RAND
CORPORATION

# RAND Model of Automated Vehicle Safety (MAVS)

## Model Documentation

Nidhi Kalra and David G. Groves

For more information on this publication, visit www.rand.org/t/RR1902

Published by the RAND Corporation, Santa Monica, Calif.

# Preface

The RAND Corporation has a long history of research on intelligent systems. Since the 1950s, with work on chess-playing computers and the Logic Theory Machine, RAND has produced objective, evidence-based research to help inform how society can harness the benefits and manage the risks of intelligent, transformative technologies. RAND's work on highly automated vehicles builds on this firm foundation. The 2009 article "Liability and Regulation of Autonomous Vehicle Technologies" and the flagship report *Autonomous Vehicle Technology: A Guide for Policymakers* in 2014 (revised in 2016) examined the policy landscape surrounding these technologies. The 2016 report *Driving to Safety: How Many Miles of Driving Would It Take to Demonstrate Autonomous Vehicle Reliability?* determined whether it is possible and practical to test-drive highly automated vehicles as a method of assessing their safety.

This report continues the line of inquiry on highly automated vehicle safety. As the technology matures, policymakers must consider how near-term choices about highly automated vehicle policies will shape the future of road safety over time. This report helps inform that policy debate. It describes an exploratory model (the Model of Automated Vehicle Safety, or MAVS) that, given a user's hypothesis about a variety of factors, calculates how many lives would be lost each year in a future with and without highly automated vehicles. Insights from the model can help users consider how policies might shape these factors to improve safety now and over time. In particular, this model was used in the 2017 report *The Enemy of Good: Estimating the Cost of Waiting for Nearly Perfect Automated Vehicles* to assess how safe highly automated vehicles should be before they are allowed on the roads for consumer use.

## RAND Science, Technology, and Policy

This research was conducted in the RAND Science, Technology, and Policy program, which focuses primarily on the role of scientific development and technological innovation in human behavior, global and regional decisionmaking as it relates to science and technology, and the concurrent effects that science and technology have on policy analysis and policy choices.

This program is part of RAND Justice, Infrastructure, and Environment, a division of the RAND Corporation dedicated to improving policy- and decisionmaking in a wide range of policy domains, including civil and criminal justice, infrastructure development and financing, environmental policy, transportation planning and technology, immigration and border protection, public and occupational safety, energy policy, science and innovation policy, space, and telecommunications.

During the development of this report and at the time of publication, co-author Nidhi Kalra's spouse served as co-founder and president of Nuro, a machine-learning and robotics start-up company engaged in autonomous vehicle development. He previously served as a principal engineer for Google's driverless car project. Neither Kalra's spouse nor the companies he has worked for had any influence on this report.

Questions or comments about this report should be sent to the project leader, Nidhi Kalra (Nidhi_Kalra@rand.org). For more information about RAND Science, Technology, and Policy, see www.rand.org/jie/stp or contact the director at stp@rand.org.

## RAND Ventures

RAND is a research organization that develops solutions to public policy challenges to help make communities throughout the world safer and more secure, healthier and more prosperous. RAND is nonprofit, nonpartisan, and committed to the public interest.

RAND Ventures is a vehicle for investing in policy solutions. Philanthropic contributions support our ability to take the long view, tackle tough and often-controversial topics, and share our findings in innovative and compelling ways. RAND's research findings and recommendations are based on data and evidence, and therefore do not necessarily reflect the policy preferences or interests of its clients, donors, or supporters.

This venture was made possible, in part, by the Zwick Impact Fund. Charles Zwick—a researcher at RAND from 1956 to 1965 who later served as both a trustee and an advisory trustee—presented RAND Ventures with $1 million and the charge to take on new and emerging policy challenges and to support top talent in their focus on these issues. Each year, RAND's president draws on this generous gift to help RAND research and outreach teams extend the impact of completed research.

Support for this project is also provided, in part, by the income earned on client-funded research, and by the generous contributions of the RAND Justice, Infrastructure, and Environment Advisory Board.

# Contents

## Figures



## Tables

# Summary


Highly automated vehicle (HAV) safety is a principal concern for the transportation industry, policymakers, and the public. Much of the concern is focused on how safe HAVs should be before they are allowed on the road for consumer use. Determining this is important because it will shape when and how HAVs are introduced into the marketplace. However, HAV safety at introduction considers only part of the picture and should be complemented by asking how safe HAVs will become over time and how policy choices made today could shape the future of road safety.

It is important to consider safety *over time* for several reasons. First, the impact of HAVs on road safety will be far greater in the future than it is today as more people adopt such vehicles over time. Second, HAV safety will improve over time through learning and further development. Near-term choices about when HAVs are introduced will shape how safe they become in the future and how quickly that improvement takes place. Third, HAVs are expected to increase travel demand because they will reduce the cost of driving. So, even if the safety of HAVs improves, the effects of a lower fatality rate, for example, may be offset by growth in the amount of travel. Fourth, the future performance of HAVs is often compared with the current performance of human drivers, but this is a moving benchmark. Without considering potential changes in safety among non-HAVs over time, these comparisons are incomplete.

The important long-term implications of the timing of introducing HAVs, the rate of their diffusion, their improvement over time, and changes in non-HAV safety are essentially invisible when the discussion of safety focuses only on the present. To enable discussions to consider safety over the long term, we have developed a simple model of how these factors may interact and result in different safety outcomes over time.

Importantly, the Model of Automated Vehicle Safety (MAVS) does not predict the pace of future diffusion or the rate of learning of HAVs. Nor does it predict how safety will evolve among non-HAVs. Predicting technological diffusion, human behavior, and future interactions of new technology in real-world transportation settings decades into the future is exceedingly difficult—and any such prediction would likely be wrong. Furthermore, such predictions may not provide the needed insights into what would lead to favorable outcomes and how policies could promote these outcomes.

Instead, given a user's hypothesis or scenario about key factors, MAVS estimates safety outcomes in a future with and without HAVs. This helps users explore how different assumptions lead to different outcomes and which factors may lead to greater or lesser safety. Such insights can help policymakers and the public understand how policies can shape these factors to improve safety. The model can also offer insights into how to balance HAV safety at the point of introduction and the goal of improving safety over time. In particular, this model was

used in the report *The Enemy of Good: Estimating the Cost of Waiting for Nearly Perfect Automated Vehicles* (Kalra and Groves, 2017) to assess how safe HAVs should be before they are allowed on the roads for consumer use.

This report documents MAVS. The model calculates how HAVs could improve or worsen road safety (measured in fatalities, injuries, crashes, or other metrics) over time depending on the following user-defined assumptions about non-HAVs and HAVs:

- the change in travel demand in a future without HAVs
- the change in safety performance of non-HAVs, including such features as advance driver assistance systems that provide some automated functions
- the timing of HAV market introduction and the rate and level of diffusion and use over time
- the safety performance of HAVs at the time of introduction, the rate of improvement in safety as a function of miles traveled, and the maximum possible safety performance
- the upgradeability of the existing HAV fleet to keep pace with the state of the art in HAV performance.

From these factors, MAVS uses simple representations of technology diffusion and learning to calculate safety outcomes each year in a future with and without HAVs. This model is not intended to include all of the complex factors and relationships that could shape HAV safety, as doing so would complicate the model and obfuscate the key factors driving safety without adding accuracy (because those factors are also deeply uncertain). Rather, the model is intended as a foundation for exploration that enables users to expand their understanding of HAV safety from introduction and over time.

# Acknowledgments

We would like to acknowledge Anita Chandra, Marjory Blumenthal, and James Anderson for their advice and support from the very beginning of this work. We are enormously grateful to Kenneth Kuhn at RAND and Kara Kockelman at the University of Texas for their reviews. We thank Charles Zwick for his generous support of RAND through the Zwick Impact Awards, without which this report and related materials could not have been produced.

# Abbreviations

| | |
|---|---|
| ADAS | advanced driver assistance system |
| ADS | automated driving system |
| DDT | dynamic driving task |
| HAV | highly automated vehicle |
| MAVS | RAND Model of Automated Vehicle Safety |
| NHTSA | National Highway Traffic Safety Administration |
| ODD | operational design domain |
| VMT | vehicle miles traveled |

**Chapter One**

# Introduction

Motor vehicle crashes are a public health crisis in the United States and around the world. In 2015, 35,092 people lost their lives in such crashes in the United States, an increase of 7.2 percent from 2014, and 2.44 million were injured, an increase of 4.5 percent from 2014 (National Highway Traffic Safety Administration [NHTSA], 2016a). 2016 was even deadlier with 37,461 fatalities (NHTSA, 2017). U.S. motor vehicle crashes can pose economic and social costs of more than $800 billion in a single year (Blincoe et al., 2015). Moreover, more than 90 percent of crashes involve human errors (NHTSA, 2008), such as driving too fast and misjudging other drivers' behaviors, as well as distraction, fatigue, and alcohol impairment.

Many are looking to highly automated vehicles (HAVs)—vehicles that drive themselves some or all of the time—to mitigate this crisis.[1] But there is recognition that this technology, too, may pose risks to safety. HAV safety is a key concern for the transportation industry, policymakers, and the public.[2] Much of the concern is focused on how safe HAVs should be before they are allowed on the road for consumer use. Determining this is important because it will shape when and how HAVs are introduced into the marketplace. However, to address the issue properly, stakeholders must also ask how safe HAVs will become over time and how policy choices made today could shape road safety in the future.

It is important to consider safety *over time* for several reasons. First, the impact of HAVs on road safety will be far greater in the future than it is today, given that widespread diffusion of essentially any technology takes time. This is particularly so for automotive technologies because fleet turnover can take decades (Litman, 2017). So, even if HAVs are very safe or very unsafe in the initial years, the actual difference they will have on overall road safety will be very small. The impact will be far greater many years from now, if and when many more people use HAVs.

Second, HAV safety is expected to improve over time (Musk, 2015; Kalra, 2017). The machine-learning algorithms that govern HAV perception, decisionmaking, and execution rely largely on driving experience to improve. Therefore, the more (and more-diverse) miles that HAVs drive, the more potential there is for improving the state of the art in HAV safety performance. This means that near-term choices about when HAVs are introduced will also

[1] We use the term *HAV* to refer to vehicles that fall into Levels 3, 4, or 5 of SAE International's automated vehicle taxonomy (SAE International, 2016). We elaborate on this and other definitions and discuss differences in terminology in Chapter Two.

[2] For instance, Congress has held many hearings on HAVs, and of HAV safety is consistently a priority in statements from policymakers. See, for example, Walden (2017) and Collins (2016).

shape how safe they become in the future and how quickly that improvement takes place. For instance, a policy that requires HAVs to be nearly perfect before they are allowed on the roads could actually prevent them from achieving near perfection in any practical time frame because it would deny them the driving experience necessary to reach that level of performance. However, if HAVs are allowed on the road even when there are widespread safety concerns in the short term or the technology is demonstrably unsafe, there may be little market for the technology. The public may even demand that HAVs be removed from the road, resulting in much slower diffusion and delays in gaining the experience needed to improve.

Third, HAVs are expected to increase travel demand because they will reduce the cost of driving (Anderson et al., 2016; Fagnant and Kockelman, 2015). So, even if the safety of HAVs improves, the effects of a lower crash rate may be offset by increased travel. This is not a shortcoming, per se, as it shows that people can have greater mobility, which has significant economic and social value, with less risk. But the discussion of lives saved, injuries reduced, and crashes avoided is incomplete if the increase in travel is not included.

Fourth, the future performance of HAVs is often compared with the current performance of human drivers. While human driver safety is a natural benchmark, it is a moving benchmark (NHTSA, 2016a). On the one hand, infotainment applications—either built into a vehicle or on users' mobile devices—may increase distracted driving and crash rates. On the other hand, advanced driver assistance systems (ADASs) may compensate for and improve safety, thus decreasing crashes. Without considering potential changes in safety performance among non-HAVs over time, these comparisons are incomplete.

The important long-term implications of the timing of introducing HAVs, the rate of their diffusion, their improvement over time, and changes in non-HAV safety are essentially invisible when the discussion of safety focuses only on the present. To enable discussions to consider safety over the long term, we have developed a simple model of how these factors may interact and result in different road safety outcomes over time.

Importantly, the Model of Automated Vehicle Safety (MAVS) does not predict the pace of future diffusion, the rate of learning of HAVs, or how safety will evolve among non-HAVs. Predicting technological diffusion, human behavior, and future interactions of new technology in real-world transportation settings decades into the future is exceedingly difficult—and any such prediction would likely be wrong. Furthermore, such predictions may not provide the needed insights into what would lead to greater safety and how policies could promote greater safety.

Instead, given a user's hypotheses about those key factors, MAVS estimates how safety would compare in a future with and without HAVs. This helps users explore how different assumptions lead to different outcomes and which factors may lead to greater or lesser safety. Such insights can help policymakers and the public consider how policies might shape these factors to improve safety. The model can also offer insights into how to balance HAV safety at the point of introduction and the goal of improving safety over time.

This report documents MAVS. The model calculates how HAVs could improve or worsen road safety over time depending on the following user-defined assumptions about HAVs and non-HAVs:

- the change in travel demand in a future without HAVs
- the change in safety performance of non-HAVs
- the timing of HAV market introduction and the rate and level of diffusion and use over time
- the safety performance of HAVs at the time of introduction, the rate of improvement in safety as a function of miles traveled, and the maximum possible safety performance
- the upgradeability of the existing HAV fleet to keep pace with the state of the art in HAV performance.

From these factors, the model uses simple representations of technology diffusion and learning found in the literature to assess road safety in a future with and without HAVs.[3] The examples shown throughout the report of potential future outcomes are strictly meant to illustrate how MAVS works and should not be interpreted as the authors' predictions of the future.

