




Building Trust in Analytics: What Role Should Insurers Play?

As hidden bias in data and models comes under greater scrutiny, the integrity of analytic processes used by insurance companies is being called into question. As the insurance industry relies on ever more sophisticated analytics, we must proactively work to improve processes and build trust in outcomes. This is more than just having the right technical elements in place. It also requires transparency, communication, and a clear demonstration that “fairness in data” is at the heart of corporate culture.



Over the past decade, insurers have embedded increasingly sophisticated analytics into their systems and processes. The industry is enjoying the benefits of easier access to large volumes of data, new information sources, and machine-learning techniques that, together, enable cheaper, faster, and more accurate data collection and analysis than traditional methods.

These benefits are penetrating — and in some cases transforming — all aspects of the insurance product cycle, including:

- Product design
- Pricing and underwriting
- Sales and distribution
- Customer service
- Claims management
- Loss prevention

Unintended Consequences

However, as analytics becomes more complex and widely used, important questions arise about its impact on fairness and financial inclusion. More specifically, there is growing concern that analytics is embedding and perpetrating hidden biases in insurance data and algorithms, potentially leading to outcomes that disadvantage certain groups in society through higher premiums, worse claims outcomes, and poorer service.

In the US, for example, evidence suggests that the historical practice of “redlining” — a discriminatory practice of segregating cities by ethnicity and income for financial purposes — continues to pervade current data used by insurers as well as other financial institutions.

In the 1920s and 1930s, color-coded maps were created to identify levels of risk in lending and insurance, with predominantly black and poor areas categorized as high risk (and highlighted by red lines). The impact of redlining was devastating for these communities, leading to lower investment, a lack of housing development, and reduced access to private capital, which in turn led to inhibited economic growth in these areas.

The practice of redlining no longer exists, but the economic damage continues today. This raises important questions for insurance companies. Although insurers are prohibited from using income and race to influence decisions, are other data points leading them to the same outcome? We know, for example, that ZIP codes, house maintenance, local crime rates, and other factors are used in underwriting decisions, and we also know that such data may in part be linked to the historical practice of redlining.

The challenge is not unique to insurers. Amazon, for example, reportedly abandoned an AI-driven computer program designed to vet applicants’ resumes after it emerged that the program had taught itself that male candidates were preferable. This was due to sampling bias — most historical resumes were from men due to male dominance in the tech industry. As a result, however, the program penalized applications that included the word “women” or mentioned women’s colleges.

These are just two examples of how historical inequalities or data biases can lead to potentially unfair outcomes. However, bias in data comes in many forms and is all around us:

“It can creep into the decision-making paradigm in subtle ways, whether it is the subjectivity of human judgement, prejudice, historical inequities baked into the data, or faulty algorithms. As a result, models can sometimes result in unreliable and unfair decisions by uncovering correlations in the data that are only reflections of those biases and historical inequities.”¹

Artificial Intelligence: Exacerbating the Challenge

Bias in data and models is not a new problem; even the most basic data analysis can draw false conclusions and trigger unintended consequences. However, the growing use of artificial intelligence (AI) is augmenting this challenge and triggering greater scrutiny into modeling techniques.

In particular, the complexity of AI-enabled processes requires special attention in the following areas:

What is fair and ethical?

There are many definitions of what constitutes “fair and ethical” use of data.

Several groups have put forward their own definitions in an insurance context. The European Insurance and Occupational Pensions Authority (EIOPA), for example, concludes that “Ethics is thought to mean approaches that are fair based on international and national recommendations, standards and treaties, and of course legislation.”

Some definitions consider fairness specifically in the context of artificial intelligence (AI). The National Association of Insurance Commissioners (NAIC), for example, states that “AI actors should... avoid proxy discrimination against protected classes.

AI systems should... avoid harmful or unintended consequences.” Although it is impossible to find an undisputed definition, the key for insurers is the ability to explain and justify their decision-making process based on a consistent understanding of fairness.

