Potential Crash Rate Benchmarks for Automated Vehicles

Noah Goodall

Virginia Transportation Research Council, USA, noah.goodall@vdot.virginia.gov

Pre-print version. Final version in *Transportation Research Record* (2021), https://doi.org/10.1177/03611981211009878

Abstract
Most automobile manufacturers and several technology companies are testing automated vehicles on public roads. While automation of the driving task is expected to reduce crashes, there is no consensus regarding how safe an automated vehicle must be before it can be deployed. An automated vehicle should be at least as safe as the average driver, but national crash rates include drunk and distracted driving, meaning that an automated vehicle that crashes at the average rate is somewhere between drunk and sober. In this paper, automated vehicle safety benchmarks are explored from three perspectives. First, crash rates from naturalistic driving studies are used to determine the crash risk of the model (i.e., sober, rested, attentive, cautious) driver. Second, stated preference surveys in the literature are reviewed to estimate the public's acceptable automated vehicle risk. Third, crash, injury, and fatality rates from other transportation modes are compared as baseline safety levels. A range of potential safety targets is presented as a guide for policymakers, regulators, and automated vehicle developers to assist in evaluating the safety of automated driving technologies for public use.

Disclaimer: *The views and opinions expressed in this article are those of the author and do not necessarily reflect the official policies or positions of any agency of the Commonwealth of Virginia.*

INTRODUCTION

The task of driving has become increasingly automated in recent decades, particularly since the start of large-scale public road testing in 2010 (*1*). Vehicles with computer-control of combined lateral and longitudinal movement over a range of environments are referred to as autonomous or automated vehicles. Most vehicle manufacturers and several technology companies are developing automated vehicles, and these developers have conducted 6.5 million miles of supervised autonomous testing on public roads in California as of November 2019.

The safety of these vehicles remains an open question. As advanced automated vehicles are not yet sold to the public, safety assessments have relied mostly on mandated reporting of autonomous testing mileage and crashes to the California Department of Motor Vehicles (*2*). These studies have found that when equivalent crash definitions are compared, automated vehicles crash at lower rates than the general public (*3*, *4*). Automated vehicle testing in California, however, requires the use of a trained safety driver to monitor the vehicle and take control as needed. These drivers are trained to anticipate collisions and unexpected vehicle behavior, and to take control back from the automated driving system. As these safety drivers made 159,840 unplanned takeovers over 3,669,472 miles since between September 2014 and November 2018 (*5*), it is probable that crash rates from testing with trained safety drivers would not translate to general public drivers who may be distracted, tired, and unfamiliar with the technology.

There are as yet no defined safety requirements that automated vehicles must meet prior to being tested or sold to the public. Applications to test new automation technologies are generally handled on a case-by-case basis at the state level. Some automation features would likely be regulated at the national level, similar to how conventional vehicles in the United States are regulated by the National Highway Traffic Safety Administration (NHTSA), which is obligated to order a manufacturers to remedy defects that present an "unreasonable risk" (*6*) of crashes, death, or injury. In these cases, manufacturers would also be responsible for remedying defects, and dealers would be prohibited from selling vehicles with known defects (*7*). Vehicles can be defective even if they comply with NHTSA's performance standards if they present "unreasonable risk of accidents" (*8*). In California, the state's Department of Motor Vehicles (DMV) regulates automated vehicle testing within state boundaries, and requires that vehicles be insured, bonded, and in most cases monitored by a licensed driver directly employed by the AV developer (*2*). Although no regulations in the U.S. at the state or federal level prohibit AVs that exceed predetermined crash, injury, or fatality rates, a known safety benchmark could be used to begin to define "unreasonable" risks that specific AVs should not present. Private sector AV developers may also find crash rate targets useful when evaluating and communicating the safety of their automated driving systems.

Were AV safety benchmarks to be expressed as crash, injury, and fatality rates, what might those benchmarks be? Several benchmarks have been proposed in the literature and by commentators. One answer is that AV must at least be as safe as the average driver (*9*). While this seems reasonable, it fails to take into account the average driver's crashes that involve unacceptable and illegal driving behavior such as drunk driving and speeding. An automated vehicle that is safe as the average driver is, statistically speaking, somewhere between drunk and sober driving in terms of safety. Society does seem to have accepted these risks to some extent—speed governors could eliminate speeding (*10*) and compulsory alcohol detection systems would restrict drunk driving (*11*), yet neither is mandatory for new vehicles. But there is a clear difference between prohibition and aggressive enforcement, and laws against speeding and drunk driving are still enforced. Introducing an automated vehicle with a crash rate equivalent to an occasional law-breaking human driver seems unlikely and irresponsible.

Others have proposed higher standards for automated vehicle safety. Sparrow and Howard argue that AVs that are safer than average *with respect to other road users* should be made mandatory (*12*). Kalra and Paddock, in calculating the number of miles traveled to establish statistical significance of safety improvements, used a 20% reduction in average fatal crash rate as an example (*13*). Commentators have recommended safety improvements by factors varying from 2 to 100, while others use the baseline of a mature, skilled driver (*14*). A widely referenced number is NHTSA's estimate that driver error is a contributing factor in 94% of crashes (*15*), implying that automated vehicles free from human error might reduce the crash rate by a similar amount. To achieve a 94% crash reduction, however, human driving would have to be eliminated, as automated vehicles would still be struck by at-fault human drivers. Beyond the impracticality of banning human driving in the near future, others have pointed out how human error rate and automated vehicle crash reductions are not directly comparable (*16*).

