



# The Arbitration Review of the Americas

2019

**Damages Overview**

# The Arbitration Review of the Americas

2019

---

The Arbitration Review of the Americas 2019 provides an unparalleled annual update – written by the experts – on key developments. It is their wealth of experience and knowledge – enabling them not only to explain law and policy, but also to put theory into context – which makes the report of particular value to those conducting international business in the Americas today.

---

**Generated: February 8, 2024**

The information contained in this report is indicative only. Law Business Research is not responsible for any actions (or lack thereof) taken as a result of relying on or in any way using information contained in this report and in no event shall be liable for any damages resulting from reliance on or use of this information. Copyright 2006 - 2024 Law Business Research

# Damages Overview

**James Searby**

FTI Consulting

## Summary

PROBABILITY AND THE LAW: A COMING COLLISION?

PROBABILISTIC TECHNIQUES ARE NOW A FEATURE OF EVERYDAY LIFE

BIG DATA AND PROBABILISTIC TECHNIQUES ARE A POTENTIALLY POWERFUL COMBINATION

CONCLUSION

## PROBABILITY AND THE LAW: A COMING COLLISION?

The assessment of evidence has been a human process throughout legal history: evidence is weighed, determinations are made about its truthfulness, its importance to the case is assessed and the trier renders a decision having regard for whether the evidence presented meets the required standard. In principle, the concept of probability is not new to the legal profession. Standards of proof such as 'balance of probabilities', 'reasonable certainty' or 'beyond reasonable doubt' all express different probabilistic thresholds, each higher than the last, that the evidence presented in the case must meet before a tribunal can be persuaded of a claimant's case.

The law has always recognised that complete knowledge of the facts is often impossible and that human decision-making is fallible. For those reasons, standards of proof are not absolute. In the absence of certainty, the best that can be hoped for is an assessment that compares an understanding of the facts to an understanding of the law. Given the degree of uncertainty that may be present in legal disputes, most legal determinations require triers to deal in relative likelihood, or probability.

The purpose of this article is to consider – in an introductory and discursive way – the interaction between uncertainty, probabilistic techniques and the assessment of economic loss. To the extent that the examples given touch on questions of law, they are intended for illustration, rather than expressing a legal opinion.

## PROBABILISTIC TECHNIQUES ARE NOW A FEATURE OF EVERYDAY LIFE

It may help to begin with an example that contrasts certainty with uncertainty.

Traditional computing of the kind relied on by banks, airlines and law firms for, say, transaction or document processing, assumes the presence of certainty. Computer programs of this type follow a set of instructions that contain absolute rules about how to process data, depending on whether the data meet defined criteria. For example, such programs could be designed to calculate salaries or pay taxes at the end of each month, to issue airline tickets if the customer had shown proof of identity and provided payment or to select documents or emails matching specific criteria from among thousands or millions of similar documents. These types of computer programs are immensely useful for automation, the performance of repetitive tasks more rapidly and at lower cost than any person could manage. They are unable, however, to manage uncertainty, rendering them unhelpful for the almost infinite range of human tasks that require judgement.

Computing today has entered a new era in which programs explicitly seek to surmount the problems associated with uncertainty by using probabilistic techniques. Those techniques seek to determine if data meets criteria well enough for an action to be carried out. In other words, instead of determining absolutely whether data fully satisfies defined criteria, programs are now called upon to assess how close a match is, and if it is sufficiently close, then to carry out a specific task.

Such evidence-based, probabilistic determination (sometimes referred to as artificial intelligence or AI) is at the heart of the algorithms that drive much of modern life. Google weighs what it 'knows' about your browsing history and selects the advertisement to serve you by matching your browsing history against the requirements of the advertiser. Facebook weighs what it 'knows' about your friends and their interests in selecting the news stories and

other items to place in the feed of your home page. The software in a driverless car, detecting a white object on the road in front, weighs what it 'knows' about white objects and tries to determine if it is a plastic bag, a dog, a concrete block or a child, before deciding what to do.

In some sense, the computer is following a similar path to a trier of facts. It compares the facts (eg, the white object in the road), to a set of principles (eg, how the white object moves, what shape it is), to determine an action (eg, to brake or proceed), subject to the degree of certainty it assesses.

