IBM Watson Health™

WHITE PAPER

The IBM Cost of Care Model

Using claims data to help predict healthcare resource consumption

Contents

The healthcare industry is in the midst of unprecedented reform

Not only are organizations dealing with increased pressure to provide quality healthcare, reduce costs and improve outcomes for patient populations, they are also faced with changes in the reimbursement system. These changes involve moving from a volume-based system to one that is value-based. This shift can have a significant impact on employers, health plans and providers, as reimbursement will be tied to the appropriate utilization of healthcare services in addition to the achievement of high-quality outcomes within their populations.

Another aspect of reform is the moving of financial risk from payers to providers. Providers are being expected to take on more risk for their population's utilization of healthcare resources. We are starting to see penalties involved when patients are hospitalized for avoidable admissions or readmissions, and in the future, we believe there will also be bonuses for ensuring patients receive the appropriate care.

An additional consideration is the number of patients within a population who have multiple conditions and/or are chronically ill. This subpopulation can have a significant impact on overall costs and therefore may require special attention. If action is not taken, these individuals could incur far more expenses over a long-term period than the cost of an appropriate and timely intervention.

We believe predictive models are no longer a "nice to have," but rather a cornerstone of operational necessity and critical to managing risk.

All of these important changes can compel those involved to look at how they do business and make decisions through a new lens. The ability to mitigate risk and subsequently manage costs becomes paramount, and understanding a patient population can be the first step in better managing its risk.

Although there are several avenues in which organizations can maneuver to make the most of this developing landscape, we think the use of predictive analytics is becoming key—by offering insights to identify and engage in "actionable" events.

From using predictive models to identify high-risk patients who may benefit from care management programs, to risk-adjusting provider and payer populations to guide the setting of provider reimbursement and premium rates, we believe predictive models are no longer a "nice to have," but rather a cornerstone of operational necessity and critical to managing risk.

The IBM Cost of Care methodology consists of a suite of risk-adjustment models that can be used to help predict two types of expenses: concurrent (retrospective) and prospective.

The Cost of Care Model from IBM is an example of how predictive analytics can inform improved decision-making in innovative ways.

Overview of IBM Cost of Care Model

The IBM® Cost of Care Model uses administrative commercial claims to help predict healthcare resource consumption for a patient population.

The model was built from the IBM® MarketScan® Commercial Database. This database—derived from a pool of data spanning more than 218 million covered lives since 1995—includes information on enrollment, outpatient services, inpatient services and hospital admissions and prescription drug usage.

There are two key components to the risk-adjusted cost of care model:

1. Classification of patient clinical conditions into disease categories
2. Determining risk scores, which help predict relative healthcare expenditures

Classification of clinical conditions into disease categories

The model applies the IBM® Disease Staging® classification system to describe patients' clinical conditions and severity of illness. IBM Disease Staging is a categorization schema, which describes the type and severity of specific diseases and conditions. It considers the clinical characteristics of a disease as described by clinical findings and diagnosis codes, without consideration of the treatment or resources consumed during the care process. The underlying assumption is that the clinical progression of most diseases can be described by transition to increasingly greater complications and poorer prognoses.

Using objective clinical criteria, Disease Staging is designed to define discrete points in the course of individual diseases that are clinically detectable, reflect severity and possess clinical significance for prognosis and treatment.

The staging criteria identify a maximum of five stages based on the natural progression of the specific disease in the absence of any intervention or treatment.

