, with benefits that accrue smoothly in tandem with contributions, would be unlikely to adversely affect teacher quality.[31] At the very least, then, if the move would do no harm, then considerations of equity and transparency alone would argue for it. Moreover, there is reason to believe that such a move could enhance teacher quality for many of the reasons that FAS systems fail to do so. First and foremost, plans that tie benefits to contributions, rather than arbitrary combinations of age and service, recognize and accommodate the heterogeneity of teachers. Under FAS plans, the empirical evidence clearly shows the pull and push incentives forcefully concentrate retirements around those arbitrary age and service combinations. Account-based plans, by contrast, recognize that individual teachers may optimally choose short or long careers and, importantly, facilitate mobility.

Second, if human resource goals are to include rewarding longevity, CB plans can do so more efficiently. As we have seen, FAS systems do not reward longevity consistently, due to entry age effects, as well as the "push" incentives, and, most importantly, due to the often cliff-like nature of the reward to extra service at arbitrary points. CB systems can reward longevity far more rationally, by designing employer contributions to rise smoothly (no cliffs) with years of service, to accommodate teacher heterogeneity. For example, under Kansas' Tier 3 CB plan (Schmitz 2016; KPERS 2017), the nation's first such plan covering teachers, the employer match (to the employee's contribution of 6 percent) rises from 3 percent of pay for years 1 to 4 of service to 6 percent of pay for years 24 and beyond.[32]

One could go a step further and tie the employer match to the evidence regarding the productivity impact of experience. Vigdor (2008) has laid out the case for constructing such an "evidence-based" salary scale. As he summarizes the evidence, the first few years of experience matter most for teacher productivity, so a corresponding salary scale would ratchet up wages more quickly in the early years of teaching and, subsequently, greatly reduce the degree of back-loading in the salary schedule.[33] As he pointed out, the account-based pension reforms discussed in Costrell and Podgursky (2008)—the same ones we continue to endorse here—would also reduce the back-loading in compensation. In the same way that Vigdor constructed a salary scale that mirrored the evidence on productivity growth, one could readily do the same for the employer match in a CB plan for the retirement benefit component of compensation. McGee and Winters (2017) illustrate for two large school systems how the benefit accrual from a CB plan with a

31. See the survey by Koedel and Podgursky (2016), as well as the recent papers by Ni and Podgursky (2016), McGee and Winters (2016), and Roth (2017).

32. As stated earlier, these are notional contributions or "retirement credits" to individual account balances. The individual employer normal cost rate (averaged over one's years of service) for providing these credits would range from 3.0 percent to 4.7 percent, with interest credits equal to the assumed discount rate of 7.75 percent. However, the interest credit formula generates a lower average rate. Under the actuarially assumed interest credit of 6.25 percent, the range in individual employer normal cost rates is only 3.0 percent to 3.7 percent. Either way, the range is much narrower than under FAS systems, as we have seen.

33. Grissom and Strunk (2012) present evidence that reduced back-loading in the salary schedule is in fact associated with better school performance.

uniform employer match would bring the trajectory of total compensation into closer alignment with the pattern of productivity growth. More generally, the optimal plan design for the CB match—constant or gently rising over specified years of service—would be integrated with the salary structure. If the salary structure amply rewards experience over the relevant intervals, then there would be less rationale for a rising CB employer match—a uniform match should do; conversely, if the salary schedule is flat over intervals of rising productivity, then a gently rising CB employer match might be called for.

Finally, the transparency of the rewards under a CB or other account-based system may enhance the efficiency of the embedded incentives. Although FAS systems embed very strong financial incentives for certain behaviors, if they are fully understood, the evidence seems to suggest that teachers respond more to "focal points" (such as the age of retirement eligibility) that are strong, but imperfect, proxies for the actual financial incentives (see, for example, Kim 2016). Such behavior can lead to excessive uniformity of retirement decisions on the one hand, and mitigated response to costly incentives on the other hand. A smoothly varying system of employer matches, transparent to all, and tightly tied to evidence-based human resource goals could more efficiently advance those goals. In short, a CB or other account-based system, tying benefits directly to contributions, offers a more effective and equitable vehicle for delivering transparent and deliberate rewards to meet the goals of teachers and employers than the seemingly arbitrary system of cross-subsidies that are embedded in traditional FAS plans.