But unlike judges and arbitrators, the computers making these decisions do not know anything. Rather, they are driven by algorithms that compare the data they have – whether your browsing history or the white object in the road, to other data – such as the browsing history sought by advertisers or other white objects in its database – before deciding whether the 'match' is good enough and acting accordingly.

Such a match is not a matter of certainty: it is unlikely that the advertiser is looking for somebody with exactly your browsing history, and the shape of the white object is probably not exactly the same as the shape of the white objects the algorithm has encountered previously. Each of these processes therefore requires an assessment of probability: how close enough is good enough to trigger a match that allows an action by the machine to occur?

Much of the time, the decision of the algorithm will be eerily accurate, lending people to conclude that computers really do know what is going on.<sup>2</sup> Many readers, however, will be familiar with other instances in which the computer made the wrong decision, leading to an outcome that was comic or tragic, rather than eerie.

A wrong decision made by an algorithm usually stems from two main sources: either the code itself can make inappropriate determinations, or the data on which it was trained was not representative of the data it was likely to encounter in operation. Working out which of those is the case, however, is not a simple matter for three reasons.

First, many AI routines today are 'black boxes' that make (millions of) decisions about what to do but often keep no records of how those decisions were reached. The absence of an audit trail can frustrate efforts to understand why the computer decided to serve up inappropriate advertisements or to determine that the child crossing the road was really a plastic bag. Exact recreation of the circumstances that caused, say, a car to fail to brake may not be possible because perfect replicas of specific children, dogs and blowing plastic bags are unattainable. Further, the amount of data processed from moment-to-moment is so large (as is the number of individual decisions made by the program) that it may not be practical to store a replica of what is used to support every decision made. Put another way, it is hard to recreate the facts.

Second, understanding what an algorithm 'knew' is complicated by its interaction with the data used to train its decision-making. This 'training data' informs the criteria that the algorithm uses to determine if the data it encounters in a 'live' environment meet the criteria for action or not. If the data used is not representative of the task for which the algorithm is designed, problems are likely to follow. In one well-publicised case, facial recognition software exhibited what appeared to be racial bias in its decision-making because it was trained on data that did not include sufficient non-white faces.<sup>3</sup> The data sets used for this purpose are very large and constantly evolving with each encounter between the algorithm and new data, so it can be hard to understand how the criteria used for decision-making

were formed (did we use the right data?) and, harder still, whether the criteria for action were appropriately assessed (did we analyse the data correctly?).

Third, even if it can be determined what the algorithm was doing and how it reached its conclusion, there remains the question of the appropriate threshold of certainty to apply to the computer's finding. Clearly, the appropriate standard of proof should depend in part on the potential consequences of the decision: driverless cars should be subjected to higher standards than advertising servers. That observation, however, tells you nothing about the threshold needed for any given algorithm, whose 'facts' may rely on different sensors and whose 'law' may rely on different training data from those used in other algorithms. Some algorithms might also be better 'judges' than others, even given the same 'facts' and 'law'.

The decision about standards of proof is complicated by the need for algorithms to make determinations that are non-binary. The consequences of running over a white object range from the trivial (it is a plastic bag) to the unfortunate (it is a dog) to the dangerous (it is a concrete block) to the tragic (it is a child), suggesting that different thresholds might be needed if there is a possibility that a detected object is a person or an animal or a thing and, if a thing, probably what kind of thing.

### **BIG DATA AND PROBABILISTIC TECHNIQUES ARE A POTENTIALLY POWERFUL COMBINATION**

But what does any of this have to do with damages? I have previously written about Monte Carlo simulation and the use of probabilistic techniques in the assessment of damages. That article was written in response to scepticism that I had heard expressed about whether statistical techniques could ever meet legal standards of proof requiring, say, reasonable certainty. In my view, the answer remains a qualified yes. Legal processes are well-adapted to coping with uncertainty and there is no reason in principle why tribunals should reject conclusions drawn on the basis of statistical inference.

There are three main qualifications to the use of probabilistic techniques in an assessment of damages, each of which has parallels with the issues raised with AI algorithms above.

The first is that the appropriate use of probabilistic techniques requires calibration. As with the training data used to train, say, facial recognition software, the data used in simulations must be fit for the purpose to which it is being put. If a simulation of future prices relies on a history of, say, weekly prices from the last three years, it is likely to reach a different conclusion to a simulation trained on daily data over six months or monthly data over 10 years. Which of these sets of data is the right one to use is not a simple question. Likewise, knowing when to stop adding data may be a growing problem: when it is simple to add ever more data, knowing when to stop (because additional data does not add to an understanding of the problem) requires expert statistical assessment.

