

3rd Law Commission Consultation

HORIBA MIRA Response

Richard Hillman

CHAPTER 4: SELF-DRIVING AND HUMAN INTERVENTION

Consultation Question 1 (Paragraph 4.114)

18.1 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?

Agree

A good point was raised by WMG in paragraph 4.1.1.0, re. the need to monitor for permutations that are out of the ODD, so that a transition demand can be given. However, consideration also needs to be given to the immediacy of the situation. For example, the presence of rain that is heavier than the vehicle is designed to deal with could generally be addressed via a handover demand as it would not be necessary for the driver to react within less than the handover time. Indeed, there is an argument for handover times being flexible in such circumstances, as there is unlikely to be sufficient urgency that stopping in lane after 10 seconds becomes the most favourable option.

On the other hand, the presence a particular hazard that is outside the ODD and needs an immediate response (e.g. a pedestrian on the motorway) would generally make it unfeasible to rely on handover, as in many permutations the collision would have already occurred before the handover time is complete, and even if this were not the case, it would be challenging for the driver to react to a safety-critical scenario immediately after the 10 seconds. This means that the system could not safely delegate responsibility to the user, and therefore the vehicle must be engineered (and verified) to be able to handle the situation safely.

This highlights the need differentiate between ODD elements that allow time for a gradual handover, i.e. those that unfold progressively, from those that require an immediate response due to an immediate ODD excursion requiring an immediate response. Where it is possible for the latter to occur, it should not be possible for a level 3, 4 or 5 vehicle to completely exclude that permutation from the ODD.

Consultation Question 2 (Paragraph 4.115)

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

In general, applying the principles of 'inclusive design', i.e. designing human-machine interfaces to be appropriate for those with impairments, has been found to be beneficial to all users, even those without impairments, by forcing interfaces to be clear, concise and intuitive. In this particular case, mandating a haptic warning in addition to visual and audible warnings would not just allow alerts to be provided to those with hearing loss, but would help attract the attention of all users and provide a greater sense of urgency such that users don't wait until the transition time is up before removing attention from their other activities. We therefore fully support the use of inclusive design to ensure that automated driving systems are designed to be appropriate to as wide a base of users as possible.

It is also worth noting that a small number of road users have medical exemptions from wearing seatbelts, and therefore the use of seatbelts to provide haptic warnings, as used as an example in the consultation document, may not always be sufficient. Brief pulses of harsh braking to provide an uncomfortable jerk or the use of haptic actuators within the seat would be alternative options.

An alternative option would be to only provide audio and visual warnings but to make it the legal responsibility of the user to ensure they take whatever steps are necessary to remain receptive, e.g. by not playing music loud enough to render warnings inaudible. Under this principle, it could be argued that users with impaired hearing would have a responsibility to regularly check for visual warnings to compensate for lack of receptiveness to audible ones (bearing in mind that visual warnings don't have the ability to 'grab' attention and rely on some level of monitoring, unless extremely obtrusive warning lights are used). However, it is difficult to envisage regular visual checks being maintained over a long journey, and this therefore doesn't appear to be a reasonable expectation to place upon users.

CHAPTER 5: HOW SAFE IS SAFE ENOUGH?

Consultation Question 3 (Paragraph 5.118)

18.3 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?

Agree, but with caveats. There should be an internationally agreed standard method (or set of methods) used for the technical assessment of the vehicle, including what constitutes good practice in terms of development processes and how to conduct the physical and virtual testing. The development costs and complexity of the systems make international cooperation and standardisation essential.

Where individual countries would be able to apply discretion is in making judgements on what test results are acceptable for their population, bearing in mind factors such as:

- the risk level of current road traffic;
- the potential economic and social gains that AVs would be expected to bring;
- the prevalence of rare edge cases that the vehicle can't react safely to (which would therefore remain as residual risk), and;
- the willingness of the population to engage with new technology

All these considerations are inherently both political and regional in nature, and therefore best decided by individual governments.

The complexity of AVs and their operating environments, combined with the unprecedented and novel approaches that will be needed to demonstrate safe AV behaviour (which will bear little resemblance to current type-approval testing) means it will be difficult to set prescriptive requirements for what constitutes acceptability, at least in the short term prior to experience being gained and industry converging upon established norms. As such, the approval process may bear more resemblance to the approaches used in medical trials, with the regulator working closely with developers throughout the process and with expert panels making judgements on whether the trial can be expanded to a wider scale, this progressive expansion being continually monitored.

Consultation Question 4 (Paragraph 5.119)

18.4 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.

Paragraph 5.36 implies a 'competent, careful' driver is one that doesn't make mistakes that result in accidents (in addition to also being able to mitigate for mistakes of others to a high standard). This is not necessarily true; all drivers make mistakes, even competent ones, the difference is in the frequency and severity of the mistakes. It is inevitable that AVs will make some mistakes and cause some accidents, however low or high the number, so setting such a strict threshold to be better than a driver who never causes collisions (which applies to both the 'careful and competent' and the 'does not cause fault accident' categories as described in the consultation) would not be practical and would make it virtually impossible for AVs to ever be permitted. Therefore, if this metric were to be used, it must be recognised that 'good' is a long way short of 'perfect'.

It should also be considered what constitutes 'competent and careful' driving - should it represent the top 1%, 5% or 20% of human drivers, for example? In practice, it would be very difficult to identify performance thresholds for each percentile.

The selection of a metric depends upon whether it is used to assess performance in each individual scenario, or assess overall performance when the vehicle is subjected to a vast range of scenarios. The latter 'holistic' analysis of the data is vital, as it is inevitable that some residual risk will have to be accepted due to individual scenarios that the system does not perform well in.

Note that the threshold selected (e.g. 'competent' vs 'average') for the analysis of individual events and for the overall performance would not necessarily have to be the same. For example, a metric derived from RSS could be used to assess whether the response to an individual incident was in line with a competent driver, and then the overall success rate could then be assessed against the performance of the average motorist to decide whether the vehicle is ready for deployment.

