

Summary: Testimony of Dr. Philip Koopman

- The mobility, economic, and safety benefits promised by manufacturers **will only materialize if self-driving vehicles can be made safe and reliable at scale**. The industry is not there yet, and it will easily be 5 years (or perhaps 25 years) before the technology matures. Issuing overly-permissive regulations prematurely might serve the short-term desires of industry, but not the public interest.
- **Current regulations do not place any meaningful limits** on the ability of companies to develop the technology for fully autonomous vehicles. Moreover, there are significant incentives for companies to cut corners on safety, to the detriment of other road users. The chaos and harm unfolding on our public roads make it clear that manufacturers have pretty much free reign. State preemption of local regulation has been making things worse instead of better. Federal preemption of State abilities to regulate driving behaviors during a multi-year wait for federal regulations will be even worse.
- The race that counts for automated vehicles is not deployment. Rather, what matters is the **race to create reliable and safe vehicle automation technology**. To win, the Federal Government should require companies to **adopt their own industry's safety standards and best practices**, just as proposed by US DOT in late 2020. Additionally, manufacturers should be required to be **transparent about crash data** and other incidents, as well as be **held accountability for safety issues** caused by their computer drivers. Other life-critical uses of computer systems such as aviation have found that conforming to their very own industry's standardized best practices is the only viable way to get safety.
- **The economic impact that matters for the next 5-10 years, or longer, is partially automated systems** (so-called Level 2/2+ features) that require driver oversight. This technology is a convenience feature and does not contribute to safety. (Documented safety improvements come from other technology features, not the self-driving features.) Vehicles with poor driver monitoring and poorly conceived driver interfaces are prone to causing injury and death to road users, putting some US-made vehicles at a disadvantage in overseas markets. These partially automated features should be included in automated vehicle regulations.

Written Testimony of Dr. Philip Koopman

IDC Subcommittee Legislative Hearing: July 26, 2023

Thank you for inviting me to testify.

My name is Philip Koopman. In my early career I served as a US Navy submarine officer in the Cold War, and spent several years working in industry, including automotive supplier work. Then I became a professor at Carnegie Mellon University, where I have been working on self-driving car safety for more than 25 years. I am the primary author of the UL 4600 standard on autonomous vehicle safety, and have written two books about self-driving car safety as well.

I have worked with automotive industry partners such as General Motors and helped the US Military improve the safety of their automated ground vehicles. I also have extensive experience with safety-critical computer applications in other industries including rail transportation, chemical processing, industrial controls, electrical power systems, aviation, and agriculture.

My first experience with self-driving car safety was when I worked on the Carnegie Mellon NavLab project, which had just driven an early self-driving car [2800 miles to San Diego 98% hands off the steering wheel – back in 1995](#). The industry has since spent perhaps \$100 Billion on that much more difficult last 2%. There has indeed been impressive progress, especially on city streets. But even after 25 years, we still have a long way to go on reliability and safety.

The title of this hearing is “Self-Driving Vehicle Legislative Framework: Enhancing Safety, Improving Lives and Mobility, and Beating China.” I’ll address each of these topics.

Enhancing Safety

Beyond the need for effective regulation, a central lesson learned in safety in all the industries I have worked with is that safety only happens if there is transparency and accountability. Supporting both federal and state government ability to collect and publish safety data for testing and deployment is crucial. Independent oversight of safety via a combination of proactive regulatory oversight, independent

assessment, and access to public courts is essential to ensure long-term safety of any complex technical system, including automated vehicles.

There is no independently vetted data showing automated driving features improve safety. To be sure, driver support features such as [Automated Emergency Braking \(AEB\) do improve safety](#), even while there is room for further improvement for pedestrian protection at night. But the story changes once steering is automated, either with or without a human driver. There is no evidence yet that automated driving will improve the tragic US road fatality rate. One certainly hopes fatality rates will improve with this technology. But it is an aspirational goal, not a certainty.

You will hear industry proponents claim that [human drivers are terrible, so of course computers will be essentially perfect](#). Computers might never drink and drive, or text and drive. But computers are notoriously brittle when encountering something they have not been trained on. We have seen [driverless robotaxis run over charged fire hoses, drag downed power lines down the street, and stop in a way that blocks emergency vehicles](#). It used to be said that having so many different sensors meant they would never do something stupid like run into a stopped bus. [But that happened this year in San Francisco](#). The reality is that safe software is very difficult to create, and we should expect computer drivers¹ to make mistakes. The question is how many they will make, and how severe they will be.

