10-24-2016

Cross-Subsidization of Teacher Pension Normal Cost: The Case of CalSTRS

Robert M. Costrell
University of Arkansas, Fayetteville, costrell@uark.edu

Josh B. McGee
Manhattan Institute

WORKING PAPER SERIES

Cross-Subsidization of Teacher Pension Normal Cost: The Case of CalSTRS

Robert M. Costrell, University of Arkansas
Josh McGee, Manhattan Institute and Laura and John Arnold Foundation

October 24, 2016; revised January 11, 2017

EDRE Working Paper 2016-17

CROSS-SUBSIDIZATION OF TEACHER PENSION NORMAL COST: THE CASE OF CALSTRS

Robert M. Costrell, Professor of Education Reform and Economics, University of Arkansas
Josh McGee, Senior Fellow, Manhattan Institute;
and Vice President of Public Accountability, Laura and John Arnold Foundation
October 24, 2016, revised January 11, 2017

Corresponding author:
Robert M. Costrell
200 Graduate Education Building
University of Arkansas, Fayetteville, AR 72701
479-575-5332, costrell@uark.edu

ABSTRACT: Under traditional defined benefit pension plans, annual contributions are set at a uniform percentage of pay to fund accruing benefits. That normal cost rate masks wide variation in the cost of individual benefits, generating an extensive and non-transparent pattern of cross-subsidization. We provide a comprehensive analysis of cross-subsidies in employer contributions across all entry and exit ages. The gains and losses of winners and losers must add up to zero, and we explain why they do not in some previous work, which claims that nearly all teachers are winners in the California Teachers Retirement System. To the contrary, we find about two-thirds of all entrants are losers. The losers earn benefits with an average annual employer cost of 0.8 percent of pay, vs. 5.7 percent for the winners. In effect, this is a system of widely varying, but non-transparent employer matches to the employee contribution, unlike a retirement account plan with a uniform match.

KEYWORDS: teacher pensions

JEL CODE: H75

ACKNOWLEDGMENTS: The authors thank CalSTRS and Flick 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. McGee did not write this report in his official capacity at 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.

CROSS-SUBSIDIZATION OF TEACHER PENSION NORMAL COST: THE CASE OF CALSTRS

I. INTRODUCTION AND SUMMARY

Under traditional defined benefit (DB) plans, employees and employers make uniform annual contributions at a given percentage of pay to fund benefits as they accrue for the workforce. That normal cost rate embeds within it individual cost rates that are not explicitly calculated, but vary widely, by age of entry and exit. This generates an extensive pattern of cross-subsidization, as the contribution rate deviates – positively or negatively – from the individual cost rates. Moreover, these cross-subsidies are unknown to the individual, as the uniform contribution masks the variation in individual costs. In both respects – cross-subsidization and non-transparency – this differs from retirement account plans (defined contribution or cash balance, actual or ideal) where the contribution rate for each individual equals the cost of his or her benefit. This paper develops the methodology to render transparent the individual costs and patterns of cross-subsidization for DB plans and illustrates with plan data from the California State Teachers' Retirement System (CalSTRS).

Consider the standard actuarial calculation of the normal cost rate: entry age normal, or EAN. This method is designed to determine the contribution rate (employer plus employee) which, collected over the careers of all members of any entering cohort, would cover benefits earned by that cohort. Since individuals enter and exit at varying ages, there are two types of averaging that take place: across ages of exit and ages of entry. Although each individual's normal cost rate is implicitly calculated, based on age of entry and exit, these rates are not reported – it is only the overall average that is reported and set as the common contribution rate.

Each individual is thus led to believe that a uniform contribution is made on his or her behalf by the employer (in addition to the employee's own contribution). However, those contributions are heavily redistributed, with some employees receiving no benefit from these contributions and others receiving benefits that cost far more than was contributed on their behalf. In fact, for many individuals, the employer portion of their normal cost is net negative. More generally, there is an extensive pattern of cross-subsidization of the employer contribution between employees of different entry and exit ages. It is that cross-subsidization that funds retirement benefits for individuals whose joint contributions would otherwise be insufficient.

This paper is not addressed to the unfunded liabilities that plague so many systems – what might be termed intergenerational cross-subsidies (Backes, Goldhaber, Grout, Koedel, Ni, Podgursky, Xiang and Xu, 2016). Rather, we take as given the actuarial assumptions to analyze the cross-subsidies that are built into the system, when those assumptions are met.

It is well-established that traditional teacher pension systems generate winners and losers, as the value of individual benefits exceed or fall short of the contributions made on their behalf by substantial amounts (Costrell and Podgursky, 2010a,b; Costrell and McGee, 2010; McGee and Winters, 2013). However, this view has been challenged in an analysis of CalSTRS that finds nearly all active teachers are winners (Rhee and Fornia, 2016). The main feature of this critique is a shift in the analysis from winners and losers among *entrants* to *active* teachers, leaving aside the losses of those who have left.[1] The problem is that without the early leavers, the winners cannot be funded: the gains and losses do not add up.

The methodological contributions of the present paper are: (i) to supplement previous analyses of winners and losers based on the metric of pension wealth (a stock variable) with a

[1] "Winners and losers" are defined by Rhee and Fornia relative to a CB plan described below. They also (rightly) emphasize the full distribution of entry ages (rather than, say, age 25), in a state with many older entrants.

more accessible metric – annual normal cost rates (a flow variable); (ii) to comprehensively analyze cross-subsidies across all entry and exit ages; and (iii) to clarify that gains and losses of winners and losers must offset one another, and this requires a focus on entrants (including early leavers), rather than actives. Specifically, we show that gains and losses measured in pension wealth must be weighted by shares of the entering cohort to add up to zero, while those measured in normal cost rates, must be weighted by entrants' shares of the cohort's lifetime earnings.

Empirically, we offer new results for CalSTRS. There are features of the benefit formula that mute the cross-subsidies compared to other teacher pension plans. Even so, we find the cross-subsidies to be widespread and substantial. Over all entry and exit ages, the losers (net negative cross-subsidy) comprise 66 percent of entrants, and they account for 45 percent of entrants' lifetime earnings. The cost rate to fund their benefits averages 2.7 percentage points below the common contribution rate (i.e., the average normal cost), while that of the winners is 2.2 points above. Stated differently, the losers earn benefits with an average annual employer cost of 0.8 percent of pay (and for many of them the rate is actually negative), vs. 5.7 percent for the winners. In effect, this is a system of widely varying, but non-transparent employer matches to the employee contribution, unlike a retirement account plan with a uniform match.

II. NORMAL COST RATES AND CROSS-SUBSIDIZATION

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 (from which the average normal cost rate is constructed below). We calculate this rate to generate a stream of contributions, from entry to exit, sufficient to fund the individual's benefits. It is straightforward to see that the rate must equal the ratio of the present value of benefits to the present value of earnings, where both are evaluated at (or back-dated to) entry. Formally, if we denote B_{es} as the present value of benefits for an individual of type (e,s), and similarly W_{es} as the present value of earnings, then the individual normal cost rate is $n_{es} = B_{es}/W_{es}$. *This is the rate that, applied to that individual's annual earnings over his or her career, would prefund his or her benefits.*[2] These cost rates are not explicitly calculated under the EAN method, but they are implicit in the EAN algorithm to calculate the uniform contribution rate, as shown below.

Individual cost rates, n_{es} vary widely with e and s. 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 benefit accruals under traditional DB formulas (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 present value, B_{es}, is lower, but so is that of their earnings, W_{es}. For CalSTRS, as discussed below, the latter outweighs the former, so the normal cost rate, B_{es}/W_{es}, is higher for later entrants.

Normal Cost Rates for Entry-Age Cohort

The EAN algorithm does not actually start with cost rates for specific individuals (since one's exit age is not pre-determined), but rather with the entry cohort of any given age, e. The first step in the algorithm is to calculate the contribution rate that would fund the benefits of that entry-age cohort. Members of that cohort will exit at different ages, so their cost rate, call it n_e, will, in effect, be an average of n_{es} across the range of s. Formally, let $p_{s|e}$ denote the probability

[2] Note that n_{es} can be calculated solely from "provisions of the plan [and] the actuarial assumptions employed in projection of benefits and present value determinations" (Lamenzo, 2008, p. 35).

of exit at age s, given 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):

(1) $n_e = \sum_s p_{s|e} B_{es} / \sum_s p_{s|e} W_{es} = \sum_s n_{es} p_{s|e} W_{es} / \sum_s p_{s|e} W_{es} = \sum_s n_{es} \left(p_{s|e} W_{es} / \sum_s p_{s|e} W_{es} \right),$

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 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 present value of earnings for each exit age s among entrants of age e, ($p_{s|e}W_{es}/\sum_s p_{s|e}W_{es}$).

<u>Overall Normal Cost Rate, Across Entry Ages</u>

The second step in determining the uniform normal cost rate applied to all employees, call it n*, is to average the entry-age-specific normal costs, n_e, across all entry ages. The weights we use here are, again, shares of the entering cohort's present value of earnings. These are the weights that make the cross-subsidies add up to zero, so they are appropriate for this paper's focus on the pure cross-subsidies.[3] The uniform normal cost rate, using these weights, is:

(2) $n^* = \sum_e n_e (p_e W_e / \sum_e p_e W_e),$

where p_e is the share of each entry cohort of age e and $W_e = \sum_s p_{s|e} W_{es}$ is the present value of earnings for the average entrant of age e. Using (1) to take (2) a step further, we have:

(3) $n^* = \sum_e \sum_s n_{es} \left(p_e p_{s|e} W_{es} \right) / (\sum_e \sum_s p_e p_{s|e} W_{es}) = \sum_e \sum_s n_{es} (p_{es} W_{es}) / (\sum_e \sum_s p_{es} W_{es}),$

This shows n* is the weighted average of individual normal cost rates n_{es} across both entry ages e and separation ages s, where $p_{es} = p_e p_{s|e}$ is the share of type (e,s) in the entering cohort, and the

[3] These are not the weights used in the standard algorithm. That algorithm applies n_e to each individual's pay in the current workforce and divides by the current payroll. We would argue, however, that the weights used here are more consistent with the first step of the EAN method (determining entry-age-specific normal costs) and with the method's overall intended purpose: to generate contributions that will fund the benefits of each entering cohort.

weights on n_{es} are that type's share of the cohort's present value of earnings. Substituting $n_{es} = B_{es}/W_{es}$ into (3) shows that n* equals the ratio of the present value of benefits to that of earnings: the same relationship for the individual, the entry-age cohort, and the entry cohort as a whole.

