Mitigating Volatility of Retiree Health Valuation Results

By Jeff Petertil, ASA, FCA, MAAA & Justin Petertil

While actuarial practice may have a long history of using smoothing methods to help plan sponsors manage cost volatility in pension plans, the history of smoothing methods in other post employment benefits is shorter and less extensive. This is not due to a lack of volatility, however; retiree health plans in particular have shown substantial year-to-year swings in accounting liabilities. Nonetheless, there is little action or concern for sponsor funding of the liabilities through asset accumulation, so smoothing that has occurred has been according to accounting rules (corridors for gain/loss recognition) or by adjusting assumptions, without fanfare or public acknowledgement that smoothing was the intention. This paper explores some sources of the volatility in retiree health valuations, noting how the volatility arises and might be mitigated.

Volatility of retiree health accounts has causes in common with those that affect pension account volatility. This paper, however, discusses aspects of actuarial valuation that are more problematic for retiree health benefits than for pensions. We examine three influences that have had significant effect on retiree health valuation volatility—changes in benefit level, fluctuations in cost level, and alterations in eligibility for, or duration of, benefit—and focus attention on the last two. A high profile example of volatility is reviewed, problem areas are identified, including some also seen in the pension field, and volatility of retiree health cost is explored. The differing measurement approaches taken by authorities (accounting, tax and regulatory) to account for the influences are discussed, along with their strengths and weaknesses. Such approaches have not captured the economic values inherent in many of the employer/employee situations and have, indeed, led to some of the reported volatility. Recommendations to capture that value and manage volatility are suggested, the effectiveness of the recommendations is measured, and the ramifications investigated.

Retiree Health Benefit Volatility in the Headlines

The most extensive recent evidence that volatility from period to period is a problem in the retiree health area came in late March 2010, with surprise announcements by major companies of significant new expenses. Days after the Patient Protection & Affordable Care Act was passed by Congress and signed into law by the President, several corporations announced financial charges of hundreds of millions of dollars (for one corporation, a billion dollars) related to the law. For a law that was not supposed to increase costs much, this was a shocker, and political gamesmanship soon followed. But when the smoke cleared, the reported blows to corporate profits were found to involve the valuation of retiree health benefits and be legitimate accounting entries with relatively little impact on corporate cash flow. Whether it was a mountain or a molehill depended on one's perspective, but there was no question that the result was financial reporting volatility.

Public explanations for the volatility from those accounting surprises in the first quarter of 2010 nominally involved unintended consequences and changes in the tax code. The three influences covered in this paper were not mentioned. While not denying those nominal factors, a close reading of this paper and the facts in that situation may lead a reader to conclude a closer connection. While there has been some achievement in limiting systemic cost volatility of retiree health benefits, not all measurement approaches have been a success. The aspects of historic volatility discussed here have been managed differently and with different success.

Sources of Retiree Health Benefit Volatility

Actuarial Standard of Practice No. 6, Measuring Retiree Group Benefits, notes three key elements in developing a valuation model—known plan provisions, the population covered, and the current and projected claim cost. An actuary dealing with retiree health benefits soon encounters year-to-year movement in claim costs and plan provisions. As to population covered, plans will also be investigating alternatives in eligibility or duration of the retirees' coverage and occasionally implementing such changes. The reasons behind each type of change vary and are worth examining for their role in the volatility and for ability to be managed. The

changes may have some parallels with pension valuations, but usually arise from different situations and are handled differently.

The first of these changes is that of plan provisions, or benefit level—cost-sharing arrangements between retiree and plan sponsor may shift each year, as may the insurer or managed care vendor. The retiree faces new levels of deductible or copayment or changes in plan maximums related to out-of-pocket payments, annual payment amounts, or lifetime reimbursements. These changes come about for various reasons, as plan sponsors react to changes in the medical delivery system, to overall plan cost levels, and to changes in the Medicare program, which is often the primary insurer for a retiree or dependent. All such changes throw into question the applicability of past data to future costs and increase the potential for volatility in present values of long-term cost projections. The volatility example from March 2010 cited above was actually smaller in magnitude than the one that came about when the Medicare Modernization Act was passed in 2003. Those changes to what Medicare covered for prescription drugs reconfigured much of the retiree coverage that was secondary to Medicare and resulted in larger changes; the 2010 change, in comparison, was just a tweaking of the tax code. But smaller changes by plan sponsors that effect plan reimbursement levels happen more often with retiree health plans than is true for pension plans, and such changes may have more financial impact.

A second set of changes that does not have an immediate parallel in pension valuations is the change in the medical claim cost level, the main component of benefit cost. A retiree health plan may find it pays $X per participant one year, $.9X the next year and then $1.2X the following year. It is well known that there is a great deal of volatility in group claim costs from period to period. What claim level then should be projected forward? The consensus soon became that no matter what initial claim cost level was used in a valuation, there was a need to model increasing costs for the future. Thus, health care cost trend became a necessary assumption, but one that could also change from year to year. These modeled increases, however, presented a number of thorny challenges as to the cost effects on all the provisions of the benefit plan, those mentioned above and others (such as the retiree contribution).

Beyond these two major sets of changes—benefit level and claim level—that can have a significant impact on actuarial cost, there is a third set tied to the population covered, the eligibility for the benefit, and the duration of the benefit. Many retiree health plans are unilaterally rescindable.[1] The plan sponsor may, if it has the legal right, reduce the benefit in ways small and large, to the extreme of terminating the plan and rescinding the promise. While completely rescinding the benefit is not common, a more incremental approach (often with the long term effect of complete termination) is frequently employed. The plan sponsor may change the period for which a retiree is eligible (e.g., from lifetime to 10 years, or only to age 65) or the terms of eligibility (e.g., from 10 years of work service to 20 years of work service). New hires may be declared not eligible. The periodic contribution required from the retirees (sometime referred to as the retiree premium) may be increased by 100 percent or more, either by fiat or in reaction to an upper limit (usually referred to as a cap) placed on what the sponsor will pay. Closing a plan to those hired after a certain date has parallels in the pension world, but the pension equivalent of freezing a plan does not exist in the retiree health world. The increments in a pension are obviously dollar amounts; understanding a freeze point is somewhat uncomplicated since the plan provisions state what is vested and when. This is almost impossible with a health care indemnity plan, where vesting is rare and the record-keeping involved would present a daunting adjudication challenge. (Consider the problems involved in freezing medical payment levels for a future procedure at dollar levels from, say, the year 2003.) Instead, the indemnification may be adjusted at the benefit provision level, but the real change is with retiree premium increases or the wholesale dropping of

[1] In an earlier paper by one of the authors, Measuring Terminable Postretirement Obligations, published in the North American Actuarial Journal in January 2005, the word “terminable” was used to describe this feature. A commentator pointed out that, while that term might be applicable, it could confuse those familiar with pensions, where a plan can be terminated without the benefits being terminated. This would be almost impossible with a retiree health benefit plan, but we have chosen to replace “terminate” with “rescind” in numerous places for this paper. The concepts as applied to retiree health plans are identical.

coverage. Changes of this third type have generated considerable historic volatility in present values and the associated measurements.

The Limited Scope of this Paper

There are numerous elements involved in triggering volatility and there are numerous responses to managing those elements and smoothing otherwise volatile results. Only a few are covered here; for example, no attention will be given to asset smoothing, primarily because there are few assets dedicated to pre-funding retiree health benefits, but also because any smoothing of asset values would have very close parallels in the pension area and be well covered by experts in that area. For similar reasons, there will be no commentary on net plan financials, since the lack of funded assets means that most all that is said about plan liabilities also holds for smoothing net positions. The lack of funding also has led us to confine our study to results using the projected unit credit (PUC) actuarial cost method, since that is prescribed for FASB accounting, the most common reason for actuarial valuations, given the lack of funding valuation activity. Also, the important point is not how present values get attributed but how present values are determined.

