Carnegie Mellon University

August 24, 2023

Jessica Herron
Legislative Clerk
Subcommittee on Innovation, Data, and Commerce
House Committee on Energy and Commerce
2125 Rayburn House Office Building
Washington, DC 20515-6 I 15

Re: Philip Koopman's Responses to Additional Questions for the Record

Dear Ms. Herron:

I want to thank the Subcommittee for inviting me to appear before it on July 26, 2023 to testify at the hearing on the topic of Self-Driving Vehicle Legislative Frameworks.

Pursuant to the Rules of the Committee on Energy and Commerce, I am attaching my answers to additional questions for the record.

Thank you again for your help, and please let me know if you have any questions.

Sincerely,

Philip Koopman, Ph.D.
Associate Professor, ECE Department, Carnegie Mellon University
(Speaking as an individual, and not as a representative of the University)

Attachment—Additional Questions for the Record

Prof. Philip Koopman
Wednesday, July 26, 2023 Hearing: Self-Driving Vehicle Legislative Framework

The Honorable Janice Schakowsky

Several serious crashes involving cars equipped with autonomous technology, which is unregulated, have already occurred, many of which have been subject to investigation by the National Transportation Safety Board (NTSB) and the National Highway Traffic Safety Administration (NHTSA). In addition, The *Washington Post* reported last month that according to NHTSA data, there have been 17 fatal incidents, five serious injuries and 736 crashes involving Tesla vehicles operating in Autopilot mode since 2019. Furthermore, according to data collected by NHTSA's Standing General Order (SGO) 2021-1 requiring manufacturers to report certain crashes involving vehicles equipped with automated driving systems (ADS) or SAE Level 2 ADAS, there have been approximately 330 crashes involving ADS and 1040 with ADAS. These include 25 crashes resulting in a fatality.

In addition, San Francisco transportation officials have documented numerous dangerous incidents involving AVs operating in the city. These events include:

- **Interfering with emergency response operations including 18 incidents documented by the San Francisco Fire Department in which AVs put firefighters and the public at risk.**
- **Making planned and unplanned stops in travel lanes that have interfered with transit service and blocked traffic.**
- **Intrusions into construction zones where City employees were working.**
- **Obstructions caused by AVs having to interpret and respond to human traffic control officers.**
- **Erratic driving.**

1. **Dr. Koopman, The NTSB has identified numerous serious safety issues involving automated driving systems. In addition, NHTSA is investigating a number of frightening crashes involving Tesla vehicles including with first responder vehicles. Test vehicles in San Francisco continue to be involved in concerning incidents such as failing to obey commands by law enforcement and blocking traffic.**
    - **Can you please expound upon the current substantial safety concerns associated with these vehicles?**

Response: There are substantial safety concerns with both types of vehicles covered by this question: those with human drivers and those without in-vehicle human drivers. Since the time this question was posed there has been a documented injury collision between a robotaxi and a fire truck in San Francisco, which has been attributed to the failure of the robotaxi to yield as required to an emergency response vehicle with activated emergency lighting and siren.

SAE Level 2 & 3 vehicles:

For vehicles with human drivers, the NHTSA investigation into Tesla crashes with emergency response vehicles is still ongoing, having documented 14 crashes/fires, 15 injuries, and one fatality as of June 2022 regarding crashes involving first responders and road maintenance vehicles. (NHTSA Engineering Analysis EA 22-002 https://static.nhtsa.gov/odi/inv/2022/INOA-EA22002-3184.PDF) A concerning statement in the NHTSA document is: "On average in these crashes, Autopilot aborted vehicle control less than one second prior to the first impact." This is a clear indication that driver monitoring capabilities in these vehicles are inadequate to ensure the driver is paying sufficient attention to achieve acceptable safety outcomes. There is no reason to believe that companies continuing to blame drivers as a way to deflect criticism of the technology will improve future safety outcomes.

