

RESPONSE OF THE FACULTY OF ADVOCATES

to the



Joint Preliminary Consultation Paper of the Law Commission and Scottish Law Commission on Automated Vehicles

Introduction – Scope, assumptions and terminology

1. We note that underlying the Joint Preliminary Consultation Paper is the SAE analysis of six levels of vehicle automation. We agree that this provides a proper basis for the analysis in the Paper, in particular because the focus of the SAE analysis is on how, functionally, the automated vehicles operate, rather than upon the technological means which cause them so to function. However, in order to give clarity to the discussion, it is helpful first to consider some issues of a technical nature.
2. “Intelligent” automated control systems have in common that they require the use of computers which gather data, including input from sensors, process those data and then perform actions resulting from that processing. In terms of function, such systems range all the way from the operation of warning systems where those are present at level 0 to the full operation of the vehicle with no input required from the user in any circumstances at level 5. Such control systems are presently usually based upon algorithms. An algorithm is a process or set of rules to be followed in calculations or other problem-solving operations. Such algorithms are represented by a structured series of logical steps each capable of being answered “yes” or “no”, the structure of which can be visually represented by means of a flow chart. For example, cruise control in a vehicle will involve the algorithmic processing of inputs from speed measuring devices and an output which automatically adjusts the speed of the vehicle. It is possible, and, insofar as level 4 and level 5 systems presently exist, usual for the automatic processes to be based upon algorithms. However, the thrust of research in the industry is to develop systems which,

rather than (as at present) executing pre-programmed processes, instead make their own autonomous “decisions”.

3. The latter system uses processes which enable a computer to “learn” from itself and make decisions which are autonomous. Such systems are, in effect, neural networks. For example, the Google Go program (like earlier chess-playing programs), is programmed in the first instance by human programmers in such a way as to enable the system to be “educated” by feeding it with a dataset (in this case a dataset of previous actual Go games) from which it is able internally to derive appropriate responses to moves made during a game of Go by its opponent. The larger and the more reliable the dataset, the better the outcome. The next generation of such systems is typified by Google Go Zero, which instead of being educated by means of a dataset, was programmed with the basic rules of the game of Go, and then was set to play itself until it became adept at playing the game. It is, in effect, a self-educating system. Modern cancer diagnosis software is likewise trained on datasets and from its own experience. The result is that it is significantly better at detecting early-onset cancers than humans are. However, it is a feature of such systems that their internal “reasoning” processes tend to be opaque and impenetrable (what is known as the “black box” phenomenon) – the programmers may be unable to explain how they achieve their outcomes. With conventional software, on the other hand, it is always possible to explain the algorithms and examine the source code: errors ought to be capable of detection. Classical AI follows a precise step of logical rules (algorithms) whereas the behaviour of neural networks may only be described statistically (stochastical behaviour). The ability to describe such a trained behaviour and even reverse engineer it (i.e. explain which input caused the trained behaviour) is an area of intense research. With level 4 or 5 automation, control of the vehicle could be based upon algorithms, though research is seeking to employ control systems which make use of neural networks.
4. Further, the system driving an automated vehicle may not be entirely self-contained. We note that the Consultation Paper appears to assume that the Automated Driving system will consist of onboard sensors which will gather data which is then relayed to an on-board computer which will make the driving “decisions”, and which will then control the driving of the vehicle. However, it seems likely that such vehicles may also require (or benefit from) data received from other systems, particularly those of other vehicles on the road and from items within the road infrastructure itself. For instance, communication between

vehicles may allow automated driving systems in a number of vehicles to cooperate with each other or with traffic control systems to optimise traffic flow and minimise risk of accidents. For instance, net average speeds on motorways may improve significantly if vehicles some distance from an accident or obstruction modify speeds and merge lanes. In such cases, the ability of vehicles to interoperate appropriately with each other and any systems forming part of the road infrastructure will form a vital part of the authorisation process. For such interoperation, network connections will be required. If such systems are in place, a failure of the systems to interoperate might be considered equivalent to a human driver failing to comply with traffic signs or driving without due care and attention. Consideration will have to be given to how such breaches, which are only possible by automated vehicles and are not generally covered by offences applicable to human drivers, are regulated, detected, investigated and (if appropriate) prosecuted.

5. Furthermore, the tasks requiring to be performed by an autonomous vehicle are highly complex and models of automated driving system have been proposed where the processing of the data and the output of the control commands occur not in an on-board computer, but in a remote server, connected by a mobile data signal to the vehicle in real time. In these circumstances, though the SAE levels are extremely helpful in focussing discussion, it requires to be borne in mind that the same apparent level of functionality may be obtained in two different vehicles, the one using conventional algorithms and the other through means of a neural network which may either be on board the vehicle (which is therefore wholly autonomous) or which may be resident in a remote server requiring constant connectivity. This would then raise the question of whether the regulatory framework needs to recognise the possibility of malfunction where connectivity cannot be maintained, even though with an ADS based upon algorithms, and also with a self-contained onboard neural network, connectivity is not required.
6. In describing these various systems, we do not consider it helpful to use the expression “artificial intelligence” to describe them. The term is imprecise. It is defined in the Oxford English Dictionary not as a process, but as a field of study:

“the study of how to produce machines that have some of the qualities that the human mind has, such as the ability to understand language, recognize pictures, solve problems, and learn.”

ISO Standard 2382:2015 defines it as:

“the capability of a functional unit to perform functions that are generally associated with human intelligence such as reasoning and learning.”

However, the ISO standard does not add anything to an understanding of the issues and it should be borne in mind that in common parlance “artificial intelligence” can end up being used very loosely indeed, describing, in sales literature, almost any kind of computer system. We therefore welcome the avoidance of the use of that term in the Joint Consultation Paper, though we do have some further comments below on the use of that expression in Chapter 9 of the Paper.

7. We note also the terms of section 8(1)(a) of the Automated and Electric Vehicles Act 2018: “a vehicle is 'driving itself' if it is operating in a mode in which it is not being controlled, and does not need to be monitored, by an individual.” The reference to “mode” is helpful as, of course, although a vehicle may have type approval as, say, a level 5 automated vehicle, yet be being operated manually.
8. Against that background, we should make the following observations as to the scope of the consultation as outlined in Chapter 1 of the Joint Consultation Paper:
 - (a) Level 0 to 3 vehicles are not considered to be automated vehicles, and so fall outside the scope of the consultation.
 - (b) A level 5 vehicle is, in our view, plainly an automated vehicle when it is being operated as such. A level 4 vehicle is, in our view, when being operated as such within its ODD, in effect a level 5 vehicle. When it is not in its ODD, it ceases to be such and reverts to level 3 or below (depending on how it is designed to function outside the ODD). Conceptually, the distinction between level 4 and level 5 vehicles relates not to the level of automation, but rather to the design domain within which that functionality is available.
 - (c) A level 3 vehicle (which the paper refers to as being in the “mushy middle”) is, nonetheless, an automated vehicle when it is operating in that mode. The conditionality relates not to a physical ODD (as with level 4) but to operating conditions with which it may be beyond the capacity of the automated system to cope. The scope of the ODD of a level 4 vehicle is clear and predictable (i.e. all

motorways, or certain designated motorways), but in the case of the level 3 vehicle, may be uncertain and unpredictable.

