

AI AND DISPUTE RESOLUTION  
WORSHIPFUL CPOMpany POF ARBITRATORS  
MASTER'S LECTURE  
(Slightly edited)

1. I am a relatively recent convert to IT. Despite having a science degree, I had neither an email address nor a mobile phone when I became a Judge in 1996. Even a compulsory judicial IT course failed to bring me into the fold. It was only when, to my children's intense amusement and disbelief, the Lord Chief Justice, under a misapprehension that as the youngest member of the Court of Appeal and the only member with a science degree, I would be very much up on the subject, appointed me as Judge in charge of IT in 2005, that I belatedly embarked on a relationship with the electronic world. A hasty crash course over Easter ensued.
  
2. That was, frighteningly, 20 years ago, and you do not need me to tell you how much has happened in the IT world since then. As a technical measure, at that time the industry was celebrating the fact that a transistor of less than 100 nanometres had been manufactured for the first time, whereas today, a transistor of two nanometres is in use. On a more capitalist measure, Nvidia's share price, now around \$135 was 17 cents in 2005; Apple, now around \$209 was 40 cents in 2005. More relevantly for today's purpose, in 2005 coders were the essential experts, now, as Alex Carp, co-founder of Palantir has recently said, "coding is much less important": AI machines are starting to code, and will probably soon be doing it faster and more reliably than humans – and anyway, now they can effectively understand and speak English, one may wonder how much longer will they need coding.

3. This last example reflects the big change which everyone is talking about, namely the advent of Artificial Intelligence, which, rightly or wrongly, having been treated as a possible or even likely development sometime in the future, was seen, again rightly or wrongly, as being a reality with the launch of ChatGPT on 30<sup>th</sup> November 2022.
4. I say “rightly or wrongly” because it depends what you mean by Artificial Intelligence. According to some experts, AI is a misnomer, and should be described as “Machine Learning”. That is a useful expression because it can be said to describe a wider range of machines than AI, which according to some definitions, we have not yet got.
5. Whether or not Artificial Intelligence is the right expression, Google, defines it as involving “*computers and machines that can reason, learn, and act in such a way that would normally require human intelligence or that involves data whose scale exceeds what humans can analyse*”. That could be said to apply to the earliest computers and word processors, and even the first calculators. But these items were as it were tucked under most peoples’ cognitive belts, before they had really heard or thought about AI.
6. It may well be that the first time many people became properly conscious of the concept or existence of AI machines was in 1997 when IBM’s Deep Blue beat the world chess champion, Gary Kasparov. This was an example of limited AI which involves what is often called a reactive machine: i.e. one which reacts to stimuli according to pre-programmed rules. In a sense I suppose such a machine is like a glorified three dimensional specialised calculator.

7. We have now moved to a higher level, that is machines with limited memory: unlike reactive machines, machines with limited memory can be trained with new data. At any rate at the moment such machines, like reactive machines, can only carry out specific tasks, based on their programming and training. This basically involves amassing an immense amount of relevant information experiences, insights and even in some cases emotions, and then providing the machine with some sort of programmed guidance as to how analyse that material to produce a specific type of outcome. Such programmed guidance can be on a supervised basis which involves consciously feeding the machine with information with a view to producing specific outcomes, or on an unsupervised basis, which simply involves feeding the machine with information, and very general guidance, and leaving it to the machine to work out the outcomes.
8. The distinction is well illustrated by reference to another sedentary sport, the Game of Go, which was long thought by many to be too complex or mega-variable for a computer to master when Deep Blue beat Kasparov. However, over a number of years, Deep Mind developed AlphaGo, which by 2016 was beating world champion Go players. It had been trained to find its moves based on knowledge acquired through supervised machine learning, which involved extensive training, based on actual games of Go, some played by humans, others by computers. Within a year or so, Deep Mind's newly developed AlphaGo Zero, was massively outplaying AlphaGo. Remarkably AlphaGo Zero was trained on a wholly unsupervised basis, i.e. without any input of actual games: effectively the only input was the rules of Go, whereafter the machine taught itself how to win by playing endless games against itself, and it took 40 days to reach a level where it could reliably outplay AlphaGo.