Under any method with multiple assumptions, valuation of a complicated plan at a later measurement date will inevitably result in gains or losses stemming from actual experience being different than assumed. If this experience gain or experience loss is recognized immediately at that later date, it will contribute to volatility. Immediate recognition accounts for what has already happened, including changed perceptions of the future and, while there may be reason for spreading gain/loss recognition over time to limit volatility of reported figures, we will leave those to other commentators. All examples and measurements of volatility used here are based on immediate recognition of plan experience, except for brief consideration of a 15-year amortization.

The prominence of option pricing concepts in valuation discussions of recent decades has made measurement of volatility increasingly important. That focus is on volatility implied by option prices for an asset, rather than the volatility of historic measures. This paper addresses volatility coming from actuarial valuation models for long-term retiree health benefits, not implied volatility derived from market prices (there aren't any) and option pricing models. As interesting as such models might be, we are most concerned here with modifying the actuarial measurement model currently in use. It is the volatility that has arisen in the past from that model that pushed the headlines and challenges the actuary, although that past volatility might play little part in appropriately valuing, with an options pricing model, a retiree health benefit containing embedded sponsor options.

That history, of course, does affect the modeling of what we might expect to happen in the future and the uncertainty associated with that expectation. For each of the three possible sets of changes, the past experience does tell us enough to anticipate what may happen in the future and mitigate the risk that valuation results come as a surprise or generate constant volatility. Let's look at each of the three to see what volatility management has been attempted and whether that management might be improved.

Changes in Benefit Level

Benefit level or plan provisions, for this purpose, are defined to exclude eligibility or duration and long-term changes to retiree premium rates. Rather, they define the indemnification rules for a given year, such as services covered, deductibles and other cost-sharing amounts, and administrative procedures. If there are changes, they are announced once a year at enrollment time and there is a general expectation among employees, retirees and plan sponsor that each year will see some incremental change, in response to the health care economy.

This paper will not examine closely the volatility associated with changes in benefit level, because we think this is less susceptible to actuarial solutions and because we think the status quo has adapted adequately to the problem presented. Adjustments to benefit level, however, often tie to the other two sets of changes. For instance, an increase in deductible may be prompted by an increase in claims paid and, in turn, should result in a change in what is reimbursed, if not a change in amounts of claims paid.

It is possible for a health actuary, with sufficient data and time, to closely model the financial effect of such a change for not only the initial year of the change but also all later years. Such a procedure is time consuming, however, and not certain to add insight or accuracy, in part because the following year may see a new set of adjustments to plan provisions and need for remodeling. In response to criticism of its 1989 proposal to require such modeling, FASB included the idea of a "substantive plan" into FAS 106 (now ASC 715-60), and most actuarial modelers have embraced this for long-term projections of costs. The "substantive plan" concept recognizes that the plan as understood for the long-term by the involved parties may differ from the extant written plan provisions. If the plan has shown by past practices that increases in one or another provision are a consistent part of its cost-sharing policy, then the substantive plan includes that practice and such increases can be projected for the future. The other side of this coin is that prospective changes not consistent with past practice cannot be included and projected unless they have been fully communicated to all the participants. Thus, the significant but infrequent plan change is not part of the substantive plan (and for the purposes of this paper, belongs to the third set of changes to be discussed).

The substantive plan concept is a way to smooth valuation results on the front end of the valuation. What would otherwise be a series of small changes to the mathematical model, which would likely be different each year, can be summarized as single cost-sharing practice, consistent across the years. This paves the way for a single measure of the value of a year's benefit that only needs to be trended from year to year. The valuation assumption for that purpose is the initial per capita claims cost, which will be discussed next. In our view, the substantive plan concept adequately manages volatility of this first set of changes, in benefit level, as defined above.

Claim Cost Projections

In looking at the claims experience of a certain population is it quite common for there to be considerable volatility from period to period (year to year, month to month, etc) in per capita claims. Especially in smaller populations and pre-Medicare groups, changes of 20 or 30 percent can be standard and larger fluctuations occur often. When determining an initial claims per capita cost to project for valuation, it is desirable to smooth this volatility and prevent large variations from overly impacting the APBO, creating wild up and down swings from one valuation to the next. Using a normative database is a means of doing that, but it ignores the fact that the group being valued can be consistently different from the norm for many, many years without ever reverting to the norm. The normative approach not only may prove volatile, but also may prove misleading. A better approach is to use claims specific to the plan without giving full weighting to the most recent period of available claims.

One way of doing this is to average or weight various experience periods to create a best fit line out of the scattered data. In the "Modeling Initial Per Capita Health Care Rates" section of ASOP 6, brief mention is made of the option of weighting and this paper will expand on its applications. There are multiple variables to deal with in the process. The number of periods included, their weights and trending can all have significant impact on the resulting initial claims assumption. These variables can be combined to come up with many different outcomes. Actuarial judgment needs to be used and other leading factors should be taken into account. These could include the impact of a few large claims or claimants on the experience volatility, general health trends, plan changes, census makeup and changes, and more.

Here is an example of the impact of weighting. The included Table A shows these comparisons. For simplicity, trending and variability of the weights have been left out for now. This is a scenario showing the effects of using two to five periods in computing initial claims costs, instead of only one period of experience. Each period is given equal weight, i.e., if using five periods of claims experience, each period gets a 20 percent weight.

To generate claims experience, we choose an average per capita amount and then allow it to vary by as much as 50 percent up or down, using a random number generator. The resulting change in per capita claims from period to period has a standard deviation of 41 percent. When these same per capita's are weighted using the current period and the previous period to find the initial claims assumption, the standard deviation of the change in these two period weighted claims assumptions drops to 20 percent. So, this results in about half the amount of volatility as using only one period of experience. Using three periods of weighting the standard deviation is 14 percent, four periods of weighting it is 10 percent, and five periods it is 8 percent.

If the claims experience is not as volatile, then the impact of weighting is not quite as necessary, but it still does have a positive impact. If the claims experience is allowed to vary by 10 percent up or down, then the standard deviation of the change in per capita claims using one period of experience is 8 percent. Using weighting of two periods it is 4 percent, three periods 3 percent, four periods 2 percent and five periods 2 percent.

This method of generating hypothetical claims experience is a little crude, since it is more likely that the variation from period to period in per capita claims would fall on a bell curve of some sort. To test the impact of this we also ran a scenario where half of the claims were allowed to vary by 10 percent up or down, one quarter were allowed to vary by up to 50 percent, and the last quarter allowed to vary by up to 80 percent, thus making outliers less likely. Of course, this is still a simple stepwise function rather than a true curve. Infinite scenarios like this could be tested using various steps, a normal distribution and other methods. The main point would remain the same, using more than one claims experience period to determine an initial cost helps, in most cases and on average, to decrease the period-to-period volatility of the initial cost assumption.

Giving equal weight to more than one period of claims experience has the predictable effect of lessening volatility. Sometimes, giving equal weight to each period is not desirable. Especially when trend is involved and four or five weighting periods are being used, the period furthest back could have a large effect on the results. Judgment needs to be used in selecting how many periods of weighting to use and how much weight to give to each. In the chart below we also include a variety of weighting scenarios. When one period (usually the most recent) is emphasized more than others the volatility starts to creep back up, no matter how many weighting periods are used. Often the most recent period is also the most reliable and/or accurate indicator of future experience though, so the increased chance of volatility in using it more heavily in the weighting may be offset by it being a superior indicator.