The second is that statisticians may themselves disagree whether the real world populations of data used in a probabilistic assessment actually fit one of the mathematically-defined distributions, such as the familiar 'bell curve' shape of the normal distribution. The good news on this point is that the ever-growing collection of data of all kinds is helping to answer questions about distributions in ways that were previously impossible. The volume of data generated by manufacturing plants, driverless cars, online purchasing, delivery records and everything connected to the Internet of Things will reveal a great deal about human behaviour – including commercial behaviour – that was previously only guessed at.

The third is that the outcome of the process should have an underlying logic to it. As the saying goes, 'correlation does not imply causation'. Users of statistical data therefore need to continue to apply common sense to their findings, lest they become beguiled by statistical relationships into drawing false conclusions.

It may serve to offer an example. Let us say that a distributor forces through a price change on certain product lines. Affected retailers contest the price change on account of the alleged harm that it does to their business. An assessment of the loss caused to the retailers relies, first, on the direct effects that the price increase has on sales volumes (an economic phenomenon referred to as the 'price elasticity of demand'). To the extent that the higher revenue per unit is more than offset by lower volumes, there will be a loss of profits, and vice-versa. In addition, however, there may be indirect effects. Lost sales volumes of the products whose price has increased may lead to lost sales of ancillary or complementary products that customers purchased at the same time, tending to increase the loss. It may also cause customers to switch their purchase to competing products sold by the retailer, potentially mitigating some of the loss from sales of the affected products.<sup>4</sup>

Assessments of price elasticities and cross-elasticities of demand have historically relied upon small data sets that are statistically unreliable. Intuition has tended to fill in the gaps in the evidence. Today, however, it would be possible to calculate extensive cross-elasticities of demand for a potentially wide range of products given several years of sales and pricing data from a large retailer. Nevertheless, the data in question would still only be a sample – the 'true' cross-elasticities might be different – and the data is likely to contain quite a lot of 'noise', for instance associated with different income levels or demographics around different stores.

The availability of such a large volume of detailed data, however, seems likely to open up possibilities for analysis that is beyond ordinary human capabilities. An appropriately-trained AI algorithm might well find patterns of gains and losses in the data that were unanticipated, even by a person acquainted with the business and the products. As discussed above, however, such evidence might need to be weighed to ascertain whether it met the legally-defined thresholds of certainty. It will also need to pass common-sense tests: even if there is a high correlation between, say, changes in the prices of facial creams and demand for motor oils, that does not prove that the demand for one is causally related to the demand for the other.

## CONCLUSION

This article advances the view that tribunals ought to gain comfort in assessing evidence using probabilistic techniques in assessing economic damages, as a general proposition, simply because with the application of new AI algorithms to an ever-growing volume of data, it is likely to become much more common. Until relatively recently, courts and tribunals often considered discounted cash flow (DCF) calculations to be unreasonably speculative. Today, however tribunals have become more comfortable with adopting them – where appropriate – recognising, perhaps, that DCF calculations are the most usual way for investors to assess projects and investments.

It is possible that increasing familiarity with the probabilistic techniques that underpin modern algorithms will lead future courts and tribunals in the same direction as algorithms and data sets become more pervasive and robust than they are today. When the use of probabilistic inference becomes a common phenomenon of which people have daily, direct

experience, it may be that the techniques will lose some of their mystique. It may also help to bear in mind that what is going on in statistical inference is a comparison of collected facts to principles, and applying standards of proof to the comparison, a process that mimics the essence of legal (and much other) decision-making.

#### Notes

[1](#) The views expressed in this article are those of the author and not necessarily the views of FTI Consulting, its management, its subsidiaries, its affiliates, or its other professionals.

[2](#) A more accurate description might be that the accuracy of the algorithms used today has led people to become aware of the surveillance that they have been under for a considerable period already.

[4](#) The effect of a change in price of one product on demand for a second product is known as a cross-elasticity of demand.



---

**James Searby**

[james.searby@fticonsulting.com](mailto:james.searby@fticonsulting.com)

---

200 Aldersgate, Aldersgate Street, London EC1A 4HD, United Kingdom

**Tel:** +44 20 3727 1000

<https://www.fticonsulting.com/>

**Read more from this firm on GAR**