Insurance Risk and Its Impact on Provider Shared Risk Payment Models

January 2018

Insurance Risk and Its Impact on Provider Shared Risk Payment Models

AUTHORS Juliet Spector, FSA, MAAA
Cory Gusland, FSA, MAAA
Carol Kim

SPONSOR Health Section

CONTENTS

ACKNOWLEDGEMENTS

The authors wish to thank the Society of Actuaries for sponsoring this paper. In particular, the authors thank the Project Oversight Group for their assistance in preparing this paper. The Project Oversight Group is comprised of Daniel Bailey, Matthew Harker, Amanda Holland, Patrick Kinney, David Petruzzellis, Daniel Pribe, Brandy Sneed, and Coleen Young. In addition, the authors gratefully acknowledge the assistance of Rebecca Owen, SOA Health Research Actuary, throughout the course of writing this paper. Thanks also to Steven Siegel, SOA Research Actuary, and Barbara Scott, SOA Sr. Research Administrator, for managing this effort and helping to bring it to completion.

I. EXECUTIVE SUMMARY

More pressure is being placed on providers to assume financial risk in their contracting arrangements. As more providers take on downside risk, they expose themselves to insurance risk, which is the risk associated with the unknown and unpredictable variation in utilization and cost of services. It includes both the random and nonrandom factors that cause the best estimate of expected incurred claims to differ from actual claims.

A challenge facing providers and payers is to design shared risk payment models that incentivize providers to deliver efficient, high-quality care without assuming too much insurance risk from payers. Transferring the appropriate amount of insurance risk from payers to providers can be particularly challenging when stakeholders' understanding of concepts varies and there is a lack of common terms or limited objective measurement techniques.

Providers and payers should be measuring insurance risk when assessing the viability of a risk-sharing contract and creating appropriate risk contracting parameters to mitigate that risk. These pricing exercises should be approached with appropriate actuarial rigor. Performing stochastic simulations is one technique that can help shared risk participants answer the following questions:

- What is the best estimate of our future performance?
- What is the likelihood of savings?
- What is the range of possible outcomes?
- What is the risk of loss?
- How can altering payment model design impact the randomness and range of the results?

A shared risk arrangement or other value-based payment model may specify parameters for the claims cost target, risk corridor, patient attribution, risk adjustment, stop-loss, care management and included services. Small changes in any of these parameters may have a big effect on a provider's risk exposure. And two providers with seemingly similar underlying populations may experience very different results, even if they choose similar parameters. As a result, part of the feasibility study of a risk contract should involve measuring insurance risk and modeling the impact of these parameters. Providers and payers can then consider whether the contract or arrangements meet their needs by considering the five questions above in the context of their organizations' risk appetites and strategic goals. Data-driven decisions will ultimately lead to more successful, sustainable arrangements and more appropriate contracts for all stakeholders.

II. INTRODUCTION

All payment arrangements, whether standard fee-for-service (FFS) or value-based, involve financial risk[1] (as well as financial opportunity), and no one payment structure is the best in all circumstances. Because financial risk is organization-specific and difficult to generalize, we will use the same framework used in our prior paper to categorize risk into four main types: performance risk, technical risk, utilization risk and insurance risk.[2]

Performance risk relates to inefficiency and suboptimal quality of the delivery of health care services. Most models will assume some savings in claims cost dollars (i.e., reimbursement to the provider) due to care management; however, these savings may never be realized if care management is implemented poorly or not at all.

Underlying every provider risk arrangement is a contract. *Technical risk* is the risk of inappropriately structuring technical elements of the contract to match the covered population and provider-specific circumstances. Models with too much technical risk are not easy to implement or monitor. Elements that contribute to technical risk include attribution methodologies, cost target development, choice of trend assumptions, risk adjustment and so on.

Utilization risk refers to how the payment model is affected by the known changes in utilization. For example, low utilization results in lower payments to providers in FFS environments where providers are reimbursed by payers for each service they provide. The impact of changes in utilization (volumes) on provider profitability depends on the relationship between payments and operating costs (variable costs).

Insurance risk is the risk associated with the unknown variation in the utilization and cost of services. Insurance risk includes random variation but also variation that cannot easily be predicted, such as changes in acuity level. This category encompasses all risk that cannot be categorized as performance risk, technical risk or utilization risk.

A challenge for providers and payers is to design shared risk payment models that incentivize providers to deliver efficient, high-quality care without assuming too much insurance risk from payers. Transferring the appropriate amount of insurance risk from payers to providers can be particularly challenging when stakeholders' understanding of concepts varies and there is a lack of common terms or limited objective measurement techniques. We will address some of these challenges in this paper.

In Section III, we introduce insurance risk and build the case for its relevance in provider shared risk payment models. We then shift our focus in Section IV to outlining practical methods of measuring insurance risk and provide examples. Section V describes several common examples we encounter when helping providers evaluate shared risk payment model designs.

Much of the inspiration for this paper comes from frequently asked questions from providers, which actuaries are well-positioned to answer. These questions include the following:

1. What is the best estimate of future performance?
2. What is the likelihood of savings?
3. What is the range of possible outcomes?
4. What is the risk of loss?
5. How can altering payment model design impact the randomness and range of the results?

---

[1] *Risk* is loosely defined as exposure to harm or loss.

[2] Juliet M. Spector, Brian Studebaker and Ethan J. Menges, "Provider Payment Arrangements, Provider Risks, and Their Relationship with the Cost of Health Care," Society of Actuaries, October 2015, https://www.soa.org/research-reports/2015/2015-provider-payments-arrangements-risk.

These questions will be asked more frequently as providers take on more downside risk. In addition, providers may need to consider how they can best prepare for downside risk. Where appropriate, this may include establishing or increasing reserves.[3] The techniques presented in this paper may be useful to measure the potential magnitude of downside risk.

This paper provides illustrations using a two-sided shared risk arrangement. In this arrangement, we will assume that there is a per member per month (PMPM) claims cost target (i.e., benchmark). The provider then will share in either savings in the performance period (when claims costs are lower than the benchmark) or losses in the performance period (when claims costs are higher than the benchmark). In some cases, the shared risk arrangement will have a risk corridor around the claims cost target where the provider will not share losses or savings. All examples presented in the discussion will focus on "total cost of care" shared risk payment models using a commercial (i.e., non-Medicare, non-Medicaid) population. We note that the concepts apply as well to other value-based payment models (such as episode-based payments and partial capitation) and populations (including Medicare and Medicaid). We selected this payment model structure and population because they currently represent one of the most frequently seen value-based contracting arrangements in the industry.

The illustrative case studies presented in this report may indicate the potential advantages of choosing certain parameters in a risk contract over others. However, it is important to note that other organizations implementing similar payment models could achieve materially different results, in both magnitude and direction. There is no "one-size-fits-all" shared risk arrangement. Also, a robust analysis is predicated on having good-quality data. Stakeholders should endeavor to maintain credible data sources and reconcile data when transferring the information between both parties.

[3] Reserves are set aside to cover losses not yet paid. Projections are based on many assumptions; therefore, calculations of reserves should not be calculated without the help of appropriate professionals.

