CAMBRIDGE
UNIVERSITY PRESS

Cross-subsidization of teacher pension benefits: the impact of the discount rate

Robert M. Costrell*

Department of Education Reform, University of Arkansas, Fayetteville, AR 72701, USA
*Corresponding author. E-mail: costrell@uark.edu

Abstract

This paper builds on previous work (Costrell and McGee, 2017*a*, *Education Finance and Policy*) on the redistribution of teacher pension benefits, as measured by the wide variation in *individual* normal cost rates by age of entry and exit, and the associated cross-subsidies. The further steps taken here are: (i) to examine the impact of the discount rate on the degree of redistribution, and the analytics behind it; and (ii) to identify the distribution of the market value of the pension guarantee. Using the example of the California State Teachers' Retirement System, I find that: (i) high-assumed returns substantially understate the extent to which costs are redistributed for back-loaded plans; and (ii) the market value of the pension guarantee is highly concentrated among long-term teachers.

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.[1] This can be illustrated by the variation in the annual cost to fund each individual's benefits. These inequities are masked by a *uniform* contribution rate for pensions. For example, the annual joint contribution (employer and employee) for pension normal costs may be 15% of each teacher's salary. These contributions are designed to fund the future retirement benefits as they are earned,[2] for the system *as a whole*. Such contributions are conceptually comparable with those made to account-based plans, such as 401(k)'s, 403(b)'s, and cash balance (CB) defined benefit (DB) plans. However, unlike account-based plans, the annual cost of benefits for *individual* teachers under traditional pension formulas may deviate widely from this overall average (Costrell and McGee, 2017*a*; Costrell, 2018*c*; Costrell and Fuchsman, 2018). It is well-established that the benefits for early leavers are of much lower value than for those who retire at the 'sweet spot;' contributions by or for the former effectively include a cross-subsidy to the latter. The further step taken here is to examine the impact of the discount rate on this system of redistribution. As is well-known, a drop in the discount rate raises the overall normal cost rate − this is regularly reported in valuation reports when the assumed rate of return is reduced. However, the impact on the distribution of *individual* normal costs has not been previously explored, and this is the subject of this paper.

---

[1]Much of this line of research dates to Costrell and Podgursky (2008, 2009, 2010*a*, 2010*b*).

[2]In addition, the employer makes payments for the unfunded liability – benefits earned in the past, but not funded. This is a very large problem, but is not the subject of this paper. The inter-generational 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, 2018*a*, *b*). For an analysis that integrates inter- and intra-generational cross-subsidies, see Costrell and McGee (2017*b*).

There are two possible interpretations and motivations for this inquiry. First, since the discount rate is the expected return on the investment portfolio, under current public sector practice, this analysis considers how the cross-subsidies vary as the expected return is reduced in funding plans to better align with evolving market conditions. The gap between assumed returns and actual returns in recent times has led to rising and persistent unfunded liabilities. Consequently, public plans are widely, if belatedly, gradually cutting their assumed return.[3] In doing so, the plans' normal cost rates, previously held artificially low by high-assumed returns, are rising. In this paper, we examine whether the measured degree of redistribution of pension benefits has also been held artificially low by high-assumed returns.

A second motivation for examining the impact of the discount rate is to assess the distribution of the value of the pension guarantee to members of defined benefit plans. Since the expected return includes a risk premium, netting that out of the discount rate (i.e., using a risk-free rate) tacks on a measure of the market value of the guarantee (Brown and Wilcox (2009), Novy-Marx and Rauh (2009), and, relatedly, Biggs (2011)). We know that the guarantee raises the value of benefits, as measured by the overall normal cost rate (Richwine and Biggs, 2011). In this paper, I examine whether it raises the individual values uniformly or whether – and how much – it tilts the distribution of benefits and increases the cross-subsidies.

Analytically, the impact of a lower-discount rate (for either reason given above) on the degree of redistribution is not immediately obvious. The effect on individual normal costs varies with the age of exit for multiple reasons, and the impact does not vary monotonically. Overall, however, using the illustration of the California State Teachers' Retirement System (CalSTRS), I find that a lower-discount rate raises the degree of redistribution significantly.[4] Put differently, although the cross-subsidies embedded in the funding plan can be substantial, these understate the redistribution of benefits as the expected return is reduced. When evaluated with the risk-free rate, to assess the market value of the pension guarantee, I find the value of that guarantee is highly concentrated and very large indeed for career teachers, but non-existent for short-termers.

The structure of the paper is as follows. First, I will review the basic math of individual normal costs and the associated cross-subsidies embedded in the funding plan, under any given discount rate. I illustrate with the case of CalSTRS, using the plan's current expected return. I then compare the magnitude of the cross-subsidies with those under the higher expected return held until recently, and possibly lower-assumed returns in the future. I then examine the math behind the effect of the discount rate on individual normal cost rates and how that effect varies, to better understand the distributional impact in question. Next, I consider a risk-free rate; although it is unlikely to be used for funding purposes under current public sector practices, it does illustrate the extremely wide variation in individual values of the benefit guarantee. Finally, I will briefly discuss potential policy implications, limitations of this work, and extensions to it.

2. Individual normal cost rates and cross-subsidization

Pension plans calculate the normal cost rate at the aggregate level, to fund a cohort's benefits as they accrue. Individual cost rates, based on age of entry and exit are implicitly embedded within the calculation (Costrell and McGee, 2017a, Appendix), but they are not publicly reported. Specifically,

---

[3]The average assumed return for state and local plans has been trimmed from 8.1% in 2000 to 7.4% in 2017 (Census of Governments, as reported by Center for Retirement Research at Boston College, Public Plans Database: National Data, 2018 and reproduced in Boyd and Yin, 2018). Meanwhile, although the average assumed return for private plans was also about 8% in 2000, it has fallen much more sharply, to about 4.5% by 2012 (last year reported in Boyd and Yin, 2018).

[4]CalSTRS is one of the largest public pension plans in the USA, so it is an important example in its own right. It is perhaps of more general interest, representative of other public plans, because, like them, it is back-loaded through its final average salary (FAS) structure. Back-loading is key to the analysis below. For example, similar results obtain for the impact of the discount rate on the redistribution under the Massachusetts Teacher Retirement System and the Arkansas Teacher Retirement System.

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), one can identify an individual normal cost rate, n_{es}, as a constant percent of salary over one's career. This rate is calculated to generate a stream of contributions sufficient to fund the individual's future benefits. That is, the present value (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) and B_{es} is the PV of benefits (both evaluated at entry). It immediately follows that the individual cost rate is the ratio of the two:

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

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, analogous to contributions to a retirement account.

