

Law Commission:

Automated Vehicles CP3

Comments provided by Five AI Limited

22 March 2021

CHAPTER 4: SELF-DRIVING AND HUMAN INTERVENTION

Consultation Question 1 (Paragraph 4.114)

We provisionally propose that:

- (1) a vehicle should not be classified as self-driving if, with the ADS engaged, the user-in-charge needs to monitor the driving environment, the vehicle or the way it drives;
- (2) it is nevertheless compatible with self-driving to require the user-in-charge to respond to a clear and timely transition demand which:
 - (a) cuts out any non-driving related screen use;
 - (b) provides clear visual, audio and haptic signals; and
 - (c) gives sufficient time to gain situational awareness;
- (3) to be classified as self-driving, the vehicle must be safe enough even if the human user does not intervene in response to any event except a clear and timely transition demand.

Do you agree?

Yes.

Agreed with the qualification that, in the situation described in (2) (where an ADS is designed with the requirement that a user-in-charge respond to a clear and timely transition demand), for the vehicle to be classified as “self-driving” it should also be a requirement that the vehicle be ‘safe enough’ even if:

- (a) the user-in-charge does not intervene in response to a timely transition demand; or
- (b) it is not possible for the vehicle to issue a ‘timely’ transition demand.

In other words, to be classified as “self-driving”, the ADS should either be able to carry out an appropriate minimal risk manoeuvre or an appropriate failure mitigation manoeuvre in response to circumstances where the ADS is designed to request that the user-in-charge intervenes, but the user-in-charge has not yet intervened. This is to ensure that the vehicle remains sufficiently safe until the user-in-charge intervenes, and to take into account the possibility that the user-in-charge may not intervene. If the vehicle is not able to respond in a sufficiently safe way in such circumstances, it would be necessary for the user-in-charge to monitor the driving environment for the vehicle to be ‘safe enough’, and hence the vehicle would fall outside the definition of “self-driving” under (1).

For legal clarity a vehicle should either meet the definition or not; having a situation where the vehicle could fall in between would lead to uncertainty.

We agree that the transition demand must be clear and timely. For the reasons discussed in CP3 paragraphs 4.83 - 4.93, we consider what criteria the transition demand must meet to be sufficiently clear and timely would merit further study in order to set an evidence based benchmark. It seems possible that the appropriate criteria may vary depending on the nature of the environment (e.g. more complicated traffic situations, difficult weather conditions). We also agree that the effectiveness of the transition demand should be monitored in practice. To retain flexibility, we suggest the criteria that the transition demand must meet in order to be clear and timely is not separately elucidated within the definition of what is required for a vehicle to be classified as 'self-driving', but instead either is dealt with by way of separate guidance and/or forms part of the assessment of whether or not the vehicle is 'safe enough'.

The requirement at 2(a) appears to contain the inherent assumption that the user-in-charge would only be permitted to use a screen that is controlled by the vehicle. If the vehicle does not control screen use, it would not be possible for the transition demand to comply with that criterion (or alternatively that criterion would need to be limited to cutting out only screen use controlled by the vehicle). Accordingly, it would not be possible to rely on this being an essential part of the transition demand.

Consultation Question 2 (Paragraph 4.115)

We welcome views on whether self-driving features should be designed to ensure that they can be used by people with hearing loss.

We agree that self-driving features should be designed to be accessible to all in so far as practically possible. It will need to be established what combination of features in a transition demand is sufficient to ensure the user-in-charge registers the transition demand; whether visual and/or haptic cues are sufficient.

CHAPTER 5: HOW SAFE IS SAFE ENOUGH?

Consultation Question 3 (Paragraph 5.118)

We provisionally propose that the decision whether a vehicle is sufficiently safe to “safely drive itself” should be made by the Secretary of State, as informed by advice from a specialist regulator.

Do you agree?

Other.

Given the potential considerable benefits of AV when harnessed correctly, we consider it important that the decision making process that determines whether or not a vehicle (i.e. the specific hardware and software in a defined ODD) is sufficiently safe to “safely drive itself” be an evidence-based decision making process, and that the regime that is established be one that encourages AV manufacturers to bring AV systems to the UK market, whilst ensuring high standards of safety.

A regulatory regime ought to be stable, clear, transparent, consistent and fair, among other matters. There is a distinction between a decision as to whether a particular vehicle meets the requisite standard of safety in relation to the specific characteristics of the deployment proposed (for example, motorway driving in fine weather), and the decision as to what standard of safety must be met in general in relation to a deployment of that type (for example, motorway driving in fine weather). Given the nature of the technology, it may not be practicable to separate the two. However, the first decision by its nature is primarily a technical one (by which we include assessing the processes and procedures specified in the safety case to determine whether, taken together with the technical capabilities and limitations of the vehicle, the vehicle, together with the processes around it, is sufficiently safe to “safely drive itself”). In making this decision, there is limited scope for other considerations, such as policy matters, to legitimately come into play (an example of a legitimate exception to this would be a national security consideration). In contrast, while the decision on the general standard of safety to be met should also be largely evidence driven, it may be legitimate to include other considerations, such as policy considerations. If the Secretary of State is the final decision maker, we consider it would be beneficial to clearly set out the factors that the Secretary of State can take into account in making the decision, their different weighting, and the extent to which the Secretary of State has discretion, and that there be the potential for an independent and impartial review of the decision. Consideration should be given to the scope and format of evidence required from the appointed regulator(s), as a minimum, in order to provide advice to the Secretary of State.

An assessment of safety requires specialist knowledge. Regardless of who makes the decision, we agree it is essential that the decision be informed by the advice from a specialist regulator of sufficient technical competence. The appointment of one or more ‘specialist regulators’ for this role should include measures to ensure that such regulators are sufficiently independent and competent to provide the required advice.

Consultation Question 4 (Paragraph 5.119)

We welcome observations on which of the following standards is most appropriate when assessing the safety of automated vehicles:

- (a) as safe as a competent and careful human driver;**
- (b) as safe as a human driver who does not cause a fault accident;**
- (c) overall, safer than the average human driver.**

None of the above.

The regulation of the deployment of AVs will require a paradigm shift; there is no existing system approval process that accommodates the nature of the technology.

In our experience developing an AV system that was used in public AV trials on London roads (the Streetwise project), the development of an AV system is an evolving process, where one incrementally improves the operation of the system and gains increasing confidence in it over time. Hence, we would recommend establishing a defined minimum standard of safety that is practically achievable at the outset, with the expectation that the standard of safety achieved by an AV will, and should, increase over time.

Ultimately, the aim should be to set clear blue water between the performance of an AV and the performance of a human. However, if too high a standard is set at the outset it will present an unachievable hurdle to obtain (and prove) pre-deployment. This would deny the public the productivity, transport cost and convenience gains AVs bring, and could restrict AV deployment to a limited number of domains and/or companies, reducing the benefit derived and market competition. In some cases the overall value that will be achieved through the deployment, including the ability both to improve the safety of the AV by gaining more real world experience and to obtain more evidence of the safety of the AV, may justify accepting a lower standard of evidence of safety and/or lower minimum threshold of safety at the outset that increases over the lifetime of the AV. Accordingly, why we suggest the aim should be to strike the right balance between ensuring safe deployment, whilst not imposing unreasonable obligations on AV developers that will impede, or even forestall, the development and deployment of AV technologies.

In practice, we do not consider it will be possible to make as clear cut a deployment situation as the consultation paper envisages in every case. In particular, for Level 4 systems designed to be used in a more complex ODDs, we consider the deployment process would need to be more incremental, where risks will partly be managed through limited deployment (by which we include limited by ODD, not just in numbers of AVs) and partly by putting in place processes to improve the safety of the AV over time and incrementally expand the ODD. Indeed, allowing an initially more limited deployment is a vital part of being able to gain sufficient real world experience in order to be able to improve, and gain more evidence to prove, the level of safety met by the AV.