III. DEFINING INSURANCE RISK

Insurance risk is the risk associated with the unknown variation in the utilization and cost of health care services. It includes both the random and nonrandom factors that cause the best estimate of expected incurred claims to differ from actual claims. Some drivers of insurance risk include the following:

1. Change in the attributed or covered population
2. Age, gender and acuity[4] differences
3. Number of high-cost cases versus the benchmark
4. Year-to-year variation in patient demand for services

This list includes both characteristics that influence risk (items 1 and 2) and additional variation (items 3 and 4). Unlike other forms of risk, insurance risk can be difficult to measure because it is driven by unknown and unpredictable events. However, it is important to take insurance risk into account when evaluating contracts and performance. For instance, a provider may have performed well and met all quality initiatives but still have incurred higher-than-expected claims due to unusually high utilization for a number of unforeseen reasons during the year. Taking corrective action without understanding the role of randomness and insurance variation would be suboptimal.

In many risk-sharing arrangements, providers assume financial risk for a relatively small population (e.g., 5,000 or 10,000 lives) in contrast to the much larger populations typically covered by insurance carriers. In Section IV,[5] we highlight the increase in the volatility of claims as the size of the at-risk population decreases. Unlike insurance carriers that cover larger populations, providers may not be as well-equipped or accustomed to handling this type of risk. Therefore, it may be beneficial for providers to seek guidance from actuaries, who are experts on quantifying risk, to fully understand the potential financial impact of entering into a shared risk arrangement.

Insurance risk is the combination of process risk and parameter risk.

PROCESS RISK

Process risk is the risk associated with random chance. Think of a six-sided die. We know that the probability of rolling a five would be one over six, but we would not know whether we actually will roll a five. The same is true for claims. Even if we knew the exact probability that a member would have a claim, we still cannot be certain that the member *will* have a claim. Process risk can be driven by a number of different factors, including but not limited to the proportion of healthy members versus members with chronic conditions,[6] member behavior and the demographic and risk profile of the attributed population. Though healthy members often have much lower average health spending than those with chronic conditions, the former's costs can be subject to more random variation. For instance, a healthy person may incur large health care costs after an unpredictable event such as a car accident or a sports-related injury. Acute events such as these comprise a large portion of health care spending. In a prior study we performed on national data from 2014, we estimated that as much as half of health care spending is attributable to acute events.[7] Even if we knew all the relevant member information for constructing accurate probabilities, we would still be unable to predict what would actually happen.

---

[4] The amount of medical services needed to take care of a patient based on his or her diagnosis.
[5] See Figure 5 to see the potential impact on a provider's probability of savings or loss due to a reduction in population.
[6] This issue may be considered parameter risk depending on how the chronic population is defined.
[7] We ran a proprietary algorithm on our 2011 Consolidated Health Cost Guidelines™ data and bucketed claims amounts into the following categories: full onset chronic (11%), early onset chronic (15%), complex episode (12%), single events (47%) and other (15%).

PARAMETER RISK

Parameter risk is the risk associated with using imperfect information to assess probabilities. We may think we know the mean and variance of the underlying distribution, but it is almost certain that we do not. When modeling health expenditures for a population, claims costs are often estimated using prior claims experience. However, these historical claims are also subject to variation and therefore may not entirely explain the underlying distribution in future periods. There may be changes in legislation, provider reimbursement, medical technology, morbidity levels and so on. Because of this uncertainty, if we want to estimate insurance risk, we should also model parameter risk.

IV. MODELING INSURANCE RISK

Quantifying the insurance risk inherent in a risk contract will help providers and payers make informed decisions. What is the appropriate way to model and quantify that risk? How do the contracting parties know whether the population size and the underlying data used to calculate the risk are credible?

DETERMINISTIC VERSUS STOCHASTIC MODELING

Deterministic models provide a single outcome based on a set of model inputs. We can then perform scenario and sensitivity tests by varying the parameters to understand the range of outcomes. In contrast, *stochastic models* provide a range of possible outcomes for a given set of model inputs.

We will work through several illustrative examples to compare and contrast deterministic and stochastic models. Using a national medical claims data set, we extracted a single plan year of health claims experience on a 10,000-life, commercially insured population. We then projected 2015 average per member per month claims expenditures using both types of projection models.

DETERMINISTIC APPROACH

A deterministic model produces a single predicted claims cost, known as a point estimate. One commonly used method of estimating the cost of claims for a select population at some future date is an experience-based projection. Historical claims experience from a single year (or multiple years) is adjusted for population changes (among other items) and trended to a future projection year. The result is often a single point estimate, or in some instances, a few scenarios may be presented. Figure 1 shows an example of a simplified point estimate using 2014 historical claims trended to 2015 with a 7% trend.

Figure 1:
Illustrative Projection of 2015 Claims Costs (allowed[8])
Using Point Estimate for 10,000 Lives

CY 2014 Cost Per Member Per Month (PMPM)	$436
Trend to 2015	7%
CY 2015 Projected Cost PMPM	$466

Figure 1 shows how the experience-based projection, using a point estimate, provides a projected cost of $466. When looking back to our list of five questions posed at the beginning of this paper, we notice that this approach is only able to answer one of the questions: What is the best estimate of our future performance? This method does not evaluate the provider's total exposure to risk, nor does it help to explain how much of the deviation from the expected target could be due to random fluctuation. Because a range of outcomes is possible, it is helpful to understand the probabilities associated with each outcome in order to assess a contract's viability.

STOCHASTIC APPROACH

Unlike the deterministic approach, a stochastic approach provides a range of outcomes and the expected likelihoods from a single set of inputs (parameters). Actuaries use a variety of stochastic methods. One of the most commonly used, mainly because of its simplicity and ease in understanding, is the *Monte Carlo simulation*. (For a quick primer on how Monte Carlo simulations work, please refer to the first section of the Appendix.)

[8] *Allowed charges* are billed charges after discounts have been applied and include patient out-of-pocket costs.

Within the context of population heath expenditure projections, the Monte Carlo method simulates the claims cost of each member within a population using a randomly generated number. This is then repeated for a specified number of trials to simulate the expected range of results. The Monte Carlo method assumes that each member's claims costs are independent.[9]

Another stochastic model is called *bootstrapping*. (For an explanation of how bootstrapping works, please refer to the second section of the Appendix.) Bootstrapping refers to taking random samples with replacement to model variation in outcomes. This approach will be used when we discuss risk adjustment.

We created a stochastic scenario using the data underlying Figure 1. First, we used the historical claims experience of the population to develop a distribution of per capita expenses across the entire population, referred to as a claims probability distribution (CPD). A claims probability distribution is a list of claimant expenditure levels and the likelihood for each level. This is a critical model input for a Monte Carlo simulation because it attempts to explain some of the variation, or randomness, of an individual member's health expenditures.[10] We then generated a random number for each individual in the group and looked up his or her claims level in the CPD. We did this for all 10,000 members, repeating the process 5,000[11] times. We categorized the average cost from the 5,000 simulations into PMPM buckets with $5 bandwidths to create the histogram shown in Figure 2.

