

Burges Salmon response to joint Law Commission and Scottish Law Commission preliminary consultation paper dated 8 November 2018 on Automated Vehicles



1 INTRODUCTION

- 1.1 Burges Salmon is an independent UK law firm with a market-leading transport law practice across all modes include rail, road transport and highways, aviation and airports and marine and ports.
- 1.2 Burges Salmon has been actively partnering in Connected and Autonomous Vehicles (CAV) research and development projects and trials since 2014. These include the Innovate UK-funded projects VENTURER, FLOURISH, Capri, RoboPilot and MultiCAV.
- 1.3 Burges Salmon has published the following legal reports in conjunction with AXA on CAV issues relating to insurance, safety, civil and criminal liability, data and cyber-security:
 - (a) VENTURER Year 1 Report 2016¹
 - (b) FLOURISH Year 1 Report 2017²
 - (c) “Are we ready to ‘handover’ to driverless technology” (VENTURER, April 2018)³
 - (d) “Driverless cars: liability frameworks and safety by design” (VENTURER, June 2018)⁴
 - (e) FLOURISH Year 2 Report 2018⁵
- 1.4 More generally on highway safety, we produced a report in April 2018 for the RAC Foundation⁶: ‘A Highways Accident Investigation Branch - What Lessons can be learned from the Rail Industry and the Cullen Inquiry?’. Subsequent to that report the RAC Foundation received funding from DfT to run a pilot on highways accident data analysis.
- 1.5 We welcome the 3 year joint project of the Law Commission of England and Wales and Scottish Law Commission to review the legal framework for automated vehicles and we are pleased to submit our response to the preliminary consultation below.
- 1.6 Should the Law Commissions wish to explore any part of our responses further then we are happy to assist further where we can.

2 GENERAL CORE THEMES

- 2.1 Automated vehicles present a challenge to a framework of driving laws and rules which have, from their inception in the 19th Century and subsequent development, rested on the fundamental basis that ‘driving’ is solely a human activity. Application and enforcement of those driving laws focuses on individual and indivisible human actors. Applicable legal criteria and tests of expected standards or criminalised behaviours relating to vehicle ownership and use have also been constructed around the language and lenses of human behaviours and standards – reasonableness, due care and attention, careless, unfit, wanton or furious. The law has, logically, focused on individual actions looking at the behaviour of a defined individual against a defined obligation falling wholly upon the individual.
- 2.2 The introduction of automated vehicles that are truly capable of driving themselves will create a new class of non-human activity that is more analogous to automated processes of products.
- 2.3 Further, the performance of these products is likely, amongst other things, to:

¹ <https://www.venturer-cars.com/wp-content/uploads/2016/07/VENTURER-AXA-Annual-Report-2016-FINAL.pdf>

² http://www.flourishmobility.com/storage/app/media/publication/J381379_Brochure_Flourish%20Report_V14_SPREAD_S.pdf

³ <https://www.venturer-cars.com/legal-and-insurance-report-2017-18/>

⁴ <https://www.venturer-cars.com/wp-content/uploads/2018/06/Year-3-Legal-and-Insurance-Report.pdf>

⁵ http://www.flourishmobility.com/storage/app/media/FLOURISH_Insurance_and_Legal_Report_2018.pdf

⁶ https://www.racfoundation.org/wp-content/uploads/HAIB_Burges-Salmon_April_2018.pdf

- (a) Involve complex interaction between hardware and software;
 - (b) Require continuous post-supply software updating;
 - (c) Utilise and/or depend on external and/or third party services to operate (in part or in whole) such as communications networks and navigation systems;
 - (d) Require clear delineations of responsibility and handover as regards automated driving and human driving and the activities of other humans “in the loop” as part of the system; and
 - (e) Increasingly rely on advanced machine learning capabilities to develop decision-making even post-supply.
- 2.4 Consequently, automated vehicle driving does not just have features of automated processes of products but also systems and machine learning processes. The operation of a vehicle therefore begins to take on features and aspects which may be more familiar to transport systems in other modes such as rail or aviation. In those systems, the human component is just one factor within a complex set of interfaces and interdependent obligations. That is a different reality. The legal models for the regulation of complex systems (and human behaviour within them) are obviously very different to the approach used to regulate individual stand-alone behaviour
- 2.5 The regulation of safety systems is multi-factoral and inherently not one-size fits all. The safety assessment of an automated vehicle requires not just assessment of the construction and roadworthiness of the vehicle but, amongst other things, the competence and capability of its automated driving system and its specific use (defined by its variable operational design domain), operational protocols and its external dependencies.
- 2.6 An overarching point therefore is that, in considering automated vehicles, we make no assumption that automated driving and human driving should be treated the same way in legal terms. Given the conceptually different nature of the activity, treating human and automated driving the same would not be logical in many cases. That said, in other circumstances it would be logical to draw no distinction. For example, our work with Axa in VENTURER advocated strongly for a single insurer model for automated vehicles since, in the event of an accident, the effect on that third party was the same regardless of whether the vehicle was being driven by a human or not and the third party ought to have the same level of insurance protection. The Automated and Electric Vehicles Act 2018 in the event adopted the same approach
- 2.7 Finally, given the fast-changing pace of technology there should be a preference for creating a flexible safety-based framework for regulating automated vehicle systems that is able to regulate and approve different vehicles, use cases and specifications accordingly to their required permissions and conditions when approving those use cases. The legal framework needs to be able to accommodate automated vehicles which may vary from low speed pods to platooning heavy goods vehicles.
- 2.8 The objectives against which the updated legal framework needs to be tested against broad objectives to be:
- (a) Deliverable and fair, in particular it needs to reflect what is known about human/system interfaces from other transport modes and automated vehicle trials. Behaviour should not be required or criminalised that does not reflect a realistic or consistent human response (for example on hand-back)
 - (b) Adaptable and robust ;
 - (c) Identify and define the behaviours to be incentivised, dis-incentivised (and criminalised);

- (d) Governed by a mechanism which makes sure that responsibilities/competencies are clear in relation to different use cases;
- (e) Clear when using fault based attributions and non-fault based attributions for civil allocations and penal responses; and
- (f) Deliver safer outcomes overall and deal transparently with probabilistic safety concepts with clarity including societal perception issues.

Burges Salmon LLP

February 2019

3 CONSULTATION QUESTION RESPONSES

Question 1:

Do you agree that:

- 1 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?
- 2 The user-in-charge
 - Must be qualified and fit to drive
 - Would not be a driver for purposes of civil and criminal law while the automated driving system is engaged; but
 - Would assume the responsibilities of a driver after confirming that they are taking over the controls, subject to the exception in (3) below?
- 3 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

Vehicles which “drive themselves” within the meaning of the Automated and Electric Vehicles Act 2018

We have some reservations as to extending the current definition of “automated vehicles” in the context of insurance and the Automated and Electric Vehicles Act 2018 (AEVA) to underpin developing civil and criminal law. Our understanding is that the Law Commissions intend in addressing automated vehicles to mean only vehicles with functionality equivalent to SAE Level 4 and above. In other words, vehicles which do not require the intervention of a human driver for safety reasons and are able to default to a minimal risk condition by themselves if necessary. However, the AEVA in principle could arguably include SAE Level 3 vehicles.

As the Law Commissions note, the Government policy position on the definition of “automated vehicle” under the AEVA was that it did not include SAE Level 3 vehicles⁷ although the issue was not addressed directly in the legislation (notwithstanding some debate around the issue in Parliament). However, what the AEVA does say is that:

- Listing of automated vehicles by the Secretary of State is for vehicles which are “*designed or adapted to be capable, in at least some circumstances or situations, of safely driving themselves*” and “*may lawfully be used when driving themselves, in at least some circumstances or situations, on roads or other public places in Great Britain*”⁸; and
- That “driving itself” is where the vehicle is “*operating in a mode in which it is not being controlled, and does need to be monitored, by an individual*”⁹

However, the SAE definition makes clear that in Level 3 driving the vehicle is neither being controlled nor monitored by an individual (but rather the individual must be “receptive” to a handover request which is not monitoring)¹⁰. Consequently, a Level 3 vehicle manufacturer could seek listing of a vehicle as an automated vehicle under the AEVA as long as it could demonstrate that its operation was safe and lawful.

If the intention of the Law Commissions is to utilise the definition of “automated vehicles” in the

⁷ Paragraph 2.56 of the Consultation Paper.

⁸ AEVA, Section 1(1)

⁹ AEVA, Section 8(1)(a)

¹⁰ As noted in paragraph 2.9 of the Consultation Paper.

AEVA as meaning, definitively, SAE Levels 4 and above per the Government's policy, it may be prudent to ensure first that the statutory definition is not open to technical interpretations that could scope in Level 3 vehicles.

The remaining responses below, assume that by "automated vehicles" the Law Commissions **do** intend to draw a distinction between SAE Levels 3 and below and Levels 4 and above. When we discuss automated vehicles we mean the latter.

That said, even within the category of Level 4 and 5 vehicles, there could be a very large degree of differentiation between vehicle types, uses, operational design domains (ODD), etc. The SAE Levels only describe certain aspects of functionality and are otherwise quite broad.

Should all automated vehicles have a "user-in-charge"?

We consider that it is not possible to adopt a blanket approach to this and that doing so could be unnecessarily restrictive.

Fundamentally, the requirement for a "user-in-charge" should be driven by a safety case and risk assessment of the particular automated vehicle and its operation. Automated vehicles could include a range of vehicle types from accessibility pods travelling at low speed to adapted goods delivery vehicles to full size and speed automated cars or lorries. In operation, they may or they may not be expected to experience conditions outside their operational design domain. In each case and for each vehicle, these reflect different safety cases and risk profiles.