[3] MAVS keeps exogenous many trends and conditions that would influence these key factors. For example, the rate of economic growth and the price of oil will be some of many factors that affect future travel demand. Endogenizing these trends would neither improve the model's accuracy (because they, too, are deeply uncertain) nor increase the exploratory value of the model (because the results would be increasingly difficult to interpret).

**Chapter Two**

# Problem-Framing and Definitions

To evaluate the effect of HAVs on overall vehicle safety, we first must answer the following questions:

- What is an HAV and what is not?
- How do we define a future with HAVs and a future without HAVs?
- How do we measure road safety?
- How has future road safety been assessed in the literature?

## What Is a Highly Automated Vehicle and What Is Not?

The 2016 SAE International taxonomy of driving automation (2016), summarized in Table 2.1, describes six levels of automation. Consistent with the Federal Automated Vehicles Policy (NHTSA, 2016b), we use the term *HAV* to refer to vehicles that conform to Levels 3, 4, or 5 of the taxonomy;[4] in such vehicles, the automated driving system performs the entire dynamic driving task (DDT) while engaged.[5] We use the term *non-HAV* to refer to vehicles that conform to Levels 0, 1, or 2; in such vehicles, the driver performs part or all of the DDT. Table 2.1 shows each level's description taken directly from the SAE International taxonomy (in italics) and our simplified interpretation.

Common to Level 0 (no driving automation), Level 1 (driver assistance), and Level 2 (partial driving automation) is that that the human behind the wheel is responsible for some or all of the DDT, even when driver assistance systems are engaged. Level 0 consists of traditional vehicles without advanced driver assistance systems to support a human driver. Level 1 automation is widespread in the form of adaptive cruise control or automated braking, the latter of which is expected to be standard in nearly all new passenger vehicle models by 2022 (U.S. Department of Transportation, 2016). Some automakers have marketed Level 2 autonomy, which controls both steering and acceleration (or braking). Tesla's autopilot-equipped vehicles have Level 2 automation because the human driver is expected to pay continuous attention to driving even when autopilot is engaged, but such Tesla vehicles are more advanced than nearly any other Level 2 vehicle in the extent to which they can control the vehicle and perform driving maneuvers. Level 2 vehicles have also raised safety concerns because these technologies may

---

[4] Note that, in the SAE International taxonomy, the term *highly automated* would apply to Level 4 vehicles specifically, but the Federal Automated Vehicles Policy uses the term *highly automated vehicle* more broadly.

[5] SAE International (2016) defines the DDT as "all of the real-time operational and tactical functions required to operate a vehicle in onroad traffic, excluding the strategic functions, such as trip scheduling and selection of destinations and waypoints."

encourage drivers to be distracted, even if they are supposed to pay continuous attention (Trimble et al., 2014).

**Table 2.1. SAE International Levels of Driving Automation**

| Level | Name | Description |
|---|---|---|
| **Driver performs part or all of the DDT (non-HAVs)** | | |
| 0 | No driving automation | *The performance by the driver of the entire DDT, even when enhanced by active safety systems.*<br><br>The human driver is entirely responsible for driving, even if such features as active electronic stability control are available and engaged. |
| 1 | Driver assistance | *The sustained and operational design domain (ODD)–specific execution by a driving automation system of either the lateral or the longitudinal vehicle motion control subtask of the DDT (but not both simultaneously) with the expectation that the driver performs the remainder of the DDT.*[a]<br><br>The human driver is entirely responsible for driving but may be assisted by a single feature that automates steering or acceleration, such as lane-keeping and adaptive cruise control, but not both. |
| 2 | Partial driving automation | *The sustained and ODD-specific execution by a driving automation system of both the lateral and longitudinal vehicle motion control subtasks of the DDT with the expectation that the driver completes the object and event detection and response subtask and supervises the driving automation system.*<br><br>The human driver is entirely responsible for driving but may be assisted by functions that automate both steering and acceleration, such as lane-keeping and adaptive cruise control; the human driver is responsible for monitoring the environment and intervening whenever needed. |
| **Automated driving system performs the entire DDT while engaged (HAVs)** | | |
| 3 | Conditional driving automation | *The sustained and ODD-specific performance by an [automated driving system (ADS)] of the entire DDT with the expectation that the DDT fallback-ready user is receptive to ADS-issued requests to intervene, as well as to DDT performance-relevant system failures in other vehicle systems, and will respond appropriately.*<br><br>The vehicle is entirely responsible for driving in certain conditions but may request rapid intervention from the human driver as needed. |
| 4 | High driving automation | *The sustained and ODD-specific performance by an ADS of the entire DDT and DDT fallback without any expectation that a user will respond to a request to intervene.*<br><br>The vehicle is entirely responsible for driving in certain conditions and will not request intervention from the human driver. |
| 5 | Full driving automation | *The sustained and unconditional (i.e., not ODD-specific) performance by an ADS of the entire DDT and DDT fallback without any expectation that a user will respond to a request to intervene.*<br><br>The vehicle is entirely responsible for driving under all conditions and will not request intervention from anyone in the vehicle. Such vehicles may have no occupants at all. |

SOURCE: SAE International, 2016.
NOTE: Below each italicized description taken directly from the SAE International taxonomy (2016, Table 2) is our simplified interpretation.
[a] SAE International (2016) defines an ODD as “the specific conditions under which a given driving automation system or feature thereof is designed to function.”

In contrast, common to Level 3 (conditional driving automation), Level 4 (high driving automation), and Level 5 (full driving automation) is that the vehicle is responsible for the entire

DDT when the automated driving capabilities are engaged. In addition, as we discuss in Chapter Four, we would add the distinction that learning through driving experience is a central mechanism for developing and improving vehicles at Levels 3–5 and plays a much smaller role in the development of vehicles at Levels 0–2.

Level 3 (conditional driving automation) allows the vehicle to drive but request intervention from a human driver at short notice if needed. Google's test vehicles are Level 3 in that their driving is highly automated but they can request intervention as needed from a trained operator. Like Level 2 vehicles, Level 3 vehicles for consumers have raised safety concerns because of the challenge drivers face in safely taking control of the vehicle (Trimble et al., 2014). Some automakers are focusing on higher levels of automation because of these concerns (Davies, 2015; Oremus, 2016).

Many automakers are developing Level 4 vehicles, whose driving is highly automated without human intervention in many, though not all, conditions. Other automakers are developing vehicles that can be highly automated in all conditions and may not even allow for human driver control. For example, new models from Tesla will be hardware-equipped for Level 5 highly automated driving, although this feature will not be activated immediately. Tesla has indicated that it will run the Level 5 capability in a "shadow mode." This mode, which simulates but does not execute control of the vehicle, will allow Tesla to develop and test Level 5 systems under millions of real driving miles before activating it in consumer vehicles (Nishimoto, 2016). As of 2017, no vehicles at Levels 3–5 are available for consumers to lease or purchase, but pilot tests are under way with trained safety drivers behind the wheel.

## How Do We Define a Future with Highly Automated Vehicles and a Future Without?

Given the definitions of HAVs, non-HAVs, and the six levels of automation, we define a future *without* HAVs as one that has a mix of Level 0, 1, and 2 vehicles but does not have any Level 3, 4, or 5 vehicles. The future *with* HAVs has vehicles at all levels.

## How Do We Measure Road Safety?

Safety is typically measured in terms of fatalities, injuries, and property-damage-only crashes, both in absolute terms and as a rate of vehicle miles traveled (VMT). HAVs will change the future of road safety along any of these metrics and may not affect them in the same way. For example, HAVs could decrease high-probability, low-consequence property-damage-only crashes while increasing low-probability, high-consequence multi-vehicle crashes.

MAVS can be used to explore safety in any of these terms: The model equations we describe refer only to a generic performance rate, which can reflect any of these metrics, and the model equations remain the same regardless. Similarly, the model can be used to explore road safety in different locations around the world. Only a handful of parameters relating to safety performance rates would need to change to reflect different metrics of safety or geographic areas of concern. Total incidents (whether fatalities, injuries, or crashes) are calculated as the product of how much driving will occur—in terms of VMT—and how safe driving is in terms of incidents per mile traveled. Cumulative safety incidents are the sum of incidents over some number of years. For the purposes of illustrating MAVS, however, we focus on fatalities in the United States.

Since the 1970s, the number of VMT has been increasing in the United States, as shown in Figure 2.1. This trend is expected to continue (Federal Highway Administration, 2016), although diffusion of HAVs may affect this trajectory (Anderson et al., 2016; Fagnant and Kockelman, 2015; Bierstedt et al., 2014) by reducing the cost of driving and by enabling mobility for older adults, individuals with disabilities, and others who may not currently drive.

**Figure 2.1. U.S. VMT, by Year**

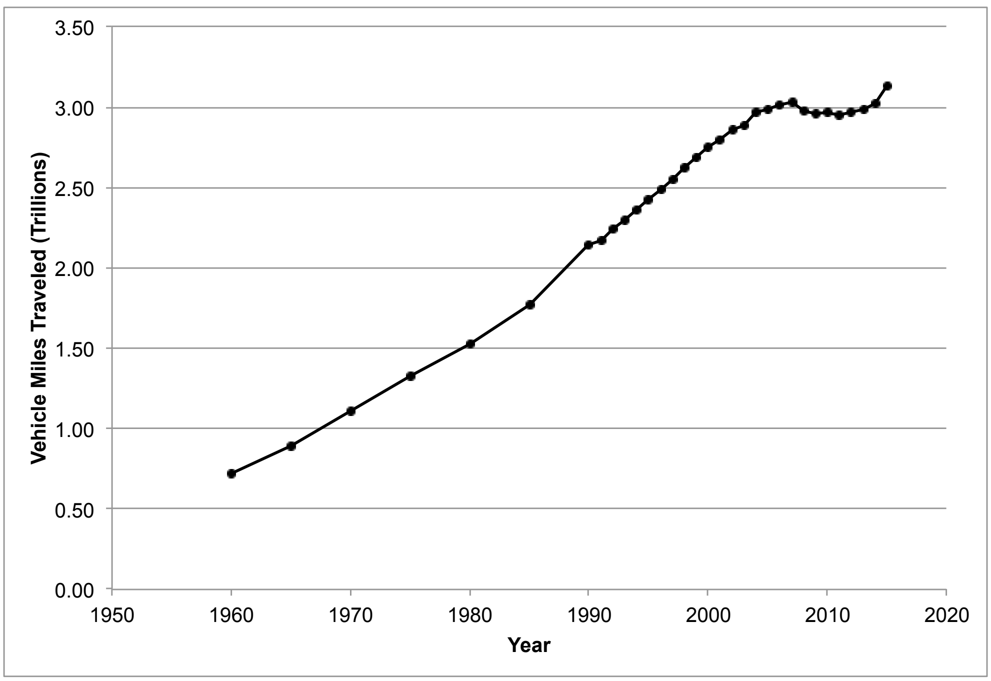


SOURCE: NHTSA, 2016a.

As shown in Figure 2.2, despite the increase in VMT, automobile fatalities have decreased over time, from a peak of just more than 51,000 fatalities in 1979 to a low of just less than 32,500 fatalities in 2011. The fatality rate has also declined generally, although there has been an uptick in recent years. While vehicles are becoming safer in general through better safety equipment (airbags, anti-lock brakes, driver assistance systems, etc.), increases in driver distractions (e.g., from mobile phones and other devices) are among the factors offsetting those gains (NHTSA, 2016a).

**Figure 2.2. U.S. Road Fatalities and Fatality Rates, by Year**

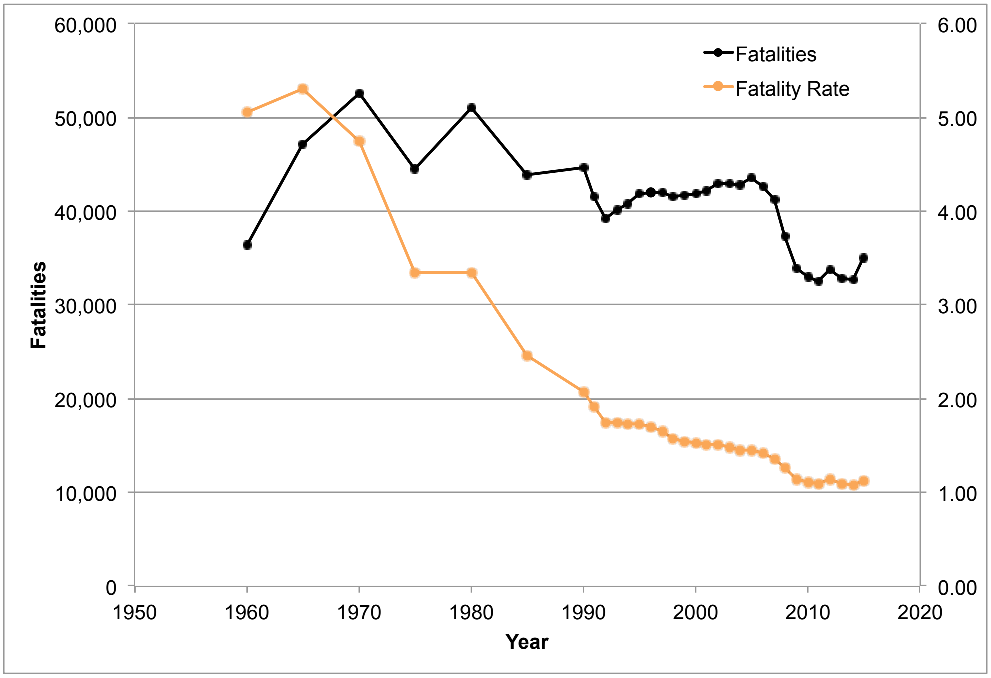


SOURCE: NHTSA, 2016a.

For comparison, Figure 2.3 shows historical police-reported road injuries and injury rates, and Figure 2.4 shows historical police-reported crashes and crash rates. These data show that injury and crash incidents and rates have decreased in general, but there has been an increase in recent years. It is important to note, however, that injuries and crashes may be significantly underreported—by 25 percent and 60 percent, respectively, according to one study (Blincoe et al., 2015). Therefore, these data are only a partial snapshot of safety.