As stated in your question, there have been many more fatalities and injuries beyond the scope of that NHTSA Engineering Analysis, and multiple recalls for features released to public roads that have later been determined to be unsafe. For example, a recent NHTSA safety recall addressed Tesla issues of: "1) traveling or turning through certain intersections during a stale yellow traffic light; 2) the perceived duration of the vehicle's static position at certain intersections with a stop sign, particularly when the intersection is clear of any other road users; 3) adjusting vehicle speed while traveling through certain variable speed zones, based on detected speed limit signage and/or the vehicle's speed offset setting that is adjusted by the driver; and 4) negotiating a lane change out of certain turn-only lanes to continue traveling straight." (NHTSA Part 573 Safety Recall Report 23V-085 https://static.nhtsa.gov/odi/rcl/2023/RCLRPT-23V085-3451.PDF)

The NTSB has issued investigative outcomes that highlight "considerable limitations in detecting hazards, as well as in maintaining an appropriate travel path" as well as the risk of driver disengagement due to automation complacency. (Source: https://www.ntsb.gov/Advocacy/safety-topics/Pages/automated-vehicles-investigative-outcomes.aspx)

While vehicles from other manufacturers will differ in their performance, the issues of automation complacency and concerns about driving performance in low probability but high consequence situations are inherent limitations to the technology. Moreover, Level 2 automated driving technology for all practical purposes is unregulated (aside from occasional recalls and the NHTSA SGO data reporting requirement).

A self-driving car bill should bring SAE Level 2 features under an automated vehicle regulatory umbrella along with SAE Level 3 vehicles. Any vehicle feature which provides sustained automation of steering control is an automated driving feature, and should be regulated as such. The safety emphasis for both Levels 2 and 3 with regard to human drivers should be: (1) requiring effective driver monitoring, (2) ensuring the human driver has adequate opportunity to recognize the need for transition and safely take over vehicle control when required, (3) forbidding the sale of immature "beta test" technology to retail customers, (4) requiring all vehicles with sustained automated steering features to prevent operation of that feature outside the manufacturer-specified Operational Design Domain (ODD), and (5) prohibiting the practice of using a "moral crumple zone" tactic to dump blame on a human driver for crashes in which

the driving automation feature did not monitor the human driver's attention adequately and/or it was unreasonable to expect an ordinary licensed driver (accounting for automation complacency) to transition safely to manual driving in time to prevent a crash.

SAE Level 3, 4, 5 vehicles:

Deployment of uncrewed driverless vehicles in San Francisco has highlighted the fact that computer drivers are imperfect, and will make mistakes. Sometimes those are mistakes that human drivers also make (running into wet concrete, hitting the back of a city bus, driving through closed roads, failing to yield to a fire truck with lights and siren activated). But they will also make some mistakes that are atypical of human drivers (not being able to move the vehicle when being shouted at by emergency responders; blocking an intersection for tens of minutes due to confusion; blocking multiple travel lanes for extended times due to loss of cellular data services).

The month of August 2023 has seen three injury crashes involving robotaxis in San Franciso, including one each from Waymo and Cruise that involved injury to a robotaxi passenger. While some (but not all) crashes might be blamed primarily on other road users, it is important to note that prudent defensive driving is also an important contribution to safety, and what matters in the end is not blame, but rather whether total harm is increased or decreased.

The vehicles have also interfered with emergency responders as documented by the City of San Francisco dozens of times. A recent crash with a fire truck responding to an emergency (with lights and sirens activated) has prompted California DMV to cut Cruise's operational fleet size by 50% pending an investigation.

These events underscore the fact that computers are not automatically safer drivers than people. Legislation should not assume that safety will automatically be improved by automated driving technology.

The NHTSA proposed framework from December 2020 (Docket NHTSA-2020-0106) would require the industry to use their own industry-written consensus safety standards before releasing uncrewed vehicles on public roads. Those standards already address the risks seen in the San Francisco mishaps and incidents. If companies had simply followed those standards, most of the news articles about the chaos on San Francisco roads would never have needed to be written. Congress should require NHTSA to proceed in a timely manner on pushing forward that framework with an update to include the latest industry consensus safety standards.

Congress should also require NHTSA to make the current SGO data reporting requirement both permanent and more robust, to include events that interfere with public safety such as blocking fire trucks responding to a scene, even if they do not involve a severe crash event.