The requirement to have a "user-in-charge" is a mitigation measure that may be necessary and/or proportionate to the hazards and risks presented by one type of automated vehicle and use case but not another. Mandating one mitigation measure regardless of actual risk profile would be a blunt regulatory tool likely to give rise to adverse or restrictive unintended outcomes.

We submit that whether or not a "user-in-charge" should be a requirement for a particular vehicle or use case should be considered as part of the risk assessed regulatory approval process which that automated vehicle should be required to clear before being deployed on roads or other public spaces. It could be a condition of use offered as part of the safety case at that stage or a condition of approval applied if deemed necessary.

It would be a sensible and useful requirement for some use cases. It is likely to be unhelpful in others.

Status of the "user-in-charge"

Notwithstanding the above, for certain automated vehicles, we can understand that there might be a need and requirement for a "user-in-charge" for safety reasons. However, within the context of a safety case or approval, humans may be placed "in the loop" in a variety of different ways and for a variety of different purposes.

As we understand it, the user-in-charge as described by the Law Commissions would essentially be a human driver on standby - hence the requirement that they be qualified and fit to drive. This describes a human in the loop who is expected, under an established protocol and safety management system, to take control of the vehicle and undertake the dynamic driving task. As we further understand it, the concept does not necessarily exclude teleoperation where the individual that may be performing this role does so remotely (i.e. a human driver on standby but not on board). We consider that there could be more clarity around this point as some CAV manufacturers and operators are actively considering and testing such teleoperation models¹¹. We note that the latest CCAV Code of Practice on automated vehicle trialling (published 6 February 2019) recognises that safety drivers in trialling may not be in the vehicle and provides guidance on additional safety assurances in such circumstances (e.g. in respect of connectivity, performance lag, etc)¹².

¹¹ <https://www.businessinsider.com/autonomous-vehicles-remote-human-control-2018-12?r=US&IR=T>

¹² https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/776511/code-of-practice-automated-vehicle-trialling.pdf

In principle, we agree that such an individual must be qualified and fit for their role. Where that role is undertaking the dynamic driving task, we are unable at this stage to say whether those driving requirements are likely to be exactly the same as manual driving requirements today. The specific technology and interfaces of the CAV system may require different or additional requirements to demonstrate competence. We are conscious that, in other transport modes, human interaction with automated driving systems is currently conducted through professional train drivers and airline pilots for example.

Where the human in the loop is not in fact tasked with the role of undertaking the dynamic driving task, it would be necessary to distinguish this from the "user-in-charge" role envisaged here. These other individuals may have a range of roles, some of which may be indirectly linked to the dynamic driving task but not directly engaged in that task. For example, humans in the loop may in certain circumstances have roles in resolving technical issues remotely or assisting the automated driving system resolve cognitive / processing anomalies. It does not appear to us that the Law Commission intends for such humans in the loop to fall within its category of "user-in-charge".

Liability of the "user-in-charge" around handover

In respect of a human in the loop who may be required to assume the dynamic driving task from an automated vehicle, we suggest that:

- When the automated driving system is driving, then the user-in-charge would not be the driver for civil and criminal law purposes as a matter of fact and that should be reflected in legislation
- The 'handover process' for an automated vehicle needs to have the liability interface mapped in each direction: human to machine handover and machine to human handover.
- In respect of "planned" handover (by which we mean a routine driving state handover from human to machine or a non-safety critical request to intervene from the machine to the human):
 - In terms of handover from human to machine, where the automated driving system has been activated within its operational design domain and in accordance with the (approved) terms of use, we suggest that ordinarily the human would cease to be the driver at the point of activation. Given the potential differentiation and confusion around operational design domain, we would expect manufacturers to make clear where and when automated driving functions can be used and/or to design in fail-safes to prevent automated functions from being available for activation unless and until the system has confirmed that the situation conforms to its operational design domain
 - In terms of handover from machine to human, the human would only become the "driver" for liability purposes once they had completed the authorised (and approved) protocol for safe handover of a dynamic driving task. We are not able within this response to map out the requirements of a "safe handover" but we have previously stressed that this should be done as part of the certification/approvals process of any automated driving system.
- In terms of "unplanned" handover scenarios:
 - Consideration should be given to scenarios where the human has been permitted to handover a dynamic driving task to an automated vehicle in a non-routine driving state manner. This may include doing so knowingly in breach of the applicable operational design domain (notwithstanding any fail-safes that may prevent this) or for example when the automated vehicle is at imminent risk of an accident which then comes to pass. In such circumstances, it would be wrong in principle to absolve the human of responsibility notwithstanding that the automated driving system had been activated.

- In terms of unplanned handovers from machine to humans, it is our view that at SAE Level 4 above then the automated vehicle should not make unnecessary requests to intervene whilst it is driving within its operational design domain and certainly not under any safety-critical circumstances. Such a vehicle should by design be able to secure its own safety and that of its passengers by achieving a minimal risk condition first. Given the human factors and other complexities around handover, such a handover would in itself import material safety risk into the dynamic driving task.
- We note the principles of 'transition' at sections 5.13 to 5.20 of the CCAV Code of Practice for automated trialling which have been set out with safety in mind and would suggest that these principles should extend not just to trialling but in due course deployment.

Please see our extensive prior literature around the aspect of handover and the risks associated with it from the VENTURER reports and findings

Human-initiated “override”

Notwithstanding the 1968 Vienna Convention¹³, unplanned human-initiated handover from an automated driving system and liabilities for any consequences of that is a particularly challenging area for legislation (which we explore further below as regards criminal law).

Given the functionality of SAE Levels 4 and above, it would seem that the safest implementation of a human “override” would be a request for the automated vehicle itself to achieve a minimal risk condition prior to handing over control of the dynamic driving task to the human. Certainly, on the other hand, the permissibility of an easy and immediate override and assumption of control would need to be carefully assessed given the increased risk of human error or inadvertent override.

Should a more immediate override be permitted to users-in-charge, the potential implications which might arise from just the one scenario of a perceived risk of accident illustrate how the civil law needs to retain flexibility in apportioning liability (see below). Much of this flexibility lies in the existing civil law approach to causation as a mixed issue of fact and law. Consequently we submit that it would not sensible for the law to legislate for a fixed rule on liability in these circumstances but rather to rely on existing concepts of factual and legal causation:

Table 1 – illustrative decision/outcomes matrix where user-in-charge perceives risk of accident

Takes action	Actual risk	Consequence	Issues
Yes	Real	Risk avoided, no damage or injury	User has caused a net positive effect. User may nevertheless have product-based remedies against manufacturer.
Yes	Real	Risk partially avoided, minor damage or injury	User has caused a net positive effect and may personally have full defence to civil liability
Yes	Real	Risk not avoided, additional harms caused	User has caused a net negative effect and may be subject to at least a

¹³ Article 8 currently provides that automated driving technologies either conform to UN vehicle regulations and if not that they will allow for override/switching off by the driver

			contribution claim
No	Real	Risk not avoided	The ability to intervene if it exists but not to have would be tested against reasonableness and causation as per other omission cases
Yes	Not real	Risk created, damage caused	User has caused wholly a net negative effect and would be subject to at least a contribution claim

Question 2:

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

For the reasons above and more, we consider that this term could be more legally precise. Specifically, our understanding is that the individual being referred to:

- Only requires the descriptor and holds this legal status when they are in fact not driving. When they are driving, they will continue to be the legal “driver” with current driving laws applicable in the usual way;
- Will not be every human in the loop but one which is mandated for the relevant vehicle to be one who may be required to undertake the dynamic driving task directly (in person or remotely);
- Is not necessarily a “user” in the sense that this is both potentially too narrow and too broad. Would a tele-operator be a "user" of the vehicle? All persons on board the vehicle would consider themselves "users" of a kind. There is sense in avoiding confusion where persons on board an automated vehicle are users but are not sure if they are or may become “in charge”. Additionally “users” of vehicles does not get across the fundamental requirement that the purpose of this role is that they may have to drive and need to be competent and fit to do that; and
- Is not in fact “in charge” (at least of the core dynamic driving function) at the point that the automated driving function is activated. Being “in charge” would suggest that such users retain responsibility and liability for the dynamic driving task even when they are not driving which should not be the case. Moreover, it is unclear what makes the individual "in charge" and whether or not that is a moral or transferrable responsibility that could pass to other users on board (if for example the 'primary' user-in-charge was incapacitated)

For automated vehicles which require and mandate such a mitigation measure, a term such as “co-driver” or “standby driver” or 'standby controller' may be preferable. These terms get across more of the core characteristics of this individual and, importantly, help clarify roles of persons on board:

- Like ‘co-pilot’ in aviation, these terms get across the concept that such automated vehicles rely on more than one driver (in effect) who will take turns to assume the dynamic driving task;
- Designation of “driver” makes absolutely clear that the purpose of this human in the loop is, at points, to drive and reinforces that they will be expected to be competent and fit to do so even when it is not ‘their turn’. It also distinguishes the individual from other humans in the loop who may be more described as ‘operators’ or ‘attendants’ who oversee some aspects of the automated driving system but are not assuming or

expected to assume control of the dynamic driving task; and

- Designation of “driver” separates this individual from other users of the vehicles. People will understand that they are either passengers or drivers (even if the co-driver or standby driver is not physically in the vehicle but is a tele-operator).

The continued use of “user in charge” in the remainder of our response is subject to our views in this section.