You will also hear industry proponents incorrectly claim that “94% of crashes are caused by human error” or some similar wording. That claim has been refuted by the [Chair of the National Transportation Safety Board](#). While humans are certainly imperfect, [average drivers have about one fatality per 100 million miles](#).² That includes the drunk, distracted, and non-seat belt-wearing drivers, in vehicles too old to have robust active safety features, driving on roads with safety issues. If we want to invest in safety, we should

¹ A *Computer Driver* is a computer that exerts sustained control of vehicle steering, and likely other aspects such as speed control. A computer driver might depend on the availability of a human driver (SAE Levels 2 & 3), or drive without the need for a human driver (SAE Levels 4 & 5).

² In the past few years that has trended worse. It is unclear if this is a temporary pandemic-related effect or a permanent shift. Regardless, we know humans are capable of better driving outcomes without automation.

address reducing impaired drivers, reducing distracted drivers, improving seat belt usage rates, improving road infrastructure safety, improving vulnerable road user safety measures, and [making public transit more attractive to use since that is already dramatically safer than private vehicles](#).

Deploying computer drivers might bring economic and societal benefits, if done with care. But automated vehicle technology is not a silver bullet for safety. Rather, we should be striving to make sure that as autonomous driving technology spends another decade or so going through its inevitable growing pains, it first does no harm to public safety.

It is important to remember that there is tremendous financial and competitive pressure to deploy the technology as fast as possible. Given a near vacuum of regulatory requirements on the safety of the computer driver itself (other than recalls, often after significant harm has been done), companies have a huge incentive to cut corners on safety and bet on getting lucky to meet their milestones and for that next round of funding. We are seeing the largely avoidable results of that pressure in the news in terms of [confused robotaxis impeding emergency responders](#), crashes, injuries, and even [continual fatalities in the case of automated steering systems](#) (so-called Level 2 automation).

We can hope that one day when computer driver technology matures it will save lives. But the reality is that in the scramble to keep their companies viable, executives might prioritize deployment milestones and getting the next feature working rather than safety. Regulations provide the guard rails and a level playing field so that American companies as well as foreign competitors can't cut corners on safety if they think that the risk/reward outcome in doing so is appealing. There is nothing inherent to this technology that guarantees that safety will be improved. If we do not explicitly require safety, we are likely to get the least amount of safety the industry thinks it can deliver while prioritizing showing progress to investors.

Industry tactics such as forcing private arbitration on terms favorable to them will exacerbate problems by reducing the expected cost of each injury or death to the manufacturer, and inhibiting public discussion of defects that might be revealed (if a technical analysis is even permitted by the arbitrator).

Federal regulators have struggled with increasingly frequent software recalls for years. A transition to computer drivers and increasing software churn caused by over-the-air software updates will completely

overwhelm their capabilities to effectively manage safety via recalls. A more proactive strategy is required, as well as more regulatory capability with sophisticated software and autonomy technology skills.

We should not let a manufactured urgency that this technology will miraculously solve the very real road fatality problem pressure us into passing laws that are unreasonably favorable to the industry. The industry should prove that it will deliver on safety before they receive favorable regulatory treatment. We should not simply assume safety will happen based on hype and promises.

Improving Lives and Mobility

Automated vehicles include those that require drivers and those that do not. The reality is that for the near- to mid-term future [the technology that will matter to the US car industry is not robotaxis, but rather so-called Level 2+ vehicles](#), that do almost all the driving, but require human oversight. These are already on our roads in large numbers, and those numbers will increase dramatically in the next decade. Unfortunately for road safety, they are [prone to abuse](#) due to [automation complacency](#). While NHTSA has, to its significant credit, included such vehicles in data reporting requirements, this technology is essentially unregulated in the US except for potential recalls in cases of egregious problems.

Cars that mostly drive while providing insufficient driver monitoring and attention management are [already causing injury and fatality crashes on our roads](#). So-called autopilot and related features are only convenience features. Automated steering control is not a safety feature, despite the marketing messaging you might have heard. Significant harm is being done in part due to [industry messaging that the cars pretty much drive themselves](#), when in fact privately owned vehicles currently on US roads do not.

