

Law Commission Consultation

Document 3

*Private Response by Richard Morris,
Innovation Lead for CAV at Innovate UK*

Introduction To Responses

Firstly, as before, my responses are personal and do not necessarily reflect the positions of my employer: Innovate UK, nor UKRI, nor other individuals in those organisations.

The consultation document, like the previous two, is an excellent piece of work which encapsulates a substantial number of issues in a coherent way. It builds on the previous documents and responses in a logical and constructive manner. It asks a number of questions in the body of the text and these are repeated at the end. Most are asked as specific questions requiring yes or no answers, some request views. I have provided views where requested and also provided views where I only partially agree to the questions posed, to explain my thoughts.

Hopefully this is regarded as helpful.

I have answered the questions in order as I read through the document. Due to the interconnections of some subjects, I address some issues in advance of some of your questions in later sections. Rather than completely rewriting my answers, I refer back to my previous responses where appropriate. My apologies if this causes any issues. Sometimes my answers have a different focus or bias in different sections which cover related areas – I hope I do not contradict myself too much!

There are some apparent contradictions in my responses. I fundamentally disagree with “Level 3” and the idea that there will be consistently safe handovers of control from the ADS to the driver with the monitoring measures available today. I therefore strongly advise against this route forward. However, I recognise that there is substantial political pressure to accept ALKS which allows this. Some of my responses reflect my opinions against it whilst some assume that “Level 3” (ALKS) will happen in practice, and that we should try and make the best of a bad concept. I hope that I am being over-cautious but have grave doubts about the ALKS regulation and in particular it’s vagueness in application.

In order to navigate the document efficiently, the Chapter titles and question titles are in **dark blue** and the questions you posed are in order in **black**. My responses are all in **purple**.

For clarity, the answers the questions are in chapter sections with each chapter starting on a new page.

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;

Agreed

(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;

Agreed – but I have severe doubts as to whether the concept of “user in charge” is viable – please see my answer to c) below.

(b) provides clear visual, audio and haptic signals;

Only partially agree. Clarity is a widely used term which is often subjective. e.g. A warning icon may be “clear” in that it is large, bright and colourful – but the image on it may not convey its meaning particularly well. The message must be unambiguous as well as clear. Similarly a “beep” may be a clear pure tone and loud – but what does it mean? What message is intended to be conveyed?

The signals must suit the specific user in charge – both in terms of the type of medium, and the level required to suit their cognitive ability.

Training may help but ideally the signal and meaning should be obvious without training being required. Perhaps the user in charge can choose a suitable combination from a selection, together with brightness or volume or haptic vibration amplitude? This area of conveyance of meaning is one in which VMs have a poor track record for ADAS functionality and we cannot afford the same mistakes to happen again with these transitions.

An alternative is to mandate a standard warning signal, with a visible indicator being defined such as those in Regulation 121; the “Identification of controls” regulation. The colour, illumination requirements and icons being specified accurately and mandated.

(c) gives sufficient time to gain situational awareness;

I disagree because “Sufficient time” is never possible to quantify. It is certainly not a simple number; neither 10 seconds nor 45 seconds, but is context and driver state dependent.

To expand on the former, it may be helpful to consider different example contexts where a transition is required:

1) Motorway junction – An ADS which may only operate on a motorway should be able to determine its approach to its intended exit junction and provide a transition demand in plenty of time – perhaps a progression of “information” to a “request” before a final “demand”. Failure to respond to the demand may not be too problematic – it might mean the ADS simply continues driving at normal speed to the next junction instead, not requiring any sort of hazardous stop, (so

neither in lane nor on the hard shoulder, both of which are dangerous, though one more than the other)

- 2) Motorway end – Again, an ADS which may only operate on a motorway should be able to determine its approach to the end of a motorway and again provide a transition demand in plenty of time – perhaps an “information” or a “request” before a “final” demand. Failure to respond to the demand **will** require the ADS to undertake some sort of hazardous stop, either in lane or on the hard shoulder or perhaps at the motorway / other road transition point.
- 3) Sudden rain squall – An ADS whose ODD allows for light rain only, may suddenly encounter an unexpectedly violent rain squall. The transition to this state may be over several seconds (less than 10) or in the low tens of seconds. Odd large drops (still within the ODD) and dark sky may be visible precursors to heavier rain a few seconds later.
- 4) Bird strike – a bird strike on a critical sensor may reduce the capability of the ADS perception system almost instantaneously. Multiple bird strikes (such as when approaching a flock of pigeons feeding on spilled grain on the road – it has happened to me!) could take out more than one sensor leading to a catastrophic perception loss.

Each of these has a different level of urgency, suggesting that different periods of time are **available** for the transition.

Now to consider the latter point – driver or user in charge “state”.

Many accidents occur when drivers are “fully engaged in the DDT” but still fall asleep at the wheel. This is known to be a particular problem in the small hours of the morning. If the UIC is not engaged, it is inevitable that in many cases he/she WILL be sleeping, whether it is “allowed” or not.

Current drowsiness detection systems may determine that a driver is asleep or has reduced cognition through drowsiness, but by then it is too late. They need to be **kept** awake, not warned when they are not. If they are not fully awake then re-engaging them may take quite some time, depending on factors which the ADS cannot measure or determine. These may include: how much sleep the person had the previous night, whether they have taken prescription medication, etc.

In summary, the “sufficient time” aspect is a meaningless concept here. The time **available** might be known under some circumstances, but certainly not all. The time **required** is almost certainly circumstance specific and almost certainly unknown. So “gives sufficient time to gain situational awareness” as a phrase in isolation might sound reasonable, but in reality it is never quantifiable so should not be used in this context. It is unhelpful at best and misleading at worst.

As a consequence, I think that the whole concept of the User in Charge is flawed since “sufficient time to hand over” is not quantifiable – therefore regulation depending upon that handover is fundamentally flawed.

(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.

Agree to the first part, but the “clear” and “timely transition demand” issues remain as above. I would say “unambiguous, timely transition request”, with the expectation that any

failure to respond in time would be fully handled by the ADS in a safe manner – and by this I mean a manner in which is at least as safe as that which a reasonable UIC would achieve.....(so definitely not stopping in lane, nor on the hard shoulder, etc.)

Do you agree?

Overall – no, not really!

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.

There is always a balance to be achieved when catering for some minorities, even when they are large minorities. Sophisticated vehicles which can cater for greater numbers of minorities tend to be more expensive than simpler ones that don't. This may make the technically superior vehicle financially inaccessible to the poor section of society.

A judgement must be made to establish a balance between the minimum level of "performance" for accessibility, including suitability for the deaf or hard of hearing, and economic accessibility.