Responsibility for establishing the safety of automated vehicles is shared among manufacturers and regulators. These entities may choose to define what constitutes a safe autonomous driving system and, more importantly, what does not. Specific safety targets allow AV developers to communicate their safety record in context and with transparency. Clear definitions of acceptable crash rates tied to achievable benchmarks with wide public acceptance allows transparent regulation and validation of automated vehicles.

The purpose of this paper is to present a range of safety benchmarks for automated vehicles. This is accomplished through a calculation of crash rates based on various driver risk factors from naturalistic driving, analysis of stated preference surveys in the literature, and a comparison of crash, injury, and fatality crash rates of other modes of transportation.

CRASH RATES CONSIDERING DRIVER RISK FACTORS

Crash rates are simply the rate at which vehicles crash, often expressed in terms of number of crashes per distance traveled. Rates can be broken down by collision angle, crash severity, speed, and a range of other factors recorded in a police report. A significant portion of crashes are due to driver error or impairment. These factors, such as excessive speeding, fatigue, and distraction should not be present in an AV, nor should they be exhibited by for-hire vehicles such as taxis and buses. Without these impairments, the baseline crash rate for an AV might be closer to an error and impairment-free model driver than the average driver.

To isolate the effect of these causal factors on crash rates, the number of crashes and distance traveled must be recorded for driving when both exposed to the factor and not exposed to the factor. While some of the factors are captured in NHTSA crash records, e.g. blood alcohol content, others are either captured with low confidence (fatigue (*17*) and distraction (*18*)) or not captured at all (emotional driving). No factors are collected for non-crash driving, making any claims about the relative safety of driving with risk factors impossible—one cannot claim that distracted driving is less safe without knowing the total number of distracted miles driven both with and without crashing.

There are other shortcomings with conventional crash databases. While crash reporting is detailed and complete for fatal crashes in many countries, injury and property damage crashes may have limited data and might never be recorded in databases. In the United States, the National Highway Traffic Safety Administration estimated that 15% of injury crashes and 24% of property damage-only crashes are never reported to police (*19*), while an additional 9% of injury crashes and 24% of property damage-only crashes are reported but not entered into databases (*20*). These estimates of missing crash records are based on telephone surveys and rely on respondents recalling crashes they were involved in over the prior year (*19*). Respondents might not recall minor crashes, further weakening the completeness of crash records. To accurately capture crash rates, both crashes and mileages would need to be recorded.

Naturalistic Driving Data

Naturalistic driving studies, where drivers are passively recorded over a period of time, allow the collection of a range of risk factors over both crash and non-crash driving. Naturalistic driving studies have the added benefit that crashes can be automatically recorded via the vehicles kinematic sensors without relying on police-reporting. Further, vehicle mileages can be recorded with much greater precision and accuracy than FHWA miles-driven estimates.

The Second Strategic Highway Research Program Naturalistic Driving Study (SHRP 2 NDS) completed data collection in 2015 and remains the largest NDS of its kind, collecting video, kinematic, and audio data from over 3,500 drivers over three years and 30 million miles (*21*). Crashes, near-crashes, other safety critical events, and baseline non-crash events were analyzed by data reductionists and classified based on observed driver error and impairment. Because this dataset captured risk factors for both crashes and baseline non-crash driving, it can be used to calculate the effect of individual risk factors on crash rate.

Calculation of Driver Risk Odds Ratios

Dingus et al. (*22*) presents the odds ratios and baseline prevalence for various risk factors based on model driving. The odds ratio is defined as crash rate when exposed to a risk factor divided by the crash risk when not exposed.

$$\text{O.R.} = \frac{\text{crash rate under risk factor}}{\text{crash rate under model driving}} \tag{1}$$

To determine the crash rate for each risk factor, the NDS needed to collect not only crashes but also the prevalence of each risk factor under non-crash driving. Due to limits in in-vehicle data storage, drivers were not recorded continuously. Prevalence of driver behaviors under non-crash driving was instead estimated from data recorded for short time intervals while the vehicle was moving. Over the course of the study, 19,732 non-crash intervals were selected based on a two-staged stratified random sampling method and subjected to the same data reduction by trained analysts as crash scenarios. From these samples, a baseline of exposure data can be developed. To determine the odds ratio from baseline data, the crashes per baseline intervals when the risk factor was present is the numerator, while the crashes per baseline intervals when under model driving is the denominator.

$$\text{O.R.} = \frac{\frac{c_{risk}}{b_{risk}}}{\frac{c_{model}}{b_{model}}} = \frac{c_{risk} b_{model}}{b_{risk} c_{model}} \tag{2}$$

In this equation, c is the number of crashes and b is the baseline mileage under the *risk* and *model* driving behaviors, respectively. Dingus et al. (*22*) provides values for odds ratios and percentage of baseline events exhibiting each risk factor. To conduct any comparative analysis, the crashes and baseline events for model driving are needed. Crashes are provided (12.3% of 905 crashes occurred under model driving), and baseline prevalence can be calculated by using a risk factor with a known odds ratio, crash percentage, and baseline percentage. As an example, driver distraction is used below, which represented 68.3% of injury and property damage crashes and 51.93% of baseline samples.

