

# A Roadmap for Autonomous Vehicles: State Tort Liability, Automobile Insurance, and Federal Safety Regulation

Mark A. Geistfeld\*

*Driver error currently causes the vast majority of motor vehicle crashes. By eliminating the human driver, autonomous vehicles will prevent thousands of fatalities and serious bodily injuries, which makes a compelling safety case for policies that foster the rapid development of this technology. Major technological advances have occurred over the past decade, but there is widespread concern that the rate of development is hampered by uncertainty about manufacturer liabilities for a crash. Apparent variations in the requirements of state tort law across the country make it difficult for manufacturers to assess their liability exposure in the national market. The patchwork of state laws and the resultant uncertainty have prompted calls for the federal safety regulation of autonomous vehicles.*

*The uncertainty seems to be the inevitable result of trying to predict how tort rules governing old technologies will apply to the new technology of automated driving. As I will attempt to demonstrate, however, well-established tort doctrines widely adopted by most states, if supplemented by two new federal safety regulations, would provide a comprehensive regulatory approach that would largely dissipate the costly legal uncertainty now looming over this emerging technology.*

---

DOI: <https://dx.doi.org/10.15779/Z38416SZ9R>

Copyright © 2017 California Law Review, Inc. California Law Review, Inc. (CLR) is a California nonprofit corporation. CLR and the authors are solely responsible for the content of their publications.

\* Sheila Lubetsky Birnbaum Professor of Civil Litigation, New York University School of Law. Many thanks to Mike Green, Andrea Matwyshyn, Bob Rabin, and Cathy Sharkey for their helpful comments. My development of this argument also benefitted from discussion at the May 2016 Roundtable on Security for Cyber-Physical Systems at the University of Pennsylvania Law School, and from a faculty workshop at NYU Law School. Special thanks to my transportation muse, Janette Sadik-Khan, who helped guide me into the driverless space. Jacqueline Cavallaro and Emmanuel Sanders provided excellent research assistance. Financial support was provided by the Filomen D'Agostino and Max E. Greenberg Research Fund of the New York University School of Law.

*The technology itself largely solves the most vexing tort problems for reasons that prior analyses have missed. Autonomous vehicles will transform the individualized behavior of human drivers into a collective, systemized form of driving. In effect, a single driver—the operating system—will guide an entire fleet of these vehicles, determining how each vehicle executes the dynamic driving task. When the fully functioning operating system causes a crash, the vehicle was engaged in systemized driving that should be evaluated through performance data for the fleet, regardless of the particular circumstances of the crash. Aggregate driving data can resolve otherwise difficult tort questions.*

*Most importantly, the manufacturer would necessarily satisfy its tort obligation regarding the reasonably safe programming or design of the operating system if the aggregate, premarket testing data sufficiently show that the fleet of fully functioning autonomous vehicles performs at least twice as safely as conventional vehicles. To avoid liability for the crash of such a vehicle, the manufacturer must also adequately warn consumers about this inherent risk. Once again, the risk involves systemized driving performance, and so aggregate driving data provide the appropriate measure. Based on these data, auto insurers can establish the risk-adjusted annual premium for insuring the vehicle. By disclosing such a premium to consumers, the manufacturer would satisfy its obligation to warn about the inherent risk of crash, eliminating this final source of manufacturer liability for crashes caused by a fully functioning autonomous vehicle.*

*The collective learning of state tort law can then inform federal regulations governing the reasonable safety of automated driving technologies. The foregoing analysis is based on tort rules that have been widely adopted across the country. States that do not follow the majority approach might reach different conclusions. To ensure that manufacturers face uniform obligations across the national market, the National Highway Transit Safety Administration could adopt two federal regulations that clearly fit within its proposed regulatory approach. Each derives from the associated tort obligations concerning adequate premarket testing and disclosure of the inherent risk of crash, respectively. These regulations would largely retain the role of tort law, because regulatory compliance would also satisfy the associated tort obligations in most states, while impliedly preempting these claims in the remaining states. The regulations would promote the federal interest in uniformity in a manner that minimizes the displacement of state tort law, thereby optimally solving the federalism problem.*

*State tort law can then supplement the federal regulations in important instances, yielding a comprehensive regulatory approach. Within this legal framework, a regulatory-compliant autonomous vehicle would subject the manufacturer to tort liability only for crashes caused by malfunctioning physical hardware (strict products liability); malfunctions of the operating system due to either programming error (same) or third-party hacking (strict liability again, with an important caveat); the manufacturer's failure to adopt a reasonably safe design or to provide adequate warnings for ensuring safe deployment of the vehicle (an ordinary products liability claim); or the manufacturer's failure to treat consumers and bystanders equally when designing the vehicle and its operating system (an ordinary negligence claim). A manufacturer would also be subject to tort liability for not complying with the federal regulations (negligence per se). The potential liabilities would not be overly uncertain. Autonomous vehicles can be regulated in a manner that ensures reasonable safety without impeding the development of this life-saving technology.*

Introduction .....	1614
I. Manufacturer Responsibility for Automated Driving Technologies .....	1624
A. Driver-Assistance Systems in Conventional Motor Vehicles....	1624
B. Automated Technologies That Eliminate the Human Driver ....	1629
II. Manufacturer Liability for the Crash of an Autonomous Vehicle .....	1632
A. Crashes Caused by Programming Bugs .....	1634
B. Crashes Caused by a Fully Functioning Operating System .....	1635
1. Product Malfunctions and the Role of Product Warnings ...	1636
2. Defective Design and the Role of Premarket Testing .....	1641
C. Crashes Causing Injury to Bystanders .....	1647
D. Satisfying Tort Obligations with Aggregate Performance Measures .....	1650
1. Adequate Premarket Testing.....	1651
2. Adequate Warnings About the Inherent Risk of Crash.....	1654
III. Manufacturer Liability for the Crash of a Hacked Vehicle.....	1660
A. The Crash of a Hacked Vehicle as a Product Malfunction .....	1663
B. Potential Limitations of Liability to Negligence.....	1669
IV. Reducing Uncertainty by Coordinating State Tort Law with Federal Safety Regulations.....	1674
A. Federal Regulations Requiring Premarket Testing and Post- Sale Updates of the Operating System.....	1678
B. Federal Regulations Requiring Product Warnings.....	1681
C. Coordination of Federal and State Law .....	1684
1. Overlap of Federal Regulatory Law and State Tort Law.....	1684
2. Federal Regulatory Law and State Tort Law as	

Supplements.....	1688
Conclusion .....	1691

## INTRODUCTION

Autonomous vehicles will be a disruptive technology. In addition to liberating humans from the task of driving, the technology will cause a migration from private car ownership to commercial car-sharing services, alter the dynamics and underlying infrastructures of urban and suburban living, and—most importantly for present purposes—substantially reduce the carnage on our roadways.<sup>1</sup>

Motor vehicle crashes in 2013 killed 32,719 people domestically while injuring another 2.3 million. These crashes caused an estimated annual economic cost of \$242 billion (or \$784 for every person in the U.S.) in addition to \$594 billion of noneconomic costs involving the decreased quality or loss of life.<sup>2</sup> The number of fatalities rose sharply to an estimated 40,200 in 2016, an increase that experts attribute in part to distracted driving.<sup>3</sup> In a detailed study of individual cases, the National Highway Traffic Safety Administration (NHTSA) found that “94 percent of crashes can be tied to a human choice or error.”<sup>4</sup>

---

1. For discussion of important social impacts likely to be caused by autonomous vehicles, see EXEC. OFFICE OF THE PRESIDENT: PRESIDENT’S COUNCIL OF ADVISORS ON SCI. & TECH., REPORT TO THE PRESIDENT: TECHNOLOGY AND THE FUTURE OF CITIES (2016), [https://www.whitehouse.gov/sites/whitehouse.gov/files/images/Blog/PCAST\\_Cities\\_Report\\_FINAL.pdf](https://www.whitehouse.gov/sites/whitehouse.gov/files/images/Blog/PCAST_Cities_Report_FINAL.pdf) [<https://perma.cc/2ZSA-AWYH>]; DANIEL J. FAGNANT & KARA M. KOCKELMAN, ENO CTR. FOR TRANSP., PREPARING A NATION FOR AUTONOMOUS VEHICLES: OPPORTUNITIES, BARRIERS AND POLICY RECOMMENDATIONS 3–10 (2013), <https://www.enotrans.org/etl-material/preparing-a-nation-for-autonomous-vehicles-opportunities-barriers-and-policy-recommendations> [<https://perma.cc/F4AY-CKNU>]; JAMES MANYIKA ET AL., MCKINSEY & CO., A FUTURE THAT WORKS: AUTOMATION, EMPLOYMENT, AND PRODUCTIVITY (2017), <http://www.mckinsey.com/global-themes/digital-disruption/harnessing-automation-for-a-future-that-works> [<https://perma.cc/DYB5-9GA9>]; Christopher Mims, *Driverless Cars to Fuel Suburban Sprawl*, WALL ST. J. (June 20, 2016), <http://www.wsj.com/articles/driverless-cars-to-fuel-suburban-sprawl-1466395201> [<https://perma.cc/88CH-PSWC>].

2. NAT’L HIGHWAY TRAFFIC SAFETY ADMIN., OVERVIEW OF NHTSA PRIORITY PLAN FOR VEHICLE SAFETY AND FUEL ECONOMY, 2015 TO 2017, at 2 (2015), [https://www.nhtsa.gov/staticfiles/nvs/pdf/NVS\\_priority-plan-June2015\\_final.pdf](https://www.nhtsa.gov/staticfiles/nvs/pdf/NVS_priority-plan-June2015_final.pdf) [<https://perma.cc/72JL-QYM8>].

3. Neal E. Boudette, *U.S. Traffic Deaths Rise for a Second Straight Year*, N.Y. TIMES (Feb. 15, 2017), <https://www.nytimes.com/2017/02/15/business/highway-traffic-safety.html> [<https://perma.cc/VL8Z-LHMU>].

4. NAT’L HIGHWAY TRAFFIC SAFETY ADMIN., FEDERAL AUTOMATED VEHICLES POLICY: ACCELERATING THE NEXT REVOLUTION IN ROADWAY SAFETY 5 (2016), <https://www.transportation.gov/AV> [<https://perma.cc/VSX9-B6J6>] [hereinafter NHTSA, 2016 AUTOMATED VEHICLES POLICY].

Autonomous vehicles would not eliminate all of these crashes,<sup>5</sup> but they should significantly enhance motor vehicle safety. According to one forecast that predicts widespread conversion to autonomous vehicles by 2040, the technology could reduce crash frequency per vehicle by 80 percent.<sup>6</sup> Another projection estimates that once 90 percent of the vehicles on the road are autonomous, 21,700 fewer domestic fatalities will occur each year.<sup>7</sup> Autonomous vehicles will save lives and prevent many more injuries, making a compelling safety case for policies that foster the widespread deployment of this technology.

The technology is developing rapidly. In September 2016, the car-sharing service Uber began using self-driving vehicles in Pittsburgh with “a safety driver in the front seat because [these vehicles still] require human intervention in many conditions.”<sup>8</sup> At the 2017 annual Detroit auto show, Waymo (formerly the self-driving car division of Google) unveiled a minivan manufactured by Fiat Chrysler and outfitted with Waymo’s self-driving equipment, illustrating the potential for strategic alliances between the technology and car-manufacturing sectors.<sup>9</sup> In canvassing industry-wide developments, one analyst found that “Tesla Motors Inc., BMW, Ford Motor Co., and Volvo Cars have all promised to have fully autonomous cars on the road within five years. General Motors Co., Daimler AG, Toyota Motor Corp., and Volkswagen AG’s Audi luxury line are pouring billions of dollars into developing autonomous

---

5. Actuaries who have re-examined the NHTSA study “found that 49% of accidents contain at least one limiting factor that could disable the technology or reduce its effectiveness. 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.” CAS. ACTUARIAL SOC’Y AUTOMATED VEHICLES TASK FORCE, RESTATING THE NATIONAL HIGHWAY TRANSPORTATION SAFETY ADMINISTRATION’S NATIONAL MOTOR VEHICLE CRASH CAUSATION SURVEY FOR AUTOMATED VEHICLES 1 (2014), [https://www.casact.org/pubs/forum/14fforum/CASAVTF\\_Restated\\_NMVCCS.pdf](https://www.casact.org/pubs/forum/14fforum/CASAVTF_Restated_NMVCCS.pdf) [<https://perma.cc/DX7X-LXHY>].

6. JERRY ALBRIGHT ET AL., KPMG, MARKETPLACE OF CHANGE: AUTOMOBILE INSURANCE IN THE ERA OF AUTONOMOUS VEHICLES 26 (2015), <https://assets.kpmg.com/content/dam/kpmg/pdf/2016/06/id-market-place-of-change-automobile-insurance-in-the-era-of-autonomous-vehicles.pdf> [<https://perma.cc/A57Z-UM8Q>]; see also Michele Bertonecello & Dominik Wee, *Ten Ways Autonomous Driving Could Redefine the Automotive World*, MCKINSEY & CO. (June 2015) (projecting reductions in motor vehicle crashes of up to 90 percent following widespread adoption of driverless vehicles), <http://www.mckinsey.com/industries/automotive-and-assembly/our-insights/ten-ways-autonomous-driving-could-redefine-the-automotive-world> [<https://perma.cc/V7ZC-VL2T>].

7. FAGNANT & KOCKELMAN, ENO CTR. FOR TRANSP., *supra* note 1, at 8 tbl. 2.

8. Anthony Levandowski & Travis Kalanick, *Pittsburgh, Your Self-Driving Uber is Arriving Now*, UBER (Sept. 14, 2016), <https://newsroom.uber.com/pittsburgh-self-driving-uber> [<https://perma.cc/5H49-LMR3>].

9. See Bill Vlasic, *Self-Driving Minivan Could Steer Car Industry*, N.Y. TIMES, Jan. 9, 2017, at B1; see also Tim Higgins, *Autos: Autonomy for Cars Attracts Suppliers’ Attention*, WALL ST. J., Jan. 10, 2017, at B6 (discussing decisions and strategies of automotive suppliers, like Delphi, to invest in automated driving technologies because “[t]he value-add is shifting toward the smarts’ of the car” as one executive put it).

vehicles.”<sup>10</sup> “[B]ased on a wave of recent developments and investments in this sector,” one forecaster recently decided to substantially increase its projections of autonomous vehicle sales; it now predicts a 43 percent annual compound rate in growth between 2025–35, culminating in sales of twenty-one million autonomous vehicles globally by 2035.<sup>11</sup> According to other forecasts, by 2035 autonomous vehicles will range from 10 percent of all new car sales worldwide<sup>12</sup> to upwards of 75 percent of all light-duty vehicle sales.<sup>13</sup>

As autonomous vehicles become more common on the roadways, the substantial reduction in the number of crashes will substantially decrease both the cost of and need for personal auto insurance. One recent projection shows that by 2040, the market for such insurance could shrink by up to 40 percent of its current size.<sup>14</sup>

Increased demand for other lines of insurance could somewhat offset the reduced need for personal auto insurance. “[C]ommercial lines could take a larger share, as the marketplace moves towards vehicle sharing and mobility on demand. As the vehicle makes more decisions, the potential liability of the . . . manufacturer will increase too.”<sup>15</sup> Any increase in the liabilities of manufacturers will presumably increase their demand for insurance covering those liabilities, further altering the overall composition of the insurance market.

The extent of liability potentially faced by manufacturers could also have a substantial impact on the emerging market for automated driving technologies. “[W]ho should be held liable for crashes involving [autonomous vehicles]? Though manufacturers, insurers, news outlets, and academics have all posed this question, they have not found easy answers.”<sup>16</sup> Consequently, “[w]hile technology is usually described as an enabler of autonomous vehicles,

10. Keith Naughton, *Here’s Where the Self-Driving Car Stands Right Now*, BLOOMBERG, QUICKTAKE (Dec. 14, 2016), <https://www.bloomberg.com/news/articles/2016-12-14/here-s-where-the-self-driving-car-stands-right-now> [<https://perma.cc/HPB2-STUB>].

11. Michelle Cullver, *IHS Clarifies Autonomous Vehicles Sales Forecast—Expects 21 Million Sales Globally in the Year 2035 and Nearly 76 Million Sold Globally Through 2035*, IHS MARKIT (June 9, 2016), <http://news.ihsmarket.com/press-release/automotive/autonomous-vehicle-sales-set-reach-21-million-globally-2035-ihs-says> [<https://perma.cc/PB34-ZJMC>]; Casualty Actuarial Society, *Actuaries Grapple with Insurance Questions on Self-Driving Cars*, INS. J. (May 16, 2014), <http://www.insurancejournal.com/news/national/2014/05/16/329422.htm> [<https://perma.cc/9G8J-KJZH>].

12. Pail Lienert, *12 Million Driverless Cars to Be on the Road by 2035—Study*, REUTERS (Jan. 8, 2015), <http://www.reuters.com/article/2015/01/08/autos-bcg-autonomous-idUSL1N0UN2GQ20150108> [<https://perma.cc/YCW2-BD7U>] (describing results from study conducted by Boston Consulting Group).

13. *Autonomous Vehicles Will Surpass 95 Million in Annual Sales by 2035*, NAVIGANT RES. (Aug. 21, 2013), <http://www.navigantresearch.com/newsroom/autonomous-vehicles-will-surpass-95-million-in-annual-sales-by-2035> [<https://perma.cc/T45S-YTGJ>].

14. ALBRIGHT ET AL., KPMG, *supra* note 6, at 27.

15. *Id.* at 28.

16. Jack Boeglin, *The Costs of Self-Driving Cars: Reconciling Freedom and Privacy with Tort Liability in Autonomous Vehicle Regulation*, 17 YALE J.L. & TECH. 171, 174 (2015).

liability is often described as an impediment.<sup>17</sup> Carmakers at an industry show described autonomous vehicles as “a future that won’t materialize . . . unless legislators around the world create a new legal framework.”<sup>18</sup> Sensationalizing this theme further, news outlets have published stories worrying about the prospect that lawsuits will “kill” the autonomous car.<sup>19</sup>

Of course, manufacturers will ultimately adopt automated driving technology—the commercial upside is too great—but substantial uncertainty about the potential scope of manufacturer liabilities could still impede the widespread deployment of autonomous vehicles. The rate at which the market converts from conventional to autonomous vehicles depends on the price that consumers must pay to adopt the new technology. For at least two reasons, systemic legal uncertainty about manufacturer liability increases the cost of an autonomous vehicle, thereby increasing price and reducing consumer demand for this technology.

Predictable risks are fundamentally different from uncertainties; the former can be calculated with reliable probabilities, whereas the latter cannot.<sup>20</sup> The difficulty of making decisions in the face of uncertainty is starkly illustrated by the evolving debate over climate change. Manufacturers face the

---

17. JOHN VILLASENOR, BROOKINGS INST., PRODUCTS LIABILITY AND DRIVERLESS CARS: ISSUES AND GUIDING PRINCIPLES FOR LEGISLATION 11 (2014), <https://www.brookings.edu/research/products-liability-and-driverless-cars-issues-and-guiding-principles-for-legislation> [<https://perma.cc/PE2Q-E6JL>]. A survey conducted by the world’s largest professional association for the advancement of technology asked more than 200 of its members to assign a ranking to six possible roadblocks to the mass adoption of driverless cars. “[L]egal liability, policymakers, and consumer acceptance were ranked as the biggest obstacles, while cost, infrastructure, and technology were viewed as the smaller speed bumps.” *IEEE Survey Reveals Mass-Produced Cars Will Not Have Steering Wheels, Gas/Brake Pedals, Horns, or Rearview Mirrors* by 2035, IEEE (July 14, 2014), [http://www.ieee.org/about/news/2014/14\\_july\\_2014.html](http://www.ieee.org/about/news/2014/14_july_2014.html) [<https://perma.cc/FXV8-DX6W>]; see also, e.g., Clifford Winston & Fred Mannering, *Implementing Technology to Improve Public Highway Performance: A Leapfrog Technology from the Private Sector Is Going to Be Necessary*, 3 ECON. TRANSP. 158, 164 (2014) (“[T]he major obstacle to motorists and firms adopting [autonomous vehicles] as soon as possible is whether the government will take prudent and expeditious approaches to help resolve important questions about assigning liability in the event of an accident, the availability of insurance, and safety regulations.”).

18. Ryan Nakashima, *Carmakers at Nevada Show: Driverless Cars Need Legal Framework*, INS. J. (Jan. 13, 2014), <http://www.insurancejournal.com/news/west/2014/01/13/316913.htm> [<https://perma.cc/GD5L-WFDW>].

19. See, e.g., Holman W. Jenkins, Jr., *Will Tort Law Kill Driverless Cars?*, WALL ST. J., Dec. 18, 2013, at A15.

20. The distinction was famously developed by Frank Knight, who argued that [r]isk [is] characterized by the reliability of the estimate of its probability and therefore the possibility of treating it as an insurable cost. The reliability of the estimate [comes] from either knowledge of the theoretical law it obeyed or from stable empirical regularities. . . . True uncertainty is to be “radically distinguished” from calculable risks: here “there is *no valid basis of any kind* for classifying instances [as required by the calculation of risk].” George J. Stigler, *Knight, Frank Hyneman (1885–1962)*, in THE NEW PALGRAVE DICTIONARY OF ECONOMICS 749 (Steven N. Durlauf & Lawrence E. Blume eds., 2d ed. 2008) (quoting FRANK KNIGHT, RISK, UNCERTAINTY AND PROFIT 225, 231 (1921)) (paragraph structure omitted).

same type of problem when trying to assess their potential liabilities for autonomous vehicles.

Like the costs of producing and commercially distributing an autonomous vehicle, a manufacturer's liability costs have a direct impact on profits. How much should the manufacturer charge for its expected liabilities? Unable to base this decision on reliable statistics, the manufacturer must instead come up with its best estimate. An estimate that is too high would artificially increase the price of the autonomous vehicle relative to other vehicles, thereby reducing aggregate demand with the attendant impact on profits. An estimate that is too low, however, would not cover actual costs and would also reduce profits. Unable to reliably predict its liability costs, the manufacturer cannot reliably predict its profits, a problem that is particularly pronounced for potentially extensive liabilities. The resultant variability in expected profits increases the risk of the underlying investment and requires higher returns to justify the added risk—an increase in the cost of capital that will be impounded into the price of the autonomous vehicle.

Manufacturers can mitigate risk by insuring against tort liabilities, which works well when the individualized liability risks of different manufacturers collectively balance out in the pool of policyholders.<sup>21</sup> But when there is a fundamental disagreement about the underlying liability rules, the uncertainty is systemic and cannot be eliminated by the pooling of individual risks within an insurance scheme. The cost of uncertainty is instead passed onto the insurer, causing it to increase the premium above the price otherwise charged for the same total amount of expected liabilities calculated with a higher degree of certainty.<sup>22</sup> By increasing either the cost of insurance or the related cost of capital for manufacturers, systemic uncertainty about liability could significantly increase prices for autonomous vehicles and unduly delay their widespread deployment.

Even if adoption of a particular tort rule eliminated this source of uncertainty, another one remains. A legal question is deeply unsettled when it could be plausibly resolved in substantially different ways. The more difficult a tort problem, the more likely that it will be initially resolved in an erroneous manner. The potential for legal error is then compounded by the need for courts to resolve this issue for each body of state tort law. As compared to a relatively “easy” problem, courts across the country are more likely to adopt different

---

21. See Mark A. Geistfeld, *Interpreting the Rules of Insurance Contract Interpretation*, 68 RUTGERS U. L. REV. 371, 383–91 (2015) (explaining why insurable risks are independent across policyholders, thereby enabling the insurer to confidently predict its expected costs of coverage by distributing the risk across the pool of policyholders).

22. See Mark A. Geistfeld, *Legal Ambiguity, Liability Insurance, and Tort Reform*, 60 DEPAUL L. REV. 539, 549–56 (2011). The uncertainty that inheres in the pricing problem then destabilizes the supply of insurance, contributing to a dynamic in which the industry cycles between “hard” and “soft” markets characterized by substantial swings in the price and availability of coverage. See *id.* at 556–64.



rules for solving a difficult tort issue, creating substantial variability within the national market. Courts will presumably correct mistakes over time, but the prospect of initial legal error and widespread disagreement creates an additional source of uncertainty for manufacturers trying to assess their potential liability in the national market.

Manufacturers now confront these forms of systemic legal uncertainty. To date, scholars have reached “the shared conclusion” that elimination of a human driver will shift responsibility onto manufacturers as a matter of products liability law, with most tort litigation involving claims for design or warning defects.<sup>23</sup> Beyond these general conclusions, “existing predictions part ways.”<sup>24</sup> How, exactly, will claims for design or warning defects be framed? Will these liabilities be common for manufacturers? Does the crash of an autonomous vehicle necessarily mean that its design is defective? What type of warning is required in these cases? On these and related matters, scholars have reached a wide range of conclusions.<sup>25</sup> “Imagine a robot car with no one behind

---

23. DOROTHY J. GLANCY ET AL., TRANSPORTATION RESEARCH BOARD, A LOOK AT THE LEGAL ENVIRONMENT FOR DRIVERLESS VEHICLES 35 (2016), [http://onlinepubs.trb.org/onlinepubs/nchrp/nchrp\\_lrd\\_069.pdf](http://onlinepubs.trb.org/onlinepubs/nchrp/nchrp_lrd_069.pdf) [<https://perma.cc/7MVJ-TYMY>].

24. *Id.* at 36.