- **Is it accurate to claim that the autonomous taxis currently operating in San Francisco are safer than human drivers in light of all the concerning incidents that have already occurred involving these vehicles?**

Response: The news from San Francisco highlights that there is more to safety than avoiding crashes.

A relentless industry narrative amounts to "human drivers make mistakes, so error-free computers will improve safety." We now have ample evidence from real-world operation that this narrative is false. Computer-driven cars also make mistakes, some of which are just as embarrassing as human driver errors – and some of which are errors that one would not expect from a human driver, such as an inability to move a drivable vehicle immediately when forcefully directed to do so by a police officer.

Waymo has seemingly credible preliminary data and a prediction methodology that claims they are having fewer minor crashes overall, with data analysis suggesting reduced predicted harm as data continues to accumulate. This is a good sign, but is a prediction nonetheless. Cruise has made similar claims that do not seem practical to evaluate for credibility given a paucity of public-facing information.

However, no company has enough data to confirm their claims that they have already reduced fatalities. At approximately 100 million miles between fatalities (San Francisco numbers vary significantly by year, but this is in the range of recent years), they have at least 95 million miles to go before it is possible to reasonably confirm a prediction of reduced fatalities based on real-world operational evidence. That leaves a lot of time for the unexpected to go wrong. To be clear, this does not mean they are unsafe, because nobody knows how this will turn out. But it does mean that it is far too early for the AV industry to declare victory on reducing road fatality rates for all companies.

Recent events highlight that safety cases need to also address public safety, and in particular the effects of blocking public roads, impeding the progress of emergency vehicles, and interfering with emergency responders a scene.

Continued safety concerns are justified. The industry needs to be more transparent with incident data. Current redaction discretion appears to be abused by companies to hide information relevant to public understanding of the risks presented by this still-experimental technology. Congress should mandate that NHTSA publish SGO reports unredacted. Safety outcomes are a legitimate public concern for this technology, and secrecy benefits only companies who are not as safe as they should be.

Many promises have been made about autonomous vehicles (AVs) bringing meaningful and lasting reductions in motor vehicle crashes and resulting deaths and injuries, traffic congestion and vehicle emissions. Additionally, claims have been made that AVs will expand mobility and accessibility, improve efficiency, and create more equitable transportation options and opportunities. However, as Transportation Secretary Buttigieg and others in the auto industry have acknowledged, these outcomes are far from certain.

2. **Dr. Koopman, The proponents of autonomous vehicles have made bold but largely undocumented predictions that AVs will eliminate crashes, reduce congestion, benefit the environment, expand mobility and accessibility, improve efficiency, and create more equitable transportation options and opportunities. How certain are you that any of these benefits will come to pass in the near future?**

Response: In the absence of regulation, we depend on the goodwill of the AV companies. It seems imprudent to assume that for-profit companies will necessarily behave in a way that reduces or foregoes profitable operations. Even if doing so is an initial a gesture of goodwill, if the business model is not fundamentally based on providing goods and services inherently aligned with promised benefits, financial and competitive pressure can too easily cause a lapse in fulfilling promises.

It is relevant to note that while state regulators have the ability to enforce some of these outcomes, we have not seen that happen in California. For example, CPUC has not required equitable access to robotaxis despite public testimony that the vehicles do not provide, for example, adequate wheelchair user access.

- Eliminate crashes: We already know this is untrue, and have seen multiple crashes with injuries with at least some fault contribution from a robotaxi. We can hope AVs will overall reduce crashes and reduce harm, but such benefits across the industry are still aspirational rather than based on accumulated evidence.

- Reduce congestion and emissions: Claims for these benefits too often lack rigor. For example, claims based on average vehicle utilization and parking might overlook peak rush hour commuting and where that peak fleet goes to park during quiet hours. Claims often assume shared vehicles despite evidence that many riders do not wish to share in current ride hailing networks. Indeed, at a recent CPUC hearing many speakers extolled the fact that a robotaxi could be ridden alone, without even a human driver, as a personal safety measure that is only attainable without ride sharing. While solo rides can bring personal safety benefits, they do not reduce congestion and emissions. It is difficult to believe that reducing the cost per mile of ride sharing will do anything other than increase demand, thereby increasing congestion, increasing net emissions, and decreasing net safety compared to policies that promote more adoption of mass transit. Reduced congestion and reduced emissions are certainly possible, but will require significant policy changes beyond the scope of simply encouraging AV technology adoption, and those policy changes have not yet happened.

