California Public Employees' Retirement System

Review of the 1997-2011 Demographic Experience Study

Produced by Cheiron

February 3, 2014

CHEIRON Classic Values, Innovative Advice

Table of Contents

February 3, 2014

Board of Administration
California Public Employees' Retirement System
Lincoln Plaza North
400 Q Street
Sacramento, CA 95811

Members of the Board:

The purpose of this report is to present our review of the 1997-2011 demographic experience study performed by the California Public Employees' Retirement System (CalPERS) Actuarial Office. This report is for the use of CalPERS in selecting demographic assumptions for its actuarial valuations.

In preparing our report, we relied on information (some oral and some written) supplied by CalPERS. This information includes, but is not limited to, the decrement and exposure data used in the experience study. We performed an informal examination of the obvious characteristics of the data for reasonableness and consistency in accordance with Actuarial Standard of Practice No. 23.

To the best of our knowledge, this report and its contents have been prepared in accordance with generally recognized and accepted actuarial principles and practices that are consistent with the Code of Professional Conduct and applicable Actuarial Standards of Practice set out by the Actuarial Standards Board. Furthermore, as credentialed actuaries, we meet the Qualification Standards of the American Academy of Actuaries to render the opinion contained in this report. This report does not address any contractual or legal issues. We are not attorneys and our firm does not provide any legal services or advice.

This report was prepared exclusively for CalPERS for the purpose described herein. This report is not intended to benefit any third party, and Cheiron assumes no duty or liability to any such party.

Sincerely,
Cheiron

William R. Hallmark, ASA, FCA, EA, MAAA
Consulting Actuary

Michael Moehle, FSA, EA, MAAA
Consulting Actuary

cc: Gene Kalwarski
Robert McCrory

EXECUTIVE SUMMARY

Scope of Assignment

CalPERS recently completed its 1997-2011 demographic experience study. Under a Letter of Engagement issued pursuant to CalPERS Agreement No. 2013-6899, Cheiron, Inc. (Cheiron) was retained to perform a comprehensive review of this experience study.

Our review included an examination of the processes used by the CalPERS Actuarial Office to analyze the data, to make a recommendation to either keep the current assumption or propose a new assumption, and to develop a new proposed assumption. We reviewed the following assumptions:

- Merit salary scale
- Disability
- Service retirement
- Pre-retirement mortality
- Post-retirement mortality

The full CalPERS experience study also includes assumptions for terminations, price inflation, wage inflation, discount rate, and payroll growth, which were not included under this letter of engagement. In addition to their analysis for each assumption, the CalPERS Actuarial Office provided the number of decrements during the experience study period and the number of exposures on which their analysis was based. We did not review the development of this underlying information.

Key Findings and Recommendations

Based on our review of the 1997-2011 demographic experience study performed by the Actuarial Office, we believe the proposed assumptions for merit salary scale, work-related and non-work-related disability, service retirement, and pre- and post-retirement mortality are reasonable, appropriate and were developed in accordance with generally accepted actuarial principles.

We suggested some technical changes that we do not believe would make a material difference in the current assumptions proposed, but should be considered for the next experience study. In particular, we recommend the use of more quantitative methods, including the use of confidence intervals for the development of certain assumptions.

Finally, we suggest some alternatives for the projection of mortality rates for the Actuarial Office and the Board to consider. We believe these alternatives have some advantages over the proposed projections, but they also add complexity and may have other disadvantages.

Quantitative Methods

For most of the assumptions studied, the Actuarial Office used quantitative methods to establish an initial proposed assumption, and then applied professional judgment to develop the final proposed assumptions. However, in some cases, quantitative measures other than actual-to-

expected ratios were not used. Instead, the proposed assumptions rely almost exclusively on professional judgment.

We recommend that in the next experience study, the Actuarial Office employ quantitative measures such as confidence intervals and R-square to aid in the decision of whether or not to change an assumption and the development of proposed assumptions.

Alternatives for the Projection of Mortality Rates

Current standards of actuarial practice require recognizing the likelihood of future improvements in member mortality. Human life expectancy has been increasing for centuries, and recent experience in the United States has shown greater than anticipated improvements for ages over 55. Every indication is that some degree of mortality improvement will continue for the foreseeable future.

There are two methods for recognizing mortality improvements: A generational mortality table and a static projected mortality table.

The best choice technically would be for CalPERS to use a generational mortality table. Such a table projects mortality improvements separately for each year of birth. A generational approach reflects the expectation that someone who is 65 years old today is likely to experience less improvement in mortality than someone who is 25 years old today. A generational table does the best job of matching the expected pattern of changing future mortality rates.

However, we understand that CalPERS is unable to use a generational table in its actuarial valuation system. Until that capability is added, a static projected mortality table must be used. Under this option, a single table is projected to some point in the future to estimate the results that would be produced if a generational table had been used.

The method proposed by the Actuarial Office to project mortality improvements is reasonable and administratively straightforward. The two alternatives we present below for consideration add some additional complexity, but improve the accuracy of the results compared to the proposed method. Whether the additional accuracy warrants the additional complexity is a question for the Board and the Actuarial Office.