(d) Even in the case of highly automated vehicles (level 4 or 5) there is the potential for the vehicles to be operated in a mode which lies anywhere within domains 0 to 3. For example, in those instances where the control system of a level 5 vehicle requires constant connectivity, in the event of connectivity being lost, the vehicle may revert to manual operation or automated operation at a lower level.

(e) A further issue arises in the case of Level 5 vehicles. Path 2 vehicles (see §2.39 of the Paper) might be developed without manual controls (and one might envisage a vehicle which has such controls which, however, might not be readily accessible in normal use). A user of the car would be no more than a passenger in it. However, it cannot be assumed that there would be full convergence between path 1 and Path 2 vehicles as, once either path arrives at level 5 automation, there may, nonetheless, be a market for a level 5 vehicle fitted with manual controls to enable a driver to take over driving for the enjoyment of it.

(f) Furthermore, it may not be appropriate to postpone to a future consultation on Mobility as a Service, issues arising from the operation of Path 2 vehicle without the possibility of manual control. Although a car without manual controls might be used to provide a taxi service, it cannot be ruled out that there may be a market for persons to purchase such vehicles for their own exclusive use. In one sense, the operating system would be providing mobility as a service, but that is in the same way that any level 5 vehicle operating by means of a remotely-accessed neural network could be regarded as providing a service, rather than being a product. The dividing lines are not sharp and clear-cut. There may be a case for sub-dividing level 5 vehicles between level 5 vehicles with the possibility of manual control and level 5 vehicles without that possibility.

(g) Although the Joint Consultation Paper is correct to focus on how automated vehicles actually function, the underlying IT is not entirely irrelevant. For example, the differing functionality amongst levels 3 4 and 5 may (no doubt amongst other factors) reflect the extent to which different models of vehicles might be driven by algorithms or by neural networks – the higher the level of automation, the more difficult it may be to drive the vehicle through the use of conventional algorithms.

This brings problems where neural networks are involved. First, if the operation of the system causes an accident, it might be perfectly possible to determine the cause through examination of the source code of a conventional system (i.e. there might be a clearly identifiable bug in the system, or one of the algorithms might be obviously flawed) but where a neural network is involved, it may be literally impossible to determine what produced the behaviour which caused the accident (the black box problem referred to above). Further, the server at the heart of the system may be located outside UK jurisdiction and/or if a level 5 system permits operation by a remote operator (for example upon failure of the guidance, leaving the vehicle in a dangerous location) then the individual undertaking the remote operation might not be within UK jurisdiction.

9. It is important, when considering the response to each of the questions in the Joint Consultation Paper to keep the foregoing factors in mind.

Q1(1): Do you agree that all vehicles which “drive themselves” within the meaning of the Automated and Electric Vehicles Act 2018 should have a user-in-charge in a position to operate the controls, unless the vehicle is specifically authorised as able to function safely without one?

10. Yes, although we have some reservations about the term “user-in-charge” (see our answer to Q2 below). It seems this question assumes that there can and should be two separate classes of self-driving vehicles: those authorised to operate with a “user-in-charge” and those authorised to operate without a “user-in-charge”. We agree.
11. In our view, a regulatory distinction between these two classes of vehicles reflects the fact that there are likely to be two generalised ways in which these vehicles will operate. One class will be broadly analogous to a privately-owned vehicle which is likely to be bought (or hired) and operated by a particular person (or defined set of persons) where there is a definite relationship of ownership or control between a specific vehicle and specific individuals; and another class which is likely to be broadly analogous to a taxi service and which may be part of a wider Mobility as a Service offering where there is no such relationship between a vehicle and individuals, or which may be purchased or leased by an owner who wishes to use it as a private vehicle, but does not wish, him or herself, to drive. We note that this distinction is reflected in the consultation in the discussion of “Path 1” and “Path 2” automated vehicles.

12. In the first class, the user of the vehicle can reasonably be expected (and may expect to be able) to take some level of responsibility for the way in which the vehicle is operated. The user may wish to drive the vehicle even when it might be capable of operating in a self-driving mode. They may expect to have full control of the vehicle when it is not self-driving. In the second class, the user of the vehicle would not have any expectation of these wider responsibilities. It seems to us necessary that the regulatory regime supports both these classes of vehicle, so the two different levels of authorisation seem appropriate.

Q1(2): Do you agree the user-in-charge: (a) must be qualified and fit to drive; (b) would not be a driver for purposes of civil and criminal law while the automated driving system is engaged; but (c) would assume the responsibilities of a driver after confirming that they are taking over the controls, subject to the exception in (3) below?

13. Yes.

14. In our view, a regulatory regime that includes a “user-in-charge” role assumes that such a person either may become a driver at their own discretion or must do so when required: such as after the vehicle comes to a safe stop condition and ceases to be self-driving (whether due to a failure or on coming to the boundary of its operating domain). If the presence of a “user-in-charge” is part of a self-driving vehicle’s authorisation, then it follows that a person filling this role must be competent to safely drive the vehicle when it is not driving itself. However, as we understand the proposals in the consultation, a “user-in-charge” would only be expected to become a driver after the vehicle brings itself to a safe stop.

15. Obviously, vehicles authorised to operate without a “user-in-charge” would require alternative arrangements for dealing with events after a safe stop. This links in with our comments relating to the role of a vehicle “operator”.

Q1(3): Do you agree that if the user-in-charge takes control to mitigate a risk of accident caused by the automated driving system, the vehicle should still be considered to be driving itself if the user-in-charge fails to prevent the accident?

16. Yes, although this should be implemented by way of a legal presumption rather than as a matter of strict liability.

17. As we understand the consultation, unlike in a SAE Level 3 vehicle with conditional automation and a “fallback driver”, a “user-in-charge” should never need to take control of

the vehicle to mitigate the risk of an accident: the vehicle should mitigate and avoid these risks itself without relying on any human intervention (at least until it has reached a safe stop condition). If a “user-in-charge” finds themselves in a situation where they need to take control otherwise than when in a safe stop condition, the vehicle has already failed to drive itself safely and an intervening driver could not be said to have caused the risk of accident.

18. However, there may be occasions where a “user-in-charge” wrongly takes control of a vehicle, perhaps due to a mistaken perception of risk; or where, having taken control, negligently or recklessly drives in such a way as to increase the risk of accident. In such cases, it should be possible for the presumption of liability on the self-driving system to be displaced. The onus should be on the ADSE to displace the presumption: it is likely that the ADSE will have access to vehicle data to allow these matters to be investigated and proven. In such a situation, of course, once the “user-in-charge” assumes control of the system, they become the driver, and should be treated as such.