Question 3:

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

We have set out the position above as regards human “override” scenarios and how the flexibility afforded by the civil law and its fact and law based approach to causation will be important to allocating liability.

As regards criminal law however, we would suggest caution against criminalising behaviour in respect of the human 'override' for a number of reasons:

- Criminalising behaviour should start with a clear understanding of the relevant culpable behaviour being targeted. Automated vehicles (Level 4+) should functionally be able to drive competently within its operational design domain so as to minimise harm in circumstances where there is a risk of serious injury. The user-in-charge does not expect an unexpected handover or requirement to intervene for safety reasons and the scenario is not comparable to the omission of a human driver in a runaway car - the automated driving system has control of the dynamic driving task. One would not expect a front seat passenger of a vehicle to attempt to take control of its wheel or handbrake if she believes that the driver is not aware of an imminent risk. Such an action might in itself be criminal behaviour; certainly one would not criminalise a failure to do so. It is unclear what would make such an omission inappropriate where the vehicle is being driven by the automated driving system
- Whilst the 1968 Vienna Convention currently mandates compliance with UN regulations and, failing that, a human override, both the regulations and the convention are subject to change. As technologies improve and become more safe, we would expect to see vehicle regulations to evolve and any associated need for and the scope of any override to recede
- Given the known issues of machine to human handover (see VENTURER trials), handover and particularly in unplanned situations is not a safety neutral activity – it creates risks of unsafe driving
- Studies exist to suggest, amongst other things that:
 - Complex decision-making under stress or made pursuant to *'flight or fight'* responses are not straightforward
 - Human drivers typically over-estimate their driving ability
 - Risk perception varies considerably between different categories of drivers
 - Safety drivers in automated vehicle trials do not have the same "feel" for or response from the road when they are not physically driving (leading to sub-optimal driving for a period even once they do take over for example)

In other words, there are significant risks of false positive interventions and each instance of false positive intervention is less safe for the driver and others around them

than if they had not intervened.

- Criminalising a failure to act is likely to exacerbate the likelihood of false positives even further since it exhorts the human to act even more conservatively and trust the automatic driving system even less. It may additionally incentivise other unhelpful behaviours (e.g. encourage users-in-charge to pay as little attention as possible since they are less likely to become “subjectively aware” of driving risks). The net effect on safety may very well not be positive.

In November 2018, Waymo reported that one of its automated vehicles had been involved in a collision which injured a motorcyclist. The accident resulted from the actions of the experienced and fully trained professional safety driver who, considering the vehicle to be at risk of one collision, took the controls and caused another. Waymo’s review suggested that the automated vehicle would have avoided the original risk completely and not caused the separate accident if the safety driver had not intervened¹⁴. This illustrates the potential risks that need to be considered around human override and adding criminal jeopardy to its exercise.

There is also significant known learning from aviation and rail accidents nationally and internationally in relation to human/system interfaces and reactions from accidents, incidents and near-misses¹⁵. Any formulation of the law in this area should follow a comprehensive review of the knowledge and resulting principles.

There is a risk of extrapolating from current vehicular (driver-centric) law which assumes continuous engagement from a specific individual.

Question 4:

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

Fundamentally this would be for manufacturers and regulators to demonstrate given the particular vehicle, automated driving system, use case, operational design domain and applicable (international) standards and type approvals, risk analysis and safety management system. Ultimately, mitigation measures required will be safety-driven.

The Law Commissions’ consultation identifies a number of known use cases being explored which would not necessarily require a user-in-charge (e.g. public transport operator services or valet parking services). If the safety case for those uses (and any others) can be made to the satisfaction of relevant regulators / authorities, it should be possible to deploy solutions without a user-in-charge or with reliance on other risk mitigations. Mandating a user-in-charge in all situations may increase risk and/or restrict future technology and use cases.

Question 5:

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

Yes - the conditions under which an automated vehicle may perform a dynamic driving task on highways and public spaces should be within the scope of the powers available to the regulators / authorities responsible for certifying the automated vehicle for deployment in the relevant use case.

¹⁴ <https://medium.com/waymo/the-very-human-challenge-of-safe-driving-58c4d2b4e8ee>

¹⁵ For example the Air France South Atlantic loss, the recent Lion Air loss and the Santiago rail derailment.

Question 6:

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

We assume the "driver" in this context is the Law Commissions' user-in-charge. That being the case, it is not clear why the relevant sections of the consultation reference SAE Level 3 vehicles.

We would support in principle the approach of Working Party 1 (WP1) of the UN's Global Forum for Road Traffic Safety. If the automated vehicle has been approved to operate without the user-in-charge being required to monitor the dynamic driving task, then in principle the user-in-charge should be permitted to undertake secondary activities as long as they do not impair whatever role he is or may be requested to undertake by the automated driving system.

There is however a lot more work required to understand in this context what are or are not secondary activities that may affect the role of a user-in-charge. If the vehicle can achieve a minimal risk condition by itself (as a SAE Level 4 vehicle should) then in practice there may be fewer restrictions. By contrast, the driver in a SAE Level 3 vehicle remains (legally) the driver at all times and secondary activities should be highly restricted. The issue of permissible secondary activities is bound into the issue of human factors where a lack of activity or periods of low cognitive load are recognised as potential issues just as much as issues of distraction are. We note that the CCAV Code of Practice (Section 4) provides guidance not just as to clarity of expectations around driver competence and behaviour but also, from the perspective, of distraction of other road users.

We are aware of the German Road Traffic Act (StVG) position (which is in principle permissive of SAE Level 3 vehicles) but it is very high level and the actual detail of what is or is not permissible by way of secondary activities appears to fall to the vehicle manufacturers to prescribe. Issues as to what activities are safe or within the actual capability of humans to undertake safely during such journeys are not transparently dealt with. It is unclear how enforcement will be undertaken.

Given high profile awareness and habit-forming driving campaigns of the past focussed on discouraging unsafe and/or unlawful human behaviours on the road (e.g. tiredness, mobile phone use, lack of due care and attention, etc), it would seem unusual for UK law to leave the matter of defining acceptable human driver secondary activities solely to manufacturers. Particularly where that includes SAE Level 3 where the UK Government position is that the human driver remains the legal driver and that the functionality is basically high functionality driver assistance (Level 2+++)¹⁶ as opposed to automation.

Any system that did rely on manufacturers to define standards of safe conduct by users-in-charge would inevitably require a greater degree of regulation than exists currently. For example, industry concerns have been raised by the likes of Thatcham as to how there is public confusion even as to how mere driver assistance features have been marketed¹⁶.

An integrated system regulation and approval approach similar to that in place for aviation and rail would be required.

¹⁶ <https://news.thatcham.org/pressreleases/carmaker-use-of-the-word-autonomous-a-danger-to-uk-roads-2537576>

Question 7:

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?

See above. Logically the question of what activities the fallback driver of a SAE Level 3 vehicle should be permitted (or required) to undertake should be addressed at the same time as the question of their authorisation for public deployment. They should be part of the same safety assessment.

Currently:

- 1 The understanding of Level 3 systems from projects such as VENTURER is that unplanned handover is highly problematic from a safety and driving performance perspective, not least as the system may request handover in safety critical situations. Beyond very narrowly defined operational design domains (such as Audi's Traffic Jam Pilot which by definition operates at low speed), it is unclear how safe such systems can be. Further research is required.
- 2 Under the legal system in the UK, the driver remains liable in civil and criminal law as the driver and is insured as the driver.

In these circumstances, the secondary activities that a fallback driver may be permitted (or required) to undertake may in practice be very limited. As stated above, there is currently insufficient research and development into what a "safe" Level 3 or above handover process would entail. Also, as stated above, the risk of distraction of a fallback driver needs to be balanced with the risks associated with low cognitive load situations (i.e. the ability of humans to focus and maintain a task with little active input over potentially prolonged periods).

Question 8:

Do you agree that:

- 1 A new safety assurance scheme should be established to authorise automated driving systems which are installed:
 - As modifications to registered vehicles; or
 - 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?

Ahead of any type approval framework for automated driving systems and vehicles, we consider the proposal of a new safety assurance scheme for modified or small series vehicles to be a logical one.

This may involve individual approvals, in which case the relevant safety agency (which could be the VCA or DVSA or a joint or new agency) should have the powers and the capability to undertake full safety assessment of an automated vehicle construction/modification, driving system, its use case and operational design domain. We would expect that the next phase of

automated vehicle testing (without safety driver) and the soon to be updated Code of Practice underpinning that will start to generate the criteria, tests and standards which might in due course be applied amongst others.

We agree with the Law Commissions' analysis of the law and in particular the potential breadth of interpretation that could be applied to the Road Vehicles (Construction and Use) Regulations 1986 but which were nevertheless not drafted with automated vehicles in mind. There is a strong public interest in the regulatory framework developing in conjunction with the technology. Safety assurance and the public acceptance this brings is the paramount factor in the success or otherwise of automated vehicles and the realisation of their potential. It is a task that the public would expect to be overseen by government and not left solely to individual vehicle manufacturers. As the Law Commissions note, the automated vehicle developers that we work with wish to see more guidance on safety standards, not less, and have an interest in ensuring a level playing field in that regard.

It is noted that the CCAV Code of Practice framework and the proposed future regime for approving advanced automation trials would essentially be laying the potential groundwork for such a regime.

Question 9:

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?

We agree with this proposal of the Law Commissions and consider that the concept will bring certainty for automated vehicle developers, authorities and the public.

The parameters of the responsibility of the ADSE would clearly need to be defined. The 2017 Principles and BSI PAS: 1885 2018 both relate to cyber security specifically but provide a useful checklist as to the key elements from design onwards through a system's lifecycle as to how the broader governance and resulting responsibilities of an ADSE might be structured.

