December 20, 2019

The Honorable Edward J. Markey
United States Senator
Suite SD-255
Dirksen Building
Washington, DC 20510-2107

Re: Responsible Driving with Autopilot

Dear Senator Markey,

This letter responds to your November 13, 2019, request for information on actions that Tesla is taking to prevent drivers from abusing or evading certain Autopilot features. We appreciate your request and are happy to share more information.

Tesla fully shares your commitment to increase occupant safety and to reduce fatalities on U.S. roads. We also share your belief in developing active safety features that promote responsible driving habits. However, we must clarify some misperceptions about our product, its design, and its safety performance.

First, throughout the purchase, user and ownership experience, Tesla clearly explains that Autopilot is not an autonomous system and does not make our vehicles autonomous. Autopilot is an advanced driver-assistance system ("ADAS") that is representative of SAE International Level 2 automation ("SAE L2"). As such, Autopilot is only designed to assist the driver in performing the driving tasks of steering, acceleration, deceleration, and lane changes. The driver must continually monitor the driving environment and be prepared to immediately overtake the vehicle controls as necessary. The driver is forced to participate through steering wheel detection. When used properly, Autopilot can greatly enhance occupant safety, but, as an SAE L2 ADAS system, the driver is ultimately responsible for the safe operation of his vehicle.

Second, every day Tesla collects anonymized data from millions of miles driven by our customers with Autopilot engaged. We capture a wide array of driving events so that our system can learn from real drivers, and we run potential software improvements in the background for analysis. Tesla analyzes and learns from the data and uses it to continuously improve Autopilot and our customers' interaction with it. The data shows that Autopilot significantly increases overall occupant safety. In the third quarter of 2019 alone, our customers registered one collision for every 4.34 million miles driven with

Autopilot engaged. By comparison in that same time period, our customers who did not have Autopilot engaged, and instead only had basic active safety features like forward crash warning ("FCW") and automatic emergency braking ("AEB"), engaged registered one collision for every 2.70 million miles driven. Thus, during this quarter our customers were 62% less likely to crash when Autopilot is used. Since we began publishing this information in 2018, the improvement has ranged from about 60-80% fewer crashes per mile driven with Autopilot engaged.

This data is directly measured across the entire fleet, not just sampled, and accounts for even mild crashes in which airbag or seat belt retractors deployed (generally, any impact greater than about 12 mph). The most recent national average from crash data from NHTSA and driving data from FHWA, shows one collision for every 498,000 miles driven (and that is based on under-estimated police-reported crashes). In comparison, these statistics show that miles driven with Autopilot engaged are significantly safer than miles driven without and overwhelmingly safer than the national average.

Tesla takes the risk of improper use or abuse of Autopilot very seriously. Making sure the driver is attentive and able to take over at any time is a cornerstone of our feature development and validation, and something we continue to improve through fleet learning, customer feedback, and over-the-air ("OTA") updates. Unlike every other automaker, we are constantly improving our vehicles through free OTA firmware updates for the benefit of our customers, including updating Autopilot performance and driver engagement requirements.

In addressing your questions below, we discuss our approach to driver attentiveness and the actions we have taken to maximize it.

1. *Does Tesla exhaustively test potential methods for evading Autopilot's safety features? If so, please describe your testing methods in detail. If not, why not?*

Yes. As mentioned, through OTA capability, we continually learn from the customer fleet about their use of and interaction with Autopilot. This data helps us to make regular and rapid improvements, introduce new features, and highlight additional areas for research. At various points, we have made major upgrades to Autopilot performance and to ensure appropriate driver engagement, described below; the entire fleet is now equipped with these upgrades.

Every major feature or improvement we deploy to the fleet goes through a rigorous development, QA, and release process. After defining specifications and requirements, we develop and validate the software using industry best practices. Validation includes initial software functional development, simulation, hardware and software in-the-loop testing, system-level and regression testing, and test track, and on-road engineering testing. Multiple departments perform cross-functional reviews of the entire process. Their reviews may include demonstrations of the specifications, risk management and mitigation, development and validation results, and compliance reviews. In addition, software updates

relating to safety also require hazard analysis, risks assessments, and failure modes and effects analysis.

After validation and reviews are complete, we often release a feature to our Early Access customer program and other small stage rollouts to gauge initial feedback and identify any unforeseen challenges with implementation. We make more refinements as necessary, and only after we are fully satisfied with performance, integrity, and safety do we finally push the feature or improvement to the entire customer fleet as part of an OTA software update. In general, the cadence for release of OTA software updates is every 4 weeks. However, there are no artificial deadlines or need to rush; if a feature or improvement is not ready for deployment, we push it to a later OTA software release.

Through these efforts, we have tailored the steering wheel torque monitoring to require human interaction and to deactivate when there is too little or too much torque applied. In practical terms in most situations, this means that a limp hand on the wheel from a sleepy driver will not work, nor will the coarse hand pressure of a person with impaired motor controls, such as a drunk driver. Additionally, most dead weights, such as the unsafe defeat devices that are marketed to trick Autopilot, may be able to trick the system for a short time, but generally not for an entire trip before Autopilot disengages.

2. *Does Tesla track or otherwise monitor the online videos that drivers are posting to share tricks for circumventing Autopilot's safety features? If so, what does Tesla do with this information once discovered? If not, why not?*

Yes. While some online videos show that there are a few bad actors who are grossly abusing Autopilot, these represent a very small percentage of our customer base. We believe that many of these videos are fake and intended to capture media attention. Nonetheless, we continually monitor for and review these videos and correlate fleet data to determine whether we can eliminate actions that lead to irresponsible and unsafe driving. All of these reviews roll into the development process described in Question #1; specific examples of OTA updates that were designed to address driver inattention are discussed in Question #3.