Separate post retirement tables for actives and retirees

The first alternative to consider is to establish separate post-retirement mortality tables for actives and retirees. Active employees are generally younger and should be expected to have lower mortality rates in retirement than current retirees. Consequently, active employees could be valued using a post retirement mortality table with a longer projection than proposed while retirees could be valued using a post retirement mortality table with a shorter projection than proposed.

EXECUTIVE SUMMARY

This approach would not significantly affect the measure of liability for CalPERS as a whole, but would more accurately assign the liability to individual employers within CalPERS. Under the Actuarial Office's proposed assumptions, there are likely to be systematic gains or losses for individual employers. This alternative would mitigate those systematic gains or losses.

Project to year of each valuation

The Actuarial Office proposed projecting mortality improvements for 20 years: seven years to get from the mid-point of the mortality experience used in the study to the mid-point of the valuations for which the assumption will be used plus 13 years for mortality improvements beyond the valuation date. As an alternative, each valuation could be performed with a mortality table projected 13 years beyond the valuation date. As a result, the initial valuation based on this experience study would use higher mortality assumptions than the proposed assumptions, but each year the mortality table would be projected an additional year so that when the next experience study is performed, there would not be an automatic increase for four years of mortality improvement. This approach better simulates the use of a generational mortality table and reduces the likelihood of a significant adjustment when an experience study is performed.

The remainder of the report provides additional detail on our analysis and recommendations, including some technical changes we suggest for the next demographic experience study.

MERIT SALARY SCALE

Summary

The merit salary scale is used to project salary increases in addition to the across-the-board wage inflation assumption of 3.0 percent. The Actuarial Office uses a logarithmic model to separately graduate the data observations during the first 8 years of service and 9 or more years of service. Some adjustments are made to smooth the connection between the models and to fit the ultimate rate selected.

The methodology used and assumptions proposed in the experience study are reasonable. We have one technical comment we suggest be considered for the next experience study, but we do not believe it would materially change the conclusions of the current experience study.

Technical Comment

Averaging Methodology

In combining multiple years of experience for the study, CalPERS staff weighted the experience for each year by the number of people present in each of those years for that particular group and entry age-service combination. Instead, the combination for different years should be weighted based on salary as it is in determining the average increase for an individual year.

The actuarial liability for a given group is largely proportional to salary. Determining the average for multiple years by weighting by count gives the same weight to the salary increase for a low-paid member as it does for a high-paid member. However, the liability being measured is much more sensitive to the salary increase for the high-paid member.
To illustrate, consider the following simple example with two scenarios.

		Pay Increase	
	Salary	**Scenario 1**	**Scenario 2**
High-Paid Member	$200,000	10%	0%
Low-Paid Member	$20,000	0%	10%

In Scenario 1, the high-paid member receives a 10% increase in pay while the low-paid member receives no pay increase. In Scenario 2, the low-paid member receives the 10% increase while the high-paid member receives no increase. In both cases, the average weighted by count is a 5% increase, which would result in a projection of total salary for both members of $231,000 ($200,000 x 1.05 + $20,000 x 1.05). In Scenario 1, however, actual pay would increase to $240,000 ($200,000 x 1.10 + $20,000 x 1.00), which is a 9% increase. In Scenario 2, actual pay would only increase to $222,000 ($200,000 x 1.00 + $20,000 x 1.10), which is a 0.9% increase.

In the study, there are usually more than two members being averaged and the disparity in pay is not as extreme as in the example, so differences are likely to be much smaller. Nevertheless, a systemic bias could emerge (with a weighting toward lower-paid employees), so we suggest that this methodology be changed in future experience studies.

MERIT SALARY SCALE

When combining salary increases over different years, weighting on an unadjusted salary gives more weight to recent experience than to experience early in the study period. Average salaries in 1997, for example, are likely to be lower than average salaries in 2011. If it is not desired to give more weight to the more recent experience, salaries for earlier years can be adjusted to the equivalent amount as of the end of the study period by increasing them for the actual wage inflation during the period.

DISABILITY

Summary

The proposed rates for disability incidence are based on graduating the observed disability rates using the Whittaker-Henderson method. The method appears to work very well where there are significant data observations, but more judgment has to be applied when the data is limited. The methods used and the assumptions proposed are reasonable, but we recommend two technical improvements for consideration for future studies.

Technical Comments

Confidence Intervals

The decisions on whether or not to propose a new assumption were largely based on the analysis of actual-to-expected ratios. However, it is not clear how far off an actual-to-expected ratio should be before an alternative assumption is warranted, particularly when some ratios are based on very few observations. Confidence intervals can provide a decision criterion that adjusts for the robustness of the data. We suggest using 90% confidence intervals on the aggregate disability incidence rate. If the assumption is outside the 90% confidence interval, consider proposing a new assumption.

The charts below confirm that in most cases, new assumptions have been proposed when the current assumption is outside the confidence interval. The gray bars represent the range of the confidence interval. The black squares represent the actual aggregate rate for the period of the experience study; the red triangles represent the current assumption; and the blue diamonds represent the proposed assumption.

In the first chart (Non-Work-Related Disability for males), new rates were proposed for Industrial, Public Agency Miscellaneous, Schools and State Miscellaneous that moved into the confidence interval. The only other group where the current assumption is outside the confidence interval is State Safety, which is just slightly below the confidence interval.