TABLE A
Standard Deviation of the Change in Initial Per Capita Claims Cost Assumption

Using Various Experience Periods and Weighting Scenarios

Variability of Hypothetical Random Generated Claims Experience

	50 percent Variability	10 percent Variability	1/2 10 percent 1/4 50 percent 1/4 80 percent
One Experience Period	41 percent	8 percent	30 percent
Two Experience Periods	20 percent	4 percent	15 percent
Even Weights (50 percent)	22 percent	4 percent	16 percent
60 percent, 40 percent	25 percent	5 percent	18 percent
70 percent, 30 percent	29 percent	6 percent	22 percent
80 percent, 20 percent			
Three Experience Periods	14 percent	3 percent	10 percent
Even Weights (33 percent)	14 percent	3 percent	11 percent
40 percent, 40 percent, 20 percent	17 percent	3 percent	12 percent
50 percent, 30 percent, 20 percent	21 percent	4 percent	16 percent
60 percent, 25 percent, 15 percent			
Four Experience Periods	10 percent	2 percent	8 percent
Even Weights (25 percent)	13 percent	3 percent	10 percent
40 percent, 30 percent, 20 percent, 10 percent	17 percent	3 percent	12 percent
50 percent, 30 percent, 15 percent, 5 percent	21 percent	4 percent	15 percent
60 percent, 25 percent, 10 percent, 5 percent			
Five Experience Periods			
Even Weights (20 percent)	8 percent	2 percent	6 percent
30 percent, 30 percent, 20 percent, 10 percent, 10 percent	10 percent	2 percent	7 percent
40 percent, 25 percent, 20 percent, 10 percent, 5 percent	13 percent	3 percent	10 percent
50 percent, 25 percent, 15 percent, 5 percent, 5 percent	17 percent	3 percent	13 percent

The source of some of the variation in claims experience has long been identified as increases in utilization and price level, which gets combined with a few other factors for what most readers know as “Trend.” If all parties are expecting an increase, the occurrence of that increase should not really be considered an example of volatility. In an actuarial valuation of retiree health benefits, we believe the trend assumption used to project forward from the initial claim cost per capita is best determined from normative data (a contrast with the avoidance of normative data we advised when selecting the initial claim assumption), including surveys for short-term trend projections, and economic indicators, including GDP limitations. This will mitigate volatility from trend, since the trend will be anticipated and can be consistent with market expectations.

But there is a second need to use trend, and that is in the analysis of claim history in determining the initial claim assumption. Our study above about the impact of weighting was constructed in the absence of trend; any variations in claim experience would primarily be due to variations in participant utilization from period to period. In reality, of course, many other factors, including those embedded in trend, would have impacted claim costs from year to year in the historical period under examination. To avoid building a source of volatility into the initial claim assumption, the historical experience needs to be adjusted for past trend, either in the weightings or with a separate trend adjustment. Either way, the adjustment to each period should be based on a relevant regional or national trend for that coverage and period, and not just what the plan itself experienced.

Economic Value and Implied Volatility

Before addressing the third and last of these major sets of changes impacting actuarial measurement, brief mention should be made of an aspect that clouds the measurement prospect but also makes it ripe for innovations. This is the fundamental unknowable nature of the value of a complicated long-term benefit that is subject to the successful operation, and occasional whim, of a plan sponsor and/or employer buffeted by events of the world and everyday life. The retiree health benefit exists and there are rules to measure it, but do those rules reflect the reality of the benefit? Retirees or covered dependents may live more than 30 years beyond retirement and the accumulated health cost can be very large. The two sets of changes already discussed—benefit level and claim level—will have an effect over those years, but even if those were transparently clear (which they are not), what financial value should be ascribed to the benefit if the employers' right to change or terminate the benefit prevails? Even as current claims for eligible retirees are being paid and plan procedures for the following year discussed by human resource specialists and actuarial consultants, the question of long-term commitment looms. Does the sponsor really have a liability? Does the retiree really own anything? And what of the benefit will be left for the current employee when he or she is 80 years old? The economic value hinges on answers to such questions, no matter what the measurement rules are. Since the decisions of plan sponsors, employees and retirees will depend as much on their perceptions of that value as on the results of the measurement models, the volatility of that economic value should not be ignored. And since the value of a rescindable benefit includes an embedded option for the plan sponsor, this paper on volatility will note that option pricing, while not a part of existing measurement rules for retiree health plans, offers a separate look at volatility.

Consider the retiree health benefit as an asset owned by the retiree, albeit an asset with a value determined by the sponsor's willingness and ability to indemnify future claim costs. This asset is part of the compensation for the employee's work service. The value of the asset is tied closely (but inversely) to the value of the option held by the sponsor to change the indemnity rules, to even terminate the benefit. Option pricing theory thus may have something to say about the volatility implied in such an asset.

Given two assets of equal current value and an option to buy each at the same price at some future strike date, the option that will be the more valuable, all things being otherwise equal, is the one attached to the asset value expected to be more volatile between now and the strike date. This volatility is prospective and in need of measurement. Measures of volatility can be grouped in two categories, based on historic volatility or implied volatility. Historic volatility is calculated from known values in the past; in the retiree health context, this is, for lack of anything else, actuarial measurements. Implied volatility is prospective and found from option prices and option pricing models. Price in this meaning involves two counterparties, a buyer and a seller. For retiree benefits, the counterparties would be the employee/retiree and the employer/plan sponsor, as they go through the long period from date of hire to date of last eligibility for benefit payment.

There is, alas, no liquid market where retiree health benefits are traded, and thus, no practical way of measuring the price or implied volatility of a particular benefit. We simply know that, if such a market existed, a key pricing ingredient for a rescindable benefit would be the underlying value of the plan sponsor, which holds an option to devalue the benefit when and if it feels the need. We can surmise that the stronger the sponsor's underlying value, the less likely it is to find value in the option to rescind and therefore the greater value in the employee/retiree's asset. In real economic terms of importance to employees, employers and retirees, it may be the benefit's implied volatility that matters, not the historic volatility, especially because the actuarial and accounting measures to date have shed little light on the sponsor's option to rescind,

as will be noted in the next section.[2] Despite this importance, a market is unlikely to develop (unless a passel of Ph.D. candidates apply a binomial lattice to the myriad possible plan design changes and find some buyers) and the R2R measurement approach outlined in following sections of this paper may be the most practical way to determine economic value.

Earlier Valuation Rules

Health cost projections have not faced the significant legal restrictions placed on defined benefit pension cost projections by laws such as ERISA and its offshoots. There were, however, rules laid down by legislative, regulatory and accounting bodies that attempted to systematize how retiree health costs are recognized. Such rules, designed to normalize the way health costs would be parceled out, often sanctioned, knowingly or not, methods that would lead to volatility that threatens the credibility of the measures. This section will examine some of those rules.

The widest reaching of the retiree health rules has been the one coming out of FASB in 1991, FAS 106. But, earlier, during the 1980s, laws concerning tax-advantaged funding imposed rules that differed from the accounting rules. Concern about the huge increases that might be anticipated in health goods and services decades hence, led to an IRS tax-advantaged contribution allowance for IRC Section 501c(9) VEBAs that limited projections to cost levels at the time of each funding contribution. Health care trend was not permitted. (Valuations also usually excluded IRC-denominated key employees, who were to be reimbursed only from separate accounts.) An earlier IRC contribution limit prohibited tax-advantaged funding of health cost that was more than one-third of the pension costs (Section 401h). The reality is that health costs do increase and either IRC method would lead to significant shortfalls in funding.

Later, other accounting bodies had variations in the FAS 106 rules that make comparison with FAS 106 useful for our paper. There were various fixes for the measurement problems that lead to volatility of the year-to-year costs and present values. Most of these were reactions to the dilemma of an uncertain future for the level of claim costs or of the sponsor's commitment to those not yet retired. For instance, if the plan includes active employees but they are not considered vested, then simply not counting any cost associated with the non-vested employees is one measurement solution (NAIC). This leads, however, to the entire cost showing up at the moment of vested eligibility, even though, for example, the individual may need 20 years of service to qualify. (Under current NAIC considerations, the non-vested exclusion is likely to be dropped.)