It would be most useful to establish a rating system which could be used in cases like this. Such a rating system would consider the costs of different technologies, their ability to overcome an issue, together with the number of people likely to benefit from the feature. So for example, a hearing aid induction loop system costs around £100. This would be an unacceptable cost for the majority who do not need it. But providing a simple, universal interface to allow fitment of an aftermarket induction loop may only cost £1. So this might be reasonable to mandate for all private AVs which may have transition requests, allowing the hard of hearing to add the aftermarket device. Similarly, a full induction loop might be worth mandating for HARPS, depending on the numbers of users (passengers) who might benefit from it. Some sort of cost/benefit rating system would be most helpful, not only for features to help the deaf, but features to help other minorities (access ramps and tie-downs for wheelchairs, additional grab handles, large signage, Braille, etc.)

I would like to add in another thought on this subject – that of reduced eye focus change speed in old age. As we get older, the speed at which our eyes can change our depth of focus reduces. Even for someone who does not need glasses, a warning indicator at a distance of less than a metre will take an old driver longer to perceive than a youngster. Similarly, returning line of sight to the road, it will take longer for old eyes to refocus back at infinity – with the additional associated risk. Putting a visible warning for a transition demand in the car, when the driver's attention is urgently needed **outside** the car is obviously becomes more of an issue for the elderly than the young. This is often not well recognised as an issue.

Head and eye tracking technologies, together with the better head-up-display technologies (which make generated icons and images appear at infinity) could attract the driver's attention directly to the hazard "through" the windscreen, rather than add an unwanted additional step of focussing on a warning light on the dashboard and the UIC trying to work out what it means.

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?

That would depend on what basis the specialist regulator gave their advice, and what advice was given. Assuming that the basis was sound and the advice appropriate, then the answer would be yes.

Presumably the alternative is that the Secretary of State takes that decision based on compliance with one or more UNECE Regulations or GB equivalent. This would depend on the content of those regulations and sufficient confidence that compliance would automatically demonstrate suitability for use on UK roads. The current ALKS regulation is far too vague in too many areas to support this route as it stands. However, an improved version, or in combination with other, new, unknown regulations may be suitable.

The latter route should be the long term goal in my opinion, to allow reasonable international trade in AVs, but I accept that we are a very long way from identifying all of the content, let alone getting international agreement.

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.

Oh dear. This discussion in Chapter 5 takes an extremely narrow view of this question which is not helpful. Various alternative views are completely missing so I will try and fill in some of the gaps for balance.

- a) as safe as a competent and careful human driver;

Pick your desired reference driver

One option would be to pick a specific “competent and careful” driver and to exhaustively characterise what he or she would do in an array of scenarios, perhaps using a very high quality driving simulator. Risk decisions made by that person could be codified into a form of Digital Highway Code (DHC), and this becomes the reference standard of what is acceptable. Alternatively a group of such drivers could be used and a composite set of risk decisions determined and codified. This would help determine a range of acceptable decisions. Unfortunately, this approach is not mentioned, but has significant advantages in that key aspects are quantified and explicit.

The UK human driver population – what level of assistance?

The proposed concepts appear to rely on the population of UK drivers and their accident records at scale. This is fundamentally flawed because the definition of what is a human driver has not been adequately determined. Chapter 2 discusses at length what constitutes an AV but completely misses what constitutes a human driver. It appears to suggest that any

vehicle which is not “fully automated” or perhaps “self driving” (take your choice!) is human driven. This completely misses the effects of automated safety intervention systems!

Let us consider Automatic Emergency Braking (AEB). This is a technology which is highly beneficial in preventing human driven vehicles having some forms of accidents. The technology is improving and more and more potential accidents are being avoided. The logical implication is that as these technologies evolve, human driven vehicle accidents will become less and less common. So what level of “assisted” human drivers are we supposed to compare AV safety against? And how?

Such safety intervention automation has been evolving over many years and can be considered as starting with Anti-Lock Braking Systems (ABS) in the 80s. These override the drivers demand for brake line pressure and cut it momentarily and repeatedly to prevent wheel lock. This is a very specific and limited form of automation.

The next evolution was Emergency Brake Assist (EBA) which overcomes many drivers’ inability or natural unwillingness to provide sufficient brake pedal force in an emergency. The time taken to move the foot from the accelerator to the brake, together with the speed of application of the brake pedal allows the EBA system to identify an emergency brake event, and the system then intervenes and applies the maximum brake line pressure to give the best chance of achieving the shortest stopping distance. This is another very specific and limited form of automation.

Then we have AEB, starting with City AEB (low speed only), Inter-urban AEB (higher speeds) and an ever expanding range of target VRUs which can be identified.

The same arguments can be applied to lane keeping systems, emergency lane keeping systems, automated cross traffic braking; both front and rear, etc.

For these reasons, the arguments in 5.37 are fundamentally flawed. It isn’t a simple choice of “pure” human driving OR full automation. Human drivers will benefit from ever-increasing intervention automation systems which will help them out when they get things wrong. Should these ADAS systems eventually prevent all accidents or incidents, or at least the vast majority, there is no possibility for AVs to be any safer!

The comparison of AV safety performance with the human driving population is, to all practical intents and purposes, impossible, or at best, totally meaningless, therefore unhelpful.

(b) as safe as a human driver who does not cause a fault accident;

Again, we could choose an individual who has never had an accident, characterise them and codified that behaviour into a digital highway code. It is highly likely that a newly qualified driver has not had an accident, but this would not necessarily mean that their limited experience would be a good reference!

The “does not cause a fault accident” choice is probably the best one, however. The description in 5.43 is entirely reasonable and best represents my opinion of the choices given. The two limitations in 5.56 and 5.59 can, and will be overcome. The rules will never be perfect and will evolve, just as road traffic rules and laws evolve as we identify their weaknesses. Rejecting the “does not cause a fault accident” choice on the basis of these limitations really is an example of perfect being the enemy of good.

(c) overall, safer than the average human driver.

Again, the statistical, population based approach is fundamentally flawed because of the issues of defining human drivers and levels of intervention automation, as above.

There is another discussion which has a place here, but seems to have been missed – one of acceptability to the driving population. There are some key aspects:

- 1) Most drivers honestly believe that they are “good” drivers, and almost all drivers will consider themselves as better than average. Logically, this cannot be true, but the issue here is one of perception, not reality. Deep down, although we all think we are good drivers, we accept that on occasion, we make the odd mistake. Not many of course, but some.
- 2) Individuals have different driving styles and risk profiles. They drive like this because for them it is comfortable, safe, reasonable and normal.
- 3) A “good” AV is therefore one that drives like “me”
- 4) A “safe” AV is one that drives like me but doesn’t make the odd mistakes that I do.

Therefore to be acceptable as “safe” to an individual driver occupant, the AV should behave like they do, but with the odd mistakes and wrinkles ironed out. Inevitably, this must mean that the ADS must adapt its behaviour to suit the driver occupant in order for them to feel safe and comfortable. It would be very interesting to understand quantitatively how different an ADS can be in terms of its behaviour, yet still be close enough to a driver occupant’s behaviour to be accepted as “good” and “safe” by that person.