DISABILITY

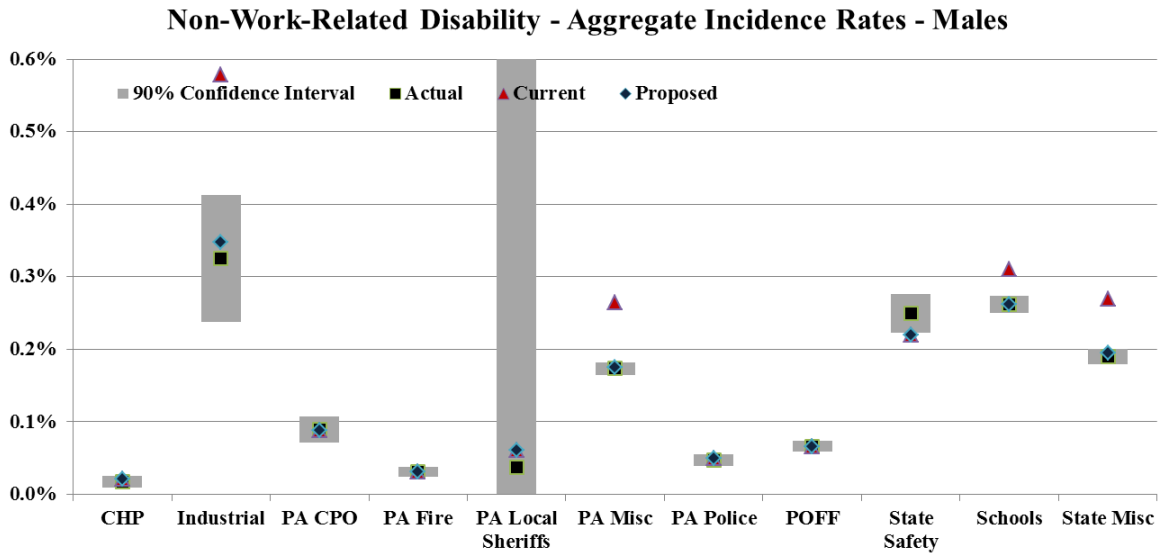

However, as shown in the chart below (Work-Related Disability for females), for CHP, Public Agency CPO, Public Agency Fire, Public Agency Police Sheriffs, and POFF, the current rates are outside of the confidence interval and the proposed rates are even further outside the confidence interval. The disability rates for males and females have been combined because the number of females is relatively small for these groups. However, the confidence intervals show that the rates for females are clearly different than the rates for males.

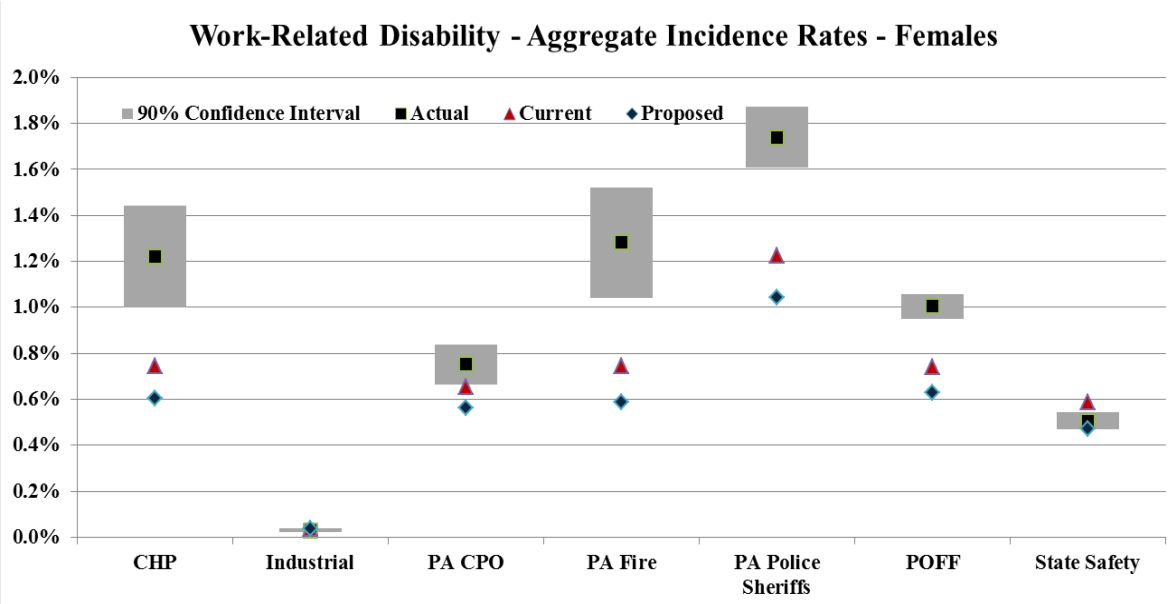

Quantitative Method to Assess Fit

The actual-to-expected ratios and the confidence intervals on the aggregate disability rates provide a reasonable indication of when the aggregate assumption is accurate or needs to change, but they don't provide any information about how well the assumptions at different ages fit the pattern of the experience. When there is sufficient data for the Whittaker-Henderson

graduation method, it automatically uses a least-squares approach to balance the fit to the data with the smoothness of the assumption curve. However, in some cases, it may emphasize the smoothness over the fit to an extent that is not ideal or the underlying polynomial may not be appropriate. The chart below shows the industrial disability experience for male POFF members, including the confidence intervals, current and proposed assumptions.