Although it could be argued that only requiring AVs to be better than the average driver means that around half of drivers would be exposed to more risk, this overlooks the fact that drivers who judge themselves to be above average have the option of driving manually and not using AVs. They will be

involuntarily exposed to risk due to AVs around them, but the risk presented by manual drivers in their vicinity would roughly equate to that of the 'average' driver, and therefore AVs that meet the 'at least as good as average' threshold would not increase their exposure to risk.

This, coupled with the benefits that adoption of AVs as soon as human-level safety is reached (as per the RAND model and white-paper) mean that requiring AVs to be at least as good as the average human is a favourable approach, although it must be acknowledged that it will result in controversy when AV incidents receive media coverage that is disproportionate to that for incidents caused by human drivers.

It should be noted that it has never been feasible to provide an accident rate comparison between a new transport system and its predecessor (e.g. railways versus canals, or aeroplanes vs railways) prior to that new transport system being deployed commercially on a wide scale. Therefore, it shouldn't be seen as feasible or reasonable to require a statistically valid comparison of accident rates for AVs and manual vehicles prior to AVs being approved. Furthermore, even if it eventually becomes feasible to make the comparison for all AVs versus all manual vehicles, it would be unprecedented to attempt to break down such a comparison by make and model, which would be necessary if such a comparison were to be used for the approval of individual AV types.

As such, the goal of safety being at least equivalent to an established statistical baseline should only be seen as a guiding philosophy in order to estimate a reasonable (yet still arbitrary) acceptance thresholds for scenario-based testing, bearing in mind that there is no mathematical means to directly interpret the test results in terms of projected real-world accident statistics.

Consultation Question 5 (Paragraph 5.120)

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

The principle of ALARP is well established within the UK, and has broad acceptance (from a technical and a legal perspective). Furthermore, as described in our answer to question 4, it would not be possible to approve purely on the basis of relative risk (in comparison to a baseline for current human drivers).

Given these considerations, we believe that demonstrating that good practice has been followed and that reasonable steps have been taken to manage each individual risk will form key component of regulatory approval.

For example, there will always be some residual risk in terms of rare edge cases that an AV can't consistently react appropriately to (whether this results in physical harm or in disruption of traffic flows). Where such cases are identified, it would be reasonable for them to remain as residual risk if (and only if) the level of the residual risk can be shown to be tolerable and it can be shown that it would be grossly disproportionate to modify the system such that the risk is eliminated. This is an example of ALARP being applied, and is also consistent with ISO PAS 21448 (Safety of the Intended Function).

There will also be a level of residual risk due to scenario permutations that are untested, either because that permutation has not been identified or because a particular factor was deemed unlikely to affect the AV (e.g. the colour of a building next to road) and therefore was not fully investigated in all combinations with other factors (e.g. lane width, bend radius, position of road sign

etc.). Such residual risk is inevitable since the real world is infinitely complex, whereas test programmes must be finite, but can be accepted if it can be shown that the test coverage is reasonable and that further exploration of parameters would be grossly disproportionate. Again, this is in line with ALARP. Regulators can assist developers by providing a database of what scenario factors need to be covered.

Consultation Question 6 (Paragraph 5.121)

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

Providing a database of what factors can be encountered within a scenario (e.g. what speed limits could be encountered, what's the most extreme camber the AV would have to negotiate, what heights road signs could appear at, as per answer to Q5) would help ensure the vehicle is developed and tested with consideration for all the challenges it could encounter. This database could include data that helps ensure all groups are treated equally, e.g. by documenting types of religious or cultural clothing that AV developers may otherwise have been unaware of, or how road crossing behaviour of visually-impaired people aided by guide dogs may differ from 'typical' behaviour, such that these permutations are engineered for. The actual engineering development itself would generally be done by the private sector and therefore not be subject to the public sector equality duty.

It should be noted that, in line with ISO PAS 21448 (Safety of the Intended Function), although residual risk can (and must) be accepted, it would not be acceptable if that residual risk is a systemic issue that would consistently manifest itself in certain permutations such that some people would be disproportionately exposed. For example, whilst it may be argued to be acceptable for an ALKS system to react appropriately in 99.99% of cases of a pedestrian stepping onto the motorway (99.99% being a relatively low threshold but arguably acceptable given the rarity of the event), it would not be acceptable if it was found that the 0.01% of failures were entirely due to the system being unable to identify a particular demographic (e.g. small children or an ethnic group).

CHAPTER 7: ASSESSING SAFETY PRE-DEPLOYMENT

Consultation Question 7 (Paragraph 7.99)

18.7 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?

Agree.

For a detailed description of how we think a scenario-based test programme could be conducted, see our white-paper 'Test Methods for Interrogating Autonomous Vehicle Behaviour – Findings from the HumanDrive Project' (<https://humandrive.co.uk/downloads/>), especially chapters 2 and 3. Section 3.1 provides an analysis of how different test environments (simulation, proving ground etc.) could be combined to make full use of their relative strengths, with the different 'layers' within the ADS being interrogated separately for a large proportion of the testing in order to control the overall volume of test cases required.

We agree that self-certification would be the most feasible approach to cover the bulk of test scenarios, with a regulatory audit providing a level of assurance that an acceptable level of diligence is being applied but not being able to assess each test case in detail due to the sheer number of scenario permutations, and therefore not being able to 'approve' each and every scenario. The audit would include confirmation that the accuracy of any simulation data used as evidence within the safety case has been validated against real-world data from physical tests.