We suggest the standard of safety be set at an appropriate tolerable level of safety taking into account the limits set around the deployment (for example, ‘as safe as reasonably practicable’ as proposed in paragraph CP3 paragraph 5.103), and that there be a target level of safety to be obtained through operation, akin to the through life continuous safety improvement process followed by NASA. As we discuss further in our response to consultation question 5, as part of the safety approval process we suggest the ADSE should present the details of the processes it will follow to ensure continual improvement of the ADS to meet that target safety level. The regulator could then assess whether or not it considered these proposals sufficient pre-deployment, and conduct a yearly performance analysis following deployment. An advantage of having a clear pre-agreed safety improvement process to follow is that it is more straightforward for the regulator to audit whether or not the safety improvement process is being followed than to determine an absolute safety level at any one point in time.

We do not agree with the suggestion in CP3 paragraph 5.103 of using a blend of the four tests; this would not be practicable. From an ADSE’s perspective it needs to know what it is aiming for, and what evidence it needs to collate to prove that. From a regulator’s point of view, it too needs to know what standard it is judging the ADS against and when the ADS has reached that standard. If more than one standard is used, what should the regulator’s decision be if the ADS satisfies one standard, but not another? Using a blend of tests could lead to inconsistent decisions and a loss of trust in the system.

Out of the three options suggested, setting the target AV safety level at the level of a competent and careful driver (at least) would seem the most appropriate target to aim for.

However, an AV is not a human and it is possible that the performance of an AV will exceed that of a competent and careful driver in some areas and not meet that of a competent and careful driver in others. If, overall, the areas in which an AV exceeds the performance of a competent and careful driver are such that overall, there is a net benefit from the deployment of the AV (and one that does not change the distribution of risk inequitably) it would seem beneficial to permit an AV's deployment. Accordingly, a better target standard might be "overall, as safe as a competent and careful driver".

The standard must also be one that can be tested by a relevant party at an appropriate level of abstraction using a consistent set of test parameters and a searching sample of test points. How objective measures of the standard can be established requires further consideration.

For example, ALKS provides a model of skilled human driving that can be used to establish an objective minimum performance threshold for the safety critical manoeuvres that are required by the ALKS regulation. This is based on data to characterise emergency braking performance by human drivers, human perception times and vehicle behaviours that humans associate with being precursors to particular manoeuvres such as a cut in. In this limited ODD, it is possible to set objective standards to represent a 'competent and careful' human driver and test explicitly whether an ADS is able to match or exceed these performance thresholds.

As the ODD complexity increases (especially with the inclusion of vulnerable road users and more varied traffic manoeuvres) it becomes more difficult to identify evidence-based objective performance boundaries that represent competent and careful human drivers. For example, traffic data is not captured at a granular enough level to determine how safe a human driver is in fine weather in London, or the relative safety level of motorways compared to urban streets. Discussions around using a positive risk balance approach suggest that collision rates for human drivers would be used as a safety benchmark: what remains unclear is what collision rate would be used to benchmark 'careful and competent' drivers and, given that injury collisions are thankfully extremely rare, how an ADS can meet such a standard when the necessary exposure to generate such collision data will not yet exist.

There are clearly some associations to be made with the approaches discussed in Chapter 10 to establish the methods by which such a standard can be assessed when the data are not yet available.

It is also noted within the consultation document that the errors of a human driver (even a competent and careful human driver) that are documented as contributory factors in collisions are not necessarily the same errors that will manifest in AVs. Therefore any reliance on existing collision statistics associated with human driving will not provide the basis for an appropriate safety benchmark for AVs.

The practical challenges of evaluating whether the performance of an AV meets the relevant standard pre-deployment are non-trivial, witness the extensive SIL, HIL, VIL ADAS feature testing conducted by OEMs before they are confident to accept the product liability associated with those features. We would recommend a subset of such measures will need

to be defined if we are to avoid public exposure to an inappropriate level of risk by insufficient standards being set, versus over-setting the standards and curtailing agile feature development and testing.

Overall, the standard of safety would seem to merit further consultation.

Consultation Question 5 (Paragraph 5.120)

We welcome observations on how automated vehicles can be made as safe as reasonably practicable.

Both physical and simulation training and testing are essential in developing an ADS and a regulator should require an ADSE to adopt and evidence a mixture of both.

Given the lack of real world collision data that will exist for new deployments of automated driving features/vehicles, it would be valuable to consider the approach used in other industries (e.g. the European railway industry) of requiring operators to implement a Safety Management System (SMS) that has explicit structures in place to monitor performance and deliver continuous improvements to safety. Regulators would be able to use such a model to 'approve' an ADSE based on a conformity assessment of its SMS followed by certification and regular audits to ensure the SMS was being used appropriately. This helps create a uniform structure across ADSEs within which it should be possible to ensure that AV deployments are operated in a manner that is 'as safe as reasonably practicable', with the emphasis placed on continuous monitoring of safety levels with continuous improvements being the target.

Information sharing is also important; if an issue is identified that information should be shared with all AV developers and all ADSEs obliged to check for that problem within their own AVs.

Since, as a matter of practicality, it seems unlikely that a regulator will be able to obtain the same depth of technical expertise or understanding of an AV system, nor have equivalent access to sensor physics models, HiL test rigs, test suites and parameters, behavioural models, test oracles and analytic tools as the ADSE who developed it, imposing some form of self-reporting obligations may be appropriate (as in the aerospace industry).

Consultation Question 6 (Paragraph 5.121)

We welcome practical suggestions for how AV regulators can fulfil their public sector equality duty.

In fulfilling the public sector equality duty, the focus should be on the behaviour of the whole system and striving for equality of outcome. I.e. regulators should not be mandating that the achieved precision (measure of false positives) - recall (measure of false negatives) for a pedestrian detector is unbiased, but that the ultimate system behaviour is unbiased.

A minimum scenario catalogue, rule set and minimum level of tests to identify unacceptable bias before deployment should be created. In doing so, consideration should be given to what characteristics of the relevant groups could potentially lead to unintentional bias (for

example, skin colour, size, certain forms of attire) and these characteristics investigated to determine which might meaningfully influence risk, in order to direct resources appropriately.

Consideration should also be given to the extent training on any particular dataset may affect the eventual behaviour of the system in relation to the distribution of risks amongst protected groups, and also other relevant groups. Fundamentally, an AV must meet the required minimum level of safety when considering all relevant characteristics that may have an effect on safety. These characteristics may reflect protected characteristics and characteristics that are not (for example, vehicle colour). A truly diverse dataset should be representative of all relevant characteristics. But equally, it must not skew the behaviour of the system in a way that is unacceptable, and to achieve unbiased system behaviour it may be necessary to train a system with a 'biased' dataset.

Particular care should be taken to ensure detection and classification of vulnerable road users, motorcyclists and cyclists so that certain fast moving or difficult to detect road users are not systematically poorly detected or their pose, motion or intention systematically subject to high error.

In addition to ensuring AV's are presented with a sufficient diversity of training data, consideration should be given to how it is possible to identify any unintended bias in practice.

Finally, it is important to establish or adopt a consistent ontology within the industry. Such ontologies are often hierarchically structured and an ADSE will need to determine the correct level of abstraction for its ADS system. If the ADS may need to perceive, plan, or act differently for e.g. adults and children, then these categories need to be distinguished as separate sub-classes in the ontology. However, whether or not a particular ADS needs to distinguish between the two in practice will depend on its design (the ADS may be designed such that it expects the worst case and the ADSE can establish that the ADS does not need to distinguish between the two in order for the system to be sufficiently safe).

CHAPTER 7: ASSESSING SAFETY PRE-DEPLOYMENT

Consultation Question 7 (Paragraph 7.99)

We provisionally propose that:

- (1) safety assessment should use a variety of techniques;**
- (2) manufacturers/developers should submit a safety case to regulators showing why they believe that the automated driving system is safe;**
- (3) regulators should:**
 - (a) provide guidelines for what is in the safety case;**
 - (b) audit the safety case;**
 - (c) prepare guidance for manufacturers and developers on preferred standards; and**
 - (d) carry out at least some independent tests.**

Do you agree?