- **Learning:** AI systems “learn” from patterns and correlations in past and current data, thereby replicating any biases that may already exist in the sample. An algorithm may produce a racist or sexist outcome, for example, because it has been trained on data from a racist, sexist world. This creates unintended consequences that are perpetuated and embedded in a wide range of processes with greater efficiency and scale than traditional analytics.
- **Proxying:** One of the benefits of AI is that it enables more accurate consumer profiling thanks to its ability to process vast amounts of data points. However, this comes with ethical consequences. Insurers are legally prohibited from using specific information about a person’s race, gender, sexual orientation, and other protected characteristics that could expose them to unfair discrimination. However, if an AI system has access to enough nonspecific data points, it can combine the information to draw specific conclusions. For insurers, it is therefore possible to unintentionally (and unknowingly) arrive at the same conclusion as if the protected characteristic was used in the first place.
- **Transparency:** As AI algorithms become more complex, they become more opaque. At the same time, there is less human involvement than traditional analytics. Complex models can be made more transparent, but it is harder to unpack the process and decipher bias resulting in an enhanced risk that the end user is blindly following the data rather than making conscious decisions about the appropriateness of the outcomes.

How Should Insurers Respond?

The ethical use of data and the fairness of outcomes are not new concerns for the insurance industry.

Writing and pricing insurance policies is inherently discriminatory: different groups of people with different inherent risk characteristics will be treated differently. An individual with points on a driver's license, for example, will be charged more for car insurance. Or a family with a sophisticated alarm system will pay less for homeowners insurance.

This "discrimination" is uncontroversial and a core part of the insurance business model, and insurers have long understood that they have a duty to consider what constitutes "fair" discrimination. Insurers thus have a long track record of implementing checks and balances to ensure that outcomes are appropriate.

However, with society — and, increasingly, regulators — taking a more proactive interest in the fairness of analytics, the insurance industry has an opportunity (1) to ensure that its checks and balances remain effective in a more complex, AI-driven world and (2) to build trust in the use of analytics more broadly.

Checks and Balances: Three Perspectives

In assessing the current effectiveness of existing checks and balances, the key question for the insurance industry is whether the use of more sophisticated analytics is leading to unfair outcomes. The answer to this question can be considered through the following three lenses.

Actuarial lens: To achieve fair insurance outcomes, methodologies must be transparent, abide by rules and regulations, and be subject to appropriate scrutiny. In insurance, responsibility for this is predominately the role of actuaries who are trained and highly experienced in the ethical application of data and models.

Pricing is tightly regulated to ensure fairness. Actuaries are experts at subjecting models and algorithms to direct scrutiny regarding the potential for discrimination.

However, advanced analytics and AI techniques are increasingly used in other areas of insurance, such as fraud detection and claims handling, where there is less regulation and actuaries have less oversight or experience. The use of advanced analytics is therefore less scrutinized and could potentially be vulnerable to unfair outcomes. Insurers may want to consider measures to bring the same rigor found in pricing to all processes.

Societal lens: While actuaries provide important safeguards to ensure that insurance processes meet required standards, it is up to other groups in society to consider whether the rules in place are satisfactory in the first place.

Consumer groups and regulators are currently raising such questions: Are current rules watertight? Do they cover all relevant areas of the insurance lifecycle? Do they protect the right groups of people? Is there enough transparency for consumers to contest outcomes?

The outcome of these debates will have lasting effects on insurance processes, and it is critical that the industry have a seat at the table to add its own value to shape future outcomes.

Business models lens: Finally, the insurance industry must also engage in its own debate about how issues of fairness relate to its effective functioning. There may be a tension between fairness and economic outcomes because achieving truly fair outcomes requires relying on a so-called "veil of ignorance" — avoiding the use of certain variables. However, ignoring those variables makes it harder to accurately assess risk, thereby potentially affecting economic outcomes.

Broadly speaking, we accept that some commercial priorities should be sacrificed in the name of fairness. For example, certain characteristics cannot be used in pricing regardless of their correlation with loss outcomes. But it is equally accepted that insurers need to profile risks with a reasonable degree of accuracy for their business model to remain viable.

As we consider how to enhance fairness, it should also prompt debate about where we as an industry draw the line between accuracy and fairness and how best to manage the tension between the two.

Building Trust

Addressing these technical issues is important, but it's not sufficient to build societal trust in the insurance industry's use of analytics. Building trust requires a more holistic approach that actively demonstrates that the industry is taking issues of fairness seriously. Even if we apply the most water-tight checks and balances, they may not be visible or relatable to non-insurance stakeholders.