In addition, and in response to the "Autopilot Buddy" defeat device we saw in some online videos, we joined the National Highway Traffic Safety Administration's ("NHTSA") efforts to remove this device from sale on the market. In late 2018, shortly after NHTSA sent a cease-and-desist letter to the manufacturer of Autopilot Buddy and a similar, rebranded product, we worked with Amazon, their primary marketplace, to remove these products from sale on grounds that they perpetuated dangerous driving and were likely illegal. We will continue to monitor the marketplace for similar defeat devices to make it harder for intentional bad actors to succeed.

3. *What actions is Tesla taking to upgrade the Autopilot system in Tesla cars currently on the road to address these now-known flaws? What actions is Tesla taking to improve Autopilot for future models of Tesla cars before they are put on the market?*

Every vehicle on the market can be used in a manner that is dangerous, illegal, and grossly against the recommended instructions of the manufacturer. No driver monitoring technologies available on the market today, including driver-facing camera monitoring, are immune from misuse. These driver misuses are not flaws attributable to the auto manufacturer, just as abuses of Autopilot technology is not a flaw attributable to Tesla. We take seriously our responsibility to educate our customers to help them develop good driving habits, and to discourage them from misuse.

To that end, we provide instructions on proper Autopilot use through our website, our delivery associates, who are all Tesla corporate employees, and our owner manual. Additionally, initial feature activation in the touchscreen warns customers they must pay attention; drivers receive a warning each time Autopilot is engaged; they receive repeated and escalating warnings whenever the driver does not maintain hands on the wheel with appropriate steering wheel torque; and finally, the driver experiences feature deactivation when they fail to maintain hands on the wheel with appropriate torque. In this way, we habitually train drivers to understand the limitations of Autopilot and how to use it properly.

As previously discussed, our customers are safer using Autopilot than not using Autopilot. Nevertheless, as mentioned before, we work continuously to improve Autopilot. We use OTA capability to learn from the customer fleet, make improvements, and introduce new features. This capability is not model- or model year-specific. Tesla vehicles receive the same OTA updates, whether vehicles are new or used, unless they are constrained by older hardware iterations, in which case they receive the same OTA updates to the extent capable of the hardware. The major change in hardware occurred in October 2016. Vehicles built from about October 2014 to October 2016 received the first generation of Autopilot hardware as standard equipment ("HW1"). Vehicles built after October 2016 to the present received second (and subsequent) generations of Autopilot hardware as standard equipment (HW2+). HW2+ includes more cameras and a new CPU for greater processing capacity. OTA updates are free of charge, and while customers may opt out of receiving them, over 98% choose to receive updates regularly and at no cost.

As part of our continuous improvement of Autopilot, four major OTA updates relating to driver attentiveness are worth noting. The first major update, introduced in late 2016, introduced our three-strikes reminder, wherein if you receive three escalating alerts in a short time then Autopilot will deactivate and not reactivate until the car is parked and turned off. At that time, we also introduced several hands-on-wheel detection warnings for different types of road events and shortened the time for warnings. The second major update, introduced in June 2018, increased hands-on alerts to ensure attentive driving with consistent hands on the steering wheel. The additional alerts, both visual and

auditory, escalated in frequency as vehicle speed increased, or immediately when Autopilot detected unusual or invalid lane lines, no drivable free space ahead, or a stationary object in the driving path. As a result, hands-on alerts became more frequent and forced bad actors to curb their habits and keep their hands on the wheel with more prevalence. The third and fourth major updates, introduced in May and December 2019, respectively, introduced new warnings for red lights and stop signs (both visual and auditory.) The purpose was to minimize the potential risk of red light- or stop sign-running as a result of temporary driver inattention. We will continue to make improvements to Autopilot.

4. *What safeguards beyond the Autopilot system is Tesla deploying, or considering deploying, to address the risks of AV technologies on public roads?*

Tesla's development of true AV features (SAE L3+) will follow the same iterative process described in this response. While we have not yet introduced an AV feature to the customer fleet, we continue to consider new improvements for Autopilot today. This includes, among others, continuing to enhance our hands-on detection for greater detection capacity. Moreover, our efforts over the past four years have enabled us not only to improve Autopilot today and AV features tomorrow, but also active safety capabilities yesterday. We made AEB standard on every Tesla vehicle since 2016, while the rest of the US auto industry still lags behind and will not make it standard on all models for several more years. Further, all Tesla vehicles since 2016 also have pedestrian- and bicycle-detecting AEB, lane departure warning, and emergency lane keeping, as standard equipment. We deployed these features to vehicles via an OTA software update in 2019, and they meet the toughest AEB standards currently in the world—the Euro NCAP 2019+ AEB standards. We were able to make these improvements through the capabilities of Autopilot hardware as well as learning from anonymized data collected from our customer fleet about how people use their vehicles in challenging situations in the real world.

* * *

Thank you for the opportunity to share comments with your office. We welcome continued collaboration and communication to alleviate any concerns that you or your constituents may have about Autopilot. If you have any questions, please contact my colleagues, Brooke Kintz (bkintz@tesla.com) or Hasan Nazar (hanazar@tesla.com), anytime.

Sincerely,

Alexandra N. Veitch
Senior Director, Government Relations & Policy