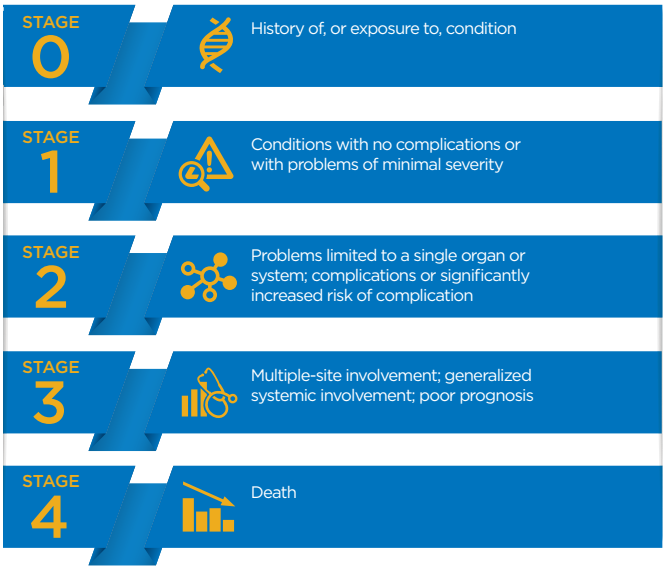

Determining a population's financial risk

The IBM Cost of Care Model consists of a suite of risk-adjustment models that estimate a population's retrospective and prospective expenses. The concurrent models use administrative claims from a historical time period to assign relative risk for the same time period. The prospective models use administrative claims from the prior 12-month period (year one) to assign risk for the subsequent period (year two).

Both the retrospective and prospective versions provide separate models for total cost and medical cost. Total cost is defined as medical cost and outpatient pharmacy cost. Medical cost is defined as only medical cost, excluding pharmacy cost. Cost is calculated as the total payment received by the provider, including payments made by both the health plan and the patient. These two categories are further stratified into three additional subcategories:

- No claim truncation
- $100,000 claim truncation
- $250,000 claim truncation

Some variations of the models were developed using data with an individual's annual cost truncated at specified levels. This method is referred to as "Windsorizing," as opposed to the exclusion of outliers from the analysis. Truncating the cost data helps to reduce the level of variation by eliminating the impact of extreme, high-cost outliers. For example, if truncation occurs at $100,000, then all patients who have dollar values greater than $100,000 will be set equal to $100,000.

The IBM Cost of Care Model is designed to estimate financial risk for both full-year and partial-year population members. All members receive a risk score, even if they do not have any diagnoses for the time period. In the case of the former, the risk score is based only on age and gender and is most relevant to the prospective model. In total, there are 12 risk models: six concurrent and six prospective, as shown below.

Risk models

Concurrent	Prospective
Total cost – No truncation	Total cost – No truncation
Total cost – $100,000 truncation	Total cost – $100,000 truncation
Total cost – $250,000 truncation	Total cost – $250,000 truncation
Medical cost – No truncation	Medical cost – No truncation
Medical cost – $100,000 truncation	Medical cost – $100,000 truncation
Medical cost – $250,000 truncation	Medicalcost – $250,000 truncation

Cost of Care Model business applications

Our Cost of Care Model estimates an individual's relative risk score for each of two types of cost—total and medical cost. These risk scores can be used in a variety of ways to help you evaluate and manage the cost of providing care to a population of individuals.

In addition to a score, the models provide "reason codes" to help explain the underlying drivers of risk. Each of these risk factors describes a medical condition and quantifies its contribution to the risk score.

We believe the ability to understand and quantify the current health of a patient population can help answer the following business questions:

How can we measure the relative cost-effectiveness of the care provided to the enrolled population?
Some populations may be "sicker" than others and will justifiably incur more healthcare expenses. Using concurrent risk scores allows organizations to take this "illness burden" into account when evaluating the relative cost incurred by a particular group. This application of risk scores, often referred to as "risk-adjusting" population expenses, is an important component in evaluating where potential unnecessary care is being delivered or where unusually high payments are being made.

The use of risk adjustments can also help in evaluating and comparing providers. Adjusting the cost of care of a provider's population for its illness burden, we think, allows for a more accurate review of healthcare utilization among providers. This process can better help identify potentially unnecessary care and atypical practice patterns among providers.

How can we improve outcomes in our patient population?
With increased pressure to provide quality healthcare and improve outcomes for patient populations, healthcare organizations, health plans and employers all need a clear understanding of the potential risk of their populations.

Risk scores allow for the stratification of a population into clusters of similar risk and to trend costs and utilization of those categories over time. This helps deliver the ability to evaluate the cost of care for specific subpopulations of high or increasing risk. In our opinion, this is the first step in evaluating the cause of a high or rising risk score: Is it due to aging, acute but manageable conditions and/or chronic conditions?

We believe this knowledge will better inform and shape intervention programs and help organizations to proactively manage high-risk, high-cost cases.