Across individuals with different entry and exit ages, (e,s), cost rates, n_{es}, vary widely. In general, for any given e, n_{es} rises with s, from some point after 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 (Costrell and Podgursky, 2009, 2010a). The variation in n_{es} with e, for any given s, is less obvious, and can go either way.[5]

Traditional pension plans levy a joint (employee plus employer) contribution rate, n, that is *uniform*, applied to all members of the cohort (of varying entry ages) throughout their careers (of varying length), calculated to fund the benefits of the whole entering cohort. Formally, denote the joint frequency of ages of entry and exit, e and s, among entrants, as p_{es}. It can be readily shown that the uniform cost rate required to fund the cohort's projected benefits is[6]:

$$n = \frac{\sum_e \sum_s n_{es}(p_{es}W_{es})}{\left(\sum_e \sum_s p_{es}W_{es}\right)}. \tag{2}$$

This is the ratio of the PV of the cohort's benefits[7] 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 n_{es} are $(p_{es}W_{es})/(\sum_e\sum_s p_{es}W_{es})$, representing the share of type (e,s) in the cohort's PV of earnings.[8]

The deviations $(n_{es} - n)$ are positive and negative, as the cost of funding any individual's benefit exceeds or falls short of the cohort's uniform contribution rate, n. They effectively comprise a system of cross-subsidies. Moreover, by the nature of averages, these cross-subsidies must add up to zero,

---

[5]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}. Thus, the ratio, $n_{es} = B_{es}/W_{es}$, can rise or fall with e, over different ranges of s, various discount rates (compare Figures 1 and 4) and benefit formulas (e.g., compare CalSTRS, below and Massachusetts, Costrell and Fuchsman, 2018). Another way of seeing the ambiguity is to note that for any given exit age, the cost rate varies with the starting pension as a percent of final average salary (FAS) and with FAS relative to cumulative earnings. For older entrants, with shorter service, the starting pension is a lower percent of FAS, but their FAS is higher relative to cumulative earnings (since it is a shorter stream), so the effect on normal cost is ambiguous.

[6]It can also be shown that n applies not simply to a single entering cohort, but to any cohort, past or present, or the full set of such cohorts working their way over time through the workforce, under a given benefit formula and set of actuarial assumptions (Costrell and McGee, 2017a).

[7]Substituting $n_{es} = B_{es}/W_{es}$ into the numerator gives $\sum_e\sum_s p_{es}B_{es}$.

[8]These are not the exact weights used in actuarial practice, but are consistent with the approach (see Costrell and McGee, 2017a, Appendix).

when properly weighted, by shares of the cohort's PV of earnings:

$$\frac{\sum_e \sum_s (n_{es} - n)(p_{es} W_{es})}{\left(\sum_e \sum_s p_{es} W_{es}\right)} = 0. \tag{3}$$

The negative cross-subsidies provided by the losers fund the positive cross-subsidies enjoyed by winners. I will illustrate the system of cross-subsidies under CalSTRS' current plan, followed by an examination of how this system of redistribution is affected by the discount rate.

2.1 Individual normal cost rates for CalSTRS

I now apply these concepts to estimate the individual normal cost rates, $n_{es} = B_{es}/W_{es}$, for CalSTRS members of all entry and exit ages, $e, s = 20, \ldots, 75$. I base the calculations on CalSTRS' actuarial assumptions (slightly modified, as explained below) and benefit formula.[9]

Benefits can be in the form of a pension or refund of employee contributions.[10] If a teacher takes the refund she forgoes any future pension and receives, instead, 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 qualifying 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 age 55 or higher. Finally, teachers who leave service and qualify for an immediate pension, may still choose the refund, although it is generally not financially prudent to do so. I assume that teachers choose the refund or pension to maximize the PV of their benefits.[11]

If a teacher takes the pension, B_{es} is the PV of the stream of pension payments, weighted by her survival probabilities, discounted to entry. The payments begin with a starting pension equal to an age-specific multiplier × years of service (YOS = $s - e$) × final average salary (FAS, last 3 years), augmented annually with a 2.0% simple cost of living adjustment (COLA). Specifically, I consider the '2% at 62' program for new hires (since 2013), with multipliers ranging from 1.16% at age 55 to 2.0% at 62 and 2.4% at 65, after 5-year vesting.[12] For example, a 25-year-old entrant working to 65 retires with a starting pension of $40 \times 2.4 = 96\%$ of FAS. This formula, together with CalSTRS' actuarial assumptions, allows us to calculate the PV of benefits, relative to the PV of wages, $n_{es} = B_{es}/W_{es}$, the annual contribution rate required to fund the benefits of an individual entering at age e and exiting at age s.

2.2 Variation in normal cost rates by age of entry and exit

Consider the variation of normal cost rates under the current assumed return, $r = 7.0\%$. Figure 1 depicts n_{es} for selected (but representative) ages of entry and all exit ages. The variation is wide, from 7.3% to 25.0% (the full range, for entry ages not shown, is 6.8% to 27.1%). The pattern for any given entry age (e.g., age 25) is depicted *along* each curve, as exit ages vary. Prior to vesting,

---

[9]The actuarial assumptions cover investment returns, wage growth, exit rates, and mortality rates. These assumptions are drawn from the 2017 annual valuation report (CalSTRS, 2018a), based on the most recent 5-year experience study (CalSTRS, 2016). (The actuarial assumptions used in Costrell and McGee (2017a) were drawn from the 2015 valuation report, based on the prior 5-year experience study.) The benefit formula is delineated in the valuation report and the member handbook (CalSTRS, 2018b). This includes the retirement eligibility conditions, age-specific multipliers, cost of living adjustments (COLA), employee contribution rate, and interest rate on refunds.

[10]I leave aside death and disability benefits, which comprise about 5% of normal cost, or 1 percentage point.

[11]CalSTRS assigns probabilities of taking the refund which may not maximize PV. Our modified assumption eliminates a precipitous drop in the individual normal cost rate upon vesting, due to suboptimal cash-outs.

[12]Vested employees who withdraw before age 55 but do not cash out must defer the pension to at least age 55 and we assume they collect then. CalSTRS assumes they defer to age 62. Our modification eliminates a discontinuity in the individual normal cost rate between ages 54 and 55 that arises for lower discount rates.
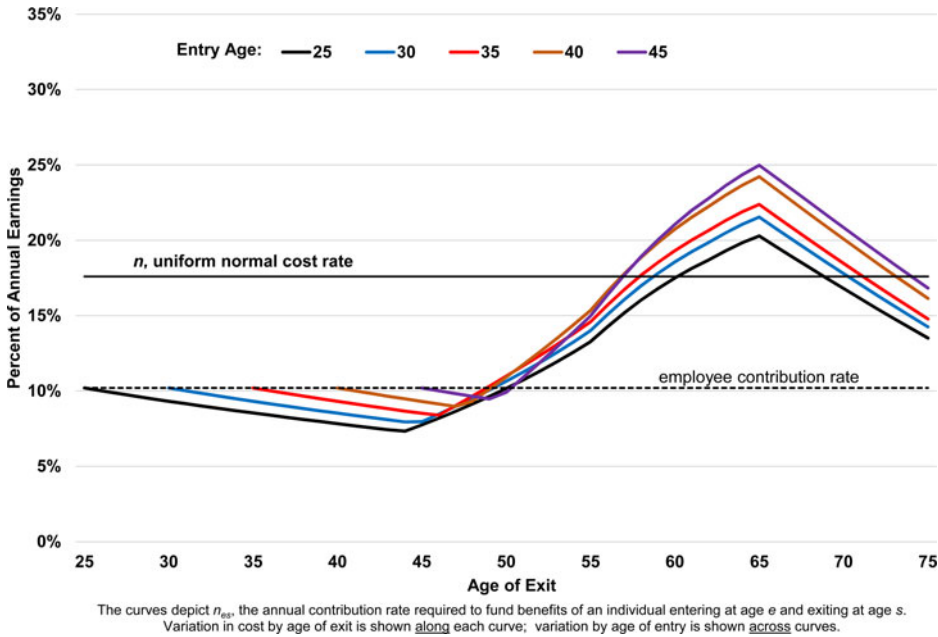
The curves depict n_{es}, the annual contribution rate required to fund benefits of an individual entering at age e and exiting at age s. Variation in cost by age of exit is shown along each curve; variation by age of entry is shown across curves.