Yes.

Agreed. In relation to (1), the safety assessment should include evidence from both physical testing and testing in simulation in conjunction with the safety case, but not be prescriptive as to the mix.

Consultation Question 8 (Paragraph 7.100)

We seek views on whether an approval authority that intends to use a scenario database as part of the testing procedure should consult road user groups on the range of scenarios to be included.

Scenario-based testing should ensure that the scenario space is adequately covered for the specified ODD. If the ODD includes exposure to specific road user groups, then consultation may help elicit examples of situations that should be incorporated within scenarios. However, the decision to include scenarios in a database should be based on risk in the intended domain (rather than the lobbying of road user groups) and what is realistically achievable and meaningful from a technical perspective (e.g. requiring an AV to detect if a cyclist had turned their head when coming up to a turn off would be unreasonably onerous technically at the present stage of AV technology development, and may not be a useful indicator that a cyclist is about to turn).

In addition, the scenarios in the database and how the many parameters in them might be sampled in combination will be valuable intellectual property, and if open access is given it could encourage some ADSEs to pre-configure their ADS to ensure tests are passed, even if the ADS would exhibit unsafe behaviour outside the specific test case, for instance. For the purposes of establishing public trust and obtaining useful feedback from road user groups it should be sufficient to publish the scope of the scenario classes covered at a meaningful level of abstraction (e.g. “cyclist making a right turn at a T-junction”).

CHAPTER 8: INITIAL APPROVALS AND CATEGORISATION – PROPOSALS

Consultation Question 9 (Paragraph 8.17)

We provisionally propose that:

- (1) unauthorised automated driving systems should be prohibited; and**
- (2) this should be subject to an exemption procedure by which the Secretary of State may authorise unauthorised systems to be used in tests and trials.**

Do you agree?

Yes.

Agree in principle. It is assumed that the intention is to retain the current Code of Practice in relation to testing carried out with safety drivers. In relation to (2), guidelines should be produced so that the information that needs to be provided, the requirements that need to be met, and the process to be followed to obtain authorisation are clear, fair and consistent.

Consultation Question 10 (Paragraph 8.25)

We provisionally propose that:

- (1) the Government should establish a domestic scheme to approve automated driving systems (ADSs) for use on roads in Great Britain (a “national ADS approval scheme”);**

- (2) manufacturers should have a free choice to apply for approval under either the UNECE system of international type approvals or through the national scheme;**
- (3) developers should be able to submit an ADS for national approval, even if they are not responsible for manufacturing the whole vehicle.**

Do you agree?

Other.

It may be helpful to undertake a cost-benefit analysis to determine what scope of domestic scheme would merit the resources that would need to be expended in establishing a scheme of such scope. The second “self-driving categorisation” step proposed in the consultation, allows any specific desired national requirements to be incorporated within that process. One would expect established manufacturers to pursue approval under the UNECE system. Once both systems are defined and available this suggests the primary take up of a domestic scheme might well be limited to smaller series and post type approval modifications. However, a separate national system could be advantageous to permit faster regulatory progress in Great Britain or could provide a more permissive or less onerous regulatory burden to ADSEs. Any variations between the domestic scheme and international scheme (particularly if the international scheme were in fact less stringent) would need to be accounted for in the second “self-driving categorisation” step.

In respect of (3), a self-driving system is highly sensitive to small changes. For example small adjustments of the control system will affect the vehicle dynamics, the position of the sensors and their field of view will affect the perception system. It should be considered whether the developer would have sufficient knowledge about the rest of the vehicle to be in a position to provide sufficient information to satisfy the requirements to obtain approval (see further our response to Consultation Question 11(3) below), as much of the intellectual property in a vehicle that would need to be provided is held by the OEM and/or Tier 1 supplier.

Consultation Question 11 (Paragraph 8.43)

We provisionally propose that:

- (1) an ADS approval scheme should be established through regulation under the Road Traffic Act 1988, without further legislative reform;**
- (2) an ADS should be defined as a combination of software, hardware and sensors, which can be installed in a “type” of vehicle;**
- (3) when an ADS is approved, the approval should be accompanied by specifications for:
 - (a) the type of vehicle in which it can be installed; and**
 - (b) how the ADS is installed within the vehicle;****
- (4) where an ADS is installed in a pre-registered vehicle, an example vehicle should be submitted to the regulator for approval of the installation.**

Do you agree?

Yes.

Agreed. These provisions would cater for smaller scale deployments using small numbers of vehicles and would allow for deployments once the development of the ADS goes beyond

the test/trial stage (where it may be exempted from these approvals). This would enable more bespoke deployments (e.g. in specific geo-restricted areas, campus locations) using small series vehicle production that could include modification of a base vehicle platform designed for such purposes.

Whilst we are conscious the consultation paper considers the vehicle 'type' in which a vehicle might be installed would require detailed discussion in each safety case (CP3 paragraph 8.34), we observe the specifications to be provided will need to cover a much broader range of matters than envisaged at (3). As we mentioned above, a self-driving system is highly sensitive to small changes. The approval scheme would need to be a true holistic verification of the vehicle and its associated parts, requiring information about the other systems within the vehicle. For example, it would require information about the fall back system, which could be produced by another developer, and how this interoperates with the ADS (for example 1oo2 or 2oo3 safety assurance fail-safe or fail-operational handovers).

In regard to (4) (sample vehicles submitted for testing) - this would need some guidance on how similar/different the example vehicle can be from others within the proposed batch submitted for type approval. For example, minor variations in equipment/trim levels or fitted OEM systems, variations in wheel/tyre sizes, etc. can all create small differences in performance. The latest WLTP emissions testing regime is a good example of this, with every OEM model and option combination requiring a submission for testing because of possible variations in the results.

Consultation Question 12 (Paragraph 8.44)

We invite observations on the appeal process in regulation 19 of the Road Vehicles (Approval) Regulations 2020, including:

- (1) how it works in practice; and**
- (2) how well it is suited to the proposed national ADS approval scheme.**

We do not have any comments on the appeal process.

Consultation Question 13 (Paragraph 8.71)

18.13 We provisionally propose that:

- (1) once an ADS has received type approval at either international or domestic level, an Automated Driving System Entity (ADSE) would need to submit the vehicle to the UK safety regulator for categorisation as able to safely drive itself;**
- (2) the safety regulator should make a recommendation to the Secretary of State for how the vehicle should be classified;**
- (3) it should be open to the safety regulator to recommend that an ADS-enabled vehicle is classified in one of three ways: as not self-driving but driver assistance; as self-driving only with a user-in-charge; or as self-driving without a user-in-charge;**
- (4) the safety regulator should only recommend classification as self-driving (either with or without a user-in-charge) if it is satisfied that:**
 - (a) an ADSE is registered as taking responsibility for the system;**
 - (b) the ADSE was closely involved in assessing safety and creating the safety case; and**

(c) the ADSE has sufficient funds accessible to the regulator to respond to improvement notices, to pay fines and to organise a recall.

Do you agree?

Yes.

(1) - (2) Agreed.

(3) Agree in principle. However, as explained in our response to consultation question 4, we would envisage deployment in some cases may take place in a gradual process, with an initial deployment in a limited ODD that expands over time. As a result, any classification would need to specify the ODD in which it applied, and it seems possible that the same system may meet different classifications in different ODDs, i.e. the approval scheme would need to be more flexible than merely those three options - in some ODDs it may only meet the standard for being a driver assistance system, in other ODDs it may be self-driving with a user-in-charge, in other ODDs no user-in-charge may be required. Consideration should be given as to whether such type approvals should have a time limit so that ADSEs are compelled to continually update the ADS (e.g. through software updates) and/or their safety case.

(4)(a)-(b) Agree in principle. However, we do not think the entity that takes responsibility for the continuing safety of the system and the entity that takes responsibility for the information presented necessarily needs to be the same in all cases; the safety regulator merely needs to be assured that the entity taking on the responsibility is in a position to do so. As a simple example, the business of the second entity may have been bought by the first. But other relationships are possible.