GASB has taken a different approach, using situational discount rates that may, depending upon the plan's funded status, include the expected return from equity assets and differ in effect substantially from FASB's match of discount rates to high-quality bond market yields. GASB rules also leave open the possibility of quite low discount rates, which to most sponsors will be less attractive for financial reporting of liabilities than it is for justifying tax-advantaged funding (and federal tax benefits are of limited advantage to state and municipal entities). The latter, however, is a consequence of the IRC not allowing very much in the way of tax-exempt investment income for retiree health benefits VEBAs. If the assets grow slowly, more funding is needed up front and a discount rate based on after-tax yields is justified.

The following table simplifies these rules, in relation to key differences.

[2] A critique for a business audience of the accounting approach by one of the authors was published in The Wall Street Journal, February 21, 1992, with the title Ignore the Retiree Health Benefits Rule.

	FAS	GASB	VEBA	401h	NAIC
Non-Vested?	Yes	Yes	Yes	Yes	No
Key Execs?	Yes	Yes	No	No	Yes
Discount Rate	Bond Quality	Low	Low	Bond Quality	Bond Quality
Health Cost Trend?	Yes	Yes	No	Yes	Yes

All these rules have, somewhat predictably, had major flaws in terms of consistency of cost. The basic objective, of preparing today for promised future payments, is advanced when goals such as sufficient contribution funding, cost predictability and intergenerational equity combine with benefit security. Each of the rules approached these goals in different ways, but none of them have the authority to guarantee benefit security. For pensions, security for the retirees was accomplished through laws such as ERISA. But potential beneficiaries receive no such strong guarantees of eventual retiree health coverage from federal laws; broken promises must be adjudicated in court one by one. Many times courts have determined that sponsors have a right to rescind the retiree health benefits, or at least make major changes in eligibility. Courts often found that there is not a clear enough exchange of active work for a retiree health benefit to be an enforceable contract. Such legal findings may have reduced the promised benefit to nothing more than a gratuity, but even if not, the result has been a general absence of benefit security. Many observers, thus, find none of the measurement methods captures the economic worth of retiree health benefits.[3]

Significant volatility may have been mitigated by the way the methods constrained the actuarial projection, but distinctions in plan design and economics were ignored. And in the case of the FASB and GASB rules for financial reporting, the volatility is not mitigated. Methods may not only ignore certain realities (health costs rise, work service may count towards vesting), they also are not reflecting the unilateral control the plan sponsor has reserved to rescind all or parts of the benefit. In contrast to these current rules, this paper advocates another method that has been introduced by one of the authors and received interest as a way to more truly represent the measurement over time of the OPEB promise and, in doing so, reduce some of the accounting volatility.

Alterations in Plan Design and the Sponsor's Right to Rescind

Modifying the valuation model to explicitly measure the impact of a Right to Rescind can mitigate the volatility of periodic plan costs that might otherwise arise from all reasons associated with that legal right. In their prescribed valuation methods, the various rule-making bodies have come up with valuation methods that include preventive ways to tamp down the costs. One did not include non-vested employees but included full health care trend, while another did not allow trend but did include all participants. But each contained the explicit assumption that the substantive plan (to borrow the FAS 106 phrase) would remain in place. Not surprisingly, every time the substantive plan was altered, the cost pattern was altered more significantly. The long-term financial impact was proportionally greater than the immediate cash impact. The PPACA change in tax provisions cited as an example at the beginning of this piece was a mild change in 2010 cash flows. Nonetheless, there were few who were not taken aback by the magnitude of the announced swing in financial cost as measured by the change in FAS 106 APBO. The volatility was great because the valuation results

[3] One way that each of these methods would result in accurate measurement, as viewed in hindsight, would be if the plan sponsor made changes to hold down costs. If, after the tax-advantaged contributions were made, the plan was changed to only pay out those amounts funded, one might find that the limited contributions were sufficient funding for the reduced plan and thus, ipso facto, measurement was accurate under IRC Section 419 (or 401h). Similarly, if the plan dropped any participant not vested at a certain time, the NAIC method would appear accurate. But if the sponsor had reserved the right to make both of those changes as well as other non-specified changes, and later did implement those changes, few would argue that the Section 419 or 401h or NAIC measurement method had anticipated those changes.

gave substantial weight to payments projected decades in the future. The modified model we propose would have shown a much smaller swing, both in absolute and percentage measures. The paper will show how that would work. The method will be compared to FAS 106 figures, but similar superior volatility mitigation can be expected in comparisons with GASB and, less dramatically, with the IRC and NAIC methods. We will trace the APBO, the service cost and a simplified annual expense.

The main effect of a valuation model that recognizes the sponsor's ability to make unilateral changes is to decrease the cost attributed to years later in the duration. It anticipates that some changes will be made to reduce plan payments in the future and reduces the present value of future payments. It transfers the cost load from the front end of the employee/retiree participation to a later time when the payments are more certain, if at that time they have not been eliminated. This lower early cost, in relation to the FAS 106 results, is suitable for methods that incorporate the right to rescind (we will refer to this as the R2R method in contrast to the FAS method used in FAS 106 and the related accounting standards). The employer/employee situation is such that any changes by the sponsor under such provisions will be to decrease benefits, not increase them, since an increase is unlikely without a new economic exchange in which the employees made additional contributions in work or money.[4]

If the benefit is later decreased by the sponsor's unilateral action, the lower R2R cost in the earlier years has been justified—although even with R2R, there may be experience gains. If, on the other hand, the benefit is not later decreased—if it is paid at the same level as if it had been guaranteed—then the attributed cost is shifted to the later years. This is appropriate, however, because the management decision to fully pay the tentative benefit took place later, in comparison with a deferred pension benefit that management guaranteed from the beginning. The economic benefit received in the later years by the sponsor for paying the OPEB benefit may be most succinctly stated as maintenance of reputation and goodwill.

The advantage of moving to an attribution method that allows closer recognition of the actual agreement between plan sponsor and employee/retiree becomes clear when volatility is examined under R2R models. An employer that has provided OPEBs but also has the ability to significantly alter the eligibility for, or level of, the benefits will be subject to variability in expense level beyond variation due to changes in claim cost experience and benefit level. Under FAS, plans which invoked the right to make eligibility changes saw major drops in liability. The resulting experience gains may have been amortized, or absorbed in the gain/loss corridor, but under immediate recognition rules now in place, these changes would have triggered great volatility.

While some employers did drop the plans entirely, they more frequently introduced incremental limitations to eligibility or payout. This section compares the volatility of the FAS actuarial model results, which assumes the certainty of the payout, with a R2R model that includes the uncertainty of that intent. Volatility under the two models will be compared in a hypothetical five-year period in which the only event that does not match the actuarial assumption is the limit of payout/eligibility. The health care benefit being valued is a lifetime benefit, valued at the beginning of each of the five years. The initial valuation assumptions are carried through the following years as valid, with the exception that the substantive plan's provision for payout, or eligibility, will be altered. This alteration restricts the projected benefit in a way that can be considered rescinding the promise, under a right to rescind with legal sanction.[5]

---

[4] Another economic point can be made from the employee perspective. If they are receiving a retiree benefit not guaranteed, they are receiving less value than if the benefit were guaranteed. As part of a rational exchange, therefore, they would be contributing less in work and/or money in the earlier years. This assumes that they recognize the retiree health benefit is offered but not guaranteed, an assumption that is discussed in more detail elsewhere.

[5] Whichever treatment is used—one that ignores the possibility of a rescinded benefit or one that incorporates the rescission possibility—there is clearly a need to disclose the other side of the story. If the balance sheet number became the R2R figures advocated here, there is still a need to indicate what the liability would be if the promise is kept. Alternatively, if FAS remains in

Our proposed model efficiently measures this, with accuracy dependent on the ability to quantify the risk that the provisions substantively in place at the measurement date will be downgraded. A plan that has a 90 percent risk of disappearing or being significantly diminished in the next ten years should show a much lower current cost than one with only a 25 percent risk (under existing measurement rules, both would show an equivalent cost with a guaranteed plan, as though the actual risk is zero percent). The trick is in quantifying the risk[6]. But the proposed method has the potential to be more accurate, while also mitigating the volatility.