Q2: We seek views on whether the label “user-in-charge” conveys its intended meaning.

19. We have some reservations about the term “user-in-charge” but agree that it is one possible way of expressing this concept. In particular, we can see the advantage in using the terminology “in charge” as giving some continuity with the existing legal concept of being “in charge” of a vehicle as a concept wider than being the driver of a vehicle.
20. In our view, there needs to be a clearly defined set of roles within the regulatory regime for automated vehicles. Any person travelling in a vehicle is a user of the vehicle, so that term is overly broad. We would prefer the well-understood terms of “driver” and “passenger” for the extreme cases of the vehicle user who is, at least for the time being, performing the dynamic driving task and a vehicle user who cannot be expected to play any active role in driving the vehicle or controlling any of its driving systems. In between these extremes are the “fallback driver” and “user-in-charge” roles. As we see it, persons performing these roles may become drivers when the advanced driver assistance (SAE level 3) or self-driving system (SAE level 4) are not driving the vehicle. The distinction between “fallback driver” and “user-in-charge” is the circumstance in which they may be required to take control and whether human intervention is relied upon to ensure the vehicle remains safe when any automated system is driving the vehicle.

- 21.** We would tentatively suggest there needs to be a further defined role for which we propose the name “operator”. This would be the person who, for the time being, had legal responsibility for the vehicle other than the execution of the driving task (the responsibility for which lies with the driver (SAE levels 0 to 2), the fallback driver (SAE level 3) or the automatic driving system (SAE levels 4 and 5). The operator would be the person who would, for example, be responsible for appropriately engaging or disengaging the self-driving system; or dealing appropriately with the vehicle after it reaches a safe stop condition, whether that is planned (at the limits of the operating domain) or unplanned (in the event of a system failure or other unsafe event). Although the operator might be present in the vehicle, we discuss above, in paragraph 8(g) circumstances in which the operator of the vehicle may not be in the vehicle, or even on the same continent. Most such remote operators will be employees or agents of the company providing the automatic system, and it would be appropriate to define “operator” so as to include legal as well as natural persons.
- 22.** Where a self-driving vehicle is a Path 1 vehicle, a type analogous to a private car, then there is a clear overlap between the “operator” and the “user-in-charge” roles. Where the self-driving vehicle is a Path 2 vehicle, there will be no “user-in-charge” and relevant duties will require to be imposed on the “operator”. In these cases, the “operator” will not be in the vehicle but can be expected to have some sort of remote connection with it, which may include the ability to remotely operate the vehicle (or, possibly, even to remotely drive it). The operator role would in those circumstances be performed by a remote individual, or by a remote server and the term “operator” should be defined so as to include the natural or legal person responsible for the employment of the remote individual operator or the functioning of the remote server.
- 23.** An illustration of this role is where a self-driving vehicle comes to an unplanned safe stop due to failure, breakdown or departure from the operating domain. In a private car example, the “operator” / “user-in-charge” would be the person responsible for the necessary action (for instance to assume the role of driver or to summon roadside assistance). In a path 2 example, the “passenger” would have no such responsibility; but an “operator” would still be needed to take necessary actions to recover the vehicle.
- 24.** Given the overlap between “user-in-charge” and “operator”, we would tentatively suggest “onboard operator” as an alternative to “user-in-charge”. Vehicles authorised to operate

without a “user-in-charge” would be those authorised to operate with a “remote operator” rather than an “onboard operator”.

Q3: We seek views on whether it should be a criminal offence for a user-in-charge who is subjectively aware of a risk of serious injury to fail to take reasonable steps to avert that risk.

25. We do not think that this should be a criminal offence. As we understand it, the defining feature of a “user-in-charge” (as opposed to a “fallback driver”) is that they are not expected to take control of the vehicle other than after a safe stop. As that is the basis on which the vehicle is authorised, then it would be wrong to attach criminal liability to a failure to do something outside the defined and understood role of a “user-in-charge”. Attaching criminal liability to intervene in a system, which should by definition require no intervention, would also have the undesirable effect of encouraging “users-in-charge” to intervene based on their own perception of risk. This perception may be wrong or (as they are not expected to maintain full awareness) based on an inadequate or flawed understanding of the driving environment. In our view, that may well create risk by leading to increased unsafe human intervention into systems operating safely.
26. If a “user-in-charge’s” behaviour in not intervening is exceptionally blameworthy, it may (in Scotland) fall within the definition of the existing common law offence of culpable and reckless conduct. In our view, this is sufficient to address extreme cases of egregious or deliberate behaviour by “users-in-charge”.

Q4: We seek views on how automated driving systems can operate safely and effectively in the absence of a user-in-charge.

27. It seems to us that this is a technical question rather than a legal one, so we express no firm views on this matter.
28. However, it seems that such vehicles would begin by operating within strictly limited operating domains segregated from other vehicles. These might be exclusive lanes on certain roads. Arguably, the Docklands Light Railway or lifts in buildings are examples of self-driving vehicles operating within very constrained operating domains.
29. The existence of the “operator” role (see Q2) is necessary as, in our view, there must always be some way of bringing a human (or corporate entity) into the overall

responsibility for the safe and effective operation of a vehicle, even if there is no “user-in-charge”.

Q5: Do you agree that powers should be made available to approve automated vehicles as able to operate without a user-in-charge?

30. Yes.

Q6: Under what circumstances should a driver be permitted to undertake secondary activities when an automated driving system is engaged?

31. As we understand it, this question relates to “fallback drivers” in systems operating at SAE level 3 rather than a “user-in-charge” in a SAE level 3 system. We are minded to say that there should be no formal relaxation of the requirements on a “fallback driver” as compared to a “driver”. We think that a clear and simple rule has significant advantages. We agree that it is helpful to important to maintain a clear distinction between driving and being a passenger.
32. We are not convinced that it will be possible to define a certain class of non-driving activities (beyond those presently permitted for drivers) which will enhance a “fallback driver’s” ability to take control on a “request to intervene” (say by preventing the fallback driver from falling asleep) rather than degrade it. Such a request could, presumably, come at any time, including when the vehicle was travelling at high speeds, around corners, or when approaching stationary traffic or a location where another vehicle may have right-of-way. We would prefer that vehicles (as some of today’s vehicles already do) include systems to detect inattention or sleepiness in drivers (or fallback drivers) and alert the driver and/or bring the vehicle to a safe stop condition.
33. However, we accept that much may depend on the circumstances in which a “request to intervene” might occur in a SAE level 3 vehicle. If it were that a “request to intervene” could only occur when the “fallback driver” may have a significant period of time (at least several seconds) to assimilate the situation and prepare to take control safely, then relaxation of the requirement may be possible. However, this seems to us to be a technical and human factor question rather than a legal one.
34. If, however, the Law Commissions were minded to recommend allowing a limited class of activity beyond those presently permitted for drivers of non-automated vehicles (such as listening to the radio, using a hands-free telephone, eating or drinking) as being

permissible for a “user in charge” in a level 3 vehicle, then such additional secondary activities should be limited to engagement with a screen placed such that the human is seated in the driving seat and facing the correct way and that, in the event that human intervention is required, the screen and sound system is programmed so as to command such intervention whilst simultaneously terminating all other activities. This would have the effect of permitting the “user in charge” (or fallback driver) to, for example, attend to emails, browse the internet and watch videos, but with the assurance that such distractions would immediately cease when human intervention is required. Secondary activities at SAE3 other than this (reading physical books or magazines, for example) in which an engrossed human could fail to observe the command signal would not be permissible. SAE3 requires the human to be on standby. A failure to observe the emergency command would bring about criminal responsibility on the part of the human, akin to RTA 1988 s.3.