All brands of AV will not have the same “vanilla” behaviour. Brands are likely to have behaviour trends and each may well have some degree of adjustability. It could be envisioned that a Mercedes sports AV could have a “Lewis Hamilton” mode, for example. As long as it still obeys the fundamental rules of the road, why shouldn’t this be the case? There would probably be a market for such driving style modes.

Consultation Question 5 (Paragraph 5.120)

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

The challenge of making AVs “as safe as reasonably practicable” is often confused by the unnecessary merging of two ADS issues: “Perception” and “Decision”.

In essence the challenge is in these two rather simplistic related phases: “Perceive what is around you” and “decide whether or not to hit it”.

This is a gross oversimplification of course, but it is fundamentally true. Most of this consultation paper rightly concentrates on the issues surrounding the “decisions” part and what could and should happen in the edge and corner cases. These are interesting, can be contentious and are an excellent area for debate.

However, the reality is that all the high profile cases of vehicle automation failure to date are rooted in a perception system failure. The perception system either failed to identify a target correctly or did so too late. If the perception system correctly identifies all of the surrounding targets and risks, then for the vast majority of the time, the decisions required are absolutely trivial: Do not roll over the leaves, don’t roll over the pedestrian, etc.

Getting the perception system to perform reliably is absolutely critical and using existing type approval testing approaches to ensure a minimum safety performance level is entirely practical and reasonable. Perception system components (Lidar, radar, machine vision, etc.) could be tested and component approved to ensure that they meet a minimum safety level in the same way as a seatbelt, a tyre or a light bulb. The perception system should be subject to a system approval. This could then be lab tested against criteria determined by

the ODD, to ensure that it can correctly identify targets and clutter at various ranges in a suitable timescale to suit the ODD. Degradations in perception by poor weather or other environment hazards (oncoming headlamp glare, dirt encrustation, etc.) could be quantified and suitable speed reduction requirements could be determined to allow continued safe operations when deployed.

The separation of “perception” from “decision” is also key to AV accident investigation. If the perception system has timely and correctly identified all of its key targets before causing an accident, then it is the “decision” part of the ADS which is at fault and the issues mentioned in this consultation document come to the fore. If the perception system did NOT correctly determine all of its targets in time, then it was at fault and perhaps this was as a result of other preceding damage, faulty manufacturing or perhaps a fundamental error in the design or testing. This could be determined from target information stored in the event data recorder (EDR)

Fundamentally, accident investigation and lessons learned would be much easier if the “perception” and “decision” functions were separated and the perception outputs recorded in an EDR.

When faced with a large complex problem, the solution is almost always to split it down into more manageable chunks. To answer the question posed “how can automated vehicles be made as safe as reasonably practicable”, my answer is to break the problem down into “perception system” and “decision system”. We can then concentrate a significant amount of effort on getting the perception standards and validation processes right using conventional approval techniques, as this will have the greater benefit in ensuring AV safety in practice. The “decision” part of ADS performance should not be ignored of course, but I believe that the emphasis should be on perception system performance.

Consultation Question 6 (Paragraph 5.121)

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

The equity and protected characteristics paragraphs 5.69 to 5.78 seem to concentrate on these issues for people outside the AV, rather than occupants. This is reasonable since the occupant version of this issue was covered in Consultation 2 – HARPS. Also, it is unlikely that the equality issues will extend to other vehicle occupants, so I will concentrate on VRUs in my answer.

I refer back to my answer to Q5 and recommend splitting the approval of “perception” from “decision” functions of the ADS. It would be wholly abhorrent to suggest that an ADS would “consciously” decide to treat, say a black person differently to a white person, or to choose to run into a wheelchair rather than a standing person in a version of the “trolley problem”. This will not happen, but in the extremely unlikely case that it did, the evidence to prove it would be in the perception system target data within the EDR.

The issue of equality and protected characteristics relates entirely to perceiving the VRU and this can be addressed by simple, controlled testing at system level, as part of a perception system approval.

In the approval process, a “worst case” meeting is held between the manufacturer and the approval authority before testing. This examines evidence put forward by the manufacturer, most likely simulation data backed up with some physical test data. A reasonable “worst case” condition for a test is agreed and the test is then undertaken. This is the existing practice.

This practice could be extended to perception system testing. Let us consider the example case in 5.47. A perception system test programme would include a child running out from between parked cars. Let us assume that the perception system in question is predominantly vision based. The manufacturer provides a bulk of evidence data which includes the following observations:

- Perception is generally poorer at night, when there are no streetlamps and the system relies on the ego vehicle headlamps.
- The vision system has a limited dynamic range ofX
- At maximum aperture setting, the focal range of the camera is.... X to Y
- The vision system is most sensitive to rainfall at a rate of X with a droplet size distribution of X

From this data, it could be postulated that the worst case would consist of (say):

A black child, dressed in black running from between black vehicles at night (fully dark) in the rain condition above, whilst on a curve which presents the headlamps of an on-coming vehicle directly onto the camera sensor. It could be made worse by a strategically placed puddle which also allows direct reflections of the oncoming headlamps from the road surface too.

The test would take place in a lab capable of providing those test conditions. The sensing system would be installed at one end and the scenario would be presented to it. The results of the tests in terms of range (distance) and the timing (from the “reveal” to the correct, confident target identification) would help dictate the maximum safe speed that the vehicle could travel when under those weather, lighting and parked traffic conditions.

Crucially, the test should be re-run in simulation, exactly mimicking the physical test. The results should be very similar and if so, the test would help validate the simulation evidence. If the results are markedly different, it shows that the models do not sufficiently represent the system, the real world or possibly both. Even if both were “good”, this would still result in a “fail”, since the simulation data upon which the worst case was derived must be flawed.

Going back to the question of how AV regulators can fulfil their public sector equality duty. Clearly they can and should make available standards which cover the practical implications of the various equality based issues, so that these can be incorporated into simulation and physical test programmes. They can be highlighted as potential worst cases under specific conditions for approval testing. Again, these can be iterated and improved as gaps are recognised.

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;

I agree that safety assessment should use a variety of techniques. I am comfortable that component testing, system testing and track testing can be developed and be useful and helpful. I am also comfortable that eventually, simulation will become sufficiently reliable to be useful and helpful.

However, I am most concerned about public road testing. Whilst I think it can be beneficial, I am unclear as to how beneficial and at what cost. There are major drawbacks which in my mind have not yet been sufficiently addressed. A simple mileage accumulation is largely meaningless. Millions of miles could be driven up and down the same section of roadway under similar conditions. This would not be helpful.