9. As we all know, machines with limited memory are not just game-players. Often described as LLMs (large language models), they are designed for natural language processing tasks such as language generation, and enjoy self-supervised learning with the benefit of a vast amount of text. Perhaps the most startling recent example is AlphaFold which can quickly carry out the electronic exercise of working out how a particular protein or other polypeptide will reshape to conform with physiological receptors – a problem which molecular biochemists can take many months to solve. This is a very important breakthrough with potentially enormous implications for pharmaceutical discoveries, albeit that it is in a somewhat niche area. A better known and more generally applicable example of such machines are of course Chat GPT and DeepSeek.

10. So we are in this second stage, machines with limited memory. The next stage, the third stage, is AGI, artificial general intelligence, or theory of mind, which is a machine which can think and act like a human. It will be far more versatile than any current machine. And, frighteningly to many (including me), it will presumably be able to outperform any human, if only because it will be able to survey and process far more information far more quickly. We are not there yet, and it is a matter of speculation among experts and non-experts when we will get there – even whether we will get there – and what the implications will be when or if we do. As Richard Susskind has put it, we are “*still feeling our way through some agonisingly difficult and unmapped terrain*” which will lead to “*questions we haven’t yet imagined relating to problems that haven’t yet arisen from systems that haven’t yet been invented*”. To me at least, this is challenging in the same way as being told by experts on string theory that there are in fact eleven dimensions, although we can only appreciate four of them because of the way in which our brains are constructed.

11. As I have been discussing, AI has already been used for carrying out what may be characterised as low level professional roles. However, certain jobs which would be seen much higher level are now better done by AI than the most expert humans. It has for some time been clear, for instance, that an AI machine can identify whether a mole on your skin is cancerous or not more reliably than the most expert human dermatologist. That diagnosis is a classic example of an exercise for which AI is well suited. One can simply train a machine by showing it tens of thousands of moles with the information whether each mole was or was not cancerous, and it then carries out the pattern-recognition exercise which a human dermatologist will carry out – but the machine will have seen far more moles than any human could, it will not forget any of them as a human would, and it will be able to scan them all far more quickly and reliably than any human could. The precise criteria which the machine employs may be a mystery, but it gets the answer right almost every time.

12. Of course, that does not mean that a machine can resolve a legal dispute. What one might characterise as a one-trick pony AI machine was not a great surprise to those in the IT world, however uncomfortable it may be for dermatologists. However, it appeared to many that legal problems and the resolution of legal disputes could rarely, if ever, be resolved by machines, as that would require AGI machines – i.e. as I have mentioned, AI machines with so-called general intelligence, which match or exceed human cognitive capabilities across many, or even all, areas of economically valuable work or cognitive tasks, which many people thought – and quite a lot of people still think – will never be achieved.

13. But the development of LLMs, of which ChatGPT is now but one example, has cast very severe doubts on such a view. These machines are designed to process natural language which have many parameters, and are trained with self-supervised learning on a vast amount of text, and which can be trained for specific tasks. And when I say “specific” tasks, the specificity can be pretty wide.
14. Thus, a couple of years ago GPT-4 passed the Uniform Bar Exam (UBE), which is the qualifying examination for practising lawyers in most US states, with a multiple-choice component, a six-essay component and a brief-writing component. And GPT-4 didn’t just pass – it passed with flying colours doing considerably better than the average score of human examinees, being in the top ten percent. And, when it comes to drafting contracts, or preparing an advice, many experienced lawyers have found, with a sense of mingled delight and horror, that Chat-GPT will do as good a first draft as many competent lawyers, and will often come up with ideas which the experienced lawyer has not had and is happy to adopt.
15. And some courts are using AI already. This is particularly true of China, where the judicial use of technology is, I think, more extensive than any other country. Chinese courts are using computers to suggest legal outcomes to judges. Thanks to IT, previously unpublished decisions are included in what is a huge reservoir of publicly available material, and software capable of analysing past cases with similar fact patterns is being used by human judges in making some decisions, for instance sentencing, for example, sentencing decisions, so that consistency of approach could be assisted. Critics complain that, although judges are permitted to disregard the recommended sentence, it is likely that the result produced by the software will have an anchoring effect. But it seems to me that the

answer to that is that anchoring is precisely what is intended, as is indicated by the description of this Chinese policy, namely “*similar results for similar cases*”, which has a ring of fairness about it. However, it can be seen as having a somewhat sinister ring as well, in that it facilitates centralised control over judicial decision making.