Some companies are positioning their cars to give time back to drivers. One company is even telling customers it will be OK to play video games on the dashboard in their soon-to-be-released vehicle. But that company [will not commit to being responsible if there is a crash while the driver is playing those games](#) – even though that driver might be subject to civil and criminal consequences for a crash. Such vehicles are not really the emphasis for this hearing, but are in fact the most important type of vehicle for at least the next decade for both US competitiveness and public road safety. More attention should be paid to the safety of partially automated driving systems that require human drivers to be involved.

With regard to truly driverless vehicles, they have the potential – and *I emphasize the word **potential*** – to improve lives and expand mobility options, but only if designed and deployed carefully. That potential can only be realized after they are safe and reliable. Right now the technology is unreliable, as news stories from San Francisco and other places tell us on a weekly basis. Nobody really knows yet how safe they will turn out to be, due to fundamental limitations of the technology. Most importantly, the “AI” used in these vehicles is good at things it has been taught, and bad at surprises. But the real world is so full of surprises, [we don’t know if we can teach the computers enough to come out acceptably safe in the end.](#)

While I truly hope the technology will achieve its potential, it is too soon to take it as a given that the potential will be reached any time soon. It might be 5 years. It might be another 25 years. It might well take another technological breakthrough that will happen who knows when. We should set policy on demonstrated capabilities rather than hopes and dreams.

Beating China

The race is not to put robotaxis on roads. That race is already over – prototype robotaxis, prototype trucks, and production partially automated vehicles are operating on the roads in the US, China, and Europe. The race that matters is to put safe, reliable cars on roads at scale. Companies must [earn the trust of the public](#) to succeed at doing that. The US and Europe have an advantage here in that we are creating the core technology, and we are writing key industry safety standards.

Exempting our own companies from safety standards might provide the illusion of progress for their technology. However, foreign competitors are taking safety much more seriously than many of the US companies due to stronger regulatory regimes at home. Those foreign companies are likely to get ahead on safety engineering to serve their home markets – while using the US as a largely unregulated testbed (with concomitant risk to constituents).

Legislative framework

Right now there are no regulations that impede the development of safe, reliable computer drivers. Companies can shop for a State that has favorable laws and operational conditions as they like. Issuing Federal Motor Vehicle Safety Standard (FMVSS) waivers for prototype vehicles for things such as

dashboard indicators due to an unnecessary connection to a human driver in the current procedures is a reasonable plan so long as [exemptions require a comparable level of safety](#), and do not create a loophole for large scale deployment of inadequately designed vehicles. Crucially, FMVSS does not cover computer driver safety.

[US DOT has already proposed an automated vehicle framework](#) (85 FR 78058, Docket No. NHTSA-2020-0106) based on requiring the industry to follow their own safety engineering practices and standards for computer driver safety. The US Federal Government should require that any automated vehicle operated on US roads will be built to an acceptable level of safety based on self-certified conformance to standards listed in that document, plus a newly issued automotive industry cybersecurity standard and a road testing safety standard.³ NHTSA should also be given a mandate to continue and expand its [current data reporting requirements](#). That will ensure that other countries are unable to exploit the fact that we have weaker automated driving safety regulations than they do to deploy unsafe vehicles and use our people as testing guinea pigs to develop their technology.

It will take years for NHTSA to put regulations into place. In the meantime, the states need to be empowered to place some guardrails on safety and ensure the accountability of manufacturers for the safety outcomes of the vehicles and computer drivers those manufacturers are putting onto public roads. A useful tool in this process would be for the federal government to [establish a duty of care for a computer driver](#) that is equivalent to that of a human driver under the same conditions, and let states apply their tort law and product liability law from that point on.

Key takeaway points

In closing, the key points I would ask you keep in mind are:

- Companies have shown they will be as opaque as possible about reliability and safety, as well as evade accountability whenever possible. They are losing the public's trust.

³ Specifically: ISO 26262 (functional safety), ISO 21448 (safety of the intended function), ANSI/UL 4600 (autonomous vehicle system safety), ISO/SAE 21434 (automotive cybersecurity), and SAE J3018 (road testing safety). Requiring a state-of-the-art Safety Management System is also important (e.g., AVSC 7202107).

- If the government does not find a way to incentivize companies to achieve reliability and safety, foreign companies that produce reliable, safe vehicles will likely win in the long term.
- The core issue is public trust. Fake it until you make it has run its course here. If we want to still have an automated vehicle industry in the future, Congress needs to act to require transparency, accountability, and adoption of the industry’s own safety standards.