Question 10:

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.

The safety assurance system must be proportionate to the hazards and risks involved in the automated vehicle and use case being proposed. That is unlikely to be satisfied by a 'one size fits all' approach to safety regulation. However, multi-tiered safety regulation is a common approach in the UK and one well known to transport regulators. For example, the way that the Office of Rail and Road regulates heavy rail differs from the way that it regulates light rail and tram which in turn differs from how it regulates heritage rail. In general, the fewer the potential hazards, the lesser the potential impacts and the lower the incidences of risk, the "lighter touch" the necessary and proportionate safety regulatory regime will be.

In CAVs, for example the regulation of low impact design pods operating at low speed and/or in confined areas and the regulation of high capacity lorry platoons travelling at high speed would have a common architecture but different depths of oversight and prescription.

It is acknowledged that, at the early stage of technology development, it may be difficult for an ADSE safety regulator to undertake full independent pre-market safety assessment of vehicles. However:

- 1 There is a material body of research and data available on the use of automated systems available from other transport modes (in particular aviation and rail) and experienced regulators in this context that can inform expectations of safety system

management even at an early stage for automated road vehicles;

- 2 The UK has funded and hosted a large number of automated vehicles trials since 2014 including a network of physical automated testbeds under Meridian. It has also initiated competitions in and trials are ongoing in respect of simulation testing for automated vehicles;
- 3 The current and future cohort of automated vehicles trials and data are already identifying and focussing attention on specific areas of safety concern for autonomous vehicles – handover, cyber-security, incident data, etc. In light of this, the safety assurance system should specifically address safety both at a general level and at the level of identified specifics;
- 4 The UK must involve the public in the consideration of safety assurance since they are those seeking to be assured of safety as much government or regulators and the non-passenger general public will be potentially more exposed to automated vehicles than rail and aviation which (on the whole) operate in protected (and segregated) spaces. The level of safety assurance required to enable and underpin an effective safety assurance scheme and automated vehicle sector is not just determined by actual risk but *perceived* risk too.

Consequently, whilst the Australian case study is useful to consider, the UK will need to conduct its own assessment on an appropriate regulatory regime based on its own evidence base and resources. That regulatory framework must however be able to deal effectively with proposed deployments ranging from single low speed mobility pods in rural settings to large scale fleet deployment in urban centres. It is anticipated that the minimum requirement is for authorities to be satisfied through active monitoring and audit that the relevant ADSE is competent, has undertaken a comprehensive safety assessment and risk management programme and has a continuing and effective safety management system. This would not therefore be “self-certification” (which, in our experience, automated vehicle and automated driving systems developers in the UK have, broadly, not been in favour of, notwithstanding legal developments in Arizona, Germany and others relying on self-certification).

The 2017 cyber security Principles and PAS 1885 provide a structured approach which could sensibly apply to safety and assurance beyond cyber, covering:

1. Ownership of organisational security at board level
2. Proportionate risk assessment including supply chain
3. Lifetime approach: aftercare and incident response
4. Shared organisational responsibility including supply chain
5. System design requirement for strength in depth approach
6. [Software security] [system integrity and safety] managed throughout its lifetime
7. Storage and transmission of data is secure + controlled
8. System resilient to attack + respond appropriately when defences or sensors fail

Again, also, it is noted that the CCAV Code of Practice framework and the proposed future regime for approving advanced automation trials would essentially be laying the potential groundwork for such a regime.

Question 11:

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

It is unclear the extent to which effective interface between automated vehicles and local conditions requires specific working arrangements between a safe assurance scheme and local agencies. At a national level, there is a clear need for automated vehicles to demonstrate that they can operate not just to international driving standards but to UK driving standards and Highway Code. This requirement may, on a technical level, be met at the safety assurance stage of seeking authority for deployment. Depending on the technical standards to which vehicles have been approved, it may be possible for this assessment to be quite light touch – in much the same way that, at a national level, the UK recognises certain foreign driver licences for driving (at least temporarily) in the UK.

The level of co-operation required between authorities and between authorities and automated vehicle developers is probably most important on the level of data.

Automated vehicle developers should know and understand every aspect of their operational design domain and how their vehicles will monitor, detect and respond to local conditions. The system (at least at Level 4) should require the vehicle to fallback to a minimal risk condition should it encounter any anomaly which may not be within its operational design domain. In those circumstances, local agencies can best improve and assist the development and performance of automated vehicles by ensuring that relevant local conditions (fixed and variable) data is made available to automated vehicles which may need to respond to it as and when they need it.

Question 12:

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.

In the context of a new safety assurance scheme it would make sense for the agency tasked with verifying safety and approving deployment onto roads to play a significant role in the safety assurance framework once those vehicles have been deployed. Indeed, and certainly in early phases of deployment of new technology, new vehicles and new models, approval of an automated driving system could be conditional on a period of enhanced monitoring and data feedback.

The safety assurance framework post-deployment is particularly important in the context of automated vehicles since, unlike conventional vehicles and hardware, key aspects of the safety case approved may perform significantly differently in the "real world" than anticipated (either due to simulation error or because the system relies on machine learning capabilities post-deployment to improve) and indeed significant performance and safety aspects may change

relatively easily in terms of planned software updates and unplanned aspects such as changes to external dependencies such as communications services or mapping software. By way of example as to significant changes that could be applied to automated vehicles post-deployment, Tesla has been able through over-the-air updates to significantly alter aspects of vehicle performance (such as radically changing braking performance via a hot fix in 2018) and its "Autopilot" system. Autopilot was itself delivered by software update to vehicles as part of the Tesla 7.0 update in 2015. Consequently automated vehicles post-deployment can quickly become materially different to the vehicle approved by the safety agency pre-deployment. To avoid pre-deployment approvals being largely academic and quickly irrelevant, the safety assurance scheme needs to span both pre and post deployment and it is anticipated that automated vehicle manufacturers would be expected to work closely with the safety agency as regards material changes to the approved safety case of the relevant vehicle. On this aspect, there is relevant learning from other transport modes and modal regulators. In the context of aviation, it is noted that the development and updating of safety-critical avionics software (which would include automated flying function) is highly regulated (at least compared to conventional development of commercial software).

The expertise and knowledge that will be centrally developed by the safety body depends on developing an understanding of the whole lifecycle and ecosystem of the automated vehicles it will be asked to approve and a continuing interaction with the vehicle manufacturer. The safety of automated vehicles should not be regulated in silos and the safety agency which approves the vehicles for use in the first place is best placed to develop that overall specialism.

The Law Commissions have identified a number of post-deployment concerns and the current roles and powers of the likes of DVSA, Advertising Standards Agency and Trading Standards. We do not go so far as suggesting that the automated vehicles safety agency should assume all these roles. However, at the very least, it would need to work closely with these organisations to address the safety concerns raised by the Law Commissions and could support those agencies in the field of automated vehicles through issuing technical guidance, codes of practice, etc as well as intervening directly if necessary (e.g. to relay complaints or make a super-complaint in its own capacity as to consumer issues or safety issues, etc).

The model for the funding of the assurance and regulatory body would also need to be established from the outset. Logically some sort of a levy would be used. However the risk of front loading costs would potentially be to discourage the introduction of systems relative to (lower cost) conventional vehicles.

Whilst traditionally seen as politically difficult, reform of conventional road use charging in the shape of fuel duty and vehicle tax may be needed in parallel. An integrated scheme for road use based around actual use may be needed. However difficult in policy terms if that is not done the UK may see schemes stall with advantage swiftly passing to other jurisdictions and economies with different and more integrated pricing structures.

Question 13:

If there is to be 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.

At this point in time, we do not understand there to be a need to mandate additional driver training for advanced driver assistance systems (ADAS). The more pressing issue, as per Thatcham's concerns, is probably to reinforce to drivers that ADAS is not autonomy. There are concerns that as ADAS gets more advanced and approaches automated driving, this line becomes blurred. Partly this may be the result of manufacturer marketing approaches. Thatcham have flagged this issue specifically with one high capability level 2 system which is being marketed effectively as Level 3 capable when it is not.

However, we additionally anticipate that (understandably) there is not a high level of

understanding in the general population of different automation levels and the key distinctions between them.

As the Law Commissions note, the key differentiator between SAE Level 4 and 3 is that, when activated validly, a Level 4 vehicle should in all circumstances be able to fallback to a minimal risk condition without human intervention – the human is not a necessary part of the safety-critical operation of the vehicle in automated driving mode. From this crucial difference flow very different approaches to human civil and criminal liability, insurance and role and expectation of the human driver.

This lack of public understanding as to basic functionality and regulation or reinforcement of messaging as to higher functionality advanced driver assistance is beginning to raise issues which have not previously been as obvious with lower level assistance features such as 'lane assist' or 'cruise control'. A number of the known SAE Level 2 accidents involving activated ADAS mode appear to involve technology being used by drivers in breach of their terms of use. 2018 saw the first example in the UK of a prosecution for dangerous driving by a driver who activated his Tesla Autopilot and moved into the passenger seat¹⁷. The AAA Foundation for Traffic Safety also commissioned research and reported that there is evidence that human drivers are starting to over-rely on ADAS features without a full appreciation of the functional limitations¹⁸. Any increasing risk of false expectations will increase safety risks.

In the first instance, the emphasis should fall on manufacturers and retailers to market ADAS functionality accurately and transparently, provide clear user advice and giving such information increased prominence. If required - whether to mandate such behaviour or to ensure a level playing field - there may be a case for regulating required consumer messaging. It is not unknown for the UK to legislate for such messaging in other areas particularly if considered to further health and safety objectives or to be otherwise in the public interest e.g. tobacco and alcohol, product labelling, financial products key information, etc.