Figure 1. Normal Cost rate, by entry age and age of exit, $r = 7.0\%$.
Note: Estimated using current CalSTRS assumptions and benefit formula for new hires, slightly modified.

and for some years beyond, the benefit is a refund of employee contributions. The normal cost rate, therefore, starts at the employee contribution of 10.2%[13]: each curve begins at the dashed horizontal line representing that rate. The cost then declines, falling slowly below the employee contribution rate. That is because the interest credit of 3.0% is below the fund's assumed return of 7.0%, with the fund retaining the difference. The contribution rate needed to cover the refund falls as this difference accumulates.

At a certain point, the pension becomes more attractive than the refund. A 25-year-old entrant reaches that point at age 44; at this age the pension would still be deferred, but exceeds in PV that of employee refunds. Beyond that point, the normal cost rate rises as the deferral becomes shorter, and then, beyond age 55, there is no deferral, but the normal cost rate continues to rise for two reasons. First, the age-specific multiplier grows.[14] In addition, as years of higher salary are tacked on to the end of one's career, FAS rises relative to career earnings, along with the annual pension payment, further raising n_{es}. Beyond age 55, each year of delayed retirement is a year of forgone pension payments, but prior to age 65, this effect is outweighed by the other factors. After age 65 the multiplier stops growing, and the normal cost declines, due to the decreasing number of years the pension will be paid. This pattern is reflected in Figure 1 along each curve, corresponding to any given entry age, e: n_{es} rises with s up to age 65 and falls thereafter. I have explained in detail this pattern of n_{es} rising with s to a peak value and then declining because it will be pertinent when analyzing the impact of a lower-discount rate.

Finally, in addition to the variation *within* entry age cohorts (i.e., along the curves depicted), Figure 1 also shows the variation *across* entry ages (i.e., the spread between curves). This widens the overall spread in individual normal cost rates, and thereby adds to the redistributive structure of such pension plans.

---

[13]The rate rose 1%, from 9.2%, effective July 1, 2018 (CalSTRS, 2018*a*, p. 1).

[14]CalSTRS is somewhat unusual in using age-specific multipliers. In other plans, with uniform multipliers, the normal cost rate rises with age of exit due to the role of age embedded in their pension eligibility conditions.

2.3 Cross-subsidy rates and the degree of redistribution

The wide variation among individual cost rates contrasts with the uniform contribution rate, n. As shown in (2), that is the weighted average of the individual cost rates, n_{es}, that will fund the benefits of each cohort, past and present, taken as a whole, under the current benefit formula and assumed return. I calculate n to be 17.6% of pay for the CalSTRS, depicted in Figure 1 as the solid horizontal line.[15] 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. For example, the extreme points depicted for n_{es}, of 7.3% and 25.0%, represent cross-subsidies of −10.3% and +7.4% of pay.[16] These cross-subsidies are built into the funding plan. For those individuals below the solid line, the plan is counting on using some or all of the employer contributions – plus, for many, some of the assumed returns on the employee contributions − to help finance benefits of those above the line.

Using the joint frequencies of entrants, p_{es}, and their shares of lifetime earnings, $(p_{es}W_{es})/(\sum_e\sum_s p_{es}W_{es})$, we can calculate a few summary statistics. Those who provide the cross-subsidies (those below the line in Figure 1) comprise 61% of entrants[17] and account for 36% of entrants' lifetime earnings[18]; those who receive the cross-subsidies are the remainder. Taken together, the losers provide cross-subsidies that total −4.8% of their lifetime earnings. That is the average cross-subsidy rate for those below the line (weighted by shares of lifetime earnings). The winners receive cross-subsidies that average +2.7%. One can readily verify the zero-sum result from (3): $0.64 \times 2.7\% - 0.36 \times 4.8\% = 0.0\%$. Thus, in all, taking absolute values of the cross-subsidies, *3.4% of total income is redistributed*, about one-fifth of the total normal cost, and almost half the employer contribution.[19]

3. The impact of the assumed return on cross-subsidies

The analysis above was based on $r = 7.0\%$, the investment return assumed since the 2017 valuation. This represents a reduction from 7.5% in 2015 and 7.25% in 2016. In reducing its assumed return from 7.5% to 7.0%, CalSTRS has, of course, raised its normal cost rate for new hires, n (about 2.1 percentage points by my estimate). The point here, however, is that the degree of redistribution also rose, from 3.0% of total income to 3.4%.

To illustrate the distributional impact of further rate reduction, consider what the system would look like if r were reduced to $r = 6.0\%$ (comparable with recent 10-year rolling averages). The individual normal cost rates are depicted in Figure 2. All the normal cost rates are increased from those depicted in Figure 1: with lower assumed investment returns, contributions must be higher to fund the benefits. This much is well-known. What is perhaps less widely understood is that a drop in the assumed return will increase the degree of redistribution embedded in the funding plan. Stated alternatively, an over-optimistic assumed return not only underfunds the plan, but also understates the degree of redistribution.

For example, the extreme points depicted in Figure 2 now represent cross-subsidies of −14.4% to +7.5%, widening the previous range (−10.3 to +7.4), especially among the losers.[20] On average, losers provide cross-subsidies that widen from −4.8% of their income to −6.9%, while winners receive cross-

---

[15]CalSTRS (2018a) calculates the normal cost rate for new hires as 17.9%. Netting out death and disability benefits, this is a bit lower than my estimate, consistent with my assumption that there are no sub-optimal cashouts.

[16]For points not depicted, the range is −10.8% to +9.5%.

[17]Rhee and Fornia (2016, 2017) argue that prior entrants who are no longer in the workforce should be excluded when counting winners and losers. But as explained in Costrell and McGee (2017a), this results in 'survivorship bias' toward winners. As a result, the losses left behind by prior leavers are excluded, such that the cross-subsidies do not sum to zero. In other words, the funding math simply does not add up.

[18]Those with lower normal cost rates (negative cross-subsidies) tend to be early leavers with shorter earnings streams, so their share of the cohort's PV of earnings is lower than their share of entrants.

[19]While the uniform employer contribution rate is 7.4%, the losers receive, on average, employer funded benefits worth only 2.6%, while those of the winners' are worth 10.1%.

[20]For points not depicted, the range widens from (−10.8 to +9.5) to (−14.8 to +8.8).