Our initial examples use a 16 percent discount rate and compare those results with results under a 6 percent discount, which stands in for a typical FAS 106 discount rate.[7] (6 percent is also the ultimate health care trend rate used in these examples.) The difference in discount rates is equivalent to a 1000 basis point risk premium on top of the 6 percent discount. As applied, this risk premium sets a 8.62 percent chance that the plan would pay nothing after that year versus a 91.38 percent chance that it would pay everything projected. These are rates compounded annually, so the 8.62 percent default probability operates on the "solvent" probability each year. After ten years, there is an almost 60 percent chance the plan will be completely gone versus a 40 percent probability it has survived fully. Other interpretations are possible. For instance, the risk premium could alternatively be seen as setting an 17.24 percent chance that the plan would pay only half of what was projected after that year versus an 82.76 percent chance that it would pay everything projected. Similarly, other combinations of change are encompassed in the model, including an interpretation of a 100 percent chance that the plan would pay only 91.38 percent of what was projected after that year. The 16 percent discount rate may be higher than what is appropriate in many plan situations, but the higher figure provides more contrast to study the differences between R2R and FAS

Let us first note that volatility is not affected by the choice of discount rates, once that choice has been made. In other words, if all other assumptions are met, R2R generates a smooth curve from initial valuation to last date-of-service or last benefit payment in the same way that FAS does, for service cost, APBO, or interest cost. R2R will start at a lower level than FAS and rise more steeply, but there will be a smooth progression over time. On the other hand, if the assumptions do not hold, the volatility level becomes of interest and that is what we examine next. Note, however, that if the assumption of a stable substantive plan (inherent in FAS) does not hold and the plan is reduced from that original level, R2R has less distance to fall than FAS. This is the clearest and most dramatic illustration of how R2R will reduce the volatility that has been witnessed in previous valuation methods.

In the proposed model, the risk premium has a compound discounting effect if it is included for more than one year. The examples are taken from a five-year period and the APBO, service cost and interest cost for each period. A first scenario assumes there are no alterations to the plan during the period; a second scenario posits a major alteration in Year 3. The difference in the Year 3 alteration, in both absolute dollars and

---

place, those sponsors that have reserved the right to alter the promise could be allowed to disclose that fact and the lower value it would imply.

[6] The risk of unilateral rescission is similar to the creditworthiness risk but not the same thing. Also, there is a different level of volatility in an actuarial cost based on a closed group without new entrants than in an open group. One of the limits we have placed on our investigation for this paper is to deal only with closed group volatility.

[7] The approach to modeling economic value discussed here is the discounting of a single projected cash flow with a discount rate appropriate to the risk of the cash flow. Our examples use 6 percent as the low-risk discount rate, although for the theory, it might be seen as a risk-free, or at least default-free, rate, representing the certainty of the payment. The uncertainty of the single payment stream is reflected in the present value through the use of a risky discount rate, 16 percent in our examples. An appendix on economic present values of uncertain future events discusses a more comprehensive approach in the modeling of retiree health benefits, which would project the uncertainty through the probability of the cash flows, rather than through the discount rate. If all of the payment uncertainty could be captured in probabilities of cash flow, the appropriate discount rate(s) would be default-free.

percentage, is compared for the two scenarios. In addition, the relation between the Year 3 change and the changes for the other years is compared for the two scenarios. (Because the effect of ignoring the Right to Rescind is almost always to overstate the liability, the actual significant exercising of that right almost always leads to lowering APBO. Thus, the effect on volatility can be shown by simply comparing percentage change and absolute amounts, without bringing to bear more sophisticated measures of volatility that adjust for ups and downs—standard deviations and the like. There are rarely "ups" with this parameter.) Finally, the absolute dollars in the expense (Service Cost Plus Interest Cost in this simple model) and APBO after the Year 3 alteration in the second scenario is compared to the absolute dollars in the expense and APBO after the Year 2 change in the first scenario. (This does not so much show the volatility of either method as it shows that the R2R model in earlier years anticipated where the value under the FAS 106 would rest once a significant change was made and recognized in Year 3. This is hardly an exact match, but shows how preferable this is to the results of the commonly used model.)

The cash flow projection of benefit payments is made with the provisions of the substantive plan - not only those benefit level provisions governing year-to-year reimbursement policy, but also eligibility and benefit duration provisions. The present value calculation is where the approach accounts for the risk that the substantive plan will be rescinded.[8] In the descriptions below, the variations in provisions for retiree eligibility and benefit duration will be referred to as "alterations" to avoid overuse of the word "change." A major proportion of the narrative in the next sections will be devoted to the first alteration, to illustrate the measures being used and their significance.

An Active Population

The first alteration imposed on the plan being valued was to limit average per capita claim reimbursements for the future to the level anticipated for Year 4. A common alteration from the 1980s to the present, this "cap" would mean that, for years beyond Year 4, per capita annual costs would not be affected by trend and, to some extent, aging. The cap would be announced at the beginning of Year 3 and thus first affect the valuation for that year, although the restriction would not cause any real change in cash flow until Year 5 at the earliest. The result for the active employee population in the FAS 106 model for Year 3 was to reduce the closed-group APBO and interest cost 55 percent and the SC 70 percent. The service cost plus interest cost was reduced 61 percent for this closed group. Under the R2R model, the APBO and interest cost were reduced 28 percent and the SC was reduced 52 percent. The service cost plus interest cost was reduced 31 percent. The next paragraph discusses these changes in valuation results in some detail, for the purpose of explaining what the R2R model accomplishes.

The alteration to the substantive plan, when announced, provides a jolt to the valuation results, taking them far off the assumed path. To the extent that the "jolt" was unanticipated, the valuation results should be considerably different; the alteration justifies the consequent one-time drop in present value (a drop signaling volatility). To the extent, however, that an impartial observer could anticipate an alteration at some point that sharply reduced the worth of the substantive plan, a jolt to the valuation results was also expected. In such a situation, a valuation method with "shock absorbers" that mitigates the volatility is a better method than one without such features. In the example in the last paragraph, continuation of the substantive plan from Year 2 to Year 3 without alteration would be expected, for this closed group of active employees, to result in an increase of APBO, a lesser increase (and possible decrease) in Service Cost, and an increase in Service Cost Plus Interest Cost, which as a percentage increase would be between the other two percentage increases. But the alteration to cap the future claim levels results in a drop in all these indicators and some volatility in

[8] A termination decrement separate from the discount rate may be more easily related to the sponsor's right to terminate a plan than an uncertainty premium embedded in the discount rate, but the uncertainty premium can be handled in most valuation systems without any modification. That is a major reason we chose the uncertainty premium assumption in modeling the effectiveness of the proposed method to mitigate volatility - conceptually and mathematically the two assumptions are almost the same.

valuation results. We suggest that, given the sponsor's legal right to alter the benefits, a valuation method (FAS) that drops the APBO 55 percent and the SC 70 percent is more volatile and less representative than one (R2R) that holds the decreases to, respectively, 28 percent and 52 percent.

This simple measure of volatility should be supplemented by a slightly more complex comparison of the actual percentage change from Year 2 to Year 3 once the alteration was in place to the expected percentage change. The R2R method expects a steeper increase, due to the deeper discounting, than the FAS method. In the example, the APBO was expected to increase 9 percent from Year 2 to Year 3 with the FAS method and 16 percent with the R2R method. The Service Cost was expected to decrease 5 percent from Year 2 to Year 3 with the FAS method and to decrease 3 percent with the R2R method. Measured this way the APBO change decreased 64 percentage points under the FAS method and 44 percentage points under the R2R method. The SC change decreased 65 percentage points under the FAS method and 49 percentage points under the R2R method. The advantage is still with the R2R method, although not as strong.