However, a smaller scale test programme should be undertaken by a 3rd party (the regulator or a technical service representing them) in line with point 3d in the question; whilst it wouldn't be possible to test more than a small proportion of scenarios and therefore wouldn't be possible to confirm safety to be statistically greater than that of existing manual vehicles, nonetheless it would allow systems that fall a long way short of acceptability to be kept from the market. This is essentially the role of type approval of current vehicle systems/components - the regulations do not allow a comparison between high-performing alternatives (which is the role of consumer programmes such as EuroNCAP), but provide a means of keeping exceptionally poor systems from the marketplace.

We agree that acceptance should be based upon a safety case as the level of technical maturity and stability is unsuitable for prescriptive regulations, and therefore developers need flexibility to put forward a bespoke safety argument. This may change in the long-term future.

Consultation Question 8 (Paragraph 7.100)

18.8 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.

It would be appropriate to consult road user groups; the more sources of information available to identify a comprehensive range of possible scenario permutations, the lower the probability of hazardous system limitations remaining uncovered due to the scenario permutation that triggers the hazard remaining undetected.

However, caution must be taken not to rely solely upon groups with a particular area of interest, as this would mean scenario permutations that are not a concern to any of the road user groups that engage with the process would not be captured. The main source of information should be through collecting and analysing real-world driving data - this could be seen as equivalent to doing multi-billion miles of validation as per the RAND 'How Many Miles...' report, but doing it once in order to extract the data for a centralised repository, as opposed to doing it for each vehicle.

CHAPTER 8: INITIAL APPROVALS AND CATEGORISATION – PROPOSALS

Consultation Question 9 (Paragraph 8.17)

18.9 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?

Broadly agree. However, the requirement for trials to gain an exemption from the Secretary of State under section 44 of the RTA should only be required for trials where there is not a conventional safety driver physically present in the vehicle and able to override via 'traditional' vehicle controls; otherwise there would be an excessively onerous burden on both trialling organisations and regulators with regard to relatively simple and safe trials. The above is in line with the guidance on 'Advanced Trials' within the Code of Practice.

Consultation Question 10 (Paragraph 8.25)

18.10 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?

Agree. However, the UK domestic scheme should look to either accept UNECE approvals for systems/ components or to mirror them closely wherever such alignment is possible. This would minimise the duplication of effort for developers who wish to use both schemes (e.g. starting off in the UK but later seeking UNECE approval for wider deployment, or starting with some systems that have already been approved under UNECE regs for a lower level of automation and then seeking to add more advanced automation with the UK) and would aid efforts towards international standardisation. As such, the two approval mechanisms shouldn't be seen as entirely separate.

Consultation Question 11 (Paragraph 8.43)

18.11 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?

Agree

Consultation Question 12 (Paragraph 8.44)

18.12 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 don't have direct experience of how the appeals process works, so cannot provide detailed feedback on its suitability. From the brief description given, it sounds reasonable, although the scale and complexity of an ADS approval will introduce new challenges (e.g. it might be difficult to include all the required areas of expertise within the appeals panel whilst also ensuring that all members are 'independent from the work area to which the appeal concerns').

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?

Disagree.

The intended level of automation (driving assistance, automated with user-in-charge, automated without user-in-charge) will be intrinsic to the testing undertaken within the type approval process. Different levels of automation would require fundamental differences in aspects such as:

- the range of scenarios that would need to be covered (e.g. all ODD permutations that could be managed by handing back to the driver at lower levels of automation would require the system response to be validated for higher levels of automation);
- the total volume of test scenarios deemed proportionate;
- the oversight of safety management processes;
- the system redundancy required to accommodate technical malfunctions;
- the assessment of the human-machine interface, and;
- the pass/fail criteria for individual test cases and for the overall test programme.

It would therefore be absolutely essential that the level of automation is defined at the point of conducting the technical approval.

It may be feasible for a vehicle that was subjected to a regulatory process for a higher level of automation to be downgraded based upon test results. However, this would mean developers incurring massive costs for performing a test programme that is grossly disproportionate to the available use cases in the deployed vehicle, and would presumably normally result in the system being revised and resubmitted rather than being deployed at a lower level of automation. Downgrading the system to a lower level of autonomy would also require some new tests to be undertaken that wouldn't have been applicable to a higher level, e.g. tests upon handovers.

It would not be possible to increase the level of automation beyond that originally targeted and tested for, and it would not be feasible to require all systems to be tested to the highest level of automation (the required investment would only be acceptable where developers have a realistic chance of being certified at that level of automation and have a business case to be able to recoup the investment). This leaves limited flexibility to adjust the level of automation in the late stages of the approval process.

The concept of the Secretary of State providing an additional, final level of approval with regard to the ability to operate safely within the particular environment presented by that jurisdiction, and according to the risk tolerance of that population, is a good solution to allow regional considerations to be factored in before an AV can be deployed, and allows a single process to be used for this regardless of whether the prior technical approval was to the UN ECE or domestic process. However, it is difficult to foresee how it would be feasible to reach this point without having defined the level of automation long before.

Consultation Question 14 (Paragraph 8.77)

18.14 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?

Agree

Consultation Question 15 (Paragraph 8.78)

18.15 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?

Partially agree.

There needs to be an appeals process as there will no doubt be contentious or marginal decisions where further examination is needed. However, the question wording refers to a categorisation decision; as per our answer to Q13, the technical assessment that the vehicle would need to be subjected to is inherently linked to the targeted level of automation, and it would be grossly disproportionate to subject all systems to an approval program rigorous enough to allow the highest level of automation such that the Secretary of State can pick and choose a level later.

Therefore, the target level of automation should be defined by the developer early in the process (ideally with regulator input to ensure that decisions are realistic and investors are not misled), with the final decision both at the type approval and the secretary of state stages being a decision on whether the vehicle provides acceptable performance at that pre-defined level of automation. In the event of failing, the developer could improve the system, appeal the decision or re-apply for approval at a lower level of automation; in the latter case, most of the technical assessment would not need repeating, although there may be some new test requirements that come into play regarding human factors of handovers, driver controls etc. for lower automation levels.