Additional Key Points

Safety technology. We can improve safety more easily without resorting to full autonomy.

- Robust deployment of active safety technologies: automatic emergency braking, vulnerable road user crash avoidance, lane departure prevention, blind spot detection, rear crash prevention, etc.
- Stronger safety requirements for Level 2/2+ vehicles (driver monitoring, curtailing misuse)
- [Clear transfer of duty of care](#) between human and computer drivers for all SAE Levels⁴
- Recognize that there are more cost effective ways to improve safety than autonomous vehicles
- Create a consensus framework among stakeholders for [how to measure “safe enough” automation](#)

Jobs and economic impact. As with all new technologies, there is opportunity for better, upskilled jobs.

- Driver assistance/partial automation will have significant market share for at least a decade
- Any transition to full autonomy will be ramped over a decade or longer
- Highly skilled safety drivers for testing will be needed indefinitely
- Highly skilled maintenance and operations technicians to maintain life-critical computer drivers
- People needed for local driving, cargo security, hazardous cargo supervision, in-building delivery
- It is more important to get things right than get a problematic federal law immediately

⁴ For example, defining during which operating modes the manufacturer is the responsible party in production Mercedes Benz vehicles that conform with [UN ECE #157 \(ALKS\)](#)

Liability. Create an explicit duty of care for a computer driver.

- Federal law should expressly state that AV manufacturer owes a duty of care to other road users
- A computer driver should have the same duty of care as a human driver would in any situation
- Manufacturer is responsible party for any breach of computer driver duty of care
- Avoid the [moral crumple zone](#): blaming human drivers for crashes initiated by technology faults
- Product liability should be a last resort – not the only option available for crash after running red light

Data Reporting. Make permanent and enhance current NHTSA data reporting requirements.

- Fatalities, injuries, significant property damage, regardless of fault contribution
- All moving violations of traffic rules, including explaining reasonable exceptional incidents
- Substantively blocking a travel lane for more than 60 seconds after a safety stop
- Hard braking events (with analysis as to whether it was phantom braking)
- Duty to proactively collect data from such events if computer driver active within 30 last seconds
- Number of hours and miles computer driver was active, to support computing failure rates
- Much more stringent limits on data redaction in narrative reports; presumption of transparency

Regulating safety. Require companies to follow their own industry's safety best practices & standards.

- Current approaches (FMVSS, VSSAs) are wholly inadequate to regulate vehicle software safety
- Industry standards are already issued to address autonomous vehicle safety
- Increase NHTSA expertise at computers and software much more quickly than current plans
- Prohibit uncrewed testing until vehicle is acceptably safe for deployment
- FMVSS exemptions should ensure equivalent safety and not apply to vehicles sold to retail customers
- Build on NHTSA/DOT existing framework from December 2020: 85 FR 78058 / NHTSA-2020-0106

Avoid complete preemption. States and localities need the ability to adapt to local situations.

- Federal government should not preempt ability of states to control and enforce their own traffic laws
- States should be able to license computer drivers and revoke licenses for reckless driving
- States and municipalities should be able to ticket breaches of traffic laws by computer drivers

- NHTSA should not be only regulator with authority over [rolling through stop signs](#), etc.
- States should be able to limit testing to reduce impact on disadvantaged and other areas of concern
- Municipalities should be able to regulate testing (e.g., avoid schools, churches, festivals, and so on)
- States and municipalities need ability to curtail reckless and disruptive testing and operations
- Local governments need ability to constrain AV operations responsive to specific local hazards
- Should not preempt property, injury, death, or product liability claims grounded in state law

Industry-Promoted Myths. Many common [industry talking points are misleading or outright false.](#)

- *False:* Claim US regulations are stifling the creation of safe, reliable computer drivers
 - The longer timeline is due to getting the technology to first work, then be safe and reliable.
- *Misleading:* Claim 94% of crashes are due to human error
 - Computers make mistakes too. Nobody knows if computers will reduce the fatality rate.
- *Misleading:* Claim product liability and tort laws are efficient and effective for computer drivers
 - Computer drivers are not a legal entity, and some state laws serve to deflect accountability
 - Product liability claims for this technology can be expensive and difficult to pursue
- *Misleading or False:* Claim that manufacturers already follow industry safety standards and practices
 - Companies do not claim to follow their own standards, and some actively reject them
- *False:* Claim that automation of steering (Level 2/2+) is a safety feature
 - Driver Assistance (ADAS) safety data is for active safety features, not automated steering
- *False:* Claim that computer drivers are inevitably safer than human drivers
 - This is pure speculation. Computers will make different mistakes than people.
- *False:* Claim current data proves that computer drivers are safer with regard to fatalities
 - The industry needs 100 million more miles (at least) to support such a claim.

