Editor's Corner

By Rick Pawelski

I recently had a conversation with a former colleague who had submitted an article for publication in *Health Watch*. He was also interested in working on additional ideas for it, but he was concerned that perhaps one had to have the right job title or the proper connections with the Society of Actuaries (SOA) to get an article published in the digital pages of this newsletter. I was surprised to hear that. In the preceding month I had been splashing noisily in the deep end of my new volunteer position as editor, trying to figure out how to keep a steady supply of quality content in the pipeline for future issues of *Health Watch*. The conversion to more frequent digital output was a move to near-constant management of that pipeline. Let's just say I have been more concerned with the supply side than the demand side.

I would like to tell all of our readers just what I told my old teammate: *Health Watch* encourages and welcomes contributed articles of quality from all contributors. It doesn't matter who you work for or who you know, if you've got an interesting and educational story to tell, then share it with us. If you need help developing an idea into a document, we can provide it. Our mission is to inform and educate the members of the Health Section, and we will energetically leverage any and all input from all corners to do so.

We members of the Health Section are also encouraged—by each other—to contribute articles to *Health Watch*, so to practice what I preach, I recently reached out to a coworker and the two of us put together an article. This was a new experience for both of us, but we wrote about what we knew, trying to make it interesting and informative. Through this process I benefitted intellectually from the effort of rooting through all I've learned and done over the last few years in an attempt to narrow it down to only the most useful, important and universal facts. You will see that article in a future issue of *Health Watch*; I hope you find it worth your time.

Over the last few months I have become not only an editor but also an author. These were not things I planned for in advance, which makes me think about skiing. Specifically, skiing through the trees. "Skiing the trees" is another thing I once thought I would never do; there's an enhanced element of risk that comes with sliding between and around trees that once seemed unbearable to me. Over time, I got a bit better at the edging, rotation, and pressure movements required to make your skis do what you want them to do, and I eventually got comfortable with the idea of heading off the trail and into the glade. I discovered that skiing the trees is an exercise in decision making—you have to decide immediately which gap between trunks you're headed for next, and that outcome will present you with another decision, another outcome requiring yet another decision, and so on. I found it a useful metaphor for professional life, where today's decision is based on the outcome of so many decisions we've already made, and we can only see so far into the glade, where tomorrow's decisions await.

The SOA volunteer experience has been that way for me. When I first followed the leadership of others I knew and became an SOA volunteer, I didn't really know what to expect. I only knew I was willing and hopefully able to point my metaphoric skis into that glade and start making turns. Along the way I've learned about actuarial science, task management, communication and other things, and now that I've begun working on *Health Watch*, I'm excited to see what lies a little farther down the slope. We've got some articles in the pipeline for future issues, and if you're interested in digital health or pharmacy pricing, you'll want to stay tuned. Beyond that, it depends on you. Like all volunteer-driven organizations, the Health Section depends on the energy, industry and creativity of its members to benefit the general good. If you're reading this, you've almost certainly got expertise and experience that will be interesting and informative for others to read about. I encourage anyone with any interest to help define our future together by contributing to *Health Watch*. ■

Rick Pawelski, FSA, MAAA, is a manager of actuarial services focused on vendor financial management for Highmark Health. He can be reached at *rick.pawelski@highmarkhealth.org*.

Copyright © 2020 Society of Actuaries. All rights reserved.

How Blockchain Technology Will Disrupt the PBM-Payer-Pharmacy Relationship

By Brian N. Anderson, Gregory Callahan and Michael DiPrima

Editor's note: This article was originally published at https://www.milliman.com/Insight/How-blockchain-technology-will-disrupt-the-PBM-payer-pharmacy-relationship, *copyright © 2019 by Milliman Inc. Reprinted by permission.*

Pharmaceutical prices and the purpose of pharmacy benefit managers (PBMs) received no shortage of attention in recent months, especially in light of both congressional and White House pressure to tackle the high cost of prescription drugs. The complex business model—one that, by nearly all accounts, is lacking in transparency—came about out of necessity from the increased and overwhelming complexity of pharmacy claims processing.[1] However, where there is complexity and unequal access to information, there is often room for increased efficiencies and opportunities for disruption.

