

Challenges and Approaches to Realizing Autonomous Vehicle Safety

Nidhi Kalra

CT-463

Testimony presented to the House Energy and Commerce Committee, Subcommittee on Digital Commerce and Consumer Protection on February 14, 2017.



For more information on this publication, visit www.rand.org/pubs/testimonies/CT463.html

Testimonies

RAND testimonies record testimony presented or submitted by RAND associates to federal, state, or local legislative committees; government-appointed commissions and panels; and private review and oversight bodies.

Published by the RAND Corporation, Santa Monica, Calif.

© Copyright 2017 RAND Corporation

RAND® is a registered trademark.

Limited Print and Electronic Distribution Rights

This document and trademark(s) contained herein are protected by law. This representation of RAND intellectual property is provided for noncommercial use only. Unauthorized posting of this publication online is prohibited. Permission is given to duplicate this document for personal use only, as long as it is unaltered and complete. Permission is required from RAND to reproduce, or reuse in another form, any of its research documents for commercial use. For information on reprint and linking permissions, please visit www.rand.org/pubs/permissions.html.

www.rand.org

Challenges and Approaches to Realizing Autonomous Vehicle Safety

Testimony of Nidhi Kalra¹
The RAND Corporation²

Presented to the Committee on Energy and Commerce
Subcommittee on Digital Commerce and Consumer Protection
United States House of Representatives

February 14, 2017

Chairman Latta, Ranking Member Schakowsky, and distinguished members of the subcommittee, thank you for the opportunity to testify on the safety and testing of autonomous vehicles. As you know, traffic crashes pose a public health crisis in the United States, and autonomous vehicles have the potential to mitigate this crisis. As a society, we want autonomous vehicles to become as safe as possible as quickly as possible so that their potential is realized and we can usher in a new era of safer transportation.

Today I would like to discuss three challenges that currently stand in the way of this vision:

1. There is currently no proven, practical way to test autonomous vehicle safety prior to widespread use.
2. There is no consensus on how safe autonomous vehicles should be.
3. Real-world driving experience is an essential ingredient for improving safety, but it also exposes people to the very safety risks we hope to reduce.

The enormous potential benefits of autonomous vehicles—not only for safety but for other transportation goals—makes it urgent that we overcome these challenges. I will describe how it is important to develop and validate methods of demonstrating safety and incorporate these methods into a regulatory framework that gradually increases the use and exposure of autonomous vehicles as increasingly stringent safety criteria are met. I will also describe how, until such a framework is developed, regulators and industry can use pilot studies and data-sharing to manage the risks of the technology while fostering its development.

¹ The opinions and conclusions expressed in this testimony are the author's alone and should not be interpreted as representing those of the RAND Corporation or any of the sponsors of its research.

² The RAND Corporation 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.

Traffic Crashes Pose a Public Health Crisis

In the United States, roughly 32,000 people are killed and more than 2 million are injured in motor vehicle crashes every year.³ Although safety has generally improved over the past several decades, 2015 saw 35,000 road fatalities, the largest percentage increase in fatalities in this country in more than 50 years. This occurred partly because Americans drove more and partly because they drove worse. While vehicles in general are becoming safer 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 that have recently offset those gains.⁴

U.S. motor vehicle crashes as a whole can pose enormous economic and social costs—more than \$800 billion in a single year.⁵ And more than 90 percent of crashes are caused by human errors,⁶ such as driving too fast and misjudging other drivers' behaviors, as well as alcohol impairment, distraction, and fatigue.

Autonomous Vehicles Present Benefits and Risks to Safety

Autonomous vehicles have the potential to significantly mitigate this public safety crisis by eliminating many of the mistakes that human drivers routinely make.⁷ To begin with, autonomous vehicles cannot be drunk, distracted, or tired; these factors are involved in 29 percent, 10 percent, and 2.5 percent, respectively, of all fatal crashes.⁸ Autonomous vehicles could also perform better than human drivers because of better perception (e.g., no blind spots),

³ Bureau of Transportation Statistics, *Motor Vehicle Safety Data*, Washington, D.C.: Research and Innovative Technology Administration, U.S. Department of Transportation, 2015, Table 2-17.

⁴ National Highway Traffic Safety Administration, “2015 Motor Vehicle Crashes: Overview,” *Traffic Safety Facts: Research Note*, Washington, D.C.: National Center for Statistics and Analysis, U.S. Department of Transportation, DOT HS 812 318, August 2016. As of January 18, 2017: <https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/812318>

⁵ Lawrence Blincoe, Ted R. Miller, Eduard Zaloshnja, and Bruce A. Lawrence, *The Economic and Societal Impact of Motor Vehicle Crashes 2010*, Washington, D.C.: National Highway Traffic Safety Administration, DOT HS 812 013, May 2015.