25. See, e.g., Ryan Abbott, *The Reasonable Computer: Disrupting the Paradigm of Tort Liability*, GEO. WASH. L. REV. (forthcoming 2017), [http://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2877380](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2877380) [<https://perma.cc/URJ4-L5FG>] (concluding that manufacturers of autonomous vehicles would be subject to strict liability under current standards and proposing that liability instead be based on a negligence standard that treats the vehicle as a person); Boeglin, *supra* note 16, at 186–87 (“[P]roducts liability suits are often prohibitively expensive and may be a bad fit for the frequent litigation that car accidents instigate.”); Kyle Colonna, Note, *Autonomous Cars and Tort Liability*, 4 CASE W. RES. J.L. TECH. & INTERNET 81, 102 (2012) (arguing that products liability will increase manufacturer costs and “hamper[] the entrance of autonomous cars into the marketplace,” thereby justifying a limitation of liability); Sophia H. Duffy & Jamie Patrick Hopkins, *Sit, Stay, Drive: The Future of Autonomous Car Liability*, 16 S.M.U. SCI. & TECH. L. REV. 453, 479 (2013) (concluding that “[e]xisting laws governing vehicles and computers do not provide the proper means to assess liability for autonomous cars” and that the owner should be strictly liable for crashes); Kevin Funkhouser, Note, *Paving the Road Ahead: Autonomous Vehicles, Products Liability, and the Need for a New Approach*, 1 UTAH L. REV. 437, 440 (2013) (arguing that products liability law is “ill-prepared” to address the potential claims involving autonomous vehicles and proposing a “no-fault compensation system that can promote the interests of manufacturers and plaintiffs alike”); Andrew P. Garza, Note, “*Look Ma, No Hands!*”: *Wrinkles and Wrecks in the Age of Autonomous Vehicles*, 46 NEW ENG. L. REV. 581, 581 (2012) (“Products liability law is capable of handling the advent of autonomous vehicles just as it handled seat belts, airbags, and cruise control.”); Julie Goodrich, Comment, *Driving Miss Daisy: An Autonomous Chauffeur System*, 51 HOUS. L. REV. 265, 284 (2013) (arguing that because autonomous vehicles provide the same social benefits as vaccines—both reduce the incidence of physical harms—legislators should consider immunizing autonomous vehicles from civil liability under a legislative scheme like the National Childhood Vaccination Injury Act of 1986, which immunizes vaccine manufacturers from civil liability for unavoidable injury); Kyle Graham, *Of Frightened Horses and Autonomous Vehicles: Tort Law and Its Assimilation of Innovations*, 52 SANTA CLARA L. REV. 1241, 1270 (2012) (predicting that the first tort suits against the manufacturers of autonomous vehicles will involve failure to warn claims because claims for defective design may require plaintiffs to “engage in a searching review of the computer code that directs the movement of these vehicles,” which is likely to be “difficult, and expensive”); F. Patrick Hubbard, “*Sophisticated Robots*”: *Balancing Liability*,

the wheel hitting another driverless car. Who's at fault? The answer: No one knows."<sup>26</sup>

The uncertainty largely stems from the complexity of driving behavior. "The road, more than simply a system of regulations and designs, is a place where many millions of us, with only loose parameters for how to behave, are thrown together daily in a kind of massive petri dish in which all kinds of uncharted, little-understood dynamics are at work."<sup>27</sup> Behavioral dynamics considerably add to the complexity of driving and help to explain why human error is the underlying cause of so many crashes. In a world of autonomous vehicles, driver error will be eliminated, but the problem of human error will remain. Computer software determines the driving performance of an autonomous vehicle. The coding of driving behavior shifts the source of error from a human driver to those who program, design, and build autonomous vehicles. Autonomous vehicles will not be perfectly safe; they will inevitably fail at times. Given the complexity of driving and the inherent limitations of coding that behavior, how can courts reliably determine whether such a

---

*Regulation, and Innovation*, 66 FLA. L. REV. 1803, 1806, 1850 (2014) (concluding that "the legal system fairly allocates the costs of injuries from innovation in robots [like autonomous vehicles] and has not unduly hindered innovation"); Dylan LeValley, Note, *Autonomous Vehicle Liability—Application of Common Carrier Liability*, 36 SEATTLE U. L. REV. 5, 6 (2013) (arguing that tort law should deem the manufacturers of autonomous vehicles to be "common carriers" that owe "the public the highest duty of care [and are] liable for even the slightest negligence"); Gary E. Marchant & Rachel A. Lindor, *The Coming Collision Between Autonomous Vehicles and the Liability System*, 52 SANTA CLARA L. REV. 1321, 1333–35 (2012) (concluding that for accidents resulting from situations that the programming did not anticipate, plaintiffs can regularly recover under products liability for defective product design, thereby "present[ing] a serious barrier for the production and development of autonomous vehicles, even if the products are socially beneficial overall"); Andrzej Rapaczynski, *Driverless Cars and the Much Delayed Tort Law Revolution* 1, 9–10 (Colum. Law & Econ., Working Paper No. 540, 2016), [http://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2764686](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2764686) [<https://perma.cc/XG3F-T3B5>] (arguing that "the advent of self-driving cars . . . is likely to force a comprehensive re-thinking of products liability," resulting in "a large-scale return to the principle of strict manufacturers' responsibility"); Bryant Walker Smith, *Automated Driving and Product Liability*, 1 MICH. ST. L. REV. (forthcoming 2017) (manuscript at 2), [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2923240](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2923240) [<https://perma.cc/DF6A-BADD>] (concluding "that the current product liability regime, while imperfect, is probably compatible with the adoption of automated driving systems").

Automated driving will pose a host of other liability issues, including hard questions about how to apportion responsibility among the manufacturer and other entities within the supply chain. See generally Daniel A. Crane et al., *A Survey of Legal Issues Arising from the Deployment of Autonomous and Connected Vehicles* (Mich. Law Pub. Law & Legal Theory Research Paper Series, Paper No. 510, 2016), <http://ssrn.com/abstract=2807059> [<https://perma.cc/22MQ-GPC9>] (providing a comprehensive survey of these issues). In the event that an autonomous vehicle crashes, however, the liability of these other entities ultimately depends on the predicate question of whether the manufacturer is subject to liability—the fundamental question addressed by this Article.

26. Keith Naughton & Margaret Cronin Fisk, *Driverless Cars Give Lawyers Bottomless List of Defendants*, INS. J. (Dec. 22, 2015), <http://www.insurancejournal.com/news/national/2015/12/22/392781.htm> [<https://perma.cc/9GRM-YJEA>].

27. TOM VANDERBILT, TRAFFIC: WHY WE DRIVE THE WAY WE DO (AND WHAT IT SAYS ABOUT US) 6 (2008).

failure—the crash of a fully functioning autonomous vehicle—was caused by a defect that subjects the manufacturer to liability?

Even if an autonomous vehicle is properly coded, its driving performance will still often be opaque to consumers—another potential source of liability. “Autonomous vehicles are composed of electronics, software, sensors, and mechanical parts. Simply by observing such a machine, a person will not intuitively know where the machine will move next.”<sup>28</sup> When the safety performance of a product is not well understood by the average user, the manufacturer is obligated to provide a warning about the associated foreseeable risks of physical harm.<sup>29</sup> How could the manufacturer of an autonomous vehicle adequately warn about the full range of risky driving behaviors across a complex operating environment? Like the potential liabilities regarding the design or coding of the autonomous vehicle, the liabilities for inadequate warning would also seem to be fundamentally uncertain due to the complexity of driving behavior.

Although the legal uncertainty manufacturers now face would appear to be substantial, it is an open question whether the uncertainty is more imagined than real. Prior legal analyses have not comprehensively examined the different reasons why an autonomous vehicle can crash. Doing so requires detailed study of the varied technologies of automated driving and how they are likely to be governed by established tort doctrines across the full range of crash scenarios. The resultant assessment of liabilities will necessarily be predictive—no settled case law addresses these exact questions—yet it is still possible to draw conclusions about the plausibility and significance of the potential uncertainties. As I will attempt to show, established tort doctrines supplemented by a few important forms of federal safety regulation provide a comprehensive regulatory approach that largely dissipates the costly legal uncertainty now looming over this emerging technology.

Much of the solution resides in the basic technology of automated driving, which simplifies the tort problem in a manner not previously identified. Driving behavior today is individualized in the basic sense that the risk of crash for each vehicle largely depends on the behavior of each driver, requiring case-by-case analysis of crashes. Autonomous vehicles will transform this individualized driving behavior into a collective, systemized form. In effect, an entire fleet will be guided by a single driver—the hardware and software that determines how this class of autonomous vehicles executes the dynamic driving task, what we will call the vehicle’s “operating system” for obvious

---

28. Harry Surden & Mary-Anne Williams, *Technological Opacity, Predictability, and Self-Driving Cars*, 38 CARDOZO L. REV. 121, 127 (2016).

29. RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 2(c) (AM. LAW INST. 1998).

reasons.<sup>30</sup> Different manufacturers will presumably deploy different operating systems, and even the same manufacturer may utilize different operating systems for different versions of its autonomous vehicles. Each vehicle with the same operating system, however, systematically executes the dynamic driving task in the same manner. Quite unlike the crashes of conventional vehicles that require case-by-case analysis of driver behavior, the crashes of autonomous vehicles with fully functioning operating systems are properly evaluated in relation to the systemized performance of the entire fleet of vehicles with the same operating system. Prior analyses have largely missed the manner in which systemized driving alters the tort inquiry, which explains why the liability issues now appear to be so difficult and riddled with uncertainty.

Due to the systemized driving behavior of autonomous vehicles, manufacturers can rely on aggregate driving data to satisfy their otherwise vexing tort obligations to design these vehicles in a reasonably safe manner and to warn about the inherent risk of crash. These determinate safety measures can then inform the federal regulation of autonomous vehicles while also clearly demarcating the future role for automobile insurance. Although the technology will have social impacts that no one can fully predict at this point, autonomous vehicles pose liability questions that can be largely resolved with a sufficiently high degree of certainty.

Part I discusses manufacturer responsibility for automated driving technologies, distinguishing between two basic types of technologies that fundamentally differ with respect to the interface between the vehicle and its human operator or driver. One type relies on a human driver as backup, requiring transitions with the attendant possibilities for behavioral errors that can cause the vehicle to crash. The associated liabilities are neither novel nor likely to be a source of significant legal uncertainty. The other type does not have this interface and instead eliminates the human driver, thereby creating a new set of legal questions. Despite the absence of established precedent, there is no doubt that manufacturers will be subject to various tort duties involving the driving performance of autonomous vehicles. A tort duty, however, does not necessarily entail liability for all crashes, which makes it necessary to determine the conditions under which manufacturers are likely to incur liability.

Part II evaluates problems that could cause an autonomous vehicle to crash. The most significant source of legal uncertainty stems from the manufacturer's potential liabilities for crashes caused by a fully functioning operating system. In these cases, the autonomous vehicle was engaged in systemized driving performance that can be evaluated with aggregate driving data for the fleet. The safety performance of each autonomous vehicle within

---

30. See HOD LIPSON & MELBA KURMAN, *DRIVERLESS: INTELLIGENT CARS AND THE ROAD AHEAD* 66–67 (2016) (adopting this term and explaining the similarities and differences between the operating system of a driverless vehicle and the operating system of a computer).

the fleet continuously improves as the operating system “learns” from its collective driving experience. Thus, the question of whether the coding or design of the vehicle is reasonably safe reduces to the question of whether the operating system was subject to adequate premarket testing. When the market is transitioning to autonomous vehicles from conventional vehicles, the requisite amount of premarket testing can be determined by comparing the operating system’s collective safety performance with the associated aggregate crash data for conventional vehicles. To satisfy its obligation to warn about the inherent risk of crash, the manufacturer can disclose a different measure based on aggregate performance data—the annual, risk-adjusted premium for insuring the vehicle—which can be feasibly derived from the fleet’s performance during the premarket testing phase.

Autonomous vehicles, however, will create a new threat. An autonomous vehicle could crash because a third party hacks into the operating system and executes commands that cause a collision. The cybersecurity of these vehicles can be compromised in other ways as well. As Part III shows, liability for these crashes could be extensive due to the prospect that courts will find this type of product performance to be a “malfunction” subject to strict liability. Although courts have not rigorously defined a product malfunction, Part III more fully develops the doctrine and shows that it provides a compelling rationale for making the manufacturer strictly liable for these crashes, with one important caveat. These liabilities are not necessarily limited to isolated incidents; hacking exploits a vulnerability in the cybersecurity of the operating system that could place the entire fleet at risk. A rule of strict liability, therefore, could generate an unpredictable, systemic form of extensive liability that would undermine market stability. This outcome depends on empirical questions that cannot be resolved at this point. If the problem exists, it should be addressed by tort doctrines that immunize safety-enhancing products—in this case autonomous vehicles—from such a rule of strict liability, subjecting manufacturers to the more limited rule of negligence liability. But unlike the prior conclusions that clearly derive from established tort doctrines, this one is much more debatable. Crashes caused by hacking generate hard problems about cybersecurity and immunity from liability, making them a plausible source of significant legal uncertainty for manufacturers.

Part IV then addresses a different source of uncertainty based on the prospect that state courts across the country will resolve these liability questions in different ways. The foregoing analysis is based on established tort doctrines adopted by most, but not necessarily all, of the states. To prevent variations in tort obligations across the country, NHTSA—the branch of the U.S. Department of Transportation with primary responsibility for roadway

safety<sup>31</sup>—could draw on the collective learning of state tort law to inform federal regulations governing the reasonable safety of automated driving technologies. These regulations would largely retain the role of tort law, because regulatory compliance would also satisfy the associated tort obligations in most states, while impliedly preempting these claims in the remaining states. State tort law would then fill in gaps. The resultant regime should largely dissipate the legal uncertainty that now looms over this developing technology, thereby reducing costs and helping to speed the emergence of automated driving and the associated reduction of motor vehicle crashes.

## I.

### MANUFACTURER RESPONSIBILITY FOR AUTOMATED DRIVING TECHNOLOGIES

Autonomous vehicles can be developed in two different ways. “The first involves gradually improving the automated driving systems available in conventional vehicles so that human drivers can shift more of the dynamic driving task to these systems. The second involves deploying vehicles without a human driver and gradually expanding this operation to more contexts.”<sup>32</sup> Each type of technology creates a different behavioral interface between the human operator and the vehicle, which in turn has different implications for the manufacturer’s legal responsibilities.

#### A. *Driver-Assistance Systems in Conventional Motor Vehicles*

Driver-assistance systems (DAS) are incorporated into conventional vehicles and are capable of taking over one or more functions of the dynamic driving task under certain operating conditions.<sup>33</sup> Examples of DAS currently

---

31. NHTSA’s legislative purpose is to “reduce traffic accidents and deaths and injuries resulting from traffic accidents.” 49 U.S.C. § 30101 (2012). To do so, NHTSA is authorized to “prescribe motor vehicle safety standards for motor vehicles and motor vehicle equipment in interstate commerce.” *Id.* § 30101(1). “Motor vehicle safety” for this purpose is defined as the “performance of a motor vehicle or motor vehicle equipment in a way that protects the public against unreasonable risk of accidents occurring because of the design, construction, or performance of a motor vehicle, and against unreasonable risk of death or injury in an accident, and includes nonoperational safety of a motor vehicle.” *Id.* § 30102(a)(8). A “motor vehicle safety standard” is “a minimum standard for motor vehicle or motor vehicle equipment performance.” *Id.* § 30102(a)(9). The regulations that NHTSA adopts are incorporated into 49 C.F.R. §§ 501–508. Within this statutory scheme, “NHTSA’s authority is broad enough to address a wide variety of issues affecting the safety of vehicles equipped with [automated driving] technologies and systems.” Stephen P. Wood et al., *The Potential Regulatory Challenges of Increasingly Autonomous Motor Vehicles*, 52 SANTA CLARA L. REV. 1423, 1501 (2012).

32. OECD INT’L TRANSP. FORUM, AUTOMATED AND AUTONOMOUS DRIVING: REGULATION UNDER UNCERTAINTY 13 (2015), [http://www.internationaltransportforum.org/pub/pdf/15CPB\\_AutonomousDriving.pdf](http://www.internationaltransportforum.org/pub/pdf/15CPB_AutonomousDriving.pdf) [<https://perma.cc/K6F8-FZ79>].

33. Bryant Walker Smith, *How Governments Can Promote Automated Driving*, N.M. L. REV. (forthcoming) (manuscript at 10), [http://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2749375](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2749375) [<https://perma.cc/NKA7-UFJA>]. DAS can be distinguished from automated emergency intervention

on the market include antilock braking systems (first introduced in 1978), electronic stability control (1995), parking-assistance systems (mid-1990s), adaptive cruise control (1999), lane-departure warning systems (2001), and forward-collision prevention systems (both long range, introduced in 2003–06, and short range, 2010).<sup>34</sup> “Other extensions of current DAS are soon to come. Examples include an assistant for collision avoidance by evasive steering, assistants for the detection of oncoming traffic and pedestrians under adverse vision (weather) conditions, or assistants for improved intersection safety.”<sup>35</sup>

These new technologies are not exposed to uncertain forms of legal liability for a clear-cut reason: “These functions improve the interface between the driver and the vehicle in such a way as to provide better control or more convenient operation but do not fundamentally alter the roles of the driver and vehicle in executing the [dynamic driving task].”<sup>36</sup> Humans are still behind the wheel, so vehicles equipped with DAS have not created liability issues fundamentally different from those posed by conventional vehicles not equipped with this technology.

To be sure, the nature of the safety problem is likely to change as DAS more fully develop. Automated driving technologies can be classified by the extent to which they reduce the role of the human driver. For example, the classification scheme NHTSA adopted ranges from no vehicle autonomy (level 0) to full vehicle autonomy under all conditions in which a human could otherwise perform the driving task (level 5). Levels 2 and 3 involve limited autonomous driving that requires the human operator to monitor conditions and assume control if necessary, and level 4 involves full vehicle autonomy only within certain operating conditions.<sup>37</sup> Both level 2 and level 3 DAS create an interface between automated driving and conventional driving—the point at which the human takes over the dynamic driving task from the automated vehicle. The switch from one driving mode to the other presents a safety problem that does not exist in conventional vehicles lacking this technology.

The sustained autonomous operation of these vehicles can make the person behind the wheel complacent or otherwise overly reliant on the DAS. What if road conditions suddenly change and require human intervention, but the driver is not sufficiently attentive to quickly take over the wheel?

---

systems that provide support to a human driver by warning of impending collisions and so on. *See id.* (manuscript at 12). The difference between the two types of safety technologies does not affect the ensuing analysis, and since each one effectively assists the driver of an otherwise conventional vehicle, both are treated as forms of DAS for present purposes.

34. Klaus Bengler et al., *Three Decades of Driver Assistance Systems: Review and Future Perspectives*, 6 IEEE INTELLIGENT TRANSP. SYS. MAG. 6, 7–9 (2014).

35. *Id.* at 10 (citations omitted).

36. CRASH AVOIDANCE METRICS P’SHIP (CAMP) AUTOMATED VEHICLE RESEARCH (AVR) CONSORTIUM, KEY CONSIDERATIONS IN THE DEVELOPMENT OF DRIVING AUTOMATION SYSTEMS 2 (2014), <http://www-esv.nhtsa.dot.gov/Proceedings/24/files/24ESV-000451.PDF> [<https://perma.cc/MZ62-TK7C>].

37. NHTSA, 2016 AUTOMATED VEHICLES POLICY, *supra* note 4, at 9.

According to one study, “[i]n our current environment, over 30% of accidents involve a behavioral characteristic that may cause the automated vehicle to be used incorrectly.”<sup>38</sup>

Manufacturers currently disagree about the best approach for addressing this problem. Google and Volvo have concluded that the safest route is to bypass levels 2 and 3 DAS altogether in favor of fully autonomous vehicles (levels 4 and 5) that by definition do not rely on DAS and instead eliminate any chance for human driving error.<sup>39</sup> Other manufacturers are trying to reduce these errors by incorporating fault-tolerant planning into levels 2 and 3 DAS. “It’s the kind of planning that can handle a certain number of deviations or errors by the person who is asked to execute the plan.”<sup>40</sup> The reasonable safety of these technologies depends on the capacity of fault-tolerant design to help ensure that the person behind the wheel will take over the driving task when necessary.

This safety issue arose in an accident that occurred in May 2016 when Joshua Brown was killed while behind the wheel of a Tesla Model S electric sedan in self-driving mode (level 2)—the first known fatal accident involving a self-driving car.<sup>41</sup> According to Tesla, “the vehicle was on a divided highway with Autopilot engaged when a tractor trailer drove across the highway perpendicular to the Model S. Neither Autopilot nor the driver noticed the white side of the tractor trailer against a brightly lit sky, so the brake was not applied.”<sup>42</sup> According to news reports, Brown was watching a Harry Potter movie at the time of the crash.<sup>43</sup>

Without directly pinning full responsibility for the crash on Brown, Tesla observed that

[w]hen drivers activate Autopilot, the acknowledgment box explains, among other things, that Autopilot “is an assist feature that requires you to keep your hands on the steering wheel at all times,” and that “you need to maintain control and responsibility for your vehicle”

38. CAS. ACTUARIAL SOC’Y AUTOMATED VEHICLES TASK FORCE, *supra* note 5, at 13.

39. John R. Quain, *Makers of Self-Driving Cars Ask What to Do with Human Nature*, N.Y. TIMES (July 7, 2016), <https://www.nytimes.com/2016/07/08/automobiles/wheels/makers-of-self-driving-cars-ask-what-to-do-with-human-nature.html> [https://perma.cc/353F-ZCTB].

40. *Programming Safety into Self-Driving Cars*, NAT’L SCI. FOUND. (Feb. 2, 2015), [http://www.nsf.gov/mobile/discoveries/disc\\_summ.jsp?cntn\\_id=134033&org=NSF](http://www.nsf.gov/mobile/discoveries/disc_summ.jsp?cntn_id=134033&org=NSF) [https://perma.cc/4HNR-A2YB]; *see also* GLANCY ET AL., *supra* note 23, at 76 (discussing the role of human factors in the design of automated driving technologies).

41. *See* Bill Vlasic & Neal E. Boudette, *Self-Driving Tesla Was Involved in Fatal Crash, U.S. Says*, N.Y. TIMES (June 30, 2016), <https://www.nytimes.com/2016/07/01/business/self-driving-tesla-fatal-crash-investigation.html> [https://perma.cc/L6QB-QX9Y].

42. *A Tragic Loss*, TESLA (June 30, 2016), <https://www.teslamotors.com/blog/tragic-loss> [https://perma.cc/P3NA-DJAR].

43. Will Oremus, *The Tesla Autopilot Crash Victim Was Apparently Watching a Movie When He Died*, SLATE (July 1, 2016), [http://www.slate.com/blogs/moneybox/2016/07/01/tesla\\_autopilot\\_crash\\_victim\\_joshua\\_brown\\_was\\_watching\\_a\\_movie\\_when\\_he\\_died.html](http://www.slate.com/blogs/moneybox/2016/07/01/tesla_autopilot_crash_victim_joshua_brown_was_watching_a_movie_when_he_died.html) [https://perma.cc/VW77-WQYW].



while using it. Additionally, every time that Autopilot is engaged, the car reminds the driver to “Always keep your hands on the wheel. Be prepared to take over at any time.” The system also makes frequent checks to ensure that the driver’s hands remain on the wheel and provides visual and audible alerts if hands-on is not detected. It then gradually slows down the car until hands-on is detected again.

We do this to ensure that every time the feature is used, it is used as safely as possible.<sup>44</sup>

Do these measures adequately ensure that the person behind the wheel remains alert? Or does some or all responsibility for the crash instead fall on the failure of Joshua Brown to manually apply the brakes in time? The issues are novel in the sense that they are not implicated by the crash of a conventional vehicle lacking level 2 or 3 DAS. The liability question, however, is not new.

Established tort doctrine already obligates manufacturers to adopt fault-tolerant product designs. As the *Restatement (Third) of Torts* explains, “instructions and warnings may be ineffective because users of the product may not be adequately reached, may be likely to be inattentive, or may be insufficiently motivated to follow the instructions or heed the warnings.”<sup>45</sup> Consequently, “when a safer design can reasonably be implemented and risks can reasonably be designed out of a product, adoption of the safer design is required over a warning that leaves a significant residuum of such risks.”<sup>46</sup> A manufacturer that does not adopt a reasonably safe, fault-tolerant design is subject to tort liability for the resultant physical harms.

This doctrine is capable of adequately addressing the safety question of how DAS design should address the foreseeable risk that the technology will lull the driver into complacency or inattention. The tort inquiry involves a cost-benefit analysis (known as the risk-utility test), which requires the product design to incorporate any safety feature costing less than the associated safety benefit.<sup>47</sup> For example, a machine manufacturer could provide a warning to consumers—“avoid contact with the exposed moving parts of this machine”—

---

44. TESLA, *supra* note 42. Tesla also provided extensive warnings about the limits of the autopilot system, several of which “apply directly to the situation apparently faced by the driver in this crash.” Chris Ziegler, *Tesla’s Own Autopilot Warnings Outlined Deadly Crash Scenario*, VERGE (June 30, 2016), <http://www.theverge.com/2016/6/30/12073240/tesla-autopilot-warnings-fatal-crash> [https://perma.cc/R33H-X73V].

45. RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 2 cmt. 1 (AM. LAW INST. 1998).

46. *Id.* A different outcome would occur under patent danger rule, which eliminates the design obligation for risks that are open and obvious, including those disclosed in the warning. “A strong majority of courts have rejected the ‘open and obvious’ or ‘patent danger’ rule as an absolute defense to a claim of design defect.” *Id.* § 2 reps. n. cmt. d(IV)(C).

47. See MARK A. GEISTFELD, PRINCIPLES OF PRODUCTS LIABILITY 164–69 (2d ed. 2011) (rigorously developing the risk-utility inquiry for fault-tolerant design) [hereinafter PRODUCTS LIABILITY].

or instead incorporate a safety feature into the design—a guard to prevent users from inadvertently coming into contact with that part of the machine. If the guard were a cost-effective method for reducing this risk, then a manufacturer that only provided a warning would be subject to liability for defective design. So, too, manufacturers cannot merely warn drivers to stay “alert” in order to take over the driving responsibilities when prompted by DAS; they must instead adopt fault-tolerant designs for DAS whenever doing so would be a cost-effective method for reducing the risk of driver error.

A reasonably safe fault-tolerant design ultimately implicates the coding of the DAS in the vehicle. For example, General Motors is planning to implement an operating system for its semi-autonomous vehicles containing software that “can detect whether a driver is dozing off or not watching the road” and then use “audible and visual alerts to grab the person’s attention. If the alerts don’t work, a representative [of the manufacturer] will activate the vehicle’s intercom and communicate with the car’s operator. If the driver still doesn’t respond, the car will pull over on the side of the freeway and stop.”<sup>48</sup>

Another coding option is to forgo DAS altogether by eliminating the human driver—creating fully autonomous vehicles (levels 4 and 5). Although manufacturers sharply differ about the desirability of this approach, the disagreement will not translate into a highly uncertain form of tort liability for vehicles equipped with DAS.