In terms of additional training available to human drivers of cars with ADAS, for the time being, it is not clear that these need to be mandated. Use of ADAS could however form part of optional advanced driving courses such as those from the Institute of Advanced Motoring or PassPlus. There are some insurers that take such courses into account to offer premium discounts so insurance incentives can be used and doubtless would be used if the data showed clear benefits to reducing claim risks.

Question 14:

We seek views on how accidents involving 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

Burges Salmon was a contributor to the position of the RAC Foundation cited by the Law Commissions in this section of the consultation which we endorse. We refer to our paper produced for the RAC Foundation: "*A Highways Accident Investigation Branch – What Lessons Can Be Learnt from the Rail Industry and the Cullen Inquiry?*"¹⁹

It is important to note firstly that, as in other modes, a Highways Accident Investigation Branch (HAIB) would not be intended to displace the existing police framework for investigating criminal acts. The HAIB would be a parallel investigation process focussed on safety so it is not a case that such a body would be an alternative to the police as the question suggests²⁰. Secondly, in

¹⁷ <https://www.bbc.co.uk/news/uk-england-beds-bucks-herts-43934504>

¹⁸ <https://newsroom.aaa.com/2018/09/drivers-rely-heavily-new-vehicle-safety-technologies/>

¹⁹ https://www.racfoundation.org/wp-content/uploads/HAIB_Burges-Salmon_April_2018.pdf

²⁰ There is a useful analysis of the different roles and policy objectives of criminal and fact finding investigations in a transport context respectively at Chapter 11 of the Ladbroke Grove Rail Inquiry Part 2 Report by Lord Cullen: http://www.railwaysarchive.co.uk/documents/HSE_Lad_Cullen002.pdf

our paper, the analysis of the RAC Foundation and the general debate around the need (or otherwise) of an independent Highways Accident Investigation Branch (like those which exist for rail, aviation and marine) is not restricted to the context of automated vehicles. However, the same principles that we have cited in our paper apply equally if not more so in the case of automated vehicles:

- 1 In complex systems, there is an identifiable safety benefit to an independent non-fault investigation body with statutory powers to require access to evidence, documents and information. The non-fault aspect and the general principle that evidence collected by the AIB would be protected from use in criminal or civil proceedings post-accident may be particularly important where there are concerns over the sharing of data or commercially confidential information;
- 2 The priority is to identify cause neutrally, identify safety lessons and make recommendations to address safety risks. This is particularly important in automated driving systems where the causes of incidents may be system level risks applicable to all automated vehicles or vehicle driving systems of the same type;
- 3 The development of expertise in the investigation of complex system safety incidents and improvement in the efficiency by which lessons are disseminated. The expertise that will be required to undertake effective investigation of such accidents is more likely to be acquired at a central level than with individual police forces for deployment nationally (especially where, for example, there may be likely to be uneven deployment of automated vehicles across the country)

For the same reason, in respect of the criminal investigations framework which the HAIB is not intended to replace, it is considered more likely that prosecution of automated vehicle crimes might be more appropriately dealt with by a specialist national unit of the police. Whether or not that would be formally part of the existing British Transport Police (BTP) is open to question but it is worth noting that the BTP is funded by levy on the rail and light industry that it polices.

- 4 To maximise the potential safety benefits of this system, obligations to report incidents could include "near miss" safety incidents which are identified by automated driving systems. System level learning could therefore be analysed and promoted by the AIB so that its role - like those for other modes - would include a significant accident *prevention* role.

Question 15:

- 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?

In principle, we are in favour of safety agencies having access to all the relevant data they require to fulfil their roles. In our view, this should include (non-personal) accident rate data from which they can assess overall safety lessons and trends. Put another way, we are unaware of any reason why a safety regulatory *should not* monitor such data (or such data could be mandated to be provided to the regulator and/or independent investigation body (see above)). This includes monitoring of ADAS incidents which may in due course evidence the need for regulator intervention (per Question 13 above).

Comparison would undoubtedly be made to existing data collected on accident rates of human drivers. As the Law Commissions note, how this comparison is undertaken and what is considered a tolerably acceptable comparison ratio is a matter for Government.

The RAC Foundation is currently running a pilot for DfT on this issue (relating to conventional

accident data analytics).

Question 16:

- 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?

We would repeat our general position above that automated driving is a conceptually different activity to human driving and that any side by side comparison is inevitably comparing oranges and apples. In addition, in terms of safety management systems, it is important not just to look at outcomes but root causes and the trends associated with those over time.

An analogy would be that in aviation, data is collated and monitored on the performance and capacity of air traffic control systems. It is also collated and monitored on pilot error. The two areas are obviously directly linked (for example if there is a communications failure between the two that is an interface issue).

Clearly societal acceptance will depend upon data being captured and analysed. A large degree of out-performance and designed redundancy of automated systems as against manual driving will in reality be needed. This does however also link in to the treatment of probability and risk within the legal model and the (very) challenging concept of tolerable/acceptable levels of adverse outcomes as part of a balancing exercise. An overly risk-averse system will greatly reduce road capacity and increase congestion; one perceived as creating unacceptable levels of risk to humans will not be societally accepted. There will not (without a frank policy debate) be anything like a level playing field between the currently high levels of risk presented to humans by other humans and the future levels which are system generated.

This is a difficult issue in policy terms. However without addressing it directly it will not be possible to have a viable roll out.

We do not have visibility on the quality of current data. It would be worth speaking however to the RAC Foundation who are looking at existing data quality. We understand that it may – for understandable reasons – be highly variable due to the absence of standardised data capture requirements and the levels of variability in the seriousness of different incidents. A combination of insurance and police data may however provide a fuller picture.

Question 17:

We seek views on whether there is a need for further guidance or clarification on Part 1 of the Automated 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 then need to retain data to deal with insurance claims? If so:
 - 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?
 - How long should that period be?

See the response to Question 1 above as regards clarification on the definition of automated vehicles under Part 1.

In principle there are no pressing concerns with the provisions on contributory negligence applying established legal concepts of causation. However, as we stated as part of the VENTURER final report, in practice we would expect vehicle and vehicle system safety by design to reduce certainly (if not eliminate) the possibility of a human negligently activating an automated driving system when not appropriate to do so or to permit the same to operate without having fully up to date safety-critical software. Given the obvious safety implications of an automated driving system which permits itself to operate otherwise, we anticipate that these aspects would fall to be considered at the approval stage as part of the safety risk assessment.

The work on FLOURISH, highlights the sheer volume of data which may be produced by automated vehicles. However, in principle, it should be possible for the automated vehicle industry and insurance industry to agree data quality and capture standards to ensure that adequate data is available in the event of an accident and the insurance industry is leading efforts to do that through forums such as the EU (as the Law Commissions note). Telematics solutions already exist for fleets for example which are able to supply a degree of incident data when accelerometers are triggered in collisions²¹.

In terms of requirements to notify police and insurers of accidents, there is no reason why current reporting responsibilities in law and under insurance policies should not continue to apply to automated vehicle accidents.

Finally, we note that the CCAV Code of Practice (Section 2 and 5.12 to 5.16) deals with the need for police and other relevant parties to be able to access data in a way that maintains forensic integrity, security and preservation of data. It mentions event data recorders capturing defined categories of data from at least 30 seconds before an incident to 15 seconds after. Consequently if such systems have been developed and are in place at trialling phases, there is no reason to expect that this requirement should be any less in any commercial deployment.

Question 18:

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?

We agree that commonly identified issues (as noted by the Law Commissions) regarding the

²¹ For example, see <https://www.geotab.com/blog/accident-reconstruction/>

application of the Consumer Protection Act 1987 in respect of:

- 1 Pure software products and updates;
- 2 The definition of producer; and
- 3 The “state of the art” defence

do merit a review / clarification given their particular pertinence as regards automated vehicles, their dependence on software and post-supply updates, their machine learning capabilities and the complex supply chains and component interactions.

It is acknowledged that this review may need to be undertaken in conjunction with the approach of supra-national bodies such as the EU or UN given the international nature of the products and their manufacturing supply chain. It is acknowledged that consideration of these issues in the context of the Consumer Protection Act 1987 may have implications going far beyond automated vehicles. Alternatively, there may be a case for considering whether or not automated vehicles should be treated as a special category of product given the safety-critical nature of its primary function to user and members of the public.

See VENTURER Year 3 report.

Question 19:

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?

None we would raise at this stage

Question 20:

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

It would assist in providing certainty to apply a similar exemption to Regulation 107 as that proposed in draft in the USA to cover cars operating in automated driving mode without humans on board. This is, for example, a potential requirement of automated valet parking which is a mooted early use case for limited Level 4 operations.

Question 21:

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

None we would raise at this stage

Question 22:

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?

Yes.

Question 23:

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))

As a matter of principle, we do agree that there needs to be a separate category of offences for a user-in-charge given their specific status and role. What new offences are necessary should be based upon a careful analysis and identification of specific behaviours, why they merit criminalisation and impact assessment.

For reasons stated above in response to Question 3, we have significant reservations as to the example offence cited.

Question 24:

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 body for investigation?
- 3 Where the ADS 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?

In principle, as with human driving, the registered keeper should be obliged to assist in identifying the relevant subject of a criminal investigation into a road traffic offence. It is very possible that, unless the registered keeper was in the vehicle at the relevant time or has direct access to the automated driving system data, the ability of the registered keeper to state with certainty if the vehicle was driving itself at the time may be limited. Consequently, we agree with Footnote 503 of the consultation paper that any obligation should prioritise information, co-operation and assistance to access data in the alternative to stating their belief.