The pharmaceutical industry has many stakeholders. The primary purpose of the PBM is to serve as an intermediary among pharmacies, drug manufacturers, and payers. PBMs manage formularies, negotiate discounts and rebates, and process and pay prescription drug claims through a web of electronic transactions.

Blockchain could potentially transform the relationship between payers, pharmaceutical manufacturers, wholesalers, and pharmacies by offering an alternative, transparent mechanism for processing, pricing, and validating prescription transactions. This approach could lead to less waste, reduced pricing variations between pharmacies, and a better app-based purchasing experience for consumers through transparency of the true cost of prescription drugs.

This article describes: (1) the origin of PBMs; (2) how claims are processed; and (3) ways in which blockchain could disrupt the PBM marketplace. The article does not address every aspect of blockchain but rather is intended to start a conversation about the potential for disruption of this emerging technology.

Bear in mind, the pharmacy benefit space has many niches. The content of this paper may apply to state and federal health plans, but the focus of this paper is on the private health care space.

THE ORIGIN OF PBMs

PBMs as we know them today originated during the early 1980s when the health insurance and pharmaceutical market grew rapidly in size and complexity because of technological advancements. As the demand for prescription drugs increased, pharmacists and insurance companies found themselves overwhelmed manually processing claims. The need for an intermediary to adjudicate prescription claims arose, and PBMs developed to fill that gap. PBMs served as financial intermediaries, negotiating lower prices with the drug manufacturers. They also promised to make processing of prescription claims easier and more efficient.[2]

The growth of PBMs during the 1980s and 1990s created a web of contracting complexities which eventually shielded patients from the actual cost of prescriptions. Utilization and prices subsequently grew at an outstanding rate. Patients struggled to make cost-based purchasing decisions, and there was no mechanism to exert market pressure on manufacturers and PBMs.

Then the Medicare Modernization Act of 2003 provided a significant boost to the PBM industry. As Medicare Part D expanded to cover prescription costs, PBMs found themselves playing a significant role. To reduce the administrative burden on insurance companies, PBMs assumed the role of identifying patients eligible for Medicare Part D as well as setting drug prices with manufacturers through discounts and rebates.

In this role, PBMs set the formulary. If a drug was not on the formulary, doctors would be less likely to prescribe it and insurers would not cover the cost. This gave the PBMs tremendous power to negotiate prices with drug manufacturers and shield consumers for the actual cost of medications. Today, there are less than 30 mid-to-large PBMs, with three major PBM companies (CVS, OptumRx, and Express Scripts) handling between 70% and 75% of the prescription claims.[3] All three of the largest PBMs own or are controlled by a major health plan.

HOW PBMS CURRENTLY PROCESS AND STORE PRESCRIPTION DRUG CLAIMS

PBMs currently use a traditional database management system (DBMS) model to adjudicate and store prescription claim data. A DBMS allows users to create and manage one or more databases. A traditional DBMS is based on a centralized client/server model, where clients communicate with a centralized server where data is stored. A simplified example of this is the way many businesses have computers for each employee but store the company's data in a centralized server facility. Users must be authenticated before they can view and edit the data, and users with enough authority can create, read, update, and delete data in a traditional client/server DBMS model. If the security on a database is compromised, a malicious user can modify, corrupt, or even delete information on a database. Additionally, the direct availability of the data to mutual third parties is generally discouraged due to security concerns. DBMS approaches are shown in Figures 1 and 2.