Figure 2: Illustrative Claims Distribution from Monte Carlo Simulation

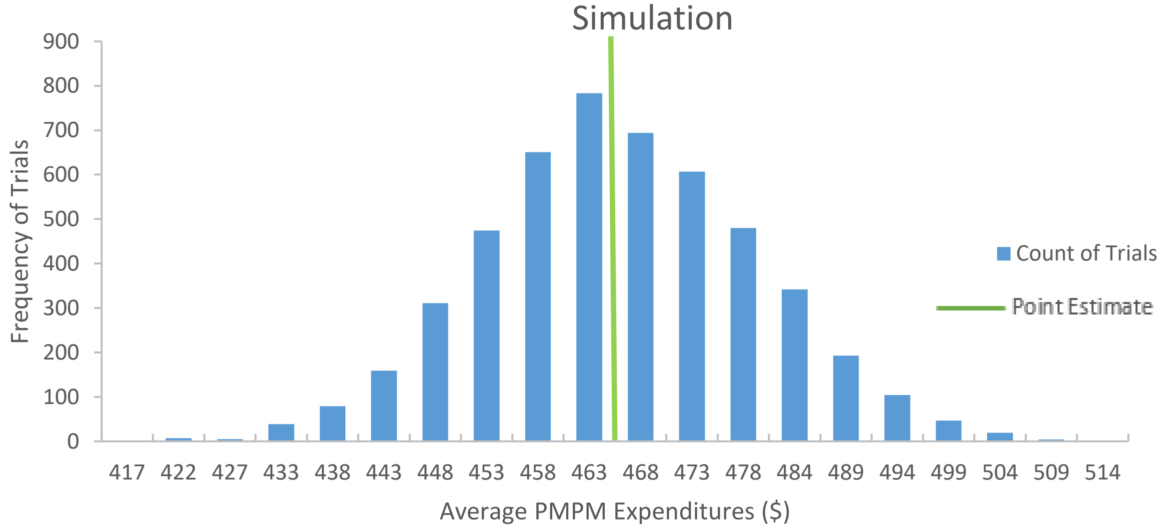

[9] This is a commonly used simplifying assumption. Claims may not be independent.
[10] Notice that in the previous example, the deterministic approach, the model input was a single data point, aggregate health expenditures. Thus, our projection did not take into account any information about the past variation in health expenditures within the population.
[11] Please see the Appendix for information about choosing the number of scenarios.

**Figure 3:
Illustrative Monte Carlo Assumptions and Summary Statistics**

Scenario Assumptions	
Population size	10,000
Stop-loss	None
Care management savings	0%
Attribution	None
Target (102% of CPD mean)	$471
Summary Statistics	
Average PMPM	$464
Coefficient of variation	2.9%
Mean absolute deviation	10.8
Probability of loss (assuming target of 102% of CPD mean)	30%
1st percentile of claims	$432
25th percentile of claims	$455
50th percentile of claims	$463
75th percentile of claims	$473
99th percentile of claims	$496

Figure 2 illustrates that the simulations show a wide range of potential outcomes. The average PMPM generated from the simulation is very close to our point estimate from the deterministic model. However, the simulation gives us a more robust set of information about the distribution of future expenditures. Going back to the original questions at the beginning of this paper, the simulation outputs provide insights into all five questions instead of just a point estimate of future expenditures. Using the information in Figures 2 and 3, we can better understand the expected range of results by looking at the shape of the graph or by calculating statistics such as the standard deviation or mean absolute deviation.[12] We can also estimate the likelihood of savings or loss for a given target. In this example, we assumed a PMPM cost target of 102% of the CPD mean, which resulted in a 30% probability of loss. To calculate the probability of loss, we look at how many of the 5,000 simulations produced an expected PMPM above the target. Because there is no risk corridor, the probability of savings equals 100% minus the probability of loss.

CREDIBILITY

Many times a provider organization may be faced with a situation where it is asked to take on risk for a population size that is only partially credible or not credible at all. Or the provider may have an estimate of how many members will enroll or attribute to a risk contract, but that enrollment or attribution may fall short of the estimate. In addition, a provider may want to subdivide its physicians into practice groups or geographic regions. At what point does the population's data lose credibility or materially increase the provider's insurance risk?

Credibility is a measure of the predictive value of the data.[13] Understanding how much predictive value a population's data has will help a provider understand if the risk contract is in line with its overall risk appetite. Credibility theory

---

[12] *Mean absolute deviation* is the average distance between the mean and each trial.

[13] "ASOP No. 25: Credibility Procedures," revision, 2nd exposure draft, Actuarial Standards Board, June 2013, http://www.actuarialstandardsboard.org/asops/credibility-procedures-3.

can be complicated, and a complete review of credibility procedures is beyond the scope of this paper.[14] One way to review the credibility of a population's data, per Actuarial Standard of Practice (ASOP) 25, is to review the confidence interval[15] and hypothesis test. Actuarial judgment is used to establish the desired level of accuracy, or confidence interval. A stochastic approach will help us develop these statistics.

Using a Monte Carlo simulation, we simulated the claims costs for groups of 5,000, 10,000 and 50,000 members, using 5,000 iterations each. Figure 4 shows a probability distribution graph of possible claims costs. Note that we are using a scatter plot with smooth lines versus a histogram so it is easier to compare the shape of each graph. Instead of graphing the frequency of trials by PMPM band, we have graphed the number of trials in a band divided by the total number of trials to approximate the probability distribution across average PMPMs.

In the example in Figure 4, we made three new assumptions: (1) the contract had an attribution criteria, "prospective 12 months,"[16] resulting in a higher PMPM; (2) the provider would achieve 2% care management savings; and (3) the contract had a $150,000 individual excess stop-loss provision. The impact of modeling these assumptions will be discussed in subsequent sections of this paper.

Figure 4
Illustrative Summary of Per Member Per Month Annual Expenditures

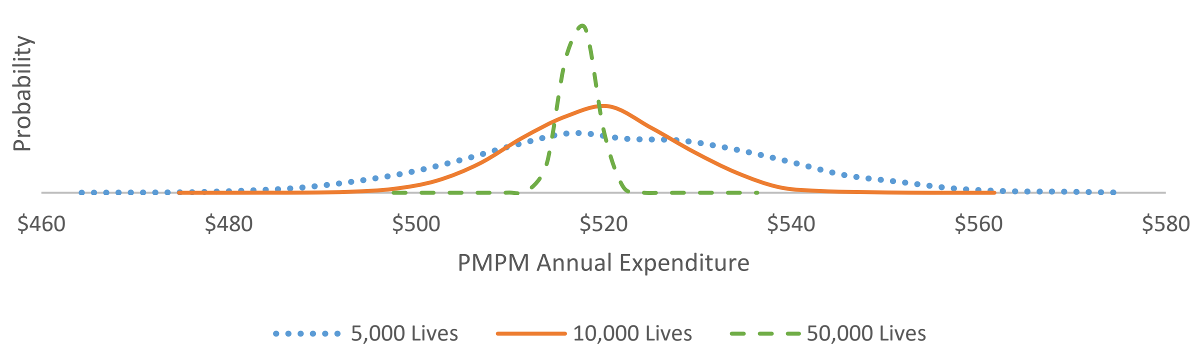

The simulation curves in Figure 4 show higher probabilities around the mean and lower probabilities for costs farther from the mean, resembling a shape similar to a normal distribution. The shape of the curve is determined by the CPD

[14] Karl Volkmar, "Long-Term Care Credibility Monograph Work Group," American Academy of Actuaries, January 2015, https://www.actuary.org/files/imce/LTC_Credibility_Monograph_08172016.pdf.
Robert DiRico, "Credibility Practice Note," American Academy of Actuaries, July 2008, https://www.actuary.org/files/publications/Practice_note_on_applying_credibility_theory_july2008.pdf
[15] *Confidence interval* is the probability of an estimate falling within an acceptable range.
[16] Our prospective 12-month attribution criteria attributed members who had had an office visit with the provider within 12 months prior to the measurement period. Typically, the PMPM claims are expected to be higher for a visit-based attributed population because members with no claims would not be attributed. Many provider risk models rely on similar types of claims-based attribution models.

and population size. The flatter the shape of the curve, the more volatility is expected in the results. As can be seen in Figure 4, the graph of the smallest population is considerably flatter that for the largest population.