In practice, while one entity may take responsibility as the ADSE, it may be difficult to assure the continuing safety of the system without a close co-operation between the vehicle operator and ADS system developer. For example, if the vehicle includes a second protection system that feeds into a path planner, and an OTA update is made to the path planner which affects how the ADS interoperates with the path planner and causes the second protection system to take over when it should not. As part of the safety approval process, the regulator should be assured that the contractual arrangements between the various entities are sufficient to ensure safety. Indeed, it may be necessary to mandate that the vehicle operator provides certain data from the field and associated information back to the developer to enable the developer to monitor and improve the safety of its system.

Regarding 4(c), in our view the safety regulator needs to be assured that the entity taking responsibility for the continuing safety of the system has sufficient resources, including access if required to sufficient funds to do so (which may be available to it in a number of ways such as insurance, as noted in CP3 paragraph 8.69). In addition to the examples in CP3 paragraph 8.69, these funds could include committed support from a third party such as the licensed operator. Furthermore, for NUIC vehicles in particular, it may be appropriate to allow the ADSE and licensed operator to agree between themselves which of them will be responsible for organizing a recall.

We support the suggestion in CP3 paragraph 8.70 that the financial standing requirement should be flexible. However, the assessment of sufficient resources (including access to sufficient funds if required) should be in the context of the proposed deployment/sales model, not just volume, as suggested in CP3 paragraphs 8.69 to 8.70. For example, a model sold to the mass market for any member of the public to use would require a different funding structure to offer adequate support compared with a limited number of units deployed via selected operators for a small scale service. Flexibility in relation to deployment/sales model will support a greater range of entrants to the market, whilst maintaining sufficient financial guarantees proportionate to the potential financial consequences.

Finally, it seems to us that a useful distinction could be made between whether or not the system was sufficiently safe and whether the ADSE was of sufficient good standing, the first being a matter of meeting a safety standard, the second being a matter for deployment. It may be helpful to require the regulator to rule separately on these, to give the ADSE the opportunity to correct any failings, or to allow another entity, sufficiently involved in the safety process to ensure the continuing safety of the system, to step in.

Consultation Question 14 (Paragraph 8.77)

We provisionally propose that a new legislative framework should provide regulation making powers to specify:

- (a) who should assess whether a vehicle is capable of self-driving;**
- (b) the procedure for doing so; and**
- (c) criteria for doing so.**

Do you agree?

Yes.

We have no comments on this question.

Consultation Question 15 (Paragraph 8.78)

We seek views on whether the new legislation should include provisions for appeals against a categorisation decision. If so, should these be similar to those in regulation 19 of the Road Vehicles (Approval) Regulations 2020?

We agree there should be provision for appeals. We have no comments on whether or not these provisions should be similar to those in regulation 19 of the Road Vehicles (Approval) Regulations 2020.

Consultation Question 16 (Paragraph 8.83)

We seek views on whether the regulator that classifies vehicles as self-driving should have the power to allow their deployment in limited numbers, so as to gather further data on their safety in real world conditions.

Agreed. An ADSE operating an appropriate risk management system would potentially propose a limited (not just in numbers, but also in ODD), iterative deployment with continuous monitoring. Accordingly, the regulator's power should not be restricted just to allowing deployment in limited numbers, but should also extend to allowing deployment in a

limited ODD. Such a deployment (indeed any deployment) should be coupled with continuing reporting requirements on the ADSE to harness the benefits of this approach.

However, equally it should be recognised that a limited deployment that extends over a long period may not be commercially viable, and hence restrict who can enter the market to those with the deepest pockets and limit competition. Whilst safety must always be the overriding factor, if the regulator is only approving a more limited deployment than the ADSE has requested, it should be incumbent upon the regulator to assess the risks in light of the value of the deployment in the round, including the commercial implications, and to support the ADSE in working towards a wider deployment, for example by giving clarity over the measures that need to be taken to achieve this.

From a business perspective, it is important that there is a clear distinction between development trials and deployment, and the approval process for each does not creep too much in either direction – for small trials operational safety and human factors will be more important for managing safety, whereas for actual deployments safety should predominantly be justified by technical arguments.

CHAPTER 10: ASSURING SAFETY IN USE

Consultation Question 17 (Paragraph 10.82)

We provisionally propose that legislation should establish a scheme to assure the safety of automated driving systems following deployment, giving scheme regulators enhanced responsibilities and powers.

Do you agree?

Yes.

We have no comments on this question.

Consultation Question 18 (Paragraph 10.83)

We provisionally propose that the enhanced scheme should give regulators the following responsibilities and powers:

(1) scheme regulators should be responsible for comparing the safety of automated and conventional vehicles using a range of measures;

(2) to do this the regulator should have power to collect information on:

(a) leading measures (instances of bad driving which could have led to harm) and

(b) lagging measures (outcomes which led to actual harm);

(3) regulators should have power to require an ADSE:

(a) to update software where an update is needed to ensure safety and continued compliance with the law;

(b) to keep maps up-to-date, where an AV relies on maps to ensure safety and compliance with the law;

(c) to communicate information about an ADS to users in a clear and effective way, including where necessary through training.

Do you agree?

Yes.

Agree in principle.

The use of leading measures may be misleading in the context of AVs if the measures collected are based on those that apply to 'bad' human drivers. For example, studies have found harsh braking events are correlated with 'bad driving' in humans, which is why many insurance-based telematics systems monitor the incidence of such events and penalise drivers who display these driving characteristics frequently. Such instances in an AV may not be a similar indicator of 'bad driving' - each instance would need to be considered in context. For example an AV would generally only apply harsh braking to avoid a potential collision (whereas in human drivers it is more often associated with poor situational and hazard awareness coupled with excessive speed). An AV may brake harshly to maintain its prescribed 'safety envelope' and other road users may take advantage of or even 'bully' AVs into braking harshly by not giving them sufficient space.

Similarly, excessive speed is another leading measure of bad human driving which is unlikely to be of value for AVs that will undoubtedly be prevented from speeding. Care will be needed to select leading measures that are relevant to AVs. Examples might include state changes such as the frequency of transition to a MRC, the use of an Emergency Manoeuvre (as specified in the ALKS regulation), or unstable lateral positioning within lane.

Accordingly, work will be required to define leading measures that are meaningful indicators that the ADS is not operating as intended. An ADSE/ADS developer could be encouraged to specify in their safety case what leading/lagging measures they will monitor as part of their safety monitoring, with the regulator then having the power to ensure the ADSE then monitors those measures.

In relation to (3), the regulator should have a general power to require the ADSE to take appropriate corrective action when an issue is identified. The precise corrective action to be taken should be a matter for the ADSE to determine. Only the ADSE will have the necessary deep technical knowledge and familiarity of the whole system to be able to determine what action is likely to correct that issue without causing regression elsewhere, it being up to the ADSE to then convince the regulator that the ADS will meet the requisite standard of safety following that corrective action. It is the ADSE that bears liability for the system. It would be undesirable to empower the regulator to require the ADSE to take a specific action (e.g. an update to the software), as if that action then caused a regression, it would raise difficult issues of which party was liable for any consequences of that regression. However, as we noted above in response to consultation question 13, it may be appropriate for the regulator to require an ADSE to continuously update its software to improve its safety in order to maintain self-driving approval.

The two actions identified at 3(a) and (b) are only two of a variety of actions that may be required to fix an issue. These two options may not be appropriate in the circumstances, or it may be possible to fix the issue using one of a number of different approaches (for example, a software update, hardware change, process change (e.g. new calibration method), an ODD change (a new route) or an infrastructure aid change) and the ADSE may consider a different approach is better. For example, the ADSE may prefer to upgrade the

sensor hardware on a vehicle rather than the software to fix an issue with an AV hesitating to merge into a lane. We do not consider it helpful to single these specific actions out; the focus should be on good problem solving processes e.g. the 8D problem solving process.