ACKNOWLEDGMENTS

This paper was presented at the 42nd Annual Conference of the Association for Education Finance and Policy. The authors thank Andrew Biggs and other participants, as well as the very helpful referees for their comments and suggestions. We also thank CalSTRS and William Fornia for making available the fine-grained data used by Rhee and Fornia (2016). Costrell gratefully acknowledges research support from the Laura and John Arnold Foundation. The views expressed here are those of the authors and should not be attributed to their institutions, data providers, or funders. Any and all errors are attributable to the authors.

REFERENCES

Aldeman, Chad. 2017. *AFT still working at cross-purposes with teachers' interests*. Available bellwethereducation.org/aft-continues-bury-its-head-teacher-pensions. Accessed 23 October 2017.

Aldeman, Chad, and Kelly Robson. 2017. *Why most teachers get a bad deal on pensions*. Available educationnext.org/why-most-teachers-get-bad-deal-pensions-state-plans-winners-losers/. Accessed 23 October 2017.

Aldeman, Chad, and Andrew J. Rotherham. 2014. *Friends without benefits: How states systematically shortchange teachers' retirement and threaten their retirement security*. Available https://bellwethereducation.org/sites/default/files/BW_PensionPaper_031314.pdf. Accessed 5 December 2018.

Backes, Ben, Dan Goldhaber, Cyrus Grout, Cory Koedel, Shawn Ni, Michael Podgursky, P. Brett Xiang, and Zeyu Xu. 2016. Benefit or burden? On the intergenerational inequity of teacher pension plans. *Educational Researcher* 45(6): 367–377. doi:10.3102/0013189X16666812.

California State Teachers Retirement System (CalSTRS). 2016a. *Defined benefit actuarial valuation as of June 30, 2015*. Available www.calstrs.com/sites/main/files/file-attachments/2015_db_valuation_report.pdf. Accessed 23 October 2017.

California State Teachers Retirement System (CalSTRS). 2016b. *Member handbook*. Available https://www.calstrs.com/sites/main/files/file-attachments/memberhandbook2016_final.pdf. Accessed 5 December 2018.

California State Teachers Retirement System (CalSTRS). 2016c. *Comprehensive annual financial report for the fiscal year ended June 30, 2016*. Available https://www.calstrs.com/sites/main/files/file-attachments/cafr2016.pdf.

California State Teachers Retirement System (CalSTRS). 2016d. *UC Berkeley study: Are California teachers better off with a pension or a 401(k)?* Available www.calstrs.com/are-california-teachers-better-pension-or-401k. Accessed 26 May 2017.

Costrell, Robert M. 2018a. Accounting for the rise in unfunded public pension liabilities: Faulty counterfactuals and the allure of simple gain/loss summations. *Journal of Pension Economics and Finance* 17(1): 23–45.

Costrell, Robert M. 2018b. The 80 percent pension funding target, high assumed returns, and generational inequity. *Contemporary Economic Policy* 36(3): 493–504.

Costrell, Robert M. 2018c. *Policy brief: Distribution of teacher pension benefits in California: A vast system of cross-subsidies*. Available www.uaedreform.org/downloads/2018/01/distribution-of-teacher-pension-benefits-in-california-a-vast-system-of-cross-subsidies.pdf. Accessed 23 January 2018.

Costrell, Robert M., and Josh McGee. 2010. Teacher pension incentives, retirement behavior, and potential for reform in Arkansas. *Education Finance and Policy* 5(4): 492–518. doi:10.1162/EDFP_a_00013.

Costrell, Robert M., and Josh McGee. 2017. Cross-subsidization of teacher pension costs: The impact of assumed market returns. Paper presented at APPAM Fall Research Conference, Chicago, November.

Costrell, Robert M., and Michael Podgursky. 2008. Peaks, cliffs, and valleys: The peculiar incentives of teacher pensions. *EducationNext* 8(1): 22–28.

