The Concurrences 7th Innovation Economics Conference for Antitrust Lawyers

Bush House, 13 June 2023

ARTIFICIAL INTELLIGENCE: THE CONFLUENCE OF TECHNOLOGY, LAW AND ECONOMICS

Sir Marcus Smith[1]

President of the Competition Appeal Tribunal

A. INTRODUCTION

1. These days there is a proliferation of books and writing relating to artificial intelligence ("AI"). A significant subset of this writing relates to the control of AI, which inevitably involves making pre-suppositions about what AI will be and how it will manifest itself. Much of this writing involves "doomsday" scenarios where the risk of a "general" form of AI taking over and displacing humanity is seriously discussed (as well it might be). A subset of the work in this area seeks to evaluate this risk and consider how AI might be given ethics or morality so as to avoid the "doomsday" scenario.

2. Two broad points can be made about this literature. First, in many cases the discussion proceeds at a high level of generality or abstraction. Secondly, the level of input from lawyers and economists is extremely low.

3. There is, I anticipate, a link between these two broad points. It is difficult to consider questions of regulation and liability without understanding clearly that which is being regulated or causing the liability. Any attempt to consider the legal aspects of AI in the abstract will be doomed either to fail or else operate in the more abstract realms of jurisprudence. That failure means that the economic potential of AI is lost because law shapes markets which shape our economic lives. If AI is to be economically significant, it needs to be sufficiently embedded in law in our society.

4. So I am going to seek to avoid these dangers of abstraction by considering the not insignificant legal and economic implications of a specific form of AI that is *(i)* reasonably easy to envisage with a high order of specificity, and *(ii)* of significant social importance. The form of AI that I propose to consider will be automated vehicles. Automated vehicles are already in development, and their arrival on the roads in non-experimental form is viewed as inevitable and fairly imminent. Indeed, there is already legislation on the statute book – the Automated and Electric Vehicles Act 2018 – that anticipates the arrival of automated vehicles. So this is an excellent instance of AI that is

[1] I am speaking in a personal capacity, and all views are my own; as is any responsibility for errors and infelicities.

likely to occur in the near future, which will clearly have significant implications and is one that we can all readily visualise.

5. It is obvious that – just as with vehicles under personal control – there will be an enormous amount of regulation regarding vehicle specification and safety standards; insurance and taxation; and driving standards. I am obviously not going to anticipate such detail. Rather, I am going to try to consider, and shed light upon, the likely future shape of such regulation and its implications.

6. Before I get onto this, it is necessary – at least in broad-brush terms – to sketch out what driverless cars actually entail in terms of nature and operation. The picture rapidly becomes quite complex, with very interesting questions arising.

B. DRIVERLESS CARS: NATURE AND OPERATION

(1) Driven cars

7. As matters stand, the liability regime regarding cars driven by human beings is relatively easy to understand. Applying a "broad-brush":

   (1) In terms of third party liability, the driver causing the injury through their negligent act or omission is liable in tort, and the driver's ability to pay such claims is underpinned by third party liability insurance that the driver is obliged to obtain.

   (2) Absent negligence, the third party has no claim, unless they could claim against the manufacturer of the car for a defective product, either under the tort of negligence or the Consumer Protection Act. Such product liability claims are rare.

   (3) If the driver injures themselves, then the driver will be covered by first party insurance (if taken out).

8. Essentially, therefore, there are generally speaking only two classes of person primarily relevant to the liability regime in the case of third parties: the operator of the vehicle – the driver – and the manufacturer or repairer of the vehicle. I disregard parties with secondary liability, like insurers.

9. Let us now turn to the likely or potential nature and operation of driverless cars.

(2) Driverless cars, but not exclusively so

10. *Ex hypothesi*, the vehicle will be driverless, and will (through a combination of hardware and software) be able to drive itself. However, I am assuming that – at least in the early phases of introduction of driverless cars – their use will not be mandatory. It will still be possible to drive one's own car, still more cycle or walk. Driverless cars are not going to eliminate the pool of potential claimants in the case of accidents.

11. It is unlikely that the services of driverless cars would cease simply at driving. No doubt such vehicles would take themselves to be serviced; and fill-up – most likely with electricity – without external prompting; and self-report safety issues. They would no doubt also deal with the administrative aspects of their regulation: ensuring proper registration, and complying with whatever other requirements the law imposed.