Figure 5 shows additional statistics related to Figure 4.

Figure 5:
Illustrative Impact of Varying Population Size

Scenario Assumptions			
Population size (lives)	5,000	10,000	50,000
Stop-loss ($)	150,000	150,000	150,000
Care management savings	2%	2%	2%
Attribution	Prospective 12 months	Prospective 12 months	Prospective 12 months
Target (100% of CPD mean)	$528	$528	$528

Summary Statistics			
Average PMPM	$519	$518	$518
Coefficient of variation	3.0%	2.1%	0.9%
Mean absolute deviation	12.4	8.6	3.9
Probability of savings	71.6%	81.7%	98.0%
Probability of loss	28.4%	18.3%	2.0%
99th percentile of loss (PMPM)	$27	$15	$1
99th percentile of loss (annual total)	$1,645,278	$1,751,251	$492,823
95% confidence interval	$488–$549	$497–$539	$508–$527
90% confidence interval	$493–$545	$500–$536	$510–$526
75% confidence interval	$501–$537	$505–$531	$512–$523

As can be seen from Figures 4 and 5, the average PMPM is similar across all population sizes. However, the volatility is very different. The 50,000-life group has significantly less variation and probability of loss than the 5,000- and 10,000-life groups. Stochastic modeling shows expected volatility across different population sizes. It also sheds light on the credibility of experience data.[17] Therefore, providers can make more informed decisions on what minimum population size they would require. In the preceding example, if the provider were looking to identify the size of the population required to keep its probability of loss below 20%, it knows that it would need a population size of at least 10,000 lives.

INCLUDING PARAMETER RISK WHEN MODELING

The last two stochastic examples used a fixed input for all scenarios in the form of the claims probability distribution. Thus, our simulations assumed that we knew the underlying probability distribution of the population with 100% certainty, which would not be possible in reality. In this example, the parameter risk is the risk that our CPD misrepresents the true underlying distribution, as was discussed in Section III. To better quantify the overall insurance risk, we might consider applying a random variable to the mean and a random variable to the adjustment factor (e.g., trend). Other parameters could be modified as well, but for this example, we will focus on the two parameters already mentioned as they account for the greatest amount of uncertainty in our projection model. Starting with the same 10,000-life example from Figure 5, and increasing or decreasing the CPD by a multiplicative scalar randomly selected over 5,000 trials, will produce the results shown in Figures 6 and 7.

Figure 6
Illustrative Claim Distribution by Percentage from Target

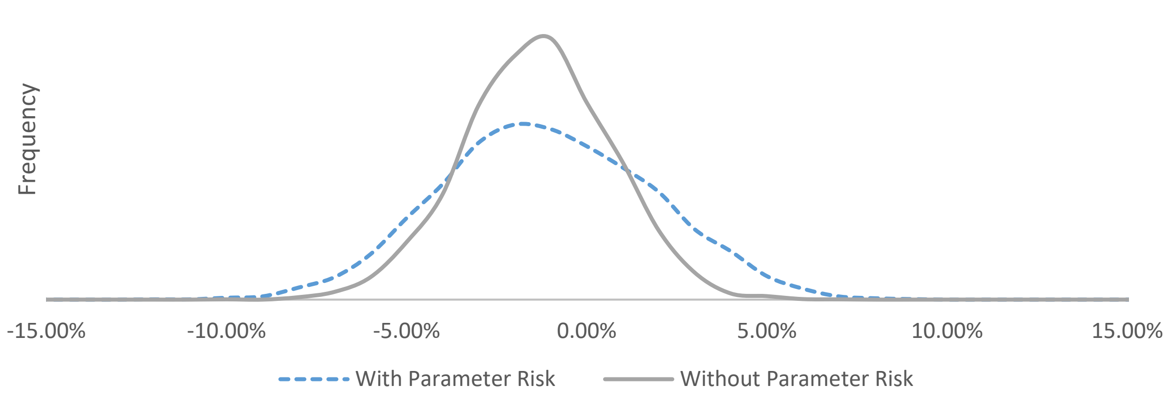

[17] If an actuary deems the population as not credible, he or she can use a limited fluctuation approach or a greatest accuracy credibility to create credibility factors and blend experience data with a larger industry subset.

Figure 7: Illustrative Impact of Modeling Parameter Risk		
Scenarios	**With Parameter Risk***	**Without Parameter Risk***
Population size	10,000	10,000
Stop-loss ($)	150,000	150,000
Care management savings	2%	2%
Attribution	Prospective 12 months	Prospective 12 months
Target (100% of CPD mean)	$528	$528
Summary Statistics		
Average PMPM	$519	$518
Coefficient of variation	3.0%	2.1%
Mean absolute deviation	12.6	8.6
Probability of savings	70.2%	81.7%
Probability of loss	29.8%	18.3%
99th percentile of loss (PMPM)	$27	$15
99th percentile of loss (annual total)	$3,230,017	$1,751,251

* Parameter risk includes a CPD scalar with uniform distribution of 3% around the mean and trend with triangle distribution, with Min/Mode/Max of 4%/7%/11%.

Figure 7 shows how introducing parameter risk increases the estimate of insurance risk. The amount of parameter risk added will depend on the credibility of the data used to create the CPD. Generally speaking, parameter risk is greater for parameters that are set based on data from smaller populations. We could model parameter risk for any of the adjustment factors we use to develop parameters where there is uncertainty (e.g., trend assumption). In some cases, historical experience for the specific attributed population does not exist. In these cases, using benchmark data may be the best option, but parameter risk should also be considered.

V. QUANTIFYING THE IMPACT: EXAMPLES

Quantifying the impact of parameters in risk contract design is important in establishing the appropriate terms for the contract. In addition, shared risk arrangements are usually predicated on providers modifying their performance to lower the cost of care and eliminate waste. It is important to sensitivity test how results will change if those provider targets are not met or are exceeded.

The following are some questions we frequently encounter from providers:

- How does the change in target affect my overall likelihood of success?
- How much can the risk corridor protect me from incurring a loss or reduce the likelihood of savings?
- What overall impact do the high claimant exclusions have on the variation in results?
- Will risk adjustment help protect me from insurance risk?
- Will excluding certain service categories help reduce volatility?

The next few sections demonstrate how stochastic modeling can be an excellent tool to help answer these questions.

TARGET

The performance target (sometimes called the financial benchmark) plays a large part in the probability and magnitude of a savings or loss for a provider in a shared risk payment model. Thus, it is a critical element of the contract and often a key point of contention in provider-payer contract discussions.