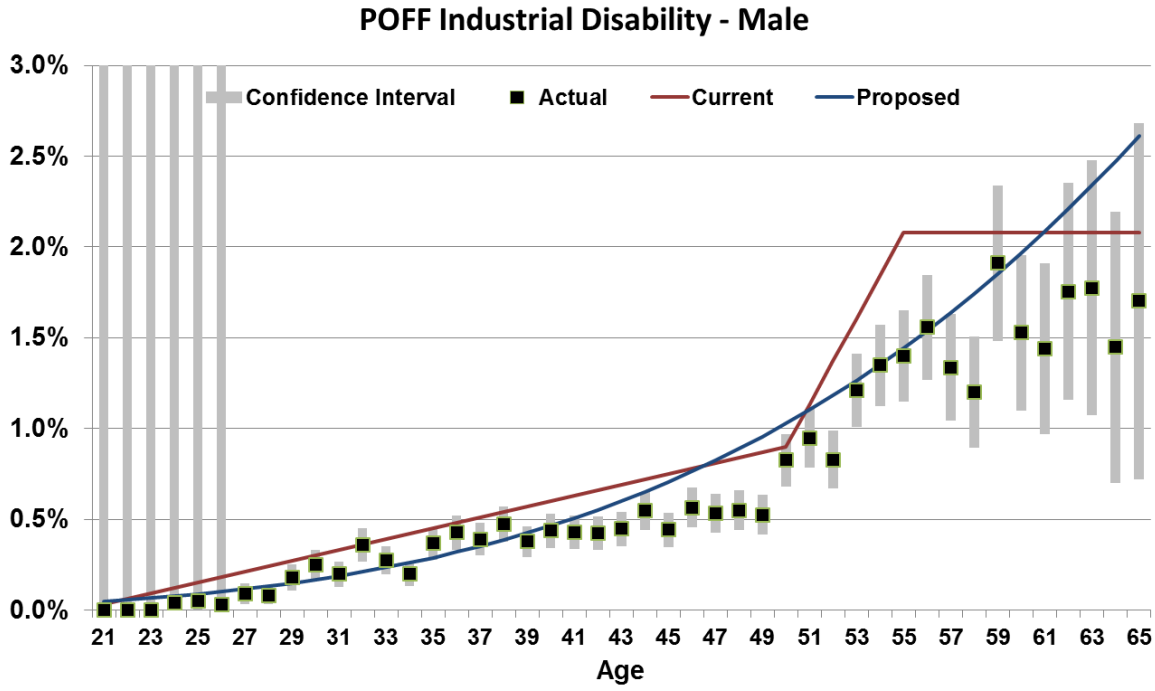

The confidence intervals at each age provide some insight as to the pattern of the experience and whether the assumption follows that pattern. In addition, the R-square statistic can be calculated to test the fit. In this case, the current assumption has an R-square of 0.939 and the proposed assumption has an R-square of 0.865, indicating that the proposed assumption does not fit the data as well as the current assumption although the A/E ratio improves from 0.707 to 0.843.

In future studies, we suggest that confidence intervals, R-square statistics, and/or other quantitative measures be used to gauge the fit of the assumption to the data.

SERVICE RETIREMENT

Summary

Proposed rates of service retirement are generally higher than the current rates. Decisions on which rates to change appear to be made primarily on actual-to-expected ratios with significant amounts of judgment applied to select proposed rates for individual age-service combinations. While the proposed assumptions appear to be reasonable, we suggest that a more quantitative approach be developed for the next study that still allows for the application of some professional judgment.

Technical Comments

The actual-to-expected ratio provides a good measure of whether the aggregate rate of retirement is reasonable, but does not provide any information on how well the assumption fits the pattern of actual retirements. Instead, the pattern of retirements has been developed and modified over the years based on the application of professional judgment to the actual rates experienced. Graduation techniques like the Whittaker-Henderson method are unlikely to work for retirement rates because the underlying pattern does not fit well with a polynomial.

Given the large amount of data and set of assumptions, we suggest that a quantitative approach should be used, at least to set initial rates, before making adjustments based on professional judgment.

To illustrate the issue and a possible approach, consider the retirement rates for School members with five to nine years of service. The chart on the following page shows the confidence interval, actual rate, current assumption, and a "Quantitative" assumption. There are no proposed changes for this group.