16. A somewhat different problem arises in the US, where the criminal courts in many states use a scoring algorithm to assess whether defendants should be granted bail pending trial (by reference to the risk of their absconding) and to assess the risk of convicted offenders re-offending, which is then used as a factor in sentencing decisions. However, a study of one of these systems established that it showed an ethnic bias against people of colour, in that black defendants who did not reoffend were nearly twice as likely to be misclassified as higher risk, compared to their white counterparts.

17. It would be interesting to see if a similar discriminatory outcome was occurring in comparable courts which did not use an algorithm. At least judging from recent newspaper reports about a Sentencing Council recommendation, it appears that it is at least possible that we have a similar discriminatory problem in this country where judges do not use algorithms. In any event, I rather suspect that the machine’s racial bias was a reflection of human bias in the programming, and so it is scarcely much of a reason for eschewing IT input into judging.

18. Having said that, this experience does emphasise that no sensible lawyer will rely on ChatGPT or any other AI/LLM system to produce a document which is not then rigorously checked, and no doubt improved, before being sent out to a client, the other side, or a tribunal. As is notorious, machines

invent and cite non-existent cases, and sometimes simply get things wrong. However, we are at an early stage of LLMs, and I think one can be reasonably confident that with the passage of time, and quite possibly not much time, these glitches may become a thing of the past.

19. However, the argument that, because of the risk of glitches, we can never leave things to a machine seems to me to be taking that point too far. Take driverless cars. In due course, it looks likely that the technology will develop so that driverless cars scarcely ever have accidents. There will then, I suspect, be a view held by some that “scarcely ever” is no good: until driverless car technology can deliver guaranteed safety, it should be eschewed. But that seems to me to be misconceived. Once it is clear that driverless cars have a provably and substantially better safety record than cars driven by humans, that should be more than good enough. (A good example of saying that the best is the enemy of the good).

20. For a long time there has been a debate as to whether AI will develop so as to be able to replace judges and arbitrators. Because emotional, societal and moral judgments appear to be based on something other than logic, there are people who think it will never be done. All the more so when one recalls that, especially in our common law system, judges develop the law, which is particularly demanding in a world which is so fast-changing, technologically, socially, culturally and morally. Is that something, many people wonder, which a machine could ever do? And, even if it is, will people trust machines to resolve their legal disputes?

21. The contention that a machine could resolve disputes as reliably as or more reliably than a human judge or arbitrator is a question which has yet to be established. So far as resolving legal issues is concerned, the chances of a



machine being able to do a first class job must be pretty good given that machines already are doing well at law exams. The point is reinforced by the good record which certain computer programs already have shown of predicting the outcome of US patent disputes and European Human Rights Court decisions.

22.As to factual disputes, I would have thought it likely that if the case turned on whose evidence was more consistent with the contemporary documentation or even whose evidence was more consistent with commercial common sense, a machine will, at least one day, similarly cope well.

23.What about cases which turn on an assessment of the witnesses? There have been many articles stretching back into the last century recording how unreliable humans, including judges, are at telling a truthful witness from a liar. And now there is evidence from research carried out in Germany which shows that AI can be programmed so as to outperform humans in detecting liars. Its record was not great – it was right two-thirds of the time; but other research suggests that humans are right just over half the time.

24.And then there is the argument that the law, especially the common law, needs to be developed by judges to adapt to the requirements of changes in social, ethical, societal, and technological changes, a particularly important feature in fast changing world. Can a machine reliably perform this function? While I readily accept that such a type of assessment seems to be currently less suited for a machine bearing in mind what machines seem able to achieve, I expect that this may well change as the technology improves. Thus, I imagine that AI will probably be able develop so that it can accurately mimic human judgments morals and emotions.