How can we reduce/minimize costs across our population?
Reducing costs and increasing efficiency is a leading priority for most healthcare organizations today, and managing patient populations can be more important now than ever before.
This can be looked at on two different fronts—the first being to analyze recent costs, and the second being to better understand and prepare for the future cost risk of the individuals in your population.

The prospective risk scores can be used to group individuals into cost categories to identify patients at risk of becoming high-cost patients next year. This information, along with the provided reason codes, can either be analyzed within a particular disease to help identify patients for disease management, or analyzed across various disease conditions to identify the need for new care management programs.

Actively monitoring a patient population and understanding why and where costs are rising, can help identify opportunities for improvement.

Are there members within our population whom we should be directing to disease management programs?
As organizations work to deliver higher-quality healthcare and improve outcomes for their patient populations with chronic illness, one avenue of pursuit is the use of disease management programs.

Being able to identify members of their population who would benefit most from participation in such a program is often the first obstacle, with the second being when to engage those members. Access to a patient's prospective risk score allows the identification of patients who are not currently at high risk, but appear to be heading in that direction. This, together with the reason codes provided by the model, can help an organization to understand what particular condition is driving that risk.

The ability to evaluate a patient's underlying cause of risk, as well as the progression of a disease, can help inform if and when a patient could benefit from participation in a disease management program.

How can we identify and manage high-volume healthcare users within our population?
As the industry moves from a volume-based payment system to one that is value-based, being able to identify and manage high-volume care users within a population becomes more important.

Patients who are high-volume users tend to be older and have multiple chronic conditions. Identifying these users is necessary to begin to manage and control their costs—allowing for intervention, whether participation in a disease management program or another type of outreach.

Once those members have been identified, there are tactics that can be employed to closely monitor these patients, which include trending their risk scores, as well as monitoring the associated reason codes. The reason codes show how each of their chronic conditions contributes to their overall risk and the potential impact of unexpected or related acute conditions.

We believe this information provides insights to help assist organizations in controlling costs and managing this subpopulation.

How can we help ensure reimbursement and premium rates are appropriate for our population?
It is important in the changing payment environment to be able to plan for future financial risk. Being able to accurately anticipate future risk, we believe, is the first step in setting premium rates based on the health risk and healthcare utilization among populations.

The prospective risk scores can be used to help organizations determine the best way to allocate payments across groups based on the expected future resource utilization of those groups. This insight can help to ensure that prospective premium rates are set appropriately.

When establishing reimbursement rates, it is important to take into account differences in the health risk of patient populations among providers in order to develop an appropriate and equitable reimbursement system.

To equitably set provider capitation payments, it is important to adjust for the illness burden and related financial risk for the assumed population. The prospective risk scores can be used to risk-adjust costs for a given population.

Model techniques, exclusions and performance

Our Cost of Care Model applies a variety of nonlinear modeling techniques to estimate individual risk scores and reason codes. The relative risk score is a measure of each member's cost risk relative to the average of the population.

For each member, all possible diagnosis codes from inpatient and outpatient claims are incorporated into the model to capture all comorbidities. The risk score represents each member's risk and is based on several variables, including:

- Age
- Gender
- Diagnosis groupings
- Severity levels

Each input to the model is assigned a "model weight" that was estimated in the development and calibration of the model. The model weights are applied to the raw values of the inputs for each member for the 12-month time period, which results in a calculated relative risk score for each member. Members with no medical claims or diagnosis codes during the 12-month period will receive a risk score based only on age and gender.

All models use the same disease classification system, but model weights will differ among the 12 models. Depending on the model, the risk score either represents a risk for total cost or medical cost for either the concurrent or prospective models. The relative risk score can be used as an indicator of the overall illness burden of the member, as well as the related financial risk.

The models also provide "reason codes" that help to explain the results by identifying the drivers of the risk. Each reason code includes the diagnosis category and severity level and quantifies its contribution to the risk score.

In our opinion, the most appropriate use of the results is to calculate mean (average) relative risks for a population of members. It is at this level, rather than the level of the individual person, that the statistical performance of the models is highest.

Methodology exclusions

All diagnostic information reported on claims from the following service type and facility categories were excluded:

- Radiology
- Durable Medical Devices (DME)
- Laboratory/Pathology
- Supplies/Devices
- Transportation

Often, diagnoses on these types of claims are either "rule-out" diagnoses or are not recorded by clinical professionals, and are therefore not considered reliable.