⁶ National Highway Traffic Safety Administration, *Traffic Safety Facts, A Brief Statistical Summary: Critical Reasons for Crashes Investigated in the National Motor Vehicle Crash Causation Survey*, Washington, D.C.: National Center for Statistics and Analysis, U.S. Department of Transportation, DOT HS 812 115, February 2015.

⁷ James M. Anderson, 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-433-2-RC, 2014; and Daniel J. Fagnant 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.

⁸ National Highway Traffic Safety Administration, 2016; and National Highway Traffic Safety Administration, “Drowsy Driving,” *Traffic Safety Facts: Crash Stats*, Washington, D.C.: National Center for Statistics and Analysis, U.S. Department of Transportation, DOT HS 812 449, March 2011. As of January 18, 2017: <https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/811449>. This does not mean that 41.5 percent of all fatal crashes are caused by these factors, because a crash may involve, but not be strictly caused by, one of these factors, and because more than one of these factors may be involved in a single crash.

better decisionmaking (e.g., more-accurate planning of complex driving maneuvers), and better execution (e.g., faster and more-precise control of steering, brakes, and acceleration).

However, autonomous vehicles likely will not eliminate all crashes. For instance, inclement weather and complex driving environments pose challenges for autonomous vehicles, as well as for human drivers, and autonomous vehicles might perform worse than human drivers in some cases, particularly at early stages of testing and deployment.⁹ There is also the potential for autonomous vehicles to pose new and serious crash risks—for example, crashes resulting from cyber attacks.¹⁰ Clearly, autonomous vehicles present both potential benefits and potential risks to transportation safety. Several challenges stand in the way of managing those safety risks and maximizing the benefits.

There Is Currently No Proven, Practical Way to Determine Autonomous Vehicle Safety Prior to Widespread Use

Road Tests Cannot Prove Safety

A road test that a person takes at the Department of Motor Vehicles assesses whether the person can perform a specific set of driving skills under regular traffic situations. Passing the test does not prove that the person will be a safe driver but is nevertheless viewed as adequate for granting a permit or a license.

Such road tests could similarly demonstrate that an autonomous vehicle can perform basic driving skills but would not be able to prove their safety. Doing so would require demonstrations of safe performance under the full range of conditions in which the vehicle is expected to drive.

Extensive Test-Driving Is Statistically Powerful but Impractical

A logical alternative is to test-drive autonomous vehicles extensively in real traffic and analyze their performance. Developers of autonomous vehicles rely on this approach to evaluate and improve their systems,¹¹ typically with trained operators behind the wheel who are ready to take control in the event of an impending failure incident. Developers can analyze the failure incident after the fact to assess what the autonomous vehicle would have done without

⁹ Lee Gomes, *Hidden Obstacles for Google's Self-Driving Cars: Impressive Progress Hides Major Limitations of Google's Quest for Automated Driving*, Massachusetts Institute of Technology, August 28, 2014.

¹⁰ Anderson et al., 2014.

¹¹ Extensive testing on public roads is essential for developing and evaluating autonomous vehicles, given their great complexity and the diversity and unpredictability of conditions in which they need to operate. In contrast, typical automobile components are significantly simpler and their operating conditions can be well defined and recreated in controlled settings, which enables laboratory testing and verification. For example, curtain-style airbags are tested to assess inflation time, fill capacity, and other responses in a range of temperature conditions and impact configurations; they are also tested in laboratory crashes to evaluate their performance in collisions. See Helen Kaleto, David Winkelbauer, Chris Havens, and Michael Smith, "Advancements in Testing Methodologies in Response to the FMVSS 201U Requirements for Curtain-Type Side Airbags," Society of Automotive Engineers International, Technical Paper 2001-01-0470, 2001.

intervention and whether it would have resulted in a crash or other safety issue. Developers have presented data from test-driving to Congress in hearings about autonomous vehicle regulation.¹²

This approach has statistical merits: Data about the number and types of miles traveled and the number of crashes, injuries, and fatalities can be used to assess safety.¹³ However, this approach is largely impractical for pre-market testing: even though the number of crashes, injuries, and fatalities from human drivers is high overall, the rate of these failures is low *in comparison with the number of miles that people drive*. Americans drive nearly 8 billion miles every day and 3 trillion miles every year.¹⁴ The 35,092 fatalities and 2.44 million injuries in 2015 correspond to a rate of 1.12 fatalities and 78 injuries per 100 million miles driven. Given that current traffic fatalities and injuries are rare events compared with vehicle miles traveled, fully autonomous vehicles would have to be driven hundreds of millions of miles, and sometimes hundreds of billions of miles, to demonstrate their reliability in terms of fatalities and injuries.¹⁵ Under even aggressive testing assumptions, existing test fleets would take tens and sometimes hundreds of years to drive these miles—an impossible proposition if the aim is to demonstrate their performance *prior* to releasing them on the roads for consumer use.¹⁶ It is important to note that safety could probably be readily demonstrated once autonomous vehicles are released on the roads and consumer use is widespread, given how much Americans drive.

