NHTSA's Implausible Safety Claim for Tesla's Autosteer Driver Assistance System

February 8, 2019

Report by Quality Control Systems Corp. • www.quality-control.us

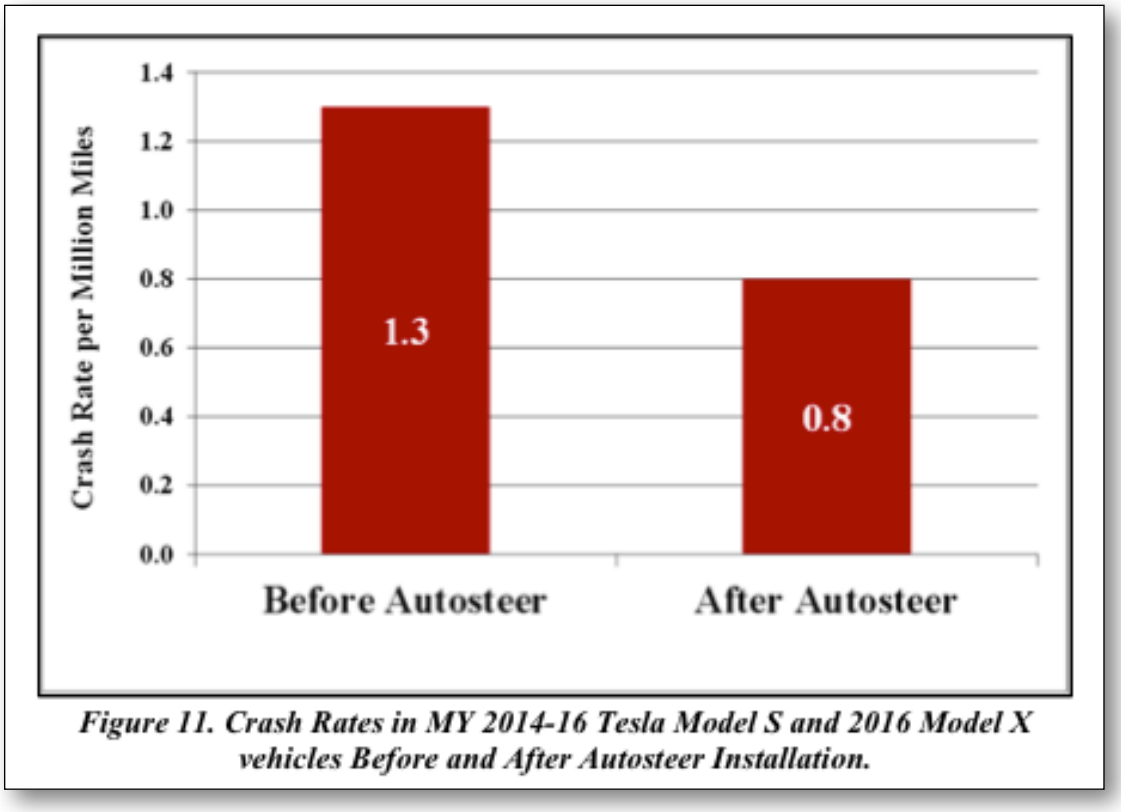

Figure 11. Crash Rates in MY 2014-16 Tesla Model S and 2016 Model X vehicles Before and After Autosteer Installation.

National Highway Traffic Safety Administration Investigation PE16-007 Docket, Report attached to the Closing Resume of Preliminary Evaluation PE16-007, Figure 11, available online at <https://static.nhtsa.gov/odi/inv/2016/INCLA-PE16007-7876.PDF>, accessed December 21, 2018.

NHTSA's Implausible Safety Claim for Tesla's Autosteer Driver Assistance System

February 8, 2019

Report by Quality Control Systems Corporation, Crownsville, Maryland

Abstract

In January 2017, the National Highway Traffic Safety Administration (NHTSA) published the remarkable claim that the airbag deployment crash rate dropped by almost 40 percent in Tesla passenger vehicles equipped with the Autopilot Technology Package following the installation of a new driver assistance system component, Autosteer. However, our replication of NHTSA's analysis of the underlying data shows that the Agency's conclusion is not well-founded.

The calculation of accurate crash rates of this type depend on reliable counts or estimates of both airbag deployment crashes as well as the mileage travelled exposing vehicles to the risk of a crash. But after obtaining the formerly secret, underlying data through a lawsuit filed under the Freedom of Information Act (FOIA) against the U.S. Department of Transportation, we discovered that the actual mileage at the time the Autosteer software was installed appears to have been reported for fewer than half the vehicles NHTSA studied. For those vehicles that do have apparently exact measurements of exposure mileage both before and after the software's installation, the change in crash rates associated with Autosteer is the opposite of that claimed by NHTSA – if these data are to be believed.

For the remainder of the dataset, NHTSA ignored exposure mileage that could not be classified as either before or after the installation of Autosteer. We show that this uncounted exposure is overwhelmingly concentrated among vehicles with the least "before Autosteer" exposure. As a consequence, the overall 40 percent reduction in the crash rates reported by NHTSA following the installation of Autosteer is an artifact of the Agency's treatment of mileage information that is actually missing in the underlying dataset.

Financial Disclosure: Quality Control Systems Corp. has no financial interest with regard to Tesla, Inc. or its Autopilot or Autosteer technologies. We have no financial interest concerning any organization, individual, or technologies associated with autonomous vehicles or advanced driver-assistance systems in competition with Tesla. We have received no financial support from any individual or entity for the pursuit of the litigation necessary to obtain the data analyzed for this report or for the writing of the report itself.

NB: The airbag deployment and mileage data on which NHTSA based its findings were apparently collected through July 8, 2016.[1] Tesla announced on October 19, 2016 substantial changes to the hardware on which Autopilot relies in vehicles produced as of that date.[2]

Acknowledgment: We are very grateful for the diligent efforts of our attorney, Mr. David L. Sobel. Without his skillful, legal representation in our Freedom of Information Act lawsuit against the U. S. Department of Transportation, this research would not have been possible.

A review of scientific research generally works best when there is a channel of communication with the authors of the original research. Unfortunately, neither the Reviewer nor the Investigator leading NHTSA's Preliminary Evaluation 16-007 of Tesla's automatic vehicle control systems responded to invitations by telephone and email to discuss our findings about the data underlying their safety claim for Tesla's Autosteer driver assistance system.

---

[1] National Highway Traffic Safety Administration Investigation PE16-007 Docket, Initial Response from Tesla (INIM-PE16007-66301), page 1, available online at <https://static.nhtsa.gov/odi/inv/2016/INLE-PE16007-66301.pdf>, accessed January 2, 2019.

[2] Tesla, Inc., The Tesla Team, "All Tesla Cars Being Produced Now Have Full Self-Driving Hardware," October 19, 2016, at <https://www.tesla.com/blog/all-tesla-cars-being-produced-now-have-full-self-driving-hardware/>, accessed January 6, 2019.