Suppose a provider is trying to determine the minimum target level that results in an acceptable amount of risk. The provider would likely want to know what the probability of loss is at each target. To answer this question, we used a Monte Carlo simulation along with the "with parameter risk" scenario from Section IV to model the probability of loss at various targets. Figure 8 displays a graph of the results. Targets are shown as a percentage of the mean of the underlying CPD. The provider plans to implement care management initiatives that we assume will result in 2% lower claims costs, thus centering the graph at 98% (for a full list of assumptions, see the "With Parameter Risk" scenario in Figure 6).

Figure 8: Illustrative Probability of Loss by Target

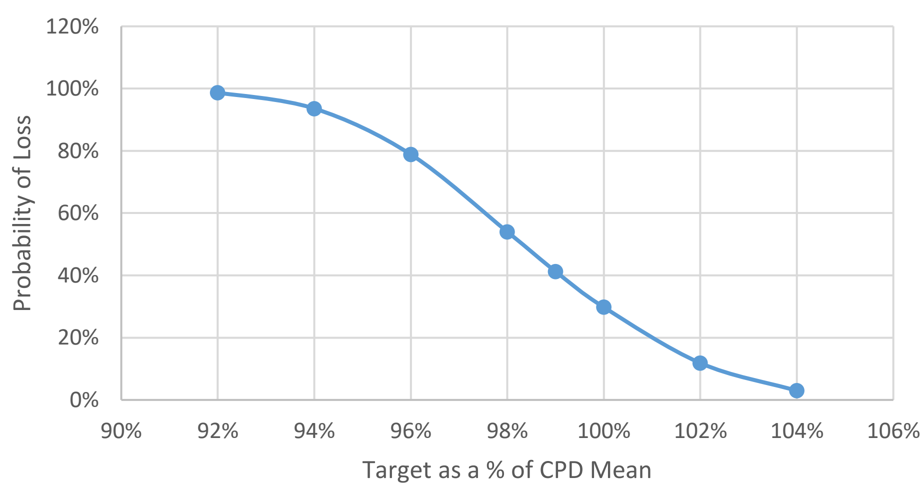

As expected, the probability of loss decreases as the target increases. What is perhaps less intuitive is the fact that the relationship between target and probability is nonlinear. The change in probability is greatest around 98% of the CPD mean where the slope of the curve is the steepest, and the curve flattens as it gets farther from 98%. The change in loss from 97% to 96% is much higher than the change from 102% to 101%. This essentially means that the provider is exposed to more risk if the target moves from 97% to 96% than if it moves from 102% to 101%. It is important to note that a 1% change in target has different implications depending on where the target is set.

To properly manage risk, the provider is also interested in the expected loss, assuming a loss has occurred. In other words, how much loss would the provider expect if it did not meet its target? For this type of analysis, conditional tail expectation is a useful statistical tool that quantifies the expected value of a distribution of risk events above a predetermined threshold (e.g., in the example that follows, we chose the 99th percentile). We also considered the expected savings if a provider experiences claims below the target and at the 1st percentile of claims. These statistics are helpful because they give providers an idea of what loss or savings they might expect due to an unfavorable or a favorable year in claims.

Figure 9:
Illustrative Probability and Expected Value of Savings/Loss across Various Target Levels
10,000 Lives
$150,000 Stop-Loss Level
2% Care Management Savings
Attribution = Prospective 12 Months

Target as a % of CPD Mean	Target Shown as a PMPM	Probability of Savings	Probability of Loss	Expected Loss Given Claims Exceed Target	Loss Given 99th Percentile Claims ($555 PMPM)	Expected Savings Given Claims Below Target	Savings Given 1st Percentile Claims ($484 PMPM)
92.0%	$486	1.3%	98.7%	$4,121,466	$8,925,556	$630,933	$789,352
94.0%	496	6.4%	93.6%	3,036,921	7,658,602	793,403	2,056,306
96.0%	507	21.1%	78.9%	2,205,851	6,391,648	1,015,646	3,323,260
98.0%	517	46.0%	54.0%	1,648,156	5,124,694	1,375,515	4,590,215
99.0%	523	58.7%	41.3%	1,422,439	4,491,217	1,638,582	5,223,692
100.0%	528	70.2%	29.8%	1,217,764	3,857,740	1,953,478	5,857,169
102.0%	538	88.2%	11.8%	892,529	2,590,785	2,701,579	7,124,123
104.0%	549	97.0%	3.0%	665,725	1,323,831	3,671,368	8,391,077

These results should also be considered in the context of risk corridors and stop-loss, which are discussed further in the following sections. As shown in Figures 8 and 9, results from stochastic modeling can provide useful insight to help providers choose a target.

RISK CORRIDOR

A risk corridor defines the minimum threshold that the savings or loss must exceed for a payment to be made. For example, if the contract claims cost target is $400 PMPM with a risk corridor of 2%, there will be no savings or loss for the provider if the average PMPM is between $392 and $408. Risk corridors can help protect both the payer and the provider from making payments due to the random fluctuation of claims.

As soon as a corridor is introduced, it often leads to questions such as, “How much does this impact the likelihood of savings or losses?” The payer and provider may have different perspectives about an appropriate corridor in such cases. Figure 10 shows the probability of a savings, loss or neither at different thresholds. This type of information can help facilitate a constructive dialogue between a payer and provider about an appropriate corridor.

Figure 10: Illustrative Probability of Shared Savings Payment across Scenarios and Minimum Savings/Loss Corridors 10,000 Lives $150,000 Stop-Loss 2% Care Management Savings Attribution = Prospective 12 Months Target (100% of mean)			
Minimum Savings and Loss Corridor	Probability of Savings	Probability of No Payment	Probability of Loss
0.0%	70.2%	0.0%	29.8%
1.0%	58.7%	21.4%	19.9%
2.0%	46.0%	42.2%	11.8%
3.0%	32.9%	60.5%	6.6%

As can be seen from Figure 10, the risk corridor can have a large impact on the probability of loss. However, even when set at a fairly high percentage (i.e., 3%), it cannot completely remove insurance risk. The problem is that at the same time it helps to minimize risk, it may also reduce the provider's opportunity to generate savings. For example, a provider that makes small incremental savings over time may not receive any share of those savings if they always fall within the risk corridor.

ATTRIBUTION LOGIC

A common goal across shared risk payment models is increased accountability of care. In many of these models, a claims-based methodology is used to determine which physician or organization is accountable for a patient or episode of care. This is often referred to as *attribution*, or *assignment*. So the "attributed" population, for which a provider may assume financial risk, is assigned based on where patients receive care, which may not be known ahead of time and will vary from year to year. This methodology reflects a stark contrast from the membership in a typical health plan. A health plan typically knows exactly who will be covered going into the year based on annual enrollment (i.e., the member's plan selection).

We previously mentioned that having an estimate of the CPD is a crucial parameter in running a stochastic model for a population's aggregate health expenditures. Using the populations in previous examples, we created CPDs for a standard covered population and a population that would be "attribution-eligible" under a commonly used primary care physician (PCP) claims-based attribution model. *Attribution-eligible* means that a member incurred a type of claim that would result in assignment to a physician (e.g., certain evaluation and management claims). We used a methodology that attributed members both prospectively (using claims from the prior year) and concurrently (using claims from the performance year) and compared these populations' CPDs with the CPDs of populations where no attribution logic was applied. Figure 11 summarizes the results. The average PMPM of an attributed population is 19% to 20% higher than that of a standard population. This is expected because attributed members must, by definition, have claims during the attribution period (and any members with no utilization are therefore excluded).