MAVS explores how the safety of non-HAVs and HAVs could evolve over a 50-year period, from 2020 to 2069. The model allows the VMT and *safety rates* (that is, incidents per VMT) of both non-HAVs and HAVs to change in this time frame based on users' inputs. The model does assume that VMT and safety rates will change smoothly over time, but, in reality, there will likely be punctuated periods during which the rates may significantly change, reflecting technological leaps or other developments. MAVS further assumes that, like many other automotive technologies (including traditional vehicles), HAVs will only become safer over time through experience, learning, and improvement.

**Figure 2.3. U.S. Reported Road Injuries and Injury Rates, by Year**

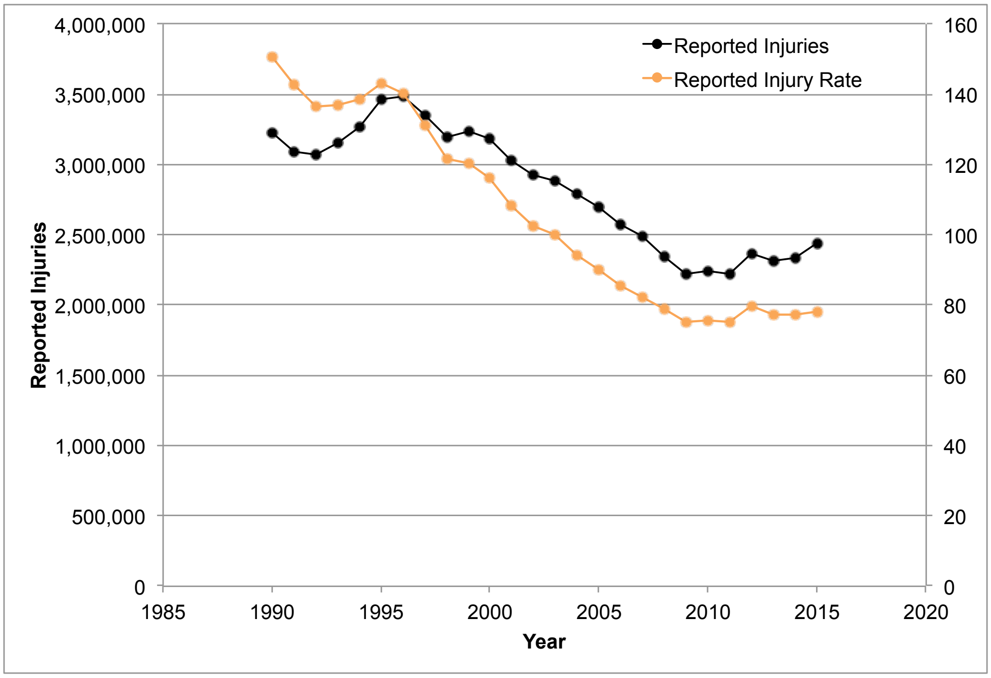


SOURCE: NHTSA, 2016a.

**Figure 2.4. U.S. Reported Crashes and Crash Rates, by Year**

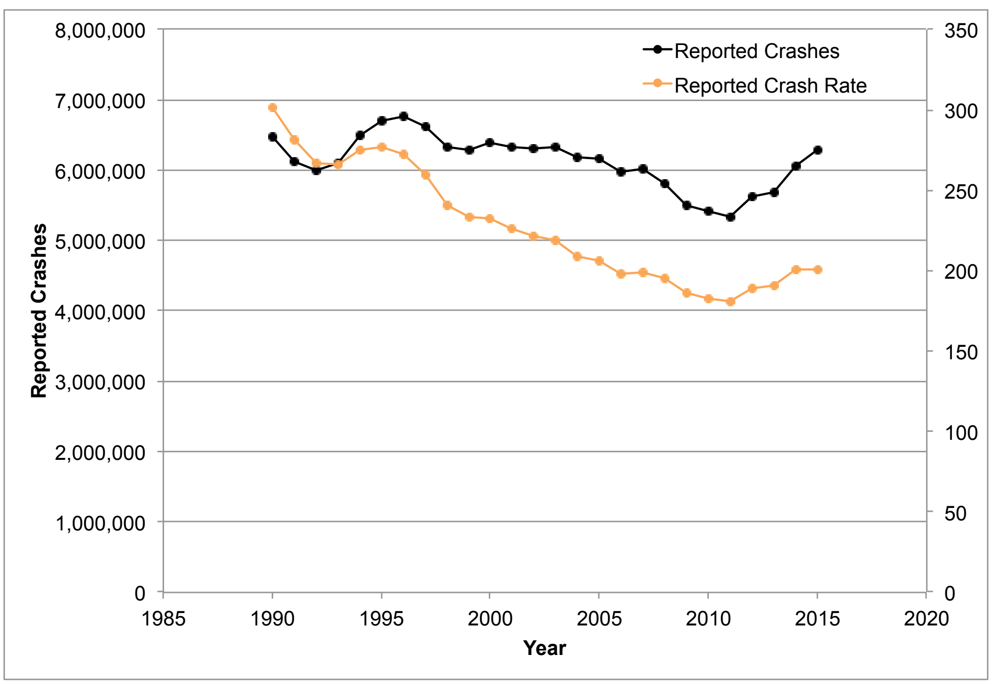


SOURCE: NHTSA, 2016a.

## How Has Future Road Safety Been Assessed in the Literature?

There are many estimates of the benefits of different types of advanced driver assistance systems or crash avoidance systems, individually and in combination (Gordon et al., 2010; Funke et al., 2011; Perez et al., 2011; Jermakian, 2011; Harper, Hendrickson, and Samaras, 2016). Vehicles equipped with these technologies are usually classified as having Level 1 or Level 2 automation according to SAE International's taxonomy and are considered non-HAVs. As one example, Funke et al. (2011) summarizes the safety potential of four crash avoidance technologies as the product of the size of the crash problem in the entire U.S. fleet (e.g., the number of annual crashes related to lane departure) and the fraction of such crashes that could be mitigated by the technology (e.g., from a lane departure warning system). There are also efforts to estimate the benefits of connected vehicle technologies, in which on-board applications use communication with other vehicles or infrastructure to improve safety—for example, for coordinating vehicle movement through an intersection (Najm, Toma, and Brewer, 2013; Eccles et al., 2012).

Rau, Yanagisawa, and Najm (2015) describes a method for identifying the types and potential number of current crashes that could be mitigated by technologies between Level 2 and Level 5 autonomy. Li and Kockelman (2016) draws on this methodology to estimate the safety benefits of a variety of connected vehicle and Level 1 and Level 2 automated vehicle technologies, assuming widespread adoption of those technologies. Going a step further than prior work, Li and Kockelman (2016) estimates both the types and severity of crashes that could be avoided by each type of technology, as well as the economic benefit of those savings. In recognition of the uncertainty in the technology performance, the authors assess benefits under three scenarios of technology effectiveness.

There are fewer estimates of the safety benefits of HAVs, and there is no consensus yet among those estimates (Winkle, 2015). Fagnant and Kockelman (2015) calculate the societal benefits of Level 4 and Level 5 HAVs across a variety of benefit categories, including safety.[6] Drawing on the findings of the National Motor Vehicle Crash Causation Survey, which found that human error accounts for 93 percent of today's crashes (NHTSA, 2008), Fagnant and Kockelman assume in their calculations that HAVs reduce crash and injury rates by 50 percent at the 10-percent market penetration rate and by 90 percent at the 90-percent market penetration rate. In contrast, the Casualty Actuarial Society's Automated Vehicles Task Force recently evaluated the findings of the National Motor Vehicle Crash Causation Survey in the context of HAVs. The task force's study found that HAVs could address about half of the accidents, while "49% of accidents contain at least one limiting factor that could disable [HAV] technology or reduce its effectiveness" (Casualty Actuarial Society, 2014).

---

[6] The paper does not explicitly define what levels of autonomy the authors include in their calculations, but we infer that they refer to Level 4 and Level 5 autonomy, not Level 3.

This literature provides important insights into how different non-HAV and HAV technologies could mitigate today's crashes. However, these insights have not yet been used to understand how safety effects might play out over time—because different technologies are adopted in different time frames, and the performance of the technologies changes as they are deployed. It is difficult to use these estimates to make such projections for two key reasons. First, the estimates generally focus on how technologies could mitigate the types of crashes that human drivers currently cause, but they overlook important ways in which new technologies could add to crashes. This could occur if technology erodes human drivers' skills or attention, technology is vulnerable to cybersecurity failures that lead to new types of crashes, or HAVs simply perform worse than human drivers, even initially (Kalra, 2017).[7] As the Casualty Actuarial Society notes, "The safety of automated vehicles should not be determined by today's standards; things that cause accidents today may or may not cause accidents in an automated vehicle era" (2014, p. 1). New safety risks are difficult to anticipate, making the full effect of many new technologies deeply uncertain.

Second, these existing estimates also compare the marginal benefits of a technology with the safety performance of current vehicles and drivers. However, the benefit of a vehicle with a particular technology at some point in the future is more correctly estimated when compared with the performance of future vehicles without that technology at that same future time, rather than with vehicles in current conditions.

The history of airbags illustrates these issues (Anderson et al., 2016; Houston and Richardson, 2000). When airbags were first introduced in the 1970s, they were designed to protect an unbelted adult male passenger and envisioned as an alternative rather than a supplement to seat belts, which were then used infrequently. Estimates of future safety benefits made at that time were based on this use case and ultimately proved to be overblown by an order of magnitude—in large part because, by the time airbags were widely deployed, seat belt use was also widespread, so the marginal benefit of airbags was much less than anticipated. Moreover, while airbags still saved many lives, the force needed to protect an unbelted adult male injured and killed many passengers of smaller stature (such as women and children) who might have otherwise survived the crash had airbags not deployed. These crashes led to improvements in airbag technology but also showed that airbags introduced new crash risks even as they mitigated existing risks.

In sum, the long-term evolution of road safety is important to understand and yet complex, deeply uncertain, and difficult to predict. This work fills a gap in the existing literature by using a simple modeling platform to explore the safety impact of HAVs under different policies and conditions. A follow-on study (Kalra and Groves, 2017) uses MAVS to better understand the factors that contribute to road safety and to identify how robust policies can help shape a safer future.

---

[7] Complicating matters, as Kalra and Paddock (2016) argues, there is no currently accepted method of assessing the safety of HAVs with statistical confidence prior to making them available for widespread use. Therefore, it is possible that stakeholders may simply not know how safe the technology is.

Chapter Three

# Modeling Safety in a Future Without Highly Automated Vehicles

As described in Chapter Two, safety as measured by the number of safety incidents is a product of VMT and safety rates (incidents per VMT). Both of these rates may change in a future without HAVs, just as they have historically. How they will change is uncertain but can be modeled using simple relationships.

## Future Safety Rates of Non–Highly Automated Vehicles

The Federal Highway Administration projects total future VMT. Its 2016 projections suggest, for example, that from 2014 to 2044, total VMT could grow between 0.53 percent and 0.65 percent annually (Federal Highway Administration, 2016). While these bounds are themselves uncertain and include potential changes from HAV use,[8] the model of compounded annual growth is useful in thinking about changes in non–highly automated VMT over time.

Using a reference point of 3.131 trillion VMT in 2015, we can model VMT in a future without HAVs using Equation 3.1.

$$VMT_x = 3.131 \times 10^{12} \times (1 + g)^{(x-2015)}, \tag{3.1}$$

where $x$ is the future year and $g$ is the compound annual growth rate in VMT chosen by the user, assuming that there are no HAVs. Figure 3.1 shows historical total VMT and two scenarios of future total VMT under two different assumptions about compounding growth over the model study period of 2020 to 2069.

[8] According to the Federal Highway Administration, its projections include a variety of uncertainties, such as uncertainty in future economic growth, vehicle use and ownership, and technology such as HAVs.