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.
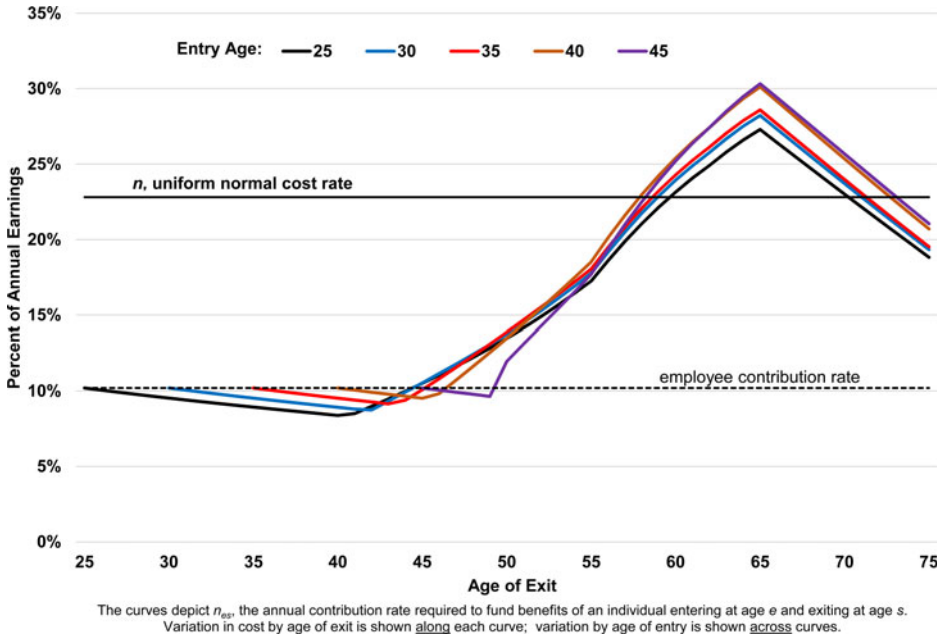
Figure 2. Normal cost rate, by entry age and age of exit, $r = 6.0\%$.
Note: Estimated using current CalSTRS assumptions and benefit formula for new hires, slightly modified.

subsidies that rise from +2.7% of income to +3.4%. The zero-sum result on cross-subsidies still holds (with winners' share of lifetime earnings now at 0.67), and, finally, taking the absolute values, we find that our measure of redistribution rises from 3.4% of total income to 4.5% (0.67 × 3.4% + 0.33 × 6.9%).

3.1 A graphical illustration of the distributional impact

Why does a cut in the discount rate increase the cross-subsidies and the degree of redistribution? On a purely mechanical level, the reason lies in the manner of variation of the impact of r on the individual normal cost rates n_{es}. If there were no variation, such that all n_{es} were to rise by the same amount with a drop in r, then n would rise by approximately that amount[21]; the cross-subsidies $(n_{es} - n)$ would remain nearly unchanged, and so would the degree of redistribution. This is not the case. The non-uniformity is seen in Figure 3, which depicts the *rise* in normal cost rates, by entry and exit age, as r drops from 7.0% to 6.0%. All the individual normal cost rates rise, but generally more so on the right side of the diagram than on the left side.[22] The uniform cost rate, n, is a weighted average of the individual rates, so it should rise by an amount that exceeds the (smaller) rise of the individual rates on the left side of Figure 3 and is generally less than the (greater) rise on the right side. And so it does: n rises by 5.2% (from 17.6% to 22.8%), while the individual normal costs rise by amounts close to zero for early departures, and up to 7% for departures at age 65.

What does this mean for the cross-subsidies? These are the gaps (negative or positive) between individual normal cost rates and the uniform rate. Since the rise in the uniform rate exceeds the rise in individual rates on the left side of Figure 3, this widens the gap (by an average of 2.1% of pay, from −4.8% to −6.9%). Conversely, on the right side, since the individual rates generally rise by more than the uniform rate, this widens the gap here, too (by an average of 0.7% of pay, from +2.7% of pay to +3.4%). Thus, on both sides, we find an increase in the magnitude (absolute

---

[21]There would be some small effect from the change in weights, induced by the change in r.
[22]I will discuss the pattern in more detail when I examine the mathematics of the impact of r in the next section.

The curves depict Δn_{es}, the rise in annual cost to fund benefits of an individual entering at age e and exiting at age s, as r falls from 7.0% to 6.0%.
Variation in Δn_{es} by age of exit is shown <u>along</u> each curve; variation by age of entry is shown <u>across</u> curves.
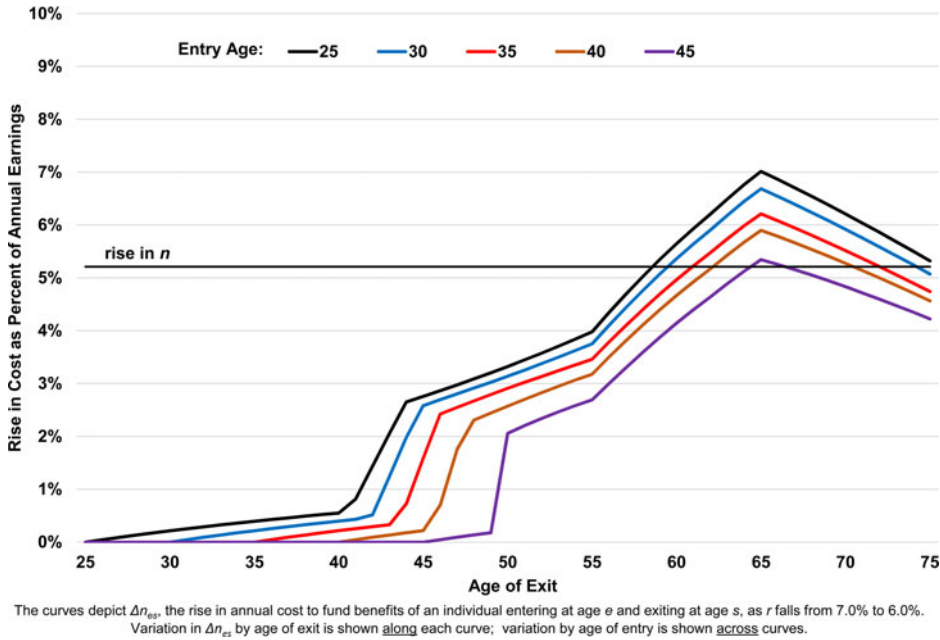
Figure 3. Rise in individual normal cost rates, $r = 6.0\%$ vs. 7.0%.
Note: Estimated using current CalSTRS assumptions and benefit formula for new hires, slightly modified.

value) of the cross-subsidies provided and received: a drop in the assumed return thereby increases my measure of the degree of redistribution.[23] Conversely, overstating the return most markedly understates the PV of long-termers' benefits and, correspondingly, understates the cross-subsidies to support those benefits.

3.2 The analytics of individual normal cost rates and the discount rate

As we have seen empirically, a drop in the discount rate raises normal cost rates for individuals of early exit ages by less than those of late exit ages. More precisely, as illustrated in Figure 3, for each entry age, the rise in individual normal cost rates increases with the age of exit, up to the point of peak normal cost. What explains this pattern? The simple answer lies in the back-loading itself of FAS pensions.[24] As I will argue, the main impact of a drop in the discount rate is to amplify that back-loading, thereby increasing the redistribution.

To elaborate a bit, before plunging into the math, as a first approximation, n_{es} rises proportionately with a drop in r. That is, a drop in r magnifies individual cost rates n_{es} more, the larger n_{es} is to begin

---

[23]Another dimension of the redistribution embedded in FAS plans is the proportion of entrants who fail to recoup the value of their own contributions (Aldeman and Johnson, 2015; Lueken, 2017). As Biggs (2018) points out, the break-even age drops with the discount rate. This can be seen geometrically from a diagram of individual normal cost rates (e.g., Figure 1). Since the cost curve crosses the employee contribution line from below, a shift up in the curve (with a cut in the discount rate), while holding the employee contribution rate constant, necessarily moves the crossing point to the left, i.e., lower break-even age and, hence, a lower number of entrants who fail to reach it.