25. After all, our emotions, morals and principles are all produced in our brains, and our brains are a system of neurons that communicate via synapses using a combination of electrical and chemical signals, with tens of billions of synapses. By contrast, an AI machine is basically an electronic system of chips or electronic circuits consisting of trillions of transistors. Just as our human ability to recall information or to carry out mathematical calculation is the result of electrical and chemical signals in our brains, so are our moral, emotional or principled assessments. Unless the human mind includes some sort of unreplicable divine spark, which I find hard to believe, then it seems to me that there must be a good chance that, just as our ability to recall and to calculate can be replicated electronically in machines, so too, I suggest, can our ability to form moral, emotional or principled judgments. Developments such as Chat GPT and DeepSeek do nothing to undermine that belief.

26. There must at least be a powerful case for saying that the fact that principles or emotions are developed by electronic neural systems rather than by human neurons would not render them any less effective or authoritative. Indeed, in practical terms it may be easier to ensure that there are no unconscious biases in a machine than in a brain. It may be argued that this analysis disregards the fundamental moral requirement that justice must be dispensed by humans. Maybe I am not sufficiently grounded ethically, but I regard that point as being not so much an abstract ethical principal, but one of practical acceptability. If justice is dispensed in a way which does not have public confidence, then that is wrong in principle. So, I think the question is more to do with public acceptability than abstract moral principle.

27. Would public opinion accept machine-based justice? Most people with a suspect mole given the option of a human dermatologist or an appropriately trained and more reliable AI machine to diagnose their condition would, I think, opt for the machine. But I doubt that that would be true when it comes to the resolution of a legal dispute – especially for the losing party, who is often described as the most important person in the case, when it comes to considering the addressee of a judgment. At any rate at the moment, there would be obvious and strong concerns about machine judges.

28. AI machines are often described as reaching “black box” conclusions – we don’t know why they reach their conclusions. But if machines can pass bar exams with flying colours, it is hard to believe that it will not be possible for machines to resolve legal disputes, and to give reasoned judgments and awards. There could still, however, be a concern that, as the machine are “black boxes”- i.e. they have thought processes which are opaque - even if they purport to give reasons, how can one tell if they are their real reasons? As the highly respected computer scientist, Stephen Wolfram, put it, when it comes to training neural networks within AI machines, “*it’s complicated in there, and we don’t understand it – even though in the end, it’s producing recognizable human language*”. And if Mr Wolfram doesn’t understand how AI machines actually come up with their answers, nobody does.

29. And, while I have talked about AI machines “mimicking” human thinking, that is more true of outcomes than reasoning. Even when one considers the reasoning in, say bar exams, the machines are producing what they are programmed to understand is human reasoning, but they do not produce it in the same way as humans. Accordingly, while we may not have a full appreciation of how a machine produces an answer to a legal dispute, we

at least know that it does so differently from a human brain. So, it may be said, it is wrong for machines to be dispensing justice in civil or criminal courts or in arbitrations if we do not know how they arrive at their results and we do know that it is in a different way from that in which humans arrive at their results.

30. However, the view that this is a reason for rejecting AI-based dispute resolution is, I suggest, very much open to question. If an AI machine is a black box, it could be said with equal force that our brain is a black box. I have already mentioned that a machine works electronically, whereas a brain works chemically and electrically. In each case, however much we may not like to admit it in the case of our brains, the output is determined by physical events involving electron movements and, at least in the case of the brain, molecular interactions. And in many ways we have a better idea how an AI machine works than how our brains work..

31. We know what goes into a machine, and we can, at least to some extent, see how a machine reaches its decisions, by looking at the wiring and working out the electronics. But we cannot now at least currently carry out the same degree of examination to see how the human brain reaches a result. And the fact that a machine arrives at its result in a different way from humans is not by any means a reason for rejecting machine-based dispute resolution. “Different” does not mean “worse” as anyone who has an understanding of diversity will appreciate. In the end, what justice ultimately requires is right answers, but at least for the moment at any rate, I accept that a fair procedure is regarded as being as important – hence the rules of natural justice, the requirement of open justice, and the need for reasons.