Consultation Question 16 (Paragraph 8.83)

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

Agreed. Pharmaceutical trials form a good reference point for AV approvals, as the complexity of the system (human body or ADS) and of the influences from the surrounding environment mean that it is not possible to make decisions purely based upon controlled experiments.

It's worth noting that such an approach to trials will present new challenges in terms of schedules and budgets; without professing to have any more than a layperson's knowledge regarding medical trials, they appear to consume a far greater proportion of the overall product development budget and require a far longer duration of testing after the technical phase of development has been completed, compared to automotive type approval. As such, AV developers would have to be ready to allocate large sums of money towards the approval programme (on top of the high cost of

engineering an ADS), and could no longer expect to bring products to market in short and predictable timeframes.

Limited scale initial deployments would not allow a statistical demonstration of greater safety than manual drivers or to form the sole basis for safety assurance, as there would be insufficient exposure to rare edge cases. However, it would allow systems that are clearly a long way below acceptable standards to be identified before further rollout increases the exposure to risk, as a pattern of frequent mistakes would soon be observed.

Furthermore, the small-scale rollout would give a good idea of how the system will perform in more typical scenarios (non-edge cases) and how the public will typically interact (e.g. whether an overly cautious system causes traffic queues behind and a rise in unsafe overtaking). This would allow public opinion to be factored into further decisions (bearing in mind type-approval testing and the creation of regulations will be a largely technocratic exercise, from which the communities affected by AVs may feel disconnected). Feedback from the early stages of the progressive rollout would allow systems to be improved, operational procedures to be updated, regulations to be amended, or in extreme cases, systems to be prohibited from further use.

As per medical trials, it would be appropriate for the process to be overseen by a panel who are able to provide oversight of ethics and trial safety, together with ultimately deciding whether the product is accepted for wider deployment. Although various 'oracles' can be used to assess whether an ADS performs adequately (e.g. RSS, as discussed in the consultation paper), it will not be possible to produce an oracle that gives a definitive judgement on all possible test cases and outcomes. As such, the oracle should be seen as a filtering mechanism to highlight concerning or borderline cases for the panel to review, thereby using partial automation to arrive at a volume of data for which human oversight is feasible. The oracle(s) used, and the standards by which panels make their judgements, should be identical for type approval and in-service monitoring to ensure systems are judged consistently.

CHAPTER 10: ASSURING SAFETY IN USE

Consultation Question 17 (Paragraph 10.82)

18.17 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?

Partially agree

There should be a scheme to monitor AVs post-deployment, although it is very important that they have been robustly assured already prior to deployment. This must be against a very similar set of requirements (e.g. any 'oracle' that determines whether the driving behaviour was acceptably safe within a scenario should be the same pre- and post- deployment), to avoid the possibility of an ADS passing during type-approval but then failing when in use.

As such, the in-service monitoring should be seen as validation of the previous assurance, rather than a separate assurance process. This is in line with existing conformity of production testing for emissions, the difference being that the principle would be applied on a wider and more complex scale. The in-service validation should look to confirm that results are acceptably similar to those

obtained/ required during type approval, but with acceptance that there will always be a tolerance between the two and that it would not be acceptable to remove a type-approved vehicle from the market unless there was a significant discrepancy or safety concern.

Consultation Question 18 (Paragraph 10.83)

18.18 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?

Partially agree

(1) Scheme regulators should monitor in service performance to ensure there are no major discrepancies relative to the original type approval results and no major safety concerns. However, as per our answer to Q17, this should be seen as validation of the previous results, and developers should have confidence that once they've obtained approval, their product won't suddenly be removed from the market unless a serious defect or the use of 'defeat devices' is uncovered.

Care should be taken in using real-world statistics for approval of a system, as even using 'leading measures', a large volume of data would be needed to draw statistically-significant conclusions, bearing in mind the extreme variation in the real world; there should not be an impression that vehicles are passing and failing due to random chance or situations that are outside the developer's control.

(2) Leading measures should be collected to allow more data to become available sooner, and lagging metrics should be collected since these are generally the metric society will be more concerned about. As large volumes of lagging data become available, the leading data will become less significant for determining the safety, but leading data should still be collected to validate (or correct) original assumptions about how leading and lagging measures are correlated (e.g. how many lane boundary excursions equate to each serious injury on average), such that leading data is able to be more effectively used in subsequent approval programmes.

A previous answer referred to the concept of doing the multi-billions of miles driving (as per the RAND "How Many Miles...." report), but doing it once to collect data on scenario permutations to support scenario-based testing as opposed to doing it for each vehicle (in practice, the data would require ongoing 'topping up' to account for changes over time). However, this data collection should also include collecting data on leading and lagging data, bearing in mind the limitations of existing data sources such as STATS19 (especially with regard to leading measures).

(3) Regulators should have the power to require the updates listed. The first in the list, correcting safety and legal compliance issues may be the most challenging, as in some cases it may not be clear whether a technical solution would be available. Where updates are available, testing should not just confirm the defect is corrected but also that other effects haven't been introduced, e.g. a correction for failing to stop for dogs in the road would not be acceptable if it causes the vehicle to become overly cautious and hold up traffic by stopping for paper bags.

Consultation Question 19 (Paragraph 10.84)

18.19 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 should also deal with cybersecurity?

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

It seems not entirely unreasonable for UK-only updates to be able to be approved by the UK regulator without involvement of the original authority that performed the type approval, although careful consideration would need to be given to the practicalities once more detail is known about the regulatory structure to work. For example, would sufficient technical information about the original approval be available to the UK regulator, in a format that could be readily understood, such that they would be in a position to assess what has been changed and what the possible implications would be? Would overseas approval authorities or the UNECE accept this, given that the updates may invalidate the approval to UNECE regs for the UK-market vehicles, with the UK therefore permitting vehicles with an invalid UNECE approval on the basis of a local approval overriding aspects of the UNECE one? Would there be confusion if further updates were needed and some aspects of the same regulation are approved by different authorities in different jurisdictions? Would the original approval authority maintain some sort of record of the local update the UK had approved, and if so, by what mechanism would the information be shared and recorded? In practice, therefore, it could prove complex and controversial to implement.