Model performance

Model performance is commonly evaluated using the statistical measure R2. R2 indicates how much of the variation in cost between members can be explained by the model. An R2 closer to 1.0 (100 percent) reflects better performance and higher predictive power.

		R2 for concurrent models		R2 for prospective models	
	Truncation level in dollars ($)	Full-year plan enrollees*	All enrollees	Full-year plan enrollees	All enrollees
Medical cost	None	0.515	0.460	0.164	0.159
	100,000	0.615	0.591	0.212	0.214
	250,000	0.612	0.587	0.200	0.209
Total cost	None	0.524	0.495	0.204	0.182
	100,000	0.609	0.593	0.279	0.249
	250,000	0.610	0.595	0.251	0.233

* Full-year enrollees have been enrolled all 12 months during the baseline period.

The concurrent models have R2 values of 0.515 to 0.615 for full-year enrollees and 0.460 to 0.595 for all enrollees (full- and partial-year enrollees). Prospective models have R2 values in the range of 0.164 to 0.279 for full-year enrollees and 0.159 to 0.249 for all enrollees.

It is common for concurrent models to exhibit higher performance than prospective models, since a concurrent model uses data from a specified time period to estimate risk for the same time period, while prospective models anticipate the healthcare resources for the next year based on current-year diagnosis information. The retrospective models have access to information on unexpected acute events. These acute events have high predictive power for the retrospective time period, but lose much power in predicting future costs.

Model output description

Risk score

The risk score reflects the expected cost for an individual relative to that of the population. The higher the risk score, the greater the financial risk. The scale of the risk score is normalized to 1.0 for the average of the large dataset from which the model was calibrated. Most users will choose to "renormalize" the scores, so that their own population average is 1.0. Under this scenario, a risk score greater than 1.0 suggests a higher risk (and cost) relative to the average of the population, and a risk score less than 1.0 suggests a lower risk (and cost) relative to the average of the population.

The expected cost for an individual can be calculated using the individual's risk score in combination with the population's average risk score and the population's average cost. It is calculated as the population's average cost multiplied by the ratio of the individual's risk score to the average of the population's risk score.

Model output

Following is an example of the output from the concurrent model with no truncation:

ID		Risk Score	Reason codes	Age	Gender
XXXXX	4	4.156	GUS10-2.04-81 GIS04-2.02-18	29	Female

In this example, the 29-year-old female member has been enrolled for four months and has a risk score of 4.156.A risk score of 4.156 indicates that she has a much higher risk (and expected cost) relative to the average member of the population experience used to estimate the model. In fact, her risk is over 4 times the average risk of that population.

The reason codes provided indicate that this member has two disease categories that contribute to her high-risk score:

- GUS10: Urinary Tract Infection – Stage 2.04: With Perinephric Abscess (moderate severity) – contributes 81 percent to her risk score
- GIS04: Anorectal Suppuration – Stage 2.02: With Anorectal Abscess (moderate severity)— contributes 18 percent to her risk score

This is an example in which an acute condition (GUS10: Urinary Tract Infection Stage 2.04—With Perinephric Abscess) is the primary driver of current risk. This condition is unlikely to affect future risk; therefore this patient will likely have a lower prospective score.

Summary

This paper has described the business and analytic needs for cost of care models, outlined some of the most typical applications, explained the modeling techniques used to create the models and provided a sample of the model output.

We believe cost of care models stand to provide today's organizations with important predictive insights necessary to manage the healthcare costs, outcomes and potential risks of both individuals and virtually any commercially insured population.

Get connected

email: watsonh@us.ibm.com

About IBM Watson Health

Each day, professionals throughout the health ecosystem make powerful progress toward a healthier future. At IBM Watson Health, we help them remove obstacles, optimize efforts and reveal new insights to support the people they serve. Working across the landscape, from payers and providers to governments and life sciences, we bring together deep health expertise; proven innovation; and the power of artificial intelligence to enable our customers to uncover, connect and act — as they work to solve health challenges for people everywhere.

For more information on IBM Watson Health, visit: ibm.com/watsonhealth

IBM Corporation
Route 100
Somers, NY 10589

Produced in the United States of America
April 2018

HPW03065-USEN-01

IBM Watson Health™