Analysis: Bad data - A risky business for insurers?

Data is everywhere: good data, partial data, valuable data, bad data. Insurers need data to measure and price risks but do they take enough care in assessing the quality and provenance of the data they are using and are their processes sufficiently robust?

The tools that are now available to consume, manage and interpret massive datasets are so powerful that there is growing concern insurers could be feeding in data that contains underlying weaknesses, such as built-in bias, that will produce distorted, misleading and potentially costly outputs. In particular, the increasing use of unstructured data sources such as social media has taken insurers into new, unfamiliar territory.

It is not just these new data sources that contain the potential to throw insurers off course. Many established players are investing in huge data mining projects to trawl through decades of legacy data, much originally collected on paper. There are vast reservoirs of valuable intelligence, especially about past claims, buried in legacy data but it needs to be handled with care.

Amazon: ditching AI for recruitment

One of the world's leading pioneers in the use of AIto build a business has ditched it when it comes to hiring new staff. Last year, it emerged that Amazon had been battling against gender bias introduced into its recruitment processes because of the way in which AI amplified the bias buried in the historic data it was asked to analyse.

In 2014 it introduced a new system to analyse the thousands of job applications it received in order to rank the candidates and only bring in those with top ratings for interviews. However, by 2015, the company realised its new system was not rating candidates for software developer jobs and other technical posts in a gender-neutral way.

That is because Amazon's computer models were trained to vet applicants by observing patterns in applications going back over ten years. Inevitably, most came from men, a reflection of male dominance across the tech sector.

This meant Amazon's system taught itself that male candidates were preferable. It penalised CVs that included references to women's educational establishments and sports clubs. Although it edited the programs to make them neutral to these particular terms it discovered this did not provide a guarantee that AI would not devise other ways of sorting candidates that could prove discriminatory.

When these problems were exposed by Reuters in the autumn of 2018 the company pleaded it always sense-checked the decisions but admitted it had abandoned AI in many of its recruitment processes, especially where the use of historic data was involved.

Much of it was processed at a time when re-keying data was commonplace and that is known to be highly susceptible to errors – 10% is frequently cited as an industry norm. It was also collected in an era quite different to the world of 2019. The rapid pace of technological, cultural and social change in the last 25 years imposes a duty of caution when using legacy data. The potential for these projects to back-fire was illustrated by the problems Amazon faced when artificial intelligence imported and amplified gender bias into some of its recruitment decisions.

The world is also growing increasingly sensitive about the use of data, especially personal data. The roll out of General Data Protection Regulations last year raised awareness of how much personal data is now in the hands of business and has made people more willing to question how they came to have it, how they are using it and what benefit the individual gets from allowing their data to be used.

Analysis of data

Insurers base decisions about who to insure, how much to charge them and how to respond to claims on analysis of data but those decisions are now subject to a level of scrutiny previously unimaginable. Insurers need to have their proverbial ducks in a row to ensure they can survive in this new world with their reputations intact.

This is all part of a complex mix of potential commercial and reputational damage arising from using bad data, asking it the wrong questions or placing too much trust in the answers. The consequences can manifest themselves in a variety of ways, says Andrew Dunkley, head of analytics at lawyer BLM. "It depends on the type of decision that data related solutions are used to take, and the extent to which we completely delegate to those solutions rather than using them to inform or advise. The bigger the decision the larger the consequences of getting it wrong – but also the greater the benefits of getting it right."

He points out: "Some of these will be commercial decisions with commercial consequences such as mispriced policies, inadequate reserving and so on."

The result of these will be reduced market share and lower margins where they occur.

"Other decisions will have human consequences – claims rejected or not settled early enough, dragging people through litigation needlessly and otherwise imposing human costs. Over the medium term these human consequences arguably represent the bigger risk for insurers, as they may have the greater potential to drive a political or regulatory response."

Insurers are not ignorant of these problems, far from it says Lou Lwin, group head of operational planning for Markerstudy. "With regards to insurers being aware, I would say that the GDPR penalties are very real and, therefore, all businesses are mindful of the risks of using poor data. Specifically, from an insight point of view, blindly trusting numbers is not something insurers are known for, quite the opposite. If you know your business, you know roughly what your outcomes should look like."

He adds: "Machine learning, AI and deep learning all provide a high level of accuracy in a very short space of time compared to traditional analysis but the end result should be recognisable."

10 big risks from bad data

1. Poor pricing decisions leading to a loss of business or taking on risks that are underpriced
2. Incorrect understanding of exposures leading to unexpected losses and under-reserving
3. Inaccurate reinsurance purchases, either too much or too little
4. Misunderstanding of accumulated risk, especially with catastophe covers and new areas such as cyber
5. Delusional exactitude – being seduced by the precision of outputs into trusting figures that should be challenged
6. Black box algorithms producing decisions that cannot be explained, leading to regulatory action
7. Reputational damage due to careless or inappropriate use of personal data
8. Importing historical bias from legacy data, possibly amplifying it
9. Confusing correlation with causation and making bad underwriting decisions as a consequence
10. Becoming too risk adverse when too little data is available, making it hard to find cover for new risks

Rigorous scrutiny

It must also be rigorously scrutinised for the sort of bias that Amazon found AI introduced into its recruitment processes, warns Charlotte Halkett, head of insurance products at Buzzvault.