<u>Cross-Subsidies, Weights, and Adding Up</u>

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

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

Thus, when cross-subsidies are measured in normal cost rates, the weights for adding up are the shares of type (e,s) in the present value of earnings among all entrants, ($p_{es}W_{es}/\sum_e\sum_s p_{es}W_{es}$).

We can also measure the cross-subsidies in present values, $(n_{es} - n^*)W_{es} = (B_{es} - n^*W_{es})$, i.e., the difference between the present value of benefits for an individual of type (e,s) and the contributions on his or her behalf. Expression (4) implies

(5) $$0 = \sum_e \sum_s (n_{es} - n^*)(p_{es} W_{es}) = \sum_e \sum_s (B_{es} - n^* W_{es}) p_{es}.$$

Thus, when the cross-subsidies are measured in present value terms, the weights for adding up are simply the shares of entrants of type (e,s), p_{es}.

To summarize, the zero-sum result holds in two forms: (i) for cross-subsidies in present value terms, applying weights equal to the shares of the entering cohort; and (ii) for cross-subsidies in normal cost rates, applying weights equal to the shares of the entering cohort's lifetime earnings. Both sets of weights are pertinent for constructing counts of winners and losers. It is important to stress, however, that in both cases the unit of analysis is the entering cohort – *not* the active workforce. That is the only way for gains and losses to add up properly.

It will also be useful to note that we can decompose individual cross-subsidy rates into those within and between entry-age cohorts: $(n_{es} - n^*) = (n_{es} - n_e) + (n_e - n^*)$. Finally, the cross-subsidy rate $(n_{es} - n^*)$ and its decomposition are identical for the total rate and the employer rate: netting a uniform employee rate out of both terms leaves the cross-subsidy $(n_{es} - n^*)$ invariant.

III. Individual Normal Cost Rates for CalSTRS

We consider the case of CalSTRS. First we estimate the (triangular) matrix of individual normal cost rates, $n_{es} = B_{es}/W_{es}$, for e,s = 20, … , 75. Since B_{es} and W_{es} are proportional to the entry wage, n_{es} is independent of it, so we can normalize B_{es} and W_{es} per dollar of entry wage. Thus normalized, the equation for W_{es} is:

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

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

Benefits can be in the form of an annuity or a refund,[4] so:

(7) $B_{es} = (prob\ refund)_{e,s-e} PV(Refund_{es}) + \left[1 - (prob\ refund)_{e,s-e}\right] PV(Annuity_{es}).$

CalSTRS assumptions give the probability of refund by entry age e and service (s-e). Refunds of employee contributions (currently 9.205 percent of pay) include interest currently at 4.5 percent. Thus, $PV(Refund_{es})$, per dollar of entry wage, is a straightforward calculation.

The starting annuity for an individual of type (e,s) is calculated under the defined benefit program for new hires (since 2013), the “2% at 62” program. This is a final-average-salary (3

[4] We leave aside disability and death benefits, which comprise about 5 percent of normal cost, less than 1 point.

years) formula with age-specific multipliers ranging from 1.16 percent at age 55 to 2.0 percent at age 62 and 2.4 percent at 65, after 5-year vesting. Thus, a 25-year-old entrant who works to 65 will retire at 40 × 2.4 = 96 percent of final average salary. Vested employees who withdraw before age 55 but do not cash out are assumed by CalSTRS to defer the pension to age 60 (when the multiplier is 1.76 percent). So a 25-year-old entrant exiting at age 50 will receive, at age 60 a pension of 25 × 1.76 = 44 percent of her final average salary (frozen 10 years earlier).

Annuity factors (including 2.0 percent simple COLA) are calculated using CalSTRS mortality assumptions for active female members, discounted to entry. The annuity factor times the starting annuity gives $PV(Annuity_{es})$ and, with $PV(Refund_{es})$, we have B_{es}, as in (7). Using W_{es} from (6), we have $n_{es} = B_{es}/W_{es}$. These calculations 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 Within and Across Entry-Age Cohorts

Figure 1 depicts normal cost rates n_{es} for selected entry ages, e = 25, 30, 35, 40, and 45,[5] and all exit ages s. The graduated multiplier mutes the variation in this system, compared to a uniform multiplier (e.g. 2 percent). Still, the variation is substantial. The lowest rate depicted is 4.4 percent, at the bottom of the lower curve, for 25-year-old entrants exiting at age 30. The highest rate is 20.4 percent, at the top of the upper curve, for 45-year-old entrants, exiting at 65.

Consider the pattern of normal cost rates *within* entry-age cohorts. This pattern is depicted in Figure 1 *along* each curve, for any given entry age, as the exit age varies. Prior to vesting, the benefit is the refund of employee contributions. The normal cost rate, therefore,

[5] These are representative of the vast majority of entrants. The proportions at or below ages 25, 30, 35, 40, and 45 are, respectively, 29%, 57%, 72%, 81%, and 88%. The full distribution, used here and below, is shown in Figure 3 of Rhee and Fornia (2016). The underlying data from CalSTRS were kindly provided to the authors by Fornia.

starts at the employee contribution of 9.205 percent: each curve in Figure 1 begins at the horizontal line representing that rate. The normal cost rate then gently declines over the vesting period, 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. The contribution rate needed to cover the refund falls as this difference accumulates, but not by much over these 5 years.

Then, at 5 years, the normal cost rate drops precipitously, as depicted for each curve in Figure 1. The reason is that one is now eligible for a deferred pension with a low normal cost, and CalSTRS assumes a sizeable probability of eschewing the refund, which is worth a lot more. Figure 2 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 for exit immediately upon vesting, since the final average salary is frozen and the pension is deferred until age 60. The solid curve – the actual normal cost – is a weighted average of the pension and refund, using the probability of taking the refund. Over the vesting period, 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.[6]

From vesting up to age 55, the normal cost rate for the pension rises as the deferral to age-60 becomes shorter, and it continues to rise beyond 55 as the age-specific multiplier grows to 65.[7] Each year of delayed retirement beyond 55 reduces the number of years to receive the pension. Prior to age 65, the growth in the multiplier outweighs this effect, but after age 65 the

[6] As Figure 1 shows, the normal cost rate for 25-year-old entrants is reduced by those who decline refunds up to age 47, and by those who continue to take refunds after age 47. (These calculations do not reflect the probability of re-entry.) See Lueken (2014) for the first analysis of sub-optimal cash-out decisions, by Illinois teachers.

[7] Interestingly, there is no discontinuity at age 55, when the pension draw becomes immediate, instead of the 6-year deferral assumed for exits at age 54. That is because the age-specific multiplier is higher for age 60 than for 55, a feature that smoothes out the accrual pattern for CalSTRS, as noted previously by Costrell and Podgursky (2009).

multiplier stops growing, and the normal cost declines. This pattern, before and after 65, is reflected in Figure 1 along each curve, corresponding to any given entry age.

In addition to the variation *within* entry-age cohorts, Figure 1 also depicts the (vertical) variation *across* entry ages. As stated earlier, this can go either way. To see why, write the normal cost rate for the pension as (annuity factor at date of exit) × (annuity/final average salary) × (final average salary/cumulative earnings at exit). At any given exit age, the annuity factor is independent of entry age. But for older entrants (with shorter service), the annuity is a lower percent of final average salary, while their final average salary is higher relative to cumulative earnings (since it is a shorter stream). In the case of CalSTRS, as reflected in Figure 1, the latter outweighs the former, so the normal cost curves are generally higher for later entrants.

Individual Employer Normal Cost

The individual normal cost depicted in Figure 1 includes the employee contribution, 9.2 percent. The individual's *employer* normal cost rate, (n_{es} – 9.2 percent), is the distance, positive or negative, from each curve to the employee contribution line. Thus, the employer normal cost rate varies from 11.2 percent, for (e, s) = (45, 65), to -4.8 percent, for (e, s) = (25, 30).

As explained above, the normal cost rate dips below the employee contribution rate throughout the vesting period, and for most entry ages, it drops further immediately upon vesting, before rising again. Consequently, for each entry age there is a range of exit ages over which the employer normal cost is negative. These ranges can be read off Figure 1. For entry age 25, the employer cost is negative up through exit age 53. The corresponding break-even points for entry ages 30, 35, 40, and 45 occur, respectively, after exit ages 51, 50, 50, and 50.

What proportion of entrants fall into these ranges of negative employer normal cost?[8] Figure 3 shows that half or more of every entry cohort aged 38 or below will exit before the employer normal cost turns positive, and for older entrants, the proportion is generally 40 percent or more. Overall, we estimate 51 percent of all entrants receive benefits that cost less than the employee contribution. These entrants, however, are early leavers, so they account for a much smaller proportion of the cohort's lifetime earnings, as shown in the bottom curve of Figure 3.

IV. WINNERS AND LOSERS WITHIN AND ACROSS ENTRY-AGE COHORTS

The wide variation among individual cost rates contrasts with the uniform rate contributed by and/or on behalf of all individuals. We calculate the weighted average normal cost rate across all ages of entry and exit, n*, as given in (3). As discussed above, we use present value of earnings as weights. This is the normal cost rate that will fund the benefits of the entire cohort if all assumptions are met. We calculate the overall normal cost rate n* to be 12.8 percent, depicted in Figure 1 as the dashed line.[9] The employer normal cost rate, 3.5 percent, is represented by the distance between that line and the solid line.