Costrell, Robert M., and Michael Podgursky. 2009. Peaks, cliffs, and valleys: The peculiar incentives in teacher retirement systems and their consequences for school staffing. *Education Finance and Policy* 4(2): 175–211. doi:10.1162/edfp.2009.4.2.175.

Costrell, Robert M., and Michael Podgursky. 2010a. Distribution of benefits in teacher retirement systems and their implications for mobility. *Education Finance and Policy* 5(4): 519–557. doi:10.1162/EDFP_a_00015.

Costrell, Robert M., and Michael Podgursky. 2010b. Golden handcuffs: Teachers who change jobs or move pay a high price. *EducationNext* 10(1): 60–66.

Costrell, Robert M., and Michael Podgursky. 2011. Reforming K-12 educator pensions: A labor market perspective. New York: TIAA-CREF Institute.

Grissom, Jason A., and Katherine O. Strunk. 2012. How should school districts shape teacher salary schedules? Linking school performance to pay structure in traditional compensation schemes. *Educational Policy* 26(5): 663–695. doi:10.1177/0895904811417583.

Johnson, Richard W., and Benjamin G. Southgate. 2015. Are California teacher pensions distributed fairly? Washington, DC: Urban Institute.

Kansas Public Employees Retirement System (KPERS). 2017. *Valuation report as of December 31, 2016*. Available www.kpers.org/valuationreport123116.PDF. Accessed 23 October 2017.

Kim, Dongwoo. 2016. Worker retirement responses to pension incentives: Do they respond to pension wealth? Unpublished paper, University of Missouri.

Koedel, Cory, and Michael Podgursky. 2016. Teacher pensions. In *Handbook of the economics of education*, volume 5, edited by Eric A. Hanushek, Stephen J. Machin, and Ludger Woessmann, pp. 281–303. Amsterdam: North-Holland Elsevier.

Lueken, Martin. 2014. Cash on the table? A behavioral analysis of refund claimants and annuitants in the Illinois teachers' retirement system. PhD dissertation, University of Arkansas at Fayetteville, Fayetteville, AR.

Lueken, Martin. 2017. (No) money in the bank: Which retirement systems penalize new teachers? Washington, DC: Thomas B. Fordham Institute.

McGee, Josh, and Marcus A. Winters. 2013. *Better pay, fairer pensions: Reforming teacher compensation*. Available https://www.manhattan-institute.org/pdf/cr_79.pdf. Accessed 5 December 2018.

McGee, Josh, and Marcus A. Winters. 2014. *Better pay, fairer pensions II: Modeling preferences between defined-benefit teacher compensation plans*. Available https://www.manhattan-institute.org/pdf/cr_90.pdf. Accessed 5 December 2018.

McGee, Josh, and Marcus A. Winters. 2016. *Better pay, fairer pensions III. The impact of cash-balance pensions on teacher retention and quality: Results of a simulation*. Available https://www.manhattan-institute.org/sites/default/files/R-JMMW-0616.pdf. Accessed 5 December 2018.

McGee, Josh, and Marcus A. Winters. 2017. How pensions contribute to the premium paid to experienced public school teachers. *Educational Researcher* 46(5): 250–258. doi:10.3102/0013189X17721906.

Ni, Shawn, and Michael Podgursky. 2016. How teachers respond to pension system incentives: New estimates and policy applications. *Journal of Labor Economics* 34(4): 1075–1104. doi:10.1086/686263.

Rhee, Nari. 2017a. *Teacher pensions vs. 401(k)s: Part 1—Look at the whole workforce, not just new entrants*. Available laborcenter.berkeley.edu/teacher-pensions-vs-401ks-part-1-look-at-the-whole-workforce-not-just-new-entrants/. Accessed 27 May 2017.

Rhee, Nari. 2017b. If someone tells you your kid's teacher would be better off with a 401(k) than a pension, don't believe it. *Los Angeles Times*, 18 July.

Rhee, Nari, and William B. Fornia. 2016. *Are California teachers better off with a pension or a 401(k)?* Berkeley, CA: UC Berkeley Center for Labor Research and Education.