A third way of measuring the efficacy of the methods, however, shows another advantage of the R2R approach. Let us posit that the new altered substantive plan is, in fact, the ultimate substantive plan, i.e., there are no more substantial alterations to jolt the valuation for the long remainder of the plan life. In that case (unlikely, but theoretically of importance), the accurate Year 3 expense after the alteration is given by the FAS method (because there is no longer an effective right to rescind). An appropriate question to ask is, "Which method better indicated this new level of expense in the year preceding the alteration?" For our simple model, Service Cost Plus Interest Cost will serve as expense. In absolute terms, and indexing the Year 3 expense as 1.00, in our example the Year 2 FAS expense was 2.55 and the Year 2 R2R expense was 1.46. Clearly, the cost impact of the alterable plan was better anticipated under the R2R method than under the FAS method.

For many purposes, annual expense will also include a gain/loss component. If the experience gain in APBO triggered by the alteration were to be immediately recognized or brought into Year 3 expense through a multi-year amortization, the contrast with the Year 2 expense would be even more favorable to the R2R approach regarding mitigating volatility. For the population and assumptions used in this example, the service cost and interest cost would be more than offset by an amortization of the experience gain over the future working lifetime, leading to an expense cost that, all other things equal, would be negative or near zero for a number of years. In more general terms, under the FAS approach, the reduction in APBO is almost 17 times the amount of the new annual Service Cost Plus Interest Cost. Under the R2R approach, the multiple is only about three and a half times. An amortization of 15 years essentially wipes out the FAS expense cost in the near years, even sending the expense into negative territory. This happens while the sponsor cash flow remains relatively constant and might still increase. Under the R2R approach, a 15-year amortization results in just a 23 percent reduction off the new annual expense, not the 112 percent under the FAS approach.

The results discussed above are based on an assumption that the right to rescind extends throughout the life of the plan. Thus, every payment is discounted to the present at 16 percent. The model could easily accommodate other interpretations of the legal right to rescind, such as having the rescindability expire once an employee has retired. In that case, the R2R discount rate would only be applied to the periods before retirement.[9] The contrast in the valuation results would not be as dramatic as above, but would be in the same direction, with the R2R model mitigating the volatility.

[9] Another twist to the R2R model could make the distinction between it and the FAS model even more dramatic. Consider the argument that if the right to rescind extends into the retirement years (an extension bolstered by some court cases and bankruptcy law), then the benefit attribution period should also extend into that period. That would spread the service cost over a longer period and reduce the APBO. The traditional view may reject the possibility that a participant could be "earning" a benefit and generating a service cost while no longer working, but during the retirement years the employer can benefit from the goodwill generated around retirees and retiree goodwill has economic value. If the employer has the right to rescind the benefit but does not do so, the

A Retiree Population

Next we examine a comparison of FAS and R2R models for a retiree group. (The comparison above was for an active employee group. In both cases, they are closed populations with no new entrants. The active group was relatively mature—the FAS APBO was thirty times the service cost and almost one-third of the APBO was for fully eligible employees. This is not too unusual for 2011 and full eligibility available at age 55.). We assumed the plan was rescindable during the retirement years but that the costs were attributed only during the working years, so there was no service cost. Again, we use an alteration of the substantive plan imposing a cap at the beginning of Year 3 that limits future per capita costs to the projected Year 4 level.

The result for the retiree population in the FAS 106 model was to reduce the closed-group APBO and interest cost 32 percent. As we assumed no service cost for retirees, the annual expense of service cost plus interest cost was reduced the same 32 percent as the APBO. Under the R2R model, the APBO and interest cost were reduced 18 percent, as was the expense cost. Again, given the sponsor's legal right to alter the benefits, the inflexible valuation method (FAS) is more volatile and less representative than our alternative (R2R).

The results are similar if the two methods are compared by measuring from Year 2 to Year 3 the difference between the actual percentage change, once the alteration was in place, to the expected percentage change. The FAS difference is 30 percent points and the R2R difference is 18 percent points. For this closed group of retirees, the expected percentage change if the substantive plan continues without alteration is a decrease of APBO, as the benefit payout exceeds the interest cost in most years. The decreases are slight, 1.6 percent under FAS and 0.1 percent under R2R. As with the active group, the FAS APBO is at a higher level than the R2R APBO in Year 2 and the expense is lower as a proportion of APBO.

Under the third measure of efficacy we proposed, however, the R2R approach does not fare as well with the retiree-only group as does the FAS approach. This measure hypothesizes there will be no more substantial alterations and then compares the Year 2 expense under either approach with the Year 3 "right answer," which is calculated by the FAS method (since, under the hypothesis, there will be no further changes arising from the right to rescind). But, in our simple model, expense is Service Cost Plus Interest Cost. The retirees have no service cost, so the expense for the retirees come down to interest cost. In this comparison, that is the Year 2 FAS APBO at 6 percent and the Year 2 R2R APBO at 16 percent. The result is that when the Year 3 expense for retirees was indexed as 1.00, the Year 2 FAS expense was 1.47 and the Year 2 R2R expense was 1.93. For this retiree-only group then, the Year 2 expense under FAS was a closer match to the altered plan Year 3 expense than was the Year 2 expense under R2R.

Nevertheless, that apparent disadvantage to the R2R approach should not be weighed heavily. For one thing, if the sponsor does have the right to rescind during the retirement years, then the expense during the retirement years probably should include some normal cost in addition to the interest cost, as noted in footnote 9. That normal cost will be lower under R2R than FAS and offset the high interest cost of R2R. Also, when the need to account for the experience gain is included, either through immediate recognition or amortization, R2R again bests FAS. Under the FAS approach and without any retiree normal cost, the reduction in APBO is over seven times the amount of the new annual Interest Cost. Under the R2R approach,

---

rationale almost certainly includes recognition that retiree resentment over losing a benefit scores against the enterprise's desire to be well thought of in the community being served. In this argument, a cost of mitigating the reputational risk to community goodwill is maintenance of a retiree benefit that is otherwise rescindable. The discussion of how to measure this cost will be left to another time, but when management is balancing reputational risk and its unilateral rescission right against retirees, there can be little doubt that something akin to a normal cost continues to be incurred during the retirement years.

the multiple is less than two. An amortization of 15 years offsets about 50 percent of the FAS expense cost in the near years. Under the R2R approach, a 15-year amortization results in a reduction of less than 10 percent off the new annual R2R expense. For a retiree-only group, an amortization period of 15 years may be too long, but the R2R advantage would be proportionally the same if comparing the same amortization periods. If the R2R amortization period were tighter than the FAS period, it would cut down some of the advantage, but would certainly not eliminate it.

But the main reason for ignoring any small problem that arises with a retiree-only group is that such groups are still relatively rare and the R2R approach is designed for plans with active employees, particularly those not yet eligible for retirement. In the foregoing sections, attention has been devoted separately to an active employee population and then a retiree population, in showing the distinct advantages of an R2R measurement approach. Most OPEB populations combine active and retired employees. The previous sections give quantitative evidence of the value of the R2R approach in mitigating volatility (among other advantages) for particular examples. We reviewed the same quantitative evidence when the populations in those examples were combined in one group. Not surprisingly, the results from the combined group fall in between those of the separate groups. The detail of the comparisons will not be given here, but the separate results shown above give clues to what might be found with other combined groups that had different proportion of retirees to actives.