There is no direct “back-to-back” comparison as road traffic is effectively random; therefore mandating such tests must be considered as fundamentally unfair for ADSEs who will be competing with each other commercially

If public road testing is used it should include a wide variety of road types, traffic types, weather conditions, etc. For an ADSE to have confidence that they have covered sufficient variety of these aspects, they would need to have a comprehensive list of such requirements which align with the ODD. The trivial requirements will be easy to cover. The more probing requirements (let’s call these “scenarios”) will be more difficult to achieve by a random drive and unnecessary mileage which does not encounter them is a waste of resources. It would be far better to assess these by simulation, with key aspects informed from real world data capture.

In short, in my opinion, public road “testing” has its place in data capture, to determine edge and corner cases, but once a comprehensive scenario set has been achieved, it is pointless as well as expensive.

Collecting scenario data from real world driving is absolutely valid – and I would consider this a key part of in-service monitoring of AVs, but also using many other data sources – such as EDRs on human driven vehicles, traffic cameras, dash cams, weather stations, etc. Real world testing has proved useful to date but its usefulness is diminishing and will eventually become unnecessary as scenario sets develop and simulation becomes more effective as a development tool.

(2) manufacturers/developers should submit a safety case to regulators showing why they believe that the automated driving system is safe;

I agree that the safety case approach is reasonable and helpful today, but it will become less so as time goes on.

Safety cases are useful for “one-offs”, for “clean-sheet-of-paper” concepts and where problems or solution technologies are radically new. We rarely build new nuclear power stations or new railways, so when we do, there is a good chance that the technologies available will be different, and safety perceptions or requirements may have changed.

Under such circumstances, safety cases are entirely appropriate.

However, when the market and solution are mature, standards and/or regulation will be more appropriate, reliable and efficient. Individual organisations need not “reinvent the

wheel” each time, with the potential risks of missing something or making an invalid assumption or conclusion.

The automotive industry puts new vehicle models into the UK market at a rate of several per week. Their operations, hazards, and risk control measures, are pretty much the same. The acceptable residual level of risk is the same across the vehicle parc, noting that some individuals may want their own, personal lower level.

Para 7.70 presents six numbered points raised by the Health Foundation. If points 1 to 4 contain anything which is substantially new or novel, then a safety case may be appropriate. If points 1 to 4 contain nothing substantially new, then the market, risks and solutions are mature. Also in this case, 5 and 6 will almost certainly be mature and so standards and/or regulation are more appropriate and should have been developed by that point.

Naturally some aspects will mature faster than others and these should be standardised earlier. Hence later safety cases are likely to draw more and more on standards, until a point is reached when safety cases are dropped in favour of a fully (or virtually fully) prescriptive regulation environment.

Determining the timescale for the transition will be challenging, but I would expect that after 4 or 5 safety cases for similar vehicle types, that should be enough to establish a standard, and after consultation and 10 or 12 similar safety cases from different sources, regulation is probably practical.

(3) regulators should:

(a) provide guidelines for what is in the safety case;

Agreed – but if it is known well enough to be very clear, it shouldn’t be a safety case; it should be a prescriptive regulation.

(b) audit the safety case;

Agreed, recognising that the auditors will soon recognise best practice and will be in a good place to write or contribute to standards and regulations.

(c) prepare guidance for manufacturers and developers on preferred standards;

Agreed, as above

and

(d) carry out at least some independent tests.

Agreed – certainly physical tests. I would also like to see the regulators have the ability and powers to run simulations using the ADSE’s model of the ADS whilst using standard models of the rest of the aspects required in scenario simulation runs. Alternatively, the regulator could work with a 3rd party to do the same.

Do you agree?

With the caveats above, yes.

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 is my opinion that the National Scenario Database (or an equivalent which may replace it), should consist of at least 2 sets of publicly available scenarios.

The first set is comprehensive and open to all to add scenarios. Road user groups can generate their own scenarios detailing the issues they envisage or detailing known incidents.

The second set is a sub-set which is used for regulatory approvals. The regulator regularly reviews the first set and assesses which scenarios, if any, are appropriate to port into the regulatory scenarios.

ADSEs would be expected to comply (successfully cope) with all scenarios in the legislative set as a minimum, but, should it be necessary in court, would have to justify why they chose not to comply with one or others in the "general set". This would put the onus on the ADSE to continually review the database and to decide what was relevant irrespective of the legislative set.

It may be appropriate to align scenarios with standard ODDs, so that it becomes relatively simple to assess relevant changes in the database.

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?

I agree with the sentiment and intent with the proviso that we need at least 2 levels of trials to be defined around which exemptions could be organised.

In essence, although not specifically stated, the current Code of Practice allows trials which retain a safety driver in the vehicle. This is a relatively low risk assuming the safety driver is able to retake control whenever necessary – in essence, a “procedurally intense” Level 3 which does not explicitly require a take-over request.

The CAVPASS team are currently considering what constitutes an “advanced trial”, which would require a higher level of confidence before allowing an exemption. Aspects which could determine whether a trial was “simple” or “advanced” could include:

- Removal of the safety driver from the vehicle; making it advanced.
- Level and type of remote supervision / monitoring; e.g. one full-time supervisor per trials vehicle, or one supervising 100 vehicles, full comms coverage with a chase vehicle and a known full coverage 5G as back-up versus relying solely on a patchy 4G coverage.
- Scale of the trial: is it just one vehicle or 1000?
- Relative risk of the trial based on the trial environment and the technologies being trialled: A trial on a lightly used rural road is probably less risky than on a busy main road in a city. Also a small, slow, lightweight, unladen cargo pod would likely do less damage than a 44 tonne truck. Similarly, a vehicle whose systems predominantly meet all type approval / C&U requirements would most likely be a lower risk than one which doesn’t comply with many or worse, none at all.

I would suggest that some thresholds need to be established for each to determine if “simple” or “advanced” or indeed if there should be more gradations. The complexity of each exemption scheme should reflect these thresholds and the associated risks.

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”);

I absolutely and enthusiastically agree!

- (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;

I agree.

(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?

My answer to this must include some observations of inaccuracies and confusions in the consultation document. These centre on definitions and misunderstandings.

A “manufacturer” is not necessarily the organisation that assembles the vehicle. As per type approval, the manufacturer is the organisation which offers the vehicle for approval, is responsible for compliance and for conformity of production. The manufacturer “owns” all of the system approvals and the overall vehicle approval.

Let us take an example. The Jaguar iPace electric vehicle was designed and developed by JLR. It was submitted for type approval by JLR and JLR is the “manufacturer”. The vehicles are built in Austria by Magna Steyr. Magna Steyr is not the “manufacturer”. Presumably, there are tight contractual requirements which port through build quality responsibility as JLR are responsible, legally, for the conformity of production, for recalls, etc.