Given the proposed structure of the initial type approval being done either domestically or internationally to UNECE regulations, and then a domestic approval being applied in a second stage that considers the specific nature of UK roads, a more practical approach might be for any updates that are required specifically for the UK market to be reassessed under the latter phase of the approval system only wherever possible. Where updates would invalidate the UNECE approval, the developer would have to go back to the original authority to certify it as a separate variant of the type. This is similar to how different versions of a type co-exist within the existing type approval system. It could result in a greater admin overhead initially, but would probably be clearer and simpler in the long run.

It is very important that the regulators have powers to both enforce and review cybersecurity updates where they are required. However, it may not always be necessary to update approvals, and would not in general require any retesting of the functional behaviour (such as scenario-based testing) of the vehicle, only testing of the cybersecurity aspect to confirm that the specific vulnerability had been addressed and no new vulnerabilities introduced.

Powers to require and approve updates should include issues due to system faults (functional safety), design limitations (safety of the intended function, SOTIF), cybersecurity, changes to road traffic laws, map updates or issues identified with human factors/the HMI. Note that some of these aspects will change over time, e.g. SOTIF may be achieved initially but then a new type of vehicle or infrastructure may be introduced that highlights a previously unobserved system limitation.

Consultation Question 20 (Paragraph 10.100)

18.20 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?

The consultation sets out the relative merits of each approach well, and in line with this, one approach is not demonstrably superior to the other in the long term. In the short term however, it may be expected that approvals will consist of a progressive rollout of partial automation within production vehicles, according to the 'path 1' route, and it seems unlikely that the initial volume of these approvals would be sufficient to justify creating an entirely new system, given that the current system provided by VCA and DVSA is acceptably effective.

Once experience has been gained in low volumes, the question should be reviewed again in the light of any issues that may come to light, at which point it may (or may not) become clear whether a new structure is required.

Consultation Question 21 (Paragraph 10.101)

18.21 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)?

Care should be taken in requiring external views to be incorporated, such as advisory committees, as it may be difficult to avoid leakage of intellectual property, and any non-public sector organisations are at risk of bias due to allegiances with particular companies or technical approaches.

There should be a means for members of the public and industry to raise concerns to the approval authority. For example, if vehicles appear to be breaking laws or acting antisocially (e.g. clogging up roads by cruising at low speeds between paid journeys), such intelligence should be able to be collected. The DVSA process for reporting vehicle defects would be an appropriate starting point, which could be further developed if necessary, as the rollout of AVs increases.

CHAPTER 11: INVESTIGATING TRAFFIC INFRACTIONS AND COLLISIONS

Consultation Question 22 (Paragraph 11.24)

18.22 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?

Agree

Consultation Question 23 (Paragraph 11.53)

18.23 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.

Do you agree?

Agree

Consultation Question 24 (Paragraph 11.54)

18.24 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?

Agree

Discretion will be essential while the technology and its applications remain relatively novel, as it may reasonably be expected that there will be safety and regulatory compliance issues, and consequences to affected parties, that are currently unforeseeable and therefore difficult to prescribe legal remedies for.

Consultation Question 25 (Paragraph 11.69)

18.25 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?

Agree

However, the remit of focussing upon "the most serious, complex or high-profile collisions" could potentially be too narrow. The selection of which cases to pay closer attention to should be based on the overall risk to the public, bearing in mind likelihood as well as severity. Therefore, less serious, complex and high-profile collisions would be worthy of investigation if there is a pattern in the high-level accident stats indicating that there might be an issue that a large number of people are exposed to.

Furthermore, if there is a pattern in the data such that the overall risk is low but the risk to a specific subgroup is disproportionately high (e.g. people journeying past a specific location that exposes a flaw, pedestrians with particular characteristics who are poorly classified by the sensors), it would be an unacceptable hazard that needs further investigation.

As such, decisions about prioritising resources for accident investigation should be risk based, rather than profile based.

Consultation Question 26 (Paragraph 11.82)

18.26 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?

Agree

Consultation Question 27 (Paragraph 11.83)

18.27 We welcome views on:

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

(1) The remit of the forum should include any aspects related to safety of AVs, and also to ensuring they operate in a manner that is considerate to other road users, the environment etc. Reviewing existing Highway Code requirements to agree more suitable requirements for AVs would be a reasonable starting point, which further discussions would no doubt branch out from. However, care should be taken not to be overly prescriptive about the priority for areas to review, and therefore

the forum should maintain a level of control over its own remit, and should commence by identifying which areas need attention most urgently. It would be appropriate for a representative of CCAV or DfT to chair the forum, to ensure neutrality and to ensure discussions represent the public interest.

(2) The forum should include representatives from developers (including start-ups, major suppliers and OEMs), consultancies involved in AV safety, DfT, CCAV, VCA and other public sector agencies involved in managing the road network (e.g. Highways England, Transport for London). These are broadly the same organisations involved in existing standards working groups, although it should be noted that contributions to such groups tend to under-represent organisations involved in the technical development of AV technology, so care should be taken to ensure a balanced view with a suitable level of expertise. Technology developers tend to get more heavily involved in regulatory working groups when they have, or will shortly have, products on the marketplace, so it may be expected that the balance of contributions will improve over time.