[24]In this section of the text, I analyze the segment of the n_{es} curves where the benefit is taken as a pension. The initial segment, where the benefit is taken as a refund is simpler to analyze, and I do so here in this note. As discussed earlier, at $s = e$, n_{es} equals the employee contribution rate, independent of r, and as s increases, n_{es} drops; it costs less to pay for the refund as the fund accumulates over time the excess of r over the interest to be paid on refunds. As r is reduced, the fund's arbitrage profit narrows and n_{es} draws closer to the employee contribution rate. Thus, the *rise* in n_{es} with a drop in r increases from zero as s increases from e, throughout the range of s for refunds. This is relatively easy to show analytically, but would be overly formalistic, since the intuition is straight-forward.

with. Thus, since n_{es} rises with s up to a peak of 65 in CalSTRS, for reasons discussed earlier, Δn_{es} also rises with s up to 65 as r drops. Formally, we can write:

$$\Delta n_{es} = n_{es}[\Delta n_{es}/n_{es}] \sim n_{es}[\Delta B_{es}/B_{es} - \Delta W_{es}/W_{es}]. \tag{4}$$

My claim is that the variation across individuals of the impact of Δr on Δn_{es} primarily reflects variation in the initial level of n_{es} – the first term on the right-hand side (RHS) – rather than variation in the proportional change of n_{es} – the bracketed term on the RHS.

More precisely, to examine the distributional impact of a drop in r, we want to understand how $-\partial n_{es}/\partial r$ varies across individuals. First, we back up and examine more closely why a drop in r raises the individual rates, n_{es}, to begin with. Let us write out n_{es}:

$$n_{es} = B_{es}/W_{es} = \sum_{a>s}(1+r)^{(e-a)}b(a|e,s)f(a|s)/\sum_{a=e}^{s}(1+r)^{(e-a)}w_{a|e}. \tag{5}$$

The numerator of n_{es} – the PV of pension benefits – is the sum for all ages a, following the individual's separation s, of the PV (dated to entry e) of her annual pension benefit, $b(a|e,s)$, weighted by her probability of survival to age a, conditional on survival to exit age s, $f(a|s)$. Similarly, the denominator of n_{es}, is the sum across all ages from her entry, e, to exit, s, of the PV of her annual earnings, $w_{a|e}$. Using (5), we can then analyze the differential form of (4):

$$-\partial n_{es}/\partial r = -n_{es} \cdot \partial\ln(n_{es})/\partial r = -n_{es}[\partial\ln(B_{es})/\partial r - \partial\ln(W_{es})/\partial r]$$
$$= n_{es}\left[\sum_{a>s}(a-e)\theta_{es} - \sum_{a=e}^{s}(a-e)\gamma_{es}\right]/(1+r), \tag{6}$$

where θ_{es} and γ_{es} are weights that sum to one. Thus, $-\partial n_{es}/\partial r > 0$ for all (e,s) since the first term in brackets is a weighted average of years since entry, for the period *after* separation, and the second term is an average of years since entry for the period *before* separation. The point here is simple: a fall in the discount rate raises all the individual cost rates, $n_{es} = B_{es}/W_{es}$, because the sequence of benefits follows in time (after separation) the sequence of earnings (before separation), so the PV of benefits, B_{es}, rises proportionally more than that of earnings, W_{es}. That is, n_{es} rises as r drops because of the non-overlapping time patterns in benefits and earnings.

We now turn to the paper's main point: how the impact of a drop in r varies across individuals, particularly by s, resulting in a greater degree of redistribution. More formally, having examined the first derivative of n_{es} with respect to r, we are now, in effect, interested in the cross-partial with respect to r and s.[25] What is the impact of s on the final expression in (6)? Figure 3 suggests that the impact of s on the first term – n_{es} itself – is decisive; the back-loading of normal costs up to the peak age of 65 shapes the impact of s on the *rise* in n_{es} as r falls.

By comparison, the impact of s on the bracketed term of (6) is complex and appears *not* to be decisive – indeed its sign is ambiguous. The complexity is readily seen by noting that s appears at four points in the bracketed term: lower and upper limits, respectively, on the two summations, and the two sets of weights within those summations. Consider the impact of s through the sums' limits. Raising s eliminates the lowest ages in the first sum and adds higher ages to the second sum. Consequently, raising s increases the average number of years since entry, *both* post-separation and pre-separation. Its impact on the *difference* between the two – the bracketed term in (6), representing the non-overlapping time patterns in benefits and earnings – is ambiguous. The impact of s on the weights θ_{es} and γ_{es} is also complex and ambiguous. Since the bracketed term governs the *proportional*

---

[25]We focus on the variation across s, leaving aside variation across e, since most of the cross-subsidies are from early to late separators, and less so from early to late entrants.

change in n_{es} as r falls, these complex and ambiguous analytical effects of s are consistent with our empirical finding that they are dominated by the *absolute* impact on n_{es}, which is the back-loading itself.

To summarize, the back-loading of benefits under traditional FAS plans, which generates much of the redistributive cross-subsidies at issue, is magnified by reductions in the discount rate. As will be discussed further in the conclusion, this suggests that, by contrast, for account-based systems that are not heavily back-loaded, the distributional impact of a cut in r will be more muted. Put differently, while a high-assumed return understates the redistribution in FAS plans, this will be much less so of account-based systems: the redistribution in such systems will be minimal, independent of the discount rate.

4. The distribution of the value of the pension guarantee

Pension benefits are free of market risk, while public pension funds typically include a sizeable component of risky assets in their portfolio. Consequently, there has been animated debate between the field of finance economics (Brown and Wilcox, 2009; Novy-Marx and Rauh, 2009; Biggs, 2011) and public pension fund actors (notably their actuaries, fund managers, and advocacy organizations) over the proper discount rate for pension liabilities. It is important to distinguish between the various uses of the discount rate, as the disagreements vary by use. These uses include: (i) financial reporting; (ii) determination of contribution rates; (iii) evaluation of plan costs, including the cost of market risk; and (iv) evaluation of benefits to the members. The critique of actuarial discount rate practice for public plans has focused most sharply on (i), financial reporting (GASB rules).[26] The critique is more hedged for (ii), the determination of contribution rates, as this is recognized as more of a political decision,[27] linked to the equally political decision regarding the risk profile of fund assets.