The winners and losers from this system (those whose benefits cost more or less than the uniform contribution rate), are those who lie above or below the n* line in Figure 1. For example, age-25 entrants who exit prior to age 61 (or after 67) are losers. They comprise 85 percent of that entry age cohort and account for 71 percent of their lifetime earnings. These individuals may be said to provide cross-subsidies to others in their entry age cohort and/or those

[8] For each entry age, the distribution of exit ages is found by constructing the "survival" rates (i.e., probability of still working), from CalSTRS' assumed exit rates for retirement or withdrawal, and then differencing that series.
[9] This is about 1.7 percentage points below CalSTRS' estimate of the normal cost rate for new entrants (netting out death and disability). This is partly due to our assumption that all entrants earn the same starting wage, regardless of age. Relaxing that assumption does not substantially alter our main findings on the cross-subsidies.

of other entry ages. Conversely, entrants older than 46 are more likely to receive cross-subsidies than to provide them. Figure 4 depicts the proportion of losers for each entry age (by number of entrants and by their lifetime earnings). Averaging across all entry and exit ages, the losers comprise 66 percent of all entrants, and they account for 45 percent of their lifetime earnings.

Comparison With Canonical Cash Balance Plan

We have defined winners and losers by how the annual cost rate to fund an individual's benefit compares with the uniform rate charged for all entrants. The partition between winners and losers is the same if we compare present values instead of annual cost rates. The present value partition is also independent of its evaluation as of date of entry or exit. For the latter, one compares the value of an individual's benefits upon exit with cumulative contributions at the uniform normal cost rate, including interest at the assumed discount rate. This is equivalent to comparing the DB plan with a canonical, fiscally neutral cash balance (CB) plan.

In a CB plan, each individual's terminal benefit is given by a (notional) retirement account balance (to be annuitized or drawn down) equal to the cumulative value of contribution credits. For the canonical, fiscally neutral plan, the contribution credit rate is uniform, equal to the DB plan's contribution rate, and interest is credited at the DB plan's assumed discount rate.[10] In such a plan, the normal cost rate is identical for each individual, regardless of age of entry or exit – it is the uniform contribution rate – so there are no cross-subsidies. The winners and losers in the DB plan vs. this CB plan are identified by comparing the cost (annual rate or present value) of individual benefits in the two plans. This partition is the same as comparing the

[10] In practice, cash balance plans can and do differ in both ways: (i) contribution rates can vary with years of service, as in Kansas' Tier 3 plan for teachers (Schmitz, 2016); and (ii) the guaranteed interest credit is typically below the actuarial discount rate.

individual cost and uniform contributions within the DB plan. Thus, in comparing the CalSTRS plan with such a CB plan (12.8 percent contribution credit and 7.5 percent interest credit), we have the same results as depicted in Figure 4: 66 percent of all entrants (with 45 percent of the cohort's lifetime earnings) would do better under the CB plan.

A Potentially Misleading Measure of Winners and Losers

The picture of winners and losers under DB that has emerged from previous studies, consistent with our findings, has been challenged by Rhee and Fornia (2016). In the case of CalSTRS, we find 66 percent of all entrants are losers, while Rhee and Fornia (p. 31) find only 21 percent of "active teachers" are worse off, compared to a hypothetical CB plan. Why is there such a difference? One reason is that their CB plan is not fiscally neutral – it assumes interest credits of 7 percent vs. the 7.5 percent CalSTRS discount rate. The main difference, however, is their focus on "active teachers." Entrants who exit early are under-represented among actives, a result known as "survivorship bias." For example, Rhee and Fornia find that although "40% of new hires leave before vesting, these leavers represent just 6% of teaching positions" (p. 5).

Figure 5 depicts our estimate of the distributions of exits among new entrants and actives, for those who enter at age 30.[11] This entry age is roughly representative, as their normal cost rate of 12.6 percent is close to the overall rate, n* = 12.8 percent. Figure 5 illustrates the point quoted above: the proportion of early leavers among the entering cohort vastly exceeds that of the active population. For example, 15 percent of new 30-year-old entrants will leave in the first

[11] For the distribution among actives, the survival rate series for 30-year-old entrants, described earlier, is brought forward one year at a time to construct the conditional survival rates beyond age x for those with x-30 years of service, up through x = 75. These are then differenced to construct the probability distribution of exit ages for those of age x and service x-30. Finally, for any given exit age these probabilities are weighted across all ages x by the distribution of those of age x and service x-30 in CalSTRS' full age-and-service table for 2014. This distribution is virtually identical to that in the steady-state where the number of 30-year-old entrants is constant from year to year.

year, but those new entrants comprise only 6 percent of all actives who entered at age 30, so only 0.9 percent of them (0.15 × 6) leave at age 30. Conversely, as Figure 5 shows, the distribution of actives is concentrated among those who will exit at age 55 or higher, especially around age 60.

The difference between the two distributions has a dramatic effect on the count of winners and losers. Figure 1 shows that 30-year-old entrants who exit before age 58 (or after age 69) are losers. From the distributions in Figure 5, we find that these losers constitute 64 percent of new entrants, but only 32 percent of actives. This result is illustrative for the whole array of entry ages, and explains the vast majority of the contrast between Rhee-Fornia's findings and ours. Figure 6 shows Rhee-Fornia's comparison of the value of benefits under CalSTRS DB vs. their hypothetical CB plan, for 30-year-old entrants. In their calculation, the losers are those who exit before age 55. We find these comprise 54 percent of new entrants, but only 16 percent of actives. Adjusting their CB calculation to 7.5 percent interest from 7.0, the DB losers comprise 58 percent of entrants and 23 percent of actives. Any way one looks at it, Rhee-Fornia's focus on actives instead of entrants drastically reduces the count of losers from DB vs. CB.

Their study concludes that "focusing on new-hire attrition is misleading" (p. 5) since 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 entrants) can leave a misleading (or at least incomplete) *financial* picture. Focusing on actives omits the cross-subsidy left behind by early leavers to help fund benefits of those who stay. Consequently, while the gains and losses *among entrants* must add up to zero (as shown in (4)-(5)), the gains outweigh the losses *among actives*.

For example, among 30-year-old new *entrants*, we calculate that about 32 percent of the cohort's employer contributions (and net pension wealth) is transferred from losers to winners.

The same result holds for a series of 30-year-old entrants, running from current entrants (all of whom are active), back to the most distant 30-year-old entrants of whom any remain active (i.e. almost all have left).[12] We partition this series of entrants into those who have left and those who remain active, and further partition them into winners and losers, in order to decompose the gains and losses among the subgroups. We find that the winners' gains, 32.1 percent of net pension wealth for the series of entrants, breaks down into 22.8 percent for those who are still active and 9.3 percent for those who have left. Conversely, of the losers' total losses of 32.1 percent, only 8.8 percent will be incurred by those currently active and 23.3 percent by those who have left. Taken as a whole, the gains and losses add up (22.8 + 9.3 – 8.8 – 23.3 = 0.0), but if we just look at the actives, they do not. The actives' net gain of 22.8 – 8.8 = 14.0 percent of net pension wealth represents the cross-subsidy left behind by the inactives: 9.3 – 23.3 = –14.0. It is that cross-subsidy that explains why the vast majority of actives will be winners.

V. HOW LARGE ARE THE GAINS AND LOSSES?

We return now to the proper unit of analysis, the entering cohort. Having previously considered the count of winners and losers (Figure 4), we now consider the magnitude of their gain and loss, as measured in normal cost rates. For individual types (e,s) these can be read off Figure 1, as the distance, positive or negative, between the normal cost rate n_{es} and the dashed line for overall n*. Thus, the cross-subsidy rate varies from -8.4 points, for (e, s) = (25, 30), to +7.6 percentage points, for (e, s) = (45, 65). Over all entry and exit ages, the cross-subsidy rate provided by losers (66 percent of all entrants) is -2.7 percentage points, and the winners receive +2.2 percentage points (these add up to zero, when weighted by their shares of lifetime earnings).

[12] We assume each cohort is of the same size. The resulting steady-state closely matches the observed age-service matrix, as stated in the note above.

As noted earlier, these are the cross-subsidy rates for both the joint normal cost (employer plus employee) and the employer cost alone. We consider the employer cost rate. The losers' cross-subsidy means they receive benefits with an average employer cost rate of 0.8 percent, instead of the overall average of 3.5 percent. Conversely, the winners' benefits carry an average employer cost of 5.7 percent. By analogy with account-based retirement systems, this is as if employers' match to the employees' contribution averaged 0.8 percentage points for two-thirds of entrants and 5.7 percentage points for one-third. Figure 7 breaks out the employer normal cost rates for winners and losers by entry age. The gap is quite striking.

Decomposing the Cross-Subsidies Within and Across Entry-Age Cohorts

The cross-subsidy rates can be decomposed into those within and across entry-age cohorts: $(n_{es} - n^*) = (n_{es} - n_e) + (n_e - n^*)$. If we net out the 9.2 percent employee contribution from n_{es}, n_e, and n^*, the decomposition is unchanged, but now represents employer costs. Figure 8 presents the employer cost by entry age, (n_e – 9.2 percent), using (1). As is typical, the rates rise, up through the most common entry ages.[13] The difference between these normal cost rates and the overall rate, given by the dashed line, represents the cross-subsidy rates *across* entry age cohorts, $(n_e - n^*)$. As can be seen, entry cohorts of age 30 or less cross-subsidize older entrants.

Within each entry-age cohort, the cross-subsidies are $(n_{es} - n_e)$. These are positive or negative, as individuals enjoy a cross-subsidy from others in the cohort, or vice versa. Sorting those of entry age e by whether n_{es} is greater or less than n^*, and aggregating $(n_{es} - n_e)$ gives the average *within* entry-age component of the total cross-subsidy $(n_{es} - n^*)$ for winners and losers.

[13] Individual normal costs generally rise with entry age, *for any given exit age*, for reasons given in the discussion of Figure 1. In Figure 8, each entry age represents a different distribution of exit ages, generally rising with later entry. Thus, an additional factor in the rise in n_e, with entry age is the *greater weight on later exits*, with higher n_{es}.