Figure 1
Centralized DBMS approach

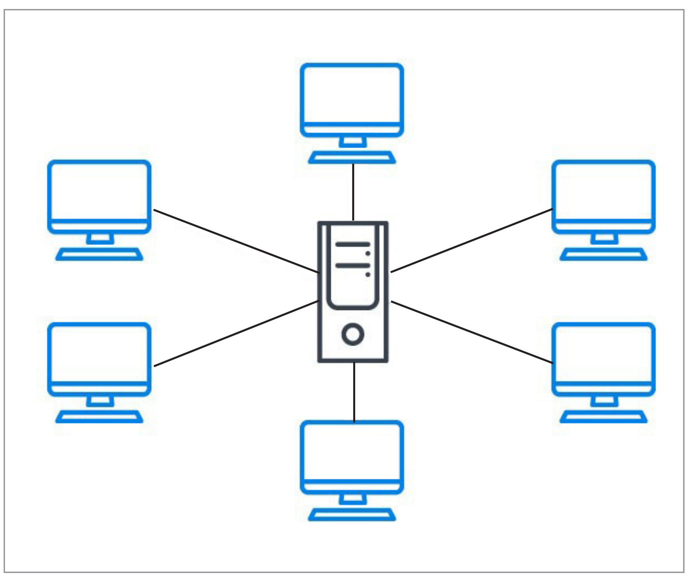

Figure 2
Cloud-based DBMS

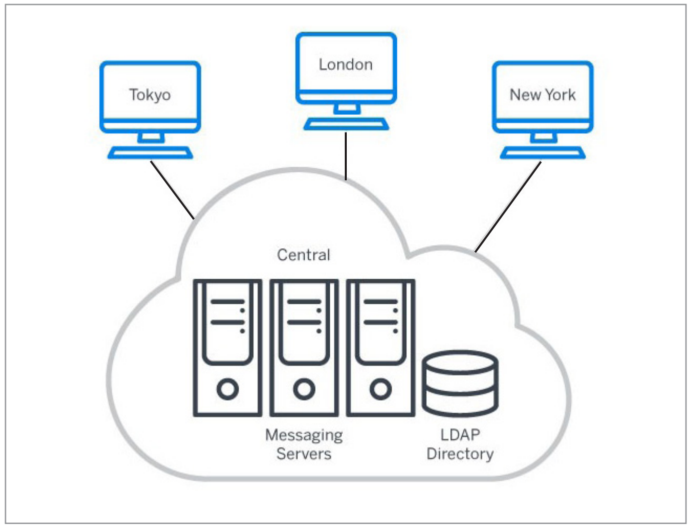

BLOCKCHAIN DBMS MODEL

Blockchain is an emerging technology that is often designed to offer a distributed ledger system to accurately validate and confirm transactions. In common configurations, each participant has access to a copy of the ledger databases. Various consensus algorithms can be utilized to make it very difficult to change the ledger history. Data that cannot be changed is called immutable, one of the most desirable characteristics in blockchain. One way to explain the concept is the "glass safe" analogy, originally coined by Fabricio Santos of Cointelegraph: imagine a bank vault filled with glass safe deposit boxes in which participants can view the contents of all the boxes but cannot access them. When a person opens a new deposit box, that person receives a unique key to that box but does not own the box—only the contents inside it. Because all box owners can see the contents of each box, the group can validate them, but individuals can access only their own.[4]

It is a common misconception that everyone can see everything on a blockchain. For a fully public blockchain, like cryptocurrencies, this is partially true. Everyone can see the first blockchain transaction ever posted, the most recent, and everything in between. The transaction itself is visible, but the real-world identity of the individuals and entities behind the transactions remains hidden. As blockchain emerges in private enterprises, it will not be open to everyone. Private blockchain designs might be a practical solution for some business problems. A drug manufacturer, for example, may participate in a private blockchain with its wholesaler in an effort to prevent fraud, waste, and abuse. There are even semi-private blockchain solutions. They can be implemented by a large agency, for example (like state and federal governments), allowing other large organizations (like auto manufacturers and insurance groups) to participate in an effort to improve vehicle safety.

In contrast to a traditional DBMS model, blockchain technology uses a decentralized model. Blockchain employs a series of nodes to validate transactions. To complete a transaction, a series of nodes must validate it before it is stored and digitally signed and encapsulated into a block of data. Once encapsulated, consensus algorithms can help protect a block from being altered or destroyed. As new transactions are generated, validated, and encapsulated, they are strung together to create a chain.

Properties of blockchain that make it useful are shown in the table in Figure 3.