Building trust therefore requires more than the correct processes. It is about transparency, clear communication, and a consistent, corporate-wide understanding of what fairness and ethics mean in the context of data and analytics. Ultimately, these are questions of culture rather than processes.

Because it has more mature experience with data ethics, the insurance industry has the opportunity to be a leader in this field. But it must take positive steps both to further improve its technical processes and to build trust in analytical systems.

Here are three areas in which the industry can take positive and tangible action.

- **Consider the whole insurance ecosystem:** Insurers are responsible for requiring sound data governance across their entire data system, including service providers, insurtechs, and data providers. Conducting due diligence of data ethics in the insurance supply chain will become increasingly expected, just as insurers are required to consider their supply chain's impact on climate change and other ESG considerations.
- **Improve transparency:** Better communication about decision-making processes and the safeguards in place will not only build trust in the use of AI. It will also demonstrate to stakeholders that the insurance industry is taking a responsible approach. Transparency has two dimensions:
 - The need for internal transparency so that processes and outcomes can be scrutinized effectively
 - The need for external transparency to ensure that regulators, consumers, and other stakeholders can understand decision-making processes at an appropriate level of detail
- **Demonstrate a culture of fairness:** Engage in corporate-wide debate about what fairness means and how ethical challenges should be approached in a consistent manner. Consider issues such as the tradeoff between fairness and transparency, and build a common understanding of what fairness looks like. Use ESG reporting as an opportunity to demonstrate cultural considerations and report on progress on objectives.

Guidewire's Role

At Guidewire, we are exploring the role of data and service providers in addressing some of the challenges described in this paper. We are particularly interested in how analytics and data can be part of the solution, helping insurers monitor bias, increase transparency, and take mitigating action if appropriate.

For example, insurers already have robust checks in place to ensure they are not using prohibited data or proxies for such data. However, Guidewire Analytics can support insurers go further by automatically assessing if the combination of permitted variables may be at risk of creating an unfair outcome. If this risk is present, it raises a flag on the insurer's system to indicate that further due diligence is required.

With business impact monitoring, we can also evaluate (based on actual results) whether business KPIs show an unexpected or unjustified disparity.

Going one step further, we could even deploy our systems to enable a comparison of outcomes by running knowingly “unethical” models in parallel with an insurance company’s real-life models. If both models deliver similar results, the real-life model may be unintentionally discriminatory.

Not only would such a system provide an additional layer of checks and balances from a neutral third party. It would also offer objective evidence to stakeholders that insurers are taking its obligations seriously — helping build trust in the industry’s use of analytics.

These ideas are in early development, and we welcome feedback and input from all of our P&C colleagues. Guidewire is not aiming to solve all the challenges identified in this paper. That would take time and industry coordination with regulators and consumers. Our goal is smaller but more tangible: to provide meaningful and achievable steps to visibly improve processes that build trust in the industry’s approach to fairness in the context of analytics.

Conclusion: Lead the Debate

Governments and society are increasingly aware of the potential ethical challenges posed by the more sophisticated use of analytics and AI. As a result, this is likely to become a more heavily regulated area, not just in insurance and financial services but in a wide range of sectors.

The industry should welcome this and take an active participation in the debate. Insurers are experts in the use of data, and the industry is more mature than most in considerations about fairness. Although there are many more challenges to overcome in this fast-moving area, the industry as a whole can be a beacon for the ethical use of advanced analytics if insurers take a proactive stance.

To achieve this, the industry must continually improve technical processes and be more proactive at building trust in the outcomes that AI-enabled systems can produce. Ultimately, analytics is just a tool — it has no normative value. It has huge potential benefits for society, but it also comes with risks.

It is up to the users of these systems to ensure that they are deployed in a way that has a positive impact. Guidewire is committed to working with its customers and partners to support this process.

1. Roosevelt Mosley and Radost Wenman, Methods for Quantifying Discriminatory Effects on Protected Classes in Insurance, Casualty Actuarial Society Research Paper Series on Race and Insurance Pricing (2022).

Guidewire is the platform P&C insurers trust to engage, innovate, and grow efficiently. We combine digital, core, analytics, and machine learning to deliver our platform as a cloud service. More than 500 insurers, from new ventures to the largest and most complex in the world, run on Guidewire. For more information, contact us at info@guidewire.com.