**Figure 3.1. Historical VMT and Two Scenarios of Future VMT**

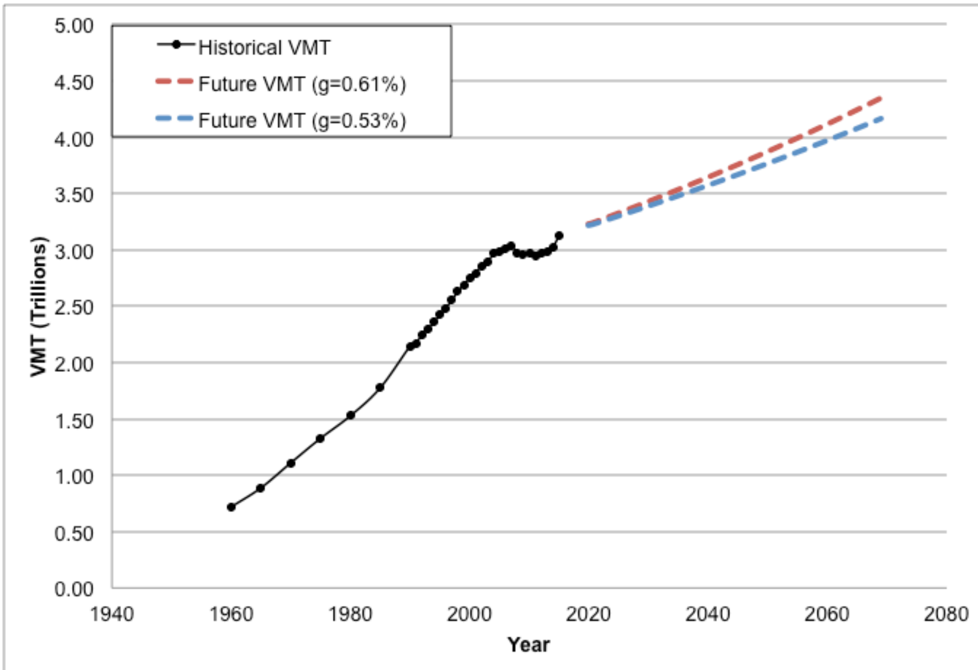


## Non–Highly Automated Vehicle Safety Rate over Time

While road fatality rates have declined significantly since the 1950s, there has been a plateau over the past decade and an increase in recent years (Bureau of Transportation Statistics, 2016). Estimates of the safety benefits of driver assistance systems described in Chapter Two suggest that non-HAVs will become safer over time. However, it is also possible that fatality rates could increase nonetheless if the safety benefit of these technologies is outpaced by a decline in driver attentiveness and skill (which will still be essential for non-HAV driving) or other maladaptive behaviors (Milakis, van Arem, and van Wee, 2017). To incorporate this uncertainty, we use a simple linear model that calculates safety rates $s_x$ (e.g., fatalities per VMT) in year $x$ based on safety rates in reference year 2015, $s_{2015}$, and given a user's projection of non-HAV safety rates in the year 2069, $s_{2069}$, as shown in Equation 3.2:

$$s_x = mx + b, \tag{3.2}$$

where

$$m = \frac{(s_{2015} - s_{2069})}{(2015 - 2069)}$$

$$b = ms_{2015} - 2015.$$

Here, $s_{2015}$ and $s_{2069}$ could be fatality, injury, or crash rates per VMT, depending on the safety metric used. Figure 3.2 shows historical fatality rates and two scenarios of future fatality rates of non-HAVs. One future scenario assumes that the fatality rate declines to 0.9 fatalities per 100 million miles in 2069, a 20-percent decrease relative to the 2015 fatality rate of 1.12 fatalities per 100 million VMT. Here, $s_{2015} = 1.12 \times 10^{-8}$ and $s_{2069} = 0.9 \times 10^{-8}$. The other future scenario assumes that the fatality rate declines to 0.7 fatalities per 100 million VMT by 2069, an approximately 40-percent decrease relative to 2015 fatality rate. Here, $s_{2069} = 0.7 \times 10^{-8}$.

**Figure 3.2. Historical Fatality Rates and Two Scenarios of Future Fatality Rates**

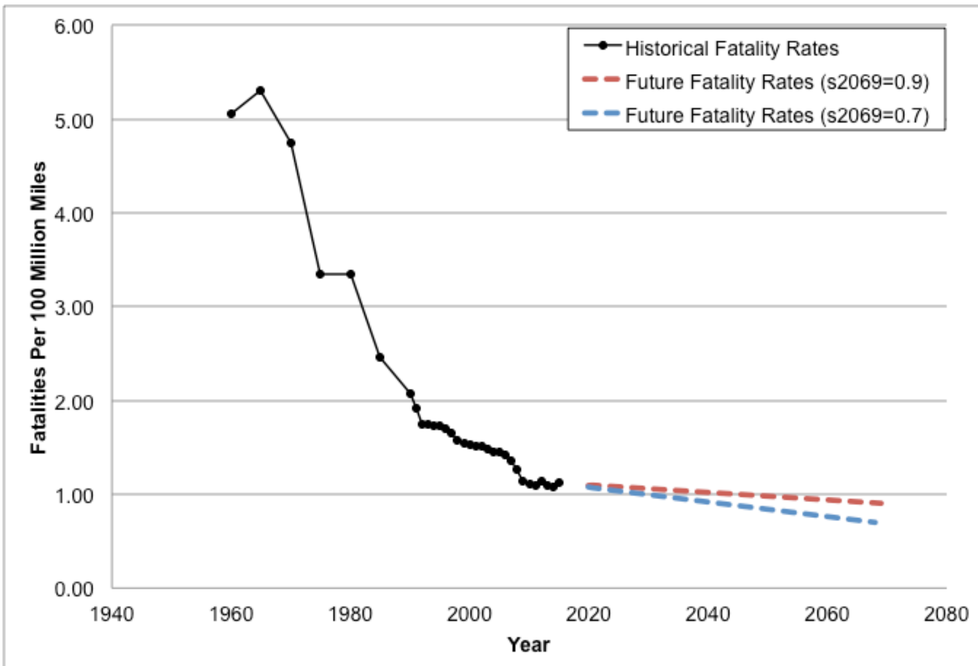


## Highly Automated Vehicle Safety Rate over Time

Calculating the number of safety events $f$ is straightforward once annual VMT and safety rates are determined using Equations 3.1 and 3.2, as shown in Equation 3.3:

$$f = s_x VMT_x \ . \tag{3.3}$$

Figure 3.3 shows historical fatalities and two scenarios of projected fatalities in a future without HAVs, using the projected annual VMT and fatality rates shown in Figures 3.1 and 3.2. In one scenario, annual fatalities increase overall because the increase in VMT ($g = 0.61$ percent) outpaces the modest decrease in fatality rates ($s_{2069} = 0.9$). In the other scenario, annual fatalities decrease because the increase in VMT ($g = 0.53$ percent) is offset by the large decrease in fatality rates ($s_{2069} = 0.7$).

**Figure 3.3. Historical Annual Fatalities and Two Scenarios of Future Annual Fatalities**

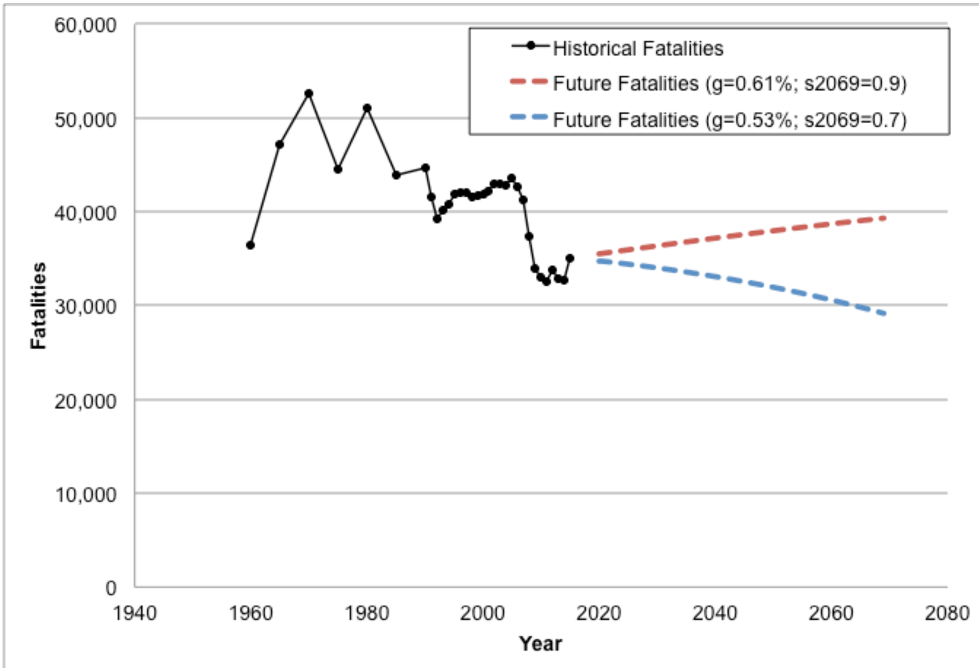


From this point forward, we call a future without HAVs the *baseline* against which to compare the future with HAVs. Note that the user of the model can adjust the specifics of the baseline. The figures in Chapter Four all assume the latter values for the baseline future ($g = 0.53$ percent and $s_{2069} = 0.7$).

Chapter Four

# Modeling Safety in a Future with Highly Automated Vehicles

The number of safety incidents in a future with HAVs is the sum of incidents among non-HAVs and HAVs. This is described in Equation 4.1.

$$f_x = VMT_{n,x}s_{n,x} + VMT_{a,x}s_{a,x}\ , \tag{4.1}$$

where

- $f_x$ is the number of safety incidents in year $x$
- $VMT_{n,x}$ is the VMT of non-HAVs in year $x$
- $s_{n,x}$ is the safety rate of non-HAVs in year $x$
- $VMT_{a,x}$ is the VMT of HAVs in year $x$
- $s_{a,x}$ is the safety rate of HAVs in year $x$.

Note that, for this formulation, $VMT_n$ plus $VMT_a$ equals total VMT in a future with HAVs. We begin by describing how we model VMT of non-HAVs and HAVs and then describe how we model the safety rates of non-HAVs and HAVs.

## Highly Automated and Non–Highly Automated Vehicle Miles Traveled over Time

The VMT of HAVs ($VMT_a$) and non-HAVs ($VMT_n$) depends on two factors: (1) when and how fast HAVs are adopted and (2) how much HAVs are driven compared with non-HAVs. These are in addition to the factors needed to model VMT in a future without HAVs.

We represent diffusion of HAVs over time as the fraction of VMT that ultimately would be attributed to HAVs and the time that it would take to reach this level of "full diffusion." We model this pattern of diffusion using a logistic function. Logistic functions (or other "S curves") are often used to model the diffusion of technology or other phenomena in which growth is initially slow, then increases exponentially, then slows to linear as saturation begins, and then stops at full maturity (Karshenas and Stoneman, 1993; Jeyaraj and Sabherwal, 2015). The logistic function reasonably approximates more-detailed scenarios of HAV diffusion that have been developed by others, such as Bansal and Kockelman (2017) and Litman (2017). A standard logistic function is described by Equation 4.2:

$$f(x) = \frac{1}{1+e^{-x}}\ . \tag{4.2}$$

We scale $x$ to describe time in years and $f(x)$ to describe HAV diffusion in terms of the *fraction of baseline VMT that is highly automated.* This requires defining four parameters:

- $A_{max}$, the largest fraction of baseline VMT that will be replaced by highly automated driving in a single year.[9] If HAVs are expected to fully replace human-driven vehicles one day, then $A_{max} = 1$. Alternatively, human-driven vehicles may persist, so $0 < A_{max} < 1$.
- $x_a$, the year in which some fraction $a$ of baseline miles are driven by HAVs.
- $x_{(1-a)}$, the year in which some fraction $1-a$ of baseline miles are driven by HAVs.

This results in Equation 4.3, the fraction $M_x$ of baseline miles driven by HAVs each year:

$$M_x = \frac{A_{max}}{1+e^{-(mx+b)}} , \qquad (4.3)$$

where

$$m = \frac{2 \ln\left(\frac{1-a}{a}\right)}{x_{1-a} - x_a}$$

$$b = -\ln\left(\frac{1-a}{a}\right) - mx_a .$$

Appendix A describes the mathematical transformation of the logistic function (Equation 4.2) into a function of the fraction of highly automated miles driven over time (Equation 4.3).

Figure 4.1 shows a diffusion curve of the percentage of baseline miles $M_x$ driven by HAVs in year $x$, in a scenario in which 0.01 percent of diffusion is achieved in year 2030, 99.99 percent of diffusion is achieved in year 2069,[10] and maximum diffusion occurs when 80 percent of baseline VMT are highly automated. In this scenario, $a = 0.0001, x_a = 2030, x_{a-1} = 2069$, and $A_{max} = 0.8$. The data points denoting the baseline years of diffusion are annotated. In year 2030, 0.01 percent of the diffusion has occurred and 0.008 percent of baseline VMT has been replaced by highly automated driving. In year 2069, 99.99 percent of the diffusion has occurred and 79.992 percent of baseline VMT has been replaced by highly automated driving.

[9] For example, if VMT is 3 trillion non–highly automated miles in a baseline future and $A_{max}$ is 66 percent, then, in a future with HAVs, 2 trillion of those miles will be replaced by highly automated driving and 1 trillion will remain human-driven.

[10] For comparison, Litman (2017) notes that many vehicle technologies require 30 to 50 years to be implemented in 90 percent or more of operational vehicles.

**Figure 4.1. Scenario of the Percentage of Baseline Miles That Will Be Driven by HAVs in the Future**

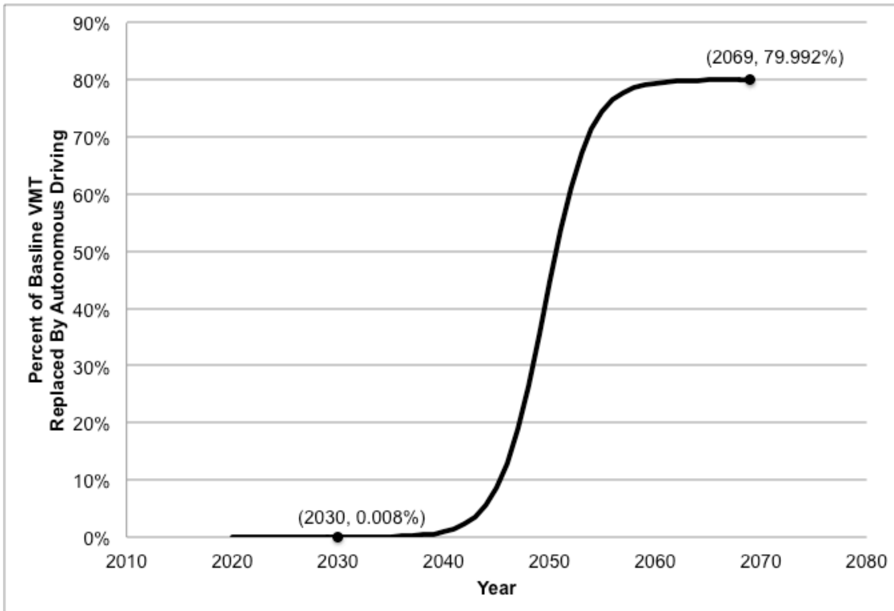


NOTE: This figure shows diffusion over time under a scenario in which 0.01 percent of diffusion is achieved in year 2030, 99.99 percent of diffusion is achieved in year 2069, and maximum diffusion occurs when 80 percent of total baseline VMT is replaced by highly automated driving.