The available data on autonomous vehicle driving performance illustrates the impracticalities of test-driving as a means of assessing their safety. As one example, Google’s autonomous vehicle fleet was test-driven approximately 1.3 million miles in autonomous mode and was involved in 11 crashes from 2009 to 2015.¹⁷ Blanco and colleagues compared Google’s fleet performance with human-driven performance.¹⁸ They found that Google’s fleet might result in fewer crashes with only property damage, but they could not draw conclusions about the relative performance in terms of injuries and fatalities. Given the rate of human and autonomous vehicle failures, there were simply not enough autonomously driven miles to make statistically meaningful comparisons.

¹² Chris Urmson, “Hands Off: The Future of Self-Driving Cars,” testimony before the Senate Committee on Commerce, Science and Technology, Washington, D.C., March 15, 2016. As of March 22, 2016: https://www.commerce.senate.gov/public/_cache/files/5c329011-bd9e-4140-b046-a595b4c89eb4/BEADFE023327834146FF4378228B8CC6.google-urmson-testimony-march152016.pdf

¹³ Nidhi Kalra and Susan 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 February 9, 2017: http://www.rand.org/pubs/research_reports/RR1478.html

¹⁴ Bureau of Transportation Statistics, 2015.

¹⁵ The number of miles depends on the metric of concern (fatalities or injuries), the performance of the technology, and the level of statistical confidence desired.

¹⁶ Kalra and Paddock, 2016.

¹⁷ Two of these crashes involved injury and none involved a fatality. Seven of the crashes did not reach a level of severity that would warrant a Department of Motor Vehicles report.

¹⁸ Myra Blanco, Jon Atwood, Sheldon Russell, Tammy Trimble, Julie McClafferty, and Miguel Perez, *Automated Vehicle Crash Rate Comparison Using Naturalistic Data*, Virginia Tech Transportation Institute, January 2016. As of February 9, 2017: http://www.vtti.vt.edu/PDFs/Automated%20Vehicle%20Crash%20Rate%20Comparison%20Using%20Naturalistic%20Data_Final%20Report_20160107.pdf

As another example, in its review of the fatal crash involving a Tesla in May 2015,¹⁹ the National Highway Traffic Safety Administration (NHTSA) found that Tesla vehicles equipped with Autopilot were involved in 40 percent fewer crashes in which airbags deployed than pre-Autopilot Tesla vehicles.²⁰ However, the investigation reports neither the statistical significance of this difference nor the number of miles driven by each group of vehicles, from which statistical significance could be computed. It is therefore not possible to judge from this evidence whether Tesla vehicles with Autopilot are indeed safer than those without.²¹

Testing in Partial Simulation Has Merits but Remains Unproven

Another approach is to test in partial simulation. For example, Tesla announced last year that all of its vehicles would be equipped with hardware and sensors to allow for fully autonomous driving. However, this feature would not be enabled and instead operate in a “shadow” mode.²² In shadow mode, the autonomous system gathers sensor data and makes driving decisions as though it were in control of the vehicle, but those decisions are not executed. Instead, the human driver maintains control of the vehicle and the *simulated* decisions of the shadow autonomous system are compared with the executed decisions of the driver in order to detect errors and anomalies.²³ In aggregate across the fleet, these errors and anomalies can be used to validate the autonomous system’s performance. This approach enables autonomous vehicle performance to be demonstrated in only weeks or months because of the scale of deployment, rather than decades or even centuries with small fleets of test vehicles. Additionally, it enables learning during market deployment—which may be crucial for achieving safety—but without putting early adopters and the public at undue risk. However, the partial simulation approach must still be analyzed and validated as a method of testing the performance of autonomous systems. Furthermore, it may not be an option for developers who do not have a non-autonomous fleet deployed, and it may be expensive.

¹⁹ Office of Defects Investigation, “ODI Resume: Automatic Vehicle Control Systems,” investigation PE 16-007, National Highway Traffic Safety Administration, January 19, 2017. As of February 9, 2017: <https://static.nhtsa.gov/odi/inv/2016/INCLA-PE16007-7876.PDF>

²⁰ There is debate about whether Tesla’s Autopilot system qualifies as an autonomous system or a very sophisticated driver assistance system.

²¹ It is worth pointing out that Tesla’s Autopilot system has been driven many millions of miles on real roadways, so the difference in performance between Autopilot-equipped and pre-Autopilot vehicles could be statistically significant. However, the data provided in NHTSA’s investigative report is inadequate for determining this. See Ken Yeung, “Elon Must Says Tesla’s Autopilot Has Driven 222 Million Miles,” *Venture Beat*, October 7, 2016. As of February 9, 2017: <http://venturebeat.com/2016/10/07/elon-musk-says-teslas-autopilot-has-driven-222-million-miles/>

²² Alex Nishimoto, “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/>

²³ This approach can also be used to test a new, pre-market autonomous technology against an earlier autonomous technology. Tesla will likely be comparing the performance of its autonomous system not only with human drivers but also with its deployed Autopilot systems.