Consider the crash of a vehicle with level 2 or 3 DAS, like the one that killed Joshua Brown. Under established doctrine, the plaintiff in such a case cannot claim that the design of the vehicle is defective because its reliance on a human driver makes it unreasonably dangerous as compared to a fully autonomous vehicle. In effect, such a claim is one of categorical liability, alleging that any product within one product category (vehicles equipped with DAS) is unreasonably dangerous as compared to those in another product category (fully autonomous vehicles). To preserve the role of informed consumer choice across product categories, courts have roundly rejected claims of this type.<sup>49</sup>

The term “category” is analytical for this purpose, defined by the ordinary consumer’s informed choice that fully accounts for the safety decision implicated by the plaintiff’s tort claim.<sup>50</sup> In deciding whether to purchase a

---

48. Gauthem Nagesh, *Business News: GM’s Eye-Tracking Tech Aims to Keep Drivers Alert*, WALL ST. J., Sept. 12, 2016, at B3.

49. Although numerous courts and the *Restatement (Third)* recognize that claims of categorical liability can be appropriate in some cases, only a few cases have imposed such liability. More searching analysis shows that the validity of such claims is limited to cases of bystander injuries in which recovery is based on the allegation that consumers unreasonably disregard the safety interests of these third parties and therefore should not be given the choice in question. See GEISTFELD, PRODUCTS LIABILITY, *supra* note 47, at 311–19. The issue of bystander liability is analyzed separately in *infra* Part II.C.

50. See GEISTFELD, PRODUCTS LIABILITY, *supra* note 47, at 125–35.

conventional vehicle equipped with DAS instead of a fully autonomous vehicle, the ordinary consumer considers the costs and benefits of the two vehicle types. The consumer can make an adequately informed safety decision based on any evident risk differentials and the warnings that manufacturers are obligated to provide about all other risks that would be material to the safety decision. The consumer's informed choice to purchase the DAS-equipped vehicle (one product category) instead of a fully autonomous vehicle (another category) forecloses liability based on the premise that consumers should not be given such a choice. Manufacturers instead are obligated to provide designs that are reasonably safe for products within the same category, eliminating any significant potential for liabilities based on the claim that a vehicle equipped with DAS is defectively designed simply because it is not a fully autonomous vehicle.

Despite the risk of driver error, strict products liability will not force manufacturers to forego further technological development of DAS. Manufacturers of these technologies must instead design the vehicle's operating system to account for the human errors that will foreseeably occur at the interface of automated driving.

#### *B. Automated Technologies That Eliminate the Human Driver*

Once automated driving technologies fully take over the dynamic driving task, a legal discontinuity occurs. A vehicle is autonomous in the sense that it can drive without human assistance (or indeed, any human in the vehicle at all). This definition of autonomy matters for tort purposes because one can incur tort liability only through the exercise of autonomous agency.<sup>51</sup> When the vehicle's occupant is no longer executing the dynamic driving task, human driving error is no longer the cause of an accident. Instead, the manner in which the vehicle executed the driving task becomes the focus of inquiry. The vehicle, however, cannot be legally responsible for its performance (it is, after all, not truly autonomous), which leads to the question of who should be responsible for the vehicle's operation: The consumer of the product (the owner and, potentially, users) or the manufacturer (and other entities in the chain of distribution)? Autonomous vehicles raise questions of legal responsibility fundamentally different from those involving conventional vehicles.

To allay consumer concerns, leading manufacturers have already recognized that they are legally responsible for the driving behavior of their autonomous vehicles. "Volvo Cars, Google and Daimler AG's Mercedes-Benz

---

51. See, e.g., MAYO MORAN, *RETHINKING THE REASONABLE PERSON: AN EGALITARIAN RECONSTRUCTION OF THE OBJECTIVE STANDARD* 21 (2003) ("[L]iability in negligence requires a minimum capacity for rational agency. . . . Because they cannot meet the threshold 'agency' requirement, children of 'tender years' (approximately 5 years and below) are typically totally immune from liability in negligence. But beyond this category, courts and commentators are divided over what is sufficient to negate the presumption of agency and thus preclude liability in negligence.").

have all pledged to accept liability if their vehicles cause an accident.”<sup>52</sup> By accepting legal responsibility for operation of their autonomous vehicles, these manufacturers are sending a signal to consumers about safety: “‘We want customers to trust we’ve done a really good job,’ said Anders Eugensson, Volvo’s director of government affairs. ‘That’s why we say if anything happens, we assume liability.’”<sup>53</sup>

Lest there be any doubt about the matter, NHTSA has ruled that Google’s self-driving car is the equivalent of a human driver for federal regulatory purposes.<sup>54</sup> The logic of this ruling readily resolves the associated tort questions, further establishing that the manufacturer will be legally responsible for the driving behavior of an autonomous vehicle.<sup>55</sup>

This tort obligation is beyond serious doubt, even though there is no established body of case law recognizing that a manufacturer incurs a tort duty for defective software.<sup>56</sup> In general, the tort duty for software designers can be limited for various reasons, most notably the economic loss rule that limits consumers to contractual remedies for intangible forms of intellectual property that have been designed for a specific purpose.<sup>57</sup> Relying on the policy rationales for strict products liability, others have argued that these reasons for limiting the tort duty should not apply to defective software that foreseeably causes physical harms.<sup>58</sup> The rationale for the tort obligation, however, is much more straightforward in the case of autonomous vehicles.

---

52. Naughton & Fisk, *supra* note 26; see also *Volvo Cars Responsible for the Actions of its Self-Driving Cars*, VOLVO CARS (Oct. 20, 2015), <http://www.volvocars.com/intl/About/Our-Innovation-Brands/IntelliSafe/IntelliSafe-Autopilot/News/Volvo-Cars-responsible-for-the-actions-of-its-self-driving-cars> [https://perma.cc/7SZA-JBG8] (“Volvo Cars will accept full liability for the actions of its autonomous cars when in Autopilot mode, making it one of the first manufacturers to take this vital step forward in the development of self-driving cars.”).

53. Naughton & Fisk, *supra* note 26.

54. Letter from Paul A. Hemmersbaugh, Chief Counsel, Nat’l Highway Traffic Safety Admin., to Chris Urmson, Director, Google, Inc. Self-Driving Car Project (Feb. 4, 2016), [http://isearch.nhtsa.gov/files/Google -- compiled response to 12 Nov 15 interp request -- 4 Feb 16 final.htm](http://isearch.nhtsa.gov/files/Google--compiled%20response%20to%2012%20Nov%2015%20interp%20request--4%20Feb%2016%20final.htm) [https://perma.cc/ZK69-EZLZ].

55. See Mark A. Geistfeld, *Tort Law in the Age of Statutes*, 99 IOWA L. REV. 957, 963–67 (2014) (discussing the common law principle that courts will defer to any legislative policy judgment that is relevant to the resolution of a tort claim) [hereinafter *Age of Statutes*].

56. “Despite the fact that discussions of liability for defective software go back more than forty years, very few cases have addressed the issue outside the financial services context.” Alan Butler, *Products Liability and the Internet of (Insecure) Things: Should Manufacturers Be Liable for Damage Caused by Hacked Devices?*, U. MICH. J.L. REFORM (forthcoming) (manuscript at 103), [http://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2955317](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2955317) [https://perma.cc/KJ6G-WP65].

57. Frances E. Zollers et al., *No More Soft Landings for Software: Liability for Defects in an Industry That Has Come of Age*, 21 SANTA CLARA COMPUTER & HIGH TECH. L.J. 745, 758–60, 764 (2005).

58. See, e.g., Butler, *supra* note 56 (manuscript at 103–04) (arguing that the tort duty can be justified because the risk of property damage is foreseeable, software defects can be remedied by remote updates, and “holding manufacturers liable for downstream harms caused by their insecure devices is well aligned with the purposes of products liability law”); Zollers et al., *supra* note 57, at

The coding or design of the operating system determines the performance of a product—a motor vehicle. Although the coding is an intangible form of intellectual property developed for a specific purpose, these are not sufficient reasons for eliminating the tort duty. If they were, then a conventional motor vehicle that performs according to engineering plans that were developed or otherwise embodied in a computer program would also be exempt from tort liability. Regardless of the form taken by a product design, manufacturers are responsible for ensuring that the design causes the product to perform in a reasonably safe manner.<sup>59</sup> In cases of physical harm, this tort duty requires the manufacturer to adopt a reasonably safe design for the operating system, an obligation that is not negated by the economic loss rule or contractual provisions that disclaim the manufacturer's liability.<sup>60</sup>

Responsibility, however, does not necessarily entail legal liability in the event of a crash. Liability depends on both the existence of a tort duty and its breach. Human drivers are responsible for their conduct behind the wheel, but that does not mean they are legally liable anytime they are involved in a crash. What, then, are the liability implications for manufacturers that are responsible for how autonomous vehicles execute the dynamic driving task?

As compared to conventional vehicles equipped with DAS technologies, the interface between the operator and an autonomous vehicle poses a relatively easy safety problem. The human operator inputs the destination information into the vehicle. That behavior could be unreasonably dangerous only if the destination requires the vehicle to travel outside its parameters for safe operation. For the near future, autonomous vehicles will be capable of safe operation only under defined environmental conditions.<sup>61</sup> Driving in a city is different from driving in the mountains. Navigating under severe weather conditions can be problematic in both places. Warning consumers about the appropriate conditions for deploying an autonomous vehicle will accordingly be required to ensure reasonable safety. By adequately instructing the occupant on the appropriate operating conditions and programming the vehicle to override instructions to operate in an unsafe environment, the manufacturer

---

782 (concluding that “the policy reasons underlying strict [products] liability are congruent with the application of the doctrine to software”).

59. See RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 2(b) (AM. LAW INST. 1998) (defining the duty to design without any limitations regarding the form of the design).

60. The economic loss rule does not apply to cases of physical harm—bodily injury or damage to real or tangible property other than the product itself. *Id.* § 21 cmt. a. Any contractual limitation of liability “do[es] not bar or reduce otherwise valid products liability claims against sellers or other distributors of new products for harms to persons.” *Id.* § 18.

61. See Lee Gomes, *Google Self-Driving Cars Will Be Ready Soon for Some, in Decades for Others*, IEEE SPECTRUM (Mar. 18, 2016), <http://spectrum.ieee.org/cars-that-think/transportation/self-driving/google-selfdriving-car-will-be-ready-soon-for-some-in-decades-for-others> [https://perma.cc/TU6B-53BL] (reporting that it might be up to thirty years before the Google self-driving car will be widely available and that until then, the technology will be incrementally introduced based on geography and weather conditions).

will satisfy its tort obligation of reasonably ensuring that a person does not make a mistake in deciding when, where, or how to deploy the vehicle.

A properly deployed autonomous vehicle can crash for many reasons. For tort purposes, a manufacturer's responsibility is limited to crashes that were proximately caused by the vehicle.<sup>62</sup> The mere fact that a vehicle was involved in a crash is not sufficient for this purpose. For example, the distracted driver of a conventional vehicle could suddenly veer into the path of an autonomous vehicle. The autonomous vehicle would not be a proximate cause of the ensuing crash—its driving behavior did not increase the likelihood that the distracted driver would veer into its path, nor did it have the opportunity to avoid that outcome.<sup>63</sup> The manufacturer's responsibility, and hence potential liability, is limited to crashes caused by the risks foreseeably created by the autonomous vehicle's driving behavior.<sup>64</sup>

## II.

### MANUFACTURER LIABILITY FOR THE CRASH OF AN AUTONOMOUS VEHICLE

Once the operator has properly deployed an autonomous vehicle, the manufacturer becomes primarily responsible for the vehicle's driving performance. For centuries, tort law has required manufacturers and other product sellers to ensure that products perform in a reasonably safe manner. Today, the most important obligation is the rule of strict products liability, which the vast majority of states adopted in the 1960s and 1970s.<sup>65</sup> This rule subjects the commercial distributor of a product to strict liability for the physical harms proximately caused by a defect in the product.<sup>66</sup> After struggling with the appropriate definition of defect—the predicate for strict liability—courts ultimately adopted a tri-partite definition that distinguishes among manufacturing, design, and warning defects.<sup>67</sup>

---

62. See RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 15 (AM. LAW INST. 1998) (“Whether a product defect caused harm to persons or property is determined by the prevailing rules and principles governing causation in tort.”).

63. See RESTATEMENT (THIRD) OF TORTS: LIABILITY FOR PHYSICAL AND EMOTIONAL HARM § 30 (AM. LAW INST. 2010) (“An actor is not liable for harm when the tortious aspect of the actor's conduct was of a type that does not generally increase the risk of that harm.”).

64. “Currently, virtually all jurisdictions employ a foreseeability (or risk) standard for some range of scope-of-liability issues in negligence cases.” *Id.* § 29 cmt. e. The foreseeability or risk standard also applies to forms of strict liability. *Id.*; see also Mark Geistfeld, *Implementing Enterprise Liability: A Comment on Henderson and Twerski*, 67 N.Y.U. L. REV. 1157, 1162–67 (1992) (explaining how proximate cause limits the rule of strict manufacturer liability for product-caused injuries). The foreseeability standard governs determinations of liability in the first instance—the subject of our inquiry—but not issues concerning the extent of damages caused by the predicate, foreseeable physical harm. See MARK A. GEISTFELD, *TORT LAW: THE ESSENTIALS* 255–68 (2008) [hereinafter *TORT LAW*].

65. GEISTFELD, *PRODUCTS LIABILITY*, *supra* note 47, at 10–17 (discussing adoption of strict products liability and its doctrinal heritage).

66. RESTATEMENT (SECOND) OF TORTS § 402A (AM. LAW INST. 1965).

67. See RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 2 (AM. LAW INST. 1998).

A manufacturing (or construction) defect exists “when the product departs from its intended design even though all possible care was exercised in the preparation and marketing of the product.”<sup>68</sup> Defects of this type occur for different reasons. Materials or component parts of the product can be contaminated or otherwise manufactured in a flawed manner due to an error in the production process; the product can be improperly assembled or constructed; or the product can be improperly packaged. Because these defects depart from design specifications, they exist only in aberrant products that would not satisfy quality-control standards. A commercial distributor of the defective product would be subject to strict tort liability in most states.

Within an autonomous vehicle, a defect of this type will not implicate the software that executes the dynamic driving task. To be sure, the vehicle’s operating system may have a programming bug caused by a typo, but that coding is still part of the operating system, making it part of the vehicle’s design. *All* vehicles with this operating system would contain the coding error, unlike manufacturing defects that affect only particular products within the entire product line.<sup>69</sup> A manufacturer’s liability for manufacturing defects will be largely limited to quality-control problems with the hardware of the operating system, including the cameras, lidar (laser scanning), radar, and other physical components of the system that do not perform according to design.<sup>70</sup>

As applied to autonomous vehicle hardware, the liability rule is no different from the one already governing defects in the physical components of conventional motor vehicles. Not only are manufacturers quite familiar with this form of liability, they can also largely control their liability exposure for manufacturing defects by adopting quality-control measures and purchasing insurance to cover the remaining liabilities. This rule of strict products liability is well established and does not plausibly contribute to the legal uncertainty that could impede the development of autonomous vehicles.

For these reasons, we can limit the inquiry to cases in which the design (or programming) of the operating system causes the autonomous vehicle to crash. The technology can also cause other types of harms, but the most significant concern for manufacturers involves potential liabilities for crashes. To identify the conditions under which a manufacturer would be subject to liability, we must consider the different ways in which the programming of an autonomous vehicle could cause a crash.

---

68. *Id.* § 2(a) (defining “manufacturing defect”).

69. *See* Hubbard, *supra* note 25, at 1854 (providing more extended discussion reaching the same conclusion); Zollers et al., *supra* note 57, at 749 (“Software can only fail for one reason: faulty design.”).

70. For “an in-depth look at the suite of hardware devices that provide data to the car’s operating system,” see LIPSON & KURMAN, *supra* note 30, at 171–96.

### A. Crashes Caused by Programming Bugs

Consider a programming error or bug in the software that causes the operating system to crash, in turn causing the vehicle to crash. In these cases, the plaintiff would not have to identify the specific coding error and could instead prove defective design solely based on the manner in which the operating system misperformed.<sup>71</sup>

According to the *Restatement (Third) of Torts*, product performance is a sufficient substitute for direct proof of defect when it “was of a kind that ordinarily occurs as a result of product defect; and . . . was not, in the particular case, solely the result of causes other than a product defect existing at the time of sale or distribution.”<sup>72</sup> Because the defect in these cases is inferred from the product misperformance, the *Restatement (Third)* calls such performance a “malfunction,”<sup>73</sup> a usage adopted by some courts and commentators.<sup>74</sup> Regardless of the label, this widely adopted doctrine subjects manufacturers to liability for product malfunctions.<sup>75</sup>

Under the formulation adopted by the *Restatement (Third)*, the malfunction doctrine is limited to “situations in which a product fails to perform its manifestly intended function.”<sup>76</sup> For example, a manufacturer manifestly intends the airbags in a vehicle to safely deploy in certain types of crashes defined by the design parameters. Consequently, courts have found product malfunctions when “an air bag fails to deploy, deploys improperly, or spews acid on an occupant.”<sup>77</sup>

Based on this definition, a malfunction would occur if a coding error caused the operating system to crash, resulting in a crash of the autonomous vehicle. The coding error prevented the operating system from performing its manifestly intended function of executing the dynamic driving task, subjecting the manufacturer to liability for the crash.

“The cause of action is one involving true ‘strict’ liability, since recovery may be had upon a showing that the product was not minimally safe for its expected purpose—without regard to the feasibility of alternative designs or the manufacturer’s ‘reasonableness’ in marketing it in that unsafe condition.”<sup>78</sup> The

---

71. RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 3 cmt. b (AM. LAW INST. 1998) (explaining that the plaintiff can recover upon proof of product malfunction without having to “specify the type of defect responsible for the product malfunction”).

72. *Id.* § 3.

73. *Id.* § 3 cmt. b.

74. David G. Owen, *Manufacturing Defects*, 53 S.C. L. REV. 851, 873 n.123 (2002) (adopting this terminology and noting that several jurisdictions use this label, although “most courts refer to it simply as a principle of circumstantial evidence”).

75. *Id.* at 882–84.

76. RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 3 cmt. b (AM. LAW INST. 1998).

77. Owen, *supra* note 74, at 876 (footnotes omitted).

78. *Denny v. Ford Motor Co.*, 662 N.E.2d 730, 736 (N.Y. 1995); *see also Soule v. General Motors Corp.*, 882 P.2d 298, 308 (Cal. 1994).



malfunction itself is sufficient proof of defect. Under the rule of strict products liability, the manufacturer is responsible for the physical harms proximately caused by a defect, even if it “has exercised all possible care in the preparation and sale of [the] product.”<sup>79</sup>

### B. Crashes Caused by a Fully Functioning Operating System

Unless the technology has been perfected, the operating system of an autonomous vehicle will be designed or coded in a manner that is not completely safe. Even if the program contains no errors or bugs, the vehicle can confront circumstances not anticipated by the coding, resulting in an execution of the dynamic driving task that causes the vehicle to crash. In these cases, the manufacturer’s liability depends on whether such a crash was proximately caused by a defect in the design of the fully functioning operating system.

As one court observed, “the determination of when a product is actionable because of the nature of its design” is one of “the most agitated controversial questions that courts face in the field of products liability law.”<sup>80</sup> Courts have disagreed about whether defects of product design should be evaluated under the consumer expectations test, the risk-utility test, or some combination thereof. After surveying the case law, the Supreme Court of South Carolina found in 2010 that “[s]ome form of a risk-utility test is employed by an overwhelming majority of the jurisdictions in this country. Some of these jurisdictions exclusively employ a risk-utility test, while others do so with a hybrid of the risk-utility and the consumer expectations test, or an explicit either-or option.”<sup>81</sup> Based on this case law, the *Restatement (Third) of Torts* adopted the risk-utility test for defective design.<sup>82</sup> But even though only a “decided minority” of jurisdictions uses the consumer expectations test exclusively, a substantial majority of states continues to recognize it in one form or another.<sup>83</sup> Numerous state supreme courts have even reaffirmed their commitment to the consumer expectations test by rejecting the *Restatement (Third)*’s framework for defective design.<sup>84</sup> The liability rule appears to be largely unsettled across the country, making it extremely difficult for

---

79. RESTATEMENT (SECOND) OF TORTS § 402A(2)(a) (AM. LAW INST. 1965).

80. *Pritchett v. Cottrell, Inc.*, 512 F.3d 1057, 1063 (8th Cir. 2008) (brackets omitted).

81. *Branham v. Ford Motor Co.*, 701 S.E.2d 5, 14–15 (S.C. 2010) (citations omitted).

82. RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 2(b) (AM. LAW INST. 1998).

83. *Branham*, 701 S.E.2d at 14 n.12, 15 (citing seventeen different states as exclusively relying on the risk-utility test).

84. See, e.g., *Izzarelli v. R.J. Reynolds Tobacco Co.*, 136 A.3d 1232, 1240–50 (Conn. 2016) (describing risk-utility test in the *Restatement (Third)* and rejecting it in favor of a “modified consumer expectations test” that limits the risk-utility test to cases of products that do not malfunction in violation of the ordinary consumer’s minimal expectations of safe product performance); *Aubin v. Union Carbide Corp.*, 177 So.3d 489, 510–12 (Fla. 2015) (rejecting the risk-utility test in the *Restatement (Third)* in favor of the consumer expectations test); *Tincher v. Omega Flex, Inc.*, 104 A.3d 328, 399 (Pa. 2014) (same).

manufacturers to assess their potential liability exposure for autonomous vehicles in the national market.

As I have shown at length elsewhere, the apparent disparities among the rules governing defective product design largely disappear once the consumer expectations test has been adequately defined.<sup>85</sup> The test must distinguish between defects attributable to product malfunctions and defects attributable to the unreasonably dangerous design of a product that does not malfunction. The liability issues involving autonomous vehicles fully illustrate the logic of this conclusion. Regardless of whether a state relies on the consumer expectations test, the risk-utility test, or some combination thereof, the liability inquiry can be reduced to two different questions: (1) whether the crash of an autonomous vehicle is a malfunction, or (2) whether a vehicle that did not malfunction nevertheless has an unreasonably dangerous or defective design. If either of these conditions is satisfied, the manufacturer would be subject to liability for crashes proximately caused by the fully functioning operating system, regardless of the label that a court applies to the liability rules.

### 1. *Product Malfunctions and the Role of Product Warnings*

If the fully functioning operating system proximately causes an autonomous vehicle to crash, could this performance—the crash itself—constitute a malfunction subject to liability? A malfunction is defined by reference to the product’s expected performance, and so the liability question depends on how courts formulate the expectation of how a fully functioning operating system should execute the dynamic driving task for an autonomous vehicle.

---

85. See MARK A. GEISTFELD, *PRODUCT LIABILITY LAW* 69–116 (2012); GEISTFELD, *PRODUCTS LIABILITY*, *supra* note 47, 37–60, 91–110. In a series of articles, the reporters of the *Restatement (Third)* have also extensively argued that the rules regarding defective product design are largely settled once one recognizes that the “overwhelming majority of cases that rely on consumer expectations as the theory of liability do so only in res-ipsa like cases” of product malfunction. Aaron D. Twerski & James A. Henderson, Jr., *Manufacturer’s Liability for Defective Product Designs: The Triumph of Risk-Utility*, 74 BROOK. L. REV. 1061, 1108 (2009) [hereinafter *Triumph of Risk-Utility*]; see also James A. Henderson, Jr. & Aaron D. Twerski, *Achieving Consensus on Defective Product Design*, 83 CORNELL L. REV. 867, 890 (1998) (“When designs malfunction, violating built-in standards, courts often explain judgments for plaintiffs in terms of the designs having ‘disappointed consumer expectations.’ However, because such cases do not involve the application of a general design standard, it would constitute error to count such cases as support for the consumer expectations test as the general standard.”). The approach that is defended by Henderson and Twerski and adopted by the *Restatement (Third)* does not adequately account for the fundamental importance of consumer expectations, unlike the approach that will be developed in text below. See generally Mark A. Geistfeld, *The Value of Consumer Choice in Products Liability*, 74 BROOK. L. REV. 781 (2009) (showing why the important doctrines of products liability can be justified by the value of consumer choice and identifying the important ways in which the overall approach adopted by the *Restatement (Third)* obscures the essential ways in which strict products liability depends on consumer expectations) [hereinafter *Consumer Choice*].

Under the *Restatement (Third)*, the relevant expectation involves the product's failure to "perform its manifestly intended function."<sup>86</sup> For the design of an operating system, "[t]he most general objective" is for the autonomous vehicle "to safely reach the specified destination."<sup>87</sup> Indeed, "[m]any manufacturers have . . . adopted targets and plans for reaching zero injuries and fatalities. Volvo Car Corporation has adopted a target of zero serious injuries and fatalities in a new Volvo vehicle by the year 2020."<sup>88</sup> As the administrator of NHTSA explained in 2016, "[f]or more than a century, safety professionals have begun with the assumption that cars would crash, and focused their efforts on reducing the damage. Today, we can see a new possibility—the possibility that we can prevent those crashes from ever occurring."<sup>89</sup> Based on this coding objective for an autonomous vehicle's operating system, any crash arguably involves a failure of the vehicle's manifestly intended function, constituting a product malfunction that subjects the manufacturer to strict liability.

This conclusion is debatable because the rule adopted by the *Restatement (Third)* "is not ideal, which reflects the difficulty of formulating a concise, general statement of the principle."<sup>90</sup> But even if the malfunction doctrine were more rigorously defined,<sup>91</sup> manufacturers would still be subject to considerable uncertainty for a different reason.

---

86. RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 3 cmt. b (AM. LAW INST. 1998).

87. Andrei Furda & Ljubo Vlacic, *Enabling Safe Autonomous Driving in Real-World City Traffic Using Multiple Criteria Decision Making*, 3 IEEE INTELLIGENT TRANSP. SYS. MAG. 4, 10 (2011); see also, e.g., Rafael Arnav et al., *Safe and Reliable Path Planning for the Autonomous Vehicle Verdino*, 8 IEEE INTELLIGENT TRANSP. SYS. MAG. 22, 23 (2016) ("The problem we want to solve is safely following a predefined route while avoiding dynamic obstacles."). In general, any "information processing system" embodies a "computational theory [that] corresponds to the goal of computation and an abstract definition of the task." ETHEM ALPAYDIN, MACHINE LEARNING: THE NEW AI 20 (2016). For autonomous vehicles, the "abstract definition of the task" includes arriving safely at the specified destination along with other factors such as minimizing the time or length of trip.