SERVICE RETIREMENT

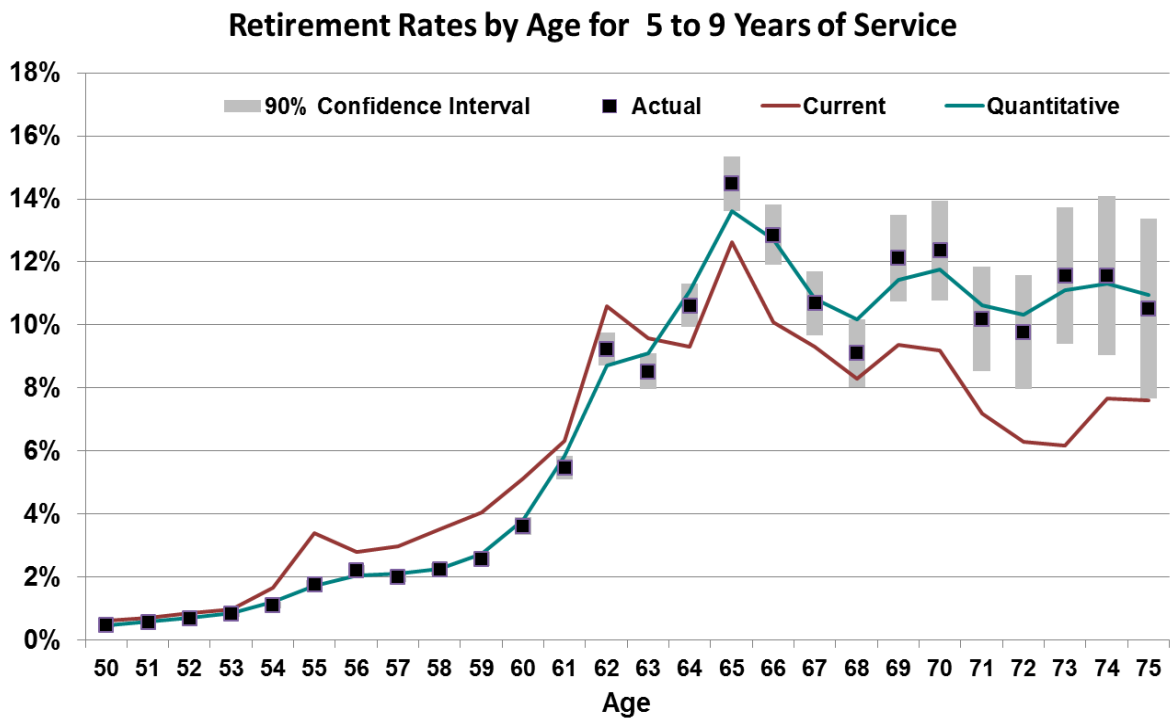

The current rate is above the confidence interval for ages 54 through 63 and below the confidence interval after age 63, with the exception of age 68. The higher rates and lower rates balance out to make the actual-to-expected ratio near 1.0, resulting in the recommendation of no changes.

The quantitative approach shown in the graph follows a simple formula to provide an initial assumption. The formula averages the current age's actual rate (50% weight) with the actual rate for the prior age (25% weight) and the actual rate for the next age (25% weight). Then, the assumption is limited to the 90% confidence interval for the current age. This initial assumption may need to be adjusted, for example, when Social Security benefits become payable or based on subsidies in the benefit formula. Other criteria could also be added quantitatively such as ensuring that retirement rates for the same age with greater service are always greater.

In addition to confidence intervals and a quantitative formula for the initial assumption, we suggest using other statistical measures such as R-square to assess the fit of the assumption to the underlying data. In this particular case, the quantitative assumption improves the R-square measure from 0.83 for the current assumption to 0.99.

The Actuarial Office was kind enough to test this quantitative assumption for the Schools group and found no material differences in the measures of liability. We suggest that a quantitative approach be adopted for the next experience study.

MORTALITY

Summary

The Actuarial Office has proposed lower rates of mortality reflecting observed improvements from the prior study and projected improvements in the future. The unprojected rates were developed based on data from 2006 to 2011 graduated using the Whittaker-Henderson method. We agree with the development of the graduated rates and find the result to be reasonable.

In the prior study, the mortality rates developed in this manner were projected using scale AA from the midpoint of the study to the year of the study, 2010. That method did not anticipate any future improvements in mortality beyond 2010. In this study, the proposed rates reflect a 7-year projection from the midpoint of the study to the midpoint of the valuations for which the mortality table is to be used, plus an additional 13-year projection for future mortality improvement beyond the valuation date. Finally, the projected improvements are based on scale BB, which was recently issued by the Society of Actuaries to replace scale AA on a temporary basis until a new projection scale is released in 2014.

Future mortality improvements are best projected using a generational mortality table. A generational table essentially creates a separate mortality table for each year of birth. This approach reflects the expectation that an 80-year-old today is more likely to die than someone who becomes 80 in 20 years. We understand that the CalPERS valuation system cannot accommodate a generational mortality table, so the result must be approximated by projecting a single table into the future. Because of this approach, liabilities for retirees are likely to be overstated while liabilities and normal costs for active members are likely to be understated. The objective is to project mortality improvements so that in aggregate, the measure of liability approximates what would be achieved with a generational table.

Given the inability to use a generational table, we believe the 20-year projection proposed is reasonable. In our technical comments, we analyze the projection in more detail and offer some alternatives for the Actuarial Office and the Board to consider.

Technical Comments

Rate of Projected Mortality Improvement

The scale used to project mortality improvements is scale BB published by the Society of Actuaries in September 2012. Scale BB was developed, based in part on CalPERS data through 2006, to replace the predominant scale in use, scale AA. The Society of Actuaries intends to release a new projection scale in 2014 (MP-2014), but released scale BB on a temporary basis because the preliminary results of their analysis found that scale AA was inadequate.