Figure 3
Uses of Blockchain

Property	Use Case	Real-World Example
Immutability	All participants agree on the same history, reducing disputes about historical events.	The country of Georgia implemented a public blockchain for land title transactions.[5]
Transparency	Transparency in a marketplace leads to better decision making by all parties.	Walmart implemented blockchain in its food supply chain to track the origin of products from multiple suppliers.[6]
Smart contracts	Transparency is not just access to the data. Smart contracts represent transparent access to business logic.	Cryptocurrencies have reduced transfer fees imposed by banks, lowering the cost of money transfers.[7]

DBMS AND PBMs: CAN BLOCKCHAIN DISRUPT THE PBM MARKET?

A traditional DBMS uses a client/server model. In contrast, a blockchain DBMS uses a decentralized model. A distributed network of nodes validates all transactions. Blockchain data is stored throughout the distributed network rather than in a centralized server. Simply put: all participants of the system (that are running a node) validate a transaction before it is committed, ensuring transparency for everyone.

The current DBMS-based ecosystem is complex, inefficient, and sometimes opaque. Relevant pharmacy claims data is stored in one centralized location (managed by the PBM), and both retail pharmacy and payer transactions must operate through that centralized database in order to process pharmacy claims. Claims that are submitted by a pharmacist for a prescription must be adjudicated and reimbursed by the PBM, which then bills the insurance company and manages the rebate process with the manufacturer. The process can take upward of 60 days.

In a blockchain model, no centralized storage database is needed. Instead, claims transactions, like the contents of the glass safe deposit boxes, can be viewed, validated, and distributed by all stakeholders simultaneously. Such a system could fuel an entirely different consumer experience. A consumer could walk into a retail pharmacy and, using a phone or card, pay for a prescription and have the relevant claims data, utilization, cost, and rebates all calculated concurrently at the point of sale using blockchain technology.

Specific benefits of blockchain applications include:

- **Drug pricing.** Under the current system, PBMs control the drug formulary in most situations. The formulary determines which drugs the insurance company will pay for and which drugs providers can prescribe for coverage. This provides an incentive for drug manufacturers to negotiate higher rebates with PBMs in order to have their drugs included on the formulary. However, these negotiations are not transparent, and there is a lag between when prices are negotiated and when health plans and payers receive the pricing information. This approach shields the consumers for price differences, which often vary significantly for both drugs and pharmacies.

  A blockchain application could enable the pricing information to be encoded into a smart contract that can then be executed at the point of sale and quickly validated by all parties. This could evolve to simplifying the pharmacy claims process to be a credit card transaction that enables the consumer to purchase products at the least expensive in-network pharmacy. The pharmacy pricing information would be available online in a manner that all other consumer products are purchased by.

- **Efficiency.** Blockchain systems use a feature called smart contracts: coding that provides conditions that must be met before a transaction can be validated and signed. Utilizing smart contracts, drug manufacturers and insurance companies could set the criteria for drug prices and coverage. Insurance companies could also instantly update information as patient coverage changed or was rescinded. Currently, when a patient goes to a doctor, a series of transactions and phone calls is initiated to update and verify insurance. When the patient presents a prescription at the pharmacy, the insurance coverage is verified again. If any of the insurance information is entered incorrectly, it may result in the PBM rejecting the claim, leading to resubmission of the claim. Many PBMs charge a fee every time they process a claim or resubmission. A blockchain system with smart contracts could apply the validations closer to the point of service, letting users iterate to an approved claim faster.

- **Durability.** Unlike a traditional centralized DBMS, a blockchain DBMS is decentralized, eliminating any single point of failure. Furthermore, a blockchain DBMS only permits transactions to be read or created, guaranteeing an audit trail for all transactions.

THE FUTURE OF PBMs AND BLOCKCHAIN

Although blockchain technology appears to hold great promise for the PBM marketplace, there are several obstacles to overcome before that promise becomes a reality, and the PBM industry is primed for an evolution.

While blockchain has demonstrated great potential for security in cryptocurrency, its true potential in health care has not yet been evaluated. Further testing and implementation would be required before a blockchain application could be used on a large scale in the PBM marketplace.