Figure 11: Illustrative Impact of Varying Attribution			
Scenarios			
Population size	10,000	10,000	10,000
Stop-loss ($)	150,000	150,000	150,000
Care management savings	2%	2%	2%
Attribution	**Prospective 12 months**	**Concurrent 12 months**	**None**
Target (100% of CPD mean)	$528	$531	$443
Summary Statistics			
Average PMPM	$519	$522	$437
Coefficient of variation	3.0%	3.0%	3.2%
Mean absolute deviation	12.6	12.5	11.2
Probability of savings	70.2%	69.4%	67.0%
Probability of loss	29.8%	30.2%	33.0%
99th percentile of loss (PMPM)	$27	$29	$26
99th percentile of loss (annual total)	$3,230,017	$3,448,277	$3,093,786

We then ran Monte Carlo simulations across all three approaches summarized in Figure 11. Our hypothesis was that an attribution-eligible population would have less variation in average claims because many of the members who did not have a PCP visit (including $0 claimants) were removed. This was, indeed, the case, although the reduction in volatility was less pronounced than we expected. Additional research may be warranted to understand the underlying variation between populations aligning with physicians as compared with more general populations under managed care. Actuaries would benefit from knowing if more readily available general claims probability distributions can be relied on for aligned populations' insurance risk analysis.

RISK ADJUSTMENT

One method to help insulate providers from insurance risk is to risk-adjust their claims targets, which will reduce parameter risk. Risk-adjusted targets are common across shared risk payment models, including both total cost of care models and episode-based payment models. While the use of risk adjustment in shared risk payment models is broadly accepted by both providers and payers, just how effective is it at mitigating insurance risk? We can answer this question using the stochastic modeling technique.

Starting with a total cost of care payment model, we ran a simulation similar to those in previous examples to measure the impact of risk adjustment on insurance risk. In this example, we simulated both the annual claims cost and the risk score for each individual in the population using the Milliman Advanced Risk Adjusters™ (MARA™) model on a concurrent basis. Rather than the Monte Carlo technique, we used the bootstrapping technique,[18] which involves taking a random sample with replacement. We believe that bootstrapping is a more suitable process than running independent Monte Carlo simulations across the two random variables. With each trial, we randomly sampled both the per capita claims cost and the risk score for each member. Each sample contained 5,000 members,[19] and we performed 5,000 iterations. We then summarized savings and losses under two scenarios: one where we adjusted

[18] Please see the Appendix for a description of the bootstrapping technique.

[19] Due to the different statistical technique (bootstrapping) and sample population, 5,000 members was chosen for modeling convenience.

claims cost for risk score and one where we did not. Figure 12 shows the results of our analysis as the probability of hitting a certain savings (–) or loss (+) percentage with and without risk adjustment.[20]

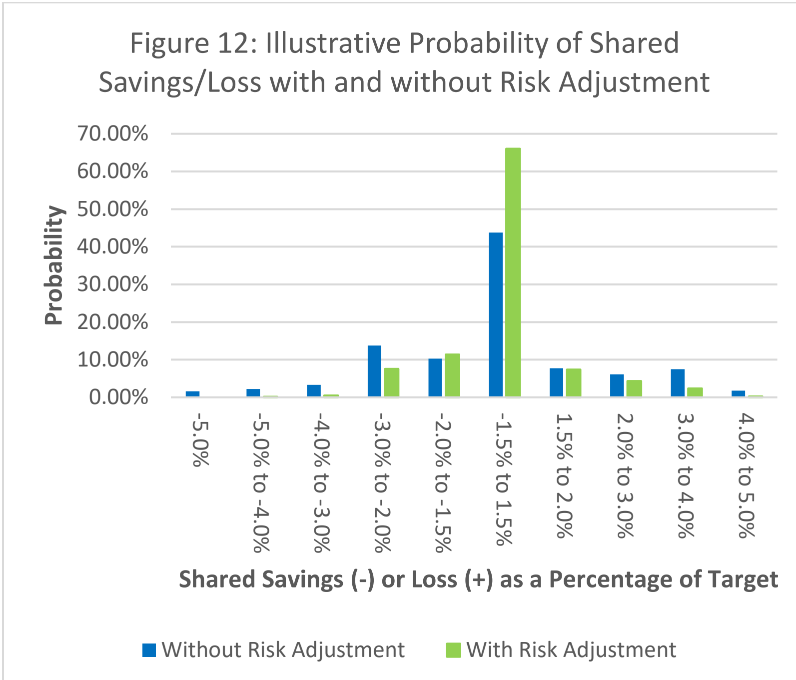

Figure 12 shows that risk adjustment materially lowers the volatility of outcomes. A higher percentage of the scenarios fall in the –1.5% to 1.5% range of savings or loss when risk-adjusted as compared with scenarios without risk adjustment (66% vs. 44%).

Figure 13 summarizes some key results and shows additional statistics. As shown here, this particular risk adjuster has reduced the overall variation of the population average expenditures compared to the cost target. Therefore, the risk adjuster has reduced the insurance risk transferred through the shared risk contract.

---

[20] Colleen Norris and Stoddard Davenport, "Risk Adjustment Techniques for Improving Value-Based Payments," Milliman White Paper, March 2016, http://www.milliman.com/uploadedFiles/insight/2016/2199HDP_20160324.pdf.

Figure 13:
Illustrative Impact of Adding Risk Adjustment

Scenarios	Without Risk Adjustment	With Risk Adjustment
Population size	5,000	5,000
Stop-loss	None	None
Care management savings	2%	2%
Attribution	None	None
Target (100% of mean)	$359	$359
Summary Statistics		
Average PMPM	$352	$352
Coefficient of variation	4.8%	3.0%
Mean absolute deviation	13.6	8.3
Probability of savings	67.8%	76.5%
Probability of loss	32.2%	23.5%
99th percentile of loss (PMPM)	$34	$18
99th percentile of loss (annual total)	$2,024,181	$1,075,829

A provider may choose to risk-adjust its performance targets if it anticipates that the performance year population will be different from the experience year population. For instance, if a provider anticipates that the performance year population will be sicker overall than the experience year population, it would increase the target proportionally to the expected cost difference. In essence, it reduces the risk of selecting an inappropriate target cost (i.e., the parameter risk).