To test whether scale BB continues to be an appropriate scale for CalPERS, we compared the rate of mortality improvement between this study and the prior study to scale BB. In the prior study, the smoothed rates of mortality were projected five years using scale AA. So, we increased the rates used in the current assumption to remove the five years of projected improvement and compared those rates to the smoothed, unprojected rates from the current

study. The charts on the following page show the comparison of annual service retirement mortality improvement rates for males and females.

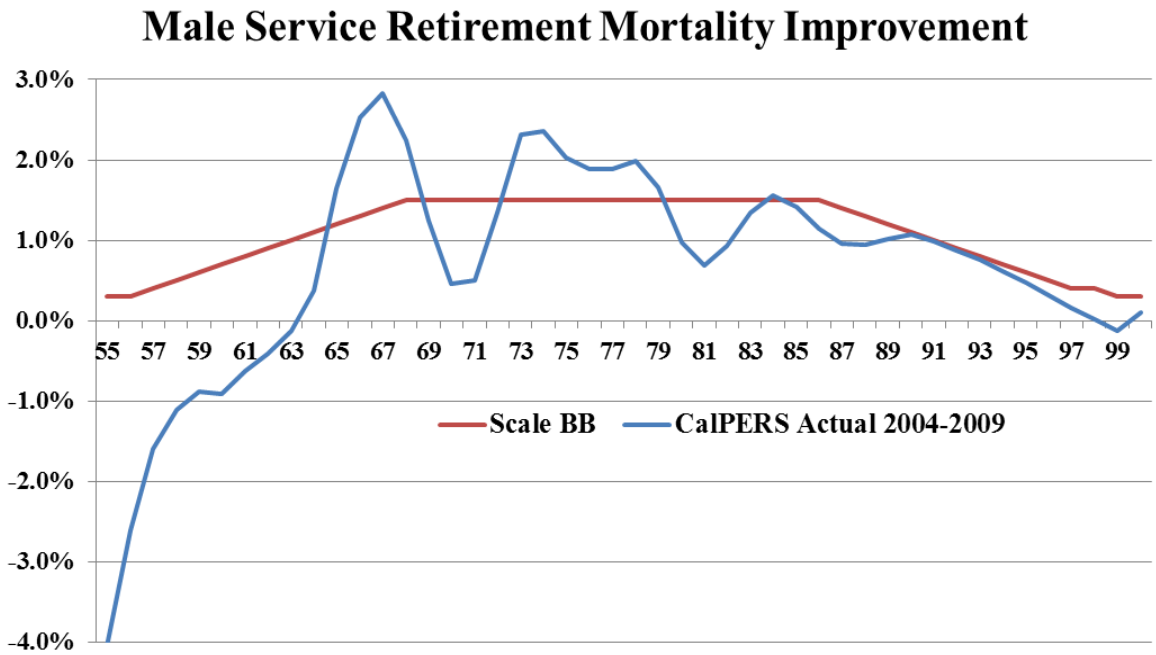

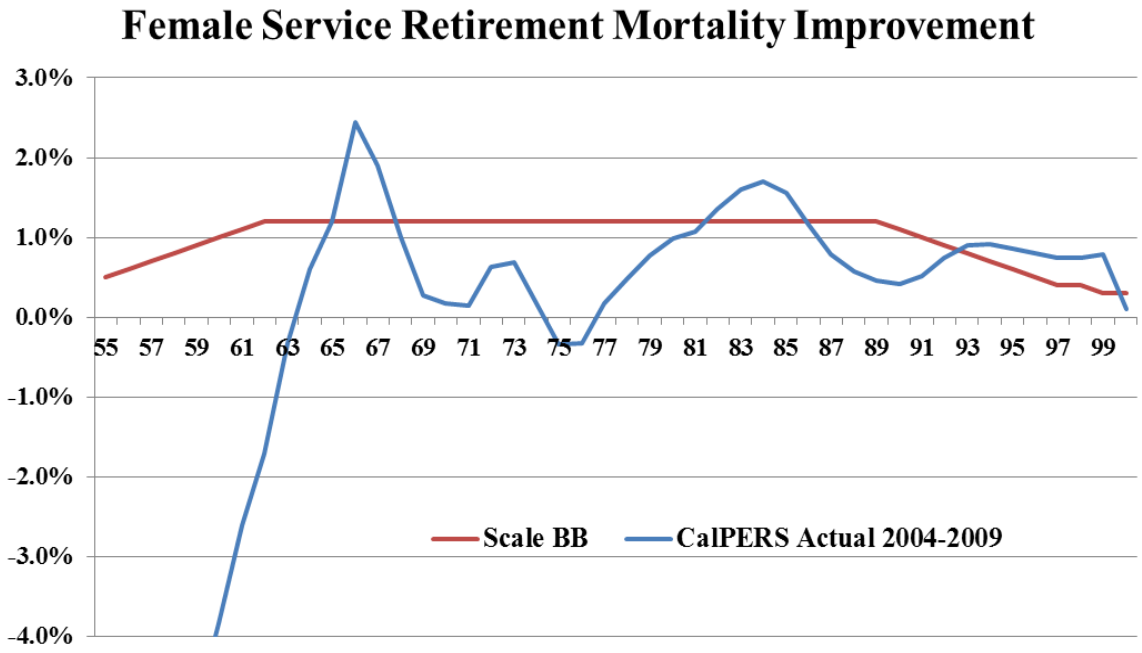

The recent experience for male service retirees matches scale BB relatively well for ages 64 and older, with some ages experiencing higher rates of improvement and others lower rates of improvement than projected by scale BB. The recent experience for female service retirees