Finance economists and public plan advocates disagree sharply on (iii), the evaluation of the plan's full costs. Advocates argue that costs are fully accounted for by discounting with the expected return on risky assets because of time diversification for immortal public plans. In other words, it is claimed, there is no cost for market risk in public plans because the law of large numbers diversifies away the variance of the *annual* return over time. Finance economists, however, have shown long ago the fallacy of time diversification, as the variance in *total* return increases over time (Samuelson, 1963). The cost of that risk may not be borne as a premium attached to the average annual contribution rate, but the variation in annual contributions (under finite amortization schedules) imposes costs on taxpayers, current or future, whether or not they are acknowledged by fund managers. Moreover, for mature public plans (where benefit outlays exceed contributions, as they now do nationally (Center for Retirement Research, 2018)), portfolios with variable annual return incur the risk of insolvency, even if the average annual return is certain (Boyd and Yin, 2018); if the period of losses precedes that of gains, the losses will be on a larger base than the gains, and the fund can deplete prematurely. This risk is also costly, but, again, it may not be factored into annual contribution rates; public plans and the taxpayers that stand behind them may, in effect, self-insure, but that does not eliminate these costs of investment in risky assets, even if they are unreported.[28]

The key point here is (iv), these costs of market risk are not borne by the plan members, the beneficiaries of the defined benefit guarantee. On this point, there is actually agreement between advocates and critics of discounting by the expected return (and of DB plans more generally). Even if they disagree on whether the costs of risky investment are borne by taxpayers or whether there is a 'free lunch' from time diversification, both sides agree that the pension guarantee is valuable to plan members (see, for example, the arguments in favor of DB plans by Rhee and Fornia (2017)). To evaluate this

---

[26]Private pension funds are governed by stricter regulations (FASB) that require lower discount rates.

[27]There may be good reasons to determine contributions using the expected return instead of the risk-free rate, as accumulating funds may tempt policy makers to raise benefits or skip contributions, rather than hold for a rainy day.

[28]Biggs (2011) shows that the risk is evaluated in the market by the value of the options that would hedge that risk, and that this is equivalent to the difference between discounting at the expected return and risk-free rate.

guarantee, finance economists argue that benefits should be discounted at something close to the risk-free rate, even if contributions continue to be determined based on expected returns. In doing so, Richwine and Biggs (2011) show how the standard normal cost rate for teachers dramatically under-states the overall benefit.

Here, I consider how the value of the pension guarantee is distributed. To do so, I take the annualized value of the guarantee to be the difference between n_{es} evaluated at a risk-free rate and that evaluated at the expected return. The former represents what the market would charge an individual of known e and s, as a percent of annual earnings, to provide an annuity following separation that replicates the pension's annual payments. By contrast, the pension fund evaluates n_{es} at the expected return – including what the market considers a risk premium – because the fund either believes its immortality diversifies away the risk, or it is willing to bear the cost of that risk (or, more accurately, to impose the cost of that risk on taxpayers). Either way, the member receives a risk-free benefit that would cost more in the market than the fund is charging. The difference between the two is the market value of the pension guarantee.[29]

Figure 4 illustrates the total value of pension benefits, including the value of the guarantee, by age of entry and exit, with n_{es} evaluated at 4.0%. The average annualized value is nearly 40%, and the spread is wide, ranging from about 10% (the value of employee contributions, for those taking refunds) to 45% or 50% for those retiring at 65.[30] Netting out the employee contribution, the value of employer-provided risk-free benefits ranges from 0% to 40%. One-third of entrants will receive employer-provided benefits worth at least 30% of pay. These are very high fringe benefit rates, but they are also highly concentrated. At the other end of the spectrum is a larger group (38% of entrants, whose benefit is the refund) for whom there are no employer-provided benefits.

My estimates for the individual market values of the pension guarantee alone are presented in Figure 5. This is the difference between the total value of the benefit, given in Figure 4 and the value of the benefit evaluated at the plan's assumed expected return, 7.0%, given in Figure 1. For those who stay to age 65, the market value of the guarantee provided by the employer is quite large, adding the equivalent of another 20–30% of pay to the 10–15% benefit covered by employer contributions and the 10% paid for by the employee herself. However, for those who receive refunds instead of a pension, the value of the guarantee is minimal – under 1 percentage point.[31] Of course, this makes sense: since the value of the guarantee is ascertained by stripping out the risk-premium from the discount rate, this will accumulate most significantly for those whose benefit is most far-removed from entry.

Conclusion

The distinguishing characteristic of traditional FAS pension plans, such as CalSTRS, is that the benefit is delinked from contributions, unlike CB or other account-based plans. Some individuals receive benefits that cost more than the contributions made by or for them, and some receive less. As in earlier work, I measure the value of individual benefits as the annual contributions required to fund them, as a percent of pay, which embed a system of hidden cross-subsidies, varying by age of entry and exit. The point of this paper is that the extent of this redistribution rises with a drop in the discount rate. Thus, if the assumed return is reduced, not only does the overall normal cost rate rise, the

---

[29]An alternative to calculating the market value of the risk-free benefit would be to evaluate the member's subjective value, using a personal discount rate. That individual discount rate often exceeds the risk-free rate, as suggested by individuals' general reluctance to buy annuities. Indeed, there is some debate as to whether personal discount rates even exceed the pension fund's assumed return, in which case members would prefer to receive more of their compensation in salaries and less in pensions (see Fitzpatrick, 2015; Goldhaber and Holden, 2018).

[30]Here we see that for any given exit age, n_{es} falls with entry age, the opposite pattern from that with higher discount rates, depicted in Figures 1 and 2, as mentioned in footnote 5.

[31]This is simply from the narrowing of the difference between the interest on refunds and the discount rate.

The curves depict n_{es}, the annualized value, as a percent of earnings, for risk-free benefits of an individual entering at age e and exiting at age s. Variation by age of exit is shown <u>along</u> each curve; variation by age of entry is shown <u>across</u> curves.
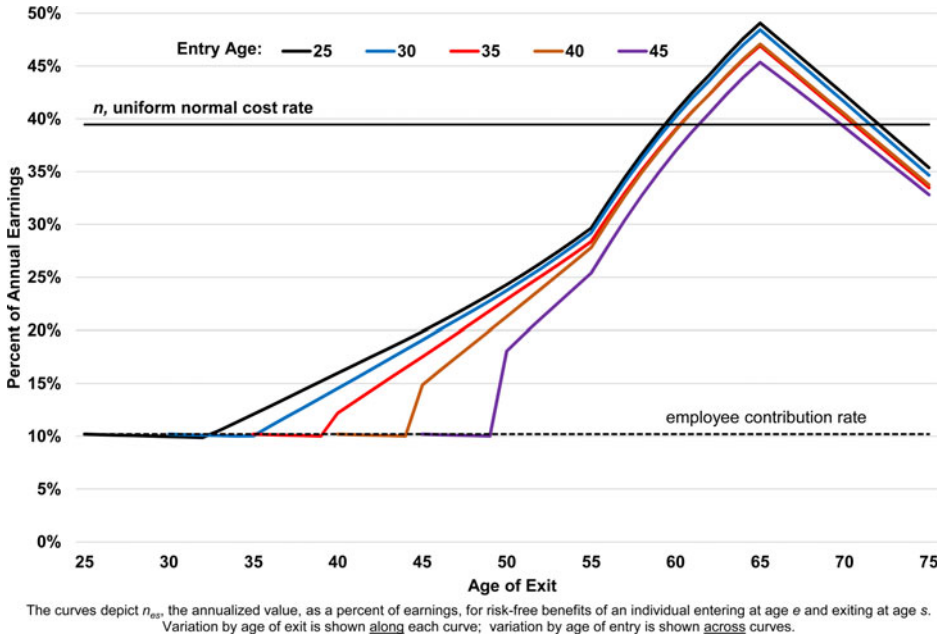
Figure 4. Annualized value of risk-free individual benefits, $r = 4.0\%$.
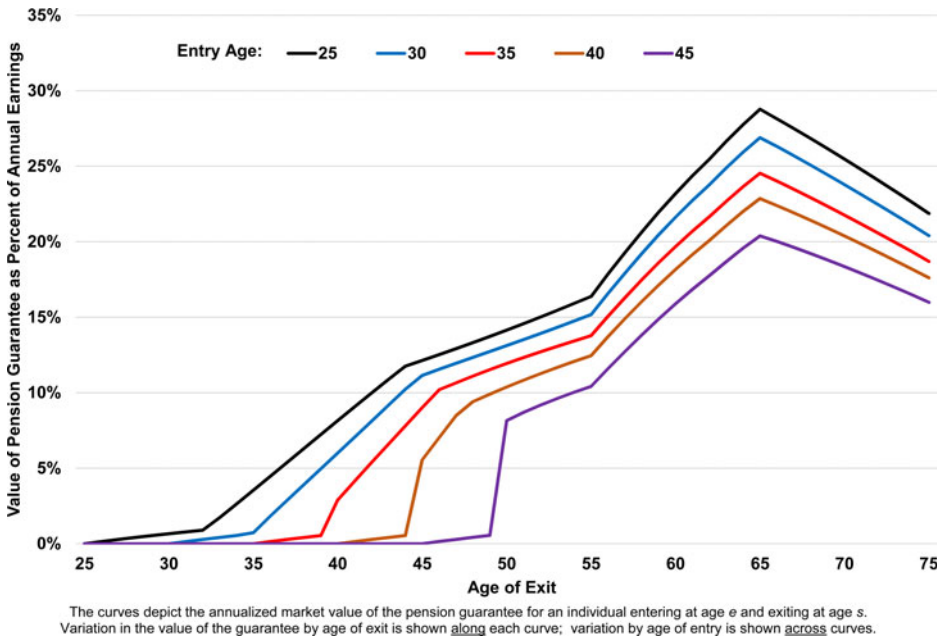*Note:* Estimated using current CalSTRS assumptions and benefit formula for new hires, slightly modified.