All members of the ecosystem—payers, providers, pharmacies, and PBMs—would need to agree on the validation and approval criteria. As we have seen in other complex capitalized industries, finding consensus among stakeholders can pose a challenge.

CONCLUSION

Industry stakeholders can agree the future should hold a more efficient, less complex, more transparent, and easily accessible integration of all health care information and technology data. Today, the electronic approach to processing pharmacy claims is still the same one used and designed back in the early 1980s.

The current ecosystem among pharmacies, payers, providers, and PBMs is neither cost-effective nor transparent to shareholders across the pharmaceutical supply chain spectrum. It is inefficient and has resulted in patients and payers absorbing the additional cost of the complex infrastructure that has resulted in higher prices in prescription drugs.

A blockchain system could increase transparency and trust among all parties involved and increase savings to the patient. Payers and pharmacies would benefit from blockchain by reducing the time they spend validating insurance coverage, making phone calls, and managing data. PBMs and other interested parties should work to evaluate a blockchain alternative now with an eye toward a more efficient drug financing system. The more efficient system would be able to handle additional technology improvements that the next generation is prepared to adopt. This includes real-time integration of health care devices, user applications, medical providers, consumerism-shopping tools, and integration to non-health care purchases. ■

Brian N. Anderson, MBA, is a principal at Milliman. He can be reached at *Brian.Anderson@milliman.com*.

Gregory Callahan is a pharmacy benefits consultant at Milliman. He can be reached at *Greg.Callahan@milliman.com*.

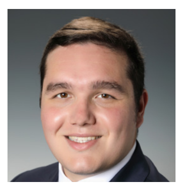

Michael DiPrima is a data analyst at Milliman. He can be reached at *Michael.DiPrima@milliman.com*.

ENDNOTES

1 Lipton, H. L., Kreling, D. H., Collins, T. & Hertz, K. C. (1999). Pharmacy Benefit Management Companies: Dimensions of Performance. Retrieved June 27, 2019, from *https://www.annualreviews.org/doi/full/10.1146/annurev.publhealth.20.1.361#_i4*.

2 Kaupp, G. S. & Lynch, S. B. (2002). To Regain Control of Drug Costs, Health Plans Must Understand the Road They've Traveled. Retrieved April 11, 2019, from *https://www.medscape.com/viewarticle/432389_1*.

3 Yu, Nancy L., Atteberry, Preston, & Bach, Peter B. (July 31, 2018). Spending on Prescription Drugs in the US: Where Does All the Money Go? *Health Affairs*. Retrieved June 18, 2019, from *https://www.healthaffairs.org/do/10.1377/hblog20180726.670593/full/*.

4 Santos, Fabricio (2016). Explaining Blockchain to Laypeople: The Analogy of the Glass Box. *Medium*. Retrieved June 18, 2019, from *https://medium.com/@lunes.platform/explaining-blockchain-11695afee9c0*.

5 Shang, Q. & Price, A. (2018). A Blockchain-Based Land Titling Project in the Republic of Georgia. *Innovations*. Retrieved June 18, 2019, from *https://www.mitpressjournals.org/doi/pdf/10.1162/inov_a_00276*.

6 Unlock. Walmart Case Study on Food Supply Chain Transparency Using Blockchain. Retrieved June 18, 2019, from *https://www.unlock-bc.com/news/2019-02-22/walmart-cse-study-on-food-supply-chain-transparency-using-blockchain*.

7 Geiregat, S. (October 2018). Cryptocurrencies Are (Smart) Contracts. *Computer Law and Security Review*. Retrieved June 18, 2019, from *https://www.sciencedirect.com/science/article/pii/S0267364918302279*.

Quantile Regression—A New Actuarial Approach to Claims Estimation

By Jason Reed

Suppose, as a health insurance underwriter, you are confronted with the following two credible groups' experience for the past two years, with all numbers trended to a future rating period as shown in Table 1.

Table 1
Trended Two-year Claim Experience for Group Rating

	Year 1 Experience	Year 2 Experience	Average
Larry's Lumber	$370	$340	$355
Bob's Trucking	$360	$380	$370

How would you set expected claims for each group in that future rating period?