88. Anders Eugensson et al., *Environmental, Safety, Legal and Societal Implications of Autonomous Driving Systems* (Nat'l Highway Traffic Safety Admin., Research Paper No. 13-0467, 2013), [www.nrd.nhtsa.dot.gov/pdf/esv/esv23/23ESV-000467.PDF](http://www.nrd.nhtsa.dot.gov/pdf/esv/esv23/23ESV-000467.PDF) [<https://perma.cc/KC4W-QLVL>]; see also Michael Aeberhard et al., *Experience, Results, and Lessons Learned from Automated Driving on Germany's Highways*, 7 IEEE INTELLIGENT TRANSP. SYS. MAG. 42, 50 (2015) ("The challenges in artificial intelligence for automated driving systems will always have their limits, but will also continuously improve until a level of intelligence is reached with which [highly automated driving] will be possible and where safety, within certain conditions, can be guaranteed."); Richard Waters, *CES 2016: Toyota Poaches Google Exec to Help Lead AI Effort*, FIN. TIMES, Jan. 5, 2016, at 15 (reporting that the "ultimate goal" of Toyota's driverless car program is to create "a car that cannot be responsible for a collision").

89. Mark R. Rosekind, Administrator, Nat'l Highway Traffic Safety Admin., Remarks at Autonomous Car Detroit Conference (Mar. 16, 2016), <http://www.nhtsa.gov/About+NHTSA/Speeches,+Press+Events+&+Testimonies/mr-autonomous-car-03162016> [<https://perma.cc/ADE8-55AH>].

90. Owen, *supra* note 74, at 883 n.195.

91. See *infra* Part III.A (providing a more rigorous formulation of the malfunction doctrine).

Rather than defining a malfunction by reference to the product's manifestly intended function, a substantial majority of states instead evaluate this issue with the consumer expectations test.<sup>92</sup> To satisfy this test, "the product must meet the safety expectations of the general public as represented by the ordinary consumer, not the industry or a government agency."<sup>93</sup> Under this test, the "crucial question in each individual case is whether the circumstances of the product's failure permit an inference that the product's design performed below the legitimate, commonly accepted minimum safety assumptions of its ordinary consumers."<sup>94</sup> At minimum, the ordinary consumer expects that a product will not malfunction. The frustration of that expectation supplies the rationale for subjecting the manufacturer to liability for product malfunctions.

To trigger the consumer expectations test, the ordinary consumer must only have well-formed expectations of the product performance in question; he or she does not otherwise have to understand the underlying technology (as with exploding airbags).<sup>95</sup> If the ordinary consumer does not have sufficient knowledge about how an autonomous vehicle will perform in any given respect, the manufacturer must adequately warn about the associated risks (an issue discussed below). The consumer, however, can still have minimum expectations of safe performance, including the expectation that the operating system will not malfunction because of a programming error or bug.<sup>96</sup> The question, therefore, is whether the ordinary consumer has minimum safety expectations about other aspects of the vehicle's performance.

When autonomous vehicles first become commercially available, the ordinary consumer presumably can expect the vehicle to perform at least as safely as a vehicle driven by a human driver. The ordinary consumer could also have a more demanding expectation, perhaps because such an assurance of safe performance is implicitly supplied by the manufacturer's statement that it will be legally responsible for the vehicle's driving performance.<sup>97</sup> As consumers gain more experience with autonomous vehicles, their expectations of safety will also change. Further technological development will make autonomous vehicles safer, and so those exceptional crashes that do occur are more likely to be deemed a malfunction that violates the ordinary consumer's minimum

---

92. See *supra* note 83 and accompanying text.

93. *Soule v. General Motors Corp.*, 882 P.2d 298, 306 (Cal. 1994).

94. *Id.* at 309; see also Henderson & Twerski, *Triumph of Risk-Utility*, *supra* note 85, at 1107–08 (finding that the "overwhelming majority of cases that rely on consumer expectations as the theory for imposing liability do so only in *res ipsa*-like situations in which an inference of defect can be drawn from the happening of a product-related accident").

95. See, e.g., *Green v. Smith & Nephew AHP, Inc.*, 629 N.W.2d 727, 742 (Wis. 2001) ("[A] condition not contemplated by the ordinary consumer[] does not inevitably require any degree of scientific understanding about the product itself. Rather, it requires understanding of how safely the ordinary consumer would expect the product to serve its intended purpose.").

96. See *supra* Part II.A.

97. See *supra* notes 52–53 and accompanying text.

expectations of safe performance. Paradoxically, the safe performance promised by the technology could generate demanding expectations of safety that subject the manufacturer to liability in the event of crash.<sup>98</sup>

The vagueness of the liability rule has important implications for our inquiry. Whether defined in terms of consumer expectations or the product's manifestly intended function, a malfunction would arguably occur whenever the fully functioning operating system proximately causes the autonomous vehicle to crash, potentially subjecting the manufacturer to strict liability for the resultant physical harms.

Framed in this manner, the rule of strict products liability is highly uncertain in application, explaining why there has been widespread concern about the potential liabilities faced by the manufacturers of autonomous vehicles.<sup>99</sup> The uncertainty is particularly pernicious in light of the massive extent of potential liability—the manufacturer could be liable for *all* crashes, creating costs that could plausibly impede the widespread deployment of this crash-reducing technology.

This uncertainty, however, can be eliminated through the manufacturer's satisfaction of an independent tort obligation. If a product creates a foreseeable risk of injury that is not adequately known by the ordinary consumer and that would be material to his or her decision regarding product use, the manufacturer is obligated to warn about the risk.<sup>100</sup> Satisfying the duty to warn does not necessarily satisfy the manufacturer's duty to adopt a reasonably safe or non-defective design.<sup>101</sup> But by satisfying the tort obligation to adequately warn consumers about the foreseeable risk of crash that is unavoidable or inherent in a safely designed autonomous vehicle, the manufacturer will also avoid liability for these crashes under the malfunction doctrine. In addition to establishing how the fully functioning vehicle is manifestly intended to perform in this respect, an adequate warning about the inherent risk of crash also apprises the ordinary consumer of how the vehicle will perform under these conditions. Having been adequately warned about the inherent risk of crash, the ordinary consumer cannot have frustrated expectations in the event that the risk materializes, thereby foreclosing liability under the consumer expectations test.

---

98. Cf. Jonathan J. Koehler & Andrew D. Gershoff, *Betrayal Aversion: When Agents of Protection Become Agents of Harm*, 90 *ORG. BEHAV. & HUM. DECISION PROCESSES* 244, 245–56 (2003) (providing results of five empirical studies finding that when a product causes the very harm that it was supposed to protect against, individuals treat these outcomes as forms of “betrayal” that trigger more negative responses than products that do not promise such protection).

99. See *supra* notes 16–26 and accompanying text.

100. See *infra* notes 165–68 and accompanying text (discussing the duty to warn).

101. RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 2 cmt. 1 (AM. LAW INST. 1998) (“Warnings are not . . . a substitute for the provision of a reasonably safe design.”); see also *supra* notes 45–48 and accompanying text (discussing this rule). Issues involving the reasonably safe design of the operating system are discussed in the next Section.

This conclusion directly follows from the implied warranty, which supplies the doctrinal foundation for the consumer expectations test and requires that products “be marketable with their true character known.”<sup>102</sup> For example, a manufacturer cannot avoid liability for manufacturing defects by warning that the product might contain such a defect. Each product either contains the defect or it does not, and so the warning would not reveal the true character of *any* product.<sup>103</sup> In contrast, an adequate warning about design-related performance conveys the true character of *every* product embodying the design. An adequate warning about the inherent risk that the fully functioning operating system can cause the autonomous vehicle to crash, therefore, would show that this particular vehicle is marketable with its true character known, absolving the manufacturer of liability for such a crash under the implied warranty and, by extension, the consumer expectations test.<sup>104</sup>

For example, the Connecticut Supreme Court recently held that “[a] cigarette that exposes the user to carcinogens and the attendant risk of cancer cannot be said to fail to meet an ordinary consumer’s legitimate, commonly accepted minimum safety expectations.”<sup>105</sup> The court, however, also “recognize[d] that a different result might be warranted in cases in which the plaintiff (or decedent) began smoking before warning labels were mandated by federal law.”<sup>106</sup> Warning labels shape consumer expectations. An adequate warning fully conveys the true character of each cigarette because the performance in question involves a design attribute of the product (the tobacco and chemical additives in the cigarette). Once the ordinary consumer has been adequately warned that smoking causes cancer, his or her minimum safety expectations would not be violated if that product use causes cancer. The same type of design-related product performance occurs when a fully functioning operating system causes the autonomous vehicle to crash,<sup>107</sup> further illustrating why an adequate warning about this inherent risk would foreclose liability under the malfunction doctrine.

---

102. William L. Prosser, *The Implied Warranty of Merchantable Quality*, 27 MINN. L. REV. 117, 128–29 (1943); see also GEISTFELD, PRODUCTS LIABILITY, *supra* note 47, at 10–19 (explaining why the implied warranty supplies the doctrinal basis of the consumer expectations test).

103. For example, a warning that a soda bottle *might* explode due to a manufacturing defect would not exculpate the manufacturer from tort liability under the implied warranty—it does not fully convey the true character of the particular product that actually exploded and injured the plaintiff. The warning for that particular product would instead have to say, “This bottle will explode if one attempts to open it.”

104. Cf. Prosser, *supra* note 102, at 144 (“If the buyer has examined the specific goods before purchase, it is of course clear that as to all visible defects he cannot expect [that the seller makes any representation that there are no such defects under the implied warranty.] The seller has said to him, in effect, ‘I propose to sell you what you see;’ and if he buys on such an offer, he cannot afterwards complain.”).

105. *Izzarelli v. R.J. Reynolds Tobacco Co.*, 136 A.3d 1232, 1249 (Conn. 2016).

106. *Id.* at 1249 n.16.

107. See *supra* notes 68–69 and accompanying text (explaining why issues involving the fully functioning operating system of an autonomous vehicle involve design).

This warning obligation will not end once the manufacturer has sold or otherwise commercially distributed the autonomous vehicle. The manufacturer can learn about product risks that were not disclosed in the warning issued to consumers at the time of sale. According to the *Restatement (Third) of Torts*, a manufacturer incurs a post-sale duty to warn existing consumers whenever it knows or should know of such a “substantial risk of harm to persons or property” that is “sufficiently great to justify the burden of providing a warning.”<sup>108</sup> The burden of a post-sale warning largely depends on the cost of communicating with consumers after the product has been sold. Consequently, “[f]or a post-sale duty to warn to arise, the seller must reasonably be able to communicate the warning to those identified as appropriate recipients.”<sup>109</sup> As applied to conventional products, such a warning obligation is ordinarily quite burdensome, explaining why many jurisdictions have not adopted the post-sale duty to warn. In this important respect, autonomous vehicles are different. The manufacturer will have a wireless connection with the vehicle, making it virtually costless to convey new warnings to consumers. The substantially reduced burden of complying with a post-sale duty to warn makes that obligation quite reasonable for autonomous vehicles as compared to conventional products. When confronted with the question of whether the manufacturer of an autonomous vehicle has a post-sale duty to warn, courts across the country are quite likely to answer in the affirmative.

Although the warning obligation will be ongoing, satisfying that obligation will enable manufacturers to avoid liability based on the claim that the crash of a fully functioning autonomous vehicle is a product malfunction that violates consumer expectations. *An adequate warning about the safe use and inherent risks of a safely designed autonomous vehicle will absolve the manufacturer from liability for crashes caused by the fully functioning operating system.*

## 2. Defective Design and the Role of Premarket Testing

Even if a product performs in accordance with the warning and does not otherwise malfunction, the consumer ordinarily has an independent expectation that the product design is reasonably safe. The warning only helps to establish the consumer’s minimum safety expectations of how the product will *actually* perform, which can differ from a more demanding expectation of how the product *should* otherwise perform. If the design causes the product to perform in an unreasonably dangerous manner, the actual performance would typically

---

108. RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 10(b)(1), (4) (AM. LAW INST. 1998).

109. *Id.* § 10 cmt. g.

frustrate the consumer's reasonable expectation of how the product should have performed, rendering the design defective.<sup>110</sup>

For example, suppose the manufacturer adequately warns consumers that a car does not have an airbag. Once the ordinary consumer has this knowledge, the fact that no airbag deploys in an accident could not be an unexpected misperformance (or product malfunction) that subjects the manufacturer to liability. While the consumer does not expect an airbag to deploy in an accident, she still reasonably expects that the vehicle would have a functioning airbag if that design feature were required for the reasonably safe operation of the vehicle. A warning that the car contains no airbag would not defeat this reasonable expectation of safety. By proving that the omission of the airbag renders the design unreasonably dangerous, the plaintiff would also show that this aspect of the design frustrates the ordinary consumer's reasonable expectations of safe product performance. Some courts call this liability rule the "modified" consumer expectations test in order to distinguish it from the (ordinary) consumer expectations test governing product malfunctions.<sup>111</sup>

So formulated, the modified consumer expectations test is substantively equivalent to the risk-utility test, a cost-benefit inquiry that requires any design modification with a disutility (or cost) that is less than the associated reduction of risk (or safety benefit).<sup>112</sup> The two tests are substantively equivalent because the tort burdens incurred by a manufacturer, including the cost of mandated safety investments and liabilities for injury compensation, are passed on to consumers in the form of higher prices or decreased product functionality.<sup>113</sup> Consequently, the risk-utility test does not account for the interests of

---

110. See, e.g., *Camacho v. Honda Motor Co., Ltd.*, 741 P.2d 1240, 1245 (Colo. 1987) ("A consumer is justified in expecting that a product placed in the stream of commerce is reasonably safe for its intended use, and when a product is not reasonably safe a products liability action may be maintained."); see also *supra* notes 45–48 and accompanying text (explaining why a product warning does not necessarily exculpate the manufacturer from liability for defective design). An exception would apply if the consumer knows of the design defect, has the option to use a non-defective product, and decides to use the defectively designed product anyway. See GEISTFELD, PRODUCTS LIABILITY, *supra* note 47, at 281–88 (explaining why the plaintiff can be barred from recovery for assuming the risk of defect if he or she makes an informed choice that depends on the same risk-utility factors as those implicated by the defect in question).

111. See *Izzarelli v. R.J. Reynolds Tobacco Co.*, 136 A.3d 1232, 1242 (Conn. 2016) ("Under the 'modified' consumer expectations test, the jury would weigh the product's risks and utility and then inquire, in light of those factors, whether a reasonable consumer would consider the product design unreasonably dangerous.") (quotation marks omitted); *Soule v. General Motors Corp.*, 882 P.2d 298, 308 (Cal. 1994) (same).

112. See RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 2 cmt. d (AM. LAW INST. 1998) (adopting "a reasonableness ('risk-utility balancing') test as the standard for judging the defectiveness of product designs," which "asks whether a reasonable alternative design would, at reasonable cost, have reduced the foreseeable risks of harm posed by the product and, if so, whether the omission of the alternative design by the seller or a predecessor in the distributive chain rendered the product not reasonably safe") (sentence structure omitted).

113. See GEISTFELD, PRODUCTS LIABILITY, *supra* note 47, at 38 n.7 (providing more rigorous support for this claim).



manufacturers and related parties: “[I]t is not a factor . . . that the imposition of liability would have a negative effect on corporate earnings or would reduce employment in a given industry.”<sup>114</sup> Excluding instances of bystander injuries—to be discussed below<sup>115</sup>—the risk-utility test only implicates consumer interests. A risk-utility test that is limited to consumer interests requires only those product designs or warnings that the ordinary consumer reasonably expects, making the risk-utility test fully congruent with the modified consumer expectations test.<sup>116</sup>

These two labels for the same liability rule have created the appearance that the two tests substantively differ and are a source of legal uncertainty for manufacturers. The appearance is misleading. The vast majority of courts across the country will use substantively equivalent liability rules to evaluate the design of a fully functioning autonomous vehicle. Any uncertainty about manufacturer liability must instead pertain to how courts will apply this rule.

In considering how the risk-utility test will apply to the crash of an autonomous vehicle, scholars have disagreed about the likely outcome. On one view, the risk-utility test as applied in the courtroom will routinely subject the manufacturer to liability for the crash of a fully functioning autonomous vehicle.

The problem is that most accidents will result from situations that the manufacturer or designer did not anticipate. This will open the manufacturer to second-guessing by the plaintiff’s expert that an adjustment would have provided a safer alternative system that would have avoided the accident in question. The manufacturer will almost always lose the cost-benefit argument, conducted in hindsight in the litigation context, when it focuses at the micro-scale between slightly different versions of the autonomous system. This is because the cost of not implementing the potential improvement will usually be severe—the loss of one or more lives or other serious injury, compared to the relatively small cost of the marginal improvement that might have prevented the accident. The technology is potentially doomed if there are a significant number of such cases, because the liability burden on the manufacturer may be prohibitive of further development.<sup>117</sup>

Others have questioned whether courts will apply the risk-utility test in this manner: “For a plaintiff to reach a jury on a design-defect claim, she may

---

114. RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 2 cmt. f (AM. LAW INST. 1998).

115. See *infra* Part II.C.

116. See GEISTFELD, PRODUCTS LIABILITY, *supra* note 47, at 44–48; see also 63 AM. JUR. 2D PRODUCTS LIABILITY § 554 (2008) (“The reasonable expectation of the user or consumer is to be determined through consideration of a number of factors, including the relative cost of the product, the gravity of potential harm from a claimed defect, and the cost and feasibility of eliminating or minimizing risk.”).

117. Marchant & Lindor, *supra* note 25, at 1334.

have to engage in a searching review of the computer code that directs the movement of these vehicles. This project may be difficult, and expensive,”<sup>118</sup> presumably preventing many plaintiffs from producing the evidence necessary to establish liability.

Although this debate would seem to show that manufacturers face a highly uncertain form of liability, it is based on a misconception of how manufacturers code the operating system. Properly conceptualized, the coding that determines the driving performance of an autonomous vehicle entails a well-defined risk-utility inquiry quite different from the foregoing specifications.

In spite of their disagreements, both sides of this debate implicitly assume that the operating system of an autonomous vehicle is pre-programmed only with rule-based or symbolic artificial intelligence, consisting of IF-THEN commands, such as “IF a pedestrian is sensed to be within 75 feet on the road ahead, THEN action X will be executed.”<sup>119</sup> The two sides disagree about how courts will apply the risk-utility test in so-called corner cases—the “unusual situations that are difficult to anticipate but can have potentially catastrophic results.”<sup>120</sup> Will courts rely on a “micro-scale inquiry” that isolates the cost or disutility of adding another pre-programmed rule that would have addressed the corner case and avoided the crash in question? Or will they instead engage in a more “searching review” that evaluates the program in its entirety and presumably recognizes that it is not reasonable to code rules for each and every corner case? Despite their different implications for the potential liabilities of the manufacturer, both formulations of the issue assume that the driving behavior of an autonomous vehicle is fully determined by a pre-programmed series of rule-based, IF-THEN commands that do not change after the autonomous vehicle first hits the road.

This assumption is erroneous for reasons that fundamentally alter the risk-utility analysis of an autonomous vehicle’s operating system.

[S]elf-driving vehicles do not primarily drive themselves based upon a series of pre-programmed computer rules about when and where to steer, accelerate, or brake. Rather, such systems typically use machine learning algorithms that have been “trained” to drive by analyzing examples of safe driving, and automatically generalizing about the core patterns that constitute effective driving from these examples.<sup>121</sup>

---

118. Graham, *supra* note 25, at 1270.

119. See LIPSON & KURMAN, *supra* note 30, at 76 (“Symbolic [artificial intelligence] involves breaking down a complex situation or task into a formal set of rules that a human programmer writes into software code.”).

120. *Id.* at 4.

121. Surden & Williams, *supra* note 28, at 148; see also ALEXANDER HARS, INVENTIVIO GMBH, TOP MISCONCEPTIONS OF AUTONOMOUS CARS AND SELF-DRIVING VEHICLES 4 (last modified Sept. 30, 2016), [http://www.driverless-future.com/?page\\_id=774](http://www.driverless-future.com/?page_id=774) [<https://perma.cc/UAJ8->

Machine learning is a data-driven form of artificial intelligence that “is a key catalyst behind recent advances in driverless-car performance and safety.”<sup>122</sup> For example, Google (now Waymo) incorporates machine learning into its self-driving cars<sup>123</sup> and “has driven almost two million kilometers on public roads with test drivers and has assembled an enormous fund of traffic situations from which its vehicles can learn.”<sup>124</sup> Rather than relying on a fixed set of behavioral rules (which characterize symbolic artificial intelligence), the operating system “learns” by adapting or changing the program to incorporate newly acquired information about the best way to execute the dynamic driving task.<sup>125</sup> Consequently, as Tesla explained in a press release addressing the first fatal crash of a self-driving vehicle, “[a]s more real-world miles accumulate and the software logic accounts for increasingly rare events, the probability of injury will keep decreasing.”<sup>126</sup>

The experience of an autonomous vehicle—or more precisely, all vehicles with the same operating system—provides the data for machine learning that enables the operating system to adapt accordingly.<sup>127</sup> As NHTSA explains, while “human driver[s] may repeat the same mistakes as millions before them, an [autonomous vehicle] can benefit from the data and experience drawn from thousands of other vehicles on the road.”<sup>128</sup>

Machine learning has important implications for how the risk-utility test applies to the design or programming of an operating system. A risk-utility examination of the coding itself is limited to rules that constrain or guide the machine learning, such as coding that instructs the vehicle to always stop at stop signs. Aside from these rules, autonomous vehicles are not “controlled by a detailed, exactly specified and in principle comprehensible software program. Instead we should conceptualize their behavior as being the result of a long and varied program of learning.”<sup>129</sup> In this respect, the programming of the

---

EX6N] (“[S]elf-driving vehicles are not programmed in the classical sense; they need to learn. It is not possible to reduce human driving decisions to a few (not even very many) IF-THEN rules.”).

122. LIPSON & KURMAN, *supra* note 30, at 197.

123. See Jeremy Hsu, *Deep Learning Makes Driverless Cars Better at Spotting Pedestrians*, IEEE SPECTRUM (Feb. 9, 2016), <http://spectrum.ieee.org/cars-that-think/transportation/advanced-cars/deep-learning-makes-driverless-cars-better-at-spotting-pedestrians> [https://perma.cc/5TTN-T86B]; Alexis C. Madrigal, *The Trick that Makes Google’s Self-Driving Cars Work*, ATLANTIC (May 15, 2014), <http://www.theatlantic.com/technology/archive/2014/05/all-the-world-a-track-the-trick-that-makes-googles-self-driving-cars-work/370871> [https://perma.cc/9WZ9-G6P2].

124. HARS, TOP MISCONCEPTIONS, *supra* note 121, at 4.

125. See generally ALPAYDIN, *supra* note 87, at 85–123 (describing the programming of artificial intelligence systems based on neural networks and “deep learning”); LIPSON & KURMAN, *supra* note 30, at 197–236 (explaining how artificial intelligence based on deep-learning software works and how it is incorporated into the operating systems of autonomous vehicles).

126. TESLA, *supra* note 42.

127. See LIPSON & KURMAN, *supra* note 30, at 17 (“As cars pool their driving ‘experience’ in the form of data, each car will benefit from the combined experience of all other cars.”).

128. NHTSA, 2016 AUTOMATED VEHICLES POLICY, *supra* note 4, at 5.

129. HARS, TOP MISCONCEPTIONS, *supra* note 121, at 5.

operating system is analogous to human behavior. “The secret to human success is ‘practice, practice, practice.’ [T]he secret to machine learning is pretty similar: repetition, repetition, repetition.”<sup>130</sup> Because the driving behavior of an autonomous vehicle is based on repeated driving or learning experience, “[l]ike with us humans, it then becomes difficult to answer the question why the car exhibits a specific behavior in a new situation: no ‘explicit rules’ have been specified; the decision results from the many traffic situations to which the [learning] algorithm had been exposed beforehand.”<sup>131</sup> Whether an autonomous vehicle “behaved” or performed reasonably in these cases, therefore, does not depend on a risk-utility examination of the coding (as prior legal analyses have assumed); the appropriate inquiry instead asks whether the operating system has had sufficient learning experience to drive the vehicle in a reasonably safe manner.

When an autonomous vehicle is first introduced into the market, its operating system would necessarily have the sufficient amount of learning if the vehicle had been subject to adequate premarket testing—a complex issue that we address below.<sup>132</sup> For now, however, the implications of the foregoing analysis are clear: *Except for rules that guide or constrain machine learning, whether the fully functioning operating system is defectively designed wholly depends on the adequacy of prior testing.*

In addition to subjecting the operating system to adequate premarket testing, manufacturers will also probably have to comply with an additional tort obligation to update the operating system so that it incorporates recent learning that enhances the safety performance of the autonomous vehicle. To date, tort law has only imposed a post-sale duty to warn on manufacturers, whereas an update of the operating system involves a post-sale duty of design modification. Redesigning a conventional motor vehicle after it has been sold requires a product recall that courts have concluded is best left to the regulatory process.<sup>133</sup> Although a recall is required in order to repair defects in the hardware of motor vehicles, it is not necessarily needed to update or redesign the operating system of an autonomous vehicle. Some manufacturers of automated driving technologies already use wireless updates for software systems in their vehicles, and NHTSA “envisions that manufacturers and other

---

130. LIPSON & KURMAN, *supra* note 30, at 206–07.

131. HARS, TOP MISCONCEPTIONS, *supra* note 121, at 4; *see also* ALPAYDIN, *supra* note 87, at 122 (explaining that because the factors for determining the behavior of the program “are not predefined but are automatically discovered during learning[,] they may not always be easy to interpret or assign a meaning to”); LIPSON & KURMAN, *supra* note 30, at 228 (“A deep-learning network is a classic example of what programmers call a *black-box architecture*, meaning it’s virtually impossible to reverse-engineer the steps the software program takes as it generates output.”).

132. *See infra* Part II.D.1.

133. RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 11 cmt. a (AM. LAW INST. 1998) (“Issues relating to product recalls are best evaluated by governmental agencies capable of gathering adequate data regarding the ramifications of such undertakings. The duty to recall or repair should be distinguished from a post-sale duty to warn about product hazards discovered after sale.”).

entities will likely update the vehicle's software through over-the-air updates or other means."<sup>134</sup> Based on the "proximity" afforded by automated driving technologies, courts will most likely conclude that manufacturers must make post-sale design modifications of this type.<sup>135</sup> Like the duty to warn, the duty to design will be an ongoing obligation for the manufacturers of autonomous vehicles, in this instance to provide software updates of the operating system.

### C. Crashes Causing Injury to Bystanders

Thus far, the analysis has exclusively focused on the tort liability of a manufacturer for physical harms suffered by consumers—the owner and any users of the autonomous vehicle. A crash can also injure third-party bystanders, such as pedestrians or the occupants of other cars. In cases of physical harm, the manufacturer's tort obligations encompass both consumers and bystanders,<sup>136</sup> and so we can limit the liability inquiry to bystander issues that are unique to the design of autonomous vehicles.

First consider the design or programming of the operating system. To "teach" the operating system how to drive, "the programmers feed the software with many traffic situations and specify the correct action for each situation," and the machine-learning algorithm then employs statistical analysis to determine the best way to achieve the desired outcomes.<sup>137</sup> What constitutes the "correct action" for situations in which the autonomous vehicle will inevitably crash and could injure consumers and bystanders? How should the vehicle's operating system be instructed to execute actions that can protect one party at the expense of another?

The issue has drawn a great deal of public attention,<sup>138</sup> in part because it implicates one version of the well-known moral dilemma called the "trolley problem." The dilemma is whether an individual should prevent a runaway trolley car from crashing into a group of people when doing so would cause the certain death of another person.<sup>139</sup> As applied to autonomous vehicles, the

---

134. NHTSA, 2016 AUTOMATED VEHICLES POLICY, *supra* note 4, at 16.

135. See Bryant Walker Smith, *Proximity-Driven Liability*, 102 GEO. L.J. 1777, 1785–88 (2014); cf. RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 11 cmt. a (justifying the lack of an independent tort duty to recall a product on the ground that "[i]f every improvement in product safety were to trigger a common-law duty to recall, manufacturers would face incalculable costs every time they sought to make their product lines better and safer").

136. See RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 1 (AM. LAW INST. 1998) ("One engaged in the business of selling or otherwise distributing products who sells or distributes a defective product is subject to liability for harm to persons or property caused by the defect.").

137. HARS, TOP MISCONCEPTIONS, *supra* note 121, at 4.

138. See, e.g., Larry Greenemeier, *Driverless Cars Will Face Moral Dilemmas*, SCI. AM. (June 23, 2016), <http://www.scientificamerican.com/article/driverless-cars-will-face-moral-dilemmas> [<https://perma.cc/CTV3-5ZGS>].

139. The original formulation of this problem involved the driver of the trolley, whose role is fully analogous to the driver of an autonomous vehicle (the operating system). See PHILIPPA FOOT, VIRTUES AND VICES AND OTHER ESSAYS IN MORAL PHILOSOPHY 19 (1978). The trolley problem has since been reformulated so that it involves a more difficult question of whether a bystander should

trolley problem creates a challenge for the design of the operating system: What is “correct action” for situations in which the autonomous vehicle can avoid hitting another car or pedestrian when doing so threatens grave injury to its occupants? Based on a series of survey questions involving variations of the trolley problem, one study found that participants approved of designs that would sacrifice the occupants of an autonomous vehicle to save others, although they would prefer not to ride in such vehicles and would be less willing to purchase one as a result.<sup>140</sup>

This issue, though complex and deeply interesting, is not novel. For example, crashes between a sport utility vehicle (SUV) and an ordinary automobile implicate the same problem. According to one government study, SUV designs in 1999 were causing nearly 1,000 “unnecessary deaths a year in other vehicles.”<sup>141</sup> “SUVs impose excessive collision damage because the height differential creates a mismatch between their structures and the protective structures of vehicles with lower ride-heights.”<sup>142</sup> To protect themselves from the increased risk of being injured while riding in a car, consumers have purchased SUVs for themselves. But as one empirical study has found, when “drivers shift from cars to light trucks or SUVs, each crash involving fatalities of light-truck or SUV occupants that is prevented comes at a cost of at least 4.3 additional crashes that involve deaths of car occupants, pedestrians, bicyclists, or motorcyclists.”<sup>143</sup> The then-administrator of NHTSA summed up the implications of these consumer choices: “The theory that I’m going to protect myself and my family even if it costs other people’s lives has been the operative incentive for the design of these vehicles, and that’s just wrong.”<sup>144</sup> Consumer choices can create incentives for manufacturers to adopt product designs that are unreasonably dangerous for bystanders, creating an “arms war” on the highways.

In cases of this type, courts have often dismissed the tort claims of bystanders by relying on consumer-choice doctrines—an outcome that does not

---

intervene to prevent the trolley from crashing into five workers by throwing a switch that would redirect the trolley onto a different track that will surely kill one worker instead. See Judith Jarvis Thomson, *The Trolley Problem*, 94 YALE L.J. 1395, 1397–99 (1985). This version of the trolley problem has attracted considerable attention but is not implicated by the programming of an autonomous vehicle because those who code the operating system are effectively drivers and not mere bystanders.

140. Jean-François Bonnefon et al., *The Social Dilemma of Autonomous Vehicles*, 352 SCI. 1573, 1574 (2016).

141. See Keith Bradsher, *Carmakers to Alter S.U.V.’s to Reduce Risk to Other Autos*, N.Y. TIMES, Mar. 21, 2000, at A1.

142. Howard Latin & Bobby Kasolas, *Bad Designs, Lethal Profits: The Duty to Protect Other Motorists Against SUV Collision Risks*, 82 B.U. L. REV. 1161, 1201 (2002).

143. Michelle J. White, *The “Arms Race” on American Roads: The Effect of Sport Utility Vehicles and Pickup Trucks on Traffic Safety*, 47 J.L. & ECON. 333, 334 (2004).

144. Danny Hakim, *Regulators Seek Ways to Make S.U.V.’s Safer*, N.Y. TIMES (Jan. 30, 2003), <http://www.nytimes.com/2003/01/30/business/regulators-see-ways-to-make-suv-s-safer.html> [<https://perma.cc/RQW5-UECQ>] (quoting Dr. Jeffrey W. Runge).

depend on whether the court was applying the consumer expectations test or the risk-utility test.<sup>145</sup> The animating idea is that consumers have the right to make an informed choice regarding product design. Various tort rules accordingly limit liability in order to foster informed consumer choice.<sup>146</sup> While salutary in other contexts, these consumer-choice rules are inappropriate for cases in which an injured bystander claims that consumers should not be given the choice in question. Why limit liability for a product design that is unreasonably dangerous for bystanders simply because consumers prefer their own safety over others? Once the safety problem has been framed in this manner (which is not typically the case), the answer seems obvious, yet courts have dismissed tort claims of this type.<sup>147</sup>

The case law accordingly provides some support for the proposition that the manufacturer can design the operating system to protect the occupants of an autonomous vehicle at the expense of bystanders, but that type of design will be vulnerable to a different tort claim. However formulated, the rule of strict products liability only supplements the default tort rule of negligence liability.<sup>148</sup>

The negligence rule provides clear guidance on how a manufacturer must design an autonomous vehicle to protect bystanders. As someone who would be foreseeably threatened by operation of the vehicle, a bystander is encompassed within the manufacturer's duty to exercise reasonable care in designing the vehicle. To satisfy this obligation, the manufacturer must give "impartial consideration" to the interests of bystanders, treating them no differently from its own interest in satisfying consumer demand for the product.<sup>149</sup> The manufacturer, therefore, must initially code or teach the operating system of an autonomous vehicle so that the "correct action" treats consumers and bystanders equally.

For example, an autonomous vehicle with sensors indicating that it is occupied by one person must swerve to avoid hitting a group of pedestrians, even if doing so would threaten grave injury to the lone passenger. If the "correct action" for the operating system were not specified in this manner, its programming or design would unreasonably privilege the interests of the single

---

145. See GEISTFELD, PRODUCTS LIABILITY, *supra* note 47, at 309–20; see generally Latin & Kasolas, *supra* note 142 (providing the first tort analysis of the SUV problem and arguing that courts have erroneously failed to appreciate the duty that automobile manufacturers owe to bystanders).

146. See generally Geistfeld, *Consumer Choice*, *supra* note 85 (describing the various tort doctrines that limit the manufacturer's liability when the ordinary consumer is able to make an informed choice about the safety matter in question).

147. In addition to the SUV problem discussed in text, courts have relied on consumer-choice doctrines to dismiss claims involving bullets and handguns. See GEISTFELD, PRODUCTS LIABILITY, *supra* note 47, at 309–20.

148. See 1 DAVID G. OWEN & MARY J. DAVIS, OWEN & DAVIS ON PRODUCTS LIABILITY § 2.1 (4th ed. 2014) (explaining why "the negligence cause of action remains a vital theory of recovery in products liability litigation").

149. RESTATEMENT (SECOND) OF TORTS § 283 cmt. e (AM. LAW INST. 1965).

occupant over the larger number of pedestrians. The manufacturer would then be subject to negligence liability for injuries suffered by the pedestrians. To avoid liability, the manufacturer must design the operating system of the autonomous vehicle to minimize the expected injuries from any given crash whether the potential victims are occupants or bystanders.<sup>150</sup>

In addition to the operating system, the manufacturer must design other aspects of the vehicle to reasonably account for bystander interests. The manufacturer, for example, may have to design the autonomous vehicle so that it can adequately communicate with other vehicles or pedestrians, signaling the driving behavior that they can expect.<sup>151</sup> Once again, the negligence standard of reasonable care (or its substantive equivalent, the risk-utility test) determines the requisite precautions. The obligation to design the autonomous vehicle to adequately account for bystander interests is not a plausible source of significant legal uncertainty for manufacturers.

#### *D. Satisfying Tort Obligations with Aggregate Performance Measures*

Well-established tort obligations can resolve otherwise vexing liability issues once we recognize how the systemized driving behavior of autonomous vehicles affects the tort inquiry. This conclusion, however, does not necessarily show that manufacturers can confidently assess their liability exposure. As we have found, adequate testing will satisfy the manufacturer's obligation to ensure that the operating system is reasonably safe and not defectively designed, but how much certainty does this doctrine provide? An adequate warning will defeat claims of product malfunction and otherwise satisfy the manufacturer's obligation to adequately warn about the inherent risk that the fully functioning operating system will cause the vehicle to crash, but how much certainty does this doctrine afford to manufacturers? Because these obligations are defined in relation to systemized driving, the manufacturer can satisfy them with identifiable aggregate performance measures of the operating system. Once again, the systemized driving behavior of autonomous vehicles can solve the tort problem.

---

150. An actor's conduct is "negligent if the magnitude of the risk [foreseeably created by the conduct] outweighs the burden of risk prevention." RESTATEMENT (THIRD) OF TORTS: LIABILITY FOR PHYSICAL AND EMOTIONAL HARM § 3 cmt. e (AM. LAW INST. 2010). In the trolley problem, the foreseeable risk created by the driving behavior involves the threatened fatal injuries to bystanders, whereas the burden of preventing that risk involves exposing the occupants of the vehicle to the fatal risk. Because the interests of bystanders are given the same weight as the interests of the vehicle's occupants, it would be negligent to design the vehicle so that it chooses to injure a larger number of individuals to prevent the same injury for a smaller number.

151. See Surden & Williams, *supra* note 28, at 163–74 (discussing how "the activities of autonomous vehicles . . . can be made more predictable through deliberate technological design decisions").



### 1. Adequate Premarket Testing

In addition to revealing programming errors that cause the operating system to crash, premarket testing will also generate the opportunities for machine learning that improves the safety performance of the operating system. To be adequate, the premarket testing must show that the autonomous vehicle can drive in a reasonably safe manner when it first becomes commercially available.

During the period when the market transitions from conventional vehicles to autonomous vehicles, there is a clear benchmark to conclusively satisfy this tort obligation. Taking into account the risk-utility trade-off between a conventional vehicle and an imperfect but safer autonomous vehicle, the fully functioning autonomous vehicle will necessarily drive in a reasonably safe manner if prior driving experience shows that the operating system at least halves the incidence of crashes relative to conventional vehicles.

To see why this performance standard would necessarily satisfy the manufacturer's tort obligation, consider how Waymo (formerly Google) conducts the premarket testing of its fleet of autonomous vehicles. Waymo teaches the operating system to learn from situations in which the human backup or "test driver" had to take manual control of the autonomous vehicle in order to avoid a crash.<sup>152</sup> Now consider how this method of testing would apply to a set of driving conditions—total miles, proportion spent on expressways, in urban areas, and so on—that makes it possible to reliably compare the safety performance of the operating system with the safety performance of conventional vehicles.

Suppose that the driving conditions in question would result, on average, in ten fatal conventional vehicle crashes, according to data. Suppose that the autonomous vehicle is equally safe, so that under these same conditions it would also cause an average of ten fatal crashes in the absence of a test driver. The circumstances involving the ten fatal crashes of a conventional vehicle (usually due to errors by the human driver) would differ from the ten unanticipated "corner cases" that cause the fatal crash of an autonomous vehicle. Consequently, the performance of the operating system cannot be evaluated by simply asking how a human driver would have responded in the case at hand. That type of inquiry would not account for the crashes caused by human drivers that are avoided by the autonomous vehicle. When the fully functioning operating system was engaged in systemized driving behavior, its

---

152. See GOOGLE, GOOGLE SELF-DRIVING CAR TESTING REPORT ON DISENGAGEMENTS OF AUTONOMOUS MODE 4–5 (2015), <https://www.dmv.ca.gov/portal/wcm/connect/df67186-70dd-4042-bc8c-d7b2a9904665/GoogleDisengagementReport2014-15.pdf?MOD=AJPERES> [<https://perma.cc/JMH2-57GQ>]; see also David Streitfeld, *Waymo to Offer Arizona Access to Self-Driving Cars*, N.Y. TIMES, Apr. 25, 2017, at B3 (describing how Waymo is expanding its premarket testing by allowing "ordinary people" to "integrate[] one [of Waymo's autonomous vehicles] into their daily lives," although each vehicle "will have a technician who can take control in an emergency").

performance in any given case must be evaluated with aggregate driving data that compares its systemic performance to that of conventional vehicles.

Based on the hypothetical data above, further testing of the autonomous vehicle would be warranted. The delayed commercial deployment would not create any safety costs, as the autonomous vehicle would not otherwise be reducing total crashes relative to conventional vehicles. Further testing, however, would create a net safety benefit insofar as the test driver would be able to avoid some of the ten fatal crashes that would otherwise occur if the unmanned autonomous vehicle confronted one of these corner cases. For example, suppose that by taking manual control of the autonomous vehicle, the test driver could avoid a fatal crash in five out of the ten cases. The further testing would create a safety cost of these five fatal crashes; it would also produce an expected safety benefit of up to ten fatal crashes that could be prevented by the operating system “learning” how to solve these corner cases. The safety benefits would exceed the costs, requiring more extensive testing under the risk-utility test even though the autonomous vehicle could otherwise perform as safely as a conventional vehicle.

As this example illustrates, the costs and benefits of more extensive testing depend on various factors, and so the requisite amount of premarket testing is an empirical question. Therefore, by adopting a set of factual assumptions that bias the risk-utility analysis in favor of more extensive testing, we can identify the relative safety performance that would conclusively satisfy the manufacturer’s obligation to subject the operating system to adequate premarket testing.

Consider once again a set of driving conditions that would result, on average, in the ordinary human driver causing a conventional vehicle to fatally crash ten times. An autonomous vehicle that halves this rate would be expected to cause only five fatal crashes under these same conditions.

The case for more extensive testing would be strongest if further testing of the autonomous vehicle under these same conditions would virtually eliminate these crashes: the test driver can always avoid crashing when the vehicle confronts a corner case, and the operating system can then learn how to avoid these crashes moving forward.<sup>153</sup> At most, then, more extensive testing would eliminate these five fatal crashes, and so this risk can be imputed to the current design of the vehicle’s operating system.

This risk must then be compared to the cost or disutility of reducing it by altering the current design via more extensive testing (and machine learning for the operating system). The delayed deployment of the autonomous vehicle determines the expected cost of additional testing. At minimum, the cost equals

---

153. Some crashes will be unavoidable either because the test driver cannot avoid a crash or because the operating system will not be able to solve the corner case. By ignoring both of these possibilities, the assumption maximizes the safety benefit that more extensive testing might attain, thereby presenting the strongest possible case for such testing.

the lost safety benefit of eliminating the five fatal crashes that a conventional vehicle's human driver would otherwise cause across the same operating conditions.<sup>154</sup>

Consequently, the minimum expected cost of more extensive testing (five fatal conventional vehicle crashes that immediate deployment of the autonomous vehicle would prevent) equals the maximum expected safety benefit or reduced risk (five fatal autonomous vehicle crashes that further learning of the operating system would prevent). Expected costs (or disutility) equal expected benefits (or risk reduction), but because the analysis so far has been biased in favor of more extensive testing, the balance at this point does not justify such testing. Therefore, more extensive testing is not required by the risk-utility test and by extension, the modified consumer expectations test.

The 50 percent threshold is derived from biased factual assumptions that present the strongest possible case for more extensive testing under current market conditions, which involve the replacement of conventional vehicles with autonomous vehicles. As the market matures and autonomous vehicles become the norm, a different baseline of comparison will be required. For now, however, the baseline for evaluating safety benefits of an autonomous vehicle can be defensibly defined in terms of conventional vehicles.<sup>155</sup> By testing the autonomous vehicle to the point at which it performs at least twice as safely as conventional vehicles, the manufacturer will conclusively show that the fully functioning operating system is reasonably safe and not defectively designed.

The Waymo self-driving car may have already attained this performance standard with respect to moderate and less severe crash events, whereas for severe crashes, the rate for the self-driving car is now about one-third lower than conventional vehicles.<sup>156</sup> But because these data are based on only 1.3 million miles of driving exposure for the fleet of autonomous vehicles, "there is currently too much uncertainty in self-driving rates to draw this conclusion with strong confidence."<sup>157</sup>

For statistical reasons (the law of large numbers), the extent of experience (or sample size) determines the reliability of the data. All else being equal, a larger sample size entails more reliable estimates. To demonstrate with 95 percent confidence that an autonomous vehicle halves the rate of fatal accidents

---

154. The cost of delayed deployment is not limited to safety concerns but also includes any lost benefits of autonomous driving, such as the time someone saves by not having to drive the vehicle. By ignoring these other lost benefits (or costs), the assumption provides the strongest case for more extensive testing.

155. For further discussion of why the safety performance of an autonomous vehicle should be evaluated in relation to conventional vehicles, see *infra* notes 268–70 and accompanying text.

156. MYRA BLANCO ET AL., VA. TECH. TRANSP. INST., AUTOMATED VEHICLE CRASH RATE COMPARISON USING NATURALISTIC DATA, at iv (2016), [http://www.vtti.vt.edu/PDFs/Automated Vehicle Crash Rate Comparison Using Naturalistic Data\\_Final\\_Report\\_20160107.pdf](http://www.vtti.vt.edu/PDFs/Automated_Vehicle_Crash_Rate_Comparison_Using_Naturalistic_Data_Final_Report_20160107.pdf) [<https://perma.cc/N538-PWS5>].

157. *Id.*

relative to conventional vehicles, the manufacturer would need to test drive the vehicle approximately 500 million miles.<sup>158</sup> Doing so could take over twenty years.<sup>159</sup> As such, this requirement is overly demanding in light of the safety benefits that would be lost during such a prolonged delay in commercial distribution.

The testing problem stems from the infrequency of fatal crashes. Crashes of other types are much more common. If the premarket testing standard were instead defined in relation to the total estimated crashes of conventional vehicles, then the fleet of autonomous vehicles would require less than 2 million miles of driving exposure to demonstrate with 95 percent confidence that it halves the rate relative to conventional vehicles.<sup>160</sup> The Waymo self-driving car project is already approaching this point.

What are the appropriate crash metrics for evaluating the relative safety performance of an autonomous vehicle? What degree of statistical certainty is required? Without clear-cut answers, manufacturers face a significant source of legal uncertainty, even though the systemized driving behavior of autonomous vehicles provides an identifiable performance benchmark—halving the crash rate relative to conventional vehicles—that would conclusively show that the operating system is reasonably safe or not defectively designed. Aggregate performance data will provide an adequately determinate measure for satisfying the manufacturer's tort obligation to test the vehicle only if courts ultimately adopt the same metrics for evaluating the reliability of the data—an issue addressed below.<sup>161</sup>

## 2. *Adequate Warnings About the Inherent Risk of Crash*

Even when subject to adequate premarket testing and properly deployed, a fully functioning operating system can still cause the vehicle to crash. The circumstances in which this might occur will be opaque to consumers, and so the underlying tort obligations require the manufacturer to adequately warn about the inherent, foreseeable risk that the fully functioning autonomous vehicle might crash.<sup>162</sup>

---

158. NIDHI KALRA & SUSAN M. PADDOCK, RAND CORP., DRIVING TO SAFETY: HOW MANY MILES OF DRIVING WOULD IT TAKE TO DEMONSTRATE AUTONOMOUS VEHICLE RELIABILITY? 7 fig.3 (2016), [https://www.rand.org/content/dam/rand/pubs/research\\_reports/RR1400/RR1478/RAND\\_RR1478.pdf](https://www.rand.org/content/dam/rand/pubs/research_reports/RR1400/RR1478/RAND_RR1478.pdf) [<https://perma.cc/RU9C-5NY4>].

159. *Id.* at 10 tbl.1 (providing different time estimates for total miles based on a fleet of 100 autonomous vehicles driving twenty-four hours a day, 365 days a year, at an average speed of twenty-five miles per hour).

160. *Id.* at 7 fig.3.

161. *See infra* Part IV.

162. *See supra* notes 99–108 and accompanying text (explaining why such a warning would defeat a tort claim that the vehicle “malfunctioned” in these cases, and why even if such performance were not a malfunction, such a warning is still required by tort law).

This type of disclosure could be particularly vulnerable to a claim that it is defective for not providing sufficiently detailed information about the circumstances that led to the crash in any given case. The tort system continues to have difficulty resolving claims that product warnings are defective for not adding more detailed information of this type.<sup>163</sup> Uncertainty about the requisite amount of detail could mean that the crash of an autonomous vehicle will routinely result in tort litigation over the adequacy of the product warning.

Litigation of this type rarely occurs for conventional vehicles because the behavior of the human driver largely determines the inherent risk of crash.<sup>164</sup> Having eliminated the human driver, the operating system of the autonomous vehicle largely determines the risk. How can the manufacturer adequately warn consumers about that risk?

Once again, the systemized driving behavior of autonomous vehicles can resolve this problem. The aggregate driving performance of the fleet provides the requisite data for auto insurers to calculate the cost of insuring the vehicle. This insurance premium is based on the inherent, foreseeable risk that the vehicle will crash, and so manufacturers can adequately warn about this risk by disclosing the premium.

“Warnings alert users and consumers to the existence and nature of product risks so that they can prevent harm either by appropriate conduct during use or consumption or by choosing not to use or consume.”<sup>165</sup> For example, suppose there is an unavoidable one-in-one-thousand risk that a safely used product will cause the ordinary consumer to suffer \$10,000 in damages. The inherent risk of injury imposes a cost on the consumer at least equal to ten dollars per product use (the expected value of the injury), so consumers should factor this ten-dollar cost into their estimate of the net benefit that they expect to derive from the product. Consumers who are unaware of the risk will not account for this cost, however, thereby inflating

---

163. Cf. RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 2 cmt. i (“It is impossible to identify anything approaching a perfect level of detail that should be communicated in product disclosures. . . . No easy guideline exists for courts to adopt in assessing the adequacy of product warnings and instructions.”); James A. Henderson, Jr. & Aaron D. Twerski, *Doctrinal Collapse in Products Liability: The Empty Shell of Failure to Warn*, 65 N.Y.U. L. REV. 265, 326 (1990) (“[N]egligence doctrine in the context of failure-to-warn litigation is little more than an empty shell.”); see also GEISTFELD, PRODUCTS LIABILITY, *supra* note 47, at 139–53 (discussing the problem of adding more detailed information to warnings and showing why it stems from the failure of current jury instructions to adequately account for the information costs that consumers must incur in order to read, remember, and follow product warnings).

164. A notable exception involves the risk of roll over that is inherent in a sport utility vehicle. “Buying an SUV involves a tradeoff. While these vehicles may do well in certain types of crashes, they also are much more likely to roll over. People should be aware of that trait when they are choosing a family vehicle.” DOT Requires Upgraded Warning Label for Sport Utility Vehicles, NHTSA 8-98, U.S. Dept. of Transp., 1999 WL 118266 (Mar. 5, 1999) (quoting Ricardo Martinez, M.D., Administrator of NHTSA).

165. RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 2 cmt. i (AM. LAW INST. 1998).

their estimate of the product's net benefit and causing them to purchase or use more of the product than they would choose if well informed. The excessive purchase or use of the product then creates excessive risk, resulting in too many injuries. To address this safety problem, tort law imposes a duty on the manufacturer to warn about any foreseeable risks of physical harm that are unknown by the ordinary consumer and would be material to his or her decision making. By satisfying the duty to warn, the manufacturer enables consumers to make informed safety decisions.

The manufacturer's general duty to warn is embodied in a set of more specific rules and standards regarding the various characteristics of an adequate warning; all serve the purpose of reducing the information costs that the ordinary consumer must incur to read, remember, and follow the warning.<sup>166</sup> For example, an adequate warning must prominently disclose the most serious risks rather than burying that information in fine print.<sup>167</sup> The ordinary consumer will not expend the time and effort to dig the information out of the fine print, rendering the warning inadequate. A more prominent warning makes this highly material information readily accessible—it reduces information costs for the consumer, satisfying the manufacturer's tort obligation.<sup>168</sup>

For these same reasons, a manufacturer can satisfy its obligation to warn about the inherent risk of crash through disclosure of the premium for insuring the autonomous vehicle. Suppose a consumer is deciding whether to purchase either Brand-A or Brand-B of an autonomous vehicle, each of which is otherwise identical except for their respective operating systems. Suppose Brand-A costs \$30,000 and has an annual, risk-adjusted insurance premium of \$2,500, whereas Brand-B costs \$31,000 and has a risk-adjusted premium of \$1,000. The consumer can readily determine that the safety decision favors Brand-B because it has a lower total cost (\$32,000 in the first year alone) than Brand-A (\$32,500). The simple price comparison enables the consumer to make good decisions about the relative risks inherent in the reasonably safe designs of different autonomous vehicles, producing a market dynamic that incentivizes manufacturers to reduce the inherent risk of crash and the corresponding cost of insurance. By enabling the ordinary consumer to make an informed safety decision about the matter, this type of disclosure satisfies the manufacturer's tort obligation to warn about the inherent, foreseeable risk that the autonomous vehicle will crash.