(3) As discussed in a previous answer, there should be a means for the public to report concerns about AV behaviour, which could work on a similar basis to the DVSA's existing process for reporting vehicle defects. This could allow more general concerns to be raised about how AVs operate as a whole, in addition to concerns about a specific type of AV. The regulator (likely to be VCA, at least initially) could bring these concerns to the forum, particularly where a pattern of complaints emerges, and this could influence the priorities of the forum's work.

Beyond this, it's difficult to see the public having direct involvement, as it wouldn't be feasible to have the public represented within discussions. Any major changes, as opposed to more algorithmic clarifications of existing Highway Code rules, could be put to the public via a consultation much like this one, although it seems unlikely that enough of the public would engage to represent anything approaching a statistically-significant sample, and many would have insufficient technical or legal knowledge to contribute, so this is unlikely to be a major influence.

On balance, therefore, it doesn't seem feasible for the public to have a direct voice in the process, as it should be led by those involved in the CAV community, but the interests of the wider public can be indirectly represented by government departments such as CCAV and DfT.

CHAPTER 12: THE USER-IN-CHARGE

Consultation Question 28 (Paragraph 12.24)

18.28 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?

Partially Agree

Maintaining direct line of sight is an important criterion, but some form of distance limitation should also be imposed within the legislation, as it would clearly be unreasonable for a user-in-charge to be located several hundred metres away on the horizon. The suggestion given previously in paragraph

12.15 of around 20m, or the distance at which a number plate can be read, seems reasonable. Such a requirement should therefore be added to the proposed text in the question.

There should also be a speed threshold above which the user-in-charge must be inside the vehicle, as adopting a remote position would only be appropriate for low speed manoeuvres such as parking. Such a clause would preclude permutations such as the vehicle driving on the motorway with no occupants but with a 'user in charge' monitoring from a second vehicle following behind. It could be argued that this is already covered by the requirement to be in "a position to operate the controls", but the intent of this clause needs further clarification - for example, is a person stood beside a vehicle that is parking in a position to operate controls, bearing in mind that they would have to enter the vehicle first? Or would being able to press an emergency stop button on a key fob or phone app be sufficient to satisfy the clause? The latter seems reasonable, but only if the speed is limited.

Consultation Question 29 (Paragraph 12.37)

18.29 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?

Agree

However, this approach is only acceptable on the assumption that the handover period would always be sufficiently long to expect a reasonable person to take control safely. It is concerning that in the ALKS consultation, there was an allowance that the handover could be less than the required minimum of 10 seconds in the event of a technical fault.

To be classed as an AV under the AEVA, the vehicle should be required to have sufficient robustness and/ or redundancy in all safety critical systems such that control is always able to be maintained for a suitable handover time (with remaining failures being sufficiently rare that the risk is acceptable, and with such failures representing illegal performance by the system such that the driver is not held responsible). If such handover times cannot be ensured, it is not fair to expect the driver to react or to penalise them for failing to do so, and therefore such systems should be regarded as level 2 assistance.

In short, therefore, the driver should only be liable if the handover time is no less than a legal minimum, and if that legal minimum is reasonable.

Consultation Question 30 (Paragraph 12.45)

18.30 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.

Holders of provisional licences should be allowed to act as a user in charge provided that they are accompanied by a qualified driver. It's difficult to imagine it will be a major component of driver training in terms of time spent in autonomous mode, as the role will mostly be passive, meaning learners will not be getting much value from their lesson! However, targeted training to deliberately trigger a handover so learners can get experience of the key part, e.g. by enabling the system shortly before the end of its geofenced region such that a handover request is sure to be imminent, or driving instructors having a button available to trigger a handover request at a time of their choosing (allowing them to select an opportune moment with no hazards present), could provide effective training.

It's difficult to predict whether responding to a transition demand would be more challenging or hazardous compared to the many hazardous scenarios that learners can encounter whilst driving manually, so there doesn't appear to be evidence to justify requiring the lesson to be with a professional driving instructor in a car with dual controls, as opposed to any licence holder in any car. It can, of course, be argued that requiring the higher bar would, on average, make this portion of driver training safer, but the same argument could be applied to any driving scenario, for which it is currently not deemed proportionate. Use of AVs on motorways would presumably be limited by the same rules that currently apply to learners driving manually.

Consultation Question 31 (Paragraph 12.53)

18.31 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?

Agree

Consultation Question 32 (Paragraph 12.59)

18.32 We provisionally propose that persons carried without a user-in-charge should be guilty of a criminal offence.

Do you agree?

Agree

With caveats about knowing a user in charge is not present but should be, in line with the next question.

Consultation Question 33 (Paragraph 12.60)

18.33 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.

Partially agree

Both clauses (1) and (2) should have the "knew or ought to have known" wording that is currently only in (2). This would be in line with the suggestion from the Law Society of Scotland in paragraph 12.56, and would address the scenarios presented in 12.55. Otherwise, the wording is appropriate.

Consultation Question 34 (Paragraph 12.66)

18.34 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.

Partially agree

The concept of recognising the driver was in control at the time seems more representative of the facts than to class the ADS as in control when it was not, so this aspect seems proportionate. However, care needs to be taken in how statistics are collected and how verdicts are reported, as it would be unrepresentative for the incident to be counted within the statistics for manual driving and not in the statistics for automated driving.

The inclusion of some form of defence based on the 'reasonableness' of the person's response is, of course, essential to avoid grossly unfair blame and penalties being applied. However, the proposed wording of "a competent and careful driver could not have avoided..." sets too high a bar - there will be many incidents where successful intervention is reasonably possible, and therefore the "could not" threshold is not reached, but where successful intervention is by no means a certainty and hence on many occasions a competent person acting reasonably would still have a collision. It would not be fair for the latter person to be punished.

An analogy would be asking a person to throw a ball at a small target 20 metres away. A competent person 'could' hit the target, even if only for a small percentage of attempts, but it would not be fair to say they were negligent on the failed attempts.

The approach of examining reasonableness, as per paragraph 12.64 point 2, therefore seems a better test to apply.