About the author:

Prof. Philip Koopman is an internationally recognized expert on Autonomous Vehicle (AV) safety whose work in that area spans over 25 years. He is also actively involved with AV policy and standards as well as more general embedded system design and software quality. His pioneering research work includes software robustness testing and run time monitoring of autonomous systems to identify how they break and how to fix them. He has extensive experience in software safety and software quality across numerous transportation, industrial, and defense application domains including conventional automotive software and hardware systems. He was the principal technical contributor to the UL 4600 standard for autonomous system safety issued in 2020. He is a faculty member of the Carnegie Mellon University ECE department where he teaches software skills for mission-critical systems. In 2018 he was awarded the highly selective IEEE-SSIT Carl Barus Award for outstanding service in the public interest for his work in promoting automotive computer-based system safety. In 2022 he was named to the National Safety Council's Mobility Safety Advisory Group. He is the author of the books: Better Embedded System Software (2010), How Safe is Safe Enough: measuring and predicting autonomous vehicle safety (2022), and The UL 4600 Guidebook (2022).

Additional information: <https://users.ece.cmu.edu/~koopman/> Contact: koopman@cmu.edu

Disclaimer: These opinions are my own and are not a statement being made by Carnegie Mellon University.

Support materials:

- How safe is safe enough for autonomous vehicles: https://youtu.be/UTdR_HE3DDw
- Autonomous vehicles and software safety engineering: https://youtu.be/oE_2rBxNrfc
- Trust and governance for autonomous vehicle deployment: <https://youtu.be/hZQyFc9ETCE>
- Starting point for State AV statutes: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4444848

Sample Q&A for Dr. Koopman

Industry Disinformation:

Q: Is it true that 94% of crashes are due to human error?

A: That is a big distortion of a statistic intended for something else. The [original source says that 94% of crashes involved a human](#). It is true that drunk driving and other misbehaviors are involved in perhaps half the crashes. But many times the human involvement is not being able to correct for an unsafe road condition or react to an equipment failure in time. The original report clearly states the statistic does not assign blame to the driver for causing the crash. NTSB and NHTSA have both made it clear the “94% human error” claim is misinformation. The reality is that computers make some of the same mistakes humans do (often for different reasons), and will make plenty of their own mistakes. It turns out that computers make errors too.

Q: Is it true that human drivers are terrible?

A: As imperfect as humans are, they are impressive at handling chaotic, unexpected situations. They drive on average up to 100 million miles between fatal crashes, and if you leave out the drunk and distracted drivers it is much better than that. Computer drivers are terrible at the unexpected, and it might take a new type of technology to fully address that issue. We don’t know yet whether computers or human drivers will be better, and we certainly cannot assume computers will automatically be better.

Q: Are computer drivers safer than human drivers?

A: We have 1 or 2 or 3 million miles of robotaxi operation now, depending on the company. At 100 million miles or more between human driver fatalities, it's another 97 million or more miles before we might confirm computer drivers are safer – assuming there are zero fatalities before then. Nobody knows whether people or computers will be better. It is likely that a human assisted by a safety-minded computer (rather than a steering convenience feature) would be safest of all for at least the next few years, and potentially much longer.

Arbitration:

Q: Does use of forced arbitration hurt public safety?

A: Yes. Safety requires transparency. Safety issues relevant to forced private arbitration matters might not be brought to the public's attention, impairing transparency. This in turn is likely to delay or reduce the chance they will be independently investigated and resolved. Additionally, if forced private arbitration is used as a tactic to reduce compensation to harmed parties, that reduces the effectiveness of tort law in incentivizing companies to behave in a responsible manner with regard to public safety.

Preemption:

Q: Can you give an example of a situation in which preemption has hurt the public interest?