doesn't match as well with improvements tending to be slightly lower than scale BB, but the rate of improvement is still reasonably close to scale BB. The mismatch under age 64 and the negative mortality improvement for both male and female service retirees is attributable to a change in methodology between the last study and the current study. In the last study, the transition between pre-retirement mortality rates and post retirement mortality rates was smoothed between ages 50 and 64, but in the current study, it was not. Consequently, the negative mortality improvement shown in the graph is primarily due to the change in methodology.

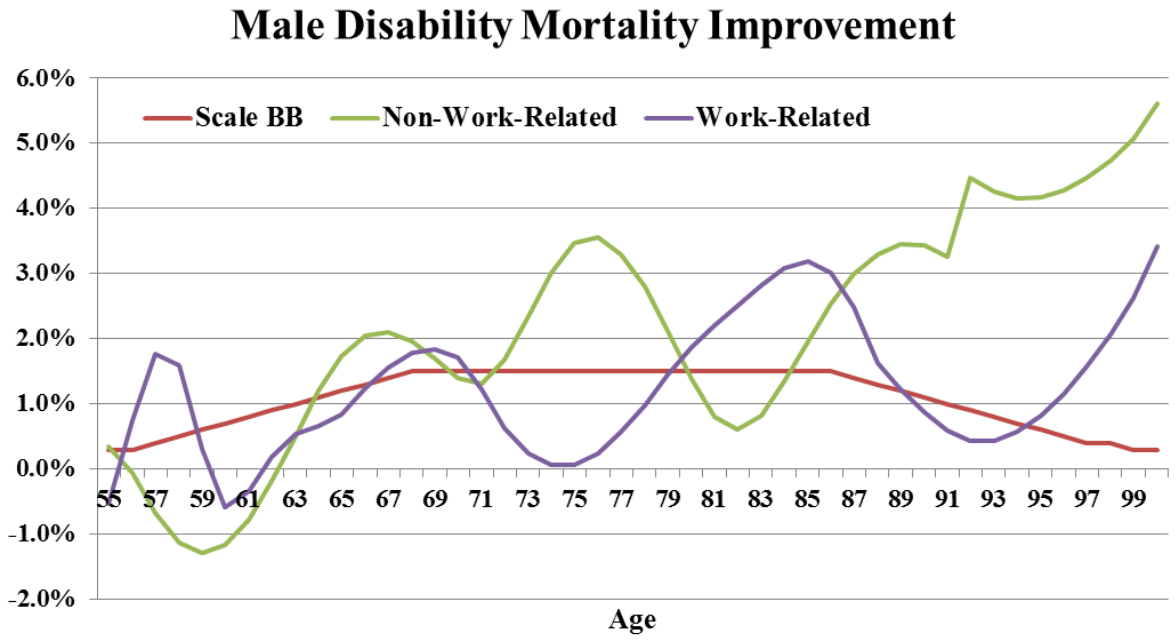

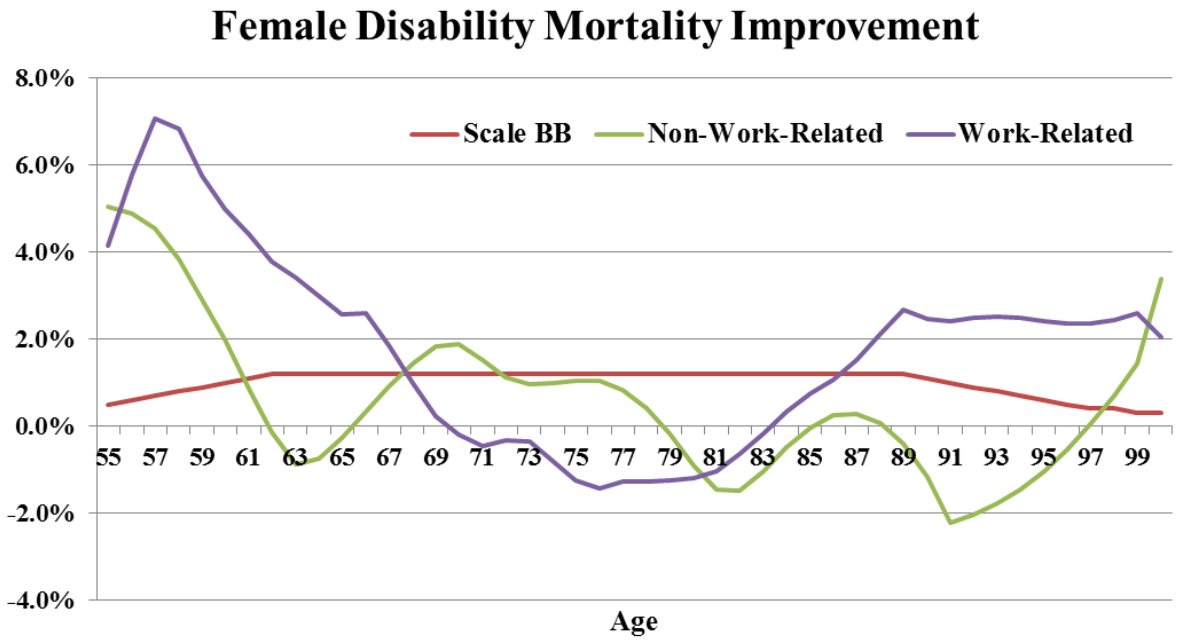

MORTALITY

For disability retirees, the match to scale BB is less clear. For males, scale BB may provide a reasonable match to recent experience for ages under 82 or 83. For the older ages, there appears to be significantly greater mortality improvement. For females under age 62, the mortality improvement appears to have been significantly greater than scale BB would project. However, from ages 67 to 85, the mortality improvement appears less than scale BB would project.

It is difficult to determine if these differences are due to less robust data, one-time effects, or ongoing effects that should be projected into the future. There is no alternative projection scale of which we are aware that would better fit the data. The alternative to scale BB is for CalPERS to create its own projection scale for disabled mortality. Given the limited data, it is not likely that a custom scale would be more reliable.

While the disabled mortality improvement trends should be monitored in the future, we believe the use of scale BB is appropriate and reasonable for this study.

Projection Period

The proposed rates include a 20-year projection, which is made up of seven years to adjust from the midpoint of the data used in the study to the midpoint of the valuations for which the mortality assumption will be used, and 13 years of additional mortality improvement.

The midpoint of the data used in the study is January 1, 2009, so the projection of seven years estimates mortality rates as of January 1, 2016. We understand these mortality rates will be used for the June 30, 2014 through June 30, 2017 valuations.

As a result of this method, and using one static table for all of these valuations, the mortality rates will tend to be somewhat conservative for the first valuation and somewhat aggressive for the last valuation of the period. If the same methodology is used for the next valuation, and mortality continues to improve, there will likely be a noticeable impact again when the mortality assumption is changed (although not as big as for this study).

The projection for 13 years of additional mortality improvement represents an estimate of the adjustment needed to match the valuation results on an aggregate basis to what would be produced with a generational mortality table. The rule of thumb employed is to project the mortality period a number of years equal to the duration of the liability. The Actuarial Office has estimated the system-wide duration to be 13 years.