The curves depict the annualized market value of the pension guarantee for an individual entering at age e and exiting at age s. Variation in the value of the guarantee by age of exit is shown <u>along</u> each curve; variation by age of entry is shown <u>across</u> curves.

Figure 5. Annualized market value of pension guarantee.
Note: Difference between n_{es} evaluated at 4.0% and 7.0% for CalSTRS new hires.

individual rates rise unevenly, tilting more toward long-termers. Conversely, if the assumed returns are over-optimistic, keeping contributions artificially low, then the redistribution of the required contributions is also understated.

I illustrate this pattern by considering the CalSTRS plan for new hires at its current assumed return of 7.0 (down from 7.5% in 2015) and a hypothetical future value of 6.0%. The range of cross-subsidies depicted in Figures 1 and 2 widens from (−10.3 to +7.4)% of pay to (−14.4 to +7.5)%.[32] The average loss of those providing the cross-subsidies rises from 4.8% of their lifetime earnings to 6.9%, and the amount of redistribution rises from 3.4% of total income to 4.5%.

As the downward drift in assumed returns continues to spread, one might ask what behavioral reactions to expect from the enhanced redistribution. We already know, from previous research,[33] that teacher retirements markedly bunch around the 'sweet spot' in the formula. The question, then, is whether a drop in the assumed return – which accentuates the cross-subsidy enjoyed at that spot – increases the bunching of retirements. Under the fully informed rational actor model, the answer would appear to be no: with no change in the benefit formula, and non-tradable pension benefits, individuals would compare the payment streams at different retirement ages using an unchanged personal discount rate, divorced from the discount rate used by the plan. Alternatively, under a behavioral model that reacts to rules-of-thumb (see, for example, Kim, 2016), it is possible that disclosure of the plan's changed valuation of *individual* benefits, previously unknown, could modestly affect teacher behavior. The public and policy makers' behavioral response, discussed below, might be more important.

Since the assumed return includes a premium for risk that is not borne by members, they enjoy an additional benefit from the pension guarantee. I evaluate that benefit at the individual level, using a risk-free rate of 4.0%, for comparison with the benefit at the assumed return of 7.0%. As a result, I find that the market value of the DB guarantee is quite large, but highly concentrated. For one-third of entrants, who stay for a full career, the guarantee is worth an additional 20–30% of pay, on top of the individual normal cost of 10–15% funded by employer contributions. By contrast, a larger number of entrants (38%) receive no employer-provided benefit, with or without the guarantee.

What are the policy implications of this analysis? At the very least, any good policy should be transparent. Where traditional FAS plans are employed, 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 the cost of their individual benefits. Moreover, as both defenders and critics of traditional plans agree, the value of a defined benefit includes the value of the guarantee; the dramatically uneven distribution of that benefit should also be disclosed to inform policy discussions. As controversies over public pension policy continue to accelerate, it would surely be a valuable piece of additional information for the public and their elected representatives – as well as the teachers – to understand the extremely high market value of the pension guarantee enjoyed over such a narrow band of beneficiaries. When expressed as an annual contribution rate required for the purchase of an equivalent risk-free annuity, the public and policy makers' understanding of the benefit can be deepened.

The result of such deeper understanding could well be to adopt policies that reduce the variation in benefits and increase the degree of risk-sharing. There seems to be little evidence that the variation in benefits under traditional plans effectively or efficiently serve the human resource goals of enhanced teacher quality as often claimed.[34] Thus, equity grounds alone would argue for mitigating the cross-subsidies embedded in back-loaded FAS plans. The most efficient way of reducing the variation is through an account-based system, such as a defined contribution (DC) or CB plan. Moreover, such plans can (and often do) incorporate human resource goals of retention more rationally than the idiosyncratic FAS plans, by incorporating employer matches that gently rise with service. In addition, traditional plans expose the public to high degrees of market risk, the consequences of which are becoming

---

[32]Similar widening of normal cost rates is found in analysis of Massachusetts' Teachers Retirement System (Costrell and Fuchsman, 2018), and Arkansas' Teacher Retirement System (Costrell, 2018*e*).

[33]For a good summary of the research, see Koedel and Podgursky (2016).

[34]See Koedel and Podgursky (2016) and Costrell and McGee (2017*a*).

increasingly onerous. For DC plans, of course, that risk is entirely borne by the members instead. CB plans offer the opportunity of shared risk.

Similarly to DC, individual CB benefits are tied to a retirement account balance (to be annuitized or drawn down), but, unlike DC, the market risk is shouldered by or shared with the employer. The account balance is the cumulative value of employee contributions and employer contribution credits (a bookkeeping entry), plus interest credits, which can be fixed or vary in part with investment returns. In the baseline, hypothetical case where the interest credit equals the assumed return and the contribution credit is uniform, the contribution credit *is* the employer-funded benefit, transparent to all – just like DC. There is no back-loading, as benefits accrue smoothly in tandem with contributions at a constant percent of pay. In practice, employer contribution credits often rise with service and the interest credit is typically below the assumed return. Nonetheless, such plans are generally far less back-loaded than FAS plans, so the cross-subsidies are far more muted.