A: The San Francisco Fire Department has logged dozens of robotaxi incidents including running over a fire hose and blocking fire trucks. City officials are powerless to intervene because of a state preemption law. The City should be able to institute a rule such as a robotaxi is not permitted to operate, for example, within two blocks of a fire, is not allowed to block a fire house garage driveway, and must keep at least two blocks away from any police response scene until these issues can be resolved. States and municipalities should not be preempted from placing specific restrictions on automated vehicles to resolve pain points. Those pain points are likely to be different than ones for human drivers, requiring different local rules.

Safety:

Q: You mentioned "acceptable safety" in your testimony. How safe is safe enough?

A: Everyone wants this technology to be as safe as possible. It will never be perfectly safe, so the question is when is safe enough. To remove the safety driver, automated vehicles (1) should conform to industry-written safety standards for autonomous vehicles to create a presumption of likely safe enough outcomes, (2) should not redistribute risk onto vulnerable populations, and (3) should not drive in a way that would be negligent if it were a human driver. Long term outcomes might be measured by comparison with

comparable human driver crash rates, including fatalities. But such comparisons have limited predictive power in the early days, where we are now.

Q: It has been said that airline safety rules are written by the blood of the victims of airline crashes. Is that going to be necessary with AVs or is there a safer way to get there through technology?

A: The current industry-written safety standards incorporate lessons learned from airline crashes. In fact, an aviation expert helped me make sure to incorporate those learnings I was writing the original UL 4600 proposal document that resulted in an industry-consensus safety standard being issued more than three years ago. Our choice is to follow the standards for automated vehicles that are already issued, or re-learn those same lessons – in blood – at those companies who decide not to follow their own industry’s safety standards.

Regulation:

Q: Is it true the regulations stifle innovation?

A: Mercedes Benz has an SAE Level 3 vehicle (driver can relax while driving and not look at the road) in operation in Europe. It happened there first and not in the US precisely because they had a regulation called [UN ECE 157](#) – basically their version of an FMVSS – that allowed certification before the vehicle went on the road. That regulatory certainty gave the manufacturer assurance because they conformed to a safety standard before deploying. In the US, companies do not have that assurance. At least for Level 3 automation, *lack* of a specific regulation is what seems to have stifled innovation here.

Q: What is the regulatory situation now in the US for autonomous vehicles?

A: Companies can shop for a state that they think has favorable rules, and many states have been eager to create very favorable rules. Any company that can get a vehicle working well enough to do road testing should find getting permission a simple matter of filling out some paperwork and providing evidence of insurance coverage in an increasing number of states. If they can manage to test for a while, say a year, without a major crash, they can likely get permission to deploy without a safety driver. Under the preferred industry rules, things are even more lax. There are no substantive impediments to deploying and maturing the technology. Even California does not require conformance to any industry safety standards. Not even

road testing safety standards. Foreign companies come to the US specifically because it is so much easier to test and operate here compared to other countries.

Q: Do you have an opinion on the role of the Federal Government in regulating AVs?

Autonomous Vehicle (AV) regulation is complicated because the historical role of the States has been to regulate drivers, while the Federal Government regulates equipment. With an AV, the driver is a piece of equipment, resulting in a potential conflict. The States should retain their ability to define acceptable and unacceptable driver behavior. Municipalities should retain their ability to modify driver behavior requirements responsive to local conditions. The Federal Government should clarify that because the computer driver is a piece of equipment, the manufacturer is responsible for its safety. The Federal Government should ensure that the equipment meets industry-written safety standards, meets government-written safety standards, and does not have faults that interfere with its ability to adhere to State and Municipal requirements for driver behavior. The Federal Government should also continue and strengthen the current “SGO” data reporting requirements. The Federal Government should not subsume the role of creating and enforcing driver behavior rules.

Q: Do we really need a massive increase in FMVSS exemptions to benefit from AV technology?

A: [Zoox has self-certified its purpose-built robotaxi to FMVSS without exemptions](#). NHTSA is currently working through the paperwork to make sure everything is in order. So at least one company seems to be happy without additional exemptions while building a vehicle comparable in configuration to the ones other companies are seeking exemptions for. Companies might think that keeping the FMVSS status quo is inconvenient, but Zoox is making the current system work given that part of [FMVSS \(the 200 series\) for crash protection safety has already been updated to permit driverless vehicles](#). Having read the exemption requests from Waymo and Cruise, [I think we need to set the bar higher](#) than where they want to set it before considering increasing the number of exemptions.