The actions that an ADSE intends to take generally to maintain safety over the lifetime of the ADS, such as the frequency or triggers for any updates, should form part of the ADSE's safety case.

In relation to providing training and information to 'users' in 3(c), we are unclear whether this refers to users-in-charge or some other category of 'user'? If it is a broader category of user, this requirement could create an excessive and unrealistic burden. For example, it may be impractical to provide training to those who hire vehicles with ADS features from a car hire company. A combination of general education (through inclusion of information or requiring use of these systems as part of the driving test), clear information and specific training where appropriate could be used.

Consultation Question 19 (Paragraph 10.84)

We welcome views on the following issues:

(1) Should scheme regulators be empowered to approve software updates that apply only within the UK, without requiring the manufacturer to return to the original type approval authority?

(2) Should the scheme also deal with cybersecurity?

(3) Are other powers needed? (Note that data is discussed in Chapter 17.)

We agree with the proposals in (1) and (2).

Consultation Question 20 (Paragraph 10.100)

Should the authority administering the scheme to assure safety while automated vehicles are in use be kept separate from type approval authorities (as is already the case)? Alternatively, should both functions be combined in a single body?

Given the technical competence required, on balance we consider it would make more sense for the functions to be combined in a single body with appropriate legal safeguards.

Consultation Question 21 (Paragraph 10.101)

What formal mechanisms could be used to ensure that the regulator administering the scheme is open to external views (such as duties to consult or an advisory committee)?

We agree it is important to avoid regulatory capture and seeking external views is helpful here, but do not have a particular view on what mechanism is used, provided it is effective and otherwise appropriate.

CHAPTER 11: INVESTIGATING TRAFFIC INFRACTIONS AND COLLISIONS

Consultation Question 22 (Paragraph 11.24)

We provisionally propose that a statutory scheme to assure AVs in-use should:

- (1) investigate safety-related traffic infractions (such as exceeding the speed limit; running red lights; or careless or dangerous driving);**
 - (2) investigate other traffic infractions, including those subject to penalty charge notices;**
 - (3) if fault lies with the ADSE, apply a flexible range of regulatory sanctions.**
- Do you agree?**

Yes.

We agree in principle with (1) and (2). However, the focus should be on investigating infractions that have a meaningful bearing on the safety of the AV (see further our response to Consultation Question 18 above), or are issues of wider public interest, rather than automatically focusing on infractions that are meaningful for human drivers. This area would benefit from further research to determine both which infractions are meaningful and which have most impact, to ensure best use of finite resources and to enable resources to be allocated appropriately between investigating different types of infractions.

In regard to (3), we agree with the qualification that regulatory sanctions should only be applied if the “fault” is an indicator that there is an overall issue with the safety of the AV. As noted in CP3 paragraph 11.22, human drivers and ADSE will not respond to penalties in the same way. In addition it is important to keep in mind the differences between human driven vehicles and AVs, and the underlying principles of the systems that regulate them. The current system for human driven vehicles is designed to moderate the driving behaviour of each driver as an individual (for example by deterring speeding through the risk of incurring points on one’s licence and possible disqualification). This concept of regulating each driver on an individual basis is simply not relevant for an AV, where one is looking holistically at whether or not AVs of that specific type meet the required standard of safety. Importing the concept of human centric regulation into the regulation and sanctions applied to AVs would create misaligned incentives.

Generally, it may be better to consider a regime like that in the aviation sector, where ADSEs are encouraged to share as much data as possible, without the threat of penalty, in order to maximise the safety of systems.

Consultation Question 23 (Paragraph 11.53)

We provisionally propose that the regulator which assures the safety of AVs in-use should have powers to impose the following sanctions on ADSEs:

- (1) informal and formal warnings;**
- (2) fines;**
- (3) redress orders;**
- (4) compliance orders;**
- (5) suspension of authorisation;**
- (6) withdrawal of authorisation; and**
- (7) recommendation of attendance at a restorative conference.**

Please select the sanctions which you think the regulator should have powers to impose, leaving blank those which you do not think the regulator should be able to impose.

Ticked all.

Any sanctions should be proportionate and applied fairly and consistently. Like in the aviation sector, the focus should be on improving safety and access to the benefits of self-driving. For this purpose, the regulatory system should encourage collaboration with the regulator and self-reporting, encouraging finite resources to be focused on what can be learnt from the incident and implementing that, rather than diverting resources to contesting liability and self-defence. When assessing proportionality, the focus should be on the conduct of the ADSE, the nature of the failure and the level of risk posed by it, rather than the actual result of the failure. This is to reflect the necessary move away from the human centric rules and regulations that apply to human drivers to a product safety led system. For example, a failure that leads to a material deterioration in an AV's ability to drive safely in the rain, a relatively common occurrence, that has by fortune not caused any fatalities at the time of discovery, should be viewed more seriously than an AV's failure to deal appropriately with a very rare event that unfortunately leads to multiple fatalities in a single event.

Consultation Question 24 (Paragraph 11.54)

We provisionally propose that the legislation should provide the regulator with discretion over:

- (1) the amount of any monetary penalty; and**
- (2) the steps which should be taken to prevent re-occurrence of a breach.**

Do you agree?

Other.

The regulator should develop a framework to ensure sanctions are proportionate and applied fairly and consistently, and there should be appropriate checks and balances in the framework developed. As noted above, it is important to foster an open culture of sharing data in order to improve safety. Penalties should not be politically motivated or influenced by extraneous factors. In setting the size of the monetary penalty, it may be appropriate to take into account the overall value of the deployment and size of undertaking. The steps to be taken to prevent reoccurrence of a breach should be set in the wider context of the continual safety improvement of the system, and not impose a disproportionate burden in one area which could be detrimental to the standard of safety overall. There should also be an effective mechanism to enable the type and level of penalty to be independently reviewed. An upper limit or limits should be set on the amount of any monetary penalty.

Consultation Question 25 (Paragraph 11.69)

We provisionally propose that a specialist incident investigation unit should be established:

- (1) to analyse data on collisions involving automated vehicles;**
- (2) to investigate the most serious, complex or high-profile collisions; and**
- (3) to make recommendations to improve safety without allocating blame.**

Do you agree?

Yes.

A specialist incident investigation unit is important for the reasons explained in the Consultation Paper.

Consultation Question 26 (Paragraph 11.82)

We provisionally propose that the UK Government should establish a forum for collaboration on the application of road rules to self-driving vehicles.

Do you agree?

Yes.

Agreed in principle. Clarifications of the existing highway code for specific circumstances could be beneficial for AVs and non-AVs alike. For example, Highway Code Rule 145 states: "You MUST NOT drive on or over a pavement, footpath or bridleway except to gain lawful access to property, or in the case of an emergency...". If the pavement were empty and it prevented deadlock in traffic, it may be considered acceptable to mount the kerb to allow an oncoming vehicle to pass. But this would not self-evidently be an 'emergency'.

It would be important that such a forum had sufficient representation from engineers responsible for developing ADS to be used in the target domain to ensure that the interpretations of road rules were realistic to implement. Also, there is a danger that if the Highway Code were to diverge into two branches to represent human drivers and AVs separately, this could lead to considerable confusion amongst those working on ADS development, amongst all road user groups and for those involved in testing, monitoring and using automated driving systems and features.

Given the diversity of UK roads, it is inevitable that rules such as the examples given in 11.76 could lead to situations where AVs are unable to proceed because of certain rule combinations. That is why there is some benefit to the existing Highway Code, which allows road users to balance risks and make informed decisions about how to proceed in a way that generates the lowest risk. Arguably AVs could improve upon human performance in this respect.

It is also worth recognising that if rules for AVs and human drivers diverge whilst there is still a mixed network of users, this could exacerbate road user frustration with AVs if they are applying rules that other road users do not recognise as correct or are simply unaware of.