$$\text{O.R.} = \frac{\frac{c_{risk}}{b_{risk}}}{\frac{c_{model}}{b_{model}}} \tag{3}$$

$$2 = \frac{0.683/0.5193}{\frac{0.123}{b_{model}}} \tag{4}$$

$$b_{model} = 0.187 \tag{5}$$

The odds ratio for all driving compared to model driving can then be calculated as:

$$\text{O.R.} = \frac{\frac{c_{risk}}{b_{risk}}}{\frac{c_{model}}{b_{model}}} = \frac{\frac{1}{1}}{\frac{0.123}{0.187}} \approx 1.5 \tag{6}$$

The all-driving injury and property damage crash rate for the naturalistic data is provided by Blanco et al. (*4*) as 20.2 crashes per million vehicle-miles traveled. (Although the SHRP 2 NDS purposely over-sampled drivers under age 25 and over age 65 (*22*), the all-driving crash rate of 20.2 crashes per million vehicle-miles used when calculating risk factor crash rates is age-weighted to represent the U.S. national driving population (*4*)). The crash rate for model driving can be obtained using the general driving odds ratio and reformulating equation (1).

$$\text{crash rate under model driving} = \frac{\text{crash rate under risk factor}}{\text{O.R.}} = \frac{20.2}{1.5} = 13.5\ \text{crashes/MVMT} \tag{7}$$

In this equation, MVMT refers to million vehicle-miles traveled. For the purposes of this paper, the crash rate under a risk factor is less useful than the crash rate when not under a risk factor, i.e., when the driver is not distracted, fatigued, etc. Crashes and baseline driving often have multiple risk factors occurring simultaneously. To calculate the odds ratio, the crashes and baseline samples where the risk factor does not occur must be isolated.

$$\text{O.R.}_{inverse} = \frac{\frac{1 - c_{risk}}{1 - b_{risk}}}{\frac{c_{model}}{b_{model}}} \tag{8}$$

For some risk factors in Dingus et al. (*22*), c_{risk} is not available. It can be obtained and substituted into equation 8 using:

$$c_{risk} = \frac{c_{model} b_{risk} \text{O.R.}}{b_{model}} \tag{9}$$

The risk factors and associated odds ratios and crash rates are shown in Table 1. These crash rates are significantly higher than 208 per MVMT reported by NHTSA for 2018 (*23*). This difference is because SHRP 2 NDS includes low-speed no-damage crashes, while NHTSA records includes only report police-reportable damage, injury, and fatal crashes.

While the crash rate for all drivers is 2020 per hundred million miles, the crash rate is reduced slightly for sober and rested driving and significantly for cautious and attentive driving. A driver exhibiting all of these qualities would be considered a model driver with a crash rate of 1347 per 100 million miles, a reduction of 33% compared to all driving.

Table 1 Crash Rates and Odds Ratios from SHRP 2 Naturalistic Driving Study

Driver Status	Crashes per 100 Million Vehicle-Miles	Odds Ratio to Model Driving
All Driving	2020	1.5
Sober	2010	1.5
Rested	2007	1.5
Non-impaired	1954	1.5
Cautious	1479	1.1
Attentive	1350	1.0
Model Driving	1347	1.0
Child in rear seat	673	0.5

STATED PREFERENCE SURVEYS

Another way to determine potential safety benchmarks is to ask users what crash or fatality rates they would find acceptable. This is referred to as stated preference, a widely used method to establish acceptable risk levels (*24*). Liu et al. (*25*) surveyed 499 residents of Tianjin, China about acceptable fatality rates in terms of miles and population for either AVs or HDVs. Even though no respondent was asked about both AVs and HDVs, answers showed a clear preference for lower fatality rates for AVs. Based on the developed logarithmic model, half of respondents felt that a fatality rate of 0.20 per 100 million miles (3.2e-9 km) was acceptable for HDVs, but for AVs, the rate should be 0.028 per 100 million miles (4.5e-10 km), a reduction of 86%. The authors compared these rates to actual fatality rates per population globally but, to maintain consistency, the U.S. 2018 figure of 1.13 fatalities per 100 million miles is used here for comparison (*23*). The AV fatal crash rate that 50 percent of respondents found acceptable represents a 98% reduction in crashes from current rates. At the same time, respondents preferred that human drivers crash 85% less than today's drivers. This suggests that survey participants might accept an AV that is only 87% safer than today's drivers. Liu et al. defines three risk levels based on survey findings: *unacceptable risk* where crash rates are higher than those of the average human driver, *tolerable risk* when AVs are 75–80% safer than human drivers, and *broadly acceptable* when AVs reduce crash rates by 99% compared to human drivers. No definition is provided for AVs that are between 0% and 75% safer than human drivers, but for now this is defined as *barely tolerable risk.*

In the survey instructions (*25*), respondents were informed that US rates are 321 traffic crashes, 124 injuries, and 1.7 deaths per billion vehicle-kilometers, but the authors do not provide a source for these figures. In fact, these rates are approximately 25% of NHTSA's 2018 U.S. rates when converted to the same base units (*23*). Using Liu et al.'s rates as actual rates, 50 percent of respondents required AVs to reduce current crash rates by 90%. The same percentage of respondents preferred a human-driven crash rate that is 27% lower than current rates. This may have introduced some bias in Liu et al.'s results, as respondents were prompted with human-driven fatality rates significantly lower than actual rates. Additionally, different cultural attitudes towards acceptable risk (*26*) may limit the applicability of these findings to the United States.