- Expanding mobility and accessibility: Human-driver ride hailing networks were promised to improve mobility and accessibility, but achieved limited results. In many regards the robotaxi industry is showing us a remake of that same movie, and we should have appropriately limited expectations for the actual benefits that will be realized. Removing human drivers seems likely to provide an incremental rather than revolutionary improvement in the near term (basically, a slightly cheaper Uber/Lyft/Taxi ride) so long as the operational areas and capital equipment costs remain high. Some day the car of the future that is in every garage might bring these benefits, but that is not any day soon.

While it is desirable to make progress to that future, current road users should not be put at increased risk against an aspirational Jetson's future that might easily be another 25 years away. (I hope it comes sooner, but we should not compromise current constituent safety against a promise that might well be that far away.)

- **Are you aware of any evidence to date that automated driving will improve the horrific fatality rate on our Nation's roads?**

Response: No. For now this remains an aspirational goal of the industry.

Claims that Tesla automated driving systems improve safety have been debunked, with a study showing an 11% increase in crash rates with autopilot turned on compared to the same vehicle with autopilot turned off in comparable operational conditions. (Noah Goodall (2023) Normalizing crash risk of partially automated vehicles under sparse data, Journal of Transportation Safety & Security, DOI: 10.1080/19439962.2023.2178566)

While robotaxi companies are making claims that they are already reducing fatality rates, those claims are based on predictions involving models with significant assumptions. For example, having the same computer driver in hundreds of robotaxis can increase the risk of a common cause failure that must be mitigated using rigorous safety engineering according to the industry consensus safety standards the companies do not want to say they follow. (For example, we hope that the change from Daylight Savings Time to Standard Time will not cause a fleet-wide failure as sometimes happens in mobile phones and other computer-based systems. But we have already seen that cell phone data network failures and power disruptions that disable traffic lights have demonstrated common-mode failure issues with these vehicles, fortunately not leading to fatalities.) We can hope companies will never make a catastrophic engineering mistake – despite not being required to follow industry consensus safety standards aimed at mitigating exactly those types of risks. But it will take a hundred million miles or more of further road data to provide any reasonable confidence that claims of improved fatality rates are in fact supported by real-world outcomes.

We can hope that this technology will improve road safety. And, with time and continuing maturation of the technology, there is every reason to expect it will, over a likely period of decades. But we should not base policy decisions on an aspirational assumption that this is necessarily true. In particular, we should not impose an elevated risk on current constituents today in hopes that things will improve tomorrow.

3. **Mr. Shapiro, President and CEO of the Consumer Technology Association, made the following assertions:**
   - **"Our adversaries and allies are moving forward without us. And, just since the last hearing we had on this in 2020, China, Europe, Japan have all moved forward and South Korea even has a target that half of its cars be self driving by 2035."**
   - **"As an economic issue, it will create 300 billion to 400 billion in revenue by 2035."**

- **Dr. Koopman, do you agree with these remarks?**
  - **Are you aware of original, independent substantiation and citations (not including media reports) for these remarks?**

Response:
The term "self driving" car in the quoted statement is ambiguous, potentially muddying the waters in this context. It is important to distinguish autonomous vehicles (SAE Level 4/5) from more generic partially automated vehicles (SAE Level 2/3, with 3 also having a fully automated driving automation mode but still requiring an available human driver).

I know of no basis for stating that any country is ahead of the US in SAE Level 2 vehicles. Indeed, Consumer Reports found that Ford and GM had the highest human-supervised automated driving (Level 2) testing scores of included US, European, and Korean brands as of January 2023. (https://www.consumerreports.org/cars/car-safety/active-driving-assistance-systems-review-a2103632203/)

Other countries and especially Germany seem ahead in the SAE Level 3 race, with Mercedes Benz saying they will introduce a Level 3 vehicle in the US this year that permits the driver to take their eyes off the road while driving. That advance was possible because of, not despite, their Level 3 regulatory framework and requirement for pre-deployment type approval according to UN ECE #157 ALKS regulations.