Figure 9 pulls this all together, depicting the decomposition of the winners' and losers' cross-subsidies into those within and across entry-age cohorts. The solid curves above and below the 0% line represent the cross-subsidy rates for winners and losers, by entry age.[14] The dotted curves net out the cross-subsidies across entry ages, (n_e - n*). That is, the gap between the dotted and solid curves represent the cross-subsidies *across* entry-age cohorts (the same gap as in Figure 8), so the dotted curves in Figure 9 represent the cross-subsidies *within* entry-age cohorts.

For example, the winners of entry-age-35 receive a cross-subsidy of +2.7 percentage points (solid blue line), of which +1.5 points come from within that entry-age cohort (dotted blue line) and +1.2 points come from other entry-ages (the gap between the lines). The losers in that cohort also gain +1.2 points from other entry ages, but it is more than offset by the cross-subsidy they provide within the cohort (-4.7 points) for a net cross-subsidy of -3.6 points. Conversely, for entry-age-25, the losers provide a cross-subsidy of -1.5 points to other entry-ages, compounding the cross-subsidy of -0.8 points within their age-cohort, for a total cross-subsidy, within and across, of -2.3 points. In general, consistent with Figure 8, the cross-subsidies across entry-age cohorts exacerbate the losses (or reduce the gains) for entrants age 30 or below, and ameliorate the losses (or enhance the gains) for older entrants. Over all entry ages, however, for both winners and losers most of the cross-subsidy (72%) comes from within the entry-age cohorts: +1.6 points out of +2.2 for the winners and -2.0 points out of -2.7 for the losers.

[14] These curves reproduce from Figure 7 the gaps between winners'/losers' normal cost and the overall average.

VI. Conclusion

The distinguishing characteristic of defined benefit pension plans is that the benefit is defined by formula, rather than contributions, unlike account-based plans. With benefits de-linked from contributions, some individuals receive benefits that cost more than the contributions made on their behalf and some receive less, effectively a system of hidden cross-subsidies. Previous work on teacher pensions has analyzed the redistribution of benefits for any given entry age, with the metric of pension wealth. The present paper adds to previous literature by analyzing cross-subsidies both within and across entry-age cohorts, and measuring them in normal cost rates – comparable with annual contribution rates for account-based plans.

We have also clarified a potential confusion in the literature, regarding the number of winners and losers between a traditional DB plan and a fiscally neutral CB plan. 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 the entry cohort. Thus, many of those entrants who leave service without earning benefits commensurate with the contributions made on their behalf are not counted, because they are no longer teaching. This approach concludes that the vast majority of active teachers are winners under the traditional DB plan, compared to a fiscally neutral CB plan.

The problem is that 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 those who are no longer in the plan. To understand how the system adds up, we must focus on entrants. The winning and losing proportions among entering cohorts can be measured by their numbers and/or by their lifetime earnings. If the measure of gains/losses is pension wealth, then the numbers of winners and losers provide the zero-sum result; if the measure is annual normal cost rates, then it is their shares of lifetime earnings that make the cross-subsidies add up.

We have illustrated the analysis with CalSTRS plan data, portraying the full pattern of cross-subsidies by age of entry and exit. Overall, we find that 66% of all entrants, accounting for 45% of entrants' lifetime earnings, receive benefits with an average employer normal cost rate of only 0.8%, compared to 5.7% for the winners. These are average cross-subsidies of -2.7 percentage points from the losers and +2.2 points for the winners. Most of these cross-subsidies are within entry-age cohorts. None of these cross-subsidies would exist at all under an account-based retirement system, such as a CB plan with uniform employer contribution rates. Even if employer contributions vary, as they do in some CB plans (e.g. with years of service), the variation is at least transparent, unlike the hidden cross-subsidies in traditional DB plans.

References

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*, August/September 2016 (Vol. 45, No. 6), 367-377.

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*, Spring 2009 (Vol. 4, No. 2), 175-211.

Costrell, Robert M. and Michael Podgursky, 2010a. "Distribution of Benefits in Teacher Retirement Systems and their Implications for Mobility," *Education Finance and Policy*, Fall 2010 (Vol. 5, no. 4), 519-557.

Costrell, Robert M. and Michael Podgursky, 2010b. "Golden Handcuffs: Teachers who change jobs or move pay a high price," *Education Next*, Winter 2010 (Vol. 10, no. 1), pp. 60-66.

Costrell, Robert M. and Josh McGee, 2010. "Teacher Pension Incentives, Retirement Behavior, and Potential for Reform in Arkansas," *Education Finance and Policy*, Fall 2010 (Vol. 5, no. 4), 492-518.

Lamenzo, James R. 2008. "Understanding an Actuarial Valuation Report," Commonwealth of Massachusetts, Public Employee Retirement Administration Commission.

Lueken, Martin. 2014. "Cash on the Table? A Behavioral Analysis of Refund Claimants and Annuitants in the Illinois Teachers' Retirement System," Ph.D. Dissertation, University of Arkansas.

McGee, Josh and Marcus A. Winters, 2013, "Better Pay, Fairer Pensions: Reforming Teacher Compensation," Civic Report No. 79 (September), Center for State and Local Leadership at the Manhattan Institute, New York.

Rhee, Nari and William B. Fornia, 2016, "Are California Teachers Better off with a Pension or a 401(k)?" UC Berkeley Center for Labor Research and Education.

Schmitz, Kirsten, 2016, "Breaking Down Kansas' New Retirement Plan for Teachers," Bellwether Education Partners, Teacher Pensions Blog, October 4, 2016.