Consultation Question 27 (Paragraph 11.83)

We welcome views on:

- (1) the issues the forum should consider;**
- (2) the composition of the forum; and**
- (3) its processes for public engagement.**

As noted in our response to consultation question 26, there are risks as well as benefits to establishing such a forum. It would be important to ensure that the forum had sufficient

representation from engineers responsible for developing ADS to be used in the target domain, otherwise it could end up recommending proposals that are technically impossible or unduly burdensome to implement. In addition, it would need to consider carefully where clarification is needed and where it may be better to leave in flexibility for road users, including AVs, to make informed decisions about how to proceed in a way that generates the lowest risk.

CHAPTER 12: THE USER-IN-CHARGE

Consultation Question 28 (Paragraph 12.24)

We provisionally propose that that the user-in-charge:

(1) should be defined as an individual in position to operate the controls of a vehicle while an ADS is engaged and who is either in the vehicle or in direct sight of the vehicle; and

(2) is not a driver while the ADS is engaged, and would not be liable for any criminal offence or civil penalty (such as a parking ticket) which arises out of dynamic driving. Do you agree?

Yes.

Agreed (although it may be better to rephrase (1) as 'an individual who is in a position to control the vehicle', to better accommodate the example in CP3 paragraph 12.14 of a user-in-charge controlling the vehicle through their phone).

It should be a requirement that the user-in-charge remain in direct sight of the vehicle - i.e. if the user-in-charge were to walk around a corner as suggested in CP3 paragraph 12.16, it should be that individual who should be liable for the failure to operate an ADS without a user-in-charge.

Consultation Question 29 (Paragraph 12.37)

We provisionally propose that following the end of the transition demand period:

(1) the user-in-charge should re-acquire the legal obligations of a driver, whether or not they have taken control of the vehicle; and

(2) if, following a failure to respond to a transition demand, the vehicle stops in a manner which constitutes a criminal offence, the user-in-charge should be considered a driver and should therefore be liable for that offence.

Do you agree?

Yes.

Agreed (and we agree with the points made in CP3 paragraphs 12.25 - 12.36 about the need to ensure a safe process for handover, so that users-in-charge do not accidentally take over charge and that when taking charge following a transition demand it can be ensured that users are ready to do so).

We are unclear why the circumstance in (2) is singled out; the user-in-charge should be liable for any offences post transition, subject to the defence mentioned (i.e. where the actions of the ADS made the commission of the offence unavoidable). The circumstance in

(2) would require evidence-based research on the appropriate amount of time to allow for a user-in-charge to respond successfully to a transition demand. Users-in-charge may receive mixed messages about the consequences of failing to respond to a transition demand (the consultation discusses how in some situations no response would not be an offence if the vehicle can safely stop itself whereas in other situations failing to respond may be a 'serious criminal offence'). If users-in-charge are fearful of prosecution, this could lead to a panicked and unsafe rush to respond to a transition demand in a manner that may lead to further risk. Care also should be taken to ensure that ADSE uses the transition demand process appropriately and that the 'fall back' response to a failure to comply is an appropriate response to minimise risk (e.g. in the case of ALKS and the potential for a transition demand if general traffic speed increases, it may be safer to continue at the maximum ALKS operational speed in the absence of a response to a transition demand rather than stopping in a live traffic lane where the general traffic speed is now substantially faster).

Consultation Question 30 (Paragraph 12.45)

We seek views on whether a person with a provisional licence should be allowed to act as a user-in-charge, if accompanied by an approved driving instructor in a vehicle with dual controls.

We do not have a particular view on this question.

Consultation Question 31 (Paragraph 12.53)

We provisionally propose that legislation should create new offences of:

- (1) using an automated vehicle as an unfit or unqualified user-in-charge; and**
- (2) causing or permitting the use of an automated vehicle by an unfit or unqualified user-in-charge.**

Do you agree?

Yes.

We have no comments to make on this question.

Consultation Question 32 (Paragraph 12.59)

We provisionally propose that persons carried without a user-in-charge should be guilty of a criminal offence. Do you agree?

Yes.

Agreed where there is the appropriate mens rea.

Consultation Question 33 (Paragraph 12.60)

We seek views on whether the new proposed offence of being carried without a user-in-charge should only apply if the person:

- (1) knew that the vehicle did not have a user-in-charge; and**
- (2) knew or ought to have known that a user-in-charge was required.**

We agree with both these requirements.

Consultation Question 34 (Paragraph 12.66)

We provisionally propose that a user-in-charge who takes over control of the vehicle:

- (1) should be considered a driver; but**
- (2) should have a specific defence to a criminal offence if, given the actions of the ADS, a competent and careful driver could not have avoided the offence.**

Do you agree? If not, we welcome views on alternative legal tests.

Yes.

Agreed. However, it may be preferable to refer to “in all the circumstances” rather than limiting the relevant circumstances to the “actions of the ADS”.

Consultation Question 35 (Paragraph 12.94)

We provisionally propose that the user-in-charge should be liable for criminal offences which do not arise from the dynamic driving task, including those related to:

- (1) insurance;**
- (2) maintaining the vehicle in a roadworthy condition (including installing safety critical software updates);**
- (3) parking;**
- (4) duties following accidents to provide information and report accidents to the police; and**
- (5) ensuring child passengers wear seatbelts.**

Do you agree?

The user-in-charge should be liable for some but not all of the offences listed.

Agree with the provisos that:

(i) where there is a licensed operator, the licensed operator should bear responsibility for installing safety critical software updates, and maintaining the vehicle in a roadworthy condition; and,

(ii) whilst the user-in-charge should bear primary responsibility for the duties following accidents, it may also be appropriate for the ADSE and/or licensed operator to be under some form of duty to report as well, where they have relevant knowledge.

In addition it would be helpful if the the user-in-charge were under an obligation to provide information and report accidents to the ADSE and/or licensed operator, or otherwise to ensure the police provide such information to ADSE and/or licensed operator to enable data to be captured and retained by the ADSE (as noted elsewhere, a large volume of data is generated by an AV in operation so this would enable the ADSE and/or licensed operator to retain richer data than that automatically captured and retained as a result of the accident).

Consultation Question 36 (Paragraph 12.95)

We provisionally propose that the legislation should include a regulation-making power to clarify those roadworthiness failings which are (and those which are not) the responsibility of the user-in-charge.

Do you agree?

Yes.

We have no comments to make on this question.

CHAPTER 13: REMOTE OPERATION: NO USER-IN-CHARGE VEHICLES

Consultation Question 37 (Paragraph 13.67)

We provisionally propose that:

(1) where an individual is exercising latitudinal and longitudinal control (steering and braking) over a vehicle remotely, that should not be regarded as a form of “self-driving”; and

(2) where lateral and longitudinal control are exercised by an ADS, all other forms of remote operation should be regulated as “self-driving”.

Do you agree?

Yes.

Agree with the qualification that for (1), only if that control is being exercised directly, in real time with no support from the ADS. Also, latitudinal OR longitudinal control should take the vehicle outside of being self-driving, it is not necessary to have both.

We welcome views on whether the current definition of when a vehicle “drives itself” under the Automated and Electric Vehicles Act 2018 should be amended to deal with some forms of remote operation which may involve a degree of “monitoring”.

We do not have a view on this question.

Consultation Question 38 (Paragraph 13.86)

We provisionally propose that:

(4) the regulation of self-driving vehicles should distinguish between an Automated Driving System Entity (which vouches for the design of the system) and an operator (responsible for the operation of individual vehicles);

(5) all vehicles authorised for use on roads or other public places with no user-in-charge should either:

(a) be operated by a licensed operator; or

(b) be covered by a contract with a licensed operator for supervision and maintenance services;

(6) it should be a criminal offence to use a NUIC vehicle on a road or other public place unless it is operated by a licensed operator or is covered by a contract with a licensed operator for supervision and maintenance services.

Do you agree?

Yes.

Agreed. However, in some instances it may be appropriate to allow either entity to take on certain responsibilities (provided that it is clear which entity is responsible) or to mandate some responsibilities are shared (for example, recall responsibilities).

Consultation Question 39 (Paragraph 13.92)

We welcome views on whether NUIC operators should be required to demonstrate professional competence through a safety management system, as set out in a safety Case.