Traditional underwriting might indicate pricing at the group's average experience or a blend of that with a block-level per member per month (PMPM). But suppose you knew the entire distribution of claim costs for each group, which would vary based on the demographic and morbidity profile of each member in the group. Then, in addition to observing actual experience, you would know what percentile of the claims distribution that experience represented, as Table 2 illustrates.

Table 2
Percentile of Claims Distributions per Group

	Year 1 Experience	Percentile of Claims Distribution	Year 2 Experience	Percentile of Claims Distribution	Average
Larry's Lumber	$370	85th	$340	70th	$355
Bob's Trucking	$360	25th	$380	35th	$370

Based on this information, you might reason that Larry's Lumber would continue to regress to its mean with lower than average trend in Year 3, but that Bob's Trucking would likely experience higher than average trend in the future period, as claims regress upward toward the mean of that group's particular claim distribution.

EVOLUTION FROM POINT ESTIMATES TO CONFIDENCE INTERVALS

Much more powerful inferences can be made from historical data when probabilities can be assigned to the range of possible future claims costs. Quantile regression, a statistical technique that involves estimating the percentiles of claims based on demographics and a group's unique morbidity profile, enables estimation with this level of precision and insight.

Like a traditional risk score, quantile regression estimates a linear function of a member's age, gender and morbidity profile. But expanding on the information in a risk score (which estimates only the mean given these predictors), it develops a distinct function at each percentile, so differing impacts of demographics and medical conditions at different parts of the distribution can be observed. For example, females may cost three times as much as males at the 50th percentile of the distribution (due to better compliance with preventive care and maternity costs) but may not have a significant difference in cost at the 98th percentile (since care for many catastrophic diseases may not vary significantly by gender).

Quantile regression adds a new dimension to the underwriting process. Traditionally one would compare normalized experience across similar groups within a block. Quantile regression takes a different perspective, allowing each group's experience to be evaluated relative to the entire distribution of possible claims outcomes *for that group*. So if a group had experience that was at the 20th percentile of expected claims given its morbidity profile, we might anticipate a regression to the mean and, therefore, higher trend than expected in a future period. This allows for

Copyright © 2020 Society of Actuaries. All rights reserved.