Figure 1. Normal Cost Rate, by Entry Age and Age of Exit, n_{es}

Estimated for CalSTRS New Entrants, using CalSTRS assumptions

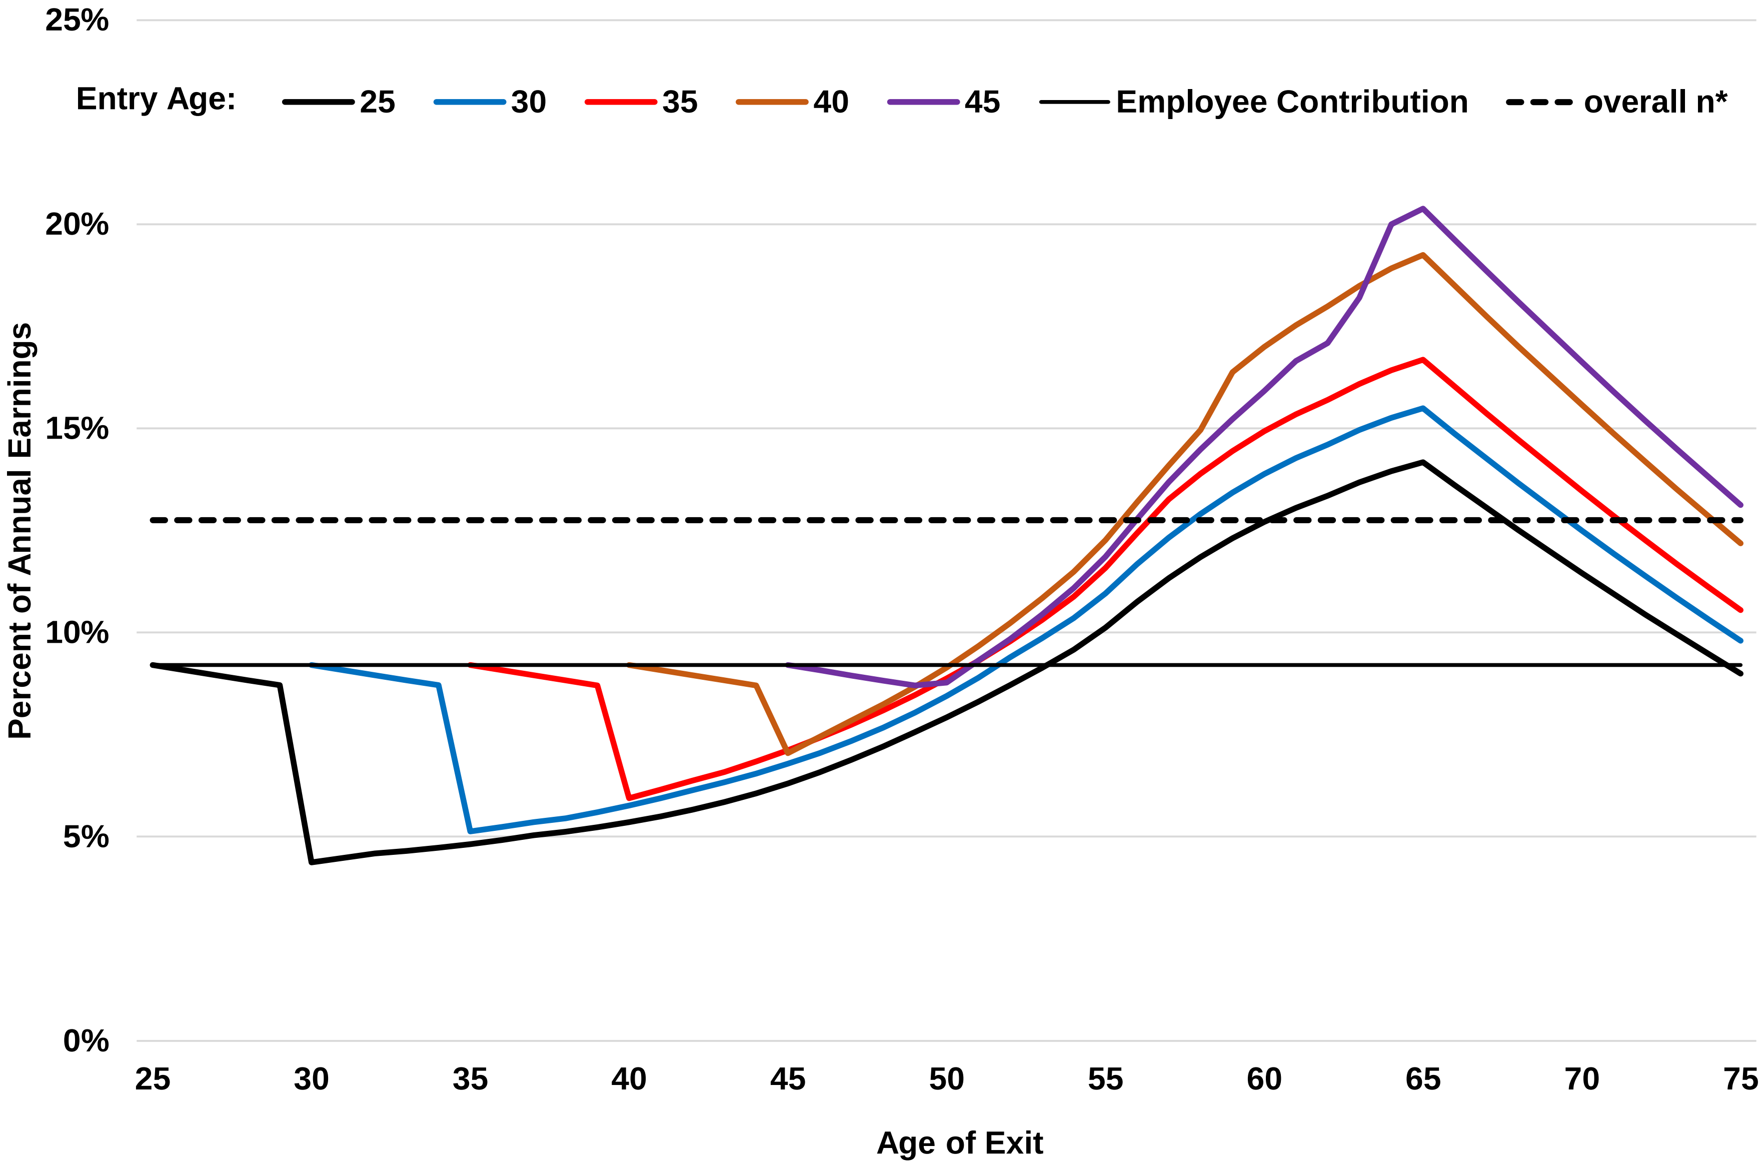

The curves depict n_{es}, the annual contribution rate required to fund benefits of an individual entering at age e and exiting at age s.