Introduction

On January 19, 2017 NHTSA published the dramatic announcement of an extraordinary reduction in the airbag deployment crash rate of MY 2014 through 2016 Tesla Model S and 2016 Tesla Model X vehicles equipped with the Autopilot Technology Package, following the installation of a new driver assistance system component, Autosteer.[3] Because of the public interest in the safety of advanced driver-assistance systems following a fatal crash in Williston, Florida on May 7, 2016 involving Autopilot,[4] NHTSA's findings immediately received wide publicity.[5]

The complete statement of the Agency's findings about Autosteer was as follows: "ODI analyzed mileage and airbag deployment data supplied by Tesla for all MY 2014 through 2016 Model S and 2016 Model X vehicles equipped with the Autopilot Technology Package, either installed in the vehicle when sold or through an OTA update, to calculate crash rates by miles travelled prior to [fn. 21] and after Autopilot installation. [fn. 22] Figure 11 shows the rates calculated by ODI for airbag deployment crashes in the subject Tesla vehicles before and after Autosteer installation. The data show that the Tesla vehicles crash rate dropped by almost 40 percent after Autosteer installation." Footnote 21 stated: "Approximately one-third of the subject vehicles accumulated mileage prior to Autopilot installation." According to footnote 22: "The crash rates are for all

---

[3] National Highway Traffic Safety Administration Investigation PE16-007 Docket, Report attached to the Closing Resume of Preliminary Evaluation PE16-007, Figure 11, available online at <https://static.nhtsa.gov/odi/inv/2016/INCLA-PE16007-7876.PDF>, accessed December 21, 2018.

[4] See National Transportation Safety Board Docket for NTSB Accident ID HWY16FH018 at: <https://dms.ntsb.gov/pubdms/search/hitlist.cfm?docketID=59989>, accessed January 28, 2019.

[5] See, for example, the same day coverage by: *The New York Times*, "Tesla's Self-Driving System Cleared in Deadly Crash"; Elon Musk (@elonmusk - Twitter), "Report highlight: 'The data show that the Tesla vehicles crash rate dropped by almost 40 percent after Autosteer installation.'"; Bloomberg, "Tesla's Autopilot Vindicated With 40% Drop in Crashes"; The Verge, "Tesla's crash rate dropped 40 percent after Autopilot was installed, Feds say"; BGR, "Report finds Tesla's Autopilot makes driving much safer."

miles travelled before and after Autopilot installation and are not limited to actual Autopilot use."[6]

Remarkably, NHTSA's announcement was not accompanied by any of the data underlying this astonishing claim. NHTSA failed even to cite the numerators and denominators of the crash rates to back up its analysis. This lack of evidence made it impossible to assess statistical confidence intervals for the reported crash rates or the statistical significance of NHTSA's finding. Judged on this basis, the Agency's findings did not meet long-established standards that would have helped the public to assess the scientific validity of NHTSA's remarkable safety claim about Autosteer.

Following the announcement that NHTSA would examine the design and performance of any automated driving systems in use at the time of the May 2016 fatal crash, Tesla insisted that: "...when used in conjunction with driver oversight, the data is unequivocal that Autopilot reduces driver workload and results in a statistically significant improvement in safety when compared to purely manual driving."[7] Yet related research based on insurance records concerning the potential safety benefit of enabling Autopilot on the Tesla Model S has had mixed results. The Highway Loss Data Institute found enabling Autopilot was associated with a (statistically significant) 13% lower frequency of collision claims, but had no statistically significant effect on other types of insurance claims, including property damage liability, bodily injury liability, claims under medical payment coverage, or personal injury protection claims.[8]

To replicate and better understand NHTSA's study, we filed a Freedom of Information Act request on February 24, 2017 for "all of the mileage and airbag deployment data supplied by Tesla analyzed by ODI to calculate the crash rates shown in Figure 11 [of the report attached to the Closing Resume of Preliminary

---

[6] National Highway Traffic Safety Administration Investigation PE16-007 Docket, Report attached to the Closing Resume of Preliminary Evaluation PE16-007, Figure 11, available online at <https://static.nhtsa.gov/odi/inv/2016/INCLA-PE16007-7876.PDF>, accessed December 21, 2018., p. 10.

[7] Tesla, Inc., The Tesla Team, "A Tragic Loss," June 30, 2016, available online at: <https://www.teslamotors.com/blog/tragic-loss>, accessed September 13, 2017.

[8] Highway Loss Data Institute, *Bulletin*, "Tesla Model S driver assistance technologies," Vol. 53, No. 4, August 7, 2018, available online at: <http://www.iihs.org/media/cb11a111-f26c-445d-a35e-afa90812bb60/gSkzrw/HLDI%20Research/Bulletins/hldi_bulletin_34.30.pdf>, accessed January 20, 2019.

Evaluation PE16-007]. In addition, we request[ed] all records related to any statistical summaries, formulas, models, adjustments, sample weights, and/or any other data or methods relied upon to calculate the crash rates shown in Figure 11."[9]

NHTSA responded by letter dated March 31, 2017, stating that "[t]he agency expects to provide a response by April 14, 2017."[10] In fact, we never heard from the Agency again until we sued the Department of Transportation on June 28, 2017 to obtain the requested data.[11] NHTSA did not even notify us that they had denied our FOIA request until July 21, 2017, well after the lawsuit was filed.[12] This decision was determined by NHTSA's judgment that compliance with our request was likely to cause Tesla "substantial competitive harm."[13]

Following the issuance of a Memorandum Opinion and Order by the Court adverse to NHTSA's position,[14] the government's attorneys released the data on November 27, 2018 – 641 days after we filed our original request and 677 days after NHTSA made its sensational claim about Autosteer.

Replication and Analysis

In the course of the FOIA lawsuit, the NHTSA investigator responsible for the crash rate calculations informed the Court that these calculations were based on data supplied by Tesla in response to an Information Request letter sent as part of

---

[9] See Complaint for Injunctive Relief, Quality Control Systems Corp. v. United States Department of Transportation at <http://quality-control.us/Quality_Control_Systems_Corp_v_USDOT.pdf>, p. 2, accessed December 21, 2017.

[10] *ibid*., p. 3.

[11] *ibid*., p. 1.

[12] See Appendix A to this report at
<http://quality-control.us/Autosteer_Report_Appendices. pdf>.

[13] *ibid*., p. 4.

[14] See Memorandum Opinion and Order, Quality Control Systems Corp. v. United States Department of Transportation, available at
<http://quality-control.us/Memorandum_Opinion_and_ Order.pdf>, accessed January 15, 2019.

NHTSA's Preliminary Evaluation 16-007.[15] This fact had not been definitively known, as NHTSA had never specifically identified the provenance of the data supporting their claim about airbag deployment crashes. And, in fact, the original Information Request to Tesla never mentions airbags.

The investigator stated that he had performed the crash rate calculations "by examining the sums of the miles driven prior to Autosteer activation, miles driven after Autosteer activation, airbag deployment events prior to Autosteer activation and airbag deployment events after Autosteer activation for all of the subject vehicles."[16] This calculation is shown in Figure A, along with the focus of our concerns about the denominators of the crash rates.