For this particular combination of actives and retirees, where the retiree APBO was about a quarter of the total FAS-basis APBO, one finding is worth mentioning. Regarding the measure that hypothesizes no more substantial alterations, so that Year 3 gives “the right answer” and Year 3 expense is indexed as 1.00, for this combined group the Year 2 FAS expense was 2.08 and the Year 2 R2R expense was 1.46. This, the reader will recall, was the one measure by which the FAS approach had an advantage for retirees. For the combined group containing those same retirees, however, the R2R approach is clearly better at approximating what the equilibrium expense will be. And when including the experience gain, the reduction in APBO is a multiple of almost 13 times the amount of the new annual Service Cost Plus Interest Cost under the FAS approach, while the multiple is only slightly more than two under the R2R approach.

Given the sponsor’s legal right to alter the benefits, the cost impact of the alterable plan was better anticipated under the R2R method than under the FAS method, which is more volatile and less representative. We performed the same set of comparisons for a second combined population, with the same alteration (cap at Year 4 level, announced at beginning of Year 3) for that different substantive plan. The results were quite similar, although this was a more mature active population.

Effect of other Substantive Plan alterations

This paper has dwelt on the effect of one alteration—capping future annual payments per capita—to show in considerable detail the effectiveness of the R2R approach in mitigating volatility of APBO, service cost and likely expense results. Several other alterations to substantive plans were also modeled. The resulting comparisons for some were uninteresting and, in hindsight, predictable. If the alteration was one that would be projected to have the same proportional effect over time, say an increase in the percentage of the plan cost paid by the retiree, the impact on the APBO and interest cost is about the same under either the FAS or R2R approaches. This was also largely the case for alterations that might be labeled as “time neutral”, such as when the spouse benefit is reduced. Nonetheless, even in these alterations, the R2R approach shows superiority on the amortization measure. In one typical case, where benefits are cut in half, an amortization of 15 years offsets about 79 percent of the FAS expense cost in the near years. Under the R2R approach, a 15-year amortization results in a reduction of only about 15 percent from the new R2R annual expense.

Two of the more interesting alterations are worth commenting upon for what they show about the accounting, the alterations and the benefit plans generally. The first is an alteration that would increase the years of service required for eligibility from 10 years to 20 years. The second is an alteration that would eliminate eligibility for anyone hired after the year 2000 (this for a valuation at the beginning of 2011, 10 years after the last hire date for eligibility).

The two alterations have very different effect, although they both target a 10-year change. Neither has the financial impact of the cap that has been described in detail throughout this section of the paper, but the differences between the two are illuminating. Eliminating from eligibility those hired in the last 10 years is most likely to effect the youngest of active workers, while raising the years of service needed for eligibility by 10 years has little effect on those employees but is hardest on employees a few years from retirement who will not be able work the years necessary to secure a benefit they may have considered already won. This second alteration would seem to be the harder enterprise decision, with the more immediate impact on cash flow and employee satisfaction. Comparative analysis with our two measurement approaches reveals much about the associated volatility and the incentives for making decisions one way or the other.

The tables below show analysis results for an active employee population from these two interesting alterations, along with the cap alteration described in detail above. For each measure shown, the lower number in the comparison between the FAS approach and the R2R approach indicates less volatility. Table 1 shows the percentage reduction in the three accounting measures from Year 2 to Year 3. Table 2 shows the reduction in percentage points from the anticipated third year percentage change to the percentage change after accounting for the alteration. Table 3 illustrates the effectiveness of the stand-in for annual expense, which is service cost plus interest cost, in three ways. Those three ways are 1) the Year 2 expense divided by the Year 3 expense, 2) the change in APBO from Year 2 to Year 3 divided by the Year 3 expense, and 3) the fifteen-year amortization of that APBO change divided by the Year 3 expense (before amortization).

Table 1 Reduction from Year 2 to Year 3 (Percentages)

Type of Plan Alteration in Year 3	FAS Approach			R2R Approach		
	APBO	Service Cost	SC + Int @ 6 percent	APBO	Service Cost	SC + Int @ 16 percent
Claims capped after Year 4	**55.0 percent**	**70.3 percent**	**60.9 percent**	**27.9 percent**	**52.3 percent**	**31.3 percent**
Service raised to 20 years from 10	**8.4 percent**	**5.6 percent**	**7.3 percent**	**11.0 percent**	**9.3 percent**	**10.7 percent**
Hires after 2000 not eligible	**17.1 percent**	**62.9 percent**	**34.7 percent**	**9.1 percent**	**66.2 percent**	**17.1 percent**

Table 2 Reduction from expected Year 3 to altered Year 3 (Percentage points)

Type of Plan Alteration in Year 3	FAS Approach			R2R Approach		
	APBO	Service Cost	SC + Int @ 6 percent	APBO	Service Cost	SC + Int @ 16 percent
Claims capped after Year 4	**64.4 percent**	**65.0 percent**	**64.6 percent**	**44.3 percent**	**49.0 percent**	**45.0 percent**
Service raised to 20 years from 10	**17.8 percent**	**0.3 percent**	**11.1 percent**	**27.4 percent**	**6.0 percent**	**24.4 percent**
Hires after 2000 not eligible	**26.5 percent**	**57.7 percent**	**38.5 percent**	**25.6 percent**	**62.9 percent**	**30.8 percent**

Table 3 Effectiveness of Service Cost + Interest in:

	Anticipating Change Yr 2 / Yr 3		Limiting Reductions APBO change		Presenting Expense Amortization	
	FAS	R2R	FAS	R2R	FAS	R2R
Claims capped after Year 4	**2.55**	**1.46**	**16.85**	**3.47**	**112.3 percent**	**23.1 percent**
Service raised to 20 years from 10	**1.08**	**1.12**	**1.97**	**1.65**	**13.1 percent**	**11.0 percent**
Hires after 2000 not eligible	**1.53**	**1.21**	**4.16**	**1.66**	**27.7 percent**	**11.0 percent**

The left-most results in Table 1 show the usual FASB/GASB viewpoint on the possible benefit alterations. The cap has the greatest cost effect, as measured traditionally, while raising the years of service has the least, with only half the impact on APBO as eliminating new hires. In terms of the expense measure, raising the service years reduces Service Cost + Interest Cost only 7.3 percent, while eliminating new hires reduces it 34.7 percent. These are also measures of volatility. The remainder of the results in the three tables shows more nuanced consequences. First, we find in the R2R results in Table 1 that, whereas R2R mitigates the volatility of two of the alterations, it seems to increase the volatility from raising the service requirement. The increased volatility is not significant (and shows up more in Table 2 than in Table 1 or Table 3), but under R2R, there is actually a greater reduction in APBO from raising the service than there is from eliminating the new hires.

The reason for this reversal is found in the alterations themselves and the different ways the two approaches handle short-term cash flows versus those that are more distant and, when rescindable, more speculative. The FAS approach gives great weight to the distant impact of eliminating the benefits paid to new hires who will not retire for decades. The experience of the last 20 years, however, has shown that employees, and the courts and the public at large, are not inclined to value those future benefits as much as the more immediate take-away from those soon to retire. This might be considered by some as short-sighted, but it is understandable as a reaction to evidence that sponsoring employers have not put sufficient (or any) funds aside for payment of such distant benefits. The lower right figures from Table 3 indicate that, with our particular assumptions, the percentage effect on expense after amortization is exactly the same, 11.0 percent, for both of these alterations.

Practical conclusions regarding the R2R alternative

This analysis shows a contrast between the R2R method and the more usual FAS 106 method that would be duplicated in comparison with other methods in use, be it the NAIC or IRC methods. The percentage change and absolute amounts demonstrate less volatility in the R2R method and the fact that the post-alteration APBO in the traditional view is closer to the pre-alteration APBO of the R2R method than the FAS 106 method shows that, in absolute amounts, the R2R is no less realistic than other methods. The R2R method is comparatively good at mitigating volatility and a moment's reflection by anyone familiar with OPEB valuations and the compounding effect of discount rates will reveal the general reason. Lifetime OPEB benefits are to be paid far into the future and, while discounting to a present value mitigates some of the effect of the long duration, incorporating the very real prospect of substantial health cost trend into a valuation has the opposite effect—it magnifies the impact of the far future. Ignoring trend may mitigate the volatility some, but it also leads to unreal projections. For rescindable plans, the most effective way to get useful year-by-year projections, but also have present values representative of economic reality, is to discount using a rate considerably higher than the usual low-risk rates of high-quality bonds.