After modeling the diffusion of HAVs in terms of percentage of baseline miles that will be driven by HAVs, we consider how VMT may change as a result of the diffusion of HAVs. HAVs are widely expected to increase miles traveled because they lower the cost of travel, can provide travel opportunities for those who may currently lack mobility, and enable zero-occupancy vehicle travel (Anderson et al., 2016; Fagnant and Kockelman, 2015; Bierstedt et al., 2014). There may also be complex effects from the use of shared HAVs (i.e., highly automated taxis) (Greenblatt and Saxena, 2015) and mode shifts from transit. The end effect of these changes is unknown, given the disruptive nature of the technology.

To reflect changes in travel demand, our model requires the user to determine how many miles of highly automated driving replace each mile of human driving, when such a replacement occurs. This is defined with the parameter $c$, where $c = 0$ indicates that there is a one-to-one replacement and no change in VMT; $c > 0$ means that each human-driven mile is replaced by more than one mile of highly automated driving, thereby increasing VMT; and $c < 0$ means that each human-driven mile is replaced by less than one mile of highly automated driving, thereby decreasing VMT.

The total VMT in a future with HAVs, $VMT_x$, is then the baseline VMT in year $x$ (as calculated in Equation 3.1) that is not replaced by highly automated driving, plus the baseline VMT that is replaced by highly automated driving, multiplied by $(1+ c)$. This is given by Equation 4.4:

$$VMT_x = VMT_{baseline,x}(1 - M_x) + VMT_{baseline,x} M_x (1 + c_f)\ . \tag{4.4}$$

As this suggests, the highly automated VMT, $VMT_{a,x}$, and non–highly automated VMT, $VMT_{n,x}$, in each year $x$ are given by Equations 4.5 and 4.6:

$$VMT_{a,x} = VMT_{baseline,x} M_x (1 + c_f) \tag{4.5}$$

$$VMT_{n,x} = VMT_{baseline,x}(1 - M_x)\ . \tag{4.6}$$

Figure 4.2 shows the resulting total VMT, highly automated VMT, and non–highly automated VMT under the assumption that HAVs will replace each mile of non–highly automated driving with 1.2 miles of highly automated driving (i.e., $c = 0.2$). Here, and in remaining figures in this section, the baseline VMT assumes compounding annual growth of 0.53 percent, as shown in one of the scenarios in Figure 3.1.

**Figure 4.2. Scenario of Highly Automated and Non–Highly Automated VMT**

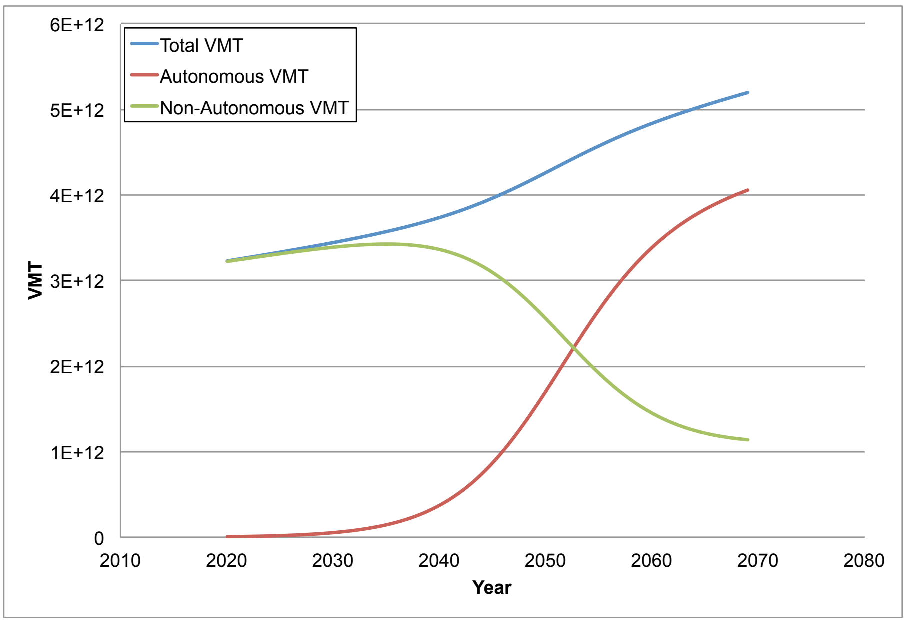


NOTE: This figure shows VMT over time under a scenario in which 0.01 percent of HAV diffusion is achieved in year 2030, 99.99 percent of diffusion is achieved in year 2069, maximum diffusion occurs when 80 percent of baseline VMT is highly automated, and HAVs replace each baseline mile with 1.2 highly automated miles.

## Non–Highly Automated Vehicle Safety Rate over Time

We assume that non-HAVs will have the same safety rate in a future with HAVs as they would in a future without HAVs.[11] In other words, $s_{n,x} = s_x$ as defined in Equation 3.2 in Chapter Three. However, the number of safety incidents from non-HAVs will be lower than in the baseline future because non–highly automated VMT will be lower.

## Highly Automated Vehicle Safety Rate over Time

We model the HAV safety rate through two safety factors: the change in the state of the art in HAV safety and the extent to which the state of the art is diffused through the existing HAV fleet.

### Change in the State of the Art in Highly Automated Vehicle Safety

Changes in the state of the art in HAV safety may emulate that of past automotive technologies, with improvements occurring in steps as developers improve designs and as supporting technological advancements are made (e.g., better and cheaper sensors). But, unlike past automotive technologies, HAV safety can also improve continuously through the very experience of driving. When a human driver makes a mistake on the road, typically only that individual can learn from that experience to improve his or her driving habits. Other drivers are unaffected. This is not the case with HAVs. HAV developers use the driving experience of individual vehicles to improve the state of the art in HAV safety. As one example of improvement, Google HAVs are requiring less-frequent disengagements, which occur when either the HAV requests intervention from the test driver or the test driver intervenes voluntarily out of a concern for safety. Google reported to the California Department of Motor Vehicles that in the fourth quarter of 2014, vehicles disengaged once every 600 miles approximately, but in the fourth quarter of 2015, the vehicles disengaged once every 2,800 miles (Google Auto LLC, 2015).[12] In 2016 overall, this dropped further to one disengagement every 5,000 miles (Waymo, 2016). As another example, Tesla first made its Autopilot technology active in Model S vehicles in October 2015 and subsequently upgraded the associated firmware in December 2015 and August 2016 to include, among other things, better safety performance based on experiences of

---

[11] A future elaboration of MAVS could forgo this assumption and create another user input. One can imagine, for example, that human driving may become worse in a future with widespread HAV use because people's driving skills and attention could degrade. On the other hand, non-HAVs may become safer through such technologies as vehicle-to-vehicle communication, which might be more effective in a future with HAVs.

[12] The disengagements reported by Google reflect cases in which the HAV software detected a technology failure and turned over control to a human operator or the operator took over control to ensure safe operation of the vehicle. There are limitations to using disengagements as a metric for safety because it is unclear what standards of risk are used and whether they are homogenous across all test drivers or developers.

existing vehicles (Tesla, undated). Tesla calls this "fleet learning"—that is, when an entire fleet is able to learn from the experiences of a few vehicles.

With HAVs, as with many other learning situations, there may be rapid gains in performance in the earlier miles driven, when there is significant room for improvement, than in later miles driven, when the remaining areas of improvement are harder to identify and achieve. Power and exponential functions are often used to model this kind of learning (rapid initially, slow later) (Newell and Rosenbloom, 1980) and have been widely used to model technological progress (Nagy et al., 2013). Here, we model the state of the art safety rate *R(t)* after *t* miles have been cumulatively driven by HAVs as an exponential decay curve. We use decay because the fatality rate drops from some initial rate at the time of introduction to some future rate.

A standard exponential decay curve is described by Equation 4.7:

$$N(t) = e^{-t} \; . \tag{4.7}$$

We can model *R(t)* on the standard exponential decay function first by scaling along the x-axis to describe the cumulative miles driven and along the y-axis to describe performance in terms of the state of the art in HAV safety rates relative to an initial benchmark safety rate of non-HAVs. Modeling *R(t)* in this way requires defining three parameters:

- $R_I$, the initial safety rate of HAVs relative to the benchmark rate of non-HAVs[13]
- $R_F$, the best safety rate of HAVs relative to the benchmark rate of non-HAVs
- $b$, the percentage of improvement achieved after driving $t_b$ miles.

This gives us the following function for $R(t)$:

$$R(t) = R_F + (R_I\text{-}R_F)e^{-t\left(\frac{-\ln(1-b)}{t_b}\right)} \; . \tag{4.8}$$

In this report, we define the *benchmark* fatality rate as the current rate of human drivers in the United States (1.12 fatalities per 100 million miles in 2015; see NHTSA, 2016a).[14] Figure 4.3 shows the fatality rate of HAVs relative to the benchmark rate of non-HAVs, assuming that the initial rate is 1.5 times the benchmark, the final rate is 0.5 times the benchmark, and it takes 10 trillion miles to reach 99 percent of the improvement from 0.9 to 0.25. Here, $b = 0.99, t_b = 10^{13}, R_I = 0.9, and\ R_F = 0.25$.

[13] This is a unitless quantity and describes the percentage of safety incidents (e.g., fatalities, injuries, or crashes) that occur per mile traveled by HAVs versus non-HAVs.

[14] In October 2017, just prior to this report's publication, NHTSA released traffic safety data for 2016 and revised estimates for 2015. NHTSA reports that that, in 2016, the fatality rate increased still further to 1.18 fatalities per 100 million VMT from a (revised) rate of 1.15 fatalities per 100 million VMT in 2015 (NHTSA, 2017). The figures in this report are based on the earlier 2015 estimate.

**Figure 4.3. Scenario of the Relative Fatality Rate of HAVs as a Function of Cumulative Miles Driven by HAVs**

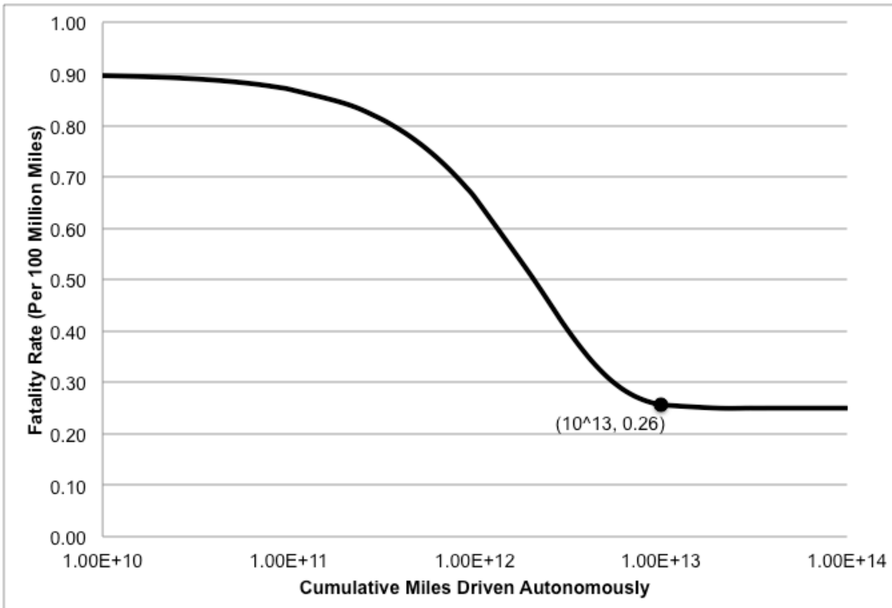


NOTE: This figure assumes that the initial HAV fatality rate is 0.9 times the fatality rate of non-HAVs (1.12 fatalities per 100 million miles), the final rate is 0.25 times that of the benchmark rate, and it takes 10 trillion miles of driving to achieve 99 percent of the improvement from 0.9 to 0.25.

Figure 4.4 shows the relative fatality rate over time, taking into account the number of highly automated miles cumulatively driven, based on the annual highly automated VMT shown in Figure 4.2. It shows how 99 percent of the performance improvement (i.e., a relative fatality rate of 0.26) is achieved in 2051 after approximately 10 trillion miles have been driven by HAVs.

**Figure 4.4. Scenario of the Relative Fatality Rate of HAVs and Cumulative Miles Driven by HAVs over Time**

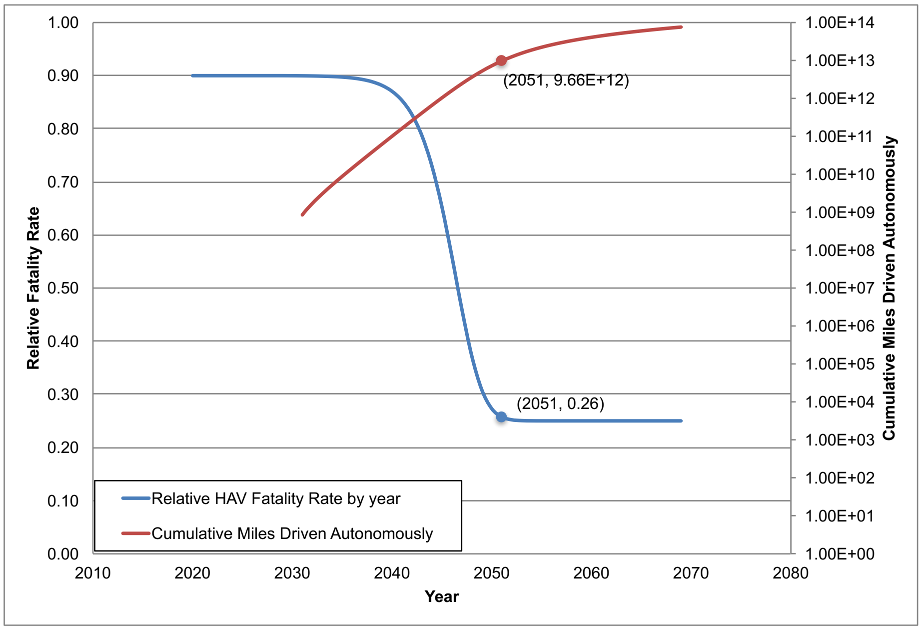


NOTE: This figure assumes all of the parameters used in prior figures.