(3) Networked

12. Although driverless cars will have individual hard- and software enabling driverless operation, it seems to me inconceivable that such vehicles will not be networked. The benefits of this are, quite simply, too great not to be realised. Thus, for instance, better routes can be selected if driverless vehicles communicate about traffic flows. Traffic flows can be improved at junctions so as to enable driverless cars to cross junctions without necessarily slowing down. Cars ahead could communicate dangers to cars behind. Cars could form convoys, allowing simultaneous braking and acceleration. In the longer term, information would be pooled regarding accidents, how to avoid them, and how to improve. One of the critical aspects of AI learning is the ability to learn through repetition. That, as it seems to me, is likely to be critical in terms of how AI delivers transport.

13. Thus, there will be two (integrated) levels of control of the driverless car. The intelligence in charge of the individual unit; and the network intelligence. Quite how the relationship would work – in terms of whether the driverless vehicle would effectively be the "slave" to the network or whether the network would simply be an information provider to a largely autonomous driverless car – is something that will very likely be technology dependent, and may vary from case to case. For the present, it is simply necessary to note the distinction between the "unit intelligence" and the "network intelligence".

14. It would probably be an error to assume a single network intelligence. Certainly, in competition terms, that would be problematic, and it may be the central legal problem that needs to be addressed, for the trend to a single platform will – I venture to suggest – be strong. Depending on how driverless cars evolve, it is quite possible that driverless cars would evolve initially on a city-wide basis, and then expand. With expansion would come competition between networks, with unit intelligences operating in the same area(s), but part of different networks. But the trend – driven by the network effect – would be "bigger is better". That is a very real problem for the competition lawyer.

B. IMPLICATIONS OF THIS "BRAVE NEW WORLD"

15. I do not have very much time, having set the stage, and so I am going to limit myself to a few points. I am going to start with the easy – what I regard as the almost inevitable – implications, ramping up to the more controversial and speculative.

(1) The "ownership" model will die

16. Ownership of driver vehicles is, no doubt, diffuse, but the model (for private persons) is essentially one of "owning" your car, even if only through a financing arrangement. Such an "ownership" model might persist for a time, but will be replaced by alternative arrangements, which will be contractual in nature. It is perfectly possible to imagine a range of services on offer to those in need of transport, ranging from the driverless taxi to the driverless car provided instantly on notice and available full-time. Having your car outside your home might cease: one can easily imagine a system where the driverless car is summoned at the press of a button on the user's mobile application.

(2) Negligence liability will end

17. The basic model for road traffic liability, these days, is negligence liability plus insurance for one's own risks, including liability risks. The one thing that we can say about our brave new world is that if – when – there is an injury, it will be a deliberate injury. Because computers don't "do" negligence.

18. Take the case where I am being driven to work in my driverless car, at vast speed because of the network, when someone decides they want to end it all by throwing themselves in front of my vehicle. I am anticipating a series of near instantaneous calculations of relative risk by the unit and network intelligences. "If I hit the pedestrian, will they die?" "If swerve to avoid, and crash into the shop to the left, what injuries will be sustained by my passenger and by third parties in the vicinity? What will the likely economic cost be?"

19. Who knows what the answer will be? It will depend on the algorithms. All I will say is that the outcome will a deliberate one, such that if the vehicle runs the pedestrian over, that will be a chosen outcome. And I don't mean choice by the suicide: I mean by the car.

20. Negligence liability is going to be confined to the outer reaches of relevance (e.g. an error in design by human developers), and the real question will be whether we have some scheme – built out of the network operation – of no-fault injury compensation of the sort contemplated many years ago by Professor Atiyah in his masterly *Accidents, Compensation and the Law*.

(3) Command and control: autonomy and agency

21. I am now moving – building on what I have said – to the more speculative and controversial.

22. I do not believe that the term "artificial intelligence" is a particularly meaningful one. Computers are never going to have – or even understand – our values, even if they can proxy them or imitate to a high level. It may very well be that AI can paint a very respectable version of something that van Gogh could have, but did not himself, paint, basing itself on the corpus of works that van Gogh did paint. But the artistic genius, what drove van Gogh to paint, what differentiated him from other impressionists, and indeed what differentiated impressionists from the schools that preceded them – that creativity does not (I think) exist in the machine in that form.

23. That is because machines do not sense in the way we do. Shylock was absolutely right, speaking of his case, when he asked "If you prick us, do we not bleed?" Bleed is precisely what machines do not do. They do not feel pain. Jealousy, and rage, and love and respect – these are all closed books to the machine. If it looks like a machine is experiencing these things, it is only because it is playing the imitation game very well.

24. I can put it less controversially, if you like: machine intelligence may be intelligence of a sort, but it is so different in quality from human intelligence that it requires a different term of art. Maybe that is what "artificial intelligence" once meant: but I don't think that is how we use the term now. I think we very lazily use the term to consider a human intelligence in the form of nuts and bolts, or silicone and software. It is better to confine

the term intelligence to humans, because at least we then know, or have an approximation to knowing, what we are talking about.