Q: Can overly lax regulations actually hurt the industry long-term?

A: My understanding is that a significant source of problems for the rail industry is rooted in events of more than a century ago. They had an under-regulation situation followed by [public backlash after a series of high profile fatality events](#). The outcomes stifled rail industry innovation for decades. Competitive pressures in any industry incentivize companies to cut corners on safety if they think they can get away with it. (We just have to look at the 737-MAX fiasco to see that dynamic in recent events.) If the manufacturers do not have reasonable requirements to adopt their own safety practices for vehicle automation, it is just a matter of time before we see public blowback after loss events. That type of scenario could undermine or possibly even collapse that industry for years or decades to come, and I think that the practices of at least some companies make that a likely outcome unless something changes.

Liability:

Q: Companies often blame liability as an impediment to deployment. How might we resolve that?

Q: How do we balance the need to innovate vs. protecting consumer rights?

Q: How do we address the patchwork of state laws governing vehicle automation?

A: Liability varies by state, because manufacturers have been lobbying for the responsible party to be anyone but them, getting a range of different outcomes in various states. The “patchwork” they complain about is largely self-inflicted. We can resolve much of this problem without a lot of disruption by explicitly providing that a computer driver has the same duty of care to other road users as a human driver, with the manufacturer being the responsible party for a computer driver. Then you should ensure that cases are not subject to forced private arbitration, but rather let the court system take it from there, making sure there is no federal preemption of state tort law and product liability law. One of the current trends that should be stopped and reversed is the state-by-state industry campaign to increase the transaction costs and difficulty of holding a manufacturer accountable for a loss by passing laws that deflect responsibility so it will be a difficult, expensive, and lengthy process to hold them accountable for harm caused by their computer driver.

China:

Q: Is China actually ahead of the US in self driving cars?

A: Right now there are robotaxis and robo-trucks on public roads in China, the US, and Europe. Recently many companies in all three regions have been walking back their ambitions because the problems of reliability and safety are turning out to be harder than they thought. However, we've seen that stronger regulatory systems in Europe have let them get ahead in Level 3 vehicles by providing regulatory certainty that they are approved before they go on the road, then bring that technology to the US. It would be no surprise if the stronger regulatory system in China results in a similar outcome over time. A weak regulatory system that does not encourage companies to emphasize safety and reliability will hurt us more than it helps us in this race, by forcing our companies to respond to investor demands for more deployment without providing regulatory pressure for them to spend sufficient resources on safety.

Cybersecurity:

Q: Does cybersecurity pose special risks for autonomous vehicles?

Q: What should NHTSA do about cybersecurity for autonomous vehicles?

A: The cybersecurity risk for autonomous vehicles is a special concern for new cars that have automated steering capabilities, not just fully autonomous vehicles. If the car can drive without the driver's hands on the wheel for a few seconds, it can be made by a malicious attacker to drive *dangerously* without a driver at all. If it has cameras, there can be privacy concerns regardless of whether it is a robotaxi or just reading traffic signs. This is not a robotaxi or robotruck problem – it is a problem for all new vehicles. The automotive industry has systematically underinvested (and I'm politely understating this) in cybersecurity since the time I was writing cryptographic algorithms for vehicle security in the early 1990s. Something we can do near-term is require conformance to the industry's own cyber security standard: ISO/SAE 21434.

Regulatory Authority:

Q: What does NHTSA need to effectively regulate this technology?

A: There are two big needs. The first is a mandate for NHTSA to require the industry to self-certify to its own industry-written safety standards for self-driving cars. Those standards have been in circulation for years, and were written by the industry, for the industry. The second is to dramatically increase staffing in the area of software vehicle features. They have started hiring in this area, but the hiring rate needs to be much more aggressive than is currently budgeted.

Traffic Safety Facts from: <https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/813428>

- NHTSA estimates 42,795 motor vehicle traffic fatalities for 2022.
- Fatality rate of 1.35 per 100 million miles (about 75 Million Miles per fatality; sometimes with more than one fatality per crash).

[2021 annual fatality numbers](#): 7,342 pedestrian, 985 pedacyclist, 6,101 motorcyclist; 11,780 speeding, 8,174 alcohol-related. 3,522 related [to distracted driving](#). (These numbers overlap categories.)