

Utilization: A Review of Two Projection Methods

By Jeremy Hamilton and Tim Kempen

INTRODUCTION

As actuaries, often our role is to quantify future risk, which may involve estimating future claims. Usually estimating future claims entails using historical data as a starting point to develop an assumption about the future. However, historical data and trends may not be enough to develop an accurate projection. Also, historical data may not have been captured as cleanly or in as much detail as we would have liked. This can make our job a bit more challenging.

Developing financial projections of long-term care (LTC) insurance utilization is no different. The assumptions and methods used to develop utilization projections can have a significant impact on estimated future claims. Also, the quality of historical data may limit what the actuary is able to do.

This article serves as a follow up to the article “Utilization: Long-Term Care’s ‘Middle Child,’” published in the December 2017 issue of *Long-Term Care News*.¹ That article covered much of the background on what utilization is, how it’s calculated from historical data, and what it means for LTC insurance products. This article focuses on two methods for using current utilization levels to develop utilization assumptions for future durations: an “average utilization” method and a “distribution” method. Each method has its own advantages and disadvantages, with the trade-off between the two (not surprisingly) being simplicity versus accuracy.

AVERAGE UTILIZATION METHOD

The more common method for projecting utilization in future durations is to simply trend the current average utilization level forward. Essentially, this method projects utilization for any given duration by multiplying the preceding duration’s utilization by a cost of care inflation assumption and dividing by the amount benefits grow by. For example, if the current level of utilization for a block of LTC policies is 75 percent, benefits inflate by 5 percent per year, and cost of care is expected to grow by only 3 percent per year, then the projected utilization for the following year would be:

$$75\% \times 1.03 / 1.05 = 73.6\%$$

Continuing with the same benefit inflation and cost of care assumptions, the projected utilization for the next five years is shown in the table in Figure 1.

Figure 1
Projected Utilization, 5% Inflation Protection

Year	Utilization
Current	75.0%
1	73.6%
2	72.2%
3	70.8%
4	69.4%
5	68.1%

Similarly, if there is no benefit inflation, the average utilization would increase each year by 3 percent and the projected utilization for the next five years would be as shown in Figure 2.

Figure 2
Projected Utilization, No Inflation Protection

Year	Utilization
Current	75.0%
1	77.3%
2	79.6%
3	82.0%
4	84.4%
5	86.9%

One advantage of using this method to project utilization is that it is fairly simple to implement. Also, to develop the starting utilization assumption, this method relies on paid claims data, which is generally readily available.

While this method may seem simple enough, there are a number of issues to consider. First is the calculation of current utilization. As the “Middle Child” article summarized, a number of nuances such as service periods, benefit payment types, care sites, etc., can affect how utilization is calculated from historical experience.

The second issue to consider is the theoretical limitations of utilization.

- For policies with no benefit inflation, will utilization increase until it hits 100 percent, or will it level off before then? Will policyholders “price shop” for care providers in order to preserve benefits as long as possible or in order to ensure their policies will cover actual expenses? The impact of policyholder behavior could consequently suppress average utilization.

The distribution method, while more complex than the average utilization method, provides a more accurate depiction of how average utilization will change over time.

- For policies where benefit inflation exceeds the projected cost of care inflation, will utilization decrease indefinitely or reach a theoretical minimum? If policyholders’ benefits far exceed the average cost of care in their area, will they gravitate toward more expensive care providers or use more home health services, because they have the policy benefits to pay for it? In this scenario, policyholder behavior could slow the decline in average utilization.
- How does plan design affect average utilization? One example of plan design that could impact utilization is “daily” versus “monthly” reimbursement. Policies with daily reimbursement may have a lower theoretical limit for home health care utilization (ex: 70 percent) than monthly reimbursement, because care may not be received every day.

A third issue to consider is that the average utilization method is simplistic in nature and relies on one average utilization assumption as a starting point. Underlying the average utilization for a group of policies are many policies with varying

levels of utilization. The change in utilization from one duration to the next for each of these policies will likely differ from the others.

Consider a group of policies with no benefit inflation, where half of the policies currently use only 50 percent of their available benefits while the other half use 100 percent. The average utilization for that group of policies would be 75 percent. Trending the average utilization forward by a cost of care inflation assumption of 3 percent will result in 3 percent higher projected utilization each year until a theoretical maximum is reached.

However, utilization for half of the policies is already at 100 percent and cannot increase in future durations. Only the utilization for the half of the policies at 50 percent utilization will increase. Thus, actual utilization will increase much more slowly than what was projected using the average utilization method.