There is no credible reason I am aware of to believe any country is ahead of the US in terms of practical, scalable SAE Level 4/5 vehicles, and especially there is no credible basis for a statement that regulatory burdens are hindering technical development in the US. Rather, companies from other countries come to the US to test here in large part because of our lax regulatory approaches to road testing safety.

Level 4/5 vehicles are less regulated in the US than in competitor countries, yet we've seen a rocky launch in San Francisco by US companies. Further institutionalizing the NHTSA "non-regulatory approach" still in effect will not help us compete. Rather it seems likely to hurt us because the ultimate competition to be won will be based on reliability and safety. The NHTSA proposed framework from December 2020 (Docket NHTSA-2020-0106) should be pursued in a timely manner to improve our competitiveness by setting a level playing field for safety for all companies, incentivizing US companies to adopt industry consensus standards that remove a competitive advantage in time to market to those who might wish to take safety shortcuts.

4. **Mr. Shapiro also asserted, "94% of crashes are because of human error. There's no other example of what it could be other than a car breaking down which is very rare…." During your remarks you mentioned that NTSB Chair Jennifer Homendy has discredited the improper use of this fact.**

   - **Dr. Koopman, is it true that the National Highway Traffic Safety Administration stated in the same document with this statistic that "[a]lthough the critical reason is an important part of the description of events leading up to the crash, it**

is not intended to be interpreted as the cause of the crash nor as the assignment of the fault to the driver, vehicle, or environment."[1]?

Response: The primary source for the relevant document with the "94% human error" quote I was referring to in my testimony seems to have been recently withdrawn. This URL https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/812115 was the official source I was aware of, but which now gives an error message of: "Sorry, it seems like we do not have this publication or it is currently unavailable" at the time of this writing.

However, a version of that document was active on August 18, 2023 per the Internet Archive, and corresponds to my personal memory of the document's contents: https://web.archive.org/web/20230818112401/https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/812115

The bottom right of that document states: "The critical reason was assigned to drivers in an estimated 2,046,000 crashes that comprise 94 percent of the NMVCCS crashes at the national level. However, in none of these cases was the assignment intended to blame the driver for causing the crash."

Figure showing that quote:

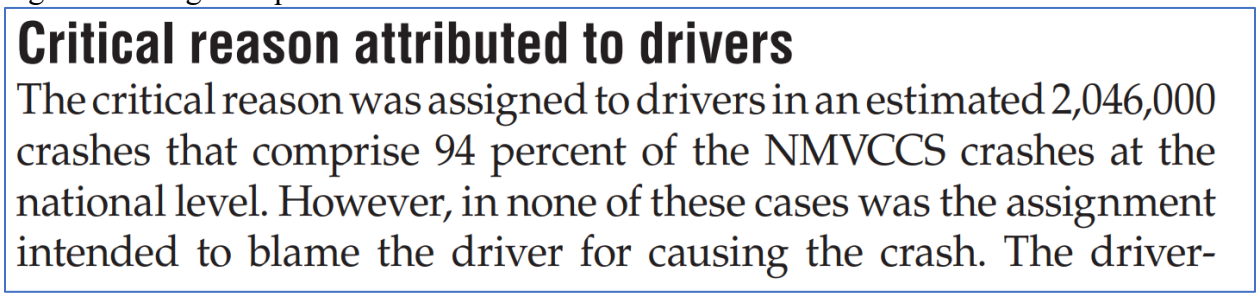

NTSB Chair Jennifer Homendy gave a press interview to debunk the "94% human error" narrative here: Wilson, "'It Ain't 94 Percent': NTSB Chair Jennifer Homendy Discusses the Role of Human Error in Car Crashes," Streetsblog USA, Jan. 31, 2022. https://usa.streetsblog.org/2022/01/31/it-aint-94-percent-ntsb-chair-jennifer-homendy-on-the-role-of-human-error-in-car-crashes