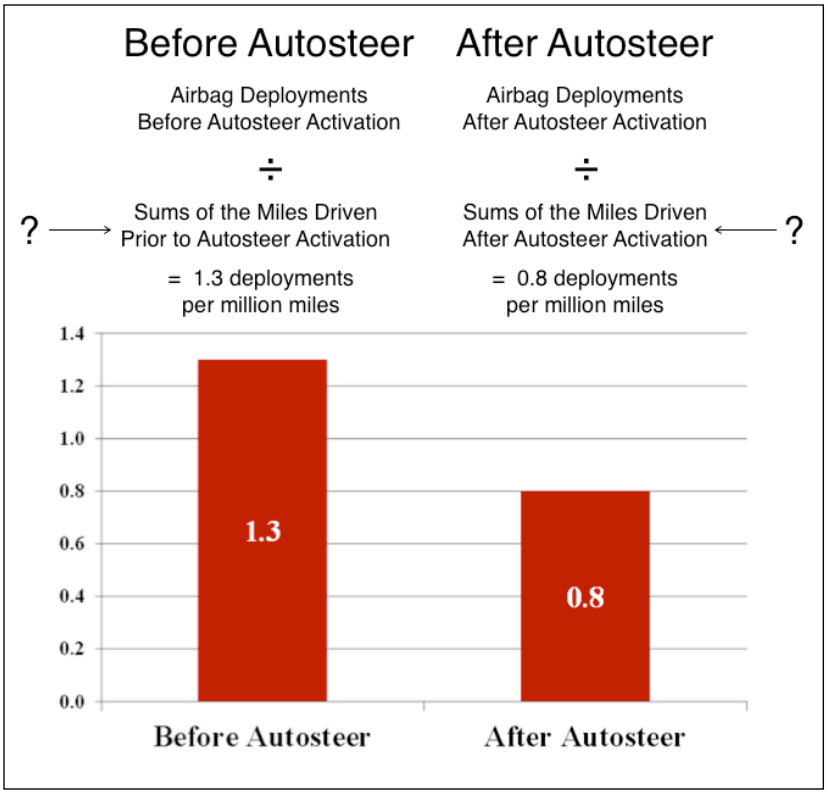

Figure A. NHTSA's Original Representation of the Method Used to Calculate Crash Risk.

[15] National Highway Traffic Safety Administration Investigation PE16-007 Docket, Information Request to Tesla (INIM-PE16007-64338), available online at <https://static.nhtsa.gov/odi/inv/2016/INIM-PE16007-64338.pdf>, accessed January 13, 2019.

[16] Quality Control Systems Corp. v. U.S. Department of Transportation (1:17-cv-01266), Defendant's Motion for Summary Judgment, Attachment 2, Declaration Exhibit A at <https://www.courtlistener.com/recap/gov.uscourts.dcd.187553/gov.uscourts.dcd.187553.10.2.pdf>, p. 4, accessed December 29, 2018.

After reviewing the Court's Memorandum Opinion and Order, NHTSA notified Tesla that the Agency intended to rescind its previous grant of confidential treatment to Tesla for the data we had requested. Through this letter, we learned for the first time that Tesla's response to NHTSA's information request for "The mileage Autosteer software was installed on the vehicle"[17] had apparently been answered instead with information about the "Previous Mileage before Autosteer Install" and the "Next Mileage after Autosteer Install."[18] (The difference between the question posed by NHTSA and the answer supplied by Tesla might be explained by the redactions in the publicly available version of Tesla's response, but it is impossible to know unless NHTSA or Tesla reveals this information.[19])

Based on this new information, we realized that replicating NHTSA's analysis might not be as straightforward as simply calculating sums of airbag deployments and exposure mileage, based on the mileage at Autosteer activation/installation, as NHTSA had represented. We also recognized that NHTSA's summarization of "miles driven prior to Autosteer activation [and] miles driven after Autosteer activation" might not actually include all of the miles driven before or after Autosteer activation. If the denominators of the calculated crash rates shown in Figure A did not count all of the exposure mileage, either before or after Autosteer, the crash rates calculated by NHTSA would be statistically biased. In order to assess the validity of NHTSA's conclusion about the reduction in crash risk following Autosteer installation, it was important to establish whether this bias actually existed and whether it might affect the "before Autosteer" crash rate more than the crash rate "after Autosteer." Figure A shows how our questions about the data concern NHTSA's crash rate calculations.

---

[17] National Highway Traffic Safety Administration Investigation PE16-007 Docket, Initial Response from Tesla (INLE-PE16007-66301), Question 1l, available online at <https://static.nhtsa.gov/odi/inv/2016/INLE-PE16007-66301.pdf>, accessed January 2, 2019.

[18] See Appendix B at <http://quality-control.us/Autosteer_Report_Appendices.pdf> for a copy of this letter.

[19] National Highway Traffic Safety Administration Investigation PE16-007 Docket, Initial Response from Tesla (INIM-PE16007-66301), available online at <https://static.nhtsa.gov/odi/inv/2016/INLE-PE16007-66301.pdf>, pp. 2-3, accessed January 2, 2019.

Once we received the data,[20] we attempted to replicate NHTSA's summaries[21] of airbag deployments as well as mileage exposure before and after Autosteer installation. We expected that the "Miles before Autosteer" exposure measure calculated by NHTSA would be equal to "Previous Mileage before Autosteer Install" reported by Tesla when mileage had been accumulated before Autosteer was installed. It was also our expectation that exposure "Miles after Autosteer" calculated by NHTSA would be equal to "The mileage of the vehicle at the last data retrieval" minus "Next Mileage after Autosteer Install" reported by Tesla when these data were not unknown, unreported, or otherwise missing for all of the vehicles studied.

In those cases where the "Previous Mileage before Autosteer Install" exactly equals "Next Mileage after Autosteer Install," it can be inferred that Autosteer was actually installed at the reported mileage in each of these two fields. Figure 1 illustrates the calculation of exposure mileage in the cohort of vehicles where the odometer data at the time of installation is known, based on this inference. (Note that the mileage of the vehicle at the last data retrieval must also have been reported for the vehicles in Figure 1.) This method of calculation of exposure mileage was applicable only to 5,714 vehicles of the total 43,781 vehicles studied, 13 percent.[22]

---

[20] Available at:
<http://www.safetyresearch.net/Library/2018-11-26%2520Redacted%2520PE16_007_PRODUCTION%2520DATA_jlq_working_file_10Jan2017%2520.xlsx>.

[21] These summaries were found in the Excel workbook <2018-11-26 Redacted PE16_007_ PRODUCTION DATA_jlq_working_file_10Jan2017 .xlsx>, worksheet "PE16_007_PRODUCTION DATA", Table 1, at row 43,784. The summary of Column AX (Miles before Autosteer) is given in this worksheet as 64,788,137. The summary of Column AY (Miles after Autosteer) is stated to be 235,880,377. We counted 86 airbag deployments "before Autosteer" in this Table and 192 deployments "after Autosteer." (We note, however, discrepancies between these summaries and another summary also presented by NHTSA in worksheet "by Mileage", Table 1: Before Autosteer 64,661,869; After Autosteer 234,612,736. The number of "Airbag Events" "Before AutoSteer" is shown to be 84; "After Autosteer", 189. These discrepancies are unexplained.)

[22] JMP statistical software, produced by the SAS Institute, was used for the purposes of data management and for the analyses presented in this report, Version 5.1.1. Cary, NC: SAS Institute Inc.; 2004.