Figure 1
10th, 50th and 90th Quantile Regressions for the Age Predictor

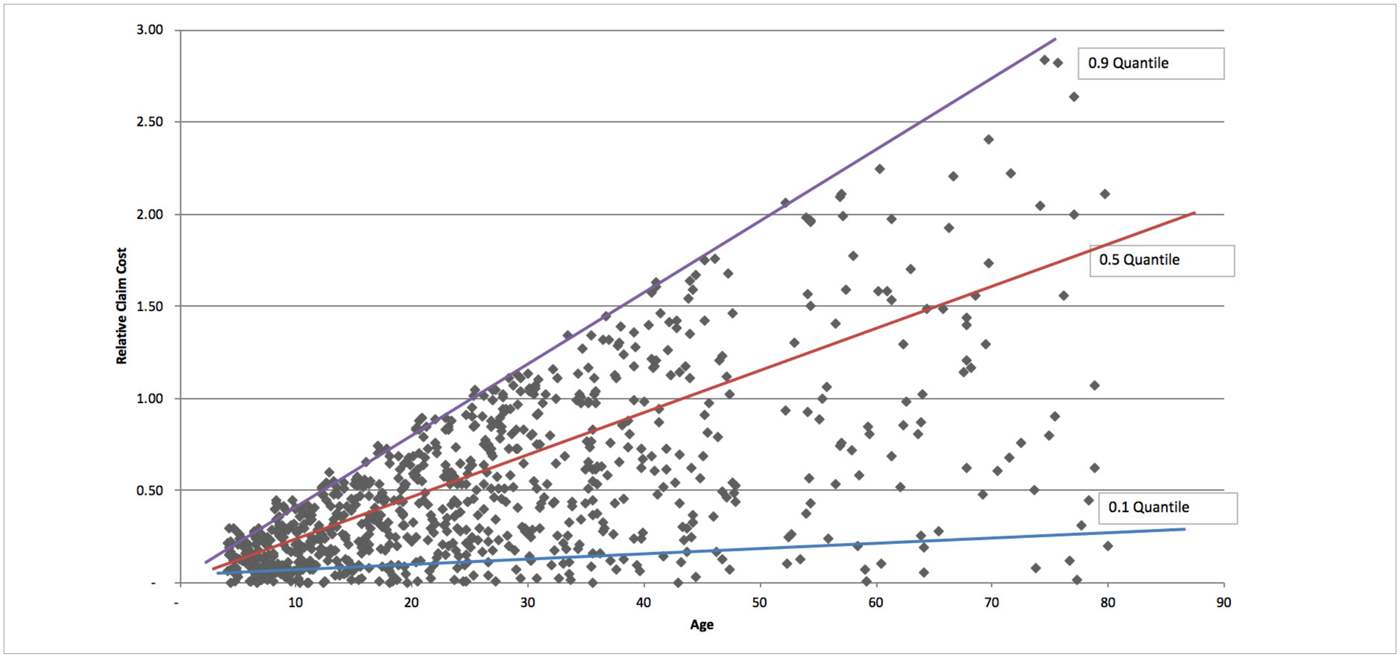

more accurate claims projections and the observation of different patterns of volatility across groups.

As a grossly simplified example, you can see from the following quantile regression equations how the impacts of age, gender and morbidity could change at different percentiles of the claims distribution:

$$y_{25th} = 0.6 * age + 1.2 * gender + 0.4 * morbidity$$
$$y_{50th} = 1.1 * age + 3.2 * gender + 4.8 * morbidity$$
$$y_{98th} = 0.3 * age + 0.5 * gender + 6.7 * morbidity$$

Note how the magnitude of coefficients changes at different percentiles, which would not be possible with traditional risk scores that just estimate the expected value.

$$\mathrm{E}(y) = 0.9 * age + 3 * gender + 5.2 * morbidity$$

Quantile regression has historically been disregarded due to computational intractability, but with the development of more efficient algorithms and increased computing power, it can now be used on large blocks of business. The simplified graph in Figure 1 shows a quantile regression on just the predictor of age.

In this one-dimensional case, quantile regression quantifies how the slopes (impact of a change in one year of age) of the best-fit line can change as we move from the 10th to the 50th and 90th percentiles, showing the differing impact of the predictor. Figure 1 suggests that the impact of increasing age is less for healthier members than for the sickest, and the slopes of the lines quantify that difference.

There are often good reasons for using traditional least squares regression and estimating only the expected value. When statistical assumptions are satisfied, the least squares estimates obtained are optimal, and it is solved with a computationally easy matrix equation. Indeed, assuming constant variance, the least squares estimates yield the "best linear unbiased estimate" of the expected value of claims, meaning that among all unbiased linear models that estimate the mean, the least squares estimate has minimum variance.

Unfortunately, the assumptions that are needed for those optimal properties to hold are strict: errors must be uncorrelated and have constant variance at all levels of the predictor. *The constant variance assumption is not satisfied by health claims*—variance at high levels of claim costs is higher than at lower levels, as Figure 1 illustrates for the age predictor.

ADVANTAGES OF QUANTILE REGRESSION

Quantile regression enables a deeper understanding relative to traditional claims estimation techniques.

1. **Quantile regression allows estimation of the whole distribution at different levels of covariates, so you can build confidence intervals.** Variability of claims at different levels of the predictors (for example, males aged 30 with hypertension versus females aged 50 with knee pain) may differ substantially. Quantile regression allows these estimates to differ by estimating different percentiles for each value of the predictors. Correspondingly, confidence intervals with different midpoints and widths can be produced that reflect those differing levels of precision.

Copyright © 2020 Society of Actuaries. All rights reserved.

Table 3
Claims Distribution Estimates Using Quantile Regression

	Percentile of Claims Distribution—Year 1				
	10th	25th	50th	75th	90th
Larry's Lumber	$337	$346	$352	$360	$375
Bob's Trucking	$345	$360	$366	$371	$377

Returning to our example with Bob's trucking, suppose our quantile regression estimates produced the distributions shown in Table 3.