Nees (*27*) performed a similar survey on AV risk acceptance. Nees noted that respondents may have had difficulty answering questions about acceptable crash given that these values are quite small and people have limitations when interpreting small numbers. Nees attempted to solve this problem by asking drivers to provide safety criteria as a function not of fatalities per vehicle-miles traveled but as a percentile of all drivers' fatality rates. In a survey of 504 participants via Amazon Mechanical Turk, respondents first ranked their own driving. Eighty-one percent of participants rated their driving as safer than the median driver. This was in agreement with previous studies such as Svenson (*28*) where 88% of U.S. drivers rated themselves as safer than the average driver. The median driver considered themselves at the 74th (73–76, 95% CI) percentile of safe drivers. When asked at what percentile an AV should be before being allowed on public streets, the median participant preferred an AV that was 8.7 (6.9–11.1, 95% CI) percentile points higher than themselves, with 66% preferring an AV safer than themselves.

The results are difficult to interpret for two reasons. First, most participants desired an AV in the 95th to 99th percentile. This may suggest not that they want an AV that was 12–16 percentile points safer than themselves necessarily, but rather an AV that is as safe as the best human driver regardless of the participant's own perceived safety. Second, there are few examples in the literature of percentile crash rates of drivers, and no public data sets from which to develop a risk profile. Most population crash studies use department of motor vehicle crash records (*29*). These studies do not control for exposure through mileage, so the "safest" drivers with the fewest crashes may not be driving at all. Other studies have estimate exposure from driver surveys, but these may suffer from respondents' failure to accurately recall past mileages (*29–32*). The SHRP 2 NDS provides individual records of crashes and mileages, but the dataset contains only 1,148 crashes for 3,546 participants, again leaving a large percentile range with perfect safety records (17). Near crashes, where vehicles made sudden movements or hard decelerations, could be used as safety surrogates (26), but this might weaken the accuracy of the findings. Many insurance companies now use GPS sensors on vehicles to track mileage. These data could be merged with crash records to create individual crash rate records, but these datasets are generally proprietary and nonpublic. If a percentile crash rate of a general population of drivers could be created, either from insurance data, NDS data, or some other source, then a desired safety percentile could be translated into a specific crash rate. This is beyond the scope of this article but may be an area for further research.

Even if individual risk profiles could be created, crashes are still rare, and many drivers will never crash in their lifetimes. Any crash risk distribution from individuals' data would have a significant percentile range with zero crashes. An AV attempting to achieve a 99th percentile safety rating might never be permitted to crash, an impossible task.

Finally, participants generally overestimated their safety in relation to the other drivers. Previous studies have shown that more than half of drivers believe themselves safer than both the average driver (*28*) and drivers within their peer group (*33*). It is not clear whether this overestimation has been factored into their stated risk acceptance. Consider drivers at the 40th percentile, who think they drive at the 70th percentile, and who want AVs to be at the 90th percentile. Do these drivers want vehicles that are at the 90th percentile for safety, or merely vehicles that are 20 percentile points above themselves, i.e., at the 60th percentile? This is a difficult question to answer from surveys and may instead require a revealed preference approach.

CRASH RATES FROM OTHER TRANSPORTATION MODES

When defining thresholds for acceptable AV crash rates, the safety of other modes of transportation may serve as useful benchmarks. Liu et al., for example, mentions rail and commercial aviation as two examples of transportation modes with broadly acceptable safety rates (*25*). In this section, crash, injury, and fatality of various modes of transportation in the United States between the years of 2008–2017 are compared. Counts of crashes, injuries, fatalities, and mileages were mostly obtained from the U.S. Department of Transportation's Bureau of Transportation Statistics (*34*).

Safety rates were compared across several exposure data. The standard crash risk measure is events per distance, with distance often expressed in terms of vehicle-miles traveled (VMT). This metric can bias against low-capacity modes such as passenger vehicles which have more crashes per mile but with fewer passengers are exposed to each crash. An alternative metric is explored in this effort, person-miles traveled (PMT), defined as the number of people traveling a given distance or vehicle-miles multiplied by average vehicle occupancy. When person-miles data was not available, it was calculated from average vehicle occupancy, often obtained from the U.S. National Household Transportation Survey (*35*).

While exposure is often expressed as a unit of time, this may bias against slower moving vehicles. Buses, for example, has higher fatality rates than commercial air travel per mile, but lower rates per hour due to their slower speeds. Alternative metrics, vehicle-hours traveled (VHT) and person-hours traveled (PHT) are used here. Vehicle-based metrics are used for crash rate, as the number of occupants is not affected. Person-based metrics are used for occupant injury and occupant fatality rates, where the number of passengers in a vehicle is inversely correlated to individual risk.

Table 2 shows crash, injury, and fatality rates for several transportation modes. Commercial air travel has the best safety rate across several metrics, with the exception of occupant fatalities per person-hour traveled where bus rates are lower. Bus and passenger rail have low injury and fatality rates, lower than walking and cycling and within range of elevators and escalators. Motorcycles have the worst safety record across all metrics.