Existing Functional Safety Standards Are Not Designed for Autonomous Vehicles

Another question is whether existing functional safety standards can provide assurances of safety.²⁴ Such discussions often raise ISO 26262, an international standard for functional safety of electrical and electronic systems in production automobiles.²⁵ But such standards as this were not designed with autonomous vehicles in mind, and applying them to autonomous vehicles is difficult for a variety of reasons,²⁶ including but not limited to the following:

- Such standards are intended for vehicles in which the human driver can ultimately correct for errors, which is inapplicable to fully autonomous vehicles that require no human intervention.
- The standards work well when system inputs and outputs can be well specified, but this is probably not possible with large amounts of diverse, high-speed data coming from vehicle sensors.
- It is difficult to apply formal methods to machine learning techniques, which are the cornerstone of rapid improvement in autonomous driving but which often result in decision rules that are difficult for humans to interpret.²⁷

This is not to suggest that functional safety standards cannot help; rather, further work is needed to adapt them to the unique challenges that autonomous vehicles pose.

In sum, the transportation industry and policymakers do not yet have a method that is both practical and sound for testing autonomous vehicle safety. The question, “How safe is an autonomous vehicle?” may be unanswerable prior to widespread use. This does not mean that their use should be prohibited; the technology has too much potential to save lives. Instead, it suggests that the race to develop autonomous vehicles needs a parallel race to develop methods for demonstrating and managing their safety.

There Is No Consensus on How Safe Autonomous Vehicles Should Be

The issue of how safe autonomous vehicles should be is worth considering, even if their degree of safety cannot yet be fully proven. Some will insist that anything short of totally eliminating risk is an unacceptable safety compromise. The argument is that it is acceptable if humans make mistakes, but not if machines do. But, again, waiting for autonomous vehicles to

²⁴ *Functional safety* is essentially the ability of a system to operate correctly in response to inputs, including recognizing and handling internal and external failures.

²⁵ International Organization for Standardization, “Road Vehicles—Functional Safety—Part 1: Vocabulary,” ISO 26262-1:2011, 2011. As of February 9, 2017: http://www.iso.org/iso/catalogue_detail?csnumber=43464

²⁶ Philip Koopman and Michael Wagner, “Autonomous Vehicle Safety: An Interdisciplinary Challenge,” *IEEE Intelligent Transportation Systems Magazine*, Vol. 9, No. 1, Spring 2017, pp. 90–96.

²⁷ *Machine learning* is a branch of artificial intelligence that enables computers to change their behavior and learn when introduced to new data. For example, a legged robot can learn to walk without being explicitly programmed to do so, by testing different leg motions, observing the resulting change in its position, and using those observations to refine its gait. These techniques can enable computers and robots to improve performance faster or reach better performance than if improvements could be made only directly by human programmers. See Nate Kohl and Peter Stone, “Machine Learning for Fast Quadrupedal Locomotion,” *Nineteenth National Conference on Artificial Intelligence*, San Francisco, July 2004, pp. 611–616.

operate nearly perfectly misses opportunities to save lives because it means the needless perpetuation of the risks posed by human drivers.

It seems sensible that autonomous vehicles should be allowed on America's roads when they are judged safer than the average human driver, allowing more lives to be saved and sooner while still ensuring that autonomous vehicles do not create new risks. An argument can be made that autonomous vehicles could be allowed even when they are not as safe as average human drivers if developers can use early deployment as a way to rapidly improve vehicle safety. The vehicles could become at least as good as the average human sooner than they would otherwise—and thus save more lives over the long term. Moreover, there might be significant non-safety benefits, such as allowing people to do more-productive things while they are in a vehicle, that could outweigh safety drawbacks.

The lack of consensus on this point is not a failure but rather a genuine expression of Americans' different values and beliefs when it comes to humans versus machines. But it complicates the challenge of developing safety benchmarks for the technology.

Real-World Driving Experience Is Needed for Safety but Poses Risks

Resolving the above challenges is urgent because real-world driving experience may be one of the most important tools for improving autonomous vehicle safety and, by extension, road safety. This is because, unlike most humans, autonomous vehicles can learn from each other's mistakes. When a human driver makes a mistake on the road, typically only that individual can learn from that experience to improve their driving habits. Other drivers are unaffected. This is not the case with autonomous vehicles. Autonomous vehicle developers use the driving experience of individual vehicles to improve the state of the art in autonomous vehicle safety. The machine learning algorithms that govern autonomous vehicle perception, decisionmaking, and execution rely largely on driving experience to improve. Therefore, the more (and more-diverse) miles that autonomous vehicles drive, especially in diverse conditions, the more potential there is for improving the state of the art in autonomous vehicle safety performance.