From a technology point of view, we would expect the automated driving system to record

definitively whether or not automated driving was activated or if it was under human control. We note that the CCAV Code of Practice provides that this is the type of data which should be captured on event data recorders even at trialling stage (Section 5.14).

Where the possibility of a criminal prosecution is being investigated and it appears that an automated vehicle driving itself was involved, we consider that the police (whether through a specialist unit or branch or not) should lead on that criminal investigation. If they consider that there is a system level issue involved which may affect safety of the automated vehicle and other such automated vehicles, there should be an obligation on the police additionally to notify the relevant safety regulator (if not already investigating) to consider regulatory investigation and, if appropriate, regulatory measures. The criminal and regulatory regime should operate in parallel and complement each other when it comes to safety-related offences with system-level implications.

Independent accident investigation would operate in parallel. As with other modes a memorandum of understanding would need to be in place between the accident investigation body and the police²².

As well as the option of the police prosecuting for criminal offences, we agree that the automated driving systems regulator should also have safety-related powers to issue improvement / prohibition notices, revoke or suspend approvals and enforce through fines. This is consistent with the powers of sector-specific safety regulators in other modes and of the general safety regulator, the Health & Safety Executive.

Question 25:

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?

In principle we would agree with this, noting that the "user-in-charge" is **not** in fact any "person able to operate the controls" but rather the mandated by law and designated co-driver or standby driver. Such a human in the loop is expected, when it is their 'turn', to assume the dynamic driving task and should be qualified and fit to do so in at least the same way that a normal driver would be. To ensure that this expectation is met if and when the user-in-charge takes control, it would be considered appropriate to criminalise certain behaviour of the user-in-charge even when not driving. The fundamental point is that the "user-in-charge" is not in fact, at any point, a mere passenger. Consequently and much like safety-critical "crew" in other transport modes, their continuing fitness for duty whilst undertaking their role is essential. On aspects such as drink or drugs, the co-pilot of an aircraft is subject to the same expectation of competence and fitness for duty as the pilot, notwithstanding that they may not be the pilot in command – the law treats them the same in this respect and indeed all aircraft crew performing aviation functions²³.

²² See for example: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/383440/ORR-RAIB-BTP-ACPO_MOU_April_2006.pdf

²³ Part 5 of the Railway and Transport Safety Act 2003

The position above as regards competence and fitness for driving would not necessarily be the case for other categories of human in the loop who may interact with the automated driving system but who are not required to assume the dynamic driving task. This consultation does not explore these other particular issues and in large part that is because the various roles that humans may assume "in the loop" are not necessarily fully known at this stage. Suffice to say that where such humans in the loop undertake safety-critical functions, we would recommend (as in the transport sector in general) that such roles should be clearly defined including mandatory qualifications, competence and fitness criteria. Where those roles do not involve driving in the conventional sense, there should be no automatic expectation that the usual driving criteria should apply to them.

Question 26:

Where the 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?

It is unclear what type of behaviour is intended to be criminalised by such an offence.

As noted above, the "user-in-charge" is not defined as any person "able to operate the controls". That would create significant confusion and litigation risk where there may be multiple persons on board such a vehicle with that ability.

If the 'mischief' being targeted by the offence is the risk of confusion as to who is the designated "user-in-charge" we consider that safety by design should mandate that this risk is minimised or eliminated in the system (see generally our VENTURER Year 3 report). By technology or system design, it should be clear from the outset of the operation of such automated vehicles who the designated "user-in-charge" is both to the system and to all passengers and the automated vehicle should automatically not permit activation without a designated user-in-charge. Circumventing this requirement of the system could then be criminalised through other (additional) offences designed to discourage 'hacking' the system.

That person (whether in the car or sited remotely) would then be subject to the obligations and potential liabilities (civil and criminal) of a "user-in-charge" in law and it would not be possible for them to evade that by climbing into the backseat or otherwise trying to disclaim that responsibility. As the Law Commissions note in the consultation paper, any other passengers in the car that are in the vehicle knowing the "user-in-charge" to be in criminal breach of his obligations are subject to potential criminal charges also.

In short, rather than creating a criminal offence to deal with a situation such as users who are all unfit to drive sitting in the back seat, it seems more sensible and safer by design to legislate or regulate appropriately to ensure that the relevant automated vehicle would never operate in the first place until a necessary component of its system (i.e. the user-in-charge) is in place just as it would not be expected to operate in automated mode if it detected that it had no brakes for example.

Question 27:

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?

In respect of insurance, given the provisions of AEVA and, upon commencement, the consequential amendments to be made to the Road Traffic Act 1988, it would seem probable that the term "use" in the context of automated vehicles would be interpreted broadly enough to cover the various users-in-charge who may be designated at various times for automated vehicles

which require one.

We agree that the position is possibly less clear as regards roadworthiness offences.

Assuming that an automated vehicle requires by law a user-in-charge and a user-in-charge has been designated for the relevant journey, we would agree that a failure by that user-in-charge to assume control (where able to do so, do so lawfully and safely and the automated vehicle permits them to) so as to remove a vehicle stopped in a prohibited place should be considered a criminal offence.

Question 28:

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

It is unclear what the human behaviour intended to be incentivised by criminalisation is in this case.

As a matter of navigational technology (such as current GPS navigation systems), it is not commonly possible to set a route deliberately which drives through prohibited places. We would expect this navigational technology directing automated driving systems to also exist in automated vehicles (reinforced by on-board sensor systems and V2X connectivity) and for the automated driving system to be designed not to accept commands contrary to the directives of that navigational software. Consequently if such an occurrence did occur, the cause may more likely be a mapping or other software error (dealt with through product liability and ADSE regulatory regimes) or an unlawful circumvention of the navigational software by the human (dealt with through other offences). At this point therefore, it is not clear that this offence would be necessary.

Question 29:

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 officer or traffic officer; and
- 3 Ensuring that children wear appropriate restraints?

Where an automated vehicle is required by law to have a user-in-charge, there is logic to placing obligations on the user-in-charge for all the above as they would apply to a conventional driver. This observation is subject to the following additional points:

- 1 Any such obligations on a user-in-charge should not absolve the ADSE of parallel (but different) responsibilities in the same instance. Indeed, in such circumstances, one would expect there to be new obligations on the ADSE to also act on data it may have received as to an accident. Telematic systems already exist by which accidents in fleets are centrally registered and acted upon. Furthermore the Law Commissions will be aware of the EU's "eCall" project whereby future connected vehicle systems will be required by EU law to automatically alert the emergency services on registering a serious accident.
- 2 The legislation may need to differentiate between an onboard user-in-charge and a tele-operator user-in-charge. The ways in practice that such user-in-charge may be able to undertake duties following an accident or comply with lawful direction on the scene will be different. That is also the case as regards children wearing appropriate restraints. In practice, the obligation to "ensure" that children wear appropriate restraints may be

limited in the case of a tele-operator to:

- (a) Making on-board announcements
- (b) Checking seatbelt sensors
- (c) Checking on-board cameras

We note that legislation in respect of coaches and certain buses carrying children require only that the driver make a mandatory announcement, audio-visual presentation or signing requesting that passengers use their seat belts.

Question 30:

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

Where an automated vehicle is not required by law to have a user-in-charge, as the Law Commissions note, there is a possibility that such vehicles may be running completely empty. Even if they are not empty but had passengers on board, it is not clear what obligations should be put onto mere passengers in such circumstances. The law does not put such obligations on passengers in other situations.

Fundamentally, it seems likely that these obligations would fall upon a combination of or all of the registered owner, licensed operator or ADSE and may be adapted to suit (see above example regarding the more limited obligation on buses and coaches to ensure seat belts are used). How the obligations are delivered in practice will depend on the available technology and systems. They may potentially involve dedicated humans in the loop (remote or despatched to site) with roles which are not associated with driving as such but for example customer care, liaison with police or providing human input into the overall automated vehicle response. This would be an example of a category of humans in the loop who are not “users-in-charge” in the manner described by the Law Commissions.

Question 31:

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

See above. The law does need reform in this area but it requires more detailed understanding of the systems that are likely to be deployed on automated vehicles with or without users-in-charge.

The legal model should be one centred around a risk based approval of different use cases. The review and approval would analyse risk scenarios. Risk is of course a combination of the severity of a specified bad outcome (hazard) and the probability/likelihood of that outcome occurring.

Use cases and technology will continue to evolve and change (probably in ways not current knowable). The system put in place therefore needs to be future-proof and therefore a system one that looks at each situation diagnostically and then mandates risk controls and mitigations targeted to the specific proposal.

The alternative in theory is to seek to predict myriad future situations and to devise a 'one-size' solution that will address all of those situations.

In practice prescriptive approaches can be successfully used for a specific and non-variable risks. They do have a key part to play in the mix of regulation and legislation. The current public policy debate about a ban on flammable cladding on high rise buildings or the need for specified blood alcohol limits for drivers are examples of that category.

However there needs to be an effective deployment of both (a) over-arching/goal setting/approvals requirements and (b) prescription (blanket provisions). Prescriptive approaches will quickly fail if used in the wrong way – in particular if applied to varying multiple scenarios. The 1972 report of the Robens Committee on Safety and Health at work analyses and is very instructive on this point. A prescriptive requirement that looks valid initially can soon become dated²⁴. Worse it can have an unintended negative effect - harming either safety or efficiency or both.