Q7: Conditionally automated driving systems require a human driver to act as a fallback when the automated driving system is engaged. If such systems are authorised at an international level: (1) should the fallback be permitted to undertake other activities? (2) if so, what should those activities be?

35. See our answer to Q6 above.

Q8: Do you agree that: (1) a new safety assurance scheme should be established to authorise automated driving systems which are installed: (a) as modifications to registered vehicles; or (b) in vehicles manufactured in limited numbers (a "small series")? (2) unauthorised automated driving systems should be prohibited? (3) the safety assurance agency should also have powers to make special vehicle orders for highly automated vehicles, so as to authorise design changes which would otherwise breach construction and use regulations?

36. Yes.

Q9: Do you agree that every automated driving system (ADS) should be backed by an entity (ADSE) which takes responsibility for the safety of the system?

37. Yes.

Q10: We seek views on how far should a new safety assurance system be based on accrediting the developers' own systems, and how far should it involve third party testing.

38. We think it unlikely that a third-party could develop sufficiently comprehensive tests objectively to test the range of likely automated driving systems. The emphasis should be on auditing an ADSE's own testing procedure as part of the approval for ADSEs.
39. The new system should be predominantly based on self-certification. Manufacturers or other entities involved in the process can go through a series of testing, overseen by a government body and with stringent guidelines / incentives to make sure the certification is fair and accurate. In the event of false certification, severe penalties should apply. Third party testing should feature, but it should be for the purpose of monitoring and encouraging compliance by the ADSE and should be random and infrequent.

Q11: We seek views on how the safety assurance scheme could best work with local agencies to ensure that is sensitive to local conditions.

40. Perhaps new departments will need to be set up which have both local knowledge and an understanding of the way that automated systems operate. Such agencies could be council led, as, arguably the best local knowledge lies with local councils. Each local council could appoint one person as a contact point who can filter up local requirements to the relevant government agency. The contact point can also liaise with other council contact points.

Q12: If there is to be a new safety assurance scheme to authorise automated driving systems before they are allowed onto the roads, should the agency also have responsibilities for safety of these systems following deployment? If so, should the organisation have responsibilities for: (1) regulating consumer and marketing materials? (2) market surveillance? (3) roadworthiness tests? We seek views on whether the agency's responsibilities in these three areas should extend to advanced driver assistance systems.

41. We favour the giving of such a role to the agency and are of the view that it should extend to level 2 systems. The present controls are piecemeal, and it would be helpful to have a "one-stop shop" for such matters. However, the main argument for giving this role to the agency is that the agency will have acquired substantial expertise from its other activities and can be expected to use that expertise in relation to these additional matters. Recent experiences with the irresponsible use by drivers of existing systems (such as sitting in the passenger seat) suggest a lack of awareness on the part of drivers of the proper function

and limitation of such systems, and there is a clear need to ensure that information provided to drivers and consumers is clear and accurate.

Q13: Is there a need to provide drivers with additional training on advanced driver assistance systems? If so, can this be met on a voluntary basis, through incentives offered by insurers?

42. The operation of advanced driver assistance systems arguably requires an additional skill set from driving conventional vehicles, centred around awareness of the limitations of such systems and the extent to which human intervention may be required and in which circumstances the need for such intervention might be required. A driver's ability to handle such systems may in part be influenced by their familiarity with and adaptability to cope with IT systems. Therefore, it does, in principle, seem desirable for all those in charge of an automated vehicle to have a basic level of competence in these matters.
43. We did consider whether there should be a formal licensing requirement, with the introduction of a new class of vehicle to be licensed to drive, by which it would be necessary to pass an appropriate test (much as drivers licensed only to drive vehicles with automatic transmission require to do in order to drive vehicles with manual transmission, or car drivers require to drive HGVs, PSVs or motor cycles.) However, we are not convinced that the task of driving an automated vehicle is so fundamentally different as to require such a burden to be placed on both drivers and the administrative resources of the DVLA. On balance, we would prefer that the acquisition of the necessary skills be encouraged through the insurance route. For example, if there were to be some readily available qualification (similar to membership of the Institute of Advanced Motorists) then Insurers could reflect such qualifications in the level of premium, or, if a more interventionist approach were required, even decline cover for driving automated vehicles unless such a qualification were obtained.
44. If the Commissions were minded to recommend a licensing obligation, then it may be that the nature of the skills to be tested would admit of testing online.

Q14: We seek views on how accidents involving driving automation should be investigated. We seek views on whether an Accident Investigation Branch should investigate high profile accidents involving automated vehicles? Alternatively, should specialist expertise be provided to police forces.

45. There may be a number of reasons why an accident may befall automated vehicles. The accident might be caused by human error, for example that the vehicle was in manual mode at the time and the accident was the fault of the driver. The accident might have happened while the vehicle was in automatic driving mode, but be due to fault on the part of the driver, for example, by inputting an inappropriate route. The accident might have happened while the vehicle was in automatic driving mode, but be due to a mechanical malfunction, such as a brake failure. The accident might, indeed, have arisen from the operation of the automated system, and the actual malfunctioning might be ascertainable by examining the relevant program, if written using conventional algorithms, but if based on neural networks may not be ascertainable. Furthermore, a failure of an automated system might require enforcement action against an ADSE, which may or may not involve a product recall and other similar measures.
46. In these circumstances, there should be a hierarchical approach, with initial investigation being carried out by the police (for which purpose the police will doubtless require special training) and the police should be equipped and trained to recognise when an automated system is implicated, but once an automated system has been identified as potentially at fault, then the investigation of that system should be by the agency responsible for the safety assurance scheme, or, at least, involve that agency. It is likely that investigation of the malfunctioning of the system (if capable of ascertainment) will require the expertise which would have been developed by the agency and be available to it. Furthermore, the agency would be better equipped to take administrative and regulatory measures against the ADSE, should such be required.