Consultation Question 35 (Paragraph 12.94)

18.35 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?

Agree

Consultation Question 36 (Paragraph 12.95)

18.36 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?

Agree

It should not be assumed that members of the public will be familiar with the list of which roadworthiness offences they are responsible for. The information should be incorporated into the Highway Code, although even this will have limited reach - few drivers read the Highway Code after passing their test. It is therefore important that the offences are chosen to be intuitive, such that it can be expected that a reasonable person would, through common sense, realise they are responsible for oversight of the required aspects. This may be difficult to do precisely in practice, as drivers vary greatly in their diligence, but efforts should be made to ensure the secondary legislation is not perceived as unfair or counter intuitive.

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

Consultation Question 37 (Paragraph 13.67)

18.37 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?

18.38 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”.

Agree

In practice, it may be something of a moot point whether systems requiring continuous remote supervision (i.e. the remote equivalent to level 2) would be classed as automated under the AEVA, since it's difficult to foresee a business case where combining the costs of developing the autonomy, developing the remote communications to allow supervision/ intervention and employing a safety

operator to continuously monitor the system would be more favourable than employing a traditional manual driver. The only exception to this would be trials of the technology in the run up to type approval.

It is also difficult to envisage a safety case where continuous remote monitoring would be the best option, other than (arguably) low speed pods where the only intervention needed would be an emergency stop.

It seems reasonable to regulate remotely supervised AVs on the same basis as fully autonomous vehicles, due to the complexity and inherent risk in relying upon such monitoring systems. The proposed new wording for the AEVA seems appropriate.

Consultation Question 38 (Paragraph 13.86)

18.39 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-incharge 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?

Agree

The proposal is much simpler and results in more consistent requirements being imposed relative to the arbitrary six-month cut-off.

Consultation Question 39 (Paragraph 13.92)

18.40 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

There should be some means of demonstrating that the operator is competent, but this competence may be divided across many individuals, so the key will be demonstrating that professional processes are in place. As such, the approval might look more similar to an ISO 9001 certification that a transport managers' exam.

Consultation Question 40 (Paragraph 13.108)

18.41 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?

Agree

These requirements should, of course, only apply to NUICS and HARPS, as for lower levels of automation that require the presence of a UIC, that UIC can and should be responsible for some of the duties listed.

Consultation Question 41 (Paragraph 13.109)

18.42 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?

Agree

Such flexibility should be used as widely as possible as it is very difficult to predict what use cases, business models and unforeseen hazards will emerge.

Consultation Question 42 (Paragraph 13.116)

18.43 We welcome views on how accessibility standards for Highly Automated Road Passenger Services (HARPS) might be developed. 18.44 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? We welcome views on what the set interval for periodically re-consulting the accessibility advisory panel should be.

Partially agree

It should be ensured that HARPS services accommodate those with disabilities where reasonably practicable, but care must be taken to ensure that the standards imposed are practical from a technical and an economic perspective. These are considerations that developers are best placed to advise upon, so it should be ensured that they have the opportunity to inform the regulations. It will be very difficult to legislate initially, as it's difficult to foresee what HARPS services will look like, so it would be worth considering a code of practice as opposed to regulations in the first instance.

It is important that industry supports mobility for those with disabilities, but there is a risk that the measures required could prevent or delay rollout. If AVs are shown to be safer than manual vehicles, this delay would result in more people being killed and seriously injured. Such fundamental safety considerations should receive even higher priority than ensuring equality, and therefore it must be ensured that accessibility requirements don't unreasonably inhibit the rollout of AVs.

Consultation Question 43 (Paragraph 13.133)

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

There are pros and cons for either option, as discussed, but the Traffic Commissioners would be less likely to be subject to regulatory capture. Furthermore, it may be expected that there will be few HARPS providers and services initially, with rollout progressively increasing, and it may be difficult to justify setting up and training a new organisation with new processes when they have a limited number of operators to deal with.

It therefore seems more proportionate to make use of the existing expertise and processes of the Traffic Commissioners initially, and then learning from early deployments can be used to update processes or to justify creating a new scheme under a new body in the future if appropriate.

CHAPTER 14: CRIMINAL OFFENCES BY ADSES AND THEIR SENIOR MANAGERS

Consultation Question 44 (Paragraph 14.107)

18.46 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?

Partially agree.

The described approach has merit, but this is true for any industry where harm can be caused by the actions of an organisation; there is no reason why legislation to fill the 'accountability gap' should be specifically targeted at the CAV industry without other industries being equally addressed. A person killed as a result of a collision caused by a CAV is no more or less dead than a person killed by a fire resulting from faulty electrical goods, a train crash, a bridge collapse or contaminated food. Or, for that matter, a traditional manually driven vehicle that has a dangerous defect. Whilst it is fair that conduct which the public would be outraged by should be prosecuted, it would not be fair to hold CAVs to a higher standard.

Therefore, whilst changes to legislation are needed, this should be done as part of the wider, cross-industry review of the accountability gap referred to in the document, and should treat all industries equally. The consultation paper highlights inconsistencies between legislation applicable to different industries, and the aim should therefore be to arrive at a more harmonised and consistent approach that applies to all industries that could reasonably be expected to cause harm, rather than separate measures to target a few high profile industries.

The legislation for other industries seems to be fairly consistent in applying a maximum sentence of two years, so it's difficult to understand why the CAV industry should be singled out for higher penalties (paragraphs 14.10, 14.17 point 6, 14.90) - again, victims of fatal CAV collisions are no more dead than victims of fatal accidents caused by other industries. Whilst there is evidence of a public appetite for harsher sentences in response to death resulting from manual driving behaviour, similar outrage exists for examples from other industries - for example, the Boeing 737 Max MCAS system faults or the installation of flammable insulation tiles on tower blocks. There may therefore be justification in raising the maximum sentence for all such cases, but increases should be consistent across industries, and should not single out CAVs.