Many of the systems which are type approved are interlinked. It isn’t possible to separate some of them. Hence the “manufacturer” takes full responsibility for ALL those systems and the overall vehicle approval.

There is no doubt whatsoever that an ADSE WILL BE THE “MANUFACTURER” if it offers the vehicle for approval with an ADS fitted.

A manufacturer may not necessarily be an ADSE, but an ADSE will always be a manufacturer.... unless.....see below:

Multi-Stage Approval

It is possible to have a multi-stage approval, where a base vehicle is approved by its “manufacturer”, but then another organisation makes changes which affect one or more of those approvals. Typical examples would be camper vans based on chassis cab platforms, minibus conversions, other special vehicles. These will all usually invalidate the Exterior Projections system approval, but may also invalidate others, such as seat belts and their anchorages, masses and dimensions, etc.

Where such changes and / or additions are made, then the organisation doing so will need to get new system approvals, ensuring that all affected approvals are adequately covered. In doing so, they become the “manufacturer” and become responsible for the overall vehicle approval.

So, an ADS developer which modifies a type approved vehicle by any changes to steering, braking, throttle control, etc. will become the “manufacturer”, if they then offer the vehicle for approval. They will also take on the in-service responsibilities associated with being an ADSE.

It is therefore quite clear that in most cases, the ADSE will become the “manufacturer”. Please refer to the e-mail conversation between Connor Champ and Matt Claydon of the VCA on this point.

The big question, which I haven’t seen asked, let alone answered yet, is what happens if the vehicle is modified and approved again, after an ADS is fitted and approved? i.e. The ADS fitment is not that last build and approval stage.

If another organisation sells a series of modified AVs, which “manufacturer” is the ADSE? Is it the one which applied the ADS and applied for the earlier vehicle approval, or is it the

subsequent organisation which has added, say, a new rear body to the vehicle or perhaps some extra seats?

So in a multi-stage approval, it is clear that the ADSE is always the “manufacturer” at the point where the ADS is approved, but we need to have a mechanism for identifying who has the ADSE operational responsibility should the vehicle be re-approved due to some non-ADS related changes and another organisation becomes the “manufacturer”. Presumably, the ADSE would remain being the body that had the ADS approved in the vehicle.

If, however there is a change to the ADS and it is re-approved, then the ADSE would become the “last” manufacturer that had the ADS approved.

So, would / could ADSE operational responsibilities be divorced from CoP and recall responsibilities in this circumstance? I think, potentially, yes.

This issue also has implications for OTA updates – What happens if there are OTA updates which are sufficient to invalidate the ADS approval and so require re-approval? This could lead to a change in both ADSE and even potentially the “manufacturer” if these are undertaken by a third body before registration.

So, to answer the question posed:

(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?

The strict answer is No, I do not agree. They may not be manufacturing the vehicle, BUT they are responsible for it meeting the all of the system approval requirements, as the developer is by definition the “manufacturer”, because it is applying for the approval.

Also, to note: An “OEM” can be either the “manufacturer”, or any supplier who supplies any component or system (hence “equipment”) into the build before registration. The use of the term is not particularly widely recognised nor defined and is not helpful in this document. Whilst the initialisation is mentioned in the BSI CAV Vocab document, the definition is not – because there is no clearly recognised definition. I would suggest you do not use the term.

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;

Agreed.

(2) an ADS should be defined as a combination of software, hardware and sensors, which can be installed in a “type” of vehicle;

I don’t completely agree with this. Referring back to my answer to Q5, I think that there is great merit in separating the “perception” and “decision” making systems. I think that the perception system should be subject to its own system approval – and that this could and should have explicit performance requirements which are physically tested and verified. The “decision” part (sometimes called the “stack” or the “AI”) should also have a system approval which would initially be based on a safety case process. This would call upon the specific perception system approval, amongst other evidence.

If the ADS is the term used to cover both aspects (perception and decision) then the answer is yes. My preference would be to call the “decision” part the ADS and to keep the perception system relatively separate.

- (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;

The ADS is a vehicle system, not a component, therefore each system approval is bound to a specific vehicle type (“type” as per type approval i.e. make, model, etc. NOT as in M1 or N1, etc.). The approval will call on a number of factors which will be specific to that vehicle type (size, weights, braking system, etc.)

and

- (b) how the ADS is installed within the vehicle;

It is completely normal and mandated to have details of how any system is installed in the vehicle is a system approval application (specifically in the “Communication document”)

- (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.

Likewise, it is completely normal and often mandated that an example vehicle is submitted to the regulator for approval of the installation as part of the system approval.

Do you agree?

With the caveat on sub-question 2, yes.

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.

I have no personal experience of the appeals process. Whilst the timescales provided seem reasonable for straightforward “mechanical” approvals, I suspect that the complexities of the ADS may require more time for “re-examination”, as this may include quantities of computer simulation runs and bespoke physical tests to assess aspects of concern.

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;

I am not entirely convinced. I appreciate that an international type approval may well not cover issues which are UK specific, but I don’t understand why a domestic approval would be similarly lacking. Surely a domestic approval would be specifically tailored to UK driving requirements? If not, why not? I agree that it should fall to a UK safety regulator should

undertake the categorisation however. I would go further. I suggest that the same process be used to formalise the ODD and hence how the vehicle may be used.

(2) the safety regulator should make a recommendation to the Secretary of State for how the vehicle should be classified;

Yes, and, as above, to formally agree its ODD.

(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;

I think the distinctions are nicely clear, but add again, that this could and should be the point at which the ODD is formally agreed and approved.

(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;

Yes

(b) the ADSE was closely involved in assessing safety and creating the safety case; and

This assumes that the safety case route is the only possibility. I think in the longer term, safety cases will be replaced by standards, which themselves become regulatory as the industry matures. I think there will be a transitional stage where safety cases are “cut and pasted” where there are chunks of accepted practice established, so would be more cautious of enforcing this para. If the ADSE takes responsibility, that should be enough, without mandating they re-invent the wheel.

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

Agreed.

Do you agree?

With the caveats above, in general, yes.

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;

Yes

(b) the procedure for doing so; and

Yes

(c) criteria for doing so.

I am not so sure that the legislative framework should be specifying the criteria. The criteria are likely to be complex and evolving, and should not be constrained by the need to change legislation. In my opinion, the legislation should identify who has responsibility to decide and maintain the criteria, not what the criteria are.

Do you agree?

In general, with the caveat for c), yes.

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?

I think that the categorisation decision will mainly be fairly trivial, if the criteria are clear, prescriptive and accessible to potential applicants. I think nuances around the ODD are more likely to cause appeals and so such an appeals process would be more usefully focused on that aspect, but also cover categorisation.

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.