“There is a naivety about how machines can have bias. Bias can become magnified if you are not on your guard. You need to go into analysis of data with your eyes open and need to compensate for inherent bias in the data you are using.”

She says this should not stop the industry from pushing the boundaries in terms of the data it uses. “There is a real tension between enhancing the customer experience and making it really easy for the customer, and generating quality data. As an industry we need to serve people and we need to focus on how people want to live so I am very passionate about looking for new data sources out of the norm for the insurance industry.”

This includes the unstructured data sources such as social media but with heavy caveats. “It is about using the data for what it is. There is a huge problem with people polishing up their lives on social media. There is huge potential for creating bias if you don’t use it carefully.”

Applying caution

This caution over the use of social media and data is widely shared. “There is an intentional bias in that data as people want to make themselves look prettier, cleverer or have more friends. We have to understand that bias and make allowances for it before we use it,” says David Edison, partner and actuary at Moore Stephens.

“It can be used but with great caution because of the prevalence of fake profiles and the lack of an adequate set of rules behind it. We focus on using it for sentiment analysis,” says Ludovic Veale, data lead at Charles Taylor Insurtech.

Veale says similar caution must be applied to the data collected through the internet of things. “Getting hold of decent IoT data is one of the biggest challenges we face, especially ensuring the chain of connectivity. It has huge potential for insurers but we need clear protocols about how we should be collecting it.”

Veale says there should be a stronger collective approach to some of the data problems by the industry. “There are shared issues around governance across the whole sector. I would like to see a common sense approach looking at shared problems and taking a collaborative approach to solving them. It is all about identifying where there can be mutually beneficial outcomes.”

While the industry searches for those mutually beneficial outcomes, it also needs to develop a greater awareness of the downsides of the race to grab and process larger and larger datasets, says Edison. This is not just an issue for the data scientists but for boards too.

“More has to be done to manage the expectations of what can be done with data, right up to board level. What you don’t want is an understanding gap between the data scientists and the board.”

This could be especially important when regulators come calling and asking for explanations of decisions based on ‘black box algorithms’ where seemingly sensible outcomes are produced from datasets but where explaining them proves hard, if not impossible. This can be a challenge for a industry used to making pricing and rating decisions based on causation rather than correlation. Crude correlation is alien territory for insurers.

“You can have the best dataset in the world but if you don’t have the right people to analyse it and understand the output you are not going to make the right decisions,” says Edison.

He says attracting these people to work in insurance is going to be a challenge: “Actuaries are having to reinvent themselves as data scientists but this means they are not so insurance focused. The big tech companies are proving attractive to these people and they are taking some of the best people away from us.”

The right processes

Without this expertise and the right processes for checking the output of the new data-driven models there is a danger that people become too trusting of the results and are often seduced by the apparent precision of the outputs – 'delusional exactitude' as it is sometimes called by data experts.

This can be caused by poor data or incorrect assumptions at the input stage, says Jarno Seegers, associate vice-president, business development at insurance sector consultants Xceedance. It can have serious implications. This was the case with the estimates insurers made of the potential losses from major earthquakes in Japan, which proved inadequate for the 2011 Tohoku earthquake that caused widespread damage and led to the tsunami that hit the Fukushima nuclear plant.

"It may be a manifestation of an incorrect hypothesis in the science. The 2011 Tohoku event was measured as a magnitude 9.0 earthquake but earlier scientific consensus had ruled out the possibility of an earthquake of that magnitude on that specific fault line. Because of a preceding data point, some CAT models therefore assumed a lower magnitude event. This was a large factor in the divergence of potential loss numbers forecasted prior to the event."

In his view this highlights the dangers of placing too much trust in AI and machine learning, and especially of not being able to offer robust explanations for the outputs.

"The use of machine learning and AI can be hampered due to the input of bad data. AI solutions need to be carefully designed to cater to imperfections in the data. Common practice is to ignore some percentage of outliers in the data or recalibrate and validate the learned behaviours. For this, AI needs to be transparent. You need to be able to understand how decisions have been taken in detail. I would not accept a black box."

Regulators and the courts have already shown they are not always impressed by the industry's use of data to make rating decisions if they feel it leads to discrimination. The Test-Achats case showed the use of gender as a proxy for other risk factors would not be tolerated, so imagine how they might react to discriminatory decisions driven by black box algorithms the industry cannot explain.

A further Achilles' heel when it comes to managing its data could lie in a tendency to place too much faith in what they have and in the ability of technology to make sense of it, says Dunkle.

"At the moment the key danger is overconfidence – delegating too much to the technology, too quickly, without first thinking through the risks involved. There is a tendency for decision makers to adopt one of two positions: assuming either that the technology cannot solve certain problems; or that data driven solutions will always outperform humans. The reality is usually messier than that.

"Another major risk is that these models are built assuming that the future will work in the same way as the past. If an organisation uses a model to automate a process, that organisation is then stuck if that model is then rendered invalid by changes in society, technology or the law."

These dilemmas will intensify as competition, especially from new entrants, drives insurers to look for new datasets they can feed into ever more powerful analytical tools. While most suggest they have an awareness of the risks, the warning that some may be going too far, too fast cannot be disregarded.