Agree. The concept of having a safety management system also aligns with other items within the consultation, such as safety monitoring.

Consultation Question 40 (Paragraph 13.108)

We provisionally propose that, irrespective of the nature of the vehicle, a licensed operator should be under a duty to:

- (1) supervise the vehicle;**
- (2) maintain the vehicle;**
- (3) insure the vehicle;**
- (4) install safety-critical updates and maintain cybersecurity; and**
- (5) report accidents and untoward events (as defined by the regulator).**

Do you agree?

Yes, a licensed operator should be subject to all the listed duties.

Agreed. However in relation to (5), consideration should be given to defining the triggers for reporting, and/or the relevant knowledge requirement. Maintaining the vehicle in (2) would include cleaning sensors and running scripts periodically to check calibration.

Consultation Question 41 (Paragraph 13.109)

We provisionally propose that legislation should include a regulation-making power by which some or all of these duties could be transferred to the registered keeper or owner, if it was shown that it was appropriate to do so.

Do you agree?

Yes.

We have no comments on this question.

Consultation Question 42 (Paragraph 13.116)

We welcome views on how accessibility standards for Highly Automated Road Passenger Services (HARPS) might be developed.

We do not have any specific comments to make on this question.

We provisionally propose that:

- (1) an accessibility advisory panel should be formed to include:**
 - (a) the Equalities and Human Rights Commission; and**
 - (b) representative groups for disabled and older persons;**
- (2) the Secretary of State should be obliged to consult with the accessibility advisory panel prior to setting any national minimum standards on HARPS;**
- (3) there should be a duty to periodically re-consult the accessibility advisory panel**

at set intervals to ensure requirements keep pace with developing evidence of technical feasibility and changing needs.

Do you agree?

Yes.

Agree - such a panel would be an important step towards unlocking the full potential of AVs for everyone. Legacy transport systems have typically developed without this influence and have then had to retrospectively improve accessibility which is always sub optimal. There is a unique opportunity for AVs (especially those developed for HARPS) to have accessibility designed in from the outset.

We welcome views on what the set interval for periodically re-consulting the accessibility advisory panel should be.

Annual re-consulting with an accessibility panel would be a sensible minimum, with it convening more frequently when there are specific HARPS developments/deployments that require its input.

Consultation Question 43 (Paragraph 13.133)

We welcome views on who should administer the operator licensing scheme.

We do not have a particular view on this question.

CHAPTER 14: CRIMINAL OFFENCES BY ADSES AND THEIR SENIOR MANAGERS

Consultation Question 44 (Paragraph 14.107)

We provisionally propose that:

- (1) it should be a criminal offence for an ADSE to omit safety-relevant information or include misleading information when putting a vehicle forward for classification as self-driving or responding to information requests from the regulator;**
- (2) the offence should apply to senior managers (where it was attributable to the manager's consent, connivance or neglect);**
- (3) the offence should not apply to more junior employees;**
- (4) the offence should carry a higher sentence if it is associated with a death or serious injury;**
- (5) the offence should be prosecuted in England and Wales by either the regulator or the Crown Prosecution Service and in Scotland by the Procurator Fiscal.**

Do you agree?

Agree in principle with (1), (2) and (5).

The provision of sufficient reliable information to the regulator is clearly important, and where there is wrongdoing it should be possible for the person responsible for that wrongdoing to be held accountable (assuming they are truly responsible). However, we are concerned that the appropriate balance is struck given the consequence of finding conduct criminal, such that conduct is only criminalised to the extent that it falls sufficiently far below the bar of what society would expect that it is proper to criminalise it. Drawing analogies with offences in other industries is helpful, but care should be taken to take into account the specific nature of

the technology and industry when importing concepts to ensure such concepts are appropriate to import and are not 'cherry picked' - being removed from the qualifications and the defences in the original offence that provided the appropriate checks and balances. An ADS is a complex product with many components and the ADS itself is sensitive to small changes. The ADSE will be reliant on information from many different sources inside and outside of the ADSE, and senior managers cannot be expected to have intimate knowledge of all the different parts of the systems. Furthermore, with hindsight certain matters can take on an importance that was not evident at the time. Any offences should be proportionate to the maleficence of the conduct. We agree with the desire in CP14.100 to encourage and open and transparent culture.

We find the issue of the liability of non-director level employees, whether senior managers or junior employees, difficult. Given the complex nature of an ADS managers will be reliant on information provided by more junior employees. It could be difficult to distinguish between who is sufficiently senior for this purpose - the distinction not being one of title or length in the job, but expected levels of responsibility. As a matter of fairness, one feels identical conduct ought to be treated identically, but employees are in a slightly different position to directors (for example, they may be less confident of speaking up). Difficult issues could arise if the senior manager or junior employee was instructed to deceive or even trained, without realising it, to deceive. Directors are in a different position. In any event, the level of seniority should be taken into account when determining the standard of conduct expected.

In relation to (4), we appreciate the importance to the public of it being seen that a person is held to account when there is a death or serious injury. However, we do not agree that linking the sentence to the occurrence of a death or serious injury is the best reflection of the culpability of the conduct. A better reflection would be linking the sentence to (i) the level to which the risk of a type of adverse incident was increased, and (ii) the severity of consequence of that type of adverse incident (i.e. death or serious injury). A person who is responsible for an offence that relates to an issue which, although widespread, was unlikely to have severe consequences but unluckily happened to cause a fatality would seem less culpable than a person responsible for an offence that relates to a widespread issue that was very likely to have severe consequences but by good fortune had not happened to cause a fatality before it was discovered.

Consultation Question 45 (Paragraph 14.108)

We seek views on the following proposed offences.

Offence A: non-disclosure and misleading information in the safety case

When putting forward a vehicle for classification as self-driving, it would be a criminal offence for the ADSE to

- (1) fail to provide information to the regulator; or**
- (2) provide information to the regulator that is false or misleading in a material particular where that information is relevant to the evaluation of the safety of the ADS or the vehicle.**

The ADSE would have a defence if it could show that it took reasonable precautions and exercised all due diligence to prevent the wrongdoing.

The penalty would be an unlimited fine.

Offence B: non-disclosure and misleading information in responding to requests

When a regulator requests specific information from an ADSE (whether before or after deployment), it would be a criminal offence for the ADSE to

- (1) fail to provide information to the regulator; or**
- (2) provide information to the regulator that is false or misleading in a material Particular where that information is relevant to the evaluation of the safety of the ADS or the vehicle.**

The ADSE would have a defence if it could show that it took reasonable precautions and exercised all due diligence to prevent the wrongdoing.

The penalty would be an unlimited fine.

Offence C: offences by senior management

Where offence A and/or offence B committed by a body corporate is proved—

- (3) to have been committed with the consent or connivance of an officer of the body corporate; or**
- (4) to be attributable to neglect on the part of an officer of the body corporate, then that officer is guilty of the offence.**

An officer includes any director, manager, secretary or other similar officer or any person who was purporting to act in any such capacity.

We see this as equivalent to offences under the Human Medicines Regulations 2012 and General Product Safety Regulations 2005, which carry a penalty of a fine and/or a maximum two years' imprisonment.

Offence D: aggravated offences in the event of death or serious injury following non-disclosure or provision of misleading information to the AV safety regulator

Where a corporation or person commits Offences A to C, that offence is aggravated where the misrepresentation or non-disclosure:

- (5) related to an increased risk of a type of adverse incident; and**
- (6) an adverse incident of that type occurred; and**
- (7) the adverse incident caused a death or serious injury.**

We see this as equivalent to the offence of causing death by dangerous driving, which carries a penalty of an unlimited fine and/or a maximum of 14 years' imprisonment.

See our comments in relation to consultation question 44.

In relation to Offences A and B, we consider it should be a requirement that the information is relevant to safety - a failure to provide information that has no bearing on safety ought not to be criminalised, and creates an offence of concerning scope which could discourage open collaboration.