Table 2 Crash, Injury, and Fatality Rates by Transportation Mode

Mode[1]	Per Unit Distance			Per Unit Time		
	Crashes / 100 M VMT	Occupant Injuries / 100 M PMT	Occupant Fatalities / 100 M PMT	Crashes / 100 M VHT	Occupant Injuries / 100 M PHT	Occupant Fatalities / 100 M PHT
Passenger cars[2]	213.9	29.9	0.28	5,946	831.7	7.9
Motorcycles[2]	546.3	375.9	20.27	13,385	9,210.5	496.6
Large trucks[3]	130.8	9.4	0.24	5,885	421.2	10.5
Buses[4]	408.4	5.1	0.015	4,655	57.8	0.18
Commercial air[2]	0.37	0.003	0.001	161	1.5	0.52
Pedestrians[6]	-	208.6	14.7	-	688.4	48.5
Cyclists[6]	-	599	8.65	-	3,893	56.2
Passenger rail[7]	-	6.7	0.034	-	155	0.8
Elevators[8]	-	993	0.74	-	1,523	1.1
Escalator[9]	-	1944	0.50	-	1957	0.51

[1] All data 2008-2017 from (*34*) unless otherwise noted. Speed and occupancy were used to convert miles to hours and vehicles to persons.
[2] Speeds and occupancy from (*35*).
[3] Trucks are defined as >10,000 lbs. Speeds average of 36 mi/hr for recreational vehicles (*35*) and 55 mi/hr Interstate Highway System speeds (*36*). Occupancy from (VM-1).
[4] Speeds from (*35*). Occupancy from (*37*).
[5] Distance and speeds from (*35*).
[6] Speeds and occupancy calculated from (*38*).
[7] 1997–2010. Speed calculated as 4.5 floors (45 feet) in 20 seconds from (*39*). All other data from (*40*, *41*).
[8] 1997–2010. Speed calculated as 0.45 m/s (*42*) and assumes average elevator distance of 10 vertical feet at a 30 degree incline for 20 feet of slope distance. All other data from (*40, 41*).

Of the modes in Table 2, buses appear to be the most applicable benchmark for automated vehicle safety. Buses used trained drivers with specialized licensing, and buses are inspected more regularly than passenger vehicles. Although buses crash at higher rates per mile than passenger vehicles, this may be due to their lower average speeds of 11.4 mi/hr to passenger vehicles' 27.8 mi/hr. Buses also travel in dense urban areas which have higher crash rates than comparative rural/suburban roads and freeways. Buses experience fewer crashes per hour, and fewer occupant injuries and fatalities across all metrics.

In evaluating AV safety benchmarks, industry and regulators must be careful to consider not only the safety rate with respect to passengers, but also the safety rate with respect to those outside the vehicle. Table 3 shows the injury and fatal crash rate for those outside the vehicle, either in other vehicle types or bicyclists and pedestrians. Non-occupants are injured in bus crashes at nearly double the rate per mile as for passenger cars. Other metrics show buses and passenger cars with similar safety metrics. Although bus occupants are at substantially less risk of death and injury than car passengers, this difference is erased for non-occupants.

Table 3 Non-occupant Injury and Crash Rates

Mode	Non-occupant injuries per 100M VMT	Non-occupant injuries per 100M VHT	Non-occupant fatalities per 100M VMT	Non-occupant fatalities per 100M VHT
Passenger cars	36.4	1013.0	0.8	20.9
Large trucks	27.2	1225.6	1.2	52.9
Buses	70.4	803.1	1.6	17.8

CONCLUSIONS AND DISCUSSION

While there is little consensus on crash rate targets for automated vehicles, some example benchmarks can be collected from other sources. Several potential benchmarks were identified in this study. Data from the SHRP 2 Naturalistic Driving Study suggests that model drivers—i.e., sober, rested, attentive, and cautious—crash 33% less than the average driver. An AV that seeks to emulate not a highly skilled human but merely one driving legally should crash at no more than this lower model rate. Surveys from Liu et al. suggest that drivers prefer an automated vehicle to crash 80% less (tolerable risk) or 99% less (broadly acceptable risk) than the average human driver (*25*). These figures are supported by findings in Nees (*27*). Others in the literature have used 20% lower crash rates as an example of acceptable AV crash risk (*13*). Crash, injury, and fatality rates from other modes provide examples of accepted risks for public transportation with stricter regulations of driver training and vehicle maintenance and inspections. Of all modes, buses provide the best comparison with AVs given their shared travel environment and similar levels of regulatory oversight. Table 4 lists a range of potential safety benchmarks for automated vehicles identified in this study. Benchmarks are expressed in terms of crashes, police-reportable (i.e., property damage or more severe) crashes, occupant and non-occupant injuries, and occupant and non-occupant fatalities. Total counts are divided by the most relevant exposure data for the metric, either vehicle-miles traveled, vehicle-hours traveled, person-miles traveled, person-hours, or more than one exposure metric.

Table 4 Potential Crash, Injury, and Fatality Rate Safety Benchmarks for Automated Vehicles

Exposure Metric	Safety Metric	Passenger cars	20% safer (*13*)	Model driving, 33% safer	Tolerable, 80% safer (*25*)	Broadly acceptable, 99% safer (*25*)	Buses for comparison
100 M VMT	All Crashes	2020	1616	1347	404	20	-
	PR Crashes	214	171	137	43	2.1	408
	Non-occupant injuries	36	29	24	7.3	0.36	70
	Non-occupant fatalities	0.8	0.6	0.5	0.2	0.008	1.6
100 M VHT	PR Crashes	5946	4757	3964	1189	59	4656
	Non-occupant injuries	1013	810	675	203	10	803
	Non-occupant fatalities	21	17	14	4.2	0.21	18
100 M PMT	Occupant Injuries	30	24	20	6.0	0.3	4.9
	Occupant Fatalities	0.28	0.23	0.19	0.057	0.003	0.015
100 M PHT	Occupant Injuries	832	665	554	166	8.3	56
	Occupant Fatalities	7.9	6.3	5.3	1.6	0.079	0.18

Metric Selection

AV developers and regulators should focus on crash metrics initially. As Kalra and Paddock have shown, fatal crashes are exceptionally rare, and any AV developer that wishes to establish with statistical significance that their vehicles are safer via on-road testing would need to accumulate an impractical number of fatality-free miles (*13*). Demonstrating that an AV has a failure rate equivalent to the national average with 95% confidence would require a fleet of 100 vehicles driving continuously for 12.5 years. The amount of required testing is substantially reduced when establishing performance for the all-crash rate. One could establish that an AV crashes at the same rate as model driving with 95% confidence with only 222,400 miles of crash-free driving. Waymo vehicles travel this distance in autonomous mode every two months in California alone (*5*).