Roth, Jonathan. 2017. *Union reform and teacher turnover: Evidence from Wisconsin's Act 10*. Available https://sites.hks.harvard.edu/pepg/PDF/Papers/PEPG17_02.pdf. Accessed 5 December 2018.

Schmitz, Kirsten. 2016. *Breaking down Kansas' new retirement plan for teachers*. Available www.teacherpensions.org/blog/breaking-down-kansas-new-retirement-plan-teachers. Accessed 23 October 2017.

Vigdor, Jacob. 2008. Scrap the sacrosanct salary schedule. *EducationNext* 8(4): 36–42.

Walsh, Mary Williams. 2017. Pension crisis in Puerto Rico fleeces young. *New York Times*, 9 March.

Weingarten, Randi. 2017. *Flawed Bellwether Education Report on teacher pensions*. Available https://www.aft.org/sites/default/files/memo_bellwether-report_052217.pdf. Accessed 5 December 2018.

APPENDIX: ENTRY AGE NORMAL (EAN) COST

We consider the standard actuarial calculation of the normal cost rate: entry age normal (EAN). This method is designed to determine a uniform contribution rate (employer plus employee), which, collected over the careers of all members of any entering cohort, would cover benefits earned by that cohort. The main point here is that although the individual normal cost rates, n_{es}, are not explicitly calculated under the EAN method, they are implicit in the EAN algorithm to calculate the uniform contribution rate.

The second point of this appendix is to show the minor difference between the EAN method and the variant used in this paper (equation 2). Because individuals enter and exit at varying ages, there are two types of averaging: across ages of exit and entry. In the text, we consider this averaging in one step, using joint frequencies p_{es} and the corresponding share for type (e, s) of the cohort's PV of earnings. The EAN method constructs the uniform contribution rate in two steps, first constructing the average across exit ages of the normal cost rate for any given entry age and then averaging across entry ages. The first step is consistent with our approach, but the second step introduces a slight deviation from PV of earnings weights.

Normal Cost Rate for Entry-Age Cohort *e*

The first step in the algorithm is to calculate the contribution rate that would fund the benefits of entry-age cohort e. Members of that cohort will exit at different ages, unknown in advance, so their cost rate, call it n_e, is constructed from probabilistic averages across exit ages s. Formally, let $p_{s|e}$ denote the frequency of exit at age s, conditional on entry at age e. Then EAN calculates, for an entrant of age e, the ratio of the expected value of benefits to that of lifetime earnings (discounted to date of entry):

$$\begin{aligned} n_e = \sum_s p_{s|e} B_{es} \Big/ \sum_s p_{s|e} W_{es} &= \sum_s n_{es} p_{s|e} W_{es} \Big/ \sum_s p_{s|e} W_{es} \\ &= \sum_s n_{es} \left(p_{s|e} W_{es} \Big/ \sum_s p_{s|e} W_{es} \right), \end{aligned} \tag{A.1}$$

where we have used $n_{es} = B_{es}/W_{es}$, in the second expression. Thus, we have shown that the individual normal cost rates are implicitly embedded in the EAN algorithm. Specifically, as the last expression shows, n_e is the weighted average of n_{es} where the weights are the shares of PV of earnings for each exit age s among entrants of age e, $(p_{s|e} W_{es}/\Sigma_s p_{s|e} W_{es})$. The normal cost rate n_e is the contribution rate from entrants of age e that would fund their collective benefits.