Figure 2. Normal Cost Rate, Entry Age 25, Pension and Refunds

Estimated for CalSTRS New Entrants, using CalSTRS assumptions

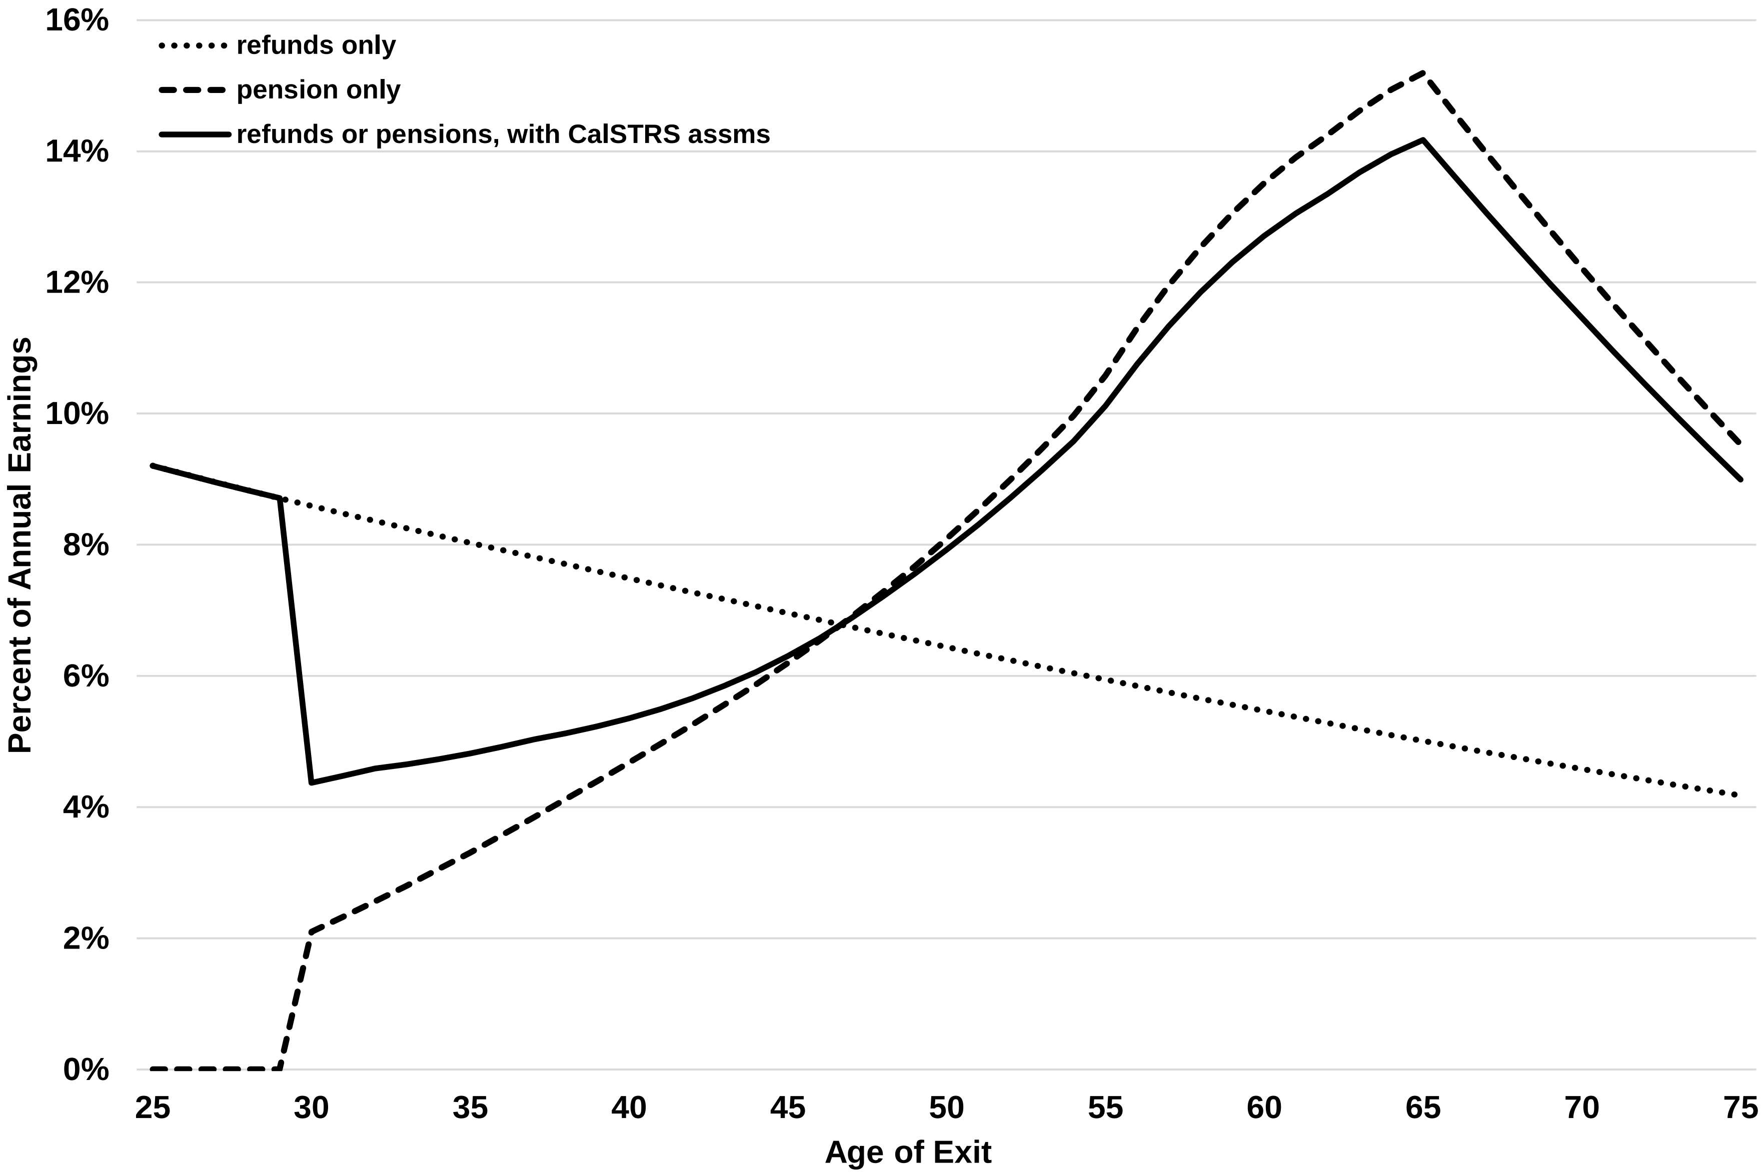

Figure 3. Proportion of Entrants with Negative Employer Normal Cost

n_{es} < employee contribution rate

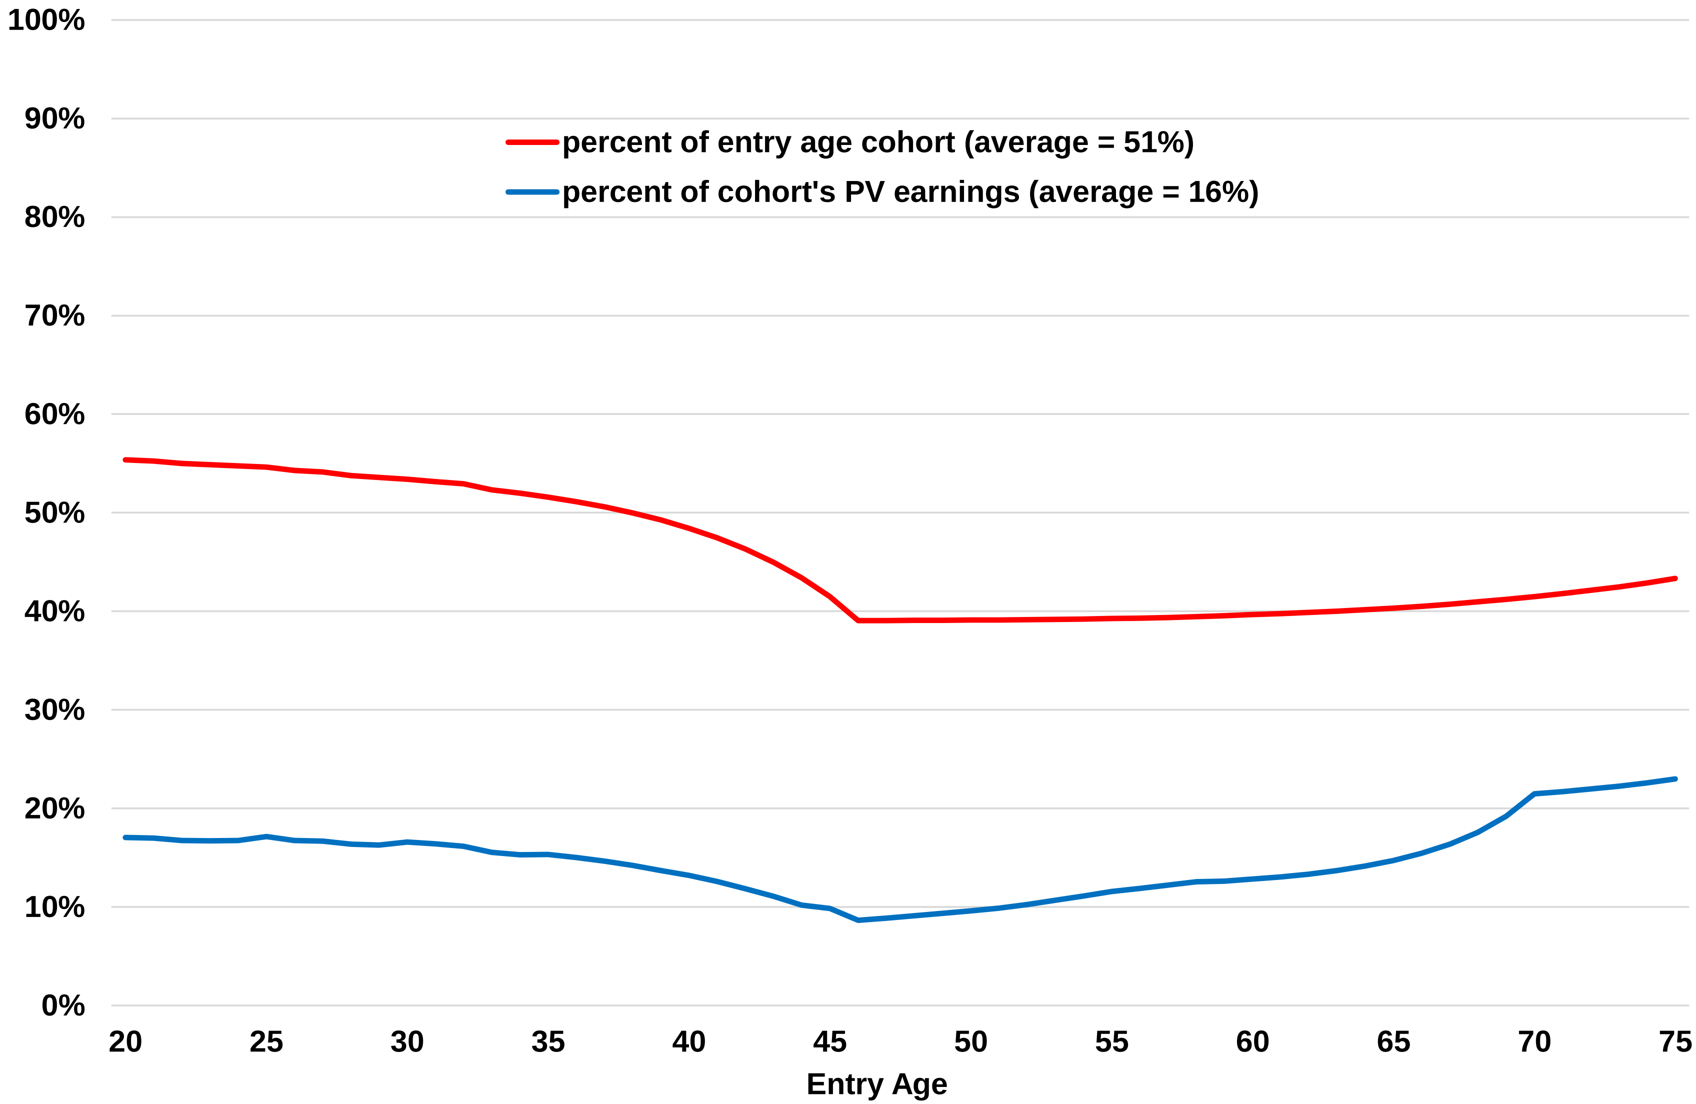

Figure 4. Proportion of Entrants with Normal Cost Below Overall Rate