32. In any event, can we be confident that a human judge or arbitrator's reasons are the true reasons – either subjectively or objectively? I must confess that at least sometimes I have wondered about the reasons I have given for a decision. When there is an argument over the meaning of a provision in a contract or statute, the answer is often really based on how the words strike you, and I sometimes felt that the reasons I gave were retrospective justification for my instinctive reaction. Karl Popper showed that scientific discoveries normally involved an inspired or instinctive idea, for which the scientist then develops a sort of retro-fitted explanation. So, too, I suspect that many judicial decisions are based on what the Judge feels is the right answer which he or she then seeks to justify in an appropriately reasoned judgment. Some people feel that this is almost shocking, but I see nothing wrong with it. One of the main reasons for requiring judges to give reasoned judgments is to make them justify their decisions not just to the parties, the appellate courts, and the public, but to themselves. Any good judge or arbitrator will have had the experience of realising that their instinctive – or even carefully thought out - view was wrong once they came to try and justify it with reasons. So, the requirement that a judge gives reasons also produces more reliable outcomes.

33. (But not everyone is a fan of reasoned judgments. 250 years ago, Lord Mansfield, one of the greatest English Chief Justices, was asked for advice by an army officer, who had been appointed governor of an island in the West Indies, where he would have to administer justice. Lord Mansfield said: “Decide promptly, but never give any reasons. Your decisions may be right, but your reasons are sure to be wrong”.)

34. Another, at least equally important ground for requiring judges to give reasons is that the ultimate decision can be tested. We can interrogate and

test human reasoning up to a point, and we do. But, unlike humans, machines' decision-making can be rigorously and reliably tested without violating their dignity, and we can get accurate responses when we do so.

35. However, people generally trust and are used to human decisions; that is not true of machine decisions, by which, as I have mentioned, most people would be very uncomfortable to be conclusively categorised, assessed or judged.

36. But humans are highly adaptable, so I suspect that people could get used to machine-made decisions. Even assuming that it could be shown that machine-justice is more reliable, faster and cheaper than human justice, I have little doubt but that to start with people would require the ultimate decision-maker to be a human. Initially, perhaps the machine's view would be advisory, and decision would be that of the human judge. If that was to prove satisfactory, the next stage might be that the decision-making would be joint, so that for instance, the role of the human judge would be to check the opinion of the machine. And, finally, if it became clear that the human judge was no more than a rubber stamp, machine judges might be allowed out into the decision-making world on their own.

37. It might appear at first sight that machine judges will be better than human judges because they will be consistent in their approach. One of the consistent criticisms of litigation and arbitration is that one judge, arbitrator or jury will reach a different decision on the same set of facts and issues as a different judges, arbitrator or jury would have reached, because of differences of view as to how to assess competing versions of events or reliability of witnesses, or how to analyse, apply or develop the law. Such differences are said to risk bringing the law into disrepute, both as a matter

of principle and because it makes outcomes more uncertain and therefore legal advice more unreliable or more equivocal.

38. It will be possible to programme judge machines so that they all have the same approach to such questions – or maybe there will simply be one judicial AI machine which judges all cases, and it will presumably be consistent. But will we find that our belief in diversity, which I briefly mentioned earlier, means that what we currently perceive as a disadvantage is actually something we value? I suspect that the answer may be yes – at least initially. If so, this may defer or even undermine the advent of machine justice, or of course diversity may somehow be accommodated in the machines' (or machine's) programming in some ingenious way.

39. That particular thought only occurred to me rather recently, and it is the sort of point which underlines how unsafe it is to predict the future of AI taking over judicial and quasi-judicial functions. Unanticipated human issues, as well as unexpected technical issues, are bound to arise. However, even taking into account that such issues will arise, it seems to me that there is a substantial possibility of human judicial and arbitral functions being replaced by machines.

40. So far I have been concentrating about dispute resolvers, judges and arbitrators, but what about lawyers who advise and represent clients? Many of those who predict which jobs are most and least at risk from AI suggest that judges are more at risk than practising lawyers. I expect that this is because clients often provide their lawyers with a mass of irrelevant information in unstructured form which omits important facts, and the lawyer then has to do a massive editing and investigating job. By contrast, at least in many cases, judges are presented with a well-structured argument, with the relevant facts included and irrelevant facts excluded –

although many may say that betrays a very rosy-tinted description of what goes on in litigation and arbitration.