It is highly likely that an approval process for any currently known CAV use case will mandate very similar requirements in areas like response post-accident and responding to the authorities to achieve required safety and societal outcomes. However the detailed specifics will need to vary solution by solution.

An approvals and licensing process (with the burden on the applicant to show how the need will be met) is therefore a more effective (and in reality probably the only sustainably viable approach).

Question 32:

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.

We agree that there is a:

- 1 Potential gap between human offences of causing death or injury by driving and the Section 22A offence of causing danger to road-users, as regards wrongful interference with automated vehicles causing death or injury
- 2 It is a gap which unlawful act or gross negligence manslaughter or corporate manslaughter could eventually fill but only ever in respect of deaths
- 3 We note that there is an existing and separate aggravated offence already which may apply where the wrongful interference amounts to hacking under the Computer Misuse Act and that causes or creates a significant risk of death or serious injury intentionally or recklessly

Given the potential for interference with an automated vehicle causing death and/or injury, the importance of being seen to punish acts causing death and injury, the need to reinforce the messaging on such dangerous prohibited behaviour and reassuring the public and users, there is a strong case for considering a new criminal offence. As existing offences around causing death and injury by driving are inextricably linked to human drivers and behaviour, there is logic to creating a new freestanding offence.

²⁴ For an extreme example of prescription in legislation not aging well we note in passing that S72 of the Highways Act 1835 - currently under active discussion with the global CAV technology community in the context of pavements and driving - is structured as follows "If any person shall wilfully ride upon any footpath or causeway by the side of any road made or set apart for the use or accommodation of foot passengers; or shall wilfully lead or drive any horse, ass, sheep, mule, swine, or cattle or carriage of any description, or any truck or sledge, upon any such footpath or causeway; or shall tether any horse, ass, mule, swine, or cattle, on any highway, so as to suffer or permit the tethered animal to be thereon..."

Question 33:

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

In addition to the offences identified within Chapter 7 of the Consultation, it is key in this context to review the applicability of the Health & Safety at Work etc Act 1974 ('HSWA') and the regulations made pursuant to it both for its use (and potential unintended consequence).

There are general offences under Sections 2-6. Sections 2 and 3 in particular impose an overarching obligation upon any employer (incorporated or non-incorporated) to control risks generated by the undertaking of the employer so far as it is reasonably practicable to do so²⁵. S2 relates to risks to employees and S3 to risks to others.

The offences apply to any activity by an employer within its own undertaking. The potential application to the widest range of activities is therefore considerable. This was not the original intent of the legislation²⁶. In practice, through regulatory resource constraints, the enforcement of the HSWA across different activities outside the workplace context varies. Heightened public sensitivity around an area or event can however lead to it being used – on a one off or ongoing basis - as a 'catch-all' charge for a range of new situations relating to risk²⁷ not within its original legislative contemplation.

The HSWA regime underpins UK safety legislation in heavy and light rail transport systems and some aspects of bus operations. Aviation safety is less centred on it although the HSWA regime applies to many aspects of (for example) activities at airports.

Uniquely a reverse criminal burden of proof applies to those offences. The organisation charged with a criminal offence under either S2 or S3 must prove at trial that it has done everything reasonably practicable to reduce risks to a level as low as reasonably practicable. That results in a current conviction rate of over 90% for HSWA prosecutions.

HSWA offences permit unlimited fines for corporate defendants and imprisonment for individuals convicted. Fines have recently been significantly increased by formal Sentencing Guidelines that the courts are mandated to use. HSWA offences now regularly result in large six and seven figure fines.

The reasonable practicability test is not defined in the HSWA or other statute. Regulatory guidance²⁸ from the HSE²⁹ (adopted formally as the position of the UK in an infraction action taken to the ECJ by the European Commission on the point)³⁰ has advanced a bespoke interpretation of the phrase based upon a sentence in a 1949' historic Court of Appeal Judgment (*Edwards v NCB*)³¹. That interpretation states that the criminal law requires that the defendant must show that it has taken every action to reduce risk except for actions which are

²⁵ This is referred to as the 'SFAIRP' or 'ALARP' (as low as reasonably practicable) requirement.

²⁶ See the Robens report giving rise to the legislation at Chapter 5

²⁷ For example: <http://news.bbc.co.uk/1/hi/england/merseyside/7129004.stm>

²⁸ See for example: <http://www.hse.gov.uk/risk/theory/alarplance.htm>

²⁹ The formulation was first created by HSE in the context of a nuclear power station new build (Sizewell B) so its initial wish in that context to build in extreme caution in formulating its construct is perhaps understandable. The same formulation has however been rolled out in further guidance to apply to every area of activity that is directly or indirectly work-related

³⁰ Commission of the European Communities v United Kingdom, ECJ Case C-127/05, on 14th June 2007, reported at [2007] ICR 1393 (also at [2007] IRLR 720)

³¹ *Edwards v National Coal Board* [1949] 1 All ER 743 (CA)

grossly disproportionate in terms of the 'time trouble or expense' involved to reduce the risk.

Other and subsequent cases including House of Lords and Supreme Court authorities³² talk of a proportionality test rather than a gross disproportion test. However, despite debate and some controversy on the point for the reasons below, the 'gross disproportion' sentence still forms the basis of the issued regulatory guidance on the meaning of reasonable practicability

However the adoption of the word '*grossly*' of course raises difficult issues with direct legal and operational consequences for those responsible for any system or activity. In particular:

- Whether actions which are disproportionate (but not *grossly* disproportionate)³³ to the risk involved are (or are not) legally required by the criminal law;
- If disproportionate actions are required by law then an action can then be non-negligent in civil law (and so for insurance purposes) but still be a criminal act giving rise to imprisonment or very substantial fine.
- Regulatory guidance requires the probability of an adverse event to inform risk assessment. However at the point of enforcement, prosecution and trial, there is little - if any - weight given to probability. There tends to be a reversion to a more 'hazard-oriented' mindset and approach³⁴. That approach tends to be hindsight based.

The HSWA is not habitually enforced in the context of driving itself. The jurisdiction to do so exists, for example in relation to fleet condition or employee driver fatigue issues. There are occasional prosecutions in egregious vehicle maintenance and/or driving cases.

The reason for raising the above in the context of CAVs and the response to Consultation Question 33 is:

- If the HSWA legislation and regulations regime is to apply then on one level it already provides the criminal law recourse contemplated.
- If its use is contemplated however, the scale of future CAV systems and operations raise fundamental issues of practicality, scale and application which will need to be consciously addressed if harmful unintended outcomes are to be avoided. See below.
- Concerns would arise in particular around;
- The unavoidable use of probability factors in the design and use of CAV systems. To exist these systems will in fact need a legal framework which explicitly recognises and respects not only 'conventional' probabilistic risk assessment but also Bayesian/algorithmic decision making which is of a much higher order of complexity. For the historic and legacy reasons above relating to its evolution, the application and enforcement of the HSWA comes from a very different place. Designing and operating CAV systems by necessity is a pre-event and risk-based process; enforcing the HSWA is a post-event (hindsight) and hazard-oriented process.
- If that critical dichotomy is not thought through then (statistically inevitable) accidents are likely to lead to prosecutions that will not distinguish between very different levels of fault.
- A perverse situation could - for example - arise where overall fatalities drop by a very significant percentage but that beneficial trend in road safety outcomes is accompanied by a sharp increase in the number of HSWA prosecutions and fines in relation to road use.
- On a practical level how and by whom would HSWA offences be investigated and

³² Including *Marshall v Gotham Co Ltd* [1954] AC 300, HL and *Baker v Quantum Clothing Group Limited and others* [2011] UKSC 17

³³ The conscious distinction between disproportionality and gross disproportionality is used elsewhere in statute – see Criminal Justice Act s76(5A)

³⁴ See for example *R. v Board of Trustees of the Science Museum* [1993] 3 All ER 853

enforced. There will be implications for the structure and resourcing of HSE, Police and prosecutors.

- The Corporate Manslaughter and Homicide Act 2007 is predicated on gross breach of duty and therefore has a high threshold of culpability built in. However there is a direct linkage within Section 8 (Factors for the Jury). Section 8(2) of that Act requires that
- *“The jury must consider whether the evidence shows that the organisation failed to comply with any health and safety legislation that relates to the alleged breach, and if so—*
- *(a) how serious that failure was;*
- *.(b) how much of a risk of death it posed”*
- *HSWA offences are often prosecuted in parallel. The reverse burden and 'gross disproportion test' makes HSWA defences difficult. This is clearly rather important in the context of S8(2) of the 2007 Act.*
- There clearly needs to be a criminally enforceable regime that sanctions (potentially seriously on a scale proportionate to the degree of fault) behaviours by organisations falling below a defined standard in the design and operation.
- In the context of CAVs there are already calls for punitive response. Many of those calls are however simplistic. It would be tempting to see the HSWA regime as an existing sledgehammer which can be deployed to crack the perceived nut of *'holding big organisations to account'*. However for the reasons above such a mindset is likely both to impair seriously the effective use of CAV systems and their anticipated societal benefits (including reduced fatality and serious injury rates) and be unfair to those directly involved, acting responsibly who have to programme in a degree of risk to create any functional system.
- The actual offences put into place to ensure proper accountability therefore need to reflect how systems operate and what behaviours are required and expected in their design and use. The current law does not do that in practice.

Question 34:

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?

See above in respect of the offence of interfering with automated vehicles so as to cause death or injury.

In other respects we agree with the Law Commissions that existing criminal offences probably would be interpreted to include similar offences that could be committed regarding automated vehicles. In other words, many of the interference-type activities in respect of automated vehicles would already be criminal offences and should be obviously dangerous.