Limitations of a Single Metric

This study investigated crash risk as a safety target for automated vehicles, defined as crashes of all severity levels per distance traveled. This metric is a useful starting point for industry and regulators who may be beginning to define safety for AVs using transparent benchmarks, but by no means should it be the end. Relying on a single benchmark is ill-advised. Automated vehicle safety improvements should be equitable, and benefits to one road user class should not come at the expense of another. An AV that kills 1,000 fewer car occupants but 100 additional pedestrians may not be acceptable, even though 900 net lives are saved. AVs should improve the safety of everyone who interacts with them, and these benefits should be distributed fairly. Crash rates at different severity levels should also be evaluated: fewer low-risk crashes do not excuse more severe crashes, and likewise fewer fatalities may not justify a two-fold increase in severe debilitating injuries. The ethics of risk is a mature field that can address some of these issues (*43*, *44*).

Similarly, AVs should focus on vehicle-miles and vehicle-hours as exposure data to avoid the temptation to put riders in the back seat and artificially inflate person-based exposure. As far as between vehicle-miles and vehicle-hours, both metrics should be used in combination. Using vehicle-hours avoids a bias towards desolate freeway driving which accumulates miles with little risk, while vehicle-miles avoids the unlikely but conceivable scenario where vehicles idle in parking lots, accumulating hours. A metric such as *vehicle-moving-hours* or *vehicle-miles-hours,* might serve as alternative exposure metrics.

Finally, different vehicle types may require their own metrics and benchmarks. In one example, several manufacturers operate low-speed automated shuttles the drive at speeds no greater that 15 mi/hr on streets with speed limits no greater than 25 mi/hr (*45*). At these speeds, vehicles will experience few fatal crashes per mile, and even fewer per hour when compared to passenger vehicles. Because of their low weight, however, shuttles may have fewer non-occupant fatalities per hour when compared to buses. Automated shuttles may require safety targets based on bus crash rates but customized for shuttles' unique operating environments and characteristics. Automated trucks operating at (relatively dangerous) high speeds on (relatively safe) controlled access highways may also require customized safety benchmarks.

Future research should investigate alternative metrics to measure safety. While this study presented metrics in terms of crash, injury, and fatality rates, others might focus on predictability, lateral and longitudinal accelerations, near-misses (*46*), inverse time-to-collision (*47*), or other crash surrogate measures. These alternative metrics might provide a more robust picture of driving risk without requiring crashes as samples. Researchers should also apply findings from risk ethics to automated vehicle crash risk. A few questions that require answers include under what circumstances an increase in crash risk to one road user group can be permitted at the expense of another group, how might AVs distribute risk among occupants and non-occupants, and how non-crash driving situations such as lateral positioning within a lane can affect permissible crash risk to others (*48*, *49*).

ACKNOWLEDGEMENTS
This work was sponsored by the Virginia Department of Transportation. The views and opinions expressed in this article are those of the author and do not necessarily reflect the official policies or positions of any agency of the Commonwealth of Virginia.

AUTHOR CONTRIBUTION STATEMENT
N.J. Goodall was sole author and is responsible for the entirety of this paper.

REFERENCES

1. Beiker, S. History and Status of Automated Driving in the United States. In *Road Vehicle Automation* (G. Meyer and S. Beiker, eds.), Springer International Publishing, Cham, 2014, pp. 61–70.
2. California Department of Motor Vehicles. Title 13, Division 1, Chapter 1, Article 3.7 – Testing of Autonomous Vehicles. https://www.dmv.ca.gov/portal/uploads/2020/06/Adopted-Regulatory-Text-2019.pdf. Accessed May 15, 2019.
3. Teoh, E. R., and D. G. Kidd. Rage against the Machine? Google's Self-Driving Cars versus Human Drivers. *Journal of Safety Research*, Vol. 63, 2017, pp. 57–60. https://doi.org/10.1016/j.jsr.2017.08.008.
4. Blanco, M., J. Atwood, S. Russell, T. Trimble, J. McClafferty, and M. Perez. *Automated Vehicle Crash Rate Comparison Using Naturalistic Data*. Virginia Tech Transportation Institute, 2016.
5. Boggs, A. M., R. Arvin, and A. J. Khattak. Exploring the Who, What, When, Where, and Why of Automated Vehicle Disengagements. *Accident Analysis & Prevention*, Vol. 136, 2020, p. 105406. https://doi.org/10.1016/j.aap.2019.105406.
6. 49 U.S.C § 30102(a)(8) (2012).
7. Smith, B. W. Automated Vehicles Are Probably Legal in the United States. *Texas A&M Law Review*, Vol. 1, 2014, pp. 411–521. https://doi.org/10.2139/ssrn.2303904.
8. 49 U.S.C. §§ 30118-30120.
9. Stewart, E. Self-Driving Cars Have to Be Safer than Regular Cars. The Question Is How Much. *Vox*. https://www.vox.com/recode/2019/5/17/18564501/self-driving-car-morals-safety-tesla-waymo. Accessed Jul. 29, 2020.
10. Várhelyi, A., and T. Mäkinen. The Effects of In-Car Speed Limiters: Field Studies. *Transportation Research Part C: Emerging Technologies*, Vol. 9, No. 3, 2001, pp. 191–211. https://doi.org/10.1016/S0968-090X(00)00025-5.
11. Radun, I., J. Ohisalo, S. Rajalin, J. E. Radun, M. Wahde, and T. Lajunen. Alcohol Ignition Interlocks in All New Vehicles: A Broader Perspective. *Traffic Injury Prevention*, Vol. 15, No. 4, 2014, pp. 335–342. https://doi.org/10.1080/15389588.2013.825042.
12. Sparrow, R., and M. Howard. When Human Beings Are like Drunk Robots: Driverless Vehicles, Ethics, and the Future of Transport. *Transportation Research Part C: Emerging Technologies*, Vol. 80, 2017, pp. 206–215. https://doi.org/10.1016/j.trc.2017.04.014.
13. Kalra, N., and S. M. Paddock. Driving to Safety: How Many Miles of Driving Would It Take to Demonstrate Autonomous Vehicle Reliability? *Transportation Research Part A: Policy and Practice*, Vol. 94, 2016, pp. 182–193. https://doi.org/10.1016/j.tra.2016.09.010.
14. Shladover, S. E., and C. Nowakowski. Regulatory Challenges for Road Vehicle Automation: Lessons from the California Experience. *Transportation Research Part A: Policy and Practice*, Vol. 122, 2019, pp. 125–133. https://doi.org/10.1016/j.tra.2017.10.006.