41. But the basic point is a good one. And it explains why AI is often seen as a more immediate threat to judges and arbitrators than to practising lawyers. And even with AI judges, the outcome of a case may depend at least in some cases on what facts and what arguments are presented and how they are presented, which suggests a continuing role for lawyers. However, in due course, I suspect that it may turn out that machines are as good at selecting and presenting facts and arguments as they are at resolving or deciding issues. Indeed, one future possibility, or even likelihood, is what some call Agentic AI, which involves autonomous AI systems which can communicate with each other and reach decisions on disputes or answers to problems without any human involvement. Or, it may turn out that a single machine can identify all the facts and arguments for itself and then resolve them.

42. Having considered the future for judges, arbitrators and litigation lawyers in the shadow of AI, it is right to consider mediators. I would have thought that they were rather more secure than judges or arbitrators, or even litigation lawyers. Because mediation is more flexible, less rule-bound, and more instinctive than litigation or arbitration, it would, I think, be less susceptible to being subject to machine learning with reliable output. I would predict therefore that the life expectancy of human-chaired mediation is significantly greater than the life expectancy of human-chaired litigation or arbitration. However, consistent with my comments about the brain and AI, I think the time will come when machine-chaired mediation will happen. At that stage, I suspect mediation and arbitration/litigation may well merge, as machines will be able to decide



the best outcome for the parties, and that may be a classic dispute resolution solution or a mediation solution. For the same reason, mediation may serve to prolong the life of litigation lawyers.

43. A final category of people I should mention is expert witnesses. The pressure which the legal system puts on them is heavy: resolving the conflict between duty to the court and loyalty to the paying client has an easy answer in theory, but it is self-evidently hard in practice. AI may make it easier for experts to explain to their clients why they are taking a line which does not suit their clients: indeed I suspect that the internet is already helping on that score. More generally, I suspect that AI may lead to the possibility in at least some cases of machines replacing humans. A judge or arbitrator may well think that a machine's answer to a problem will not be influenced by loyalty to a client. Or the issue may simply be referred to a single machine. But the airy assumption that a machine will not be biased is called into question by the US experience with algorithms used for sentencing to which I have referred. In the long run, however, I fear that AI will mark the end of expert witnesses just as much as of anyone else concerned with litigation.

44. If, as I fear may well happen, machines will one day take over the roles of lawyers and judges, arbitrators and mediators, what jobs will be left for humans? It may be that the arts, literature, poetry, painting, sculpture, composing, theatre, concerts, and the like will still be provided by humans, but I am reluctantly doubtful about that. Computers are already producing works of art and writing which are attractive, and I am doubtful whether many people will want to read a second-rate book or listen to a mediocre piece of music because it was written by a human rather than a first rate product simply because it was made by a machine. On the other hand when

it comes to sport and athletics, it is hard to believe that a contest between machines will be of nearly as much interest as one between humans. So, maybe we should be encouraging our children and grandchildren to excel at sports rather than intellectually. Such a conclusion is supported by the fact that, as I understand it, human competitive chess and Go continue unabated, despite the triumph of machines in those fields: the use of AI has merely introduced a new way of cheating.

45. Anyone tempted to make gloomy prediction about AI replacing humans should bear in mind the Luddites, textile workers 200 years ago, who went around smashing new machinery as they feared it would displace their jobs and render them redundant. They failed in their fight to stop progress, and their fears for their jobs proved well-justified. But although they lost their jobs, the very technological developments which caused their unemployment led to the creation of new, previously unheard of jobs, and in the end the Luddites and their descendants were generally better off than they would have been if they had retained their old jobs. Maybe we will find that is what happens to the lawyers who lose their jobs to AI. And there again maybe not.

46. But that observation can be said to involve speculating about AI beyond my brief which is limited to dispute resolution. And more specifically, it can also be said with justification that I have been looking at the speculative future rather than focussing on the present and the more immediate future. And for the final part of this talk I will spool back a little and look at the present position.