As one example of improvement already taking place, Google reported to the California Department of Motor Vehicles that in the fourth quarter of 2014, its vehicles disengaged approximately once every 600 miles, compared with once every 2,800 miles in the fourth quarter of 2015.²⁸ 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 then August 2016 to include, among other things, better safety performance based on experiences of existing vehicles. Tesla calls this “fleet learning,”—that is, when an entire fleet is able to learn from the experiences of each deployed vehicle in that fleet.

²⁸ Google, *Google Self-Driving Car Testing Report on Disengagements of Autonomous Mode*, December 2015. As of February 9, 2017: https://www.dmv.ca.gov/portal/wcm/connect/dff67186-70dd-4042-bc8c-d7b2a9904665/google_disengagement_report.pdf?MOD=AJPERES. The disengagements reported by Google reflect cases in which the autonomous vehicle 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. It is difficult to use 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.

Yet having autonomous vehicles learn from real-world driving experience presents its own problem: Learning in real-world settings implies risks to early adopters or to other road users, from which late adopters would benefit. This is analogous to the risk of allowing teenage drivers on the road: They may not be good drivers yet, but they need experience to become good drivers. However, until then, they pose risks to themselves and to others. We have policies in place to try to limit risks from inexperienced young drivers, such as a minimum driving age and restrictions on learner's permits. Those policies seek to balance the goal of long-term improvement with the need for near-term experience.

The same will be true for autonomous vehicles. Choices made now about when and how autonomous vehicles are introduced will affect not only safety in the near term but also how quickly the vehicles improve, at what near-term cost, and how safe they ultimately become in the future. For instance, a policy that requires autonomous vehicles to be nearly perfect before they are allowed on the roads for widespread use could 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 autonomous vehicles are allowed on the road even when there are widespread safety concerns in the near term or when the technology is demonstrably unsafe, there may be little market for the technology. The public may even demand that autonomous vehicles be removed from the road, resulting in much slower adoption and delays in the experience needed to improve.

The History of Airbag Technology and Regulation Is Instructive

The history of airbag technology development and regulation illustrates the challenges of inaccurate assessments of safety impacts, of conflicting views about what constitutes adequate safety, and of the risks of learning through experience.

Airbags were initially introduced in the early 1970s in higher-end vehicle models.²⁹ At that time, seat belt use was low and airbags were marketed as alternatives, rather than supplements, to seat belts. Using the logic that airbags would make seat belts unnecessary, NHTSA initiated efforts in the 1970s to pass regulations requiring airbags in all U.S. automobiles. Such regulations met with significant resistance from most automobile manufacturers, who did not want either the responsibility or the liability for the losses resulting from crashes and did not believe these safety features would sell. It was not until 1991 that regulations passed requiring the use of airbags in model years 1999 and later.³⁰

The pre-market methods used to assess airbag safety significantly overestimated their effects. In 1977, NHTSA estimated that airbags would save on the order of 9,000 lives per year and based its subsequent regulations on these expectations.³¹ While airbags have certainly saved

²⁹ Murray Mackay, "Liability, Safety, and Innovation in the Automotive Industry," in Peter W. Huber and Robert E. Litan, eds., *The Liability Maze: The Impact of Liability Law on Safety and Innovation*, Washington, D.C.: Brookings, 1991, pp. 191–223.

³⁰ Public Law 102-240, Intermodal Surface Transportation Efficiency, 1991.

³¹ Kimberly M. Thompson, Maria Segui-Gomez, and John D. Graham, "Validating Benefit and Cost Estimates: The Case of Airbag Regulation," *Risk Analysis*, Vol. 22, No. 4, August 2002, pp. 803–811.

many lives, they have not lived up to original expectations. NHTSA calculates that airbags saved a total of 8,369 lives in the 14 years between 1987 and 2001,³² and in recent years saved between 2,000 and 3,000 lives annually.³³ The true effects of airbags were understood only after they were widely used, not before.

Airbag technology improved with real-world use, but that improvement came, in part, because of the lives that were lost in vehicles using early versions of the technology. It became evident that first-generation airbags posed a risk to many passengers, particularly more-vulnerable passengers, such as women of small stature, the elderly, and children. NHTSA determined that 291 deaths were caused by airbags between 1990 and July 2008, primarily because of the extreme force necessary to meet the performance standard of protecting the unbelted adult male passenger.³⁴ Despite their overall benefits, there was a backlash against airbags and airbag regulations because many viewed them as saving healthy adults at the expense of children and the elderly.³⁵ Although, overall, airbags have saved tens of thousands of lives, the road to that success is winding.