## Diffusion of the State of the Art in Highly Automated Vehicle Safety

One of the factors affecting future safety is whether the safety performance of existing HAVs will improve as the state of the art for HAVs improves. For traditional automobiles, improving the safety of vehicles that have already been deployed is difficult, because safety improvements typically involve improvements to the vehicles' physical design. These improvements are generally made to new models but not to the existing fleet, so improvements are realized principally when consumers purchase new vehicles and retire old ones.[15] In contrast, safety improvements can be made across large portions of the existing HAV fleet not only through vehicle replacement but also through software performance upgrades that can occur remotely and therefore often. Already, some automakers have begun providing over-the-air updates for both ADASs and HAVs (Zhang, 2016), and this is expected to become a standard feature in the future (ABI Research, 2016).[16]

---

[15] An exception is when there are recalls in which the safety of significant portions of the conventional fleet are improved (e.g., the Takata airbag recall). Recalls generally eliminate flaws in the basic function of the vehicle and generally improve fleet safety only very slightly, because the risks involved in recalls are generally very small. For example, the faulty Takata airbags are present in more than 40 million cars and, as of January 4, 2017, have led to 11 deaths (Consumer Reports, 2017).

[16] We have not incorporated learning of ADAS in modeling safety performance of non-HAVs (i.e., Level 1 and Level 2 autonomy) because, in many cases, these safety technologies can be adequately developed in laboratory settings, so the opportunity for and impact of software-only safety improvements may be less.

We can explore the impact of safety performance improvements (through both vehicle turnover and upgrades) in the HAV fleet by calculating two bookend cases. In the best case, the entire HAV fleet is upgraded to have state-of-the-art safety performance. This could occur if vehicles are replaced often, are improved continuously during their lives, or both. In the worst case, we assume no safety improvements because once an HAV is on the road, it stays on the road indefinitely and its safety never improves.[17] The future is somewhere between these bookends, and we can explore the impacts of the level of fleet improvement by interpolating between them.

Consistent with prior equations, we use the following variables:

- $VMT_{a,x}$, the miles driven by HAVs in year $x$, as defined in Equation 3.1
- $H$, the benchmark non-HAV safety rate
- $R_x$, the state-of-the-art HAV safety rate of vehicles released in year $x$, relative to the benchmark safety rate (this is $\mathrm{R}(\sum_{\mathrm{i}=0}^{\mathrm{x}} \mathrm{M_i})$ in Equation 4.8)
- $f_{a,x}$, the total safety incidents in year $x$ from HAVs.

With these variables, the number of incidents caused by HAVs when assuming *perfect* improvement of the existing fleet is

$$f_{a,x,perfect} = HR_x M_x \ . \tag{4.9}$$

The total number of incidents with no improvement to the existing fleet is the number of incidents in the prior year ($F_{x-1}$) plus the product of the additional miles traveled in the current year that are attributed to new vehicle models ($VMT_{a,x} - VMT_{a,x-1}$) and the state-of-the-art safety rate of the new vehicle models ($HR_x$) (Equation 4.11):

$$F_{a,2020,none} = HR_{2020} M_{2020} \tag{4.10}$$
$$F_{a,x,none} = F_{x-1} + (VMT_{a,x} - VMT_{a,x-1}) HR_x \ . \tag{4.11}$$

Figure 4.5 shows the HAV fatalities with perfect fleet improvement and no fleet improvement, under the same scenario assumptions in prior figures. In the case with perfect fleet improvement, HAVs cause approximately 10,000 annual fatalities by year 50. In the case with no fleet improvement, there are approximately 13,800 fatalities caused by HAVs in year 50, or about 3,800 more fatalities annually. Over the 50 simulated years, the difference in these scenarios translates to approximately 83,000 fewer fatalities in the case with perfect fleet improvement.

---

[17] This scenario of no fleet improvement is, in practice, not possible, because HAVs will be retired and replaced with new vehicle models even if they are never upgraded during their time on the road. Therefore, this case is best thought of as a useful computational bookend rather than a bookend of the actual future of fleet improvement.

**Figure 4.5. HAV Fatalities with Perfect and with No Fleet Improvement**

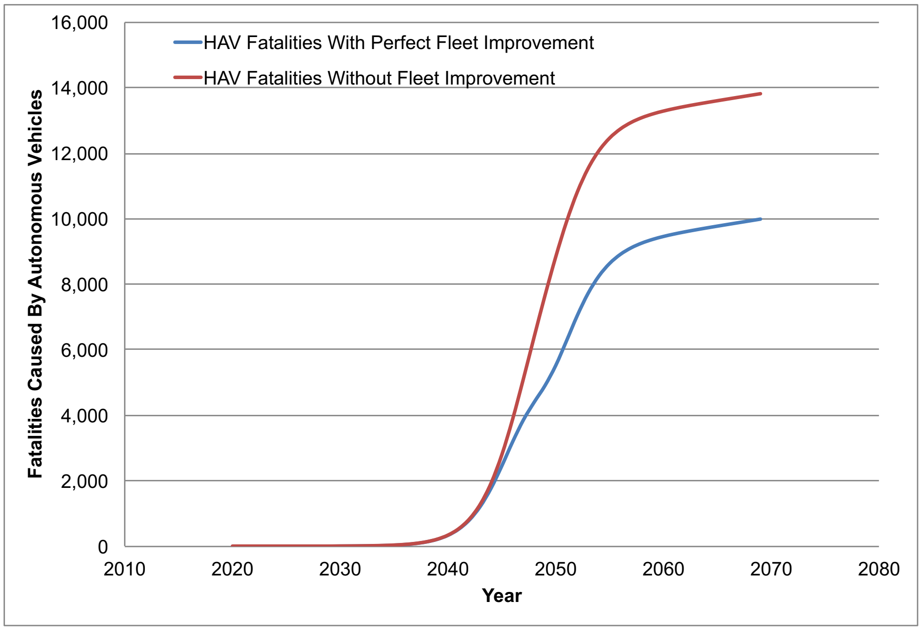


NOTE: This figure assumes all of the parameters used in prior figures.

The real fleet improvement will be within these bounds. The fleet will be improved somewhat because of software upgrades and turnover even among HAVs. However, older HAVs cannot indefinitely keep up with the latest models because, for example, some improvements will involve changes to the hardware that cannot be retrofitted. This middle ground is uncertain but can be explored by weighting the safety incidents in a given year between these extremes by some weight *w* between 0 and 1, as shown in Equation 4.12:

$$F_{a,x} = wF_{a,x,perfect} + (1-w)F_{a,x,none}\ . \tag{4.12}$$

The safety rate in a future with HAVs can thus be numerically determined as

$$s_{a.x} = \frac{F_{a.x}}{VMT_{a,x}}\ . \tag{4.13}$$

Figure 4.6 shows the total fatalities in a future with HAVs, assuming a fleet improvement weight *w* of 0.2.

**Figure 4.6. Total Fatalities in a Future with HAVs**

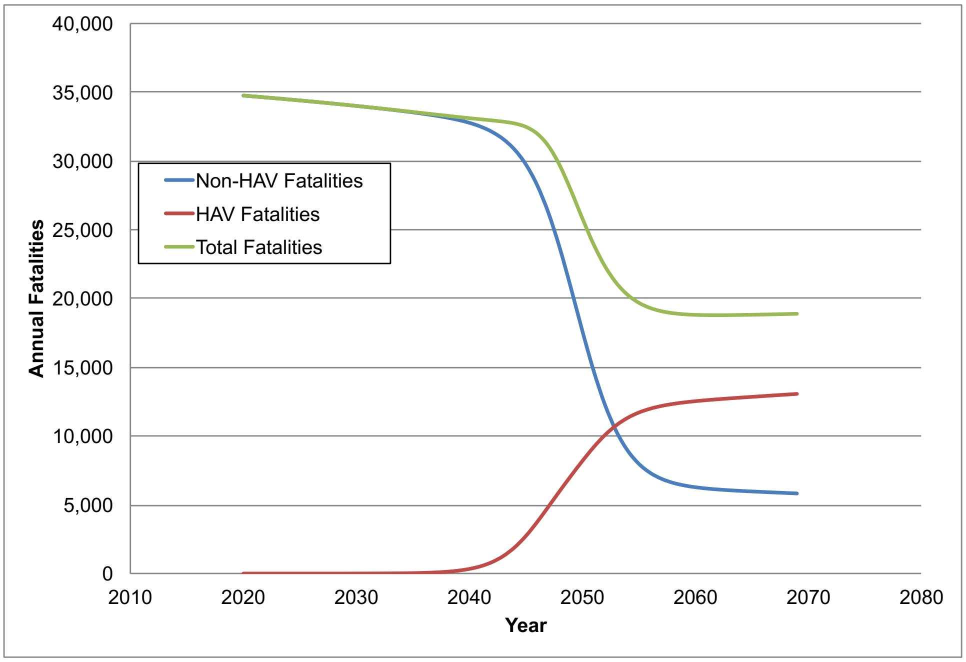


NOTE: This figure assumes all of the parameters used in prior figures.

## Chapter Five

# Model Exploration and Illustration

MAVS allows users to see the potentially diverse effects of diffusion, learning, and other characteristics of vehicle safety over time. To illustrate, we explore fatalities in a future without HAVs and in three different scenarios of HAV diffusion and performance, as defined by parameters in Table 5.1. These scenarios are strictly meant to demonstrate how the model can be used to compare assumptions and should not be interpreted as our predictions of how the future of road safety will unfold.

**Table 5.1. Parameters in Three HAV Scenarios**

| User Input | Variable | Scenario 1: Worse Than Benchmark | Scenario 2: Just Better Than Benchmark | Scenario 3: Nearly Perfect |
|---|---|---|---|---|
| In what year is 0.01% of diffusion achieved? | $x_{0.02}$ | 2025 | 2035 | 2045 |
| In what year is 99.99% of diffusion achieved? | $x_{0.98}$ | 2055 | 2065 | 2075 |
| In a future with HAVs, what is the maximum percentage of baseline miles that will be driven by HAVs? | $A_{max}$ | 80% | 80% | 80% |
| How much will HAV VMT change as a result of HAV use? | $c$ | 10% | 10% | 10% |
| What is the initial HAV fatality rate relative to benchmark? | $R_i$ | 1.2 | 0.8 | 0.2 |
| What is the final (i.e., best) HAV fatality rate relative to benchmark? | $R_f$ | 0.2 | 0.2 | 0.2 |
| How many miles would it take to reach a 99% reduction in the rate? | $t_{0.99}$ | 1 trillion | 1 trillion | 1 trillion |
| How upgradeable is the existing HAV fleet? | $w$ | 0.5 | 0.5 | 0.5 |

The baseline future without HAVs reflects a compounding annual growth rate of 0.6 percent with a 30-percent decline in fatality rate relative to benchmark (i.e., a fatality rate of 0.784 per 100 million miles). The first HAV scenario describes a future in which HAVs are adopted early, before they reach performance parity with non-HAVs. The second scenario describes a future in which HAVs are adopted in the medium term, once their performance is just better than current non-HAVs. The third scenario describes a future in which HAVs are adopted still later, only once they are significantly safer than current non-HAVs. These scenarios differ in the timing of diffusion and the performance at the time of diffusion, but we have kept as many other parameters constant across the scenarios, such as the increase in VMT as a result of HAV use and the upgradeability of the fleet.

These scenarios could be the result of different regulations (e.g., government readiness to allow HAVs at different levels of safety performance), consumer preferences (e.g., consumer willingness to use HAVs with different levels of safety performance), or other factors (e.g., developer willingness to introduce technology at different performance levels). The aim in this illustration is not to assess how the scenarios came to be but to assess how MAVS calculates the effect on road safety outcomes should the scenarios occur.

Figures 5.1, 5.2, and 5.3 show how the model can be used to understand and compare outcomes in different scenarios. Figure 5.1 shows the differences in the timing of the diffusion of HAVs among the three HAV scenarios, shown as the annual highly automated VMT over time. Figure 5.2 shows the differences in the state of the art in HAV fatality rates among the three scenarios, described as the fatality rate of HAVs relative to the benchmark human-driver fatality rate (1.12 fatalities per 100 million miles). These figures show that even though the initial safety of HAVs is poor in Scenario 1, earlier diffusion allows safety to improve rapidly and outpace safety performance in Scenarios 2 and 3.

Figure 5.3 shows that Scenarios 2 and 3 both reduce fatalities every year relative to the baseline future and, over a 50-year period, achieve a net reduction in fatalities of 390,000 and 190,000, respectively. That Scenario 2 saves more lives than Scenario 3 is not surprising, because the savings are assumed to begin much earlier, thereby allowing widespread diffusion of life-saving technology to occur sooner. In contrast, Scenario 1 increases the lives lost in early years, albeit by small amounts. Yet, over this 50-year period, this scenario saves the most lives overall—a total of 596,000—because, through learning, it achieves better performance sooner than in the other scenarios.[18]

This illustration is not to suggest that early diffusion of lower-performing technology will unfold as it has in this scenario, or that any of these scenarios will indeed occur. As just one example, public backlash in the event of a high-profile crash could stall the adoption of the technology. The purpose here is instead to show how the future may unfold in complex and important ways that are not obvious when attention is paid only to safety at the time of HAV introduction. Such information might raise different policy questions, such as how policies can be crafted to encourage rapid learning or how trade-offs should be made between small short-term risks that may lead to large longer-term gains.

[18] In the final few years, however, Scenario 1 again has slightly higher annual fatalities than the other two scenarios. This is because, although both VMT and the state of the art in HAV safety performance are the same in each case, the fleet is not perfectly improved, and some older HAVs with poorer performance remain on the road.