Q15: (1) Do you agree that the new safety agency should monitor the accident rate of highly automated vehicles which drive themselves, compared with human drivers? (2) We seek views on whether there is also a need to monitor the accident rates of advanced driver assistance systems.

47. Yes to both.

Q16: (1) What are the challenges of comparing the accident rates of automated driving systems with that of human drivers? (2) Are existing sources of data sufficient to allow meaningful comparisons? Alternatively, are new obligations to report accidents needed?

48. Meaningful comparisons may be difficult. For example, every single accident might not be known about so as thereby to allow comparison to take place. Many people might have a minor bump and not feel it is worth reporting. Comparisons might lead to a skewed understanding of safety because simple statistics will not show the causes of accidents.
49. Existing sources of data are not sufficient. The challenges identified above are problems which likely exist in current statistical analysis. New obligations to report accidents seem problematic. How is an accident defined? This can be a subjective concept: what one person perceives to be an accident requiring notification, another might perceive as a minor scratch and not needing to be notified. A requirement to report every single incident, no matter how trivial would impose an undue burden on users, and might be in danger of swamping the system. If a line is to be drawn between notifiable and non-notifiable accidents, it may be difficult to define where that line should be drawn. It should also be borne in mind that the creation of such an obligation would entail a new reporting requirement on all drivers, and not just drivers of automated vehicles.

Q17: We seek views on whether there is a need for further guidance or clarification on Part 1 of Automated and Electric Vehicles Act 2018 in the following areas: (1) Are sections 3(1) and 6(3) on contributory negligence sufficiently clear? (2) Do you agree that the issue of causation can be left to the courts, or is there a need for guidance on the meaning of causation in section 2? (3) Do any potential problems arise from the need to retain data to deal with insurance claims? If so: (a) to make a claim against an automated vehicle's insurer, should the injured person be required to notify the police or the insurer about the alleged incident within a set period, so that data can be preserved? (b) how long should that period be?

50. The existing statutory provisions are reasonably clear. However, they may need to be revisited in the event that there is imposed a requirement for notification to the insurer or police as prerequisite to the retention of data, as discussed below.
51. Causation should be left for the courts. Hard and fast rules can be restrictive and not practical to follow. Leaving causation to the courts will allow for the flexibility of examining facts on a case by case basis. However, we should qualify that answer under reference to the discussion in response to question 18, below.

- 52.** We do not see any issue in principle regarding gathering and retention of data by the ADSE. Such processing is likely to be regarded as lawful under at least paragraph 1 (b), (c) and (f) of article 6 of the GDPR. We do, however, see practical problems with the retention of data, especially where the system has not detected that an accident has occurred. It is one of the requirements of the GDPR that data is not retained for longer than is necessary, but the answer to the question of how long is necessary causes practical problems. An accident might give rise to a claim for damages for personal injuries. Such a claim would be subject to the three year limitation period for personal injuries, but, of course, because of the need for the appropriate actual or constructive knowledge before the period begins to run, and the ability of the Court to permit an action to be brought after the expiry of the period, the need for data to be retained could extend considerably beyond three years from the date of the accident. If the accident caused only property damage, then it would be the five-year prescriptive period which would operate. Again, there may be a delay to the start of the running of the prescriptive period, and claims might remain potentially exigible until the end of the long negative prescriptive period of 20 years. The demands which such a state of affairs would place upon storage capacity are obvious.
- 53.** In these circumstances, there is a practical case for imposing a requirement for notification within a limited period, but we have concerns that this may operate unfairly; for example, in the case of personal injuries where the injured party might have suffered what they at first think is a trivial injury, in respect of which they do not consider it worth claiming but which thereafter turns out to be more serious. It might be that this problem is mitigated by the present discussions referred to in the paper as to the limitation of the scope of the data to be retained.
- 54.** However, the nature of the scheme proposed by the Commissions is for liability to rest in the first instance with the insurer of the vehicle (who may then seek relief as appropriate). As between the injured party and the insurer of the vehicle, there are two areas in which data become relevant: first, to establish the presence and the behaviour of the offending vehicle at the time and, second, to assist in establishing the actual cause of the accident. In the case of the former, if the data are not available, it should still be possible for the injured party to establish that they were injured by the vehicle in the same way as presently this is proved – statements of witnesses, photographs etc. In the case of the

second, this may not be the issue which it appears, depending on the approach which is taken to product liability, upon which we comment below.

55. If a time limit for notification is imposed, it ought to be sufficient to require notice to either the insurer or the police (who would have the responsibility for notifying the ADSE). We do not see that automatic notification to the police in every case would be proportionate. We do, however, observe that the injured party (and possibly also the police) might not know or have any means to knowledge of who the insurer of an automated vehicle may be, and that (as envisaged in section 2(2) of the 2018 Act) there may be no insurer involved. In these circumstances, arrangements should be devised to ensure that any data retention request can effectively be made directly to the ADSE.

Q18: Is there a need to review the way in which product liability under the Consumer Protection Act 1987 applies to defective software installed into automated vehicles?

56. The 1987 Act implements the EU Product Liability directive. The directive is presently under review by the European Commission and is the subject of consultation. The consultation is in two parts, with liability for “Artificial Intelligence” systems being separated out. Whatever the outcome of the negotiations for withdrawal of the UK from the EU, it is likely that, given the international nature of the motor industry, the outcome of the Commission's deliberations will have a material bearing on the regulatory regime in the UK.
57. Under the existing legal regime, strict liability arises when damage is caused by a “defective product” (section 2(1) of the 1987 Act). By section 3(1), “there is a defect in a product for the purposes of this Part if the safety of the product is not such as persons generally are entitled to expect.” In theory, therefore, the issue is straightforward – the focus is on the manner in which the vehicle functions, and not the cause of any (mal)function. Persons generally are entitled to expect that a self-driving vehicle will not collide with and injure them. However, in reality, the situation is much more nuanced. Causation does require to be proved, and the tendency has been for courts to seek to identify what exactly has “gone wrong” in order to establish the necessary causal connection. If there are only two possible causes of an accident, of which one leads to liability on the part of the defender, and the other does not, then elimination of the one which does not, ineluctably leads to the conclusion that the cause was the one which gives rise to liability. However, where there is a possibility of more than two causes only

one of which gives rise to liability, then the court will not draw the inference of liability because the pursuer has eliminated only one of the other possible causes. This applies even if the pursuer has eliminated an identified cause, but there remains the possibility of another and unidentified cause. In this regard, we refer to the Speech of Lord Brandon of Oakbrook in *Rhesa Shipping Co SA v Edmunds* [1985] 1WLR 948 at pp 955 – 956 where he refers to the comment by Sherlock Holmes to Dr Watson, “How often I have I said to you that, when you have eliminated the impossible, whatever remains, however improbable, must be the truth?” He continued

‘In my view there are three reasons why it is inappropriate to apply the dictum of Mr Sherlock Holmes, to which I have just referred, to the process of fact-finding which a judge of first instance has to perform at the conclusion of a case of the kind here concerned.