The full history of airbags reveals a complex interaction of technology, policy, and social behavior. Therefore, that history may be instructive in how we approach autonomous vehicle safety. First, it may be appropriate to have modest expectations of autonomous vehicle performance and impact. Few anticipated that use of seat belts would rise as much as it has and that airbags would eventually be used more as a *supplement* than a substitute for seat belts, thus limiting their life-saving potential. Similarly unexpected developments are likely to arise with autonomous vehicles. For example, driver assistance technologies could significantly improve safety, lessening the marginal benefits of full autonomy. As another possibility, human behavior could change to undermine the benefits of the technology in unforeseen ways. Perhaps drivers of vehicles or pedestrians will behave more recklessly with the expectation that autonomous systems will avoid the consequences. While this example is speculative, the point remains: We don't know what we don't know. And there will almost surely be unexpected detours along the way.

Second, the question of how the benefits and risks of autonomous vehicles are distributed is important. Even if the technology saves lives overall, it may increase safety for some while reducing it for others, as airbags did. If this distribution of risk is inequitable or exacerbates

³² Donna Glassbrenner, *Estimating the Lives Saved by Safety Belts and Air Bags*, Washington, D.C.: National Center for Statistics and Analysis, National Highway Traffic Safety Administration, Paper No. 500, undated. As of February 9, 2017: <http://www-nrd.nhtsa.dot.gov/pdf/nrd-01/esv/esv18/CD/Files/18ESV-000500.pdf>

³³ National Highway Traffic Safety Administration, "Lives Saved FAQs," Washington, D.C.: National Center for Statistics and Analysis, U.S. Department of Transportation, DOT HS 811 105, December 2009. As of February 9, 2017: <https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/811105>

³⁴ National Highway Traffic Safety Administration, *Special Crash Investigations: Counts of Frontal Air Bag Related Fatalities and Seriously Injured Persons*, Washington, D.C.: National Center for Statistics and Analysis, U.S. Department of Transportation, January 1, 2008. As of February 9, 2017: <http://www-nrd.nhtsa.dot.gov/Pubs/AB0108.PDF>

³⁵ David J. Houston 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.

existing inequities in transportation safety, this raises ethical questions. There may be a backlash against the technology, particularly if the harms are viewed as avoidable.

Finally, the airbag experience is instructive because there was a long lag between when the technology was developed and when it was became widespread. The first airbag patents were issued in the early 1950s; by the early 1970s, airbag-manufacturing companies existed and the technology was reasonably mature, yet they were in very few cars.³⁶ It was not until 1999, nearly 20 years later, that airbags were ultimately required in all cars. Thousands of deaths might have been prevented by airbags in the interim. Developing the technology was a necessary but not sufficient condition for the technology to be widely used.

Recommendations

Feasible, Sound Methods of Testing Safety Are Urgently Needed

There is a clear and essential role for sound policy in addressing the above challenges. First, feasible and sound methods of testing safety are urgently needed. These methods can be developed and proposed by industry, researchers and academics, or federal regulators. Importantly, it is not enough to simply propose or develop testing methods. The methods need to be validated rigorously, objectively, and independently both to assess their statistical soundness, relation to existing safety and performance standards, and engineering and social considerations and to build confidence in the methods among diverse stakeholders.

These Methods Should Be Built into a Flexible Regulatory Framework

Second, it is not enough for testing methods to simply exist. They should be built into a regulatory framework that defines what safety performance thresholds must be met in order to put autonomous vehicles on the road. Importantly, conventional regulatory methods may not work. Currently, there are two distinct phases to bringing automotive innovations to market: pre-market development and market deployment. In the pre-market phase, the technology's performance is tested in simulations, laboratory settings, and limited real-world experiments. Once an innovation is judged safe and effective and meets regulatory requirements, the market deployment phase begins. The innovation is made available in the automotive market where it may be incorporated into millions of new vehicles very quickly. Little performance data are collected in the market phase, except through consumer feedback or in crash reports, and improvements to vehicle technologies are usually introduced through new vehicle sales rather than through improvements to vehicles already in the fleet.

This paradigm poses challenges to bringing autonomous vehicles to the market. First, as described, conventional pre-market testing methods make it difficult to judge whether these technologies are sufficiently safe. Second, the two-phase approach overlooks the fact that autonomous vehicles can use real driving experience to improve performance, not just for an

³⁶ Jerry L. Mashaw and David L. Harfst, *The Struggle for Auto Safety*, Cambridge: Harvard, 1990.

individual vehicle but for the entire fleet, in real time by using over-the-air updates of vehicle software. This may be essential for achieving high levels of safety.

For these reasons, the framework should balance the need for real-world autonomous vehicle testing with the need to protect the public from undue risk. For example, a lower threshold of safety may be acceptable for early demonstration projects intended to improve autonomous vehicle performance in controlled environments. A higher threshold of safety would be warranted for widespread consumer use in uncontrolled environments. Developing the appropriate thresholds of safety should be informed by research on levels of tolerable risk in automotive and other environments, historical safety regulation and performance criteria, stakeholder values and preferences, and an assessment of how different safety levels could shape the arc of transportation safety over the long term. Thus, establishing these thresholds requires a combination of objective evidence and subjective values.