25. There is also an important converse point:

(1) Even using the term AI as it is used now, "human intelligence outside a human being", I do not consider my "driverless network" to be an artificially intelligent network.

(2) It is too limited to amount to a general intelligence, and too functionally constrained – in my judgement – to be intelligent at all.

(3) You can call it "AI" if you like. What I am saying – AI or not – is that there is a regulatory need and significant economic potential and/or disruption in even this limited form of technology. It needs to be harnessed even if it is not "intelligent".

26. My suggestion is that the label is pointless, and perhaps worse than that – distracting. But, since it is common, and convenient, I will continue to use it as shorthand for machines that do clever things and which may look, if we anthropomorphise them, intelligent.

27. The relevant distinction, in my judgement, is between autonomy or agency, and being someone else's tool. There will be many cases where extremely sophisticated tools are used by someone to deliver a good or service. I may use a legal database to find relevant law. You – if you are extremely unwise – may use Chat GBT to craft a passage in your skeleton argument or an article you are writing. The point is that in these cases, the machine is not responsible. It is a tool that is used by someone who is responsible. It – the machine – has no agency. It is not autonomous. We will have AI in a meaningful sense when it is autonomous, when it operates without the kind of human control that exists in the present day.

28. So far, I have been assuming – and I am sure that you all have too – that the owner and controller of our driverless network will be a corporation, controlled by some Steve Jobs type figure of vaulting ambition and technical genius. But must that be so? This is where the lawyer – and I will come to the economist in a moment – can really add value.

29. Is it worth exploring the idea of a new kind of legal person, other than the corporation, that can exist and operate autonomously from the natural person? There are two questions here:

(1) First, it is possible?

(2) Secondly, is it desirable or worthwhile?

(4) Autonomy: is it possible?

30. The point about human beings is that they are capable of being regulated through rules, and our system of governance and regulation rests on this basic fact. To take HLA Hart's point, an alien observing conduct at traffic lights would see the causality. A red light, the traffic stops. A green light, and it proceeds. But the alien would not necessarily understand the why of the causality, which is not that there is some barrier preventing the traffic, merely the respect for the law. Our alien would have to be remarkably

embedded in our psychology to understand that in many cases, we choose to follow the law; and our choice to do so is informed by many factors, including as to the risk of non-compliance and its consequences.

31. I doubt very much whether computers get this. Maybe I am doing neural networks a disservice, and I am wrong about this. But I wonder if machines actually can differentiate between that which they can do (drive through a red light) and that which they ought to do (stop and allow other traffic to pass).

32. The first – rather mundane and practical question – is whether one can, in fact, encode legal norms so that they are not built into the machine's architecture – they do not describe the machine's operation – but are external and referred to by the machine as binding.

33. The point of externality is that third parties – by which I mean the legislature – can then legislate for machines generally or for machines in a specific sector. In short, without human intervention at the rule-obeying level, we have the rule of law for the machine. The machine is capable of autonomous operation, but subject to law. There is no need to control the corporation controlling the machine (that scenario is the machine as a tool, and not autonomous); rather, the law controls the machine as an autonomous subject of the law, with minimal human intervention.

34. Now, one might say that corporations themselves are autonomous subjects of the law, able to contract (for example) without human intervention. On the purely superficial level, you might be right. There are contracts between corporations, and people and corporations, all the time. But on the fundamental level, you would be very wrong. Corporations have shareholders, and directors, and debtors. They operate, largely, at human direction, and are really a very sophisticated tool for human interaction. They just happen to have legal personality. Coase's theory was that the market economy has two modes of organisation: *(i)* the market; and *(ii)* the firm. The corporation is simply a manifestation of the latter case.

35. In other words, the chain of control in a corporation moves in stages: *(i)* corporations can generate legal effects in their own right, and are, in their own right, subject to law; but *(ii)* they are controlled and directed by human beings who *(iii)* are themselves subject (directly) to a legal framework of corporate law (e.g. directors duties, etc). When I speak of autonomy, what I am aiming at is a cutting out of the middleman, the human controller sitting between the corporation and the law. What I am envisaging is the "AI" corporation acting autonomously and directly in accordance with the rule of law but without significant or frequent human guidance or intervention.

36. So there is a technical question as to whether legal norms can be encoded so as to guide the machine. The second, rather more philosophical question, is whether even if the law could be encoded in this way, it would comprehensible to the machine. Obeying laws is actually quite a nuanced thing. If I, as a car-owner, have to MOT my car on the 1 June, I'll probably only start thinking about complying with this obligation in the course of May. If something stops me getting an MOT in time, I will deal with that eventuality. The machine, I suspect, will compute the risks, and see that there is a finite and calculable risk that even on 2 June of the previous year, when the car has been successfully MOT'd the day before on 1 June, that some circumstances may prevail to prevent the car returning for next year's MOT in a year less one day's time. I think there is a real risk

that the machine will not allow the freshly MOT'd car to leave the garage until next year's MOT is successfully completed. And so on…for next year's MOT never comes.