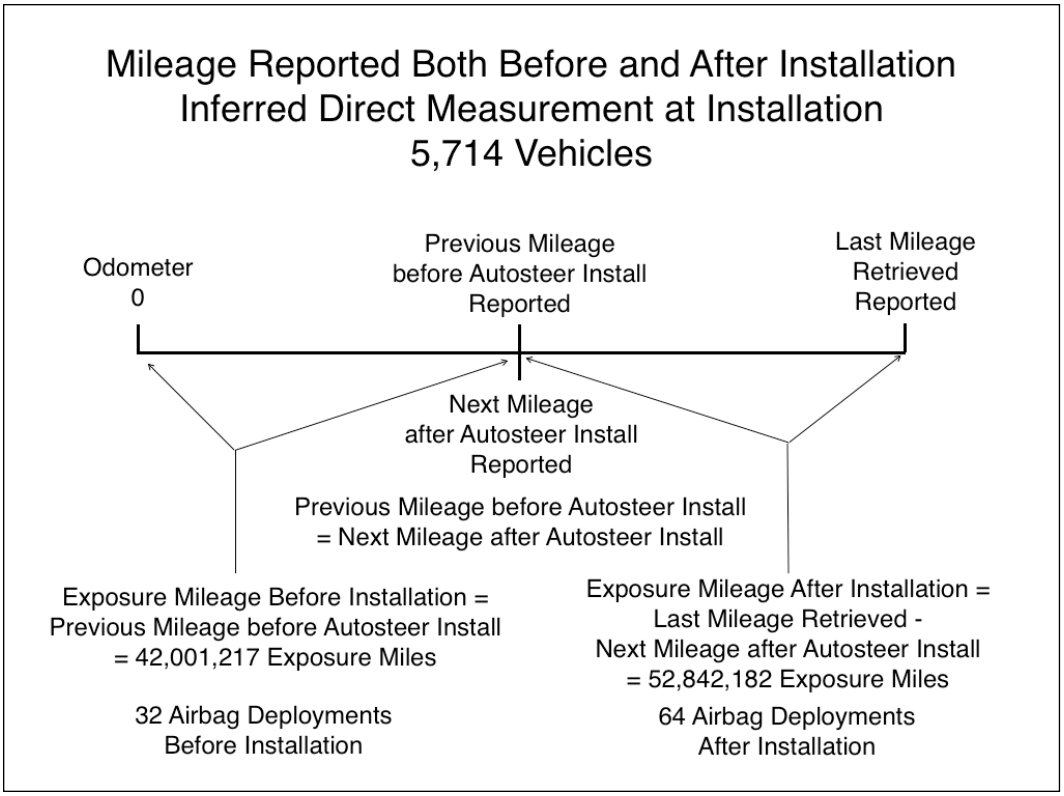

Figure 1. Exposure Mileage Calculation; Inferred, Direct Measurement at Installation.

Based on the data for crash rate numerators and denominators shown in Figure 1, the resulting calculations reveal a 59 percent increase in the airbag deployment crash rate from 0.76 per million miles of travel to 1.21 per million miles of travel following the installation of Autosteer. As explained below, this result is particularly important because it is the only vehicle cohort in the study with complete information for both before and after Autosteer crash rate calculations. Before and after comparisons of the resulting crash rates are unbiased by missing data for exposure mileage because there are no missing data in this subset of the data. This finding is the just the opposite of that claimed by NHTSA for the larger set of vehicles they studied.

We used logistic regression to measure the practical and statistical significance of Autosteer to this apparent difference in crash rates. Because the data do not record the mileage at which an airbag deployed, we employed a method that transformed the dataset of 5,714 vehicles into two equal sized segments, "before" and "after" Autosteer. (In the transformed dataset of 11,428 observations, 3 cases have missing data where the last mileage retrieved is reported to equal the "Next mileage after Autosteer installation"). Each observation in the new dataset can be understood as a segment of exposure miles that either did or did not result in

an airbag deployment crash. Each observation contained the independent variables, total "Exposure mileage" for the segment, "Autosteer installed" (equals 1 if so, zero otherwise), as well as the dependent variable of "Airbag deployed."

The model estimated from these specific data helps to answer the question concerning NHTSA's safety claim about Autosteer, "Is the installation of Autosteer associated with a decreased risk of an airbag deployment crash, controlling for exposure mileage?" The answer is "No."

Table 1 demonstrates that Autosteer is actually associated with an increase in the odds ratio of airbag deployment by more than a factor of 2.4 (95% Confidence Interval: 1.57 - 3.8), when exposure mileage is taken into account. See Table 1.

Table 1. Results of logistic regression model estimation.

Whole Model Test

Model	-LogLikelihood	DF	ChiSquare	Prob>ChiSq
Difference	12.88407	2	25.76815	<.0001
Full	541.51572			
Reduced	554.39980			

RSquare (U)	0.0232
Observations (or Sum Wgts)	11425

Converged by Objective

Lack Of Fit

Source	DF	-LogLikelihood	ChiSquare
Lack Of Fit	9499	501.99278	1003.986
Saturated	9501	39.52294	Prob>ChiSq
Fitted	2	541.51572	1.0000

Parameter Estimates

Term	Estimate	Std Error	ChiSquare	Prob>ChiSq	Odds Ratio	Lower 95%	Upper 95%	Odds Lower	Odds Upper
Intercept	-4.637396	0.2191908	447.61	<.0001	.	-5.0847923	-4.2240864	.	.
Exposure mileage	-0.0000878	0.0000247	12.61	0.0004	0.00633291	-0.0001381	-0.0000412	0.0003497	0.09301571
Autosteer installed	0.88548641	0.2244738	15.56	<.0001	2.42416325	0.45388517	1.33675921	1.57441719	3.80668681

For log odds of 1-Yes/2-No

Effect Wald Tests

Source	Nparm	DF	Wald ChiSquare	Prob>ChiSq
Exposure mileage	1	1	12.6124729	0.0004
Autosteer installed	1	1	15.5608136	0.0001

We note also the unexpected result that, while the estimated coefficient for "Exposure mileage" is statistically significant, it has the "wrong" sign. Our surprise is based on the simple expectation that an airbag deployment crash becomes more likely with increasing exposure mileage – whether or not Autosteer is installed. For these data, however, the estimated model showed that accumulating exposure mileage lowered the odds of an airbag deployment. (See more on this topic in the Discussion section.)

Figure 2 shows the calculation of exposure mileage for the cohort of vehicles apparently assumed by NHTSA to be sold or leased with Autosteer already installed. NHTSA appears to have applied this interpretation of the data to cases where the mileage of the vehicle at the last data retrieval was reported but the "Previous Mileage before Autosteer Install" and the "Next Mileage after Autosteer Install" are both unreported. This method of calculation of exposure mileage was utilized for 14,791 vehicles of the total 43,781 vehicles studied (34 percent of all cases – but see the exception for "Data row 1" in Table 3 below).

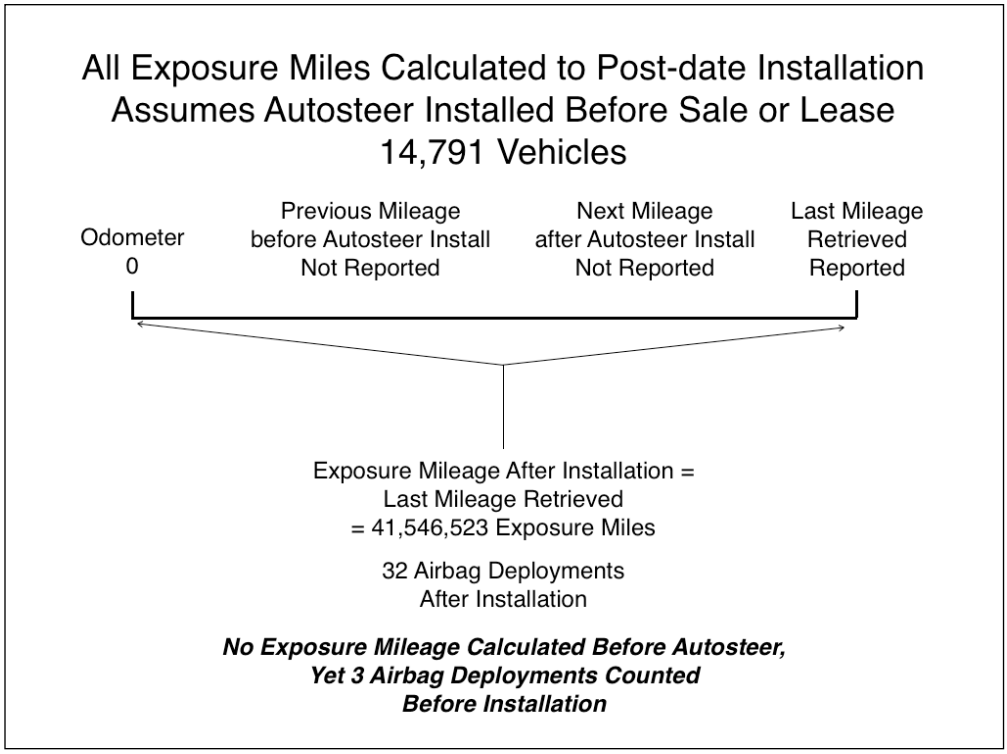

Figure 2. Exposure Mileage Calculation,
Assumes Autosteer Installed Before Sale or Lease.