Several states are already taking the first step in such a graded approach by developing requirements for testing autonomous vehicles on public roads.³⁷ Many regulations allow autonomous vehicles to be tested on public roads with no demonstration of safety as long as there is a trained human driver behind the wheel. These regulations include requirements related to registration, insurance and liability, data reporting, and so on to explicitly align developers' and the public's safety goals and to enable oversight. More-stringent requirements are in place for testing without a human driver. Pennsylvania's policy, for example, requires demonstrating that the autonomous system can fully control the vehicle under all driving conditions.³⁸ However, regulations to allow autonomous vehicles on the road for non-testing consumer use have stymied many policymakers because it is not clear how public safety can be best achieved unless the safety of autonomous vehicles is known.

Having individual states develop safety frameworks is problematic because it gives rise to a patchwork of regulations. There are also practical issues: Developing such a framework requires extensive resources and technical expertise. Such development would likely fall under NHTSA's jurisdiction as a regulating body for transportation and vehicle safety, but the framework should be developed in consultation with industry, state and local policymakers, and the public. NHTSA has already released federal policies for autonomous vehicles, but these do not specify testing methods or develop a safety framework. The policies are also guidelines for technology development and use, not requirements that automakers must follow.³⁹

³⁷ Pennsylvania Department of Transportation, *Pennsylvania Autonomous Vehicle Testing Policy: Final Draft Report of the Autonomous Vehicle Policy Task Force*, November 2, 2016. As of February 9, 2017: <http://www.penndot.gov/ProjectAndPrograms/ResearchandTesting/Documents/AV%20Testing%20Policy%20DRAFT%20FINAL%20REPORT.pdf>; and California Department of Motor Vehicles, "Express Terms: Autonomous Vehicles," undated. As of February 9, 2017: https://www.dmv.ca.gov/portal/wcm/connect/211897ae-c58a-4f28-a2b7-03cbe213e51d/avexpressterms_93016.pdf?MOD=AJPERES

³⁸ This does not constitute a demonstration of safety; instead, it functions more like a road test. See Pennsylvania Department of Transportation, 2016.

³⁹ National Highway Traffic Safety Administration, *Federal Automated Vehicles Policy: Accelerating the Next Revolution in Roadway Safety*, Washington, D.C.: U.S. Department of Transportation, September 2016. As of February 9, 2017: https://one.nhtsa.gov/nhtsa/av/pdf/Federal_Automated_Vehicles_Policy.pdf

There are existing analogues to a phased approach. The Food and Drug Administration's regulations for approving drugs are one example, where drugs must go through several phases of animal testing, human testing, and post-market monitoring.⁴⁰ Each phase requires increasingly stringent requirements to be met, while also increasing the number of people exposed to the treatment. There are many legitimate concerns with the process, including its high cost and long duration,⁴¹ as well as whether it stagnates the development of new life-saving treatments. However, there are also compelling proposals to address those shortcomings, particularly through adaptive processes designed to facilitate rapid learning and improvement over time.⁴² In contrast, there are few who would propose that the government should not regulate drugs at all or that they should be allowed on the market before being evaluated.

It is also important to acknowledge legitimate concerns about the government's ability to effectively regulate a rapidly evolving and disruptive technology. This task is daunting. But it is important to note historical precedents in which the government has been at the leading edge of safety in new technologies. The early history of aviation offers a useful example.⁴³ The Wright brothers made their first sustained flight in 1903. The National Air Mail Service began in 1918 and in 1924 inaugurated 24-hour service between New York and San Francisco. Commercial aviation began around the same time but was much riskier. The fatality rate for commercial aviation in this era was one fatality per 13,500 miles. In contrast, the Air Mail Service was nearly 60 times safer, with a fatality rate of one per 789,000 miles. According to the Federal Aviation Administration, this was due to the government's prioritization of safety and implementation of a strict safety program for pilots, aircraft, and flights. The safety record of the Air Mail Service was one of the factors motivating the commercial aviation industry to request government oversight and regulation, which subsequently began in 1926 with the Air Commerce Act and involved decades of iteration. The Federal Aviation Administration we know today was not created until more than 30 years later, in 1959. Today, the gains in aviation safety are clear: There were no fatalities among U.S. commercial air carriers in 2014, although they traveled more than 7.6 billion aircraft miles. The average fatality rate from 2011 to 2014 (the five most-recent years for which the Bureau of Transportation Statistics has data) was 0.29 fatalities per 1 billion aircraft miles.⁴⁴

⁴⁰ U.S. Food and Drug Administration, "Drug Approval Process," undated. As of February 9, 2017: <http://www.fda.gov/downloads/drugs/resourcesforyou/consumers/ucm284393.pdf>

⁴¹ Joseph A. DiMasi, Ronald W. Hansen, and Henry G. Grabowski, "The Price of Innovation: New Estimates of Drug Development Costs," *Journal of Health Economics*, Vol. 22, 2003, pp. 151–185.