Tying this to the paper's main theme, recall our finding that the distributional impact of a drop in the discount rate in FAS plans is primarily to amplify the back-loading of those plans, thereby increasing the redistribution. Thus, a CB plan that is far less back-loaded to begin with is also distributionally far less sensitive to the assumed return. Such plans embed little or no cross-subsidization, independent of the assumed return.

For example, consider Kansas' Tier 3 CB plan (Schmitz, 2016; KPERS, 2017) – the nation's first such plan covering teachers. Under the plan's provisions and assumptions, [35] I calculate that the employer normal cost rate (following vesting), ranges from about 1–3%, under the assumed return of 7.75%. These cost rates fall within a far narrower band than under traditional FAS plans, and therefore embed much lower cross-subsidies. Evaluated using a risk-free rate of 4.0%, the employer cost rate rises to about 3–5%, a much more muted rise than under traditional FAS plans, with little increase in the spread or cross-subsidies. However, as with other CB plans, the Kansas plan embeds a risk-sharing provision. The interest credit has a floor of 4.0% and a 75% share of 5-year-average returns in excess of 6.0%. Thus, the risk-adjusted interest rate is not as low as 4.0%, so the rise in normal cost rates and cross-subsidies using that rate is more muted yet.

Finally, it is worth restating the caveat that this paper is restricted to cross-subsidies within cohorts. This assumes actuarial assumptions hold, especially the assumed return, so that unfunded liabilities do not accumulate. In recent years, however, the gap between actual and assumed returns has generated large cross-subsidies across generations. Indeed, related work on CalSTRS (Costrell and McGee, 2017*b*), aimed at an integrated treatment of cross-subsidies across and within cohorts, finds that virtually all current entrants can easily be losers, by virtue of the steady state amortization payments to cover benefits of previous cohorts, if the assumed return is held somewhat higher than the actual return. It should be noted that CB plans – like other, more traditional DB plans, and unlike DC plans – can also run up unfunded liabilities and generate cross-subsidies across generations.

---

[35]The employer match (to the employee's contribution of 6%) rises in steps from 3% of pay for years 1–4 of service to 6% of pay for years 24 and beyond. The employers' actual contribution rate is less than the notional contribution credit, because the plan's assumed return of 7.75% exceeds the interest credit, assumed to average 6.25%. In addition, employer credits are not accessible to members until vesting at 5 years, and annuitization must be deferred until age 65 for 5–10 years of service, or age 55 with 10 or more years of service. In addition, the interest rate embedded in the annuity factor is the assumed return minus 2% (5.75%).

References

Aldeman C and Johnson RW (2015) Negative Returns: How State Pensions Shortchange Teachers, Bellwether Education Partners and Urban Institute.

Backes B, Goldhaber D, Grout C, Koedel C, Ni S, Podgursky M, Xiang PB and Xu Z (2016) Benefit or burden? On the intergenerational inequity of teacher pension plans. *Educational Researcher* **45**, 367–377.

Biggs AG (2011) An options pricing method for calculating the market price of public sector pension liabilities. *Public Budgeting & Finance* **31**, 94–118.

Biggs AG (2018) Costs and cross-subsidies: reconciling two strands of teacher pension analysis. RAND Education, Working Paper 1234.

Boyd D and Yin Y (2018) Investment risk and its potential consequences for teacher retirement systems and school districts. RAND Education, Working Paper 1250.

Brown J and Wilcox D (2009) Discounting state and local pension liabilities. *American Economic Review* **99**, 538–542.

California State Teachers' Retirement System (2016) Experience Analysis, July 1, 2010–June 30, 2015, Milliman, December 30.

California State Teachers' Retirement System (2018a) Defined Benefit Actuarial Valuation as of June 30, 2017, Milliman, April 9.

California State Teachers' Retirement System (2018b) Member Handbook, Sacramento, CA.

Center for Retirement Research (2018) Public Plans Database: National Data, Boston College.

Costrell RM (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**, 23–45.

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

Costrell RM (2018c) Policy Brief. Distribution of Teacher Pension Benefits in California: A Vast System of Cross-Subsidies, University of Arkansas, January 15.

Costrell RM (2018d) Cross-Subsidization of Teacher Pension Costs: The Impact of the Discount Rate, University of Arkansas, February 27.

Costrell RM (2018e) Arkansas Teacher Retirement Plan: Risks, Redistribution and Remedies, Testimony to Arkansas Legislature, Joint Committee on Public Retirement, September 11.

Costrell RM and Fuchsman D (2018) Policy Brief. Distribution of Teacher Pension Benefits in Massachusetts: An Idiosyncratic System of Cross-Subsidies, University of Arkansas, February 8.

Costrell RM and McGee J (2017a) Cross-subsidization of teacher pension costs: the case of California. *Education Finance and Policy* forthcoming. Posted online, 16 November, 2017.

Costrell RM and McGee J (2017b) Cross-subsidization of teacher pension costs: the impact of assumed market returns. University of Arkansas EDRE, Working Paper 2017–19, October 13.

Costrell RM and Podgursky M (2008) Peaks, cliffs, and valleys. *Education Next* **8**, 22–28.

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

Costrell RM and Podgursky M (2010a) Distribution of benefits in teacher retirement systems and their implications for mobility. *Education Finance and Policy* **5**, 519–557.

Costrell RM and Podgursky M (2010b) Golden handcuffs: teachers who change jobs or move pay a high price. *Education Next* **10**, 60–66.

Fitzpatrick MD (2015) How much Are public school teachers willing to pay for their retirement benefits? *American Economic Journal: Economic Policy* **7**, 165–188.

Goldhaber DD and Holden K (2018) How much do teachers value deferred compensation? Evidence from defined contribution rate choices, RAND Education, Working Paper 1238.

Kansas Public Employees Retirement System (2017) Valuation Report as of December 31, 2016, Cavanaugh Macdonald, July 10.

Kim D (2016) Worker Retirement Responses to Pension Incentives: Do They Respond to Pension Wealth? Department of Economics, University of Missouri.

Koedel C and Podgursky M (2016) Teacher pensions. Chapter 6. In Hanushek EA, Machin SJ and Woessmann L (eds), *Handbook of the Economics of Education*, vol. **5**, Amsterdam: North-Holland Elsevier, pp. 281–303.

Lueken M (2017) (No) Money in the Bank: Which Retirement Systems Penalize New Teachers? Washington, DC: Thomas B. Fordham Institute.

Novy-Marx R and Rauh JD (2009) The liabilities and risks of state-sponsored pension plans. *Journal of Economic Perspectives* **23**, 191–210.

Rhee N and Fornia WB (2016) Are California Teachers Better off with a Pension or a 401(k)? UC Berkeley Center for Labor Research and Education.

Rhee N and Fornia WB (2017) How do California teachers fare under CalSTRS? Applying workforce tenure analysis and counterfactual benefit modeling to retirement benefit evaluation. *The Journal of Retirement* **5**, 42–65.

Richwine J and Biggs AG (2011) Assessing the Compensation of Public-School Teachers, Heritage Center for Data Analysis, November 1, 2011, Washington, D.C.

Samuelson PA (1963) Risk and uncertainty: a fallacy of large numbers. *Scientia* **98**, 108–113.

Schmitz K (2016) Breaking Down Kansas' New Retirement Plan for Teachers, Bellwether Education Partners, TeacherPensions.org, October 4, 2016.