However, even at trial stages of this technology, there have been reports from the various Innovate UK funded CAV projects of persons deliberately seeking to interfere with automated trial vehicles in a way which they would plainly not otherwise contemplate in normal traffic. It is unclear why this is occurring but on one level it suggests that the objective obvious danger in doing so is potentially not as obvious as it should be. Given the novelty of automated vehicles and need to reinforce public messaging around safety, there is a case for being absolutely clear

on prohibition and criminal liability. Explicitly referencing interference with automated vehicles and automated driving systems in relevant legislation might be part of the solution.

Question 35:

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?

Yes: clarity on this point would be preferable given the novelty of automated vehicles and the need to reinforce public messaging around safety.

Question 36:

In England and Wales, section 12 of the Theft Act 1988 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 1988 be extended to any motor vehicle, even those without driving seats?

Given that "joyriding" of an automated vehicle which is not adapted for carrying persons would suggest a software hack to remotely control the vehicle or to cause the vehicle to drive itself, it is unclear whether a new offence is necessary if it is likely to be covered by the aggravated offence under Section 3ZA of the Computer Misuse Act.

If avoidance of doubt is needed then it could potentially be provided by amendment of that provision.

Question 37:

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?

For reasons discussed above, there may be an argument for ensuring that the legal position is consistent on this aspect (together with any clarifications to reinforce the "obvious" danger of interfering with such vehicles or traffic equipment as regards automated vehicles).

Question 38:

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

We are unable to comment on the technical requirements around this, however we recognise that public authorities and bodies (such as CCAV and Meridian) and government-supported bodies (such as Transport Systems Catapult and BSI) are already working with the private sector to develop standards and principles for automated driving. There are already clear channels, it appears to us, under which the necessary collaboration can be delivered. Any coded "Digital Highway Code" or formal set of "exception-handling rules" will need to be tested extensively in simulation and at least edge case testing on the road.

Question 39:

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?

Notwithstanding current legislation and rules on mounting the pavement, it is clear that there are occasions in which vehicles are permitted or expected to be able to mount the pavement. In principle, it would not be desirable to prevent an automated vehicle from doing anything which an ordinary vehicle under the control of a human driver can and may be expected or required to do.

Per the Law Commissions' consultation paper and the examples above, doing so is often permitted in circumstances of lawful authority or necessity. In road traffic legislation often the term "reasonable cause" similarly may apply for exceptions. However, these simple words mask in practice very complicated decision-making processes.

These exceptions and the resulting undertaking of the exceptional driving task are difficult to define through fixed and fast rules and algorithms. It is unclear at this stage the extent to which mounting a pavement is capable of being within an automated vehicle's operational design domain, its minimal risk condition fallback or, even if permitted to do so, whether or not the decision to do so can be made by the automated driving system where not a genuine emergency in which the safety calculation is very clear (e.g. avoiding an unexpected collision by mounting a pavement known to be clear in a rural setting). In other scenarios, the decision and reaction may need to be supplemented or directed by a human in the loop (whether or not it actually requires the human to assume the (exceptional) dynamic driving task).

We refer to our response to Q33 above also on the need for debate and clarity on how the law defines responsible behaviours around probability and algorithmic decision making - and also sanctions behaviours that do not fall within that definition.

We note in passing that Section 72 of the Highways Act 1835 relating to pavements is antiquated and clearly needs to be updated or replaced. It is not a credible legislative position to seek to engage with a global technology market on an important policy debate (the use or non-use of 21st Century technology on footways) by reference to a two century old provision expressed in very archaic language. We appreciate however that this has currently been excluded from the Commissions' terms of reference.

Question 40:

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

Per above, in principle, it would not be desirable to prevent an automated vehicle from doing anything which an ordinary vehicle under the control of a human driver can and may be expected or required to do. There is however clearly a need to understand if and how any particular automated vehicle would take such a decision and execute it and whether in whole or part reliance will be placed on a human in the loop.

Question 41:

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

The position above as regards mounting of pavements is largely repeated for speeding within accepted tolerances. However, save for safety reasons, we query whether or not, as a matter of law, speeding should be made permissive (in effect through automated vehicle approval) for automated vehicles for any other reason including purely convenience and non-safety related reasons. Within this context, we query whether or not preventing sharp braking at speed down-changes as described in the consultation paper should be considered permissible given that an automated driving system ought to have the necessary data available to it to have prior warning of such speed limit changes and adjust accordingly on approach.

Question 42:

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?

The position above as regards mounting of pavements is largely repeated for edging through pedestrians. However, we would suggest that some of issues in fact take on a heightened sensitivity and complexity because of the clear and present risk of injury to humans and the very limited circumstances when even human drivers might be permitted to edge through pedestrians.

Criminalising behaviour which is intended to impede automated vehicles on public roads and places is one part of reducing the occurrence risk of this (see above). Under no circumstances, should the law permit such events to become common occurrences.

However, when faced with a decision as to whether to 'edge' through pedestrians or not, we agree with the Law Commissions that the nuances are extremely difficult to pre-program and are inherently risky even for humans to judge. Again, in practice, the decision and reaction may need to be supplemented or directed by a human in the loop (whether or not it actually requires the human to assume the (exceptional) dynamic driving task).

The issue presents an unavoidable need to engage with the concepts of probability and intent in the criminal law and the behaviours that are to be permitted or, alternatively criminalised.

The concepts of probability and intent can work in a purely human decision making context through the deployment of the objective principles of negligence and subjective principles of recklessness or conscious act. Those are familiar and well-honed approaches.

Those same concepts do not translate to system design and operation. Attempting to legislate using them outside their parameters will not be either effective or fair. Different and more precise definitions and concepts are needed. AI and algorithmic decision making is taking the law into a new area where clean sheet thinking is needed and not the tempting - but ultimately illogical - application of traditional approaches.

Any functionality which permits an automated vehicle to edge through pedestrians (on its own initiative or under guidance of a human in the loop) must be designed to be performed in as safe a manner as possible (e.g. we would expect the ADSE to demonstrate a safe package of mitigations to accompany the functionality, not just in decision-making to undertake the manoeuvre but in its implementation including audio-visual and other warnings to members of the public of the manoeuvre such as those commonly seen of reversing industrial or commercial vehicles).

Question 43:

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?

We are aware of some of the commonly raised risks around this area as described also by the Law Commissions. We are unable to comment on some of highly technical work being conducted on this area (which is relevant generally to the whole area of data and artificial intelligence).

If risks of data bias in automated driving systems are confirmed and those risks have negative safety-related consequences, as a matter of process, it would appear to us that those relevant datasets and/or data processes ought to be at least third party tested / certified / audited at a high level to provide assurance according to recognised good practice or industry standards since they will be relevant to the safety approval of the vehicle. Generally and to the extent possible, we consider that the more that this difficult area can be brought closer to one of compliance with pre-established and recognised norms as opposed to individual bespoke assessment, the more efficient the process will be.

Question 44:

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

Ethics represent a system of moral principles but can vary greatly from person to person, organisation to organisation or society to society. We consider that the interests of the public, developers, road users and members of the public may be best served by addressing the issue of ethical frameworks at societal level, part of which is of course the relevant legal framework.

As above for data bias, this would involve a national level consideration and settlement of the key issues. From that point onwards, the issue can be more one of compliance and verifying compliance with the relevant standard or law. We anticipate in particular that this is what road users and the public would expect and is therefore crucial to public acceptance of the technology. We do not perceive any appetite to have a choice of automated vehicles differentiated by, amongst other things, their ethical frameworks.

We note that in 2017, the Ethics Commission on Automated and Connected Driving, a body of experts established by the German Federal Minister of Transport and Digital Infrastructure, issued a 20 point set of ethical rules for automated and connected vehicular traffic³⁵. These core guidelines are simple to understand, transparent and, importantly, settled at a level which maximises public acceptance and trust. From developers' points of view, it provides a clear set of guidelines against which they can verify (and if necessary have independently certified) its own framework as compliant. The requirement on developers then becomes one of certifying compliance rather than publication (although of course they may also choose to do that). Regardless of whether or not the same guidelines are agreeable in the context of the UK, we consider that the approach is a commendable one.

³⁵ https://www.bmvi.de/SharedDocs/EN/publications/report-ethics-commission-automated-and-connected-driving.pdf?__blob=publicationFile

Question 45:

What other information should be made available?

We have in our previous reports (VENTURER and FLOURISH in particular) made the case in conjunction with our partners and in particular AXA in respect of the making available of certain categories of data likely to be relevance in the event of accidents. We are aware that discussions on this aspect are ongoing internationally involving both industry stakeholders and insurers.

Care is needed to distinguish the different categories of data and the ownership use of each.

We note from the CCAV Code of Practice that CCAV encourages trialling organisations to share publishable forms of their safety case to engender greater public acceptance of the technology. A degree of safety reporting even once such vehicles reach the market could arguably continue to build public confidence and acceptance.

Question 46:

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

The Law Commissions have helpfully flagged within their preliminary consultation paper a number of areas that they intend to consult on in future (mobility as a service, etc) and we look forward to contributing to those where we can helpfully do so.

We note that the Law Commissions, generally, consider the aspect of automated vehicle data to be outside of the remit of this review project. In part, this may recognise that there are significant industry steps already being explored with authorities (at national and supra-national level) to agree protocols for data capture, protection, cyber-security, handling and processing (which we are also contributing to). For what it is worth, we consider that putting in place an adequate and effective framework for automated vehicle data is an absolutely key part of the viability, case and public acceptance of automated vehicles.