⁴² H. G. Eichler, K. Oye, L. G. Baird, E. Abadie, J. Brown, C. L. Drum, J. Ferguson, S. Garner, P. Honig, M. Hukkelhoven, J. C. W. Lim, R. Lim, M. M. Lumpkin, G. Neil, B. O'Rourke, E. Pezalla, D. Shoda, V. Seyfert-Margolis, E. V. Sigal, J. Sobotka, D. Tan, T. F. Unger, and G. Hirsch, "Adaptive Licensing: Taking the Next Step in the Evolution of Drug Approval," *Clinical Pharmacology and Therapeutics*, Vol. 91, No. 3, March 2012, pp. 426–437.

⁴³ Federal Aviation Administration, *History of Aviation Safety Oversight in the United States*, Washington, D.C.: U.S. Department of Transportation, DOT/FAA/AR-08/39, July 2008.

⁴⁴ Bureau of Transportation Statistics, "U.S. Air Carrier(a) Safety Data," Washington, D.C.: U.S. Department of Transportation, undated, Table 2-9. As of February 9, 2017: https://www.rita.dot.gov/bts/sites/rita.dot.gov.bts/files/publications/national_transportation_statistics/html/table_02_09.html

The histories of both aviation and drug regulation are long and complex. They suggest that government can and should effectively regulate industries that pose enormous benefits and risks, not only for the benefit of the public but also for the growth of the industry. They further suggest that creating effective regulation is neither easy nor fast and will require revision and adaptation over time.

In the Interim, Risk Should Be Managed with Pilot Studies and Data-Sharing

A regulatory framework will take time and depends on the development of testing methods and safety thresholds, which must also be studied and developed. In the interim, the industry should be allowed to continue development of the technology with safeguards in place to realize the benefits that the technology presents. This raises an important question: How do we enable autonomous vehicles to improve as quickly as possible while lowering the risks they pose, even while those risks may be difficult to quantify? There are several tactics policymakers could consider to accelerate autonomous vehicles' improvement.

A first step is to conduct real-world but lower-risk pilot studies of autonomous vehicles for limited uses and in constrained environments. Risk can be lowered first by operating autonomous vehicles in conditions in which crashes are less likely. This can include limiting autonomous vehicle pilot studies to areas with less-complex terrain, to routes that are well maintained and easy to navigate, to nondangerous weather conditions, or to some combination of these controls. It can also include educating communities about safe behavior in and around autonomous vehicles. Furthermore, risk can be lowered by designing and operating vehicles so that when a crash occurs, the consequences of the crash to passengers and bystanders are fewer. This could be accomplished by limiting vehicle speed, ensuring that all pilot-study passengers wear seat belts, and so forth. These strategically limited pilot studies can then be expanded as safe operation of autonomous vehicles is demonstrated.

A second consideration is sharing driving data across the industry and with policymakers. Autonomous vehicle developers already use the experiences of a single vehicle to improve the safety of their respective fleets. This improvement could occur even faster if the experiences of each vehicle in each fleet could be used across all developers to improve the entire industry.⁴⁵ There are certainly nontrivial concerns about protecting trade secrets, but these concerns could be addressed and must be balanced with the societal need for safe autonomous vehicle technology. Data-sharing should also involve regular information about miles traveled, crashes, failures, and other information that can help provide early evidence of safety and safety concerns.

⁴⁵ The kind of data that are shared must be carefully considered. Different developers have very different autonomous vehicle technology designs, which may limit the lessons that can be learned from data-sharing across fleets.

Conclusions

In sum, autonomous vehicles hold enormous promise for transportation safety, but realizing the safety benefits is not guaranteed. This is, in part, because it is difficult to know how safe autonomous vehicles are and because Americans may not agree on how safe they should be. Concerted and immediate effort should be made to develop sound and feasible testing methods and to develop those methods into a regulatory framework that balances the need for development and deployment of the technology with appropriate levels of safety at each stage of exposure. While this is taking place, policymakers should pursue ways of fostering the development of autonomous vehicles while lowering their risks.

Conflict of Interest Statement: Nidhi Kalra's spouse, David Ferguson, is co-founder and president of Nuro, a machine learning and robotics startup engaged in autonomous vehicle development. He previously served as a principal engineer for Google's driverless car project. This written testimony was carefully reviewed by subject-matter experts within the RAND Corporation; the research quality assurance team for the RAND Justice, Infrastructure, and Environment division; and the RAND Office of Congressional Relations. However, the opinions and conclusions expressed in this testimony are the author's alone and should not be interpreted as representing those of the RAND Corporation or any of the sponsors of its research.