15. National Highway Traffic Safety Administration. *Critical Reasons for Crashes Investigated in the National Motor Vehicle Crash Causation Survey*. Publication DOT HS 812 115. United States Department of Transportation, Washington, DC, 2015.
16. Koopman, P. Safe Autonomy: A Reality Check on the 94 Percent Human Error Statistic for Self-Driving Cars. Safe Autonomy, Jun 05, 2018.
17. National Highway Traffic Safety Administration. *Drowsy Driving 2015*. Publication DOT HS 812 446. United States Department of Transportation, Washington, DC, 2017.
18. National Highway Traffic Safety Administration. *NHTSA Field Crash Investigation 2018 Coding and Editing Manual*. Publication DOT HS 812 620. United States Department of Transportation, Washington, DC, 2020.
19. M. Davis and Company, Inc. *National Telephone Survey of Reported and Unreported Motor Vehicle Crashes*. Publication DOT HS 812 183. National Highway Traffic Safety Administration, Washington, DC, 2015.
20. Blincoe, L. J., T. R. Miller, E. Zalashnja, and B. A. Lawrence. *The Economic and Societal Impact of Motor Vehicle Crashes, 2010 (Revised)*. Publication DOT HS 812 013. National Highway Traffic Safety Administration, Washington, DC, 2015.
21. Antin, J. F., S. Lee, M. A. Perez, T. A. Dingus, J. M. Hankey, and A. Brach. Second Strategic Highway Research Program Naturalistic Driving Study Methods. *Safety Science*, Vol. 119, 2019, pp. 2–10. https://doi.org/10.1016/j.ssci.2019.01.016.
22. Dingus, T. A., F. Guo, S. Lee, J. F. Antin, M. Perez, M. Buchanan-King, and J. Hankey. Driver Crash Risk Factors and Prevalence Evaluation Using Naturalistic Driving Data. *Proceedings of the National Academy of Sciences*, Vol. 113, No. 10, 2016, pp. 2636–2641. https://doi.org/10.1073/pnas.1513271113.
23. National Highway Traffic Safety Administration. Traffic Safety Facts Annual Report Tables. *National Highway Traffic Safety Administration*. https://cdan.nhtsa.gov/tsftables/tsfar.htm. Accessed Jul. 13, 2020.
24. Fischhoff, B. Acceptable Risk: A Conceptual Proposal. *RISK: Health, Safety & Environment*, Vol. 5, No. 1, 1994, pp. 1–28.
25. Liu, P., R. Yang, and Z. Xu. How Safe Is Safe Enough for Self-Driving Vehicles? *Risk Analysis*, Vol. 39, No. 2, 2019, pp. 315–325. https://doi.org/10.1111/risa.13116.
26. Lund, I. O., and T. Rundmo. Cross-Cultural Comparisons of Traffic Safety, Risk Perception, Attitudes and Behaviour. *Safety Science*, Vol. 47, No. 4, 2009, pp. 547–553. https://doi.org/10.1016/j.ssci.2008.07.008.
27. Nees, M. A. Safer than the Average Human Driver (Who Is Less Safe than Me)? Examining a Popular Safety Benchmark for Self-Driving Cars. *Journal of Safety Research*, Vol. 69, 2019, pp. 61–68. https://doi.org/10.1016/j.jsr.2019.02.002.
28. Svenson, O. Are We All Less Risky and More Skillful than Our Fellow Drivers? *Acta Psychologica*, Vol. 47, No. 2, 1981, pp. 143–148. https://doi.org/10.1016/0001-6918(81)90005-6.