Though risk adjustment can be a useful tool to reduce insurance risk, care must be given in selecting an appropriate model.[21] A number of public and private shared risk payment models are being used. In modeling analyses, it may not be feasible to simulate both per capita costs and risk adjustment factors. Therefore, actuaries will need to use judgment regarding the degree to which risk adjustment may reduce the volatility. Other considerations for selecting risk adjusters is outside the scope of this paper but has been discussed in detail in another SOA study.[22]

STOP-LOSS

Shared risk payment models usually contain a stop-loss (or large claim) provision based on the notion that random catastrophic claims are outside the direct control of the provider. The incorporation of a stop-loss provision will reduce the overall volatility of the population's aggregate claims expenditures. A difficult question that often arises centers on the appropriate level at which to set the stop-loss. Sometimes the stop-loss provision is provided by the payer with whom the provider is negotiating the alternative payment contract. Other times the provider is expected to purchase a stop-loss product from an excess of loss reinsurer. In the latter case, the provider would also need to factor in the price of purchasing the excess of loss contract. While providers generally appreciate relief from the risk of catastrophic

[21] Hans K. Leida and Leigh M. Wachenheim, "Risk Adjustment and Shared Savings Agreements," Milliman Healthcare Reform Briefing Paper, January 2015, http://www.milliman.com/uploadedFiles/insight/2015/shared-savings-agreements.pdf.
[22] Geof Hileman and Spenser Steele, "Accuracy of Claims-Based Risk Scoring Models," Society of Actuaries, October 28, 2016, https://www.soa.org/research-reports/2016/2016-accuracy-claims-based-risk-scoring-models.

claimants, if the stop-loss threshold is too low, the provider may miss some of the best opportunities to manage a population's health expenditures.[23]

Adding a stop-loss provision to the contract can help reduce the risk of high-cost outliers. There are two forms of stop-loss. *Specific stop-loss* removes the claims amount in excess of the threshold for an individual or removes that individual's claims entirely. *Aggregate stop-loss* removes the claims amount in excess of the threshold if the total claims amount for the entire population reaches the threshold. By removing these high-cost amounts, stop-loss reduces the volatility of the population. This presents another situation where measuring insurance risk using a stochastic model is helpful. The curves in Figure 14 show the volatility in claims costs for a population assuming various specific stop-loss levels. A Monte Carlo simulation was used to model these scenarios. In each scenario, claims in excess of the stop-loss level were removed. For example, in the $150,000 scenario, a member with $400,000 of claims would only have the first $150,000 of claims included in the modeling analysis. In addition, we assumed that the stop-loss was internal (i.e., not purchased from an external reinsurer).

Figure 14
Illustrative Summary of Per Member Per Month Annual Expenditures

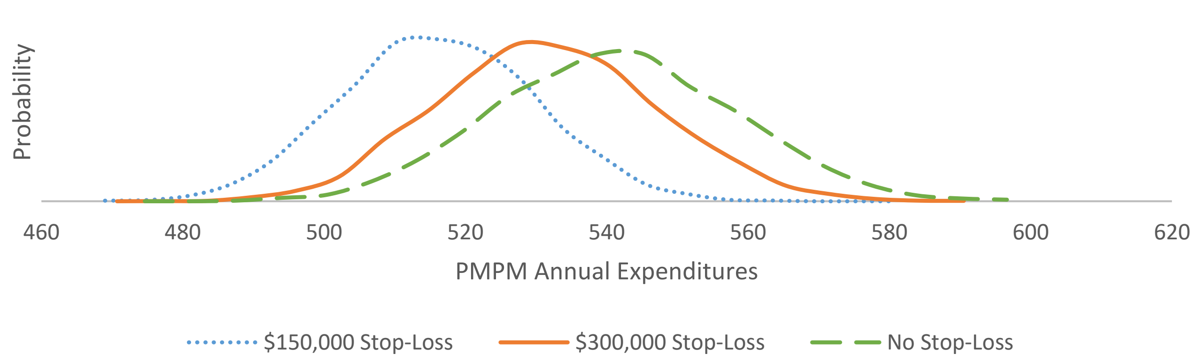

Figure 15 summarizes some key results and shows additional statistics.

[23] Population health management's impact at various levels of excess claims is an area that may benefit from further research.

Figure 15:
Illustrative Impact of Varying Stop-Loss Assumptions

Scenarios			
Population size	10,000	10,000	10,000
Stop-loss ($)	**150,000**	**300,000**	**No stop-loss**
Care management savings	2%	2%	2%
Attribution	Prospective 12 months	Prospective 12 months	Prospective 12 months
Target (100% of CPD mean after stop-loss)	10,000	10,000	10,000

Summary Statistics			
Average PMPM	$519	$530	$535
Coefficient of variation	3.00%	3.10%	3.30%
Mean absolute deviation	12.6	13.3	14.3
Probability of savings	70.20%	69.40%	68.40%
Probability of loss	29.80%	30.60%	31.60%
99th percentile of loss (PMPM)	$27	$30	$33
99th percentile of loss (annual total)	$3,230,017	$3,558,117	$3,904,217

As expected, the $150,000 stop-loss has the lowest volatility; however, the coefficient of variation and the probability of savings do not vary significantly at the different deductibles. If results are similar for a provider's specific modeling, the provider may consider choosing a higher stop-loss threshold, especially if it needs to purchase external stop-loss. However, stop-loss would still help manage the other types of risk. For example, technical risk associated with errors in a provider alignment algorithm or a change in risk adjustment methodology.

CARE MANAGEMENT

The ultimate goal of risk sharing is to incentivize providers to manage care more efficiently and reduce spending. Therefore, providers are expected to implement care management programs that help reduce utilization so claims fall below the target.[24] We modeled insurance risk under various care management scenarios. Figure 16 summarizes the results of these simulations.

[24] Providers and payers may want to consider quality measures as well to ensure care management reductions are meaningful.

Figure 16: Illustrative Probability of Shared Savings Payment across Care Management Scenarios 10,000 Lives, $150,000 Stop-Loss Level, Prospective 12 Months, 100% Target, No Corridors	
Care Management Reductions	Probability of Savings
3%	81.6%
2%	70.2%
1%	57.2%
0%	43.8%

We can see from Figure 16 that care management has a significant impact on the probability of savings. However, it is also true that even if the provider establishes programs to save on spending, it may still incur a loss that is strictly due to random fluctuation. Figure 16 shows that even with a 3% care management savings, there is an 18% chance of losses. It is also true that even if a provider does absolutely nothing to save on costs, it may still generate savings because of random fluctuation. It is tempting to attribute savings to good performance and a loss to random fluctuation. However, it is very difficult to determine the true cause. Having an understanding of the underlying random variation in claims expenditures will help a provider better assess results and past performance.

INCLUDED SERVICES

Providers and payers specify what services are included in the contract. Services that are less predictable and have a high cost, such as organ transplants, nonpreventable emergency room visits and services over which the provider may have less control (e.g., prescription drugs), are sometimes excluded. The services selected will have an impact on the underlying volatility of the population. The services that are carved out would remain the payer's risk or could be moved to a stop-loss provider (in some cases, because these services may be less predictable, it may make sense for them to be pooled with other carved out services from other risk contracts).

For our sample population, including or excluding prescription drug claims did not have a significant impact on the underlying theoretical volatility. However, prescription drugs could have a significant impact on volatility for other populations we have reviewed. Results will be different depending on the mix of services and/or the population. Therefore, it is important to perform an analysis on the specific at-risk population before drawing any conclusions.

VI. OTHER CONSIDERATIONS

The preceding sections presented our set of practical examples of using stochastic modeling in some commonly found shared risk payment model design and contracting situations. While we have attempted to touch on a few of the most common occurrences, in practice there are numerous contextual nuances that will require adapting some of the methods outlined. As we are limited in the scope and number of examples we could include in this paper, we would like to outline some additional considerations.

TECHNICAL RISK

Many of the shared risk payment models being tested are complex. Recent experience already shows that a number of unintended consequences and unforeseen unknowns arose because of the complex nature of these new payment models. In some instances, technical risk and administrative challenges could approach, or even exceed, insurance risk. This additional risk and uncertainty should be considered in addition to the volatility (or insurance risk) in most situations.