In relation to Offence C we note a number of those who fall within the definition of officer (such as a CFO or Company Secretary) would have no direct responsibility for safety or the requisite knowledge to be able to meaningfully input into this, and hence it is important the offence is delineated so that in practice it only captures those with real responsibility who ought to bear culpability, not merely those who may hold a senior position.

As noted in our response to consultation question 44, in relation to Offence D we disagree with the aggregated offence being based on causing a death or serious injury - it is the level of increased risk of an adverse event of that severity occurring which is more relevant, and more in alignment with the desired behaviours that one wishes to encourage and the undesired behaviours one wishes to discourage. An individual should not gain a criminal record for what only created a minor risk overall in context, but which unfortunately in a rare event caused a fatality or serious injury.

Also, we consider the threshold should be a “materially” increased risk, in order to impose criminal liability.

Consultation Question 46 (Paragraph 14.109)

We welcome views on whether an ADSE should be under a duty to present information in a clear and accessible form, in which safety-critical information is indexed and signposted.

Agree in principle, although the nature of some safety critical information pertinent to ADS performance may be difficult to present in a way that is accessible to all. An ADSE should, however, attempt to at least signpost and summarise more complex safety critical information so that it is digestible. On the other hand, by signposting certain information as a result the prominence of other information is necessarily downgraded, and it would be undesirable that by trying to comply with this requirement that the ADSE could be at greater risk of being committing an offence discussed in consultation question 45. Again, hindsight could play a part here.

CHAPTER 15: NEW WRONGFUL INTERFERENCE OFFENCES

Consultation Question 47 (Paragraph 15.10)

We provisionally propose that legislative amendment should clarify that the tampering offence in section 25 of the Road Traffic Act 1988 applies to anything that is physically part of a vehicle and any software installed within it. Do you agree ?

Yes.

Agree. However, it is important that the bar for the offence is set sufficiently high. For example, cleaning a camera could lead to a blurred image, but would not normally be considered tampering.

Consultation Question 48 (Paragraph 15.11)

We welcome views on whether the tampering offence should apply to external infrastructure required for the operation of the AV.

Agree. Whilst an ADS should be resilient to changes in infrastructure, attempts to deliberately interfere with safe AV operations should be an offence, whether it is through tampering with existing infrastructure or intentionally changing the presentation of infrastructure by adding artifacts that are known to interfere with perception systems, for example.

Consultation Question 49 (Paragraph 15.53)

We provisionally propose that there should be an aggravated offence of wrongfully interfering with an AV, the road, or traffic equipment contrary to section 22A of the Road Traffic Act 1988, where the interference results in an AV causing death or serious injury, in:

- (1) England and Wales; and**
- (2) Scotland.**

Do you agree?

Yes, in both England and Wales and Scotland.

We have no comments on this question.

Consultation Question 50 (Paragraph 15.55)

18.52 We provisionally propose that the appropriate mental element for the aggravated offence is intent to interfere with a vehicle, the road or traffic equipment.

Do you agree?

We find this a difficult question. Given the seriousness of the criminal offence compared to the trivial acts that could trigger it (e.g. scuffing white lines) on balance we would favour the offence requiring an appropriate mental element in relation to the consequences of the act.

Consultation Question 51 (Paragraph 15.62)

We seek views on whether an approved work defence for repair or maintenance operations authorised by a vehicle manufacturer or Automated Driving System Entity is desirable.

We consider clarity in this area is important.

CHAPTER 16: CIVIL LIABILITY

Consultation Question 52 (Paragraph 16.24)

We provisionally propose that the way the Automated and Electric Vehicles Act 2018 deals with contributory negligence and causation is:

- (1) adequate at this stage; and**
- (2) should be reviewed by the UK Government in the light of practical experience.**

Do you agree?

Yes.

We have no comments on this question.

Consultation Question 53 (Paragraph 16.32)

We provisionally propose that measures should be put in place to compensate the victims of accidents caused by uninsured AVs.

Do you agree?

Yes.

We have no comments on this question.

Consultation Question 54 (Paragraph 16.47)

We provisionally propose that:

- (1) product liability law should be reviewed to take account of the challenges of emerging technologies;**
- (2) any review should cover product liability as a whole, rather than be confined to automated vehicles; it should not, therefore, form part of this project on automated vehicles.**

Do you agree?

Yes.

We have no comments on this question.

CHAPTER 17: ACCESS TO DATA

Consultation Question 55 (Paragraph 17.65)

We provisionally propose that:

- (1) for a vehicle to be classified as self-driving, it needs to record the location as well as the time at which the ADS is activated and deactivated;**
- (2) the Government should work within the UNECE to ensure data storage systems for automated driving record these data; and**
- (3) any national system to approve an ADS should require these data to be collected, subject to safeguards.**

Do you agree?

Other.

We agree that it seems useful to record the location in certain circumstances, although it may not be necessary for all deployments. We do not agree that this should be part of the requirement for a vehicle to be classified as self-driving - rather it should be a requirement that any self-driving vehicle has an appropriate data recorder before being deployed on the road.

Consultation Question 56 (Paragraph 17.71)

We provisionally propose that legislation should impose a duty on those controlling AV data to disclose data to insurers, where the data is necessary to decide claims fairly and accurately.

Do you agree?

Yes

Agreed with the qualification that the duty to disclose (relevant) data should be framed with sufficient clarity such that it is clear when the duty arises and what data is captured by that

duty, and to ensure the requirement only requires disclosure of relevant data. This would both aid compliance and ensure that any personal data so disclosed would be brought within Article 6(1)(c) of the General Data Protection Regulation. For example, a person controlling AV data would not be in position to determine whether or not the data was “necessary to decide claims fairly and accurately”. Care should be taken to protect (and not require disclosure of) proprietary classifications and formats.

Naturally, the disclosure of data should be subject to appropriate safeguards (for example in relation to ensuring requests are properly made and that any data so disclosed is kept confidential where appropriate).

Consultation Question 57 (Paragraph 17.81)

We provisionally propose that:

(1) initially, DSSAD data from self-driving vehicles should be stored for three years;

and

(2) the issue should be reviewed in the light of experience.

Do you agree?

Yes.

Agreed, assuming DSSAD data is kept at a level that the volumes of data to be stored are not onerous (e.g. if about 1MB of data is captured per minute of driving this would equate to about 1TB per ADS over 3 years which seems sustainable).

Consultation Question 58 (Paragraph 17.95)

We provisionally propose that:

(1) when an ADSE applies for categorisation of its vehicle types as self-driving, it should present the regulator with details on how data will be recorded, stored, accessed and protected;

(2) the regulator should only categorise a system as self-driving if it is satisfied that that the ADSE has systems to abide by its obligations under the GDPR.

Do you agree?

Agreed in principle. However, we would suggest reframing the requirement at (2) as this appears to elide two distinct issues:

- whether the ADSE itself is considered to be in compliance with its obligations under the GDPR in light of what it is proposing to do. This would include demonstrating that it has appropriate technical and organisational measures in place to provide appropriate protection for any personal data in all relevant aspects of its operations, including how data is extracted from the vehicle and processed subsequently; and
- whether the software and systems of the vehicle itself are designed such that this provides adequate protection for the personal data collected. This ought to be more straightforward to demonstrate and assess. For example, it may be possible to satisfy this by compliance with relevant standards.

As data protection compliance is a matter for the ICO, it is not clear to what extent the regulator would have the appropriate expertise to determine matters of data protection. On

the other hand, specialist technical knowledge may be required to assess the measures being proposed to take account of what is, and what is not, technically possible and the extent to which personal data is required to be processed e.g. for safety. In any event, it would be important for decisions of the regulator on such matters to align closely with the approach of the ICO to avoid an ADSE having to comply with divergent regimes simultaneously. It would seem beneficial to give further thought to how the regulator and ICO could cooperate to make best use of their respective resources and expertise in carrying out the assessment.

Finally, we suggest that the issue of data protection is dealt with in a separate document to the safety case rather than within the safety case, as suggested within the accompanying text to this question, given the issues are different.