“The first reason is one which I have already sought to emphasise as being of great importance, namely, that the judge is not bound always to make a finding one way or the other with regard to the facts averred by the parties. He has open to him the third alternative of saying that the party on whom the burden of proof lies in relation to any averment made by him has failed to discharge that burden. No judge likes to decide cases on burden of proof if he can legitimately avoid having to do so. There are cases, however, in which, owing to the unsatisfactory state of the evidence or otherwise, deciding on the burden of proof is the only just course for him to take.

“The second reason is that the dictum can only apply when all relevant facts are known, so that all possible explanations, except a single extremely improbable one, can properly be eliminated. ...

“The third reason is that the legal concept of proof of a case on a balance of probabilities must be applied with common sense. It requires a judge of first instance, before he finds that a particular event occurred, to be satisfied on the evidence that it is more likely to have occurred than not. If such a judge concludes, on a whole series of cogent grounds, that the occurrence of an event is extremely improbable, a finding by him that it is nevertheless more likely to have occurred than not, does not accord with common sense. This is

especially so when it is open to the judge to say simply that the evidence leaves him in doubt whether the event occurred or not, and that the party on whom the burden of proving that the event occurred lies has therefore failed to discharge such burden.'

58. One can readily see that these observations would be highly pertinent in the case, for example, of a level 3 or level 4 vehicle, where an accident might have been caused either by a malfunction of the automatic system or by negligence on the part of the driver whilst the vehicle was under manual control or by some other possible cause (even if unidentified) not entailing malfunctioning of the automatic control system. However, what gives the issue of causation added bite is that, although it might be possible to identify an error in conventional algorithms, it is of the nature of a neural network that it is impossible to determine how the network acted as it did. The consequence is that in the case of level 3 – 5 vehicles in general, and, especially such vehicles operated by means of neural networks, an injured party might have an extremely difficult, if not impossible, task in establishing the necessary causation.
59. One way in which this could be addressed would be to introduce a statutory presumption that where a vehicle caused an accident whilst under automatic control, then the accident was caused by the functioning of the automatic system. It may, however, be appropriate to consider whether or not that presumption should be rebuttable.
60. In any event, the legislation should make explicit that the ADSE should be amongst those who are liable under section 2 of the 1987 Act.

Q19: Do any other issues concerned with the law of product or retailer liability need to be addressed to ensure the safe deployment of driving automation?

61. We cannot identify any other areas requiring attention

Q20: We seek views on whether regulation 107 of the Road Vehicles (Construction and Use) Regulations 1986 should be amended, to exempt vehicles which are controlled by an authorised automated driving system.

62. We agree with the Commissions' proposals in this regard.

Q21: Do other offences need amendment because they are incompatible with automated driving?

63. We have not been able to identify any other offences which may need amendment.

Q22: Do you agree that where a vehicle is: (1) listed as capable of driving itself under section 1 of the Automated and Electric Vehicles Act 2018; and (2) has its automated driving system correctly engaged; the law should provide that the human user is not a driver for the purposes of criminal offences arising from the dynamic driving task?

64. Yes, subject to the need to prohibit (or create criminal and/or civil liability for) unauthorised modifications to automated driving systems.

Q23: Do you agree that, rather than being considered to be a driver, a user-in-charge should be subject to specific criminal offences? (These offences might include, for example, the requirement to take reasonable steps to avoid an accident, where the user-in-charge is subjectively aware of the risk of serious injury (as discussed in paragraphs 3.47 to 3.57)).

65. Yes, though see our reservations about creating criminal liability in the specific example of failure to avoid an accident (in our answer to Q3). However, there is scope for criminal liability for a user-in-charge for non-driving tasks such as engaging a self-driving system when that might be prohibited by local rules; or in failing to remove a vehicle which has come to a safe stop (but is not parked in a suitable location) from the public highway. For vehicles without a “user-in-charge”, the same criminal liability would lie with the role we have proposed calling the “operator” (see our answer to Q2).

Q24: Do you agree that: (1) a registered keeper who receives a notice of intended prosecution should be required to state if the vehicle was driving itself at the time and (if so) to authorise data to be provided to the police? (2) where the problem appears to lie with the automated driving system (ADS) the police should refer the matter to the regulatory authority for investigation? (3) where the ADS has acted in a way which would be a criminal offence if done by a human driver, the regulatory authority should be able to apply a range of regulatory sanctions to the entity behind the ADS? (4) the regulatory sanctions should include improvement notices, fines and suspension or withdrawal of ADS approval?

66. Yes to all. There should also be a requirement to state who was the “user-in-charge” / “operator” of the vehicle at the time.

67. It should be a presumption that complying with road traffic law is a necessary part of safe operation. Breach of the law should be presumptively treated as unsafe operation and lead to urgent review of the operation of the self-driving system.

68. As a self-driving vehicle is likely to record a significant amount of information which may be useful in investigating whether there has been a breach of the law or unsafe operation, a notice of intended prosecution should probably trigger a requirement to provide access to vehicle data whether or not the vehicle was self-driving. It should also be considered whether deliberately erasing or failing to maintain vehicle data after an accident or incident should be a criminal offence.

Q25: Do you agree that where a vehicle is listed as only safe to drive itself with a user-in-charge, it should be a criminal offence for the person able to operate the controls (“the user-in-charge”): (1) not to hold a driving licence for the vehicle; (2) to be disqualified from driving; (3) to have eyesight which fails to comply with the prescribed requirements for driving; (4) to hold a licence where the application included a declaration regarding a disability which the user knew to be false; (5) to be unfit to drive through drink or drugs; or (6) to have alcohol levels over the prescribed limits?

69. Yes.

Q26: Where a vehicle is listed as only safe to drive itself with a user-in-charge, should it be a criminal offence to be carried in the vehicle if there is no person able to operate the controls?

70. Yes.

Q27: Do you agree that legislation should be amended to clarify that users-in-charge: (1) Are “users” for the purposes of insurance and roadworthiness offences; and (2) Are responsible for removing vehicles that are stopped in prohibited places, and would commit a criminal offence if they fail to do so?

71. Yes. This liability should also be extended to cover the role of “operator” as we have described it (see our answer to Q2).

Q28: We seek views on whether the offences of driving in a prohibited place should be extended to those who set the controls and thus require an automated vehicle to undertake the route.