Use of a higher discount rate equates with discounting risky cash-flow streams with rates of return that include equity premiums as an add-on. Cash flows that are not guaranteed, and with little or no collateral, are equity risks. This is basic discounted cash flow analysis. How to determine an equity premium troubles those more comfortable with bond rates, but it should not be a major deterrent. As noted in our example, the

starting point is the probability of defaulting on the full promise of the substantive plan and making that an add-on to a market-consistent default-free rate. For many a corporation, this would lead to a discount rate in close proximity to the targeted internal rate of return. To the extent that discount rate for OPEBs becomes a public disclosure (as with an accounting standard), a link with internal rate of return provides a good story line in response to criticism about tolerating default—if the corporation maintains its targeted internal rate of return, it will also maintain its OPEB promise and not default on it.

While each enterprise would set its own discount rate, and one solution for an individual entity in mitigating volatility would be a very high discount rate, the possibility of a myriad of high discount rates for OPEBs when all other discount rates are much lower should not be taken too seriously. The main implied regulators would be public disclosure and the resulting reluctance of any one enterprise to be seen as an outlier on the high side, given the obvious link between the OPEB discount rate and the probability of default on other obligations of the enterprise. Market analysts would be quick to pick up on one company using a 20 percent discount rate when all others are in the neighborhood of 10 percent to 12 percent. A lesson along these lines can be found in the history of OPEB accounting and the use of health care trend rates. While there was initial concern about how companies would choose health trend rates, after a few years an “acceptable” range became known.

Summary

As the spring 2010 business headlines showed, current financial reporting and disclosures for retiree health benefits can exhibit financial volatility that exaggerates underlying economic realities. Actuaries should be aware of the sources of the resulting problems and alert to misleading conclusions. To aid in mitigating the year-to-year volatility in measuring present values of retiree health promises, in this paper, we target two sources of that volatility. For the claim cost problem, we propose a weighted average of experience and comment on the ways that might be approached, using an increasing realistic model for the initial claim cost assumption and trend assumption. For rescindable benefits and the unilateral change problem, we propose an uncertainty premium operating by way of a higher discount rate and comment on the way that would be approached. The important point is not how present values get attributed but how present values are determined.

Appendix Modeling present values of payments, certain and uncertain

In the latter parts of this paper, we contrast the results of determining present values using two different discount rates—a low-risk rate standing in for high quality fixed-income investment yields and a high- risk rate standing in for the uncertainty we feel should be associated with future payment of rescindable benefits. Our examples use 6 percent as the low-risk discount rate, although for the theory, it might be seen as a risk-free, or at least default-free, rate. This appendix note point out when the use of a default-free rate would be appropriate in the modeling of retiree health benefits even if rescindable benefits were being valued.

The approach to modeling economic value discussed in the main body of the paper is the discounting of a single projected cash flow with a discount rate appropriate to the risk of the cash flow. A more comprehensive approach, however, would estimate the uncertainty not through the discount rate but through the probability of the cash flows. Take an example where a payment of amount X from enterprise A to person B will happen if and only if event Q happens at time T. If Q does not happen at T, the transaction has zero value. In this simple case, the alternative value is zero and any positive value depends on Q happening at T. Determine the probability of that happening, multiply it by X to get the expected value, and discount it from time T at the default-free rate to get the present value.

Why is the discount rate the default-free rate? Because the payment is contingent only on event Q happening. If Q happens, payment is certain; if Q does not happen, it is certain there will be no payment. As far as the parties to the transaction are concerned, the only uncertainty, or risk, in the payment is whether Q happens. Since the formula to determine the present value places all the payment risk on the probability of Q happening, no risk should be placed in the discount rate. In other words, the payment will not default, so the default-free rate must be used. (Volatility is associated only with the event Q, not the payment.)

If the payment is certain, but the amount X is at risk, probabilities about the amount come into the formula. The appropriate discount rate, however, is still the default-free rate. If the payment is certain, but the timing is uncertain, probabilities about the timing enter the formula, but the formula discount rate is still the default-free rate appropriate to the timing. In stochastic models where amounts and timing are uncertain but comprehensive probabilities have been assigned to them, the appropriate discount rates are the default-free rates.

When the payment is not certain, but a probability of a zero payment can be assigned, the default-free rate can be used in the formula. To put it another way, when the probability of default can be assigned, the formula should still use a default-free rate, because once the probability is assigned, there is no longer uncertainty *within that portion of the model*. For mathematical completeness, the discounting should assume certainty. In real life financial situations, achieving such complete probabilities is usually tedious at best and often impossible. We fall back to assigning probabilities of which we are uncertain and amounts which verge on “best guess.” Even in such cases, if the probabilities add to 100 percent, the formula should use default-free rates. So, if we had the ability to assign probabilities to all the various plan design change events and their timing and financial impact, the discount rates by which the present value of the benefit plan would be found are default-free rates.

With something as complicated as predicting the amount and timing of health care payments, the possibilities are large in number, each with a probability quite small (although they would add to 1, if they could all be identified). Such modeling would take us beyond the binomial latticing and option pricing theory mentioned in the body of the paper as unlikely from a practical standpoint.. Instead, we fall back on a deterministic model with a single cash flow at any given point in time and, knowing that cash flow stream is uncertain and not 100 percent probable, discount with a rate above the default-free rate. That is the only way the deterministic short-cut can come up with an answer the same as the more complicated and careful probability

analysis. This discount rate is said to carry a risk premium, the difference between it and the default-free rate. If there is a probability that the payment will be zero and that is not reflected in the deterministic cash flow, the risk premium can be large.

Let's construct an example that fits the assumptions used in the examples in the paper. In a situation where $1000 is due in a year, with the amount and timing certain, and a default-free bond that matures for $1000 in a year is available at a price of $943.40, then absent any other costs of the transaction, the present value is $943.40, implying a default-free rate of 6.0 percent. If the timing of the payment is certain, but the amount is uncertain, the appropriate discount rate is still the default-free rate. For instance, the payment will be made in a year, but it may be $1000 or it may be $1500. The appropriate probabilities must add to 100 percent - say, 70 percent for $1000 and 30 percent for $1500 - in which case, the appropriate discount is 6.0 percent, giving a present value of $1,084.91

This remains true even if one of the two payments is zero; although discounting a zero payment at any rate still gives zero as a result. For our example, if the payment is to be made in a year, it will have an 8.62 percent chance of default and a 91.38 percent chance of $1000. The expectancy in a year is $913.80 (8.62 percent x 0 + 91.38 percent x 1000) and since the calculation encompasses all the probabilities, the appropriate discount rate is the certainty rate of 6.0 percent and the present value is $862.08. But the deterministic model used in the FAS and R2R valuations does not put the probability of default in the payment stream; it only has the payment projection from the substantive plan, in this example, $1000 in one year. To determine the economic value, default needs to be taken into account. Since the probability of default is not in the cash flow projection of $1000, the uncertainty must be reflected in the discount rate of the model, to arrive at the appropriate present value. The risky discount rate is found by solving the equation, $ 862.08 = $1000 / (1+d). (Yes, we teach actuarial students that the proper equation is 1-d, but in the financial world it seems to be 1+d.) The risky discount rate is 16.0 percent, so that the risk premium is 10 percent, or 1000 basis points. That is the discount rate necessary in the discounted cash flow analysis of the deterministic payment stream that is not certain. Actually finding the appropriate risky discount rate can come from a study of market prices, or pulling a number out of thin air, or a combination of the two, which is where our number was hatched.