---

166. See GEISTFELD, PRODUCTS LIABILITY, *supra* note 47, at 142–49.

167. See, e.g., *Jones v. Amazing Prods., Inc.*, 231 F. Supp. 2d 1228, 1248 (N.D. Ga. 2002) (holding that a jury could find the defendant's warning to be defective because the "cautionary instruction" at issue was "buried in the middle of a long paragraph, in a very small print size").

168. According to a leading formulation, an adequate warning "must (1) be designed so it can reasonably be expected to catch the attention of the consumer; (2) be comprehensible and give a fair indication of the specific risks involved with the product; and (3) be of an intensity justified by the magnitude of the risk." *Pavlidis v. Galveston Yacht Basin, Inc.*, 727 F.2d 330, 338 (5th Cir. 1984).

Mandated disclosures of this type (“two-price schemes” involving the base price of a product plus a separate price for the insurance costs) have been proposed in other contexts.<sup>169</sup> The disclosures are not limited to “two prices,” however. For example, the disclosure should be further refined to separate the premium for insuring the autonomous vehicle itself from the premium for insuring against personal injury. One premium would apprise consumers about the inherent risk of property damage, and the other would inform them about the inherent risk of bodily injury, thereby enabling consumers to estimate any expected injury costs (like pain and suffering) that the insurance premium does not cover.<sup>170</sup> As empirical studies have found, disclosures of this type help uninformed consumers make better decisions.<sup>171</sup>

Compare this disclosure to a general warning that when confronted by unanticipated conditions, the fully functioning operating system can cause the autonomous vehicle to crash. Having digested this general warning, the consumer would *not* rely on it to alter the driving behavior of the vehicle—the operating system would still be in full control across the conditions in question. Instead, the consumer would need to figure out how the warning translates into the inherent risk that the vehicle will cause a crash resulting in injury. The consumer would then presumably purchase insurance to cover many of these losses. At this point, the consumer could finally determine his or her total costs for the vehicle (purchase price plus insurance premium plus uninsured losses).

The consumer could make this identical decision in a much simpler and more accurate manner by relying on the manufacturer’s disclosure of the annual, risk-adjusted premium for insuring the vehicle. By substantially

---

169. Mark Geistfeld, *The Political Economy of Neocontractual Proposals for Products Liability Reform*, 72 TEX. L. REV. 803, 821–34 (1994) [hereinafter *Neocontractual Proposals*]; Mark Geistfeld, Note, *Imperfect Information, the Pricing Mechanism, and Products Liability*, 88 COLUM. L. REV. 1057, 1063–72 (1988); Alan Schwartz, *Proposals for Products Liability Reform: A Theoretical Synthesis*, 97 YALE L.J. 353, 407 (1988) (proposing for further study a pricing scheme whereby firms must “quote two prices to consumers,” one including the cost of manufacturer liability and the other excluding this cost); see also 2 AM. LAW INST., REPORTERS’ STUDY: ENTERPRISE RESPONSIBILITY FOR PERSONAL INJURY 522 (1991) (discussing two-price schemes and concluding that they merit further study because by shopping and comparing relative differences between risk-adjusted insurance premiums, “uninformed consumers would learn much about the risks they face”).

170. Federal data regarding motor vehicle crashes provide estimates of both the economic and noneconomic costs of crashes. In 2013, for example, the noneconomic costs of crash were about 2.5 times greater than the economic costs of crash. See *supra* note 2 and accompanying text. Based on these data, the ordinary consumer in 2013 could estimate that her expected noneconomic injury costs would be 2.5 times greater than the insurance premium covering economic losses.

171. See Susan K. Laury & Melayne Morgan McInnes, *The Impact of Insurance Prices on Decision Making Biases: An Experimental Analysis*, 70 J. RISK & INS. 219, 221–30 (2003) (reporting results of experimental study which found that the disclosure of insurance costs improved the decision making of uninformed consumers); see also Mark Andor et al., *Consumer Inattention, Heuristic Thinking and the Role of Energy Labels* (U.S. Ass’n for Energy Econ. Research Paper Series, Working Paper No. 16-287, 2016), [http://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2795579](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2795579) [<https://perma.cc/DG73-EKQ3>] (summarizing results of an empirical study finding that mandated energy labels on electrical appliances increase consumer attention to operating cost and reduce consumer reliance on other salient methods of relative valuation).

reducing information costs relative to a general warning that ideally should otherwise lead the ordinary consumer to make the same safety decision, the manufacturer's disclosure would satisfy its tort obligation to warn about the inherent risk of crash.

An instructive analogy is provided by warnings about the inherent risk that a prescription drug will cause an injurious side effect. A general warning that the drug could cause a side effect would not be adequate if the manufacturer has more specific information about the likelihood and consequences of the side effect.<sup>172</sup> So, too, a general warning that the autonomous vehicle might cause a crash would not be adequate; the manufacturer could provide more detailed information about the inherent risk—the separate insurance premiums for bodily injury and property damage—that would enable the ordinary consumer to make a more informed safety decision.

To function in this manner, the warning must disclose an insurance premium that is adequately adjusted to account for the risk that the fully functioning operating system of a properly deployed autonomous vehicle will cause a crash. The information that insurers require to calculate such a premium should be feasibly attainable.

In order to provide the basis for a risk-adjusted premium, the vehicle's prior crash experience must be a reliable indicator of the risk now being insured.<sup>173</sup> Due to the manner in which machine learning improves the safety performance of an autonomous vehicle, its current capabilities will exceed or otherwise be no worse than its prior capabilities. A vehicle's prior crash

---

172. See, e.g., *Martin v. Hacker*, 628 N.E.2d 1308, 1312 (N.Y. 1993) (“A warning for a prescription drug may be held adequate as a matter of law if it provides specific detailed information on the risks of the drug. . . . Always bearing in mind that the warning is to be read and understood by physicians, not laypersons, the factors to be considered in resolving this question include whether the warning is accurate, clear, consistent on its face, and whether it portrays with sufficient intensity the risk involved in taking the drug.”) (citations omitted); see also *supra* note 168 (describing characteristics of an adequate warning for ordinary products).

173. Automobile insurers rely on two different methods for tailoring the premium to the risk characteristics of the individual policyholder, each of which is identical in the context of autonomous vehicles. First, the premium for automobile insurance can reflect the expected performance of certain safety features of the vehicle based on the prior loss experience of vehicles equipped with those features. For example, insurers discount premiums for vehicles containing devices such as anti-collision systems, anti-lock brakes, anti-theft systems, daytime running lights, and passive restraints. See *Duffy & Hopkins*, *supra* note 25, at 478. These safety features predictably reduce the cost of automobile accidents and provide an actuarial basis for reducing the premium. A different method for establishing risk-adjusted premiums relies on the policyholder's prior loss experience. See NAT'L COUNCIL ON COMPENSATION INS., ABCS OF EXPERIENCE RATING 2 (2015) (“Experience rating takes the average loss experience and modifies it based on the individual's own loss experience.”), [https://www.ncci.com/Articles/Documents/UW\\_ABC\\_Exp\\_Rating.pdf](https://www.ncci.com/Articles/Documents/UW_ABC_Exp_Rating.pdf) [<https://perma.cc/BCE3-JZ7A>]. Because the “driver” of an autonomous vehicle is the software and associated hardware devices of the operating system—all of which are safety features of the vehicle—the experience rating of these vehicles is no different from feature rating.



experience, therefore, provides a reliably conservative measure of the risk now being insured.

If the insurer has a sufficiently large sample size, it can establish risk-adjusted premiums for the vehicle.<sup>174</sup> For reasons previously discussed, manufacturers must subject autonomous vehicles to adequate premarket testing.<sup>175</sup> The crash experience of an entire fleet of vehicles during this period will presumably generate the requisite amount of data.<sup>176</sup> The actuarial problem will then be simplified over time as the operating system gains more driving experience and generates more performance data. Autonomous vehicles will provide a trove of big data that insurers can use to establish risk-adjusted premiums for each type of vehicle.<sup>177</sup>

Indeed, many automobile insurers have already adopted “usage-based” plans that rely on devices installed in the vehicle to monitor the policyholder’s driving behavior.<sup>178</sup> Based on the information collected from these devices, insurers determine “car insurance prices not only on proxy-based traditional models [such as age], but also on real driving habits and driving behaviour of policyholders (for example, distance driven, speeding, harsh braking, etc.).”<sup>179</sup> For largely the same reasons that insurers can now tailor premiums to more closely match the risk characteristics of individual drivers, they will also be able to establish risk-adjusted premiums for insuring different types of autonomous vehicles by relying on the prior crash experience of their respective operating systems.<sup>180</sup>

---

174. For example, workers’ compensation insurance must be experience rated in some states when the employer has a sufficient volume of claims experience. See NAT’L COUNCIL ON COMPENSATION INS., *supra* note 173, at 4–10 (also providing a detailed illustration of how past claims experience affects an employer’s premium for workers’ compensation). In contrast, the premiums for medical malpractice liability insurance are not ordinarily experience rated because the prior claims experience of an individual physician does not reliably predict her current risk of being subject to malpractice liability. See Mark Geistfeld, *Malpractice Insurance and the (II) Legitimate Interests of the Medical Profession in Tort Reform*, 54 DEPAUL L. REV. 439, 444 (2005).

175. See *supra* Part II.A.

176. Cf. *supra* note 174 (describing the data that insurers now rely on to establish risk-adjusted premiums for workers’ compensation insurance).

177. Cf. Rick Swedloff, *Risk Classification’s Big Data (R)evolution*, 21 CONN. INS. L.J. 339, 342–43 (2014) (explaining why the more extensive collection of data allows automobile insurers to “price auto insurance to better reflect the risks posed by the drivers”).

178. Clint Boulton, *Auto Insurers Bank on Big Data to Drive New Business*, WALL ST. J. (Feb. 20, 2013), <http://blogs.wsj.com/cio/2013/02/20/auto-insurers-bank-on-big-data-to-drive-new-business> [<https://perma.cc/V4ZK-TUGS>].

179. Lea Pogarcic Mataija & Caroline Van Schoubroeck, *Telematics Insurance: Legal Concerns and Challenges in the EU Insurance Market 4* (KU Leuven’s Research Council, Project COMPACT C24/15/001, 2016), [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2889408](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2889408) [<https://perma.cc/D8SG-MYS3>].

180. As a leading insurer recently concluded in a report on autonomous vehicles, an “aspect of increased computerisation, especially in the case of cars, is that insurers can take advantage of data facilities already present in the vehicle to use a more telematics based approach to premium pricing. This could allow better matching of exposure to premiums, and more individually tailored policies.” GILLIAN YEOMANS, LLOYD’S, AUTONOMOUS VEHICLES: HANDING OVER CONTROL:

By disclosing the annual, risk-adjusted premium for insuring the autonomous vehicle, the manufacturer would adequately warn consumers about the inherent, foreseeable risk of crash. The premium is based on the collective crash experience of the fleet, which in turn adequately reflects the inherent risk that each vehicle within the fleet will cause a crash. Once again, the systemized driving behavior of autonomous vehicles provides a determinate performance measure that satisfies the manufacturer's tort obligation, eliminating this form of liability as a plausible source of significant legal uncertainty.

\* \* \*

Autonomous vehicles can crash for various reasons, ranging from hardware problems to “corner cases” that the operating system had not previously encountered and addressed. Comprehensive analysis of these different crash types has shown that established tort doctrines provide relatively clear answers to the liability questions, with one notable exception. Although the reasonably safe performance of a fully functioning autonomous vehicle would be conclusively established if the vehicle halves the incidence of crashes relative to conventional vehicles, important questions about the requisite testing conditions cannot be conclusively resolved, creating a significant source of legal uncertainty. To complete our assessment of the most significant potential liabilities and associated uncertainties faced by manufacturers, we must consider one remaining type of crash.

### III.

#### MANUFACTURER LIABILITY FOR THE CRASH OF A HACKED VEHICLE

Like other products that make up the “internet of things,” motor vehicles equipped with automated driving technologies might be accessed or hacked by unauthorized third parties. “In recent years researchers have demonstrated hair-raising hacks that make it possible to take over the brakes, engine, or other components of a person's car remotely—forcing the auto industry to take security more seriously.”<sup>181</sup> Unless an autonomous vehicle is secure from cyberattack, a third party could gain unauthorized control by hacking into the operating system. The hacker could then subject the owner to a “ransom” demand to make the vehicle fully operational once again. “It is also feasible that driving could be maliciously interfered with, causing a physical danger to passengers. There is potential for cyber terrorism too—for example, a large-scale immobilisation of cars on public roads could throw a country into

---

OPPORTUNITIES AND RISKS FOR INSURANCE 21 (2014), [https://www.lloyds.com/~media/lloyds/reports/emerging\\_risk\\_reports/autonomous\\_vehicles\\_final.pdf](https://www.lloyds.com/~media/lloyds/reports/emerging_risk_reports/autonomous_vehicles_final.pdf) [<https://perma.cc/7SV2-PVZT>].

181. Tom Simonite, *Your Future Self-Driving Car Will Be Way More Hackable*, MIT TECH. REV. (Jan. 26, 2016), <https://www.technologyreview.com/s/546086/your-future-self-driving-car-will-be-way-more-hackable> [<https://perma.cc/9FEW-8JV6>].

chaos.”<sup>182</sup> In the event that hacking or some other type of cyberattack causes a vehicle to crash, attention will inevitably turn to the question of whether the manufacturer is liable.<sup>183</sup>

For well-established reasons, the manufacturer’s tort obligations encompass the cybersecurity of the vehicle. Hacking, unfortunately, is common these days, creating a risk of illegal third-party conduct that obligates the manufacturer to protect consumers from the foreseeable harms.<sup>184</sup> The duty in this regard is no different from the one businesses already face as they attempt to prevent hackers from gaining unauthorized access to confidential consumer data, such as social security numbers and credit card information.<sup>185</sup> Another relevant analogy comes from the duty of landlords to protect their tenants from foreseeable criminal attacks like burglaries and so on.<sup>186</sup> The difficult question in this context is not whether a manufacturer has a duty to protect its motor vehicles against illegal cyberattacks like hacking but rather the substantive content of that duty. What are the manufacturer’s obligations with respect to crashes caused by cyberattacks?

As we have found, the manufacturer can satisfy its tort obligations for the crash of a fully functioning autonomous vehicle through adequate premarket

---

182. LLOYD’S, AUTONOMOUS VEHICLES, *supra* note 180, at 16.

183. Cases in which the hacked vehicle is rendered immobile and subject to a “ransom” demand implicate the economic loss rule, as the only damage is to the product itself. For discussion of the limited circumstances in which consumers can recover for the pure economic losses caused by defective products, see generally Mark A. Geistfeld, *The Contractually Based Economic Loss Rule in Tort Law: Endangered Consumers and the Error of East River Steamship*, 65 DEPAUL L. REV. 393 (2016).

184. See *Identity Theft Resource Center Breach Report Hits Near Record High in 2015*, IDENTITY THEFT RES. CTR. (Jan. 25, 2016), <http://www.idtheftcenter.org/ITRC-Surveys-Studies/2015databreaches.html> [<https://perma.cc/ZP3W-U297>] (“The number of U.S. data breaches tracked in 2015 totaled 781.”).

185. See, e.g., *In re Sony Gaming Networks & Customer Data Sec. Breach Litig.*, 996 F. Supp. 2d 942, 966 (S.D. Cal. 2014) (holding that “the existence of [the general tort duty governing data breaches is] well supported by both common sense and . . . Massachusetts law”). Like numerous others, the court in this case then limited the general tort duty with the economic loss rule, *id.* at 967, a limitation that is not applicable to a manufacturer’s general tort duty to protect consumers against foreseeable risks of physical harm, such as the bodily injury and property damage caused by a hacked autonomous vehicle.

186. Consider the reasons why a landlord owes a tort duty to tenants to protect them from the foreseeable risk of third-party criminal acts:

[T]he landlord has control over common areas, has superior means for providing security, and derives commercial advantage from the relationship. The landlord also has an ongoing contractual relationship with the tenant, and the lease itself could allocate responsibility for exercising care. Because the landlord usually is in a better position than individual tenants to exercise control over common areas and, with respect to individual units, to provide locks and other security devices, imposing a duty on the landlord replicates the result that might be reached if landlords and tenants with similar bargaining power addressed this matter.

RESTATEMENT (THIRD) OF TORTS: LIABILITY FOR PHYSICAL AND EMOTIONAL HARM § 40 cmt. m (AM. LAW INST. 2010). This reasoning fully applies to the provision of cybersecurity by the manufacturers of autonomous vehicles.

testing and an adequate warning about the inherent, foreseeable risk of crash.<sup>187</sup> Each of these safety measures depends on the systemized driving behavior of autonomous vehicles—and the resultant aggregate statistics based on prior driving performance—to measure current safety performance. Unlike the prior crash experience of an autonomous vehicle’s fully functioning operating system, the incidence of prior hacks is not a sufficiently reliable predictor of future attacks. Adequate premarket testing and a warning about the inherent risk of crash would not insulate the manufacturer from tort liability for the crash of a hacked vehicle.

During the premarket testing phase, an autonomous vehicle is in limited use and presumably less appealing for hackers. Once the vehicle has been commercially distributed, the prospects for illicit gain considerably increase, making the vehicle a substantially more attractive target for hackers. Even if a vehicle was not hacked during the premarket testing phase, it could still be vulnerable once it has been widely distributed in the market.

After the vehicle has been commercially distributed, its hacking history still does not reliably translate into the current risk of cyberattack. As was true in the premarket testing phase, the prior history could merely reflect the relative inattention of powerful hackers. Moreover, the varied “computers, sensors, and other components” required for autonomous driving “will expand the possible entry points for attackers and the things they can do—for example, self-driving cars rely on laser scanners and other sensors, which could be made to send false data.”<sup>188</sup> Like the arcade game Whac-A-Mole, each time the manufacturer patches the operating system to protect against vulnerabilities hackers have previously exploited, the range of other attack points could enable hackers to pop up somewhere else by exploiting a different vulnerability. The ability of hackers to exploit vulnerabilities in the past does not necessarily predict future attacks that exploit different vulnerabilities.

Because the vehicle’s prior hacking history is not a reliable measure of the current threat, safety measures based on prior performance would not satisfy the manufacturer’s tort obligations with respect to cybersecurity. We return, then, to the question of whether the manufacturer would incur liability for the crash of a hacked vehicle.

An operating system that has not been reasonably designed to withstand cyberattacks would be defective under the risk-utility test, subjecting the manufacturer to liability for the resultant crashes.<sup>189</sup> Absent proof of a defect, the crash of the hacked vehicle would not result in liability.

Alternatively, the defect could be proven under the malfunction doctrine. The manufacturer obviously did not intend for the vehicle to be controlled by

---

187. See *supra* Part II.

188. Simonite, *supra* note 181 (reporting on presentation made by Stefan Savage at a cybersecurity conference).

189. See *supra* notes 111–16 and accompanying text (discussing proof of defective design).

an unauthorized third party, and the ordinary consumer presumably expects that the operating system, not a hacker, determines the vehicle's driving behavior. The performance of a hacked vehicle would seem to be a product malfunction that would subject the manufacturer to strict liability in the event of a crash.<sup>190</sup> This form of liability could be particularly potent for plaintiffs and worrisome for manufacturers.

This liability question, however, is not so easily resolved. Courts and commentators have not adequately defined the necessary attributes of a product malfunction.<sup>191</sup> The approach instead seems to be, "we know it when we see it." To address the most vexing liability issue involving cybersecurity, we need to more fully develop the malfunction doctrine.

#### A. *The Crash of a Hacked Vehicle as a Product Malfunction*

Complex, interrelated systems govern the performance of a motor vehicle with automated driving technologies, and rigorous quality control is necessary to identify and solve problems that can cause the vehicle to crash. A failure of quality control is also often the root cause of a product malfunction, such as the manufacturer's failure to detect or protect against contaminants that cause a food product to "malfunction" by being unfit for human consumption. Due to the apparent similarities between the crash of an automated vehicle and other types of product malfunctions, these crashes will undoubtedly place pressure on courts to more clearly articulate the attributes of product performance that constitute a malfunction subject to strict liability.

One of the paradigmatic examples of product malfunction—the exploding bottle of soda—illustrates why the doctrine requires more rigorous specification. The exploding bottle spawned the modern rule of strict products liability along with the contaminated food cases.<sup>192</sup> What does the ordinary consumer reasonably expect in these cases? Like the manufacturer, the consumer knows that systems of perfect quality control are either prohibitively expensive or simply unattainable. Some soda bottles will inevitably have undetected problems that cause them to explode (just as food will sometimes be contaminated). Since the consumer knows and therefore reasonably expects that perfect quality control is not ordinarily attainable, it is unclear why an exploding soda bottle (or contaminated food) is a malfunction that frustrates his or her minimum expectations of safe product performance.

---

190. See *supra* Part II.B.1 (explaining why the malfunction doctrine is defined in terms of the product performance either intended by the manufacturer or expected by the consumer).

191. See *id.* (discussing ambiguous nature of the malfunction doctrine).

192. An exploding soda bottle was at issue in *Escola v. Coca Cola Bottling Co.*, 150 P.2d 436, 437 (Cal. 1944). The concurrence by Justice Traynor ultimately persuaded the California Supreme Court and others to adopt strict products liability. See Mark Geistfeld, *Escola v. Coca Cola Bottling Co.: Strict Products Liability Unbound*, in *TORTS STORIES* 229 (Robert L. Rabin & Stephen D. Sugarman eds., 2003).

The puzzle can be reframed by reference to the ordinary consumer's reasonable expectations of product design. Instead of expecting *perfectly* safe designs, the consumer only expects a design to be *reasonably* safe.<sup>193</sup> Such a design still creates an inherent risk of injury; reasonable safety is not absolute safety. The same expectation applies to systems of quality control. Instead of expecting perfection, the ordinary consumer only expects that the soda bottle has passed reasonably safe, though imperfect systems of quality control. For an exploding soda bottle to frustrate consumer expectations, the misperformance must be attributable to the manufacturer's failure to exercise reasonable care in quality control. Consumer expectations accordingly justify a tort rule no different from ordinary negligence liability, just like the equivalent expectation of reasonably safe design justifies the negligence-based risk-utility test. What, then, justifies the rule of strict liability for the exploding soda bottle?

The rationale for strict liability is based on the difficulty of enforcing the manufacturer's obligation to adopt reasonably safe systems of quality control.<sup>194</sup> What is the full range of reasonably safe measures that a manufacturer could adopt to ensure the quality of a mass-manufactured product like bottled soda? The various measures are either complex (the incorporation of quality-control systems into the manufacturing process) or cannot be independently evaluated with reliable evidence (as with visual inspection by employees). The expectation of reasonable quality control, therefore, generates important safety obligations that the consumer cannot adequately enforce. A manufacturer that does not take such a required precaution will be able to avoid negligence liability, reducing its financial incentive to incur this costly safety investment in quality control.<sup>195</sup> Due to the difficulty of enforcement, the negligence rule does not adequately protect the consumer's expectation that the manufacturer will employ reasonable quality-control systems.

A rule of strict liability solves this evidentiary problem and thereby enforces the expectation of reasonably safe quality control. As Oliver Wendell Holmes explained, "the safest way to secure care is to throw the risk upon the

---

193. See *supra* notes 111–16 and accompanying text (explaining why the ordinary consumer reasonably expects that a product design conforms to the risk-utility test).

194. Historically, courts have also invoked a loss-spreading rationale for strict liability, but this is not sufficient because it would justify absolute liability for all product-caused injuries, not merely those caused by defective products. See, e.g., *Cafazzo v. Cent. Med. Health Servs., Inc.*, 668 A.2d 521, 527 (Pa. 1995) (rejecting strict liability based solely on the defendant's "ability to pay plaintiffs and ability to charge others" because such a rationale "would result in absolute rather than strict liability").

195. As one court observed:

It is not doubted that due care might require the defendant to adopt some device that would afford [reasonable protection against the injury suffered by plaintiff.] Such a device, if it exists, is not disclosed by the record. The burden was upon the plaintiff to show its practicability. Since the burden was not sustained, a verdict should have been directed for the defendant.

*Cooley v. Pub. Serv. Co.*, 10 A.2d 673, 677 (N.H. 1940).

person who decides what precautions shall be taken.”<sup>196</sup> Rather than have the court make the safety decision based on the available evidence in the case at hand, strict liability “throws” that decision on the manufacturer. To minimize the sum of its safety expenditures and expected costs of strict liability, the manufacturer will take any safety precaution costing less than the associated reduction in expected liability (injury) costs—the same type of safety decision required by the risk-utility test.<sup>197</sup> Under these conditions, strict liability restores the manufacturer’s financial incentive to exercise reasonable care by eliminating the evidentiary barriers to recovery that inhere in the negligence standard. Recognizing as much, the ordinary consumer can reasonably expect compensation for the exploding soda bottle because that form of (strict) liability is necessary for adequately enforcing the manufacturer’s underlying obligation to adopt reasonably safe systems of quality control. The ordinary consumer can reasonably expect the manufacturer to guarantee that the soda bottle will not explode and is otherwise fit for its intended purpose.