The opposite is true for a group of policies with benefit inflation that exceeds the cost of care inflation assumption. Of the policies that currently have 100 percent utilization, it is unknown what impact benefit inflation or cost of care inflation will have on utilization. Those policies where the current cost of long-term care services exceeds policy benefits may continue to have 100 percent utilization in the future. For example, if the cost of care is \$200 per day but the daily benefit of a policy is only \$100, utilization will remain at 100 percent until the daily benefit catches up to the actual cost of care. However, if the cost of care is equal to or just above the daily benefit, current utilization will be 100 percent but will fall below 100 percent relatively quickly. Thus, it is difficult to project how utilization will change for the 100 percent utilizers. As with the group of policies with no inflation, the actual change in average utilization is likely to be slower than what is projected using the average utilization method.

DISTRIBUTION METHOD

The distribution method, while more complex than the average utilization method, provides a more accurate depiction of how average utilization will change over time. As the name suggests, the distribution method relies on using a distribution of the underlying utilization rates for a group of policies, rather than a singular, average utilization as a starting point.

With the average utilization method, paid claims are subject to reimbursement limits, which prevents utilization from exceeding 100 percent even if billed charges exceed the maximum benefits available. As stated earlier, it is unclear what impact cost of care and benefit inflation will have on the 100 percent utilizers. The distribution method addresses this issue by calculating utilization differently. The observed utilization is developed by dividing actual billed charges, rather than paid

Figure 3
Projected Utilization, No Inflation Protection

Original Utilization Bucket	Bucket Weight	Projection Year				
		Current	5	10	15	20
0% - 20%	5%	10%	12%	13%	16%	18%
20% - 40%	5%	30%	35%	40%	47%	54%
40% - 60%	10%	50%	58%	67%	78%	90%
60% - 80%	20%	70%	81%	94%	100%	100%
80% - 100%	60%	90%	100%	100%	100%	100%
Average		75%	84%	88%	91%	93%

benefits, by the maximum possible benefit. This results in a distribution that includes utilization rates above 100 percent.

MECHANICS OF THE DISTRIBUTION METHOD

In general, the distribution method requires a three-step approach:

1. Calculate observed utilization for each claim.
2. Develop a distribution of observed utilization.
3. Project average utilization by applying benefit inflation and cost of care assumptions to the distribution.

The first step of calculating utilization for each claim can be a bit more difficult than with the average utilization method, because it relies on billed charges rather than paid claims. As stated earlier, billed charges are used in order to develop a distribution with utilization rates above 100 percent. Because companies often do not track billed charges, paid claims could still be used to develop the distribution for utilization below 100 percent. However, other data sources and judgment would be needed to expand the distribution above 100 percent for the 100 percent utilizers in the second step.

The second step is to develop a distribution of utilization. Practical considerations, such as the credibility of historical data, should be taken into account when determining the level of precision of the distribution. A continuous distribution might be difficult to develop so a discrete distribution may be more appropriate. In practice, this refers to the number of buckets the current utilization is split into. For example, when using 5 percent increments, data would be divided into buckets of 0 percent to 5 percent, 5 percent to 10 percent, etc. In the absence of actual charge data, utilization near 100 percent would need to be extrapolated, using actuarial judgment, to develop the upper tail of the underlying actual charge distribution.

In addition to the level of precision to use for bucketing utilization, the actuary will need to determine what characteristics the utilization distributions should vary by (this also applies to the average utilization method). Obvious characteristics for using separate distributions would be care setting and inflation type. But gender, daily benefit amount, benefit period, and claim incurral age could also be considered, among others.

Another consideration for setting the starting distribution would be the impact of trend over the experience period. Unless a company has robust experience, the experience period used to set the starting distribution will likely span several years. This experience would need to be adjusted for historical benefit and cost of care trend to be used as the starting distribution.

The third step is to project each bucket within the distribution separately using the appropriate projection factors for cost of care inflation and benefit inflation. The average utilization for a given duration is then calculated as a weighted average of the resulting utilizations for each bucket and the weights for each bucket. The table in Figure 3 illustrates the resulting average utilization calculated using the distribution method for a hypothetical starting distribution with 3 percent cost of care inflation and no benefit inflation. As the higher utilization buckets reach 100 percent, the increase in average utilization slows.

COMPARISON OF TWO METHODS

As you can see from the examples provided earlier, the distribution method requires a bit more effort than the average utilization method, especially if billed charge data is not available. Often times the distribution method will produce substantially different utilization rates than the average utilization method. However, there may be instances where the difference between the two methods is immaterial (or nonexistent when the cost of care trend and benefit inflation rates are equivalent.) The shape of the distribution and starting average

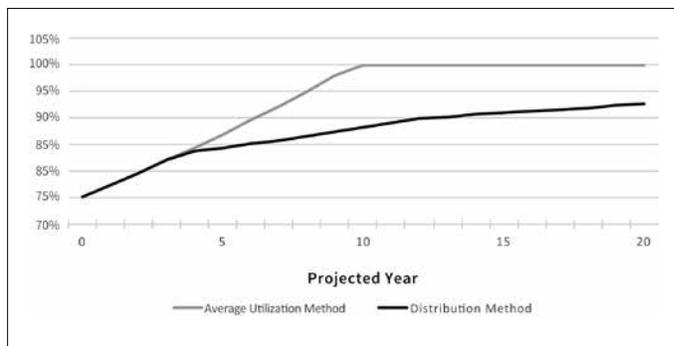
utilization will influence how large the difference in utilization rates is.