Figure 2 demonstrates a minor, but obvious, sign of trouble with NHTSA's interpretation of these data. All of the mileage accumulated in these vehicles' lifetime is assigned to the "after Autosteer" group (because the "Previous Mileage before Autosteer Install" and the "Next Mileage after Autosteer Install" are both unreported), even though three airbag deployments have still been counted before Autosteer is installed. Because of this problem, the "before Autosteer" crash rate is inflated to a very small degree because NHTSA has counted no exposure mileage for these three cases. However, this minor problem points to a assumption that is more problematic for calculations of mileage exposure. That is, simply because the data are missing for the "Previous Mileage before Autosteer Install"

and the "Next Mileage after Autosteer Install," NHTSA's method of calculation assumes that all of the exposure mileage must belong to the "after Autosteer" category. The three airbag deployments without any exposure mileage in the "before Autosteer" category show this is not the case. As the following figures illustrate, it is likely that this problem is more serious in other vehicle cohorts.

Figure 3 demonstrates NHTSA's calculations for the vehicle cohort with all of the basic mileage data reported but where Previous Mileage before Autosteer Install is less than Next Mileage after Autosteer Install. Apparently Autosteer has been installed in these vehicles sometime after lease or sale but the exact mileage at the time of installation was not recorded or preserved. We refer to the accumulated mileage between Previous Mileage before Autosteer Install and Next Mileage after Autosteer Install as the "exposure mileage gap."

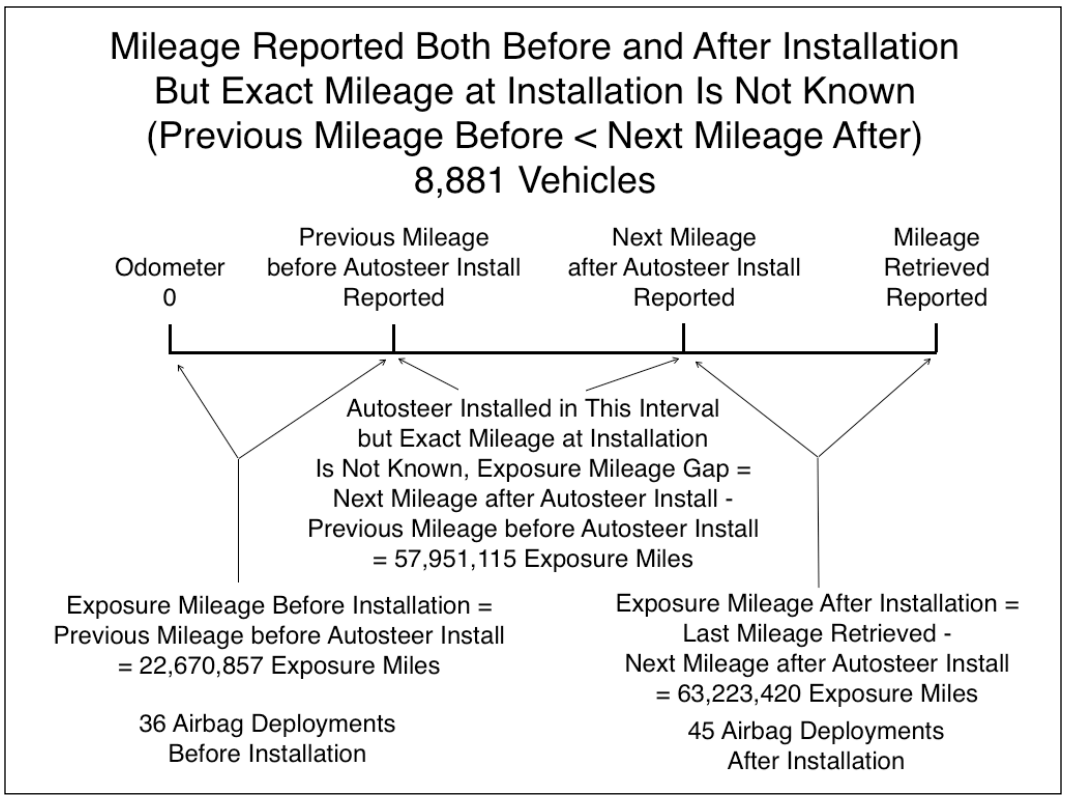

Figure 3. Exposure Mileage Calculation, Exact Mileage at Autosteer Installation Not Reported in the Data.

8,881 vehicles in NHTSA's data (20.3 percent of the total) have an exposure mileage gap, some portion of which properly belongs to the before Autosteer exposure mileage and the balance to the after Autosteer mileage exposure. For these cases, NHTSA's summaries of the exposure mileage both before and after instal-

lation result in estimates of the true exposure that are statistically biased downward – resulting in crash rates that are somewhat too large.

With the airbag deployment numerators fixed at 36 deployments before Autosteer and 45 deployments after installation, it is especially problematic for the crash rate comparisons that the total exposure mileage gap is 2.6 times as large as the entire accumulated exposure mileage before Autosteer installation. At the same time, this gap is about eight percent smaller than the total the exposure mileage NHTSA calculates after Autosteer installation.

Theoretically, if the actual mileage at installation within the exposure mileage gap were randomly distributed at some point in the gap interval for each vehicle, the "before Autosteer" crash rate would be reduced to a greater degree than the resulting reduction in the "after Autosteer" crash rate. In practice, we do not understand the nature of the mechanism or process that generates the mileage exposure gap but it is not random. Figures 3A and 3B show that a calculation of airbag deployment crash risk in this cohort is differentially biased when the problem of the exposure mileage gap is ignored.

For the cohort of 8,881 vehicles in Figure 3, we arranged the data in sort order by "Miles before Autosteer" (lowest to highest). With the data in this order, we then calculated the cumulative "before Autosteer" exposure mileage and cumulative the exposure mileage gap up to and including each observation in the dataset and plotted both. The results are shown in Figure 3A.