29. Peck, R. C., and J. Kuan. A Statistical Model of Individual Accident Risk Prediction Using Driver Record, Territory and Other Biographical Factors. *Accident Analysis & Prevention*, Vol. 15, No. 5, 1983, pp. 371–393. https://doi.org/10.1016/0001-4575(83)90015-5.
30. Peck, R. C., R. S. McBride, and R. S. Coppin. The Distribution and Prediction of Driver Accident Frequencies. *Accident Analysis & Prevention*, Vol. 2, No. 4, 1971, pp. 243–299. https://doi.org/10.1016/0001-4575(71)90042-X.
31. Lourens, P. F., J. A. M. M. Vissers, and M. Jessurun. Annual Mileage, Driving Violations, and Accident Involvement in Relation to Drivers' Sex, Age, and Level of Education. *Accident Analysis & Prevention*, Vol. 31, No. 5, 1999, pp. 593–597. https://doi.org/10.1016/S0001-4575(99)00015-9.
32. Massie, D. L., P. E. Green, and K. L. Campbell. Crash Involvement Rates by Driver Gender and the Role of Average Annual Mileage. *Accident Analysis & Prevention*, Vol. 29, No. 5, 1997, pp. 675–685. https://doi.org/10.1016/S0001-4575(97)00037-7.
33. Horswill, M. S., A. E. Waylen, and M. I. Tofield. Drivers' Ratings of Different Components of Their Own Driving Skill: A Greater Illusion of Superiority for Skills That Relate to Accident Involvement1. *Journal of Applied Social Psychology*, Vol. 34, No. 1, 2004, pp. 177–195. https://doi.org/10.1111/j.1559-1816.2004.tb02543.x.
34. United States Department of Transportation. National Transportation Statistics. *Bureau of Transportation Statistics*. https://www.bts.gov/product/national-transportation-statistics. Accessed Jul. 27, 2020.
35. FHWA (Federal Highway Administration). National Household Travel Survey. *U.S. Department of Transportation*. https://nhts.ornl.gov/. Accessed May 24, 2019.
36. Federal Highway Administration. Table 3-8. Average Truck Speeds on Selected Interstate Highways: 2009 - Freight Facts and Figures 2010 - FHWA Freight Management and Operations. *U.S. Department of Transportation*. https://ops.fhwa.dot.gov/freight/freight_analysis/nat_freight_stats/docs/10factsfigures/table3_8.htm. Accessed Jul. 28, 2020.
37. Federal Highway Administration. Table VM-1 - Highway Statistics 2016. *Office of Highway Policy and Information*. https://www.fhwa.dot.gov/policyinformation/statistics/2016/vm1.cfm. Accessed Jul. 28, 2020.
38. American Public Transportation Association. *2020 Public Transportation Fact Book Appendix A*. American Public Transportation Association, Washington, DC, 2020.
39. Lee, Y., T. S. Kim, H.-S. Cho, D. K. Sung, and B. D. Choi. Performance Analysis of an Elevator System during Up-Peak. *Mathematical and Computer Modelling*, Vol. 49, No. 3, 2009, pp. 423–431. https://doi.org/10.1016/j.mcm.2008.09.006.
40. McCann, M. *Deaths and Injuries Involving Elevators and Escalators*. Center for Construction Research and Training, Silver Spring, Maryland, 2013.
41. National Elevator Industry, Inc. Elevator And Escalator Fun Fact. *National Elevator Industry, Inc.*

http://www.neii.org/presskit/printmaster.cfm?plink=NEII%20Elevator%20and%20Escalator%20Fun%20Facts.cfm. Accessed Jul. 2, 2020.

42. May, J. If We Want to Increase Escalator Capacity, Why Don't We Just Run the Things Faster? | CityMetric. *CityMetric*. https://www.citymetric.com/horizons/if-we-want-increase-escalator-capacity-why-don-t-we-just-run-things-faster-1844. Accessed Jul. 28, 2020.
43. Hicks, D. J. The Safety of Autonomous Vehicles: Lessons from Philosophy of Science. *IEEE Technology and Society Magazine*, Vol. 37, No. 1, 2018, pp. 62–69. https://doi.org/10.1109/MTS.2018.2795123.
44. Hansson, S. O. *The Ethics of Risk: Ethical Analysis in an Uncertain World*. Palgrave Macmillan, Houndsmills, Basingstoke, Hampshire ; New York, NY, 2013.
45. Cregger, J., M. Dawes, S. Fischer, C. Lowenthal, E. Machek, and D. Perlman. *Low-Speed Automated Shuttles: State of the Practice*. Publication FHWA-JPO-18-692. John A. Volpe National Transportation Systems Center, Cambridge, MA, 2018.
46. Guo, F., S. G. Klauer, J. M. Hankey, and T. A. Dingus. Near Crashes as Crash Surrogate for Naturalistic Driving Studies: *Transportation Research Record*, 2010. https://doi.org/10.3141/2147-09.
47. Kiefer, R. J., D. J. LeBlanc, and C. A. Flannagan. Developing an Inverse Time-to-Collision Crash Alert Timing Approach Based on Drivers' Last-Second Braking and Steering Judgments. *Accident Analysis & Prevention*, Vol. 37, No. 2, 2005, pp. 295–303. https://doi.org/10.1016/j.aap.2004.09.003.
48. Goodall, N. More Than Trolleys: Plausible, Ethically Ambiguous Scenarios Likely to Be Encountered by Automated Vehicles. *Transfers*, Vol. 9, No. 2, 2019, pp. 45–58. https://doi.org/10.3167/TRANS.2019.090204.
49. Ori, M. Why Not Road Ethics? *Theoria*, Vol. 86, No. 3, 2020, pp. 389–412. https://doi.org/10.1111/theo.12248.