37. It's a slightly absurd example, but I do not think that we can take for granted that following laws is as easy as we humans think it to be. I think there is a literality to machines that may be problematic.

38. So we may have an awful lot of thinking to do about whether autonomy is even possible. But let's assume that it is. Is it desirable? Why bother at all with autonomy?

(5) Autonomy: is it desirable?

39. I'm sticking with my driverless car network. I am stressing that it is a network and networks trend to monopoly. It is easy to see why. Three reasons, at least.

40. First reason. Networks often require common standards: think railways; telecommunications; electricity distribution; payment systems. Standards imply a degree of size and investment. Separate networks operating to common standards are easy to combine (with economies of scale to the producer).

41. Second reason. Even when common standards are not required, the consumer tends to desire them, because it is more convenient and the consumer gets more choice: think insurance provision through price comparison websites; streaming of movies and television and music and books. Whenever I watch television, my competition law self watches on with horror as my couch potato self rages at the difficulty of finding the programme I want because it's not on Apple TV, Amazon, Netflix but some upstart company looking to compete. My couch potato self is the reverse of charitable when I am looking for this programme. That is not good, given that we like contested markets and "upstart" competitors.

42. Thirdly, costs and pricing. We all know – because our economics professors taught us so – that the entrepreneur will produce until Marginal Cost MC equals Marginal Revenue MR. MC = MR, one of the staples of micro-economics, and like many staples more or less totally useless when it comes to the real world.

43. The reason? Whilst MC = MR may be true, price does not equal MR minus MC, even to the most rabid competition lawyer. Why? Because all costs are not marginal costs. The entrepreneur must cover these at a minimum – but what about all the sunk costs in establishing the network? The 10,000 software developers it took to write the code for our unit and networked intelligences in our driverless cars? The 10,000 lawyers it took to get regulatory approval for the project? These costs will be incurred before a penny of revenue is seen, and have to be recovered.

44. So, the cost of market entry in the case of networks is typically high, so there will be few competitors and the market – to use Baumol's phrase – will not be contestable.

45. We therefore trend to monopoly, and the good old Chapter II prohibition of abuse of dominance for excessive pricing might – or might not – ride to the rescue. It might not ride to the rescue for many, many reasons, but one will be because cost – and therefore what is an "excessive price" – can be a remarkably difficult concept. Sticking, still, with our networked driverless cars, let's think why:

(1) Let's suppose the costs directly attributable to setting up a successful network in Cambridge are £1 billion. These costs are going to have to be recovered from the consumers, but the first consumer is not going to be hit with a bill of £1 billion. The costs to be recovered will be spread, and spread in the form of a fairly s/ophisticated pricing structure. Immediately, we are going to have some difficulty in tracking when the entrepreneur has recovered the £1 billion and has moved – looking only at these direct costs – from the red into the black.

(2) But what about failed costs? Suppose an unsuccessful attempt to establish (by our entrepreneur) a driverless network in London, which failed at great cost because London was too big? How are these costs to be recovered? What about a return on capital? £1 billion does not grow on trees. What about our entrepreneur's desire to expand into nearby Letchworth Garden City? This will involve significant additional direct costs, and no doubt our entrepreneur will keep the same pricing system throughout, raising interesting questions of subsidisation and – potentially – margin squeeze.

46. My point is that a very significant gap between marginal cost and price can be justified, but whether the price is just or truly justifiable is very difficult to assess. And so, the profit motive – the desire to maximise profits – becomes a problem, because it becomes very hard to constrain.

47. Is it here that AI brings an advantage? Let us suppose that the technical problems I have mentioned <u>can</u> be overcome, and (at some point in time) the entrepreneur has to cede control of their undertaking to the AI, taking a defined return (calculated by the AI) going forwards. The AI – no profit motive there – will run the system at cost, the return to the entrepreneur being a defined cost to the network that the AI will calculate. This is not unlike agreements to build bridges, where the builder who funds the project is entitled to charge a levy for a certain number of years at certain defined rates.

48. In the past, we have (as a nation) debated the relative virtues of nationalisation and privatisation, and have ended up with a regulatory compromise. Suppose the AI corporation is one better than this compromise, in that is acts in a competitive market without the profit motive. The implications of this clearly have to be considered, and all I am doing is posing the question.