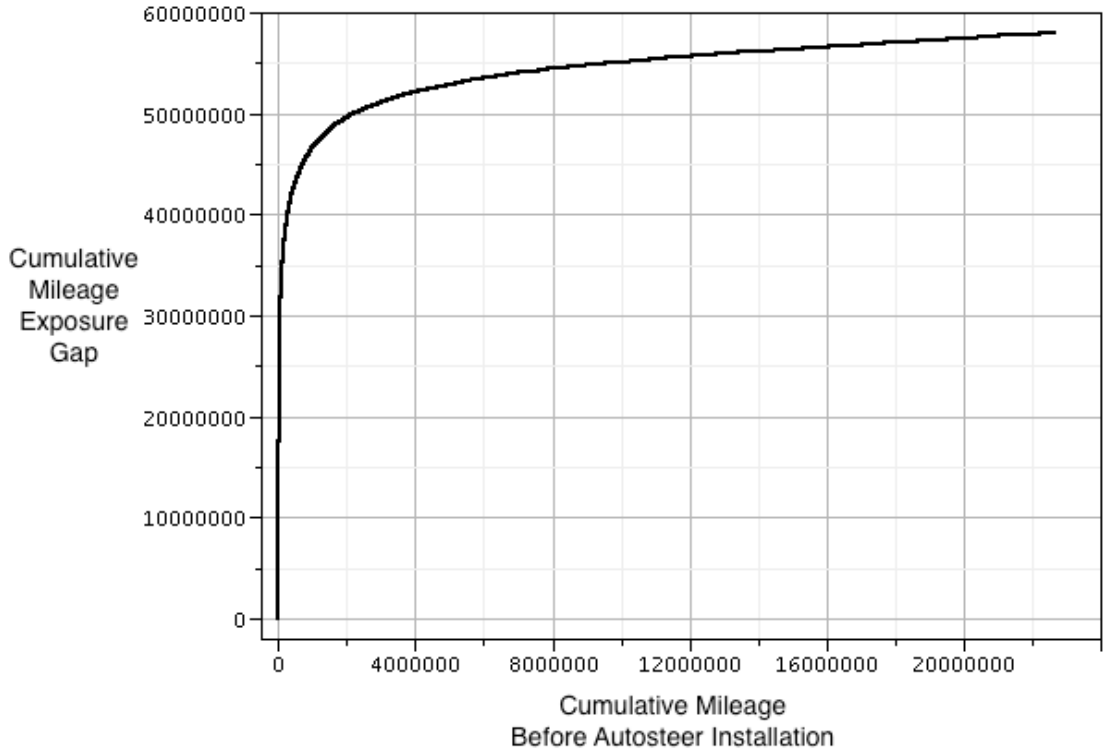

Figure 3A. Cumulative Mileage Exposure Gap by Cumulative Mileage Before Autosteer Installation for the 8,881 Vehicles in Figure 3.

Figure 3B shows the analogous plot for the cumulative exposure gap by the cumulative mileage after Autosteer installation.

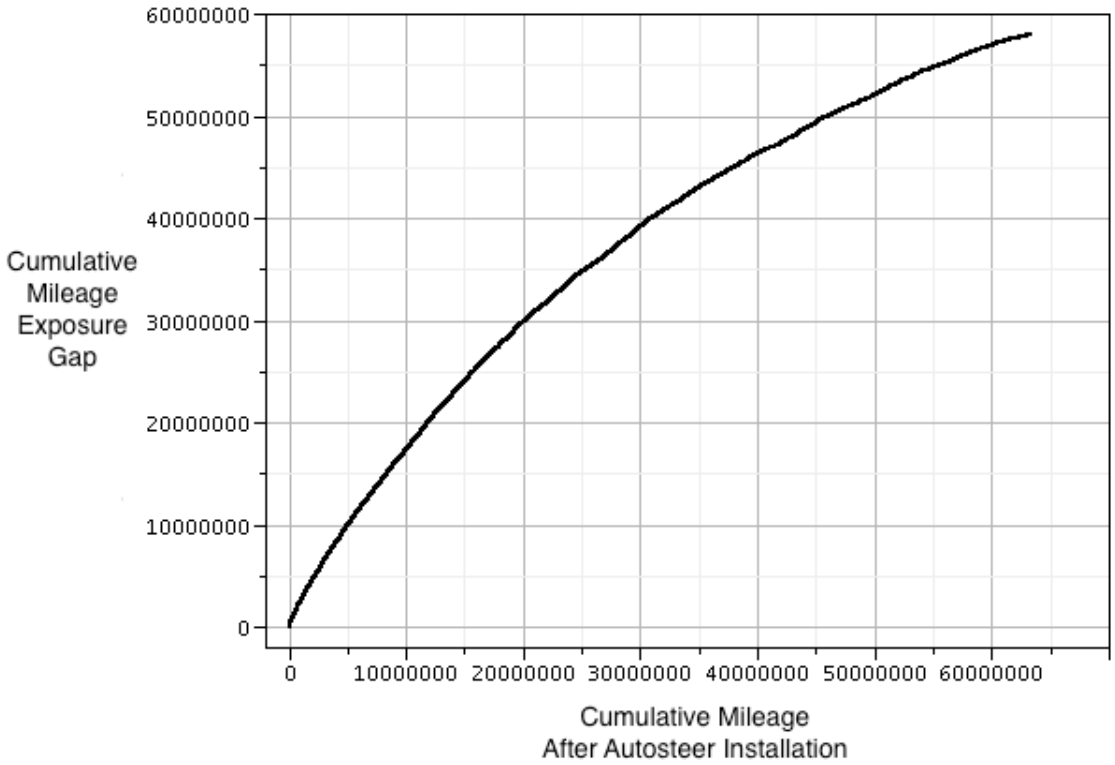

Figure 3B. Cumulative Mileage Exposure Gap by Cumulative Mileage After Autosteer Installation for the 8,881 Vehicles in Figure 3.

Comparing Figures 3A and 3B, it is clear that the mileage exposure gap is overwhelmingly concentrated among vehicles with the least "before Autosteer" exposure. This indicates a considerable undercount of actual exposure mileage before Autosteer installation in contrast to the exposure mileage after Autosteer installation for this cohort. The result is a differential bias that inflates the calculated crash risk "before Autosteer" to a greater degree than the calculated risk "after Autosteer."

This problem would be less important were it not true – as Figure 2 shows it is – that NHTSA counts airbag deployments for vehicles without any corresponding exposure mileage at all in its crash rate comparisons. Figure 4 below demonstrates how extreme this problem becomes when the only upper bound on the exposure mileage gap is the Next Mileage after Autosteer Installation.

Figure 4 shows how NHTSA accounted for exposure mileage in the class of vehicles where the exact mileage at Autosteer installation is unreported but where the "Next Mileage after Autosteer Installation" is present in the data. There are 14,260 vehicles in this category, nearly one-third of the entire study population.