I am not convinced regarding this approach, as it reflects similar weaknesses to relying on “miles driven” to determine safety, (see first part of answer to Q7). A small number of vehicles doing the same drive cycles tells you very little more than a single vehicle doing the same. In pharmaceutical trials, the human users are the variables. With a multiple CAV trial, it is the environment which provides the variables, not the vehicles – the vehicles should be the same as each other. I agree that limiting the deployment numbers might have some benefits in limiting exposure, hence overall risk, but those that are deployed should be deployed so as to exercise the system across the whole intended ODD. I would prefer to see varied operations over different roads and conditions which reflect that ODD before large scale roll-out. That may be fixed road environment (layout, features, fixed hazards) or temporal aspects (traffic, weather, etc.) A suitable coverage of the ODD should be a much better indicator than simply the number of vehicles.

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?

Yes

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;

I agree in principle, but have some reservations about the conventional vehicle “reference”. Of course, it isn’t a vehicle – it’s a human driver. Conventional vehicle driving decisions are made by the human drivers, not the vehicles. So which human drivers? Please see my answer to Q4 to see the options.

(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)

Yes – but also the power to mandate that ADSEs collect the data on behalf of the regulator. The regulator is not in a position to collect actual data themselves in practice.

and

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

Yes – and again, the power to mandate that ADSEs collect that data too.

(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;

Yes, without doubt, although what happens if the ADSE ceases trading? Does the vehicle become unusable if the software is not updated? Will we have the same issues as Microsoft Windows XT being unsupported but on a more dramatic scale? It would hardly be economic sense nor environmentally friendly to scrap a perfectly good vehicle simply because its software could not be updated? Worse, might this be a tactic to cease support on older AVs? An ADSE might deliberately become insolvent then re-emerge as a new entity without the costly update burden. This would also providing more pressure to buy a new vehicle rather than keep an old one? The dubious practices of smart phone suppliers and major software vendors provide indications on how ADSEs might behave regarding legacy updates. We should seek to learn from previous regulatory errors, and prevent this before it starts.

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

Yes, where it is required. Some AV entities are working more towards the ability to cope with basic maps rather than high fidelity 3D point clouds. The cost of lidar scanning routes is large, and whilst practical for short distances, is not likely to be sustainable on a large scale, at least in the medium term. The evolving expectations to cope with sudden changes suggests that in future, AVs will become less reliant on maps, and use them more as humans do – just for simple route finding.

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

Definitely agree. I go further and also suggest that all ADSEs be constrained to use standardised language regarding specific features and actions. This could be rolled into an update to the Identification of Controls regulations.

Do you agree?

Overall with the caveats above, yes.

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?

Yes. I see approval being a 2 stage process for international approvals. The first stage is simple and universal and can be undertaken by any international approval authority. The second stage is nation specific and relates to operation on the roads. This would be undertaken by the national authority for the territory and would relate to aspects which are important to that nation state but which are not covered in the international approval. Changes which affect the base, international approval would have to go back to that authority, whereas changes which affect local national approval would have to go to that authority.

A GB approval process could (and I would suggest should) roll these together if deployment was only in GB.

(2) Should the scheme should also deal with cybersecurity?

Yes – we don't want more bodies doing similar things – keep them together where possible. Different cyber weaknesses and attacks may well require different responses and mitigations. Splitting the schemes and roles only will provide confusion at the overlaps.

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

I suspect more powers would be required around incident and accident investigations, to require ADSEs to support investigations, to roll out learning to each other (through the regulator or otherwise).

I would also like to see powers to use specific cases and experience to generate new scenarios against which new approvals would need to comply, or against which ADSEs would need to provide OTA updates to cover. I mentioned this in my reply to Q7.

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?

I think it is perfectly possible to keep 2 separate bodies. One is looking at whether the vehicle meets the stated requirements on introduction (and that “introduction” could be when the vehicle is in service). The other is checking to see if the systems continue to do what the ADSE said they would. However, this might lead to some confusion which would need to be clarified.

In my opinion, an OTA update which altered some aspect of performance under type approval, must be checked and (where appropriate) approved by the type approval authority. The safety assurance body must understand what that change consists of. The safety assurance body may have other reasons for mandating OTA updates and these might be combined. However, I think it possible to keep the VCA and the DVLA roles separate.

Do we want to keep them separate is a different question and it may be that the benefits of combining outweigh the drawbacks. I am not in a position to give a suitable answer to that, but I suspect that it is not desirable. Change causes instability and inefficiency. We should try and keep what we have if possible.

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)?

I think that an advisory committee would be a good start, with representation from a wide group of stakeholder organisations, together with publicly available minutes of meetings. There should also be a clear mechanism for outsiders to raise issues which the committee can discuss, so that individuals or unrepresented organisations can have a say and see the responses or outcomes.

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?

Yes to all.

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?

Yes to all.

I would however offer one note of caution. The section discussing warnings refers to the FCA and the finance industry. This is hardly an example which inspires confidence in the approach. The finance industry is largely despised by the general public and the FCA is widely considered as a toothless bystander, given the lack of visible redress after the financial crisis. This example would appear to suggest (whether correctly or incorrectly) that warnings carry no weight whatsoever. It may be that some FCA warnings are effective in practice, but that is not the public perception so probably not helpful to use as an example of good practice.

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?

Yes

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?

Yes but I would go further and suggest that the specialist incident investigation unit should investigate ALL AV collisions, not just serious, complex or high profile.

Minor collisions may occur which could help identify problems which might cause more severe incidents later. In my experience, police accident reports can be very variable in quality and cannot be relied upon to determine cause of accidents in even simple cases. I would not be happy entrusting AV accident investigations to the vagaries of the many police investigation teams. We do not expect a rapid rollout of AVs, so hopefully there will not be too many cases to investigate in the early days, nor later if learning from investigations can be shared.

Also, who would decide what was serious, complex or high profile? My concern is that this would also fall to the police, who would naturally want to keep control themselves. So in practice, the AV investigation unit may get hardly any cases, and therefore hardly any central experience. Once again, the opportunity to learn would be dispersed around the police forces and lost.

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?

Yes – I do agree that this would be helpful.

Consultation Question 27 (Paragraph 11.83)

18.27 We welcome views on:

(1) the issues the forum should consider;

Before directly answering the question, I would like to pose another and try to answer that:

What is the intended output of the forum?

My answer is that it should be an ever evolving set of ever more detailed rules which ease the safe deployment of AVs. In essence, the output will quantify the requirements in a useful way – it will become the “Digital Highway Code”(DHC)!

So, to answer your question, I think a suitable starting point for issues to consider would be to work through the existing Highway Code but with an AV perspective, but also be ready to include new issues raised which do not appear in the existing Highway Code.

(2) the composition of the forum;

ADS developers, ADSEs, road safety interest groups – but it should be chaired by DVSA.

and

(3) its processes for public engagement.