For example, cells with a uniform distribution will likely have a greater difference in projected utilization between the two methods than cells that have a more compact distribution. Likewise, cells with very low current utilization may not see materially different projected utilization for many years, whereas cells with high average current utilization may see differences between the two methods much more quickly.

The graph in Figure 4 compares the projected average utilization using each method and assuming the same distribution, starting utilization, and cost of care inflation assumption as Figures 2 and 3.

The average utilization method assumes that utilization increases by a constant factor until an upper limit is reached (in this case, 100 percent in duration 10). The distribution method projected utilization increases at a much slower rate because the portion of the distribution at or near 100 percent utilization no longer increases above 100 percent when projected forward. The resulting slower growth in utilization using the distribution method relative to the average utilization method leads to approximately 12 percent lower utilization in duration 10. The difference in utilization then slowly grades off over time. However, this difference in utilization does not translate to an equal difference in claims costs. The implied reduction to claims costs from the lower utilization may be partially offset by an extension of benefits.

Figure 4
Project Average Utilization, No Inflation Protection



IMPACT OF COINSURANCE

The distribution method can also be useful in analyzing the utilization impact of coinsurance features where X percent of actual cost is reimbursed, subject to a daily benefit maximum cap. To measure the savings associated with the coinsurance feature, the distribution of charges is more meaningful than an average charge. The table in Figure 5 illustrates the calculated impact of 10 percent coinsurance on a cohort with 80 percent average utilization using the two methods. For simplification, this example assumes half of the cohort has 60 percent utilization and half has 120 percent billed charges utilization (100 percent paid benefits utilization).

As expected, the resulting utilization using the average utilization method is 10 percent lower ($8/80 = 10$ percent) after

Figure 5
Utilization Impact of 10% Member Coinsurance

	Average Utilization Method		Distribution Method	
		60% Bucket	120% Bucket	Weighted Average
Distribution Weight	100%	50%	50%	
Average Charge	\$80*	\$60	\$120	\$90**
Daily Benefit (DB)	\$100	\$100	\$100	\$100
Insurer Paid Amount (utilization)	\$80	\$60	\$100	\$80
Member Coinsurance	\$8	\$6	\$12	\$9
Remaining Charge	\$72	\$54	\$108	\$81
Insurer Paid After Coinsurance (utilization)	\$72	\$54	\$100	\$77

* Average utilization method uses paid benefits subject to daily benefit cap.

** Distribution method uses average billed charges.

implementation of the 10 percent member coinsurance. The distribution method, however, accounts for the impact of the maximum daily benefit. The 10 percent member coinsurance effectively has no impact on benefits paid for the 120 percent utilizers, because charges net of coinsurance still exceed the maximum daily benefit. Consequently, the 10 percent member coinsurance only reduced utilization 3.75 percent ($3/80 = 3.75$ percent) when using the distribution method.

POLICYHOLDER BEHAVIOR

A question that often comes up when discussing future utilization is how policyholders will behave when it comes time to use their benefits. Will they price shop and look for care providers that fit within their policy benefits? Or will they look to preserve their benefits as long as possible? A byproduct of the distribution method is the underlying analysis of billed charges, which can shed some light on policyholder behavior.

A distribution of billed charge utilization that shows a high concentration near 100 percent, but very few above 100 percent, could indicate that policyholders price shop and actively look for care providers that cost less than their policy benefits. If this pattern persists when looking at multiple time periods, this may further strengthen the hypothesis of policyholder price shopping.

Similarly, if the distribution of billed charge utilization does not change much over time, and is not concentrated near 100 percent, this could indicate policyholders attempt to preserve benefits.

CONCLUSION

The trade-off between simplicity and accuracy is something actuaries will always need to consider, particularly as the long-term care insurance industry matures and pays more attention to projected utilization. While in some instances the additional degree of accuracy of the distribution method versus the average utilization method may be minimal relative to the added complexity, in other instances the distribution method may produce materially different results. The analysis of utilization using distributions can also enhance an insurer’s understanding of policyholder behavior and assist in the pricing of new product features. ■



Tim Kempen, FSA, MAAA, is a consulting actuary at Milliman. He can be reached at tim.kempen@milliman.com.



Jeremy Hamilton, FSA, MAAA, is an actuary at Milliman. He can be reached at jeremy.hamilton@milliman.com.

ENDNOTES

- 1 Bergeson, M. & Emmert, M. (December 2017). Utilization: Long-Term Care’s “Middle Child.” *Long-Term Care News*. <https://www.soa.org/Library/Newsletters/Long-Term-Care/2017/december/ltc-2017-iss46.pdf>.