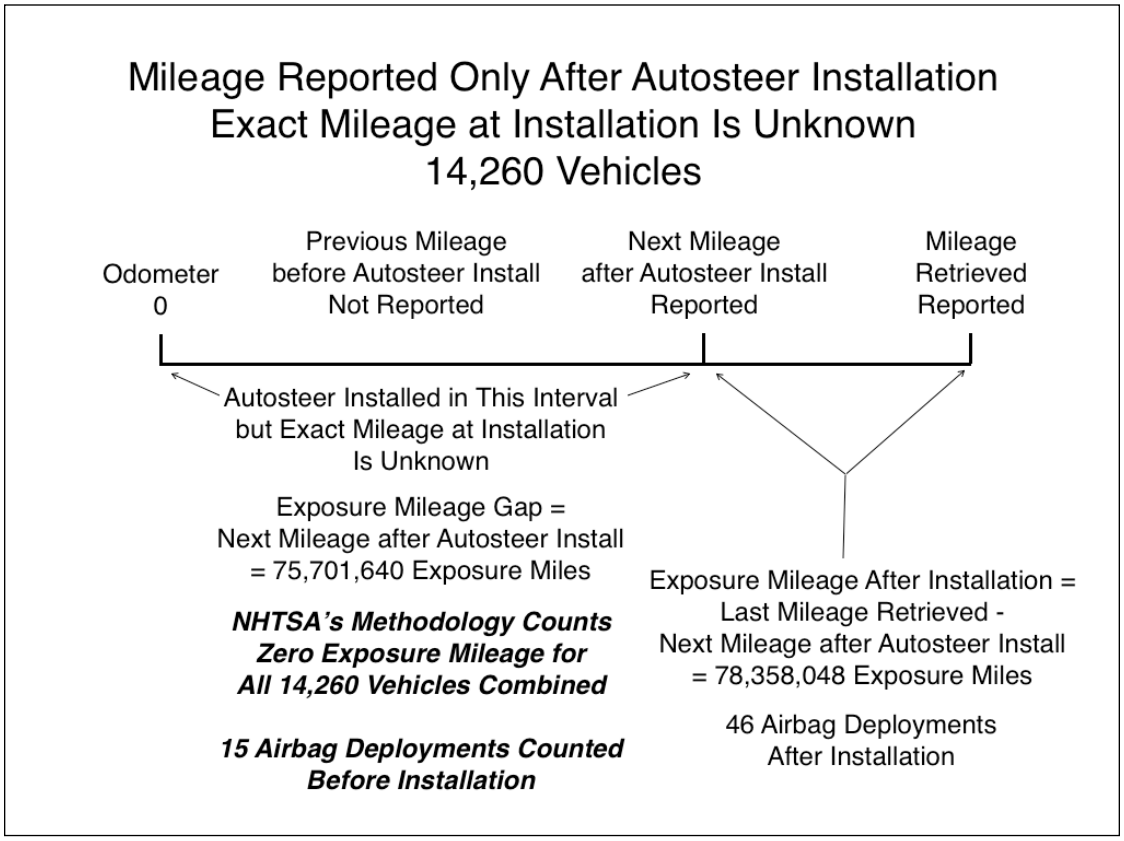

Figure 4. Exposure Mileage Calculation, Exact Mileage at Autosteer Installation Is Unknown, but Next Mileage after Autosteer Installation Is Reported.

Figure 4 shows that 78,358,048 exposure miles were included by NHTSA in its "after Autosteer" installation crash rate calculation for this group of 14,260 vehicles. This is a fractional undercount of the "after Autosteer" exposure mileage since the actual Autosteer installation took place at some prior point.

In contrast, there are no exposure miles at all included by NHTSA in its "before Autosteer" crash rate calculation for this group, even though the 15 airbag deployments NHTSA counts "Before Autosteer" represent 17 percent of the total 86 reported in the data. The exclusion of this entire group from the calculation of the overall exposure mileage before Autosteer installation, while counting all of the airbag deployment crashes in the numerator, biases the calculation in a way

that differentially inflates the "before Autosteer" crash rate compared to the inflation of the "after Autosteer" crash rate.

Table 2. Relative distributions of airbag deployments and exposure mileage in Figures 1 through 4.

Figure	Vehicle Count	Deployments before Autosteer	Miles before Autosteer	Deployments after Autosteer	Miles after Autosteer
Figure 1	5,714 (13%)	32 (37%)	42,001,217 (65%)	64 (33%)	52,842,182 (22%)
Figure 2	14,791 (34%)	3 (3%)	0 (0%)	32 (17%)	41,546,523 (18%)
Figure 3	8,881 (20%)	36 (42%)	22,670,857 (35%)	45 (23%)	63,223,420 (27%)
Figure 4	14,260 (33%)	15 (17%)	0 (0%)	46 (24%)	78,358,048 (33%)
Totals (including vehicles in Table 3)	43,781	86	64,788,137	192	235,880,377

Table 2 documents that NHTSA counted 18 airbag deployments "before Autosteer" (20% of the total) but counted no corresponding exposure mileage "before Autosteer" for these vehicles. This points to the same deficiency in NHTSA's approach to summarizing exposure mileage that has already been discussed.

Table 3 below shows how airbag deployments and exposure mileage were summarized by NHTSA for the balance of the study population not included in Figures 1 through 4 above.

Table 3. NHTSA's accounting of airbag deployments and exposure mileage for the balance of vehicles not shown in Figures 1 through 4.

Information Available	Vehicle Count	Deployments before Autosteer	Miles before Autosteer	Deployments after Autosteer	Miles after Autosteer
Last Mileage Retrieved, Reported; Previous Mileage before Autosteer Install, Reported; Next Mileage after Autosteer Install, Not Reported	34	0	54,239	0	Missing (Zero)
Last Mileage Retrieved, Not Reported; Previous Mileage before Autosteer Install, Reported; Next Mileage after Autosteer Install, Reported	7	0	60,918	2	-65,267
Last Mileage Retrieved, Not Reported; Previous Mileage before Autosteer Install, Reported; Next Mileage after Autosteer Install, Not Reported	3	0	906	0	Missing (Zero)

Information Available	Vehicle Count	Deploy-ments before Autosteer	Miles before Autosteer	Deploy-ments after Autosteer	Miles after Autosteer
Last Mileage Re-trieved, Not Re-ported; Previous Mileage before Autosteer Install, Not Reported; Next Mileage after Autosteer Install, Reported	4	0	Missing (Zero)	0	-25,195
Last Mileage Re-trieved, Not Re-ported; Previous Mileage before Autosteer Install, Not Reported; Next Mileage after Autosteer Install, Not Reported	85	0	Missing (Zero)	2	Missing (Zero)
Exception: Data row 1	1	0	Missing (Zero)	0	Missing (Zero)
Exception: Data row 24363	1	0	Missing (Zero)	1	666

Some minor issues in NHTSA's interpretation of Tesla's data are apparent in Table 3 which do not affect NHTSA's overall before and after Autosteer crash rate conclusion in important ways. For example, Table 3 demonstrates that NHTSA has actually subtracted exposure mileage in the "after Autosteer" category for eleven vehicles. This occurs when the last mileage retrieved is not reported but the "Next mileage after installation" is reported. In these cases, NHTSA subtracted the reported mileage from zero.

The next to last row in Table 3 documents the case for the first vehicle in the dataset: the last mileage retrieved was reported, but no other mileages were. We don't know why NHTSA would not have applied the method of calculation shown in Figure 2 for all other vehicles with mileage information of this type.

The last row in Table 3 documents that, for this particular case, NHTSA's assigned the exposure mileage after Autosteer installation to equal the "Next Mileage after Autosteer Install" measurement although the data for the last mileage retrieved was not reported.

Discussion

NHTSA's finding that the airbag deployment crash rate for Teslas dropped following the installation of Autosteer would have been even more dramatic if more of the Autosteer installation mileage data had been missing. The Agency's treatment of missing or unreported mileage data in its calculations of exposure mileage as though the mileage were non-existent is not justifiable. This problem affects more than half the dataset (see Figures 3 and 4).

Although the Agency stated that "Approximately one-third of the subject vehicles accumulated mileage prior to Autopilot installation,"[23] our analysis reveals that 28,904 of the total 43,781 vehicles accumulated exposure mileage before Autosteer installation, based on non-missing values of either "Previous Mileage before Autosteer Install" or "Next Mileage after Autosteer Install." That is twice as many vehicles as NHTSA reports. Unfortunately, NHTSA counts exposure mileage in only 14,639 of these 28,904 vehicles – about half – and disregards the rest. This results in the inflation of the overall "before Autosteer" airbag de-

[23] National Highway Traffic Safety Administration Investigation PE16-007 Docket, Report attached to the Closing Resume of Preliminary Evaluation PE16-007, available online at <https://static.nhtsa.gov/odi/inv/2016/INCLA-PE16007-7876.PDF>, p. 10, fn. 21, accessed December 21, 2018.

ployment crash rate reported by NHTSA, but to a degree that can't be known with certainty.