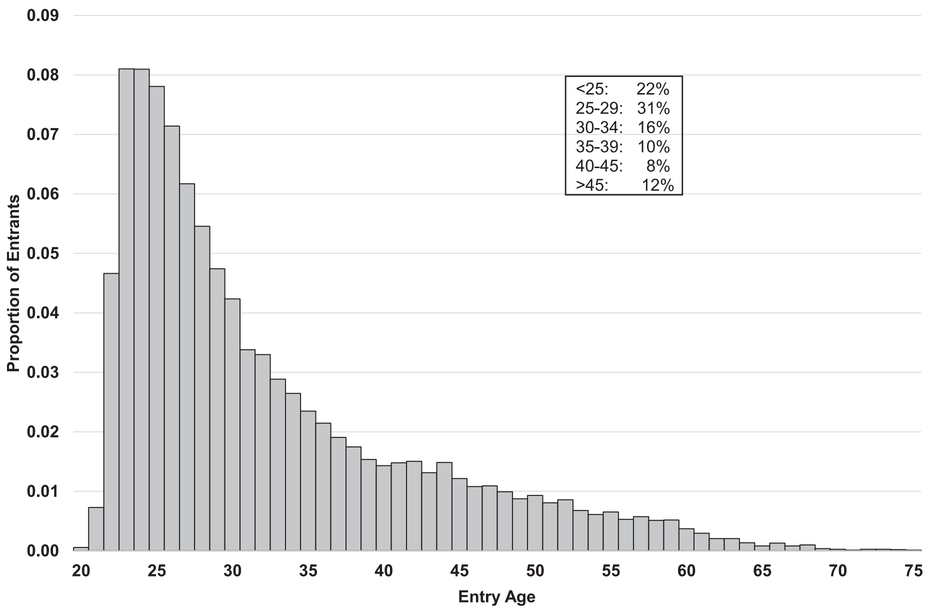

Note: CalSTRS 2014 data, courtesy of William Fornia. This is the marginal distribution of entry ages, $p_e = \Sigma_s\, p_{es}$, where the joint frequencies, p_{es}, are those given in the text.

Figure A.1. Distribution of Entry Ages

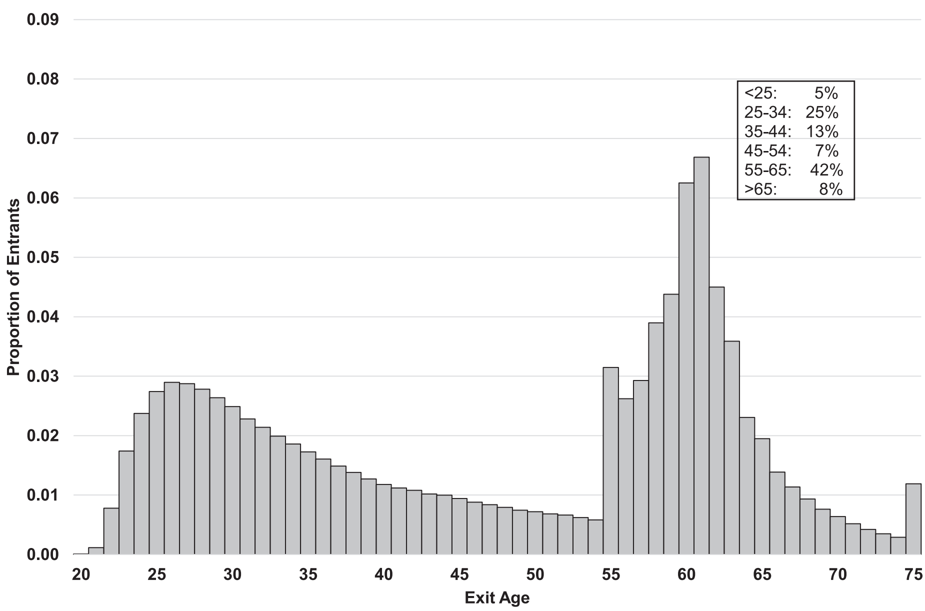

Note: Calculations from CalSTRS assumptions. This is the marginal distribution of exit ages, $p_s = \Sigma_e\, p_{es}$, where the joint frequencies, p_{es}, are those given in the text.