I think that regular updates of the “DHC” should be made publically accessible (and be used for GB approvals) and the public may be allowed to propose issues for discussion with

feedback given to those who do. These proposals could perhaps be made using an on-line “form”, etc.

CHAPTER 12: THE USER-IN-CHARGE

Consultation Question 28 (Paragraph 12.24)

18.28 We provisionally propose 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;

I have to disagree on two levels:

As per my answer to Q1 c), I am very uncomfortable regarding the requirement for a UIC to be mandated to take over the driving task within an ill-defined, impractical or unrealistic timescale, which undermines the overall UIC concept.

I also have reservations regarding the phrase “direct sight”, which I raised in response to the second consultation document. It is important that a responsible individual (UIC, if you must) be able to see not only the vehicle, but also what is in the direction of travel of the vehicle for a sufficient distance to bring the vehicle to rest. The drone analogy is misleading. A drone operator will almost always be underneath the drone and the drone is small – so it is relatively easy to be sure that the operator can see any obstructions that the drone may impact. An AV is likely to be bigger and the operator in the same plane – that is to say “ground level”. If the operator is behind the vehicle, he probably has little view of what the vehicle might hit, even if is in direct sight / direct line of sight of the vehicle. This should not be acceptable. It is the visibility of the potential impact objects which are the priority, not the visibility of the AV.

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

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;

No, I disagree. Please see my answer to Q1 c)

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?

No, I disagree. If the ADS causes the vehicle to stop in the manner which constitutes a criminal offence, then the ADS and the ADSE are at fault. In my opinion, the ADS should put the vehicle into a suitable, safe and legal minimum risk condition. I fundamentally disagree with the route ALKS is taking on this issue and I expect a backlash of complaints when such

ALKS block motorways by stopping in lane and causing accidents. I would much prefer us to avoid this entirely predictable hazard.

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.

This is an interesting point. The learner is not “in charge”, the instructor is. Therefore logically the learner could not be the “user in charge”. I agree with the concept of learners being trained on such systems if such systems are deemed legal and appropriate (which I don’t – see above).

There are benefits to mandating dual controls for learners, but also drawbacks. Learner drivers may use non-dual control vehicles with an unqualified but competent instructor (often a parent) to gain practice before a test. The cost implications of mandating a dual control vehicle and in practice, a qualified instructor, may reduce that important practice phase and lead to poorer results in driving tests.

Perhaps it would be appropriate to licence drivers who have been trained on ALKS and L3 AVs differently – and those that are licenced to be UICs must have had such dual control ALKS / L3 training.

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?

Yes – please see the final para in my response to Q30 above.

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?

Yes, if the issues surrounding UICs as per my answer to Q1 c) are resolved.

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.

Yes – this seems reasonable.

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.

Probably yes. I am concerned by the phrase “a competent and careful driver could not have avoided the offence”. Who is this driver? Who determines what they could or could not have done? Shouldn’t the word “would” replace “could”? This changes the emphasis from what might have physically possible to what would have been probable.

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?

As an interim solution (some might consider it a “short term bodge”) probably yes, with the exception of the mechanics of parking, which whether on public roads or private, is a dynamic driving task.

In the long term, I would expect all of the other aspects to become automated, with the vehicle itself refusing to move if 1), 2) and 5) were not satisfied.

Where the parking referred to relates to leaving the vehicle in a prohibited place, I would suggest this would depend upon what mode the vehicle was in and whether the UIC was responsible. If completely autonomous valet parking was used, then the UIC should not be liable. If the parking is manual, or the parking was automated but the UIC present, then the UIC should be liable.

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?

I do agree, but caution warning for issues which are not immediately obvious and which would only be expected to be discovered at a service or MOT. If a vehicle fails an MOT, the human driver is not normally prosecuted for driving an unsafe vehicle up to that point, but they are expected to rectify the defect. A UIC should have similar leeway.

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”;

Agreed

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?

In general yes, but not completely under all circumstances. Consider the case where the AV encounters a flooded road and stops, unsure whether it can progress through the flood. It requests assistance from a remote operator. The remote operator can see that the water level is quite low and advises the AV to proceed, which it does at low speed. The AV then hits a large submerged pot-hole which was invisible to it and very difficult for the remote operator to perceive. It rolls over in the flood. Who is at fault? The AV was self-driving but doing so on advice from a human. I would suggest that some or most of the responsibility would lie with the remote operator and little, if any with the ADSE.

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”.

I think that because of cases like the above, some amendment may be required.

Consultation Question 38 (Paragraph 13.86)

18.39 We provisionally propose that:

(1) 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);

Yes, I agree in principle.

(2) 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;

Yes

or

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

I agree that the licensed “operator” must be licenced for supervision and remote operation where required, but I do not agree the maintenance services must be bundled in too. Automating software updates is not difficult – phones and laptops do it regularly. An AV could, and perhaps should “check-in” to see if it is running the latest software on a periodic basis – perhaps every trip.

Other maintenance need not be regulated. Safety critical systems are maintained on many UK vehicles outside the VMs dealer network and almost no garages are licenced anyway. Why need this change? I don’t believe it does need to change. Licenced monitoring, supervision and remote driving – Yes. Licenced maintenance – No.

(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?

With the exception of the maintenance requirement as I mention above, yes.

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.

This is slightly contentious. I agree that professional competence should be demonstrated. The safety management system should also meet a certain standard. I am less sure about tying this to “a” safety case. Firstly, I see safety cases as a short term interim step whilst industry norms are established. After that, safety cases can be phased out and replaced with standards. Also, there could be multiple safety cases for different types of AV in a fleet and the deployment characteristics of different services may mean different operational safety cases. I propose that a NUIC operator should be licenced on a “module” basis, where competence in certain modules allows activities covered by those modules. One or more individuals would need to show professional competence for each module that the operator is licenced for. This would also overcome the drawbacks of a single transport manager. The modular skill and experience array could be distributed across staff working in those areas.

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;

Yes, and possibly the occupants if the vehicle is a HARPS or if the owner / registered keeper wanted this.

- (2) maintain the vehicle;

No – As mentioned above, there are no significantly new complications that require a more onerous maintenance regime than a conventional vehicle. Software could and should be updated automatically.

- (3) insure the vehicle;

No – this could be transferred to the registered keeper or owner as appropriate. An AV could and I propose should be designed such that it will not move if there is not a valid insurance policy in place. This need not rely on a communications link at every start-up. The vehicle could use knowledge of its insurance policy expiry date to determine if the policy had expired, with a further check when comms was available, to check if there had been any changes.

(4) install safety-critical updates and maintain cybersecurity;

No - If the ADSE is responsible for creating such updates and patches, the ADSE should be responsible for an automated process to keep the vehicle up-to date.

And (5) report accidents and untoward events (as defined by the regulator).