For that reason, the results of the logistic regression model estimated in Table 1, showing a significant rise in the airbag deployment crash rate following the installation of Autosteer, are particularly important. Because the mileage information about this segment of the population is complete and does not depend on statistically biased approximations, it might be supposed that these data are the most accurate to be found in the entire dataset. However, we are not in a position to judge the trustworthiness of the data per se.

The Table 1 model shows that the likelihood of airbag deployment decreases with increased exposure mileage (both before and after the installation of Autosteer). This may be an indication that even the data for this cohort are not sufficiently accurate to assess the effect of the installation of Autosteer on crash risk with confidence. However, an alternative explanation may be that in the "before Autosteer" time period an airbag deployment results in taking the vehicle off the road for some period of time to effect repairs, thus lowering the exposure mileage. This constraint is also present in the "after Autosteer" period where there is also the possibility that the deployment crash may total the vehicle.

These considerations also demonstrate why the subset of vehicles actually equipped with Autopilot and in service both before and after Autosteer is not necessarily the best choice on which to stake a risk analysis about Autosteer. Since Autopilot was eventually enabled on all of these vehicles, it is uncertain how the Autosteer comparison is affected by restricting the analysis to those vehicles that have survived potential crash exposure long enough to have been equipped with the Autopilot option. An unknown number of vehicles that were scrapped before Autopilot became available are ignored, potentially biassing the results. (We note here, however, the remarkable fact that 83 of the 86 vehicles with airbag deployment crashes before Autosteer installation were still in service after these crashes and accumulated exposure mileage after Autosteer was installed, according to these data.)

It has been previously noted that the before and after Autosteer comparison of crash rates is potentially confounded by other vehicle features, such as Blind Spot Warning, Automatic Emergency Braking, and Forward Collision

Warning.[24,25] All of these systems were enabled prior to Autopilot/Autosteer being made available and were standard equipment. In contrast, some of the mileage exposure "before Autosteer" does not benefit from these safety features. We note also that by asking Tesla for information only about vehicles equipped with Autopilot, NHTSA missed the opportunity to study changes over time in the crash rates of the same Model S and Model X vehicles in the same model years that were never equipped with Autopilot – a very natural control group.

The choice of the research design is as important as the choice of the data. If the mileage at which the airbag deployment crashes in this study occurred was available to Tesla, it should have been possible to employ a hazard (or survival) analysis. Such an analysis would have been a more informative and efficient study design than the design NHTSA actually used. But it is still unclear how Tesla would have known about all or nearly all vehicles involved in airbag deployment crashes occurring in the field. This is particularly true for high-speed crashes in rural areas that could make post-crash, over-the-air transmission of airbag deployment and odometer data difficult or impossible.

NHTSA's study never addresses the possibility that some airbag deployment crashes might be missing from its analysis. A recent statement by Tesla that it had introduced a "completely new telemetry stream ... to gather the most critical fleet-wide statistics from the exact moment a crash-related event is detected by our system"[26] suggests that achieving complete, detailed crash coverage may have been less straightforward at the time the data were collected for NHTSA's investigation.

---

[24] Timothy B. Lee, *Ars Technica*, "Sorry Elon Musk, there's no clear evidence Autopilot saves lives, The feds just threw Tesla under the bus on Autopilot safety," available at: <https://arstechnica.com/cars/2018/05/sorry-elon-musk-theres-no-clear-evidence-autopilot-saves-lives/>, accessed May 5, 2018.

[25] Highway Loss Data Institute, *Bulletin*, "Tesla Model S driver assistance technologies," Vol. 53, No. 4, August 7, 2018, available online at: <http://www.iihs.org/media/cb11a111-f26c-445d-a35e-afa90812bb60/gSkzrw/HLDI%20Research/Bulletins/hldi_bulletin_34.30.pdf>, accessed January 20, 2019.

[26] Tesla, Inc., The Tesla Team, "Q3 2018 Vehicle Safety Report," available at: <https://www.tesla.com/blog/q3-2018-vehicle-safety-report>, accessed January 8, 2019.

Conclusion

The importance of this research goes well beyond the specific issues addressed in our statistical analyses. The larger question is whether the field experience of autonomous vehicles and advanced driver-assistance systems will be fairly and transparently assessed by the public officials charged with insuring the public's safety while this technology is "beta-tested" on public roads. The litigation record in our case[27,28] documents the resources which both NHTSA and Tesla were willing to commit to prevent public scrutiny of this taxpayer funded research, based on a fear of competitive harm to Tesla.

The actual methodology applied by NHTSA to Tesla's data was not adequately explained by the Agency when its claim about Autosteer was announced. The implausibility of the dramatic change in crash rates following the installation of Autosteer would not be understood even now if we had not pursued lengthly litigation against the Agency to obtain the data that were supposed to support its remarkable safety claim.

Before a judgment was reached in our case, the government changed its position and released the data we had requested.[29] Tesla did not publicly object. Yet the long delay in complying with our FOIA request has served to turn NHTSA's tenuous safety claims about Autosteer into an established Agency policy to tolerate crashes involving Autopilot. So far, NHTSA has shown no interest in reopening their investigation of Tesla's automatic vehicle control systems on the basis of additional, serious crashes associated with the use of Autopilot.

Until October 19, 2018, the record of our FOIA lawsuit reflects NHTSA's unwavering focus on the potential competitive harm to Tesla that the Agency believed would ensue from complying with our FOIA request. The potential physical harm to consumers was ignored. In fact, public policies preventing the replica-

---

[27] Available at: <http://foiaproject.org/case_detail/?title=on&style=foia&case_id=30938>, accessed January 8, 2019.

[28] Available at: <https://www.courtlistener.com/docket/6081569/quality-control-systems-corp-v-us-department-of-transportation/>, accessed January 8, 2019.

[29] Indeed, by May 2018, NHTSA began to describe its own assessment of crash rates associated with the installation of Autosteer as "cursory." See, Timothy B. Lee, *Ars Technica*, "Sorry Elon Musk, there's no clear evidence Autopilot saves lives, The feds just threw Tesla under the bus on Autopilot safety," available at: <https://arstechnica.com/cars/2018/05/sorry-elon-musk-theres-no-clear-evidence-autopilot-saves-lives/>, accessed May 5, 2018.

tion of publicly funded auto safety research by independent researchers are harmful both to manufacturers and to consumers because such policies are anti-scientific. Efforts to hide the crash record will impede progress in achieving whatever safety benefits advanced driver-assistance systems might ultimately bring. The recorded crash data retrieved from the fatal crash in Williston, Florida, shows that extensively detailed data, including odometer and vehicle speed data, could be made available to advance independent research in this field.[30]

A very substantial fraction of the public simply doesn't trust autonomous driving technologies.[31] Given the scarcity of scientifically reliable, publicly available data about the safety of these systems, why should they? It ought to be very concerning to the proponents of advanced driver-assistance systems that an international, comprehensive, open, and trustworthy surveillance system for casualties and property damage associated with the use of these technologies on public roadways is nowhere in sight.

---

[30] Jordan Golson, *The Verge*, "Read the Florida Highway Patrol's full investigation into the fatal Tesla crash: A meticulous dissection of last summer's Autopilot fatality" available at: <https://www.theverge.com/2017/2/1/14458662/tesla-autopilot-crash-accident-florida-fatal-highway-patrol-report>, "The data logs from the Model S begin on page 60 [of part 3 of the Florida Highway Patrol report], showing the kind of information that investigators can retrieve from that model of car," accessed February 5, 2019.

[31] American Automobile Association, "Fact Sheet, Vehicle Technology Survey - Phase IIIB", May 2018, available at: <https://publicaffairsresources.aaa.biz/download/10980/>, accessed January 10, 2019.