Figure A.2. Distribution of Exit Ages

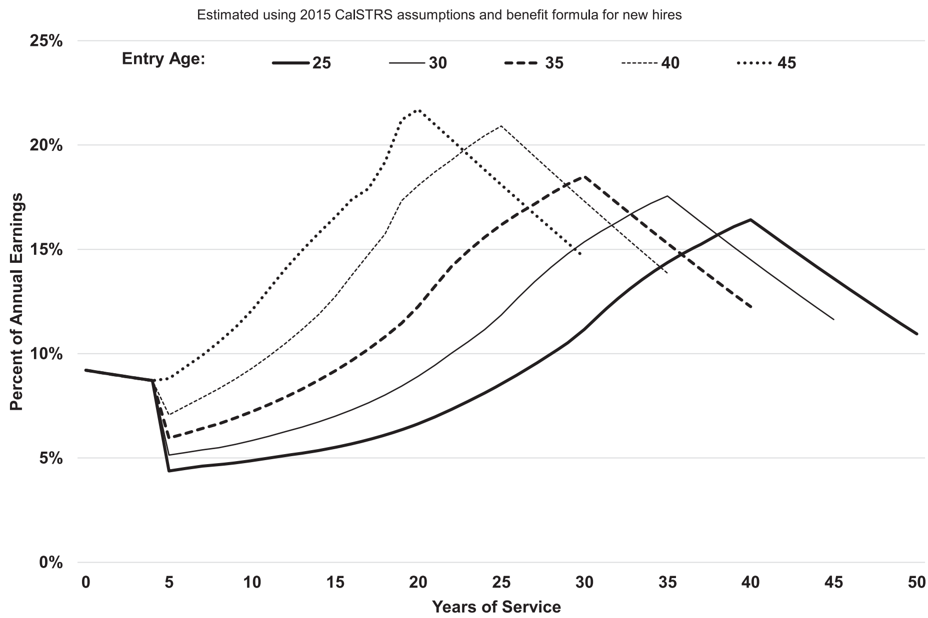

Note: The curves depict n_{es}, the annual contribution rate required to fund benefits of an individual entering at age *e* and serving (*s* — *e*) years. Variation in cost by years of service is shown along each curve; variation by age of entry is shown across curves.

Figure A.3. Normal Cost Rate, by Entry Age and Years of Service

Overall Normal Cost Rate, Across Entry Ages

The second step in determining the uniform cost rate is to average the normal cost rates, n_e, across all entry ages. The weights that would be conceptually consistent are, again, shares of the entering cohort's PV of earnings. The uniform normal cost rate, using these weights, is

$$n^* = \sum_e n_e \left(p_e W_e \Big/ \sum_e p_e W_e \right), \tag{A.2}$$

where $p_e = \Sigma_s p_{es}$ is the marginal frequency of entry age *e* and $W_e = \Sigma_s p_{s|e} W_{es}$ is the PV of earnings for the average entrant of age *e*. Substituting into equation A.1, and noting that $p_{s|e} = p_{es}/p_e$, we find text equation 2. That is, using PV of earnings weights sequentially, first among entrants of age *e*, and then across entry ages, is equivalent to using the joint weights $(p_{es} W_{es})/(\Sigma_e \Sigma_s p_{es} W_{es})$ in one step. This is the uniform contribution rate, applied to all entrants, across entry ages, that would fund the cohort's collective benefits.

However, this is not quite the standard EAN algorithm. That algorithm applies n_e from equation A.1 to each individual's pay in the current workforce and divides by current payroll. Formally, this is equivalent to

$$n^{EAN} = \sum_e n_e \left(p_e^{active} average(w_e) \Big/ \sum_e p_e^{active} average(w_e) \right), \tag{A.3}$$

where $average(w_e)$ is the average wage of those currently active who entered (now or in the past) at age *e*. Note how the weights on n_e differ from those in equation A.2. They are the shares of current payroll for actives of entry age *e* vs. the shares of all entrants' PV of earnings for those of age *e*. The weights used in equation A.2 seem more consistent with the first step of the EAN method and with the method's overall intent: to generate contributions that will fund the benefits of each entering cohort.[34] In practice, however, the difference seems quantitatively minor. We calculate $n^{*} = 14.1$ percent for CalSTRS using equation A.2 [or equation 2, directly] and $n^{EAN} = 13.9$ percent, using equation A.3.

34. They are also the weights that make the cross-subsidies sum to zero, so they are appropriate for this paper's focus.