No - Again, I think the ADSE should be responsible for ensuring that this happens automatically. If we are not expecting an operator to watch the vehicle on-screen all the time, the detection system for identifying such events will need to be automated. If that is in place, why introduce a manual step which could, and most likely will, on occasion, fail?

Do you agree?

On the whole, no. Only supervision, monitoring and when required, remote operation.

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?

Yes, I do agree, as above, when those duties can be automated or at least checked automatically, or where there is no obvious reason to change from current practice. As above, these could be: Maintenance (existing arrangements should suffice), Insurance (an interlock could prevent uninsured operation), Installation of updates and patches (should / must be automated) and accident / incident reporting (should / must be automated).

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.

As I mentioned in my response to the second consultation document, I would like to reiterate the point regarding potential conflicts of accessibility. It is possible, and some say likely, that AVs will be used for “on-demand” services. In such circumstances, not all vehicles available need to have all accessibility features, but enough need to sufficient combinations to suit the service intended – and I propose that there are guidelines or standards as to proportions of vehicles with certain accessibility features depending on the fleet size and the deployment demographics.

E.g. If you have a large family and require an 8 seat taxi rather than a 5 seater, you would naturally expect to have to book a vehicle with that specification. I think it fair that some accessibility requirements which impose substantial design changes (such as wheelchair accessibility) can and should be treated in the same way. Some accessibility features (such as clarity of signage, and perhaps braille controls, could be mandated on all vehicles without unnecessary compromises.

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?

Yes to all, as long as the term “disabled” is interpreted in the widest sense.

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

I would suggest that the panel is consulted annually, but the panel can accept input from members or members of the public at anytime.

Consultation Question 43 (Paragraph 13.133)

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

I am not in a position to provide a helpful answer to this question.

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?

I agree with all of the above, with two caveats.

- a) I think the first statement should be "...it should be a criminal offence for an ADSE to knowingly omit safety-relevant information..... " since including an unintended omission would not be in the spirit of what is being proposed.
- b) That the extent should not be limited to "classification as self-driving" but also include the formal claims and description of the ODD.

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.

I agree with the above with the caveat that, as in my answer to Q44, there should be explicit reference to the ODD. When putting forward the vehicle for classification as self-driving, I suggest that this must also include a full description of the ODD. I think that in practice, some ADSEs may be optimistic about the scope of ODD that their vehicle can cope with. I think that it should be an offence to claim a greater scope of ODD than that which can be satisfactorily demonstrated. I don't think that comes through sufficiently strongly in Offences A and B as written.

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.

I definitely agree. This could, perhaps be a subject for a standard or guideline.

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 ?

Yes, I agree, but please see answer to Q48 below.

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.

I am not sure if we need this. If para 18.49 were extended to include data (but reworded to encompass false or damaging data within the vehicle), we may not need it as such. With respect to infrastructure in this context, we are only interested in that interference which causes an issue for the vehicle, and it will, I think, always be the case that it will involve false or potentially damaging data or information being put onto or into the vehicle. If that is the offence, the mechanism (whether via infrastructure, a plug-in dongle or by some other electromagnetic corruption) becomes immaterial. I suggest that the offence be reworded to encompass “install, provide false or corrupt data in the vehicle with intent to do harm”.

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?

Although I do not consider myself an expert in this area, both seem reasonable.

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?

I do not feel sufficiently expert to make an informed choice or comment regarding this question.

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.

I do not feel sufficiently expert to make an informed choice or comment regarding this question.

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?

I do not feel sufficiently expert to make an informed choice or comment regarding this question.

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?

Yes, this would appear reasonable.

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?

Yes, both statements appear reasonable.

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?

Yes to all points.

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?

Yes.

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?

No. I think it can and should be longer. I would suggest 10 years with possibility of extending further in the light of greater experience. The costs of storing data are ever decreasing. The limited DSSAD data is not that expensive to store today and will become less-so in future. We should assume that AVs will be serviced at regular intervals so there is ample opportunity to down-load data from the vehicle (as a worst case, on an annual basis) so vehicle data storage capacity should not be an issue. Cloud storage of text data is not so expensive. It is a poor excuse to use cost as the reason for not storing data for longer than 3 years.

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;

Yes, but I would go further.

In my opinion, the ADSE should have the responsibility to collect data on behalf of the regulator (see my answer to Q7 regarding data collection during in-service monitoring to allow scenario generation, also my answers to Q18 sub questions 2a and 2b)

I also disagree with some of the fundamental presumptions implied in this chapter. There is an implication that incidents (impacts) can only be detected by accelerometer data on the

vehicle. This is both over-simplistic and untrue. An airbag deployment is a crude means to determine an impact. Airbag deployments have inherent safety risks so are only used in impacts which justify the risks of injury caused by deployment against the risk of injury if they are not deployed. The risks of pyrotechnic seatbelt pretensioner deployments are much lower, so there are some impacts which may cause their deployment but not airbag deployment. Reversible electric seatbelt pretensioners can be deployed before an impact if the control module considers there is a significant risk that an impact might occur shortly. Triggers for this may be the activation of the ABS system, activation of the ESP system, activation of the AEB system, a sudden increase in the rate of roll of the vehicle or the exceedance of an absolute roll angle. Any of these existing systems, and possibly others, could and I propose should, be used to trigger the storage of a comprehensive, short duration “event” data set, either for accident or near-miss investigation.

Similarly, it should be possible to use the AV’s own perception system to determine an impact or near miss far more accurately than from simple accelerometer data. Even basic analysis of mass market surround camera view systems could be used to determine if objects approach the AV very closely and this could be a trigger.

There is evidence to suggest that audio analysis of incidents can also be used to identify impacts with high accuracy. This has been used in Japan on traffic lights, to give rapid automated warnings of incidents to traffic controllers. This was before computer vision systems could be programmed to detect similar incidences.

The technologies available and already deployed to identify impacts or near misses are many and varied. They will develop and improve further if they are mandated.

In addition to details on how data will be recorded, stored, accessed and protected, I propose that the ADSE should provide a comprehensive explanation of how it would comply with a new, explicit requirement to detect accidents and near misses of its vehicles. It should also provide details regarding what data it will store on-vehicle and why that is sufficient to enable accident investigation by the new accident investigation body. As a minimum, I propose that this include all target data from the perception system and 360 degree surround video throughout the pre and post incident period.

It has also been suggested that the ADS provide an electronic / digital version of the “commentary drive” required by in advanced driving tests and by police drivers. This would explain what the system considered to be key aspects and why it made the decisions that it did. This data stream could also be mandated within the data recorder.

I would expect the vehicle to attempt to down-load incident data as soon as practicable after the incident, and this data should be analysed by the ADSE and reported to the regulator as and when required.

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

Do you agree?

I agree but would also require that the ADSE and the vehicle have the systems to detect all impacts and near misses and to report them to the regulator.