72. Yes, subject to clarity on what “set the controls” may mean. In our view, this liability should attach to the role of “operator”. The user of a taxi-type automated vehicle might select a prohibited destination on a control system or app by which a journey is requested. But in this type of vehicle, the user is simply a passenger and should not be liable for the vehicle being driven into a prohibited place. Furthermore, we note the use of the word “thus” in the question. We consider it essential that causation be established. It may be that the default settings of the system would not cause the vehicle to drive in a prohibited place, and that,

if the vehicle did drive in that place, it must have been as a result of the operator overriding the controls. However, satellite navigation systems are not fool proof, and one of the members of the committee drafting the present response recollects a journey in Durham where the satellite navigation system suggested that the vehicle drive, first, along a road open only to pedestrians and cyclists and then, in rapid succession, to drive up a staircase. Of course in a level 0 to level 2 vehicle, the driver should exercise responsibility over whether to follow such a route, but where the vehicle is at level 4 or 5, the question arises of whether the driver should intervene to over-ride the automatic controls in a such a situation and, in the case of a path 2 vehicle, it is even possible to do so.

Q29: Do you agree that legislation should be amended to state that the user-in-charge is responsible for: (1) duties following an accident; (2) complying with the directions of a police or traffic officer; and (3) ensuring that children wear appropriate restraints?

73. Yes, subject to our view that these liabilities should attach to the role of “operator” whether onboard or remote (as appropriate to the authorisation under which the vehicle operates).

74. Liability for the duties following an accident would, as with present legislation, only apply where the accident was reasonably within the knowledge of the operator.

Q30: In the absence of a user-in-charge, we welcome views on how the following duties might be complied with: (1) duties following an accident; (2) complying with the directions of a police or traffic officer; and (3) ensuring that children wear appropriate restraints.

75. In our view, liability to comply with these duties should attach to the role of “operator” even where the operator is remote from the vehicle. In some cases, the vehicle itself may have a role (for instance, in not operating until it detects that all passengers are appropriately restrained) but the ultimate liability should lie with the operator.

76. We would suggest that it is a condition of the authorisation of a self-driving vehicle that there is some form of automatic accident detection system and a specified set of automatic post-accident actions (which may include automatic notification to the authorities and the preservation of pre- and post-accident vehicle data for subsequent analysis and investigation).

Q31: We seek views on whether there is a need to reform the law in these areas as part of this review.

77. We tentatively suggest that these issues will largely be dealt with by technical measures within self-driving systems forming part of any authorisation scheme. For instance, it could be specified that self-driving vehicles will not operate if they detect that passengers are not appropriately restrained. Equally, a route by which the authorities can alert the remote operators of vehicles without a “user-in-charge” may need to be part of an authorisation scheme.

Q32: We seek views on whether there should be a new offence of causing death or serious injury by wrongful interference with vehicles, roads or traffic equipment, contrary to section 22A of the Road Traffic Act 1988, where the chain of causation involves an automated vehicle.

78. Yes. It is arguable that such interference may, in Scotland, fall within the definition of murder, culpable homicide or culpable and reckless conduct. Nonetheless, our view is that a specific offence may provide greater clarity.

Q33: We seek views on whether the Law Commissions should review the possibility of one or more new corporate offences, where wrongs by a developer of automated driving systems result in death or serious injury.

79. We agree that such corporate offences might be appropriate.

Q34: We seek views on whether the criminal law is adequate to deter interference with automated vehicles. In particular: (1) Are any new criminal offences required to cover interference with automated vehicles? (2) Even if behaviours are already criminal, are there any advantages to re-enacting the law, so as to clearly label offences of interfering with automated vehicles?

80. Whilst we are generally of the view that criminal legislation should not be expanded where existing laws can cover the conduct in question, we accept that automated vehicles may be very heavily reliant for their safety on ancillary systems, such as sensors, software and communication systems and that such systems may be fitted in the vehicles themselves or as part of the road and traffic infrastructure. For that reason, we can see there is a good argument for re-enacting the law to deal with unauthorised interference with vehicles or infrastructure.

Q35: Under section 25 of the Road Traffic Act 1988, it is an offence to tamper with a vehicle's brakes "or other mechanism" without lawful authority or reasonable cause. Is it necessary to clarify that "other mechanism" includes sensors?

81. See our answer to Question 34 above.

Q36: In England and Wales, section 12 of the Theft Act 1968 covers "joyriding" or taking a conveyance without authority, but does not apply to vehicles which cannot carry a person. This contrasts with the law in Scotland, where the offence of taking and driving away without consent applies to any motor vehicle. Should section 12 of the Theft Act 1968 be extended to any motor vehicle, even those without driving seats?

82. We express no view on this as it is a matter of the law in England and Wales.

Q37: In England and Wales, section 22A(1) of the Road Traffic Act 1988 covers a broad range of interference with vehicles or traffic signs in a way which is obviously dangerous. In Scotland, section 100 of the Roads (Scotland) Act 1984 covers depositing anything on a road, or inscribing or affixing something on a traffic sign. However, it does not cover interfering with other vehicles or moving traffic signs, even if this would raise safety concerns. Should section 22A of the Road Traffic Act 1988 be extended to Scotland?

83. Arguably, the offence of culpable and reckless conduct already deals with such behaviour; however, given the reliance of automated vehicles on ancillary systems and infrastructure, interference with such systems should be clearly unlawful. See our answer to Question 34 above.

Chapter 9 – General Considerations

84. In responding to questions 38 to 45, we are mindful of the distinction referred to above, between systems which operate through the use of conventional algorithms, and those which depend upon neural networks. In Chapter 6, the portmanteau term "artificial intelligence" is used to describe both types of system, without differentiation. This has the effect of eliding what are quite different types of output of the operation of the systems. In both, the vehicle behaves in a particular way, but, in the former, that is as a result of the working through of pre-programmed algorithms, whereas, in the latter, the machine performs internal processes which may be opaque and produces results which can be described only stochastically. In advanced neural networks, it may be possible for the system to exercise an independent judgement as a human does, albeit not in the same way, but we are still a long way from realising such systems. In responding to the following questions, we have these matters in mind.

Q38: We seek views on how regulators can best collaborate with developers to create road rules which are sufficiently determinate to be formulated in digital code.

85. This is a not unrealistic ambition. Clear and determinate road rules very much lend themselves to incorporation into algorithms such as those currently employed in automated vehicle technology. Beyond endorsing the objective, we leave detailed comment on the means of achieving that objective to others, for example, programmers and IT engineers, better qualified to comment.

Q39: We seek views on whether a highly automated vehicle should be programmed so as to allow it to mount the pavement if necessary: (1) to avoid collisions; (2) to allow emergency vehicles to pass; (3) to enable traffic flow; (4) in any other circumstances?

86. The ultimate objective should be to create automated systems which can make decisions based upon the exercise of judgement as humans do. However, it is difficult to see how a system based upon algorithms could be programmed to mount a pavement where that involves the making of a judgement. In reality, this is a less extreme manifestation of the trolley problem. An algorithm could be devised to allow mounting of the pavement where two conditions exist: (a) there is a reason (e.g. risk of collision, presence of emergency vehicle etc.) to do so and (b) it is safe to do so (e.g. no pedestrians are present on the pavement); but if condition (b) is not fulfilled then all that a “dumb” algorithm might reasonably be programmed to do would be not to mount the pavement. It is unlikely to be in a position to make a judgment whether it would be more appropriate to collide with an obstruction on the road or one on the pavement, though the algorithm may allow for a hierarchy of possible outcomes depending on the nature of the obstruction on the road and how many pedestrians are present on the pavement.