The concern about this rule of thumb is that the duration varies depending on what measure of liability is being used. The duration of retiree liabilities is usually in the range of 10 to 12 years while the actuarial liability for active employees usually has duration of 15 to 18 years, and the normal cost usually has duration of 20-25 years. So, the appropriate projection period depends on what is being measured. The estimate of 13 years is reasonable for the actuarial liability of the system as a whole.

Alternatives to Consider

MORTALITY

There are two alternatives to the proposed method that the Actuarial Office and the Board may want to consider. Both have advantages and disadvantages compared to the proposed methodology

Separate retiree and active mortality tables

The first alternative is to adopt separate post retirement mortality tables for actives and retirees with different projection periods. Since the liability for retirees has a lower duration, the projection period would be shorter and would reduce the measure of retiree liability compared to the proposed rates. At the same time, the projection period for active employees would be longer, reflecting the longer duration of the actuarial liability and normal cost for active members.

The advantage of this approach is that it would more accurately estimate the liability that would have been calculated with a generational mortality table for each employer in CalPERS. By adopting a single post retirement mortality table as currently proposed, employers with a greater proportion of active employees are likely to experience actuarial losses.

The disadvantage of having two separate tables is the addition of a little more complexity. In addition, when a member actually retires, it would change the post retirement mortality table used to value their benefits and create an actuarial gain.

Project to year of each valuation

The current proposal is to project the mortality table from the midpoint of the mortality experience (January 1, 2009) to the midpoint of the valuations for which it will be used plus an additional 13 years. We understand that this mortality table will be used for the June 30, 2014, 2015, 2016 and 2017 actuarial valuations, so mortality improvements are being projected seven years (from January 1, 2009 to January 1, 2016) plus an additional 13 years for improvements after the valuation date. This means that the first valuation will have mortality improvements projected further beyond the valuation date (14.5 years) than the last valuation for which the mortality table is used (11.5 years).

A second alternative to consider is to project a constant 13 years beyond the valuation date. That is, instead of projecting seven years to the midpoint of the valuations and then another 13 years, just project to the valuation date and then another 13 years. The advantages of this approach are:

- A shorter projection period for the first valuation requiring less of an immediate adjustment,
- Small adjustments each of the following years,
- A smaller expected adjustment when the next experience study is performed, and
- Better approximation in each valuation of the use of a generational mortality table.

MORTALITY

The primary disadvantage is the additional complexity of changing the mortality table with each valuation, particularly because under current policy the valuation mortality is also used for member benefit calculations. Either the factors for member benefit calculations would need to be updated every year, or the factors would need to be based on a different mortality table than the valuation.