$n_{es} < n^*$

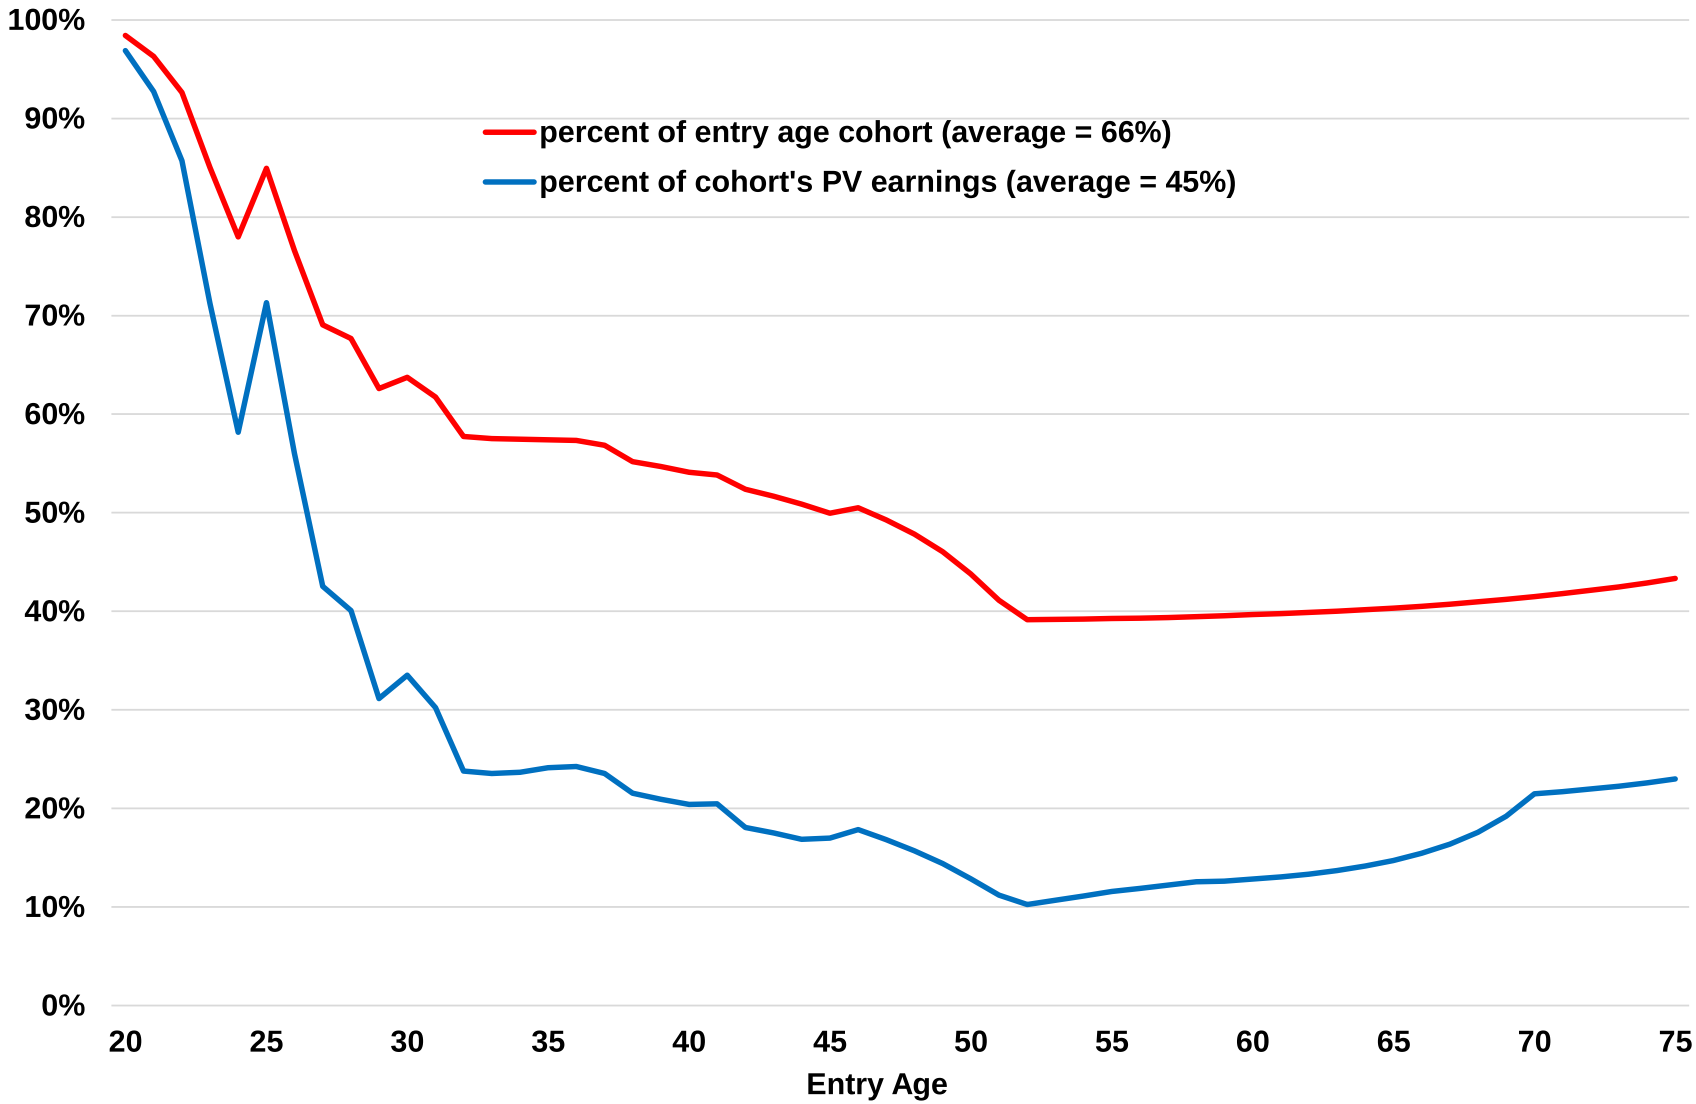

Figure 5. Distribution of Exit Age, Among those who Entered at Age 30

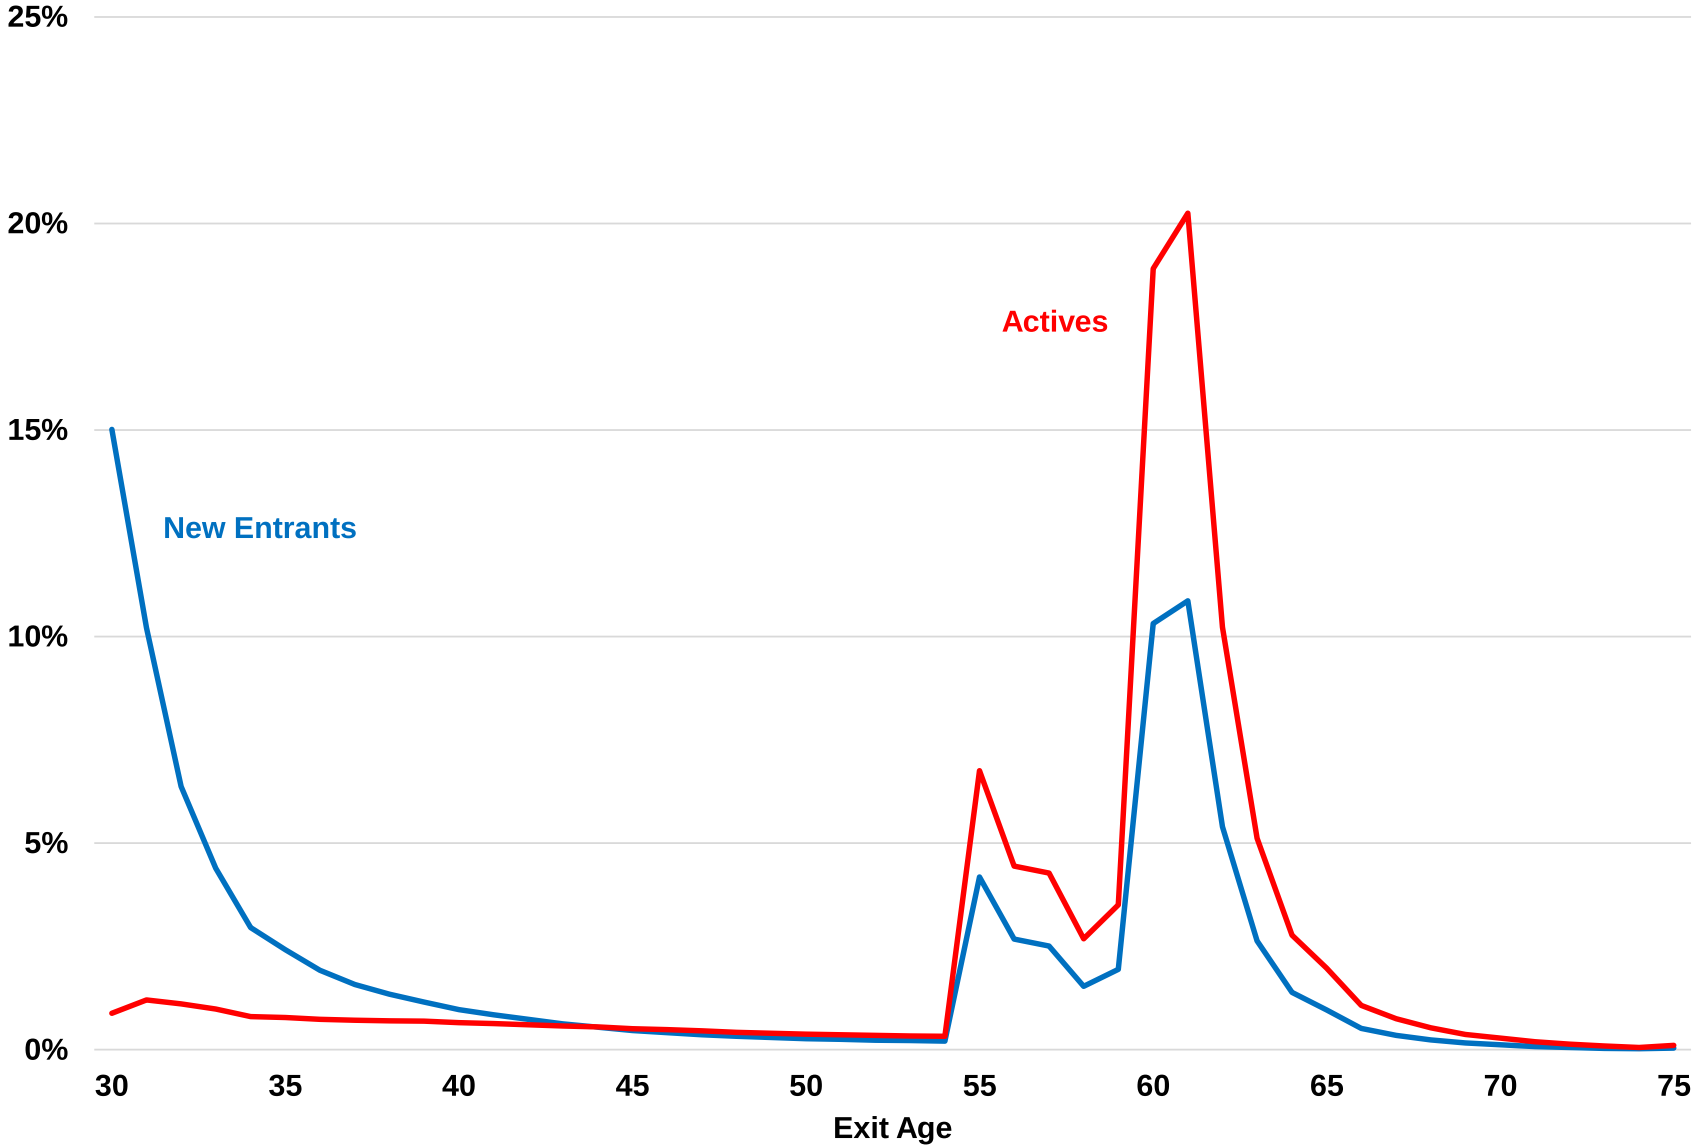

Figure 6. Rhee-Fornia, Value of Defined Benefit and Cash Balance Plans

Entry Age 30 -- Female (Rhee-Fornia, Appendix B, and authors' adjustment for CB @ 7.5%)