47. I do not want to dwell on specific details relating to dispute resolution, such as disclosure or production of documents. Such issues are important not merely in themselves, but also because they show how the advent of IT

both creates problems (the plethora of documents) and solves them (the use of search terms being an obvious example). They are also significant because they show how adaptable we humans, and in particular we lawyers, are, and have to be. However, such topics have been the subject of talks by senior judges and experienced lawyers, and, even if I was fully up-to-date, I would not be able to add anything useful to what they say. I would prefer to end by looking at dispute resolution in the shorter term more broadly.

48. Because of the high and growing cost of litigation, the pressure to find cheaper ways of dealing with smaller disputes has been increasing. Various reforms since the late 1990s have sought to address this problem in the courts across the UK, with some, if limited, success. On-line dispute resolution through online judges or online mediators can be said to represent a “quick and dirty” solution. But, given the level of legal costs, we currently have no alternative to facing a stark choice between disproportionately expensive litigation or quick and dirty resolution. I would unhesitatingly go for the latter. If it costs £20,000 to bring or defend a claim for £5,000, that is simply a denial of justice: even a moderately well-off person would normally give up, and if the claim was fought, the costs consequences would mean that even winners would walk away feeling that they had not had justice.

49. If we can maximise the use of AI in resolving such disputes that would be a real step forward – provided that it was acceptable to the public. Justice being seen to be done is currently at least part of the rule of law, so public acceptability is vital. A point which it is uncomfortable to make is that, as many of these small cases turn on whose evidence the court believes, the outcome is anyway very often something of a gamble. As I have already

mentioned, even experienced and sensible judges often find it very hard to decide which witness is telling the truth, and that is particularly true in small cases, where there is very often no contemporary documents - notes of meetings, email negotiations, draft agreements, which help to show which party's version of events is correct. Even judges who like to believe that they can identify a liar (which I don't think they can) have to face the problem that most witnesses who aren't telling the truth are not liars, but have either misremembered or have persuaded themselves that what they want to believe is what happened.

50. So having machines, even if they are imperfect, resolve small cases quickly and cheaply seems to me to be a very attractive possibility. To begin with, it will of course involve reduced human input, in the form of judicial or quasi-judicial human supervision and management, but if that involves a reduction of human input, that alone is worth doing. But in due course, the human involvement in the decision-making process may be reduced to almost nothing or very little.

51. How it will work on the lawyers' side remains to be seen. But I suspect that we can leave it to the market. There are already apps which help parties involved or contemplating involvement in litigation, and they will develop in accordance with IT developments and market demand.

52. On one view at least, this will not apply to more substantial cases, where the costs of the current way of litigating are not disproportionate. However, a problem which may arise if there is such a two-tier system is how younger lawyers will be trained. That may not be a problem for solicitors' firms, where young lawyers can be trained up through having junior roles on big cases, but for young barristers, a dearth of small cases in court may deprive them of essential experience. A partial answer to this may be found in

recent pronouncements by the President of the Supreme Court and the Lady Chief Justice, strongly encouraging leading counsel in substantial cases to give their juniors the opportunity of dealing with a witness or two, and/or making submissions on a point.

53. It may well be true that, to those people who are living through it, any time since 1750 or so has felt like an era of transition, but, with the introduction of IT into our lives half a century ago, and the advent of AI now, we are in a particularly acute time of change. As lawyers, especially lawyers in the UK which rightly prides itself as the global dispute resolution centre, we should be particularly astute to ensure that we make, and are seen to make, appropriate use of IT generally and AI in particular. Despite my gloomy, if tentative, prediction that AI may bring about the death of legal practice as we know it in the long term, we owe it to ourselves, our profession, our clients and potential clients and indeed to the rule of law, to make the best use that we can of AI. Given its fast rate of development, that is quite a challenge, but it is one which I am sure that judges, arbitrators, mediators, dispute resolution lawyers and expert witnesses in this country will meet with determination and success – and, in the hope, which I share, that my gloomy prediction is wrong.

David Neuberger

March 2025