Modern courts invoked this reasoning to justify the ancient rule of strict liability for the sale of contaminated food.<sup>198</sup> These cases subsequently influenced others, like those involving exploding bottles of soda.<sup>199</sup> This case law was then restated into the rule of strict products liability.<sup>200</sup> As the *Restatement (Third) of Torts* explains, strict liability applies to construction or manufacturing defects and serves an “instrumental function of creating safety incentives” greater than those in a negligence regime “under which, as a practical matter, sellers may escape their appropriate share of liability.”<sup>201</sup>

---

196. OLIVER WENDELL HOLMES, JR., *THE COMMON LAW* 117 (1881).

197. See *supra* notes 111–16 and accompanying text (describing the safety decision embodied in the risk-utility test); GEISTFELD, *PRODUCTS LIABILITY*, *supra* note 47, at 61–62 (providing more rigorous demonstration of this conclusion).

198. In tort cases involving the sale of contaminated food, as the Texas Supreme Court observed, “a rule which would require proof of negligence as a basis of recovery would, in most instances, by reason of the difficulty of making such proof, be equivalent to a denial of recovery.” *Jacob E. Decker & Sons, Inc. v. Capps*, 164 S.W.2d 828, 834 (Tex. 1942). After discussing the difficulties faced by a plaintiff in trying to prove that a defendant failed to exercise reasonable care in distributing contaminated food, the court concluded that these evidentiary difficulties justified a rule of strict liability: “Such a rule would seem to be more desirable because it permits the placing of the ultimate loss upon the manufacturer, who is in the best position to prevent the production and sale of unwholesome food. It stimulates and induces a greater degree of precaution for the protection of human health and life than does the rule of ordinary care.” *Id.*

199. As Justice Traynor observed in his influential concurrence arguing for strict products liability, a negligence regime does not adequately solve the safety problem because “[a]n injured person . . . is not ordinarily in a position to refute [the manufacturer’s evidence of reasonable care] or identify the cause of the defect, for he can hardly be family with the manufacturing process as the manufacturer himself is.” *Escola v. Coca Cola Bottling Co.*, 150 P.2d 436, 441 (Cal. 1944) (Traynor, J., concurring).

200. See *RESTATEMENT (SECOND) OF TORTS* § 402A cmt. b (AM. LAW INST. 1965) (discussing how the rule of strict products liability evolved from the sale of contaminated or “corrupt” food and drink).

201. *RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY* § 2 cmt. a (AM. LAW INST. 1998); see also GEISTFELD, *PRODUCTS LIABILITY*, *supra* note 47, at 82–84 (explaining why the

Taken to its extreme, this reasoning would justify a rule of strict liability that is not limited by the requirement of defect, commonly called absolute liability.<sup>202</sup> Like issues pertaining to systems of quality control, proof that a design is defective or unreasonably dangerous is often complicated and difficult, so why not solve that evidentiary problem by eliminating the requirement of defect? Absolute liability would obligate the manufacturer to pay for all injuries foreseeably caused by the product, giving it a financial incentive to make cost-effective investments for reducing these product risks, including those inherent in the product design. Why doesn't the ordinary consumer reasonably expect to receive tort compensation for all product-caused injuries?

By channeling each and every product-caused injury into the tort system, absolute liability would be excessively costly for consumers. The cost of injury compensation through the tort system is considerably higher for consumers than the cost of indemnification through other types of mechanisms, like health insurance.<sup>203</sup> Either way, the consumer incurs these costs (either by paying increased product prices to cover the manufacturer's tort liabilities, or by paying an insurance premium covering those same injuries). Absolute liability would increase total insurance costs for consumers by an amount that would significantly exceed any safety benefit that the rule would otherwise provide across the full range of product cases.<sup>204</sup> The high cost of tort recovery explains why consumers do not reasonably expect manufacturers to provide compensation for all product-caused injuries. As one court put it, "[n]o one wants absolute liability where all the article has to do is to cause injury."<sup>205</sup>

---

various rationales for strict liability that are invoked by the *Restatement (Third)* are all defensibly reduced to the protection of the ordinary consumer's reasonable expectations of product safety).

202. See *Absolute Liability*, BLACK'S LAW DICTIONARY (10th ed. 2014) ("Absolute liability is often distinguished from strict products liability, which limits strict liability to injuries caused by a product defect.").

203. Even if consumers were guaranteed injury compensation for all product-caused injuries, this "tort" insurance would not obviate the need for them to purchase other forms of insurance. Not all health problems are attributable to product injuries, for example, and so consumers would still have to purchase health insurance. Tort insurance, therefore, can be duplicative or otherwise cause problems of coordination with these other forms of "ordinary" insurance, thereby increasing total costs for consumers. Moreover, to obtain the tort insurance, the consumer must incur considerable legal expenses. In contrast, the coverage supplied by ordinary insurance is usually triggered by the fact of loss (like medical expenses for health insurance), which is easy to prove (submitting bills) and does not ordinarily require legal representation. The limited scope of coverage that tort insurance supplies, coupled with its costs of legal representation, largely explain why even in a tort regime of absolute liability, the cost per dollar of coverage supplied by tort insurance would substantially exceed the cost of ordinary insurance for the ordinary consumer. See GEISTFELD, PRODUCTS LIABILITY, *supra* note 47, at 64–67 (providing data and more extensively discussing reasons for the cost differential).

204. See Mark Geistfeld, *Should Enterprise Liability Replace the Rule of Strict Liability for Abnormally Dangerous Activities?*, 45 UCLA L. REV. 611, 639–46 (1998) (relying on a heuristic empirical assessment to show that a rule of strict liability for all injuries proximately caused by a business enterprise is unlikely to reduce risk by an amount that would minimize costs for consumers).

205. *Phillips v. Kimwood Mach. Co.*, 525 P.2d 1033, 1036 (Or. 1974) (en banc).



To minimize total costs for consumers, strict liability must be limited to cases in which it offers the greatest safety potential. The mere fact that a product has caused injury does not signal a significant safety problem; instead, strict liability only applies to injuries caused by a *defect* in the product.<sup>206</sup> Something must be demonstrably wrong with the product to trigger the rule of strict liability. With this framing of the tort problem it becomes apparent why an exploding soda bottle creates an inference of defect—a malfunction—that subjects the manufacturer to strict liability.

If the manufacturer knew that a particular bottle of soda would explode when used by a consumer, it would not sell the bottle in the first instance. The explosion of a bottle, therefore, provides circumstantial evidence that the product is defective. This evidence implicitly relies on the fact that knowledge of the performance would induce the manufacturer to pull this product from the market, or equivalently, that it would be forced to do so because consumers would not buy the product if similarly informed. Indeed, this definition of defect is entailed by the implied warranty rationale for strict products liability, which requires that products “be marketable with their true character known.”<sup>207</sup>

Given this inference of defect, the negligence standard of reasonable care would be an undesirable method for resolving the liability question. To be sure, proof of negligence is easy in some cases. The question, however, is whether negligence is hard to prove across the entire category of cases. As previously discussed, the cost and complexity of the negligence inquiry into systems of quality control would often enable the manufacturer to avoid liability as a practical matter. Consequently, the undeniable problem with this aspect of the product’s performance—established by the manufacturer’s or consumer’s presumptive response if they had known about the malfunction—is best addressed by subjecting the manufacturer to strict liability.<sup>208</sup>

A product does not necessarily malfunction simply because it caused injury, so this doctrine is not a rule of absolute liability. For example, after a conventional automobile has been involved in an ordinary crash, the manufacturer would not usually pull this make of the vehicle from the market, nor would the informed consumer forego purchase of the vehicle. The crash of

---

206. See GEISTFELD, PRODUCTS LIABILITY, *supra* note 47, at 61–67.

207. Prosser, *supra* note 102, at 128–29; see also RESTATEMENT (SECOND) OF TORTS § 402A cmt. m (AM. LAW INST. 1965) (“There is nothing in this Section which would prevent any court from treating the rule stated as a matter of ‘warranty’ to the user or consumer. . . . [To avoid confusion, it] is much simpler to regard the liability here stated as merely one of strict liability in tort.”).

208. Cf. RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 3 cmt. b (AM. LAW INST. 1998) (limiting the malfunction doctrine “to situations in which a product fails to perform its manifestly intended function, thus supporting the conclusion that a defect of some kind is the most probable explanation”). Because the inference of defect is supplied by the manufacturer’s *presumed* response to the performance in question, it does not implicate the rules that prohibit plaintiffs from introducing evidence about a manufacturer’s subsequent remedial measures or *actual* response to a safety problem. See, e.g., FED. R. EVID. 407.

the vehicle, therefore, was not a malfunction, because the presumptive response by neither the manufacturer nor the ordinary consumer provides a circumstantial inference of defect.

Under this formulation, a motor vehicle's failure to properly function due to hacking would be a malfunction subject to strict liability; the reasons are no different from those that apply to malfunctions caused by programming errors or bugs.<sup>209</sup> In hindsight, the manufacturer would presumably redesign the operating system to address the hacking vulnerability or eliminate the bug. The performance of a vehicle that crashes for either reason accordingly creates an inference of defect—a malfunction—that provides a defensible basis for (strict) liability that obviates the need for a complex negligence analysis of the vehicle's hardware and software systems.

The vulnerability of an autonomous vehicle to third-party hacking or other programming errors depends on a wide array of hardware components—the engine, steering wheel, brakes, sensors, and so on—and their interrelated software systems.<sup>210</sup> Given this complexity, “[t]he reality is that driverless cars will certainly suffer from software failure. The open question, however, is how much failure is acceptable.”<sup>211</sup> Does a reasonably reliable operating system function 99 percent of the time? 99.99 percent? The operating system's design can also include modular subsystems with redundancies that enable the vehicle to operate safely when a component malfunctions.<sup>212</sup> Which particular systems are required to make the vehicle reasonable safe? Under the negligence rule, plaintiffs would have to prove what reasonable care requires within a technologically complex and evolving environment. This evidentiary burden is comparable to, if not greater than, the burden faced by a consumer trying to prove that a soda manufacturer failed to adopt reasonably safe systems of quality control in the case of an exploding bottle.

Due to the safety problems that would be predictably created by an under-enforced rule of negligence liability, the failure of an operating system to

---

209. See *supra* Part II.A (explaining why crashes caused by a programming error or bug would seem to be obvious examples of malfunctions subject to strict products liability). For these same reasons, if the operating system were designed with symbolic, rule-based artificial intelligence, the crash of the vehicle would be a malfunction because the manufacturer would presumably take that system off the market so that it could modify the program to incorporate a new pre-programmed rule that would address this problem. By contrast, an operating system based on data-driven, machine learning is a borderline case. On the one hand, the operating system is designed to account for the problem, and so the manufacturer would not stop using the operating system following such a crash. On the other hand, the operating system is modified in the sense that it learns to solve the problem, and such a modification is arguably tantamount to pulling the prior design from the market.

210. For example, “[e]ach attached bit of hardware has a special software program called a *driver* that enables that bit of hardware to speak with the rest of the operating system it is installed onto. Driver problems are another major cause of system failure.” LIPSON & KURMAN, *supra* note 30, at 99.

211. *Id.* at 100.

212. See *id.* at 104 (“Driverless cars need an operating system that’s highly modular and redundant, similar to those that guide airplanes.”).

perform in its intended manner due to either a computer bug or third-party hacking provides an inference of defect—a product malfunction—that justifies strict liability. The liability would give the manufacturer the necessary financial incentive for ensuring the reasonable reliability of the operating system. This rule of strict liability only channels a limited number of crashes into the tort system and does not approach the rule of absolute liability that courts have uniformly rejected. For the same reasons that apply to crashes caused by programming bugs, the manufacturer will be subject to strict liability for crashes caused by hacking under the malfunction doctrine or its equivalent, the ordinary consumer expectations test.<sup>213</sup>

### B. Potential Limitations of Liability to Negligence

A manufacturer can avoid strict liability for a product malfunction by adequately warning about the performance in question.<sup>214</sup> A manufacturer's disclosure of the risk-adjusted insurance premium, for example, adequately warns consumers about the inherent risk that the fully functioning operating system will cause the vehicle to crash.<sup>215</sup> Consequently, the materialization of the risk—the crash itself—cannot be a product malfunction that frustrates consumer expectations. For these same reasons, an adequate warning about the risk of hacking would foreclose claims of strict liability based on product malfunctions, limiting manufacturer liability to negligence or the failure to adopt reasonably safe systems of cybersecurity as required by the risk-utility test.<sup>216</sup>

It is a separate question whether the manufacturer can adequately warn about the risk of hacking. Unlike the inherent risk that the fully functioning autonomous vehicle will crash, the manufacturer cannot reliably determine the current risk of cyberattack.<sup>217</sup> At best, the manufacturer can only warn consumers that the operating system *might* be hacked.

For some courts, this warning might be enough, whereas others could easily disagree. If merely warned that the vehicle might be hacked, the ordinary consumer would not obviously deem the occurrence of hacking to be an expected product performance rather than a product malfunction. The warning, after all, says very little about the risk. What is the likelihood that the vehicle will actually perform in this manner? Is the vehicle particularly vulnerable to

---

213. See *supra* notes 85, 91–95, 110–11 and accompanying text (explaining why the ordinary consumer expectations test is limited to malfunctioning products, with the modified consumer expectations test—or its substantive equivalent, the risk-utility test—governing the issue of whether a non-malfunctioning product is nevertheless defectively designed).

214. See *supra* Part II.B.1.

215. See *supra* Part II.D.2.

216. See *supra* notes 100–03 and accompanying text (explaining why a warning only defeats claims of product malfunction but does not otherwise satisfy the manufacturer's independent duty to design the product in a reasonably safe manner).

217. See *supra* notes 187–89 and accompanying text.

this safety problem as compared to competitors? A warning that the vehicle *might* be hacked would not adequately answer these questions. Courts could conclude that such a warning would not preclude a jury from finding that the crash of a hacked vehicle is a product malfunction.

A malfunctioning product, however, is not necessarily subject to strict liability. Under the widely adopted *Restatement (Second) of Torts* section 402A rule of strict products liability, comment *k* exempts “[u]navoidably unsafe products” from strict liability.<sup>218</sup> Because section 402A restates a body of case law based on product malfunctions,<sup>219</sup> the comment *k* exemption from this rule of strict liability presumably relates to malfunctions of “unavoidably unsafe” products.

Initially, most courts “embraced the rule of comment *k* without detailed analysis of its language.”<sup>220</sup> The only examples of “unavoidably unsafe products” provided by comment *k* involve drugs and vaccines. “While comment *k* could be read to apply to other products, it does not really give us any examples or suggest other areas where the policy balancing is precisely the same. For this reason, the courts and most commentators have assumed that comment *k* relates to pharmaceuticals.”<sup>221</sup>

The policy balancing that justifies the immunity in comment *k* is not necessarily limited to drugs and vaccines. The immunity is based on the policy conclusion that strict liability could disrupt the supply of drugs and vaccines, thereby limiting the potential for these products to promote public health and safety.<sup>222</sup> Like drugs and vaccines, autonomous vehicles are safety-enhancing products, and so the question is whether they are also “unavoidably unsafe” products that should be immunized from strict liability.

The liability issues involving contaminated blood demonstrate the rationale for comment *k*. According to one of the founders of strict products liability, Justice Roger Traynor of the California Supreme Court, blood is a “classic example” of an “unavoidably unsafe product” under comment *k*.<sup>223</sup> Donated blood, whether used in transfusions or blood products, has transmitted diseases such as AIDS and hepatitis, causing widespread injuries among

---

218. RESTATEMENT (SECOND) OF TORTS § 402A cmt. k (AM. LAW. INST. 1965).

219. Michael D. Green, *The Unappreciated Congruity of the Second and Third Tort Restatements on Design Defects*, 74 BROOK. L. REV. 807, 812 (2009) (footnote omitted); see also *supra* notes 194–98 and accompanying text (explaining how the rule of strict products liability evolved from cases of contaminated food and exploding soda bottles).

220. *Brown v. Superior Court*, 751 P.2d 470, 476 (Cal. 1988).

221. Victor E. Schwartz, *Unavoidably Unsafe Products: Clarifying the Meaning and Policy Behind Comment K*, 42 WASH. & LEE L. REV. 1139, 1141 (1985).

222. The ensuing argument is drawn from GEISTFELD, PRODUCTS LIABILITY, *supra* note 47, at 170–86.

223. Roger J. Traynor, *The Ways and Meanings of Defective Products and Strict Liability*, 32 TENN. L. REV. 363, 367 (1965).

hemophiliacs and other users of blood products.<sup>224</sup> Contaminated blood departs from the product specifications of pure blood, much like contaminated food departs from the product specifications of wholesome food. In both instances, the contaminated product malfunctions or does not perform in its expected or intended manner. Such a product would not be marketable with its true character known, rendering it defective. Without the exemption afforded by comment *k*, the sellers of contaminated blood would incur strict liability for these malfunctions, the conclusion that a few courts have reached.<sup>225</sup>

As compared to most products, those in the blood-products market would be exposed to a substantially greater amount of strict liability. In large part, the increased liability stems from new blood-borne diseases that cannot be detected at the time of sale.<sup>226</sup> Moreover, the risk of contaminated blood cannot always be reduced to more ordinary levels once tests become available for detecting the virus or other contaminants.<sup>227</sup> The rate and number of injuries caused by contaminated blood are far greater than the rate and number of injuries caused by malfunctions of other products, such as an exploding bottle of soda.

Because contaminated blood has caused thousands of injuries, strict liability would have a devastating effect on the financial viability of the blood-products industry. In a class-action lawsuit filed by hemophiliacs infected with HIV, Judge Richard Posner concluded that defendant manufacturers might easily have been “facing \$25 billion in potential liability (conceivably more).” Such liability would “hurl the industry into bankruptcy,” and with it “a major segment of the international pharmaceutical industry.”<sup>228</sup>

The bankruptcy of blood suppliers and other pharmaceutical manufacturers would create social problems qualitatively different from those created by the bankruptcy of other product manufacturers, like the suppliers of

---

224. For an excellent description of the factual context and litigation history regarding HIV-contaminated blood, see Eric A. Feldman, *Blood Justice: Courts, Conflict, and Compensation in Japan, France, and the United States*, 34 LAW & SOC'Y REV. 651 (2000).

225. See *Cunningham v. MacNeal Mem'l Hosp.*, 266 N.E.2d 897, 903 (Ill. 1970) (applying strict liability to sale of blood contaminated by hepatitis virus that was not reasonably detectable at the time of sale); see also *Cnty. Blood Bank, Inc. v. Russell*, 196 So.2d 115, 116 (Fla. 1967) (recognizing claim for breach of implied warranty against a blood bank that sold blood contaminated with the hepatitis virus). Each of these decisions was subsequently overruled by blood-shield statutes discussed in text below.

226. For example, HIV entered the blood supply in the 1970s. The test for detecting HIV in blood was not available until 1985. At that time, “the rate of infection among donors in San Francisco was found to be 1 in 2,632.” Michael J. Miller, *Strict Liability, Negligence and the Standard of Care for Transfusion-Transmitted Disease*, 36 ARIZ. L. REV. 473, 480 (1994) (citation omitted). By the late 1980s, almost half of America's 20,000 hemophiliacs were HIV-positive and there were about 29,000 other individuals who were HIV-positive because of blood transfusions. Feldman, *supra* note 224, at 669.

227. Almost ten years after the development of a test for detecting the presence of HIV in blood, the “chances of being infected by HIV through blood products [was about] 1 in 68,000 units transfused,” causing an estimated “90 cases of transfusion transmitted AIDS a year.” Miller, *supra* note 226, at 479–80 (citations omitted).

228. *In re Rhone-Poulenc Rorer Inc.*, 51 F.3d 1293, 1300 (7th Cir. 1995).

soda. Blood and other pharmaceutical products are necessary for public health and safety. Blood products save lives. Despite the beneficial safety incentive that it otherwise creates, strict liability can have a self-defeating safety effect by significantly disrupting the supply of these life-saving products.

Based on this rationale for comment *k*, most courts have exempted blood suppliers from strict liability.<sup>229</sup> The exemption is legislatively enshrined in blood-shield statutes. Virtually all states now have statutes protecting hospitals and blood banks from strict liability for the sale of contaminated blood.<sup>230</sup> These statutes “reflect a legislative judgment that to require providers to serve as insurers of the safety of these materials might impose such an overwhelming burden as to discourage the gathering and distribution of blood.”<sup>231</sup> Strict liability would unduly threaten an outcome that would be contrary to the safety rationale for products liability, which explains why comment *k* immunizes these “unavoidably unsafe” products from strict liability.<sup>232</sup>

The rationale for comment *k* has obvious relevance for automated driving technologies. Like blood and pharmaceutical products, autonomous vehicles promote public safety. Like blood and pharmaceutical products, autonomous vehicles are subject to strict liability under the malfunction doctrine in order to overcome the evidentiary difficulties of establishing negligence liability, thereby restoring the manufacturer’s financial incentive to adopt reasonably safe systems of quality control. As comment *k* would seem to require, a strict liability rule that serves to protect consumers from injury should yield to those same public health and safety concerns when necessary. This limitation finds further support in other tort rules.<sup>233</sup> Using this same reasoning, courts could

---

229. *E.g.*, *Kozup v. Georgetown Univ.*, 663 F. Supp. 1048, 1058–61 (D.D.C. 1987) (relying on comment *k* to reject claim of strict liability for the sale of blood contaminated by HIV), *aff’d in part, vacated in part on other grounds*, 851 F.2d 437 (D.C. Cir. 1988); *see also* RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 19 cmt. c (AM. LAW INST. 1998) (“Where legislation has not addressed the problem, courts have concluded that strict liability is inappropriate for harm caused by such product contamination.”).

230. *See* RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 19 cmt. c (AM. LAW INST. 1998) (“Absent a special rule dealing with human blood and tissue, such contamination presumably would be subject to the rule [of strict liability for a manufacturing defect]. . . . However, legislation in almost all jurisdictions limits the liability of sellers of human blood and human tissue to the failure to exercise reasonable care, often by providing that human blood and human tissue are not ‘products’ or that their provision is a ‘service’ [and therefore not subject to strict *products* liability].”).

231. *Zichichi v. Middlesex Mem’l Hosp.*, 528 A.2d 805, 810 (Conn. 1987).

232. *See, e.g.*, *Belle Bonfils Mem’l Blood Bank v. Hansen*, 665 P.2d 118, 124 (Colo. 1983) (“[T]he *raison d’etre* of strict liability is to force some hazardous products out of the market. The same rationale does not apply to blood or vaccines which are life-saving and which have no known substitutes.”).

233. *See* RESTATEMENT (SECOND) OF TORTS § 520(f) (AM. LAW INST. 1965) (providing that the social value of an activity is a factor that forecloses a finding that an activity is otherwise abnormally dangerous and subject to strict liability); Mark A. Geistfeld, *Social Value as a Policy-Based Limitation of the Ordinary Duty to Exercise Reasonable Care*, 44 WAKE FOREST L. REV. 899, 903–16 (2009) (discussing case law showing that the rationale for an ordinary tort duty can instead justify an extraordinary categorical limitation of that duty under the appropriate conditions) [hereinafter *Social Value*].

conclude that autonomous vehicles—like blood and pharmaceutical products—are “unavoidably unsafe” products with respect to the risk of hacking. In that event, manufacturers would be subject not to strict liability but to ordinary negligence liability for these malfunctions.

To be sure, this conclusion depends on questions that are hard to answer at this juncture. The immunity under comment *k* requires that the product be “properly prepared” and accompanied by a “proper warning,” and so manufacturers remain strictly liable for lapses of quality control like an inadequately sterile environment that contaminates a blood product.<sup>234</sup> To fall under the comment *k* exemption, the defect must instead threaten the entire product line with a substantial, correlated (systemic) risk that cannot be sufficiently reduced by the exercise of reasonable care. For example, HIV was undetectable when it first contaminated the blood supply, creating a systemic risk that could not be eliminated by reasonably safe methods of quality control. Is the risk of hacking analogous to the risk posed by new blood-borne diseases? Would strict liability for hacking result in extensive, largely unavoidable liabilities like those faced by the manufacturers of blood products contaminated with undetectable viruses? Or is hacking instead analogous to an ordinary lapse of quality control like inadequately sterile environments, involving a risk that can be sufficiently reduced by the exercise of reasonable care?

The disruptive effect of strict liability further depends on the extent to which manufacturers will be able to purchase insurance covering liabilities for hacked vehicles.<sup>235</sup> The availability of insurance for terrorism-related cyberattacks could be particularly problematic.<sup>236</sup> If manufacturers cannot procure liability insurance or if their liability exposure is sufficiently systemic such that it would otherwise unduly threaten bankruptcy, then there is a strong case for immunizing this type of malfunction from strict products liability.

---

234. RESTATEMENT (SECOND) OF TORTS § 402A cmt. k (AM. LAW INST. 1965) (explaining further that “many new or experimental drugs” are unavoidably unsafe if “because of lack of time and opportunity for sufficient medical experience, there can be no assurance of safety, or perhaps even of purity of ingredients, but such experience as there is justifies the marketing and use of the drug notwithstanding a medically recognizable risk”).

235. For reasons previously discussed in *supra* notes 187–89 and accompanying text, insurers will probably not be able to rely on a manufacturer’s loss experience to experience rate the premium in its entirety. The liability risk, however, may still be insurable. Rather than charging different policyholders different premiums based on their individual risk characteristics, the insurer can charge all policyholders the same premium based on the expected level of liability for the group of manufacturer policyholders. See LLOYD’S, AUTONOMOUS VEHICLES, *supra* note 180, at 18 (“As autonomous and unmanned vehicles become more commercially available, cyber risk policies will most likely be developed to suit the needs of stakeholders such as operators, systems designers, manufacturers, and infrastructure providers.”). *But see* Crane et al., *supra* note 25, at 73 (observing that “[s]uch intentionally caused losses, as with terrorism-related risks, are especially difficult for insurers to predict”).

236. For extended discussion of the reasons why these risks are difficult to insure, see generally Michelle E. Boardman, *Known Unknowns: The Illusion of Terrorism Insurance*, 93 GEO. L.J. 783 (2005).

Whether the market is vulnerable to this problem, and whether courts will respond by applying comment *k* to autonomous vehicles, is not presently clear for these and other reasons. Courts, for example, might distinguish blood and medical products from autonomous vehicles on the ground that individuals have no meaningful choice to use the former products, unlike the latter. Is comment *k* limited to life-saving products for which there is no real choice, or to life-saving products in general? As questions like this one show, cybersecurity is a potential source of systemic legal uncertainty for autonomous vehicles that is quite different from the more readily resolved questions concerning the manufacturer's liability for the crash of a fully functioning autonomous vehicle that has not been hacked.