Other factors that contribute to reduced road safety include: poor roadway design for safety, subpar maintenance, overly high speed limits, neglecting vulnerable road user protection as a vehicle safety requirement, underinvesting in vulnerable road user protective infrastructure, and policies that incentivize driving personal vehicles over mass transit. Many of the crashes blamed on human drivers might be avoided by improvements in these other areas. In other words, driver blame is often assigned if a driver fails to compensate for the results of poor policy choices, which can be the case independent of whether legitimately poor choices by drivers are also

[1] Singh, S. (2015, February). Critical reasons for crashes investigated in the National Motor Vehicle Crash Causation Survey. (Traffic Safety Facts Crash Stats. Report No. DOT HS 812 115). Washington, DC: National Highway Traffic Safety Administration.

involved. We have spent decades blaming drivers with limited results while other countries are making dramatic improvements in road safety. Perhaps it is time to try a different approach.

5. Mr. Shapiro said that motorcyclists are identifiable to AVs today.

- **Dr. Koopman, are you aware of any original, independent substantiation and citations (not including media reports) for this assertion?**

Response: NHTSA data shows that about one third of fatalities on public roads are vulnerable road users. It is essential that vehicle automation manufacturers are mindful of not only the risk of harm to vehicle occupants, but also to other road users including motorcyclists to avoid making this situation worse. However, Mr. Shapiro's statement provides no assurance that this is the case.

2022 data; source: https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/813448

All Fatalities	42,795
Motorcyclist	6,000
Pedestrian	7,345
Pedalcyclist	1,068

It is important to understand the huge safety gap between a statement that *some* motorcyclists can be identified on *most* self-driving vehicles today (the seems to be the meaning of the statement attributed to Mr. Shapiro) and a statement that *essentially all* motorcyclists can be identified by *all* self-driving vehicles (what would be expected for road safety). Mr. Shapiro's assertion has little bearing on road safety. (Per the preliminary meeting transcript, Mr. Shapiro also mentioned an absence of blind spots, which is necessary but not sufficient for safety, since the detection probability of a motorcycle not in a blind spot also matters, and might well be imperfect.) This goes back to just because humans are imperfect does not necessarily mean computer drivers will be perfect.

As an example, a scholarly study on this topic from the point of view of CCTV roadside cameras found that a commonly cited reference system ("YOLO") had an "F1-Score" (a commonly used blended performance metric) of 86% for tracking motorcycles, and a proposed improved system ("EspiNet") had an F1-Score of 94%. While 94% percent seems like a high number, when used on many motorcycle images those 6% performance issues add up. Safety typically requires performance numbers at the system level of 99.999…% with the number of 9s after the decimal depending on the specifics of the metric. As an illustration of this gap, that referenced study showed thousands of motorcycles not detected by the improved system. (J. E. Espinosa, S. A. Velastín and J. W. Branch, "Detection of Motorcycles in Urban Traffic Using Video Analysis: A Review," in IEEE Transactions on Intelligent Transportation Systems, vol. 22, no. 10, pp. 6115-6130, Oct. 2021, doi: 10.1109/TITS.2020.2997084)

That one study is not the final word, there are limitations to the results, the use case for the technology is different, and as the paper itself says public literature on this topic seems relatively sparse. So that paper does not prove there is a safety problem – but it certainly underscores that

this is an area that requires continued attention and is far from a solved problem. One would hope that companies building AVs to run on public roads do better by, for example, fusing the results from multiple types of sensors that they have proven do not suffer from common mode failures. However, I am not aware of AV companies sharing their detection performance statistics for motorcyclists. So we do not know how well such companies do, and are unlikely to know the risk presented to motorcyclists and other vulnerable road users until there is enough road experience data to see how many fatalities show up in the NHTSA SGO database over time.

The important point here is that a claim that "motorcyclists are identifiable" is not a sufficient basis for a policy decision. The required claim would be that "motorcyclists are correctly identified with extremely high probability in a way that supports reduced harm compared to encounters with human-driven vehicles." Moreover, such a claim must be backed by data rather than stated as a general assertion without evidence. I am not aware of data from the industry that supports such a statement. We simply do not know if the industry is acceptably safe at identifying motorcyclists at this time.