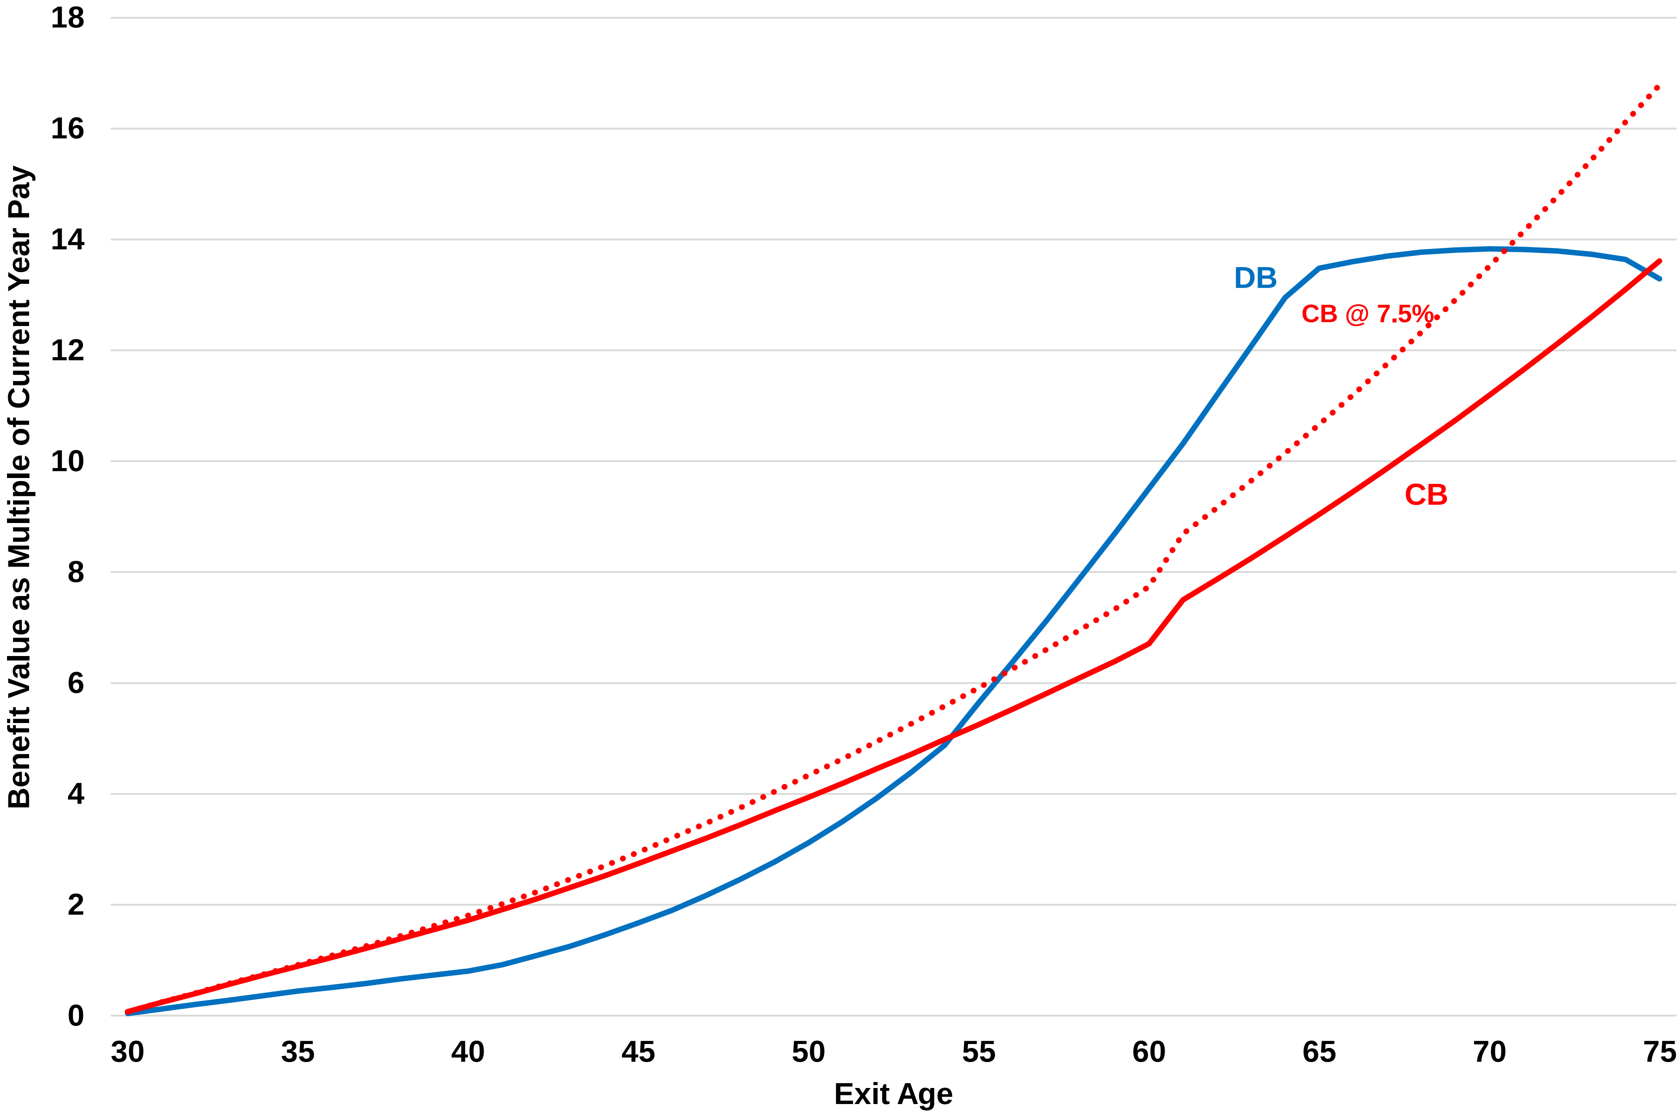

Figure 7. Average Employer NC Rate for Winners & Losers, By Entry Age

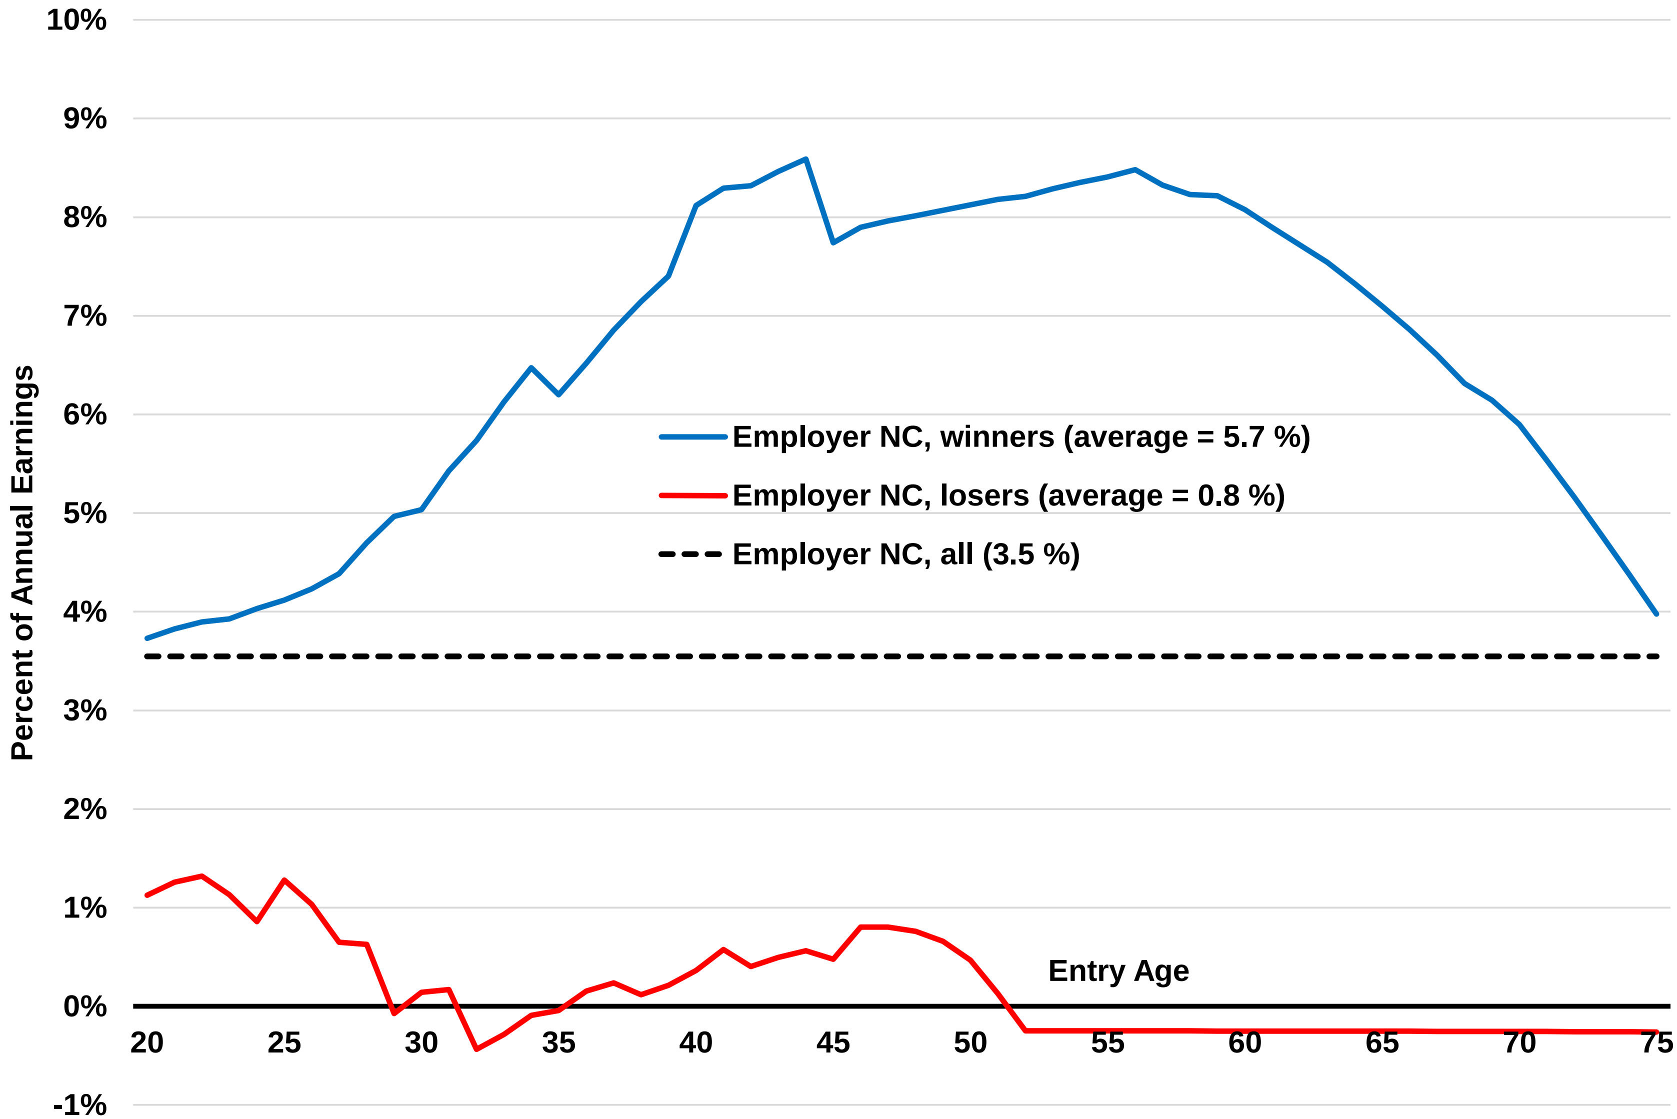

Figure 8. Employer Normal Cost by Entry Age

Figure 9. Decomposing the Cross-Subsidies Within and Across Entry Ages