For Larry's Lumber, a 50 percent confidence interval is ($346, $360) PMPM with a width of $14 PMPM. If we increase our confidence to 80 percent, the interval would be ($337, $375) with a width of $38 PMPM.

Conversely, the variability of claims in Bob's Trucking at 50 percent confidence is ($360, $371) with a width of $11 PMPM, and the 80 percent confidence interval is ($345, $377) with a width of $32. We observe that Bob's risk is lower, even though the median cost is higher.

2. **Quantile regression gives a distinct set of parameter estimates at different percentiles, so you know where they are significant and when they differ.** Using quantile regression estimates, the impact of covariates at different percentiles can be contrasted. Women may cost 3 times as much as men at the 50th percentile of claims but only 1.1 times as much at the 98th percentile. Using ordinary least squares, which has a fixed gender coefficient, would not allow the reflection of this variability.

Another application could be provider performance evaluation. For example, for each cardiologist you could determine the 10th, 20th, … 90th percentiles of cost corresponding to the panel of members attributed to that doctor, then measure the actual costs incurred by that physician against that distribution.

3. **Quantile regression enables us to observe at what parts of the distribution predictors are significant.** Using the wrong level of coefficients, or pricing with coefficients that are not really significant, can lead to mispricing across a block and anti-selection. It is critical to understand what part of the distribution is being priced to.

There are also disadvantages to the quantile regression approach. For estimating the mean of a group, traditional risk scores will be more accurate when the least squares assumptions are satisfied. Also, it is difficult to calculate these estimators correctly: a substantial amount of data is necessary to estimate many more parameters, and it must be verified that the algorithm found the "right" answer. For example, naive solutions in statistical packages sometimes lead to negative contribution to risk from some conditions. Finally, using quantile regression introduces ambiguity—it can be difficult to know whether an observed outlier at, say, the 90th percentile of the claims distribution implies that experience will regress to the mean or that a true shift in the distribution has happened.

Risk scores are a canonical example where quantile regression can be used to provide refinements to actuarial projections. Traditionally, actuarial risk scores compute the expected average contribution to future costs from age, gender and a set of conditions. So, for example, if a 55-year-old male member has stomach cancer, his risk score is a sum of demographics and conditions, such as 0.85 + 4.2 = 5.15. If instead we generated risk scores using quantile regression, we might see expected future costs as in Table 4, greatly enhancing our understanding of claim variability for this condition.

Table 4
Risk Scores Generated by Quantile Regression

Percentile	Demographic	Conditions	Total
25th	0.70	3.60	4.30
50th	0.85	4.75	5.60
75th	1.07	7.34	8.41

We could then compare the expected distribution of costs of groups of members with the given condition to the actual treatment costs by a panel of physicians to understand resource use deviating significantly from the average. Note that in practice an interaction term may be appropriate between age/gender and the condition at some percentiles of the distribution.

Copyright © 2020 Society of Actuaries. All rights reserved.

APPLICATIONS IN UNDERWRITING

Going back to our underwriting story, when broadly applied across a book of business, quantile regression can facilitate discount guidance to underwriters as well as improve operational efficiencies by identifying which groups are most likely to benefit from an underwriter's comprehensive review and which groups are straightforward and do not merit any additional attention. However, unlike traditional underwriting approaches that often focus on high dollar claims (which in many cases have resolved themselves), the focus is on groups where there are significant deviations, up or down, from expected.

The increase in computing power and development of efficient algorithms has opened up new frontiers for estimating complex probability distributions, allowing for more comprehensive and creative actuarial analysis that allows us to understand the true drivers of risk. Linear regression and its cousins allow for efficient estimation of the mean under some circumstances, but the flexibility, accuracy and empirical value of knowing the entire distribution at each level of a set of regression predictors argues strongly for a new statistical approach to actuarial claims estimation. ■

Jason Reed, FSA, MAAA, is a senior director, payer consulting, at Optum. He can be reached at *r.jason.reed@optum.com.*

Copyright © 2020 Society of Actuaries. All rights reserved.