Q40: We seek views on whether it would be acceptable for a highly automated vehicle to be programmed never to mount the pavement.

87. This very much depends on the capability of the system. There might be circumstances in which the system is sufficiently developed to be able to reach the “decision” referred to above, but, unless the algorithms had reached that degree of functionality (or the system operated on the basis of a neural network which had attained the necessary degree of sophistication), it might, on balance, be better for the default to be that the pavement should not be mounted.

Q41: We seek views on whether there are any circumstances in which an automated driving system should be permitted to exceed the speed limit within current accepted tolerances.

88. We do not favour this suggestion, especially as it is envisaged that an automated driving system could exceed the relevant speed limit *within current accepted tolerances*. In these circumstances, the default position should be that the normal maximum speed would be the applicable speed limit, with the relevant tolerance giving some leeway for a spurt of speed to avoid a potentially dangerous situation. The late notice speed restriction situation could involve a programmed response at a specified level of deceleration, short of the alarming maximum braking, even if that tolerates entry into the controlled zone at a higher (but still swiftly reducing) speed. There should not be any need for exceeding current accepted tolerances in an overtaking manoeuvre if such manoeuvres are executed only under safe conditions; and it ought not reasonably to be expected that an overtaking manoeuvre should be undertaken where it is not safe to do so. However, these observations are made upon the assumption that the automated driving systems are capable of making the necessary assessments and judgments. It is possible that, one day, neural networks might be so capable, but it is less likely that an algorithmic system could achieve the requisite degree of sophistication. In short, the rules which are to be set should depend not only on the anticipated traffic situations but also be set depending upon the technical abilities of the system.

Q42: We seek views on whether it would ever be acceptable for a highly automated vehicle to be programmed to “edge through” pedestrians, so that a pedestrian who does not move faces some chance of being injured. If so, what could be done to ensure that this is done only in appropriate circumstances?

89. We suspect this is unlikely to be possible with a system based on algorithms. Such decisions are unlikely to be practicably achievable without human judgement and intervention, or, possibly, through the use of a system based on a neural network.

Q43: To reduce the risk of bias in the behaviours of automated driving systems, should there be audits of datasets used to train automated driving systems?

90. Insofar as a system is trained by the inputting of data, we would answer this question in the affirmative.

Q44: We seek views on whether there should be a requirement for developers to publish their ethics policies (including any value allocated to human lives)?

- 91.** This opens a can of worms in relation to the regulation and control of “Artificial Intelligence” systems in general, and not only driving systems in particular. It presupposes that a neural network can be engineered to make value judgments based upon ethical premises in respect of which it has been educated. A basic level of functionality might be achievable, not only for neural networks, but also for “dumb” algorithms. As we discuss above, an algorithm might be devised which can decide between running down three pedestrians on a road or one on a pavement; but the thought experiment is often elaborated by asking the person undertaking it whether his answer would be any different if, for example, the pedestrian on the pavement were Einstein and those on the road were a gang of housebreakers. This kind of judgment would simply not be possible for a system based on algorithms, and it is difficult to see how a neural network, or even a human being, could reasonably be expected to have available to it the necessary data as to the identity of the respective pedestrians. We use the word “reasonably” because we are aware that the Chinese government is in the process of developing a system which ascribes to individuals a social credit score which measures their value to society (based, of course, upon criteria fixed by the government). It would be technically possible to implant a chip in individuals which would transmit, amongst other things, their social credit score which could then be detectable by an automated driving system, thereby making it technically possible for an automated driving system to exercise a choice between Einstein and the Housebreakers, but we cannot conceive of any circumstances whatever where such a system could be regarded as acceptable in a free, open and democratic society.
- 92.** Furthermore, there is no single right answer to the trolley problem. Most people would divert the trolley to kill the lower number, but a significant minority would not intervene. We note that the Law Commissions explicitly record that their terms of reference do not include judging between different ethical approaches. Unless it were part and parcel of a system of enforcement of an ethical system determined by society, in the form of Parliament, as being the “correct” system, it is difficult to see what value there would be to knowing that, say, Ford cars would not mount the pavement to avoid a collision whereas Tesla cars would do so. In addition, given that there is currently no correct answer to such ethical dilemmas, it does raise the question of whether, when purchasing a level 5 vehicle, the purchaser might be able to specify the ethical system with which the car is

programmed (for example, Benthamite ethics) as well as specifying the paint colour and interior trim.

- 93.** The purpose of our discussion in the previous two paragraphs is not so much to grapple with these questions as to illustrate that the technology of level 5 automated driving systems is at a relatively early stage. There are those who talk up the possibility of a fully developed neural network, capable in the real world of making ethical decisions based on the input of sufficient accurate data as being only a year or two away, whereas other commentators are sceptical of our ability to develop such a system in the near future, even though, with the development of quantum computing, it may be possible in the medium term. In these circumstances, we should suggest that it is entirely appropriate for the Law Commissions to initiate discussions on these matters, but that time may be needed to monitor developments in the field before coming to final conclusions. Any legislation which may result from the present consultation should be structured so as to permit the making of regulations dealing with these matters when the technology has developed further and public discussion on ethics admits of the possibility of consensus.

Q45: What other information should be made available?

- 94.** We cannot think of any other information which should, at this time be made available, though this could be revisited as part of the process described in the previous paragraph.

Q46: Is there any other issue within our terms of reference which we should be considering in the course of this review?

- 95.** As we noted in the Introduction, it is likely that Automated Driving Systems may be developed so as to interoperate with other vehicles, traffic lights and other road infrastructure. Failure of these elements to interoperate would prejudice the safe operation of the automated driving systems, and may represent the equivalent of a human driver failing to comply with traffic signs or driving without due care and attention. The Law Commissions may wish to give consideration to how such breaches, which are only possible by automated vehicles and are not generally covered by offences applicable to human drivers, are regulated, detected, investigated and (if appropriate) prosecuted.
- 96.** Further, as we also pointed out in the introduction, elements, or perhaps the entirety, of a vehicle's ADS may not be physically located in the vehicle but may be located on a remote server "in the cloud" with the ADS controlling the vehicle remotely. The operation of such a

system would be dependent on high-speed, high-capacity network communication. An ADS of this sort does not fit easily within a regulatory regime in which a vehicle is considered as a single entity which can be under the control of a “user-in-charge” or can be authorised or insured on an individual or type basis. The Law Commissions may need to consider whether the development of such distributed implementations of ADS is a realistic possibility and if so, whether a differently-structured regulatory regime may be necessary to support such vehicles.