The concept that non-senior employees would be exempt from prosecution seems unfair as the context needs to be taken into account - if the actions of junior staff are the result of senior management pressure, direct instructions or a toxic culture, they should not be blamed, but there will be cases where staff have sufficient responsibility such that management could not reasonably be expected to oversee their work in detail, and they commit unlawful actions entirely of their own volition. For very junior staff, such as graduates and apprentices, sufficient oversight should be in place and therefore the organisation can be held accountable, but there are a large body of engineers within any industrial organisation who are senior enough to have a high level of autonomy and low level of supervision, whilst being far below what could reasonably be termed 'senior management'. From the description in the consultation document (e.g. para 14.98), a senior manager putting their signature upon a safety case would be putting themselves at excessive personal risk given that they cannot reasonably be expected to have digested and understood the full content.

Consultation Question 45 (Paragraph 14.108)

18.47 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.

Agree

However, as per response to question 44, care needs to be taken to ensure that the requirements placed upon CAV developers, and the sentences available, are not disproportionate to those placed upon other industries.

Consultation Question 46 (Paragraph 14.109)

18.48 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.

Partially agree

Paragraph 14.105 is concerning - it would be perfectly reasonable to reject an approval application on the basis that the information provided is insufficiently clear, but the regulator should be able to make this judgement and request changes at the application time. The idea that a person could receive up to a year in prison for producing a document that is subjectively judged to be insufficiently clear (for such a decision could only be subjective) is extremely alarming!

The requirement for clarity should therefore be part of the regulatory system, but as part of the type approval process as opposed to in-service legislation. The exception would be where unclear data is deliberately used to hamper in incident investigation, as this could not be dealt with by simply rejecting an application. However, this should be legislated for carefully, and should be specific to incident investigations being provided with unreasonably inadequate data, so that developers of the original safety case during the approval process cannot be imprisoned if their work is judged unclear.

CHAPTER 15: NEW WRONGFUL INTERFERENCE OFFENCES

Consultation Question 47 (Paragraph 15.10)

18.49 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?

Agree

Consultation Question 48 (Paragraph 15.11)

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

Yes. External infrastructure is, in general, less of a safety concern, as any incorporation of offboard systems would have to be designed to include suitable system redundancy, plausibility checks on data etc. to ensure functional safety and cybersecurity are achieved - it is not possible to rely on wireless communications to the same degree as onboard systems when it comes to real-time

processing of safety critical tasks. However, there may be other tasks that could be interfered with that could have safety implications, e.g. jamming of wireless signals preventing an over-the-air software update, or tampering with map data to affect the route the vehicle takes. Furthermore, there could be instances where there is a flaw in the functional safety and/ or cyber security such that it becomes possible for a skilled (or lucky) attacker to exploit the vulnerability and directly affect the dynamic driving task.

It therefore would be appropriate for legislation to cover such attacks, particularly for any instances where the attack is (or could plausibly be) safety critical.

Consultation Question 49 (Paragraph 15.53)

18.51 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?

Agree

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?

Agree

Consultation Question 51 (Paragraph 15.62)

18.53 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.

Partially agree

As a point of interest, looking at the logic of the clause from the perspective of Boolean algebra (i.e. from a software engineer's perspective), it is unambiguous. It is an offence if the following is satisfied: NOT('lawful authority' OR 'reasonable cause'). This can be rewritten as: (NOT 'lawful authority') AND (NOT 'reasonable cause'). In other words, it can only be an offence if there is both an absence of lawful authority and absence of reasonable cause. This is in line with the preferred interpretation suggested in the consultation, and on this basis the literal interpretation of the law should be unambiguous (although in practice, this would of course depend on the way it is presented to the court in a particular case, and the way the court understands that evidence).

As noted, introducing a clause to deal with the incorrect literal interpretation within the narrow CAV context could be used as justification for taking a purposive interpretation that the law was intended to be interpreted the 'incorrect way' outside this narrow context, which would be concerning.

In the example in 15.58, there is a strong argument to suggest that there would be 'reasonable cause' for installing the update if the update was intended to, and thought to, improve performance, so even with the 'incorrect' interpretation, this wouldn't necessarily be a problem.

On balance, therefore, it appears that specifically adding the defence may not be necessary or desirable. However, this opinion is given with the caveat that it is a relatively nuanced legal issue which should ultimately be decided by legal specialists, whereas our feedback is from the perspective of technical specialists.

CHAPTER 16: CIVIL LIABILITY

Consultation Question 52 (Paragraph 16.24)

18.54 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?

Agree

Consultation Question 53 (Paragraph 16.32)

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

Do you agree?

Agree

Consultation Question 54 (Paragraph 16.47)

18.56 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?

Agree

CHAPTER 17: ACCESS TO DATA

Consultation Question 55 (Paragraph 17.65)

18.57 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?

Agree

Consultation Question 56 (Paragraph 17.71)

18.58 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?

Agree

Consultation Question 57 (Paragraph 17.81)

18.59 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?

Agree

Consultation Question 58 (Paragraph 17.95)

18.60 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?

Partially agree

The ADSE should be required to provide evidence that there is an appropriate process in place to control data, and adopting a similarly flexible approach to that described for the safety case would be appropriate at this point where many uncertainties remain about use cases and risks.

However, data management sits outside the scope of safety, and it would generally be different people with different expertise who are in a position to develop or scrutinise a data protection policy when compared to a safety case. It is therefore better to avoid conflating the two issues and to give each a separate focus within separate documents. This will also be advantageous when updates are needed to either document (or set of documents; a safety case would typically consist of many separate pieces of evidence with a 'safety argument' to show how it all fits together), as it would minimise the amount of material that would have to be reviewed and the number of people who would need to be involved in the process relative to having all the information within a single document.