**Figure 5.1. Highly Automated VMT in Three Scenarios**

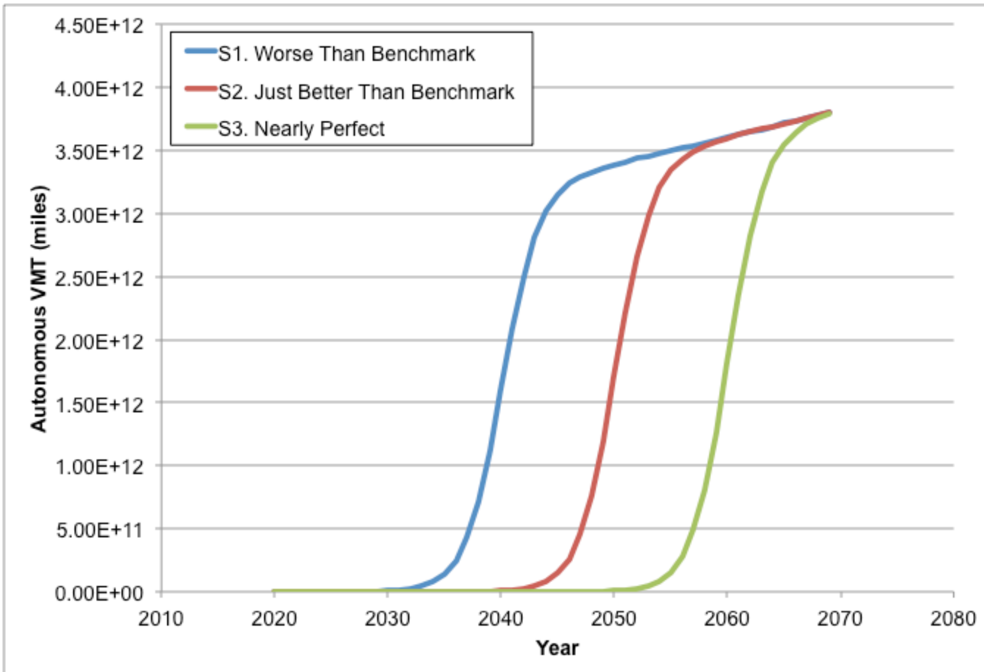


**Figure 5.2. State-of-the-Art Fatality Rate of HAVs Relative to the Benchmark Rate in Three Scenarios**

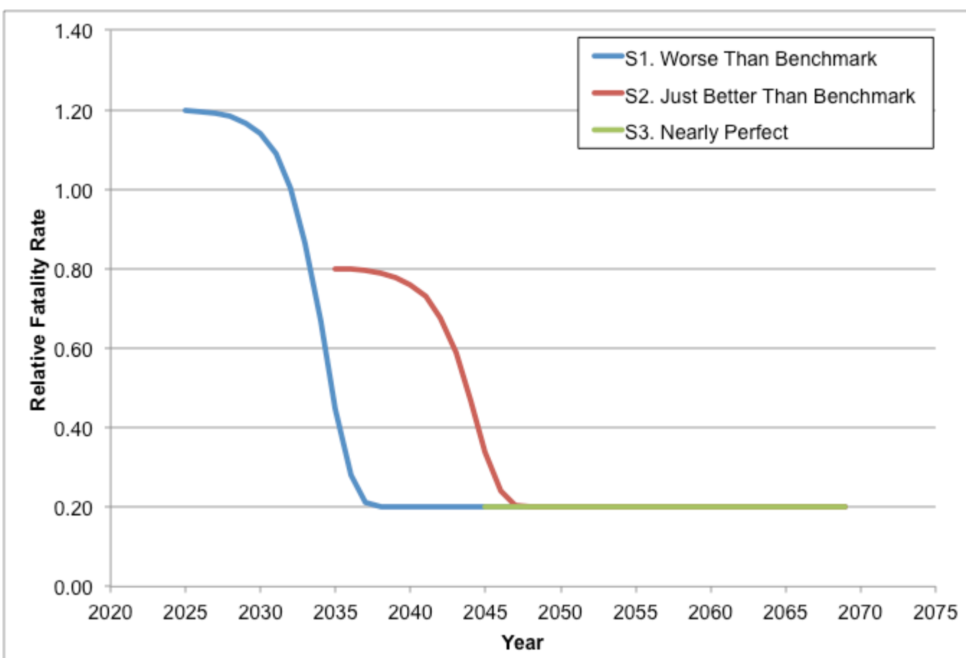

**Figure 5.3. Annual Fatalities in a Baseline Future Without HAVs and in Three Scenarios With HAVs**

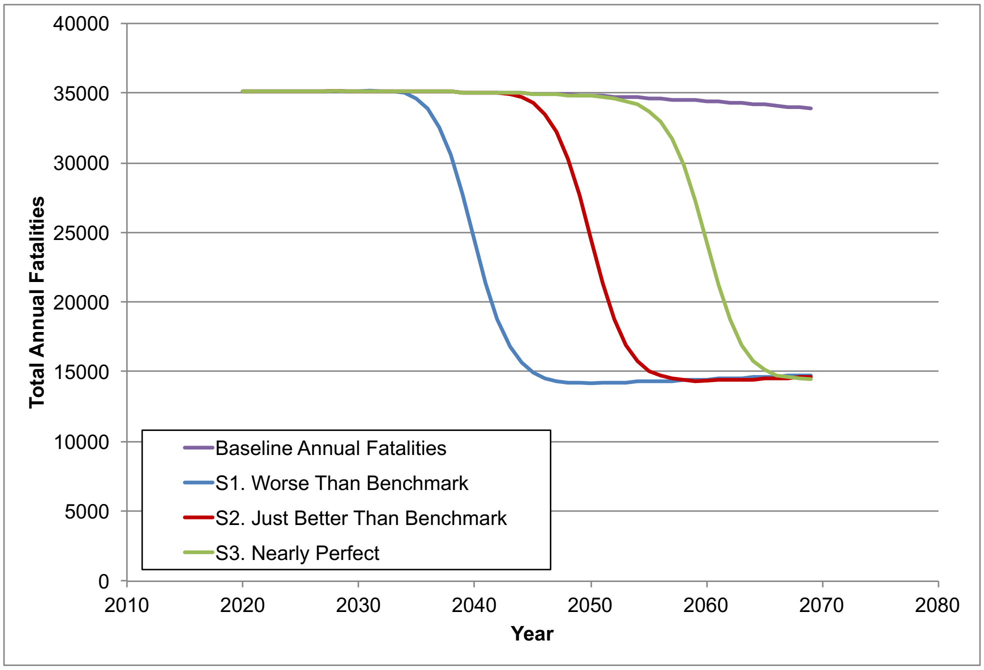

Chapter Six

# Model Summary and Conclusions

This report documents a simple model of the impact of HAV diffusion and improvement on road safety over time compared with a future without HAVs based on user assumptions about a variety of key factors. MAVS is agnostic about how road safety is defined, and it can be configured to estimate fatalities, injuries, crashes, or other road safety metrics. MAVS does not include all of the complex factors and relationships that would define HAV safety. Rather, it is intended as a foundation for exploration that enables users to expand their understanding of HAV safety from introduction and over time.

The model takes the following user-specified inputs, which are uncertain and can be shaped by social, economic, technology, and policy factors:

- **The growth in VMT in a baseline future without HAVs.** This input can be shaped by economic growth, demographic changes, and other socioeconomic factors. It can also be shaped by policies, such as taxes on fuel or the presence or absence of policies to manage transportation demand.
- **The change in safety rate of non-HAVs.** This input can be shaped by innovation in ADASs and other vehicle technologies; transportation system changes, such as maintenance of roads; and, most of all, behavior changes, such as changes in seat belt use, distracted driving, or driving while intoxicated.
- **The timing and extent of HAV diffusion.** This input will be shaped by regulatory choices, such as whether to have safety performance standards for HAVs and what those performance standards are; public preferences and appetite for the technology; and developers' willingness to introduce the technologies.
- **The impact of HAVs on transportation demand.** This input is generally expected to grow because HAVs reduce the cost of driving and create new opportunities for transportation among those who currently lack them. However, this too can be shaped by policies, such as those that manage transportation demand or that connect HAVs with transit, as well as the extent to which mobility is provided as a service in a shared economy.
- **The safety characteristics of HAVs.** This input includes the safety of HAVs at the time of introduction and the time of maximum diffusion, as well as the rate of learning that enables a transition between these endpoints. This will be a function not only of technology but also of policies that encourage or discourage data-sharing that could increase learning rates.
- **The upgradeability of HAVs.** This input will be shaped by the extent to which improvements reside in hardware versus software and by policy requirements (if any) for proving performance prior to upgrades.

With these inputs, MAVS calculates changes in VMT and safety rates over time, producing timelines of annual safety incidents (e.g., fatalities or injuries) and cumulative safety incidents under different future scenarios. The results can help inform the public and the public policy debate about how HAVs should be introduced, improved, and adopted.

**Appendix A**

# Modeling Diffusion with a Logistic Function

This appendix describes how the standard logistic function is used to model the diffusion of HAVs over time, as presented in Chapter Two. The standard logistic function is described by Equation 3.1 and has a range of $(-\infty, 0)$ to $(\infty, 1)$. It must be transformed on the horizontal axis to reflect time, defined as years from today, and on the vertical axis to reflect diffusion, defined as vehicle miles driven by HAVs.

We fist transform the horizontal axis. We choose any points on our diffusion function ($x_a$, $a$) and ($x_{1–a}$, $1–a$), where $x_a$ is the user-chosen year in which some percentage $a$ of total HAV diffusion has occurred, and $x_{1–a}$ is the user-chosen year in which some percentage ($1–a$) of total HAV diffusion has occurred, for any $0 < a < 1$.[19] The points ($x_a$, $a$) and ($x_{1–a}$, $1–a$) can be linearly transformed to ($x'_a$, $a$) and ($x'_{1–a}$, $1–a$) on the standard logistic function by solving two equations:

$$x'_a = mx_a + b \tag{A.1}$$

$$x'_{1-a} = mx_{1-a} + b\ . \tag{A.2}$$

Solving for $m$ and $b$ gives

$$m = \left(\frac{x'_{1-a} - x'_a}{x_{1-a} - x_a}\right) \tag{A.3}$$

$$b = x'_a - mx_a\ . \tag{A.4}$$

We can further solve for $x'_a$ and $x'_{1–a}$ by taking the inverse of the logistic function:

$$x = -\ln\left(\frac{1 - f(x)}{f(x)}\right)\ . \tag{A.5}$$

Because of the symmetry of the standard logistic function, this simplifies to

$$m = \frac{2\ln\left(\frac{1-a}{a}\right)}{x_{1-a} - x_a} \tag{A.6}$$

$$b = -\ln\left(\frac{1-a}{a}\right) - mx_a\ . \tag{A.7}$$

[19] We cannot choose $a = 0$ or $a = 1$ because the standard logistic function has horizontal asymptotes at 0 and 1.

We next transform the vertical axis. The axis must be scaled (but not shifted) from the range $0 < f(x) < 1$ to $0 < f(x) < A_{max}$, where $A_{max}$ is the largest percentage of baseline VMT that will ever be driven by HAVs in a single year.

The resulting function of the percentage of baseline miles that will be replaced by highly automated driving in year $x$ is defined by Equation A.8:

$$f(x) = \frac{A_{max}}{1+e^{-(mx+b)}} \,. \tag{A.8}$$

## Appendix B

# Modeling Learning with an Exponential Decay Function

This appendix describes how the exponential decay function is used to model the state of the art in HAV safety as a function of miles of highly automated travel, as presented in Chapter Four. The exponential decay function is described by Equation 4.7 and has a range of (0, 1) to $(\infty, 0)$. It must be transformed on the horizontal axis to reflect experience, defined in miles of highly automated travel, and on the vertical axis to reflect performance, defined as the safety rate relative to human drivers.

We fist transform the horizontal axis. We choose a point on our learning curve ($t_b$, $b$), where $t_b$ is the user-chosen number of miles after which some percentage $b$ of the performance improvement remains (percentage $1–b$ of improvement has already occurred), for any $0 < b < 1$.[20] The point ($t_b$, $b$) on the HAV learning curve must be mapped to the point ($t'b$, $b$) on the standard logistic function. This requires scaling only because the lower end of the range is the same in both functions ($t_0 = t'_0 = 0$).

That scaling is achieved by dividing by $t'_b/t_b$, as shown in Equation B.1:

$$N(t) = e^{-t\frac{t_b'}{t_b}} \tag{B.1}$$

We can further solve for $t'_b$ by taking the inverse of the exponential decay function:

$$x = -\ln\,(f(x)), \tag{B.2}$$

which means

$$N(t) = e^{-t\frac{-\ln\,(b)}{t_b}}. \tag{B.3}$$

We next transform the vertical axis with two parameters:

- $R_I$, the initial relative safety rate of HAVs
- $R_F$, the best expected safety rate of HAVs.

[20] We cannot choose $b = 1$, because the standard logistic function has a horizontal asymptote at 1. It is also not meaningful to choose $b = 0$, because no transformation is required at that point.

To scale, we simply multiply the numerator in the logistic function by $R_I - R_F$ and offset by $R_F$. Equation B.4 defines the resulting function of the safety rate of HAVs relative to the initial human-driver safety rate after $t$ miles of highly automated travel:

$$R(t) = R_F + (R_I - R_F)e^{-t\left(\frac{-\ln(b)}{t_b}\right)}\,. \tag{B.4}$$

# References

ABI Research, "ABI Research Anticipates Accelerated Diffusion of Automotive Software Over-the-Air Updates with Nearly 180 Million New SOTA-Enabled Cars Shipping Between 2016 and 2022," press release, March 15, 2016. As of September 8, 2017: https://www.abiresearch.com/press/abi-research-anticipates-accelerated-adoption-auto/

Anderson, James M., Nidhi Kalra, Karlyn D. Stanley, Paul Sorensen, Constantine Samaras, and Oluwatobi A. Oluwatola, *Autonomous Vehicle Technology: A Guide for Policymakers*, Santa Monica, Calif.: RAND Corporation, RR-443-2-RC, 2016. As of January 24, 2016: http://www.rand.org/pubs/research_reports/RR443-2.html

Bansal, Prateek, and Kara M. Kockelman, "Forecasting Americans' Long-Term Diffusion of Connected and Autonomous Vehicle Technologies," *Transportation Research Part A: Policy and Practice,* Vol. 95, January 2017, pp. 49–63.