OVERRELIANCE ON THE MODEL

As with all actuarial modeling, the actuary needs to resist the urge to become overly reliant on the stochastic model being used. Stochastic modeling may give a false sense of precision. While we wholeheartedly endorse the use of the techniques outlined here, we encourage actuaries to always consider the potential biases and limitations of the techniques presented.

RISK ACROSS POPULATIONS

Many providers are considering “taking on risk” across multiple populations they serve. For example, providers may have risk arrangements with multiple health plans and government payers or with the same health plan across multiple product lines. It is worth considering how the combination of these at-risk populations can reduce the overall insurance risk, just as the insurance industry has been doing since its inception. Consistency across contracts may also have the benefit of reducing technical risk. We did not provide any examples of this type of analysis, but similar techniques to those described in this paper can be used.

RISK ACROSS REVENUE STREAMS

All of the examples we have presented model the savings or losses of the payment model without taking into account the financial impact to the provider on its existing contracts and lines of business. In practice, many providers and health systems will want to measure the potential range of financial outcomes on their entire books of business.

VII. CONCLUSION

Success in provider shared risk payment models ultimately boils down to good risk management, efficient delivery of care and entering into the right shared risk contract. This means the organization must understand its exposure, volatility, probability, severity, time horizon and correlation of risk. For providers sharing increased risk, recognizing that insurance risk exists and understanding the drivers (process and parameter risk) is an important first step.

An actuary can help quantify these risks using deterministic and stochastic modeling approaches, calculating and considering credibility and understanding the risk in setting assumptions (via parameter risk). A stochastic approach allows the actuary to review key statistics such as confidence intervals, standard deviation, probability of savings and loss, expected loss and so on.

By reviewing these statistics, providers can make informed decisions on different contractual provisions such as the target, risk corridor, attribution logic, risk adjustment, stop-loss and included services. As both the payer and the provider become more comfortable with these statistics, both parties can work together to create mutually beneficial contracts that meet the risk goals of both organizations.

VIII. LIMITATIONS AND RELIANCE

This analysis was prepared on behalf of the Society of Actuaries to provide information on insurance risk in provider shared risk payment models.

This report is based on information and data from various sources, which Milliman has not audited. In preparation for writing this paper, we reviewed various published reports on provider payment models. We have not reviewed every rule, antitrust regulation or payment model regulation. A legal review of these programs may provide other insights into the potential for success of each program and/or cost approach.

Case studies presented are from an illustrative group and contract. Other actual contracts and populations would certainly produce results different from those presented in this report.

Milliman does not intend to legally benefit any third-party recipient of its work product. Even though Milliman has consented to the release of its work product to a third party, any third-party recipient of this report should not rely upon Milliman's report but should engage qualified professionals for advice appropriate to its own specific needs. The statements contained in the report are those of the authors and do not necessarily represent the views of Milliman or its other consultants.

Cory Gusland and Juliet Spector are members of the American Academy of Actuaries and meet the Qualifications Standards of the American Academy of Actuaries to render the actuarial analysis contained herein.

APPENDIX

1. MONTE CARLO SIMULATION

In general, a Monte Carlo simulation refers to any simulation using random numbers in conjunction with a probability distribution to solve a numerical problem. In the context of modeling claims costs for a population, a Monte Carlo simulation uses random numbers to generate claims costs for each member in a population. This process is repeated a number of times to produce a distribution of the population's aggregate claims costs. To perform a Monte Carlo simulation, a claims probability distribution (CPD) table is developed and the size of the population and the number of desired trials is identified. A CPD table contains possible claims costs and the probability of a member incurring that claims cost. The following steps detail the Monte Carlo process:

1. For each member, a random number from 0 to 1 is generated. This random number corresponds to a claims cost in the CPD table. The higher the probability of a certain claims cost, the more likely a member will be assigned that cost.

2. The aggregate cost for the entire population is calculated by adding all the claims costs for each member from step 1.

3. Steps 1 and 2 are repeated for a certain number of trials. For our examples, we used 5,000 trials. The results of the trials can be graphed to see the distribution. The more trials that are used, the smoother and more normal the graph will be.

A Monte Carlo simulation is a useful and robust method to model the randomness of claims costs. This method assumes that each member's claims costs are independent.

2. BOOTSTRAPPING

Bootstrapping is another technique used to model the distribution of claims costs. Instead of generating random numbers like the Monte Carlo simulation, it relies on selecting a random sample from a sufficiently large data set. To use this technique, you will need a large data set of members and their claims costs, the population size and the number of trials. Here is a detailed description of how bootstrapping works:

1. For each trial, a random sample of *x* members, where *x* is the size of the population, is selected from a large pool of data.

2. The claims are aggregated for the population.

3. Steps 1 and 2 are repeated for a certain number of trials. For our risk adjustment example, we used a trial size of 5,000. After each trial, the sample is replaced back into the original data source. Therefore, it is possible for the same member to be selected in multiple trials. The results of the trials can be graphed to see the distribution. The more trials that are used, the smoother and more symmetrical the graph will be.

Bootstrapping can be particularly useful when you do not have a CPD or it would be difficult to construct one. However, bootstrapping requires that you have a sufficiently large source of data that is representative of the actual variability in members' claims amounts. Otherwise, it would not be possible to select a large number of trials without reusing the same members repeatedly. In addition, the sample size must be large enough to be representative of the entire data set. If these criteria are not met, then the results may not be reliable.

About The Society of Actuaries

The Society of Actuaries (SOA), formed in 1949, is one of the largest actuarial professional organizations in the world dedicated to serving more than 28,000 actuarial members and the public in the United States, Canada and worldwide. In line with the SOA Vision Statement, actuaries act as business leaders who develop and use mathematical models to measure and manage risk in support of financial security for individuals, organizations and the public.

The SOA supports actuaries and advances knowledge through research and education. As part of its work, the SOA seeks to inform public policy development and public understanding through research. The SOA aspires to be a trusted source of objective, data-driven research and analysis with an actuarial perspective for its members, industry, policymakers and the public. This distinct perspective comes from the SOA as an association of actuaries, who have a rigorous formal education and direct experience as practitioners as they perform applied research. The SOA also welcomes the opportunity to partner with other organizations in our work where appropriate.

The SOA has a history of working with public policymakers and regulators in developing historical experience studies and projection techniques as well as individual reports on health care, retirement and other topics. The SOA's research is intended to aid the work of policymakers and regulators and follow certain core principles:

Objectivity: The SOA's research informs and provides analysis that can be relied upon by other individuals or organizations involved in public policy discussions. The SOA does not take advocacy positions or lobby specific policy proposals.

Quality: The SOA aspires to the highest ethical and quality standards in all of its research and analysis. Our research process is overseen by experienced actuaries and nonactuaries from a range of industry sectors and organizations. A rigorous peer-review process ensures the quality and integrity of our work.

Relevance: The SOA provides timely research on public policy issues. Our research advances actuarial knowledge while providing critical insights on key policy issues, and thereby provides value to stakeholders and decision makers.

Quantification: The SOA leverages the diverse skill sets of actuaries to provide research and findings that are driven by the best available data and methods. Actuaries use detailed modeling to analyze financial risk and provide distinct insight and quantification. Further, actuarial standards require transparency and the disclosure of the assumptions and analytic approach underlying the work.

Society of Actuaries
475 N. Martingale Road, Suite 600
Schaumburg, Illinois 60173
www.SOA.org