Bierstedt, Jane, Aaron Gooze, Chris Gray, Josh Peterman, Leon Raykin, and Jerry Walters, *Effects of Next-Generation Vehicles on Travel Demand and Highway Capacity*, Denver, Colo.: Fehr and Peers, January 2014. As of January 17, 2017: http://orfe.princeton.edu/~alaink/Papers/FP_NextGenVehicleWhitePaper012414.pdf

Blincoe, Lawrence, Ted R. Miller, Eduard Zaloshnja, and Bruce A. Lawrence, *The Economic and Societal Impact of Motor Vehicle Crashes 2010 (Revised)*, Washington, D.C.: U.S. Department of Transportation, DOT HS 812 013, May 2015. As of March 3, 2016: http://www-nrd.nhtsa.dot.gov/pubs/812013.pdf

Bureau of Transportation Statistics, "Table 2-17: Motor Vehicle Safety Data," in *National Transportation Statistics*, Washington, D.C.: U.S. Department of Transportation, April 2016. As of January 31, 2017: https://www.rita.dot.gov/bts/sites/rita.dot.gov.bts/files/publications/national_transportation_statistics/html/table_02_17.html

Casualty Actuarial Society, *Restating the National Highway Transportation Safety Administration's National Motor Vehicle Crash Causation Survey for Automated Vehicles*, Casualty Actuarial Society E-Forum, Vol. 1, Fall 2014. As of September 8, 2017: https://www.casact.org/pubs/forum/14fforum/CAS%20AVTF_Restated_NMVCCS.pdf

Collins, Susan M., "Chairman Susan M. Collins Opening Statement," hearing on "The Automated and Self-Driving Vehicle Revolution: What Is the Role of Government?" U.S. Senate, Committee on Appropriations, Subcommittee on Transportation, Housing and Urban Development, and Related Agencies, November 16, 2016. As of September 8, 2017:

https://www.appropriations.senate.gov/imo/media/doc/111616-Chairman-Collins-Opening-Statement-21.pdf

Consumer Reports, "Takata Airbag Recall—Everything You Need to Know," 2017. As of January 23, 2017: http://www.consumerreports.org/cro/news/2016/05/everything-you-need-to-know-about-the-takata-air-bag-recall/index.htm

Davies, Alex, "Ford's Skipping the Trickiest Thing About Self-Driving Cars," *Wired*, November 10, 2015. As of January 17, 2017: https://www.wired.com/2015/11/ford-self-driving-car-plan-google/

Eccles, K., F. Gross, M. Liu, and F. Council, *Crash Data Analyses for Vehicle-to-Infrastructure Communications for Safety Applications*, Washington, D.C.: Federal Highway Administration, DTFH61-06-C-00013, 2012.

Fagnant, Daniel J., and Kara Kockelman, "Preparing a Nation for Autonomous Vehicles: Opportunities, Barriers and Policy Recommendations," *Transportation Research Part A: Policy and Practice*, Vol. 77, July 2015, pp. 167–181.

Federal Highway Administration, *FHWA Forecasts of Vehicle Miles Traveled (VMT): Spring 2016*, Washington, D.C.: U.S. Department of Transportation, May 2, 2016.

Funke, James, Gowrishankar Srinivasan, Raja Ranganathan, and August Burgett, *Safety Impact Methodology (SIM): Application and Results of the Advanced Crash Avoidance Technologies (ACAT) Program*, Paper No. 11-0367, 2011.

Google Auto LLC, *Google Self-Driving Car Testing Report on Disengagements of Autonomous Mode*, Menlo Park, Calif., December 2015. As of September 8, 2017: https://www.dmv.ca.gov/portal/wcm/connect/dff67186-70dd-4042-bc8c-d7b2a9904665/google_disengagement_report.pdf?MOD=AJPERES

Gordon, T., H. Sardar, D. Blower, M. Ljung Aust, Z. Bareket, M. Barnes, A. Blankespoor, I. Isaksson-Hellman, J. Ivarsson, B. Juhas, K. Nobukawa, and H. Theander, *Advanced Crash Avoidance Technologies (ACAT) Program: Final Report of the Volvo- Ford- UMTRI Project: Safety Impact Methodology for Lane Departure Warning—Method Development and Estimation of Benefits*, Washington, D.C.: U.S. Department of Transportation, DOT HS 811 405, October 2010.

Greenblatt, Jeffrey B., and Samveg Saxena, "Autonomous Taxis Could Greatly Reduce Greenhouse-Gas Emissions of U.S. Light-Duty Vehicles," *Nature Climate Change*, Vol. 5, 2015, pp. 860–863.

Harper, Corey D., Chris T. Hendrickson, and Constantine Samaras, "Cost and Benefit Estimates of Partially-Automated Vehicle Collision Avoidance Technologies," *Accident Analysis and Prevention*, Vol. 95, Part A, October 2016, pp. 104–115.

Hayes, Brian, "Leave the Driving to It," *American Science*, Vol. 99, No. 5, October 2011, pp. 362–366.

Houston, David J., and Lilliard E. Richardson, Jr., "The Politics of Air Bag Safety: A Competition Among Problem Definitions," *Policy Studies Journal*, Vol. 28, No. 3, August 2000, pp. 485–501.

Jermakian, Jessica S., "Crash Avoidance Potential of Four Passenger Vehicle Technologies," *Accident Analysis and Prevention*, Vol. 43, No. 3, May 2011, pp. 732–740.

Jeyaraj, A., and R. Sabherwal, "Diffusion of Information Technology Innovations," *Wiley Encyclopedia of Management*, Vol. 7, 2015, pp. 1–7.

Kalra, Nidhi, *Challenges and Approaches to Realizing Autonomous Vehicle Safety*, testimony submitted to the House Energy and Commerce Committee, Subcommittee on Digital Commerce and Consumer Protection, Santa Monica, Calif.: RAND Corporation, CT-463, February 14, 2017. As of September 8, 2017: https://www.rand.org/pubs/testimonies/CT463.html

Kalra, Nidhi, and David G. Groves, *The Enemy of Good: Estimating the Cost of Waiting for Nearly Perfect Automated Vehicles*, Santa Monica, Calif.: RAND Corporation, RR-2150-RC, 2017. As of November 3, 2017: https://www.rand.org/pubs/research_reports/RR2150.html

Kalra, Nidhi, and Susan M. Paddock, *Driving to Safety: How Many Miles of Driving Would It Take to Demonstrate Autonomous Vehicle Reliability?* Santa Monica, Calif.: RAND Corporation, RR-1478-RC, 2016. As of September 8, 2017: http://www.rand.org/pubs/research_reports/RR1478.html

Karshenas, Massoud, and Paul L. Stoneman, "Rank, Stock, Order, and Epidemic Effects in the Diffusion of New Process Technologies: An Empirical Model," *RAND Journal of Economics*, Vol. 24, No. 4, 1993, pp. 503–528. As of September 8, 2017: www.jstor.org/stable/2555742

Li, Tianxin, and Kara M. Kockelman, *Valuing the Safety Benefits of Connected and Automated Vehicle Technologies*, Transportation Research Board 95th Annual Meeting, 2016.

Litman, Todd, *Autonomous Vehicle Implementation Predictions: Implications for Transport Planning*, Victoria, British Columbia, Canada: Victoria Transport Policy Institute, January 2, 2017.

Milakis, Dimitris, Bart van Arem, and Bert van Wee, "Policy and Society Related Implications of Automated Driving: A Review of Literature and Directions for Future Research," *Journal of Intelligent Transportation Systems*, Vol. 21, No. 4, 2017, pp. 324–348.

Musk, Elon, "Tesla Press Conference for the Autopilot v7.0 Software," October 14, 2015. As of December 15, 2016:
https://www.youtube.com/watch?v=73_Qjez1MbI

Nagy, Béla, J. Doyne Farmer, Quan M. Bui, and Jessika E. Trancik, "Statistical Basis for Predicting Technological Progress," *PLoS ONE*, Vol. 8, No. 2, e52669, 2013.

Najm, Wassim G., Samuel Toma, and John Brewer, *Depiction of Priority Light-Vehicle Pre-Crash Scenarios for Safety Applications Based on Vehicle-to-Vehicle Communications*, Washington, D.C.: U.S. Department of Transportation, DOT HS 811 732, 2013. As of September 8, 2017:
https://ntl.bts.gov/lib/47000/47400/47497/DOT-VNTSC-NHTSA-11-12.pdf

National Highway Traffic Safety Administration, *National Motor Vehicle Crash Causation Survey*, Washington, D.C.: U.S. Department of Transportation, DOT HS 811 059, 2008.

———, *2015 Motor Vehicle Crashes: Overview*, Washington, D.C.: U.S. Department of Transportation, DOT HS 812 318, August 2016a. As of January 18, 2017:
https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/812318

———, *Federal Automated Vehicles Policy: Accelerating the Next Revolution in Roadway Safety*, Washington, D.C.: U.S. Department of Transportation, September 2016b. As of January 17, 2017:
https://www.transportation.gov/sites/dot.gov/files/docs/AV%20policy%20guidance%20PDF.pdf

———, *2016 Fatal Motor Vehicle Crashes: Overview*, Washington, D.C.: U.S. Department of Transportation, DOT HS 812 456, October 2017. As of October 25, 2017:
https://crashstats.nhtsa.dot.gov/Api/Public/Publication/812456

Newell, Allen, and Paul S. Rosenbloom, "Mechanisms of Skill Acquisition and the Law of Practice," *Cognitive Skills and Their Acquisition*, Vol. 1, 1980, pp. 1–55.

NHTSA—*See* National Highway Traffic Safety Administration.

Nishimoto, Alex, "All New Tesla Models Will Feature Level 5-Capable Autopilot Hardware," *Motortrend*, October 19, 2016. As of January 17, 2017:
http://www.motortrend.com/news/new-tesla-models-will-feature-level-5-capable-autopilot-hardware

Oremus, Will, "Is Autopilot a Bad Idea? Why Ford, Google, Volvo and Others Think Tesla Is Wrong About Automation," *Slate*, July 6, 2016. As of January 2017:

http://www.slate.com/articles/technology/future_tense/2016/07/is_tesla_s_style_of_autopilot_a_bad_idea_volvo_google_and_others_think_so.html

Perez, M., L. S. Angell, J. Hankey, R. K. Deering, R. E. Llaneras, C. A. Green, M. L. Neurauter, and J. F. Antin, *Advanced Crash Avoidance Technologies (ACAT) Program—Final Report of the GM-VTTI Backing Crash Countermeasures Project*, Washington, D.C.: U.S. Department of Transportation, DOT HS 811 452, August 2011.

Rau, Paul, Mikio Yanagisawa, and Wassim G. Najm, *Target Crash Population of Automated Vehicles*, Paper No. 15-0430, 2015.

SAE International, *Surface Vehicle Recommended Practice: Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles*, Warrendale, Pa., J2016, September 2016.

Tesla, "Autopilot," web page, undated. As of September 8, 2017: https://www.tesla.com/autopilot

Trimble, T. E., R. Bishop, J. F. Morgan, and M. Blanco, *Human Factors Evaluation of Level 2 and Level 3 Automated Driving Concepts: Past Research, State of Automation Technology, and Emerging System Concepts*, Washington, D.C.: U.S. Department of Transportation, DOT HS 812 043, July 2014. As of September 8, 2017: http://ntl.bts.gov/lib/45000/45400/45476/812043_HF-EvaluationLevel2andLevel3AutomatedDrivingConceptsV2_1_.pdf

U.S. Department of Transportation, "U.S. DOT and IIHS Announce Historic Commitment of 20 Automakers to Make Automatic Emergency Braking Standard on New Vehicles," press release, March 17, 2016. As of January 17, 2017: https://www.nhtsa.gov/press-releases/us-dot-and-iihs-announce-historic-commitment-20-automakers-make-automatic-emergency

Walden, Greg, "Opening Statement of Chairman Greg Walden," hearing on "Self-Driving Cars: Road to Deployment," U.S. House of Representatives, Energy and Commerce Committee, Subcommittee on Digital Commerce and Consumer Protection, February 14, 2017. As of September 8, 2017: http://docs.house.gov/meetings/IF/IF17/20170214/105548/HHRG-115-IF17-MState-W000791-20170214.pdf

Waymo, *Report on Autonomous Mode Disengagements for Waymo Self-Driving Vehicles in California*, Mountain View, Calif., December 2016. https://www.dmv.ca.gov/portal/wcm/connect/946b3502-c959-4e3b-b119-91319c27788f/GoogleAutoWaymo_disengage_report_2016.pdf?MOD=AJPERES

Winkle, Thomas, “Safety Benefits of Automated Vehicles: Extended Findings from Accident Research for Development, Validation and Testing in Autonomous Driving,” in Markus Maurer, J. Christian Gerdes, Barbara Lenz, and Hermann Winner, eds., *Autonomous Driving: Technical, Legal and Social Aspects*, London: SpringerOpen, 2015, pp 335–364. https://link.springer.com/chapter/10.1007/978-3-662-48847-8_17/fulltext.html

Zhang, Benjamin, “Tesla Wants to Take Self-Driving Cars to a Whole New Level,” *Business Insider*, January 11, 2016. As of December 15, 2016: http://www.businessinsider.com/elon-musk-on-tesla-over-the-air-software-update-summon-feature-2016-1