#### IV.

##### REDUCING UNCERTAINTY BY COORDINATING STATE TORT LAW WITH FEDERAL SAFETY REGULATIONS

Aside from the issues of cybersecurity and the testing conditions required to sufficiently establish an autonomous vehicle's relative safety performance,<sup>237</sup> manufacturers will not be subject to highly uncertain forms of tort liability. This method of safety regulation, however, creates other uncertainties and the associated costs. In addition to being necessarily predictive, the foregoing liability conclusions are based on tort rules adopted by most states; some states may rely on rules that are modestly or even significantly different. How much variability can be expected? And even if courts around the country ultimately resolve the liability issues in the same manner, tort claims are usually costly to litigate and often proceed slowly through the civil justice system. The expense, time, and potential for disagreement among courts explain why an alternative regulatory approach holds so much appeal in this area.

Stepping into this void, state legislators have begun to address the regulation of autonomous vehicles.<sup>238</sup> As of early 2016, fifteen states and the District of Columbia have "passed or introduced bills related to self-driving vehicles, with California, Michigan, and Nevada likely to set the standards to be adopted by the others."<sup>239</sup>

An approach based on state regulatory law, however, creates the same troubling problem as the approach based on state tort law:

---

237. See *supra* notes 156–61 and accompanying text (identifying various issues involving testing conditions that are not clearly resolved by established doctrines).

238. The legislative activity is fully documented at *Automated Driving: Legislative and Regulatory Action*, CTR FOR INTERNET & SOC'Y, [http://cyberlaw.stanford.edu/wiki/index.php/Automated\\_Driving:\\_Legislative\\_and\\_Regulatory\\_Action](http://cyberlaw.stanford.edu/wiki/index.php/Automated_Driving:_Legislative_and_Regulatory_Action) [https://perma.cc/9JS8-DPEB] (last modified Apr. 27, 2017).

239. ALBRIGHT ET AL., KPMG, *supra* note 6, at 3.



Currently, the major difficulty is overcoming the regulatory fragmentation caused by 50 states with differing preferences on licensing, car standards, regulation, and privacy protection. Right now, car manufacturers and software developers face conflicting rules and regulations in various states. This complicates innovation because makers want to build cars and trucks for a national or international market. Greater clarity in regard to legal liability and data protection is also needed. Addressing these issues would help manufacturers implement new technologies and help to spur economic growth in transportation.<sup>240</sup>

The potential variability of state regulatory law would exacerbate the systemic uncertainty that state tort law already generates, strengthening the case for uniform federal regulation.

Keenly aware of these issues, NHTSA in early 2016 staked out a preliminary federal regulatory strategy for automated driving:

DOT and NHTSA policy is to facilitate and encourage wherever possible the development and deployment of technologies with the potential to save lives. To that end, NHTSA will use all available tools to determine the safety potential of new technologies; to eliminate obstacles that would prevent or delay technology innovations from realizing that safety potential; and to work with industry, governmental partners at all levels, and other stakeholders to develop or encourage new technologies and accelerate their adoption where appropriate.<sup>241</sup>

Though laudable, NHTSA's plan to facilitate the rapid and yet safe deployment of autonomous vehicles faces daunting problems. Like other government agencies, NHTSA has budgetary concerns. "Instead of increasing as time has progressed, the agency's funding has decreased at a time when automotive technology and the demands of investigating defects have increased."<sup>242</sup> As the interim administrator of NHTSA observed in late 2014, the Federal Aviation Administration "has close to fifty thousand employees—an order of magnitude more employees than we do. We have six hundred. . . . With more resources, we could save more lives. And each time the answer [to our request for more resources] from Congress has been no. Zero."<sup>243</sup>

---

240. DARRELL M. WEST, BROOKINGS INST., SECURING THE FUTURE OF DRIVERLESS VEHICLES (2016), <https://www.brookings.edu/research/securing-the-future-of-driverless-cars> [<https://perma.cc/M7WT-YH6E>].

241. NAT'L HIGHWAY TRAFFIC SAFETY ADMIN., "DOT/NHTSA POLICY STATEMENT CONCERNING AUTOMATED VEHICLES": 2016 UPDATE TO "PRELIMINARY STATEMENT OF POLICY CONCERNING AUTOMATED VEHICLES" 1 (2016), [www.nhtsa.gov/staticfiles/rulemaking/pdf/Autonomous-Vehicles-Policy-Update-2016.pdf](http://www.nhtsa.gov/staticfiles/rulemaking/pdf/Autonomous-Vehicles-Policy-Update-2016.pdf) [<https://perma.cc/33PP-YTPE>] [hereinafter NHTSA, 2016 PRELIMINARY REPORT].

242. Joel Finch, *Toyota Sudden Acceleration: A Case Study of the National Highway Traffic Safety Administration: Recalls for Change*, 22 LOY. CONSUMER L. REV. 472, 492 (2010).

243. Malcolm Gladwell, *The Engineer's Lament: Two Ways of Thinking About Automotive Safety*, NEW YORKER (May 4, 2015), <http://www.newyorker.com/magazine/2015/05/04/the-engineers-lament> [<https://perma.cc/YL99-452Z>] (reporting on interview conducted in late 2014).

Even if the political climate were favorable for substantially increasing NHTSA's budget, the regulatory problem would still be overwhelming. The market is already incorporating automated safety technologies into motor vehicles, and the trend presumably will only escalate.<sup>244</sup> According to the National Transportation Research Board, the large number of safety features in the wide variety of motor vehicles makes it infeasible for NHTSA to comprehensively regulate the entire product market.<sup>245</sup> The magnitude of the safety problem requires an approach based on both federal regulatory law and state tort law.

In September 2016, NHTSA addressed this issue while further clarifying its strategy for regulating highly automated vehicle (HAV) technologies.<sup>246</sup> NHTSA recognizes that “[r]ules and laws allocating tort liability could have a significant effect on both consumer acceptance of HAVs and their rate of deployment.”<sup>247</sup> In particular, “a patchwork of inconsistent laws and regulations among the 50 States and other U.S. jurisdiction . . . could delay the widespread deployment of these potentially lifesaving technologies.”<sup>248</sup> Because “a manufacturer should be able to focus on developing a single HAV fleet rather than 50 different versions to meet individual state requirements,”<sup>249</sup> NHTSA “strongly encourages States to allow DOT alone to regulate the performance of HAV technology and vehicles.”<sup>250</sup> NHTSA, however, also “confirms that States retain their traditional responsibilities for vehicle licensing and registration, traffic laws and enforcement, and motor vehicle insurance and liability regimes.”<sup>251</sup>

This proposed regulatory approach poses an obvious problem. Based on the compelling need for national uniformity, NHTSA “strongly encourages” the states to let it alone regulate the safe performance of HAV technologies, but

---

244. See *supra* Part I.A.

245. NAT'L RESEARCH COUNCIL OF THE NAT'L ACADS., TRANSP. RESEARCH BD., SPECIAL REPORT 308: THE SAFETY PROMISE AND CHALLENGE OF AUTOMOTIVE ELECTRONICS: INSIGHTS FROM UNINTENDED ACCELERATION 182 (2012), <http://onlinepubs.trb.org/onlinepubs/sr/sr308.pdf> [<https://perma.cc/C96V-A8EM>] (“NHTSA cannot be expected to hire and maintain personnel having all of the specialized technical expertise and design knowledge relevant to the growing field of automotive electronics.”); see also *id.* at 177 (“It is difficult to see how NHTSA could obtain the capacity for identifying suitable testing methods [for electronic control] in light of the wide variability in the way manufacturers design and engineer vehicle systems.”); cf. NIDHI KALRA ET AL., RAND CORP., LIABILITY AND REGULATION OF AUTONOMOUS VEHICLE TECHNOLOGIES 38–39 (2009) [hereinafter RAND REPORT] (recounting the decades-long process for regulating airbags and observing that “experiences with air-bag regulation are particularly relevant to autonomous vehicle technologies and serve to illustrate many facets of regulation”).

246. HAV is defined in terms of all technologies within automation levels 3–5. NHTSA, 2016 AUTOMATED VEHICLES POLICY, *supra* note 4, at 10; see also *supra* notes 36–38 and accompanying text (describing the classification system adopted by NHTSA).

247. NHTSA, 2016 AUTOMATED VEHICLES POLICY, *supra* note 4, at 46.

248. *Id.* at 37.

249. *Id.* at 7.

250. *Id.* at 37.

251. *Id.* at 7.

NHTSA also contemplates that the states will “retain” their tort liability regimes. Products liability is a demanding form of safety regulation involving manufacturing, design, and warning defects, so how can the states retain this liability regime while also ceding sole regulatory authority to NHTSA? To attain uniformity across the country, state tort law must somehow be adequately coordinated with the federal regulatory regime.

Apparently recognizing as much, NHTSA suggests that “[i]t may be desirable to create a commission to study liability and insurance issues and make recommendations to the States.”<sup>252</sup> This approach, however, will not be adequately reactive to the rapidly developing market. Proposals for liability reforms would require legislative approval in each state, a cumbersome process that would not provide manufacturers with sufficient guidance until such reforms have been uniformly adopted across the country. Years are likely to elapse. A different approach is necessary.

To achieve its objective of a nationally uniform regulatory regime that allows the states to “retain their traditional responsibilities” with respect to tort law, NHTSA could rely on the state rules of strict products liability to inform federal regulations. By basing federal regulations on safety measures that would satisfy the associated tort obligations of most states, NHTSA would adequately account for the important state interest in tort law. As explained more fully below, the federal regulations would attain the desired degree of regulatory uniformity for reasons that implicate both state tort law—the regulatory compliance defense—and federal constitutional law—the doctrine of implied statutory preemption. State tort law would supplement the federal regulations in important instances, yielding a comprehensive regulatory approach of the type envisioned by NHTSA.

As we have found, the existing regime poses no apparent obstacle to the ongoing development of driver-assistance systems (DAS), largely limiting the regulatory problem to the new safety issues posed by driverless vehicles.<sup>253</sup> The most important problem involves the possibility that the fully functioning operating system of an autonomous vehicle will cause the vehicle to crash. The manufacturer must first warn consumers about the conditions under which the vehicle can be safely deployed and design the operating system to correct for unsafe deployments.<sup>254</sup> To satisfy the remaining tort obligations and thereby avoid liability for such a crash, the manufacturer must subject the vehicle to adequate premarket testing and provide adequate warnings about the inherent risk that a properly deployed, fully functioning vehicle will cause a crash.<sup>255</sup> Each of these two tort obligations can form the basis of federal regulations. State tort law can then enforce the regulations and fill in gaps in the federal

---

252. *Id.* at 46.

253. *See supra* Part I.

254. *See supra* notes 60–62 and accompanying text.

255. *See supra* Part II.

scheme. The resultant regime is a composite of the state law of strict products liability and federal regulatory law, with automobile insurance providing a critical component of the information disclosure mandated by the federal regulations.

A. *Federal Regulations Requiring Premarket Testing and Post-Sale Updates of the Operating System*

NHTSA regulations are embodied in safety standards that specify minimum performance requirements for motor vehicles.<sup>256</sup> These performance standards “shall be practicable, meet the need for motor vehicle safety, and be stated in objective terms.”<sup>257</sup> “To avoid impeding . . . innovation,” NHTSA cannot prescribe “how manufacturers should meet the requirements through their product design, development, and production processes.”<sup>258</sup> Manufacturers instead certify that their products satisfy the mandated performance standards.<sup>259</sup>

Pursuant to this statutory authority, NHTSA in 2016 announced that it will consider “potential . . . regulatory action to design and implement new standards . . . to govern the initial testing and deployment of HAVs.”<sup>260</sup> The regulatory rationale is straightforward. “Essential to the safe deployment of such vehicles is a rigorous testing regime that provides sufficient data to determine safety performance and help policymakers at all levels make informed decisions about deployment.”<sup>261</sup> In addition to uncovering programming errors that might cause the operating system to malfunction, extensive premarket testing will also generate the machine learning that improves the safety performance of the operating system.<sup>262</sup> These outcomes provide a rational basis for the regulatory conclusion that adequate premarket testing is required for ensuring that the vehicle can perform in a reasonably safe manner.

As compared to the tort system, NHTSA is better situated to determine the appropriate testing conditions for evaluating the safety performance of an autonomous vehicle. What are the necessary road conditions? How many miles

---

256. See *supra* note 31 and accompanying text (describing enabling legislation for NHTSA).

257. 49 U.S.C. § 30111(a) (2012).

258. NAT’L RESEARCH COUNCIL OF THE NAT’L ACADS., TRANSP. RESEARCH BD., *supra* note 245, at 27.

259. Wood et al., *supra* note 31, at 1435.

260. NHTSA, 2016 AUTOMATED VEHICLES POLICY, *supra* note 4, at 36; see also *id.* at 70 (“Among the categories of new regulatory tools and authorities DOT might apply to regulate the safety of [highly automated vehicles] are pre-market safety assurance tools. Such tools could include pre-market testing, data, and analyses reported by a vehicle manufacturer or other entity to DOT.”). In contrast to the current regulatory approach that permits the manufacturer to self-certify about the vehicle’s compliance with safety standards, NHTSA might instead require pre-market approval or a hybrid self-certification/approval approach. *Id.* at 72–76.

261. NHTSA, 2016 PRELIMINARY REPORT, *supra* note 241, at 1.

262. See *supra* notes 121–31 and accompanying text.

should be driven on freeways and in urban conditions? How many total miles must be logged by an operating system to generate sufficiently reliable crash data? What other metrics are required for adequately measuring safe performance?<sup>263</sup> The tort system would resolve these issues through the adversarial presentation of evidence addressing the particular claims before a court. NHTSA, by contrast, would use its specialized expertise to comprehensively address these matters through the administrative rule-making process, giving it a comparative institutional advantage for determining the appropriate testing criteria for evaluating the safety performance of an autonomous vehicle.

In one important respect, however, the associated tort standard should guide NHTSA. As previously discussed, an autonomous vehicle that at least halves the rate of crashes relative to a conventional vehicle would necessarily have a reasonably safe or non-defective operating system for tort purposes.<sup>264</sup> Others have proposed such a performance standard without otherwise explaining its rationale,<sup>265</sup> including the administrator of NHTSA in 2016: “We need to start with two times better [safety performance than conventional vehicles]. We need to set a higher bar if we expect safety to actually be a benefit here.”<sup>266</sup> The rationale for this performance standard can be derived from the associated tort obligation, which does not require more extensive testing under the risk-utility test. The further pursuit of safety would be self-defeating at this point, creating disutility or safety costs (due to delayed deployment of the life-saving technology) greater than the associated risk reduction or safety benefits of more extensive testing (the improved safety performance caused by machine learning of the operating system). By requiring a fully functioning autonomous vehicle to be at least twice as safe as conventional vehicles, the federal standard would ensure that the operating system is reasonably safe for reasons the associated tort standard makes clear.

This premarket testing requirement is a minimum performance standard—the operating system can still be improved—that would obviously comply with NHTSA’s legislative mandate to “reduce traffic accidents and deaths and injuries resulting from traffic accidents.”<sup>267</sup> Because the performance standard

---

263. Cf. NHTSA, 2016 AUTOMATED VEHICLES POLICY, *supra* note 4, at 18 (“To develop new safety metrics, manufacturers and other entities should collect, store and analyze data regarding positive outcomes in addition to the type of reporting conditions listed above (event, incident, and crash data). Positive outcomes are events in which the HAV system correctly detects a safety-relevant situation, and successfully avoids an incident (e.g., ‘near misses’ and edge cases).”).

264. See *supra* Part II.D.1.

265. See LIPSON & KURMAN, *supra* note 30, at 100–02 (proposing without defending “a bolder baseline: in order to be considered legal, driverless cars must be twice as safe as the average human driver”).

266. Keith Naughton, *Regulator Says Self-Driving Cars Must Be Twice as Safe*, BLOOMBERG (June 8, 2016), <https://www.bloomberg.com/news/articles/2016-06-08/u-s-auto-regulator-says-self-driving-cars-must-be-twice-as-safe> [<https://perma.cc/99J5-XJAX>].

267. 49 U.S.C. § 30101 (2012).

satisfies a cost-benefit analysis of the safety problem (the risk-utility test), it would also comply with executive orders requiring federal agencies to justify proposed major regulations by showing that the costs of regulation are less than the benefits.<sup>268</sup>

A federal regulation of this type would also solve a difficult problem posed by the tort claim: What is the relevant baseline for evaluating the relative safety of an autonomous vehicle? To prove that the design of a product is defective, the tort plaintiff must identify a reasonable alternative design that is within the “state of the art”—a requirement that could permit the plaintiff to compare the autonomous vehicle’s safety performance to other (reasonably designed) autonomous vehicles that have already been commercially distributed by other manufacturers.<sup>269</sup> This baseline is problematic, however, because it would provide a considerable competitive advantage for the “first movers” in the commercial distribution of autonomous vehicles. Due to machine learning, the collective driving experience of a fleet already on the road will significantly enhance the safety performance of each vehicle. Without such experience to draw upon in the design of its operating system, a new entrant would often face a substantial cost disadvantage (of more extensive testing) to match the performance of other commercially available vehicles. Requiring new entrants to equal or exceed the safety performance of autonomous vehicles already on the road, therefore, could easily undermine competition in the market by entrenching the first movers, thereby suppressing the technological innovations otherwise offered by new entrants. This dynamic presumably will change as the market matures, but those are not the market conditions regulators now face. To allow new entrants sufficient access to this newly developing market, the regulatory approach should evaluate the safety performance of an autonomous vehicle in relation to conventional vehicles. The foregoing analysis is based on this approach, which is clearly within NHTSA’s statutory mandate to formulate its performance standards in a technologically neutral manner that does not impede innovation.<sup>270</sup>

NHTSA could then supplement the premarket-testing requirements with an additional post-sale obligation to update the operating system when required by concerns for safety or cybersecurity. Once a vehicle has been commercially

---

268. See Exec. Order 13,563, 76 Fed. Reg. 3821 (Jan. 18, 2011). The political history of these executive orders is recounted in RICHARD L. REVESZ & MICHAEL A. LIVERMORE, *RETAKING RATIONALITY: HOW COST-BENEFIT ANALYSIS CAN BETTER PROTECT THE ENVIRONMENT AND OUR HEALTH* 21–45 (2008). For extended argument that the tort entitlement provides a defensible basis for the cost-benefit analysis of federal health and safety regulations, see generally Mark A. Geistfeld, *The Tort Entitlement to Physical Security as the Distributive Basis for Environmental, Health, and Safety Regulations*, 15 *THEORETICAL INQUIRIES* L. 387 (2014).

269. See RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 2 *reps. n. cmt. d(IV)(B)* (AM. LAW INST. 1998) (“The term ‘state of the art’ has been variously defined by a multitude of courts. For some it refers to industry custom or practice; for others it means the safest existing technology that has been adopted for use; for others it means cutting-edge technology.”).

270. See *supra* notes 257–59 and accompanying text.

distributed, the machine learning of the fleet will substantially improve the safety performance of the operating system. Mandated updates of this type—a post-sale obligation to alter the product design in order to improve safety performance—could be implemented “through over-the-air updates or other means.”<sup>271</sup> NHTSA has statutory authority to regulate “any after-market software updates to the autonomous driving system.”<sup>272</sup> Indeed, regulations requiring post-sale updates to the operating system would complement existing regulations requiring product recalls of conventional vehicles with newly discovered safety defects.<sup>273</sup>

### B. Federal Regulations Requiring Product Warnings

NHTSA’s proposed regulatory strategy includes product warnings or instructions that enable “the human driver or operator of the vehicle to easily understand the capabilities and limitations of each HAV system.”<sup>274</sup> For example, the operator’s manual or a disclosure system integrated into the vehicle’s interface with the operator “should fully describe the capabilities and limitations of the HAV systems in each operational design domain, including operational speeds, geographical areas, weather conditions and other pertinent information. . . .”<sup>275</sup> The evident rationale for these disclosures is confirmed by tort law, which imposes the same warning obligations on manufacturers,<sup>276</sup> although NHTSA’s regulatory approach would benefit in another important respect by drawing on tort doctrine.

To inform consumers about the driving capabilities and limitations of an autonomous vehicle, the product warning must include an adequate disclosure about the inherent risk that the fully functioning operating system will cause the vehicle to crash. As we have found, the manufacturer can satisfy this tort obligation by disclosing the annual, risk-adjusted premium for insuring the vehicle, with the cost of insuring the vehicle against collision damage separated from the cost of insuring bodily injury.<sup>277</sup> Such a disclosure would also further NHTSA’s policy objectives and could be a regulatory requirement.

---

271. NHTSA, 2016 AUTOMATED VEHICLES POLICY, *supra* note 4, at 16 (envisioning these types of updates).

272. Wood et al., *supra* note 31, at 1443.

273. NHTSA possesses the authority to mandate a recall of motor vehicles that do not comply with federal regulations or are defective and unsafe. 49 U.S.C. § 30118(b) (2012). Instead of directly exercising this regulatory authority, NHTSA often relies on “voluntary” recalls that proceed without any agency involvement. KEVIN M. McDONALD, SHIFTING OUT OF PARK: MOVING AUTO SAFETY FROM RECALLS TO REASON 72 (2006) (“Historically, nearly 80% of recalls are conducted without any NHTSA involvement. The remaining 20%, again conducted voluntarily, are what insiders euphemistically call ‘NHTSA-influenced.’”).

274. NHTSA, 2016 AUTOMATED VEHICLES POLICY, *supra* note 4, at 27–28.

275. *Id.* at 25.

276. *See supra* Part I.B.

277. *See supra* Part II.D.2.

The regulatory logic of this mandated disclosure parallels the rationale for federal regulations that require the manufacturers of certain electrical appliances, such as refrigerators, to disclose the annual energy costs of their products.<sup>278</sup> Based on these mandated disclosures, a consumer considering two different refrigerators can easily compare both their purchase price and energy costs. A refrigerator selling for \$900 with annual energy costs of \$100, for example, will be significantly more costly over a five-year period than a comparable refrigerator selling for \$1,000 with annual energy costs of \$50. The consumer can easily determine that the brand with the lower retail price will be more expensive over time because of its higher energy costs, bolstering the incentive for manufacturers to make refrigerators more energy efficient.

A similar market dynamic would be produced by the mandated disclosure of an autonomous vehicle's annual, risk-adjusted insurance premium. All else being equal, the vehicle with the lowest risk-adjusted insurance premium would be safer than competitors. By comparing prices, consumers could easily compare the relative safety performance of different autonomous vehicles, thereby incentivizing manufacturers to improve upon the safety performance of their vehicles in order to lower the associated insurance costs.<sup>279</sup>

NHTSA has already adopted this type of regulatory strategy. Under the New Car Assessment Program (NCAP), NHTSA provides safety ratings (from one to five stars) for certain newly manufactured vehicles "based on their success in frontal and side crash tests and in rollover resistance tests."<sup>280</sup>

High NCAP scores are widely used by vehicle manufacturers in advertising to demonstrate to potential buyers the safety attributes of the vehicles they produce. Thus, through providing the public with objective information on the relative safety performance of new vehicles, NCAP has been successful in achieving its purpose of creating consumer awareness of those differences, thereby creating market forces that prompt vehicle manufacturers to make added safety improvements to their vehicles.<sup>281</sup>

A regulation requiring manufacturers to disclose the annual, risk-adjusted premium for insuring their autonomous vehicles would produce this same market dynamic by providing consumers with "objective information on the relative safety performance of new vehicles." Indeed, NHTSA already provides consumers with information about the relative cost of insuring different makes and models against collision damage to the vehicle.<sup>282</sup> By excluding data about

---

278. See 16 C.F.R. § 305 (2008).

279. See *supra* notes 168–71 and accompanying text.

280. Wood et al., *supra* note 31, at 1437.

281. *Id.* at 1437–38.

282. NAT'L HIGHWAY TRAFFIC SAFETY ADMIN., COMPARISON OF DIFFERENCES IN INSURANCE COSTS FOR PASSENGER CARS, STATION WAGONS, PASSENGER VANS, PICKUPS, AND UTILITY VEHICLES ON THE BASIS OF DAMAGE SUSCEPTIBILITY 1 (2014) (providing information on "vehicles' collision loss experience in relative terms, with 100 representing the average for all



personal injury, this information “does *not* indicate a vehicle’s relative safety for occupants.”<sup>283</sup> Under the proposed regulatory approach, manufacturers would supply that information by disclosing the vehicle’s risk-adjusted premium covering bodily injury, thereby enabling consumers to directly compare the safety performance of one vehicle with another. Regulations could require manufacturers to supply this information through a new NCAP safety-rating category on the label that manufacturers must now affix to the side window of new automobiles.<sup>284</sup>

When coupled with premarket testing requirements showing that the autonomous vehicle performs at least twice as safely as conventional vehicles, a mandated disclosure of this type would further NHTSA’s policy objectives of “ensur[ing that] these technologies are safely introduced (i.e., do not introduce significant new safety risks), provide safety benefits today, and achieve their full safety potential in the future.”<sup>285</sup> The premarket testing requirements only help to ensure that the vehicle does not introduce significant new safety risks and provides safety benefits today; they do not otherwise incentivize manufacturers to improve upon this minimum performance standard once the vehicle has been introduced into the market. A regulation requiring disclosure of the annual, risk-adjusted insurance premium would give manufacturers a sufficient incentive to further improve the vehicle’s safety performance in order to reduce the premium and enhance the vehicle’s competitiveness within the market.

This approach depends on manufacturers sharing performance data with the insurance industry, which is another pillar of NHTSA’s proposed regulatory framework: “The data generated from [premarket testing] activities should be shared in a way that allows government, industry, and the public to increase their learning and understanding as technology evolves but protects legitimate privacy and competitive interests.”<sup>286</sup> By sharing the relevant crash data about the safety performance of an autonomous vehicle, the manufacturer will enable the insurance industry to set risk-adjusted premiums for the vehicle. The mandated disclosure of those premiums will then help the public to better understand how the HAV technology translates into safe driving behavior. Data

---

passenger vehicles”), <https://www.nhtsa.gov/sites/nhtsa.dot.gov/files/2014-comparison-insurance-costs-812039.pdf> [<https://perma.cc/R6NY-SGXB>].

283. *Id.*

284. NHTSA is authorized to adopt new NCAP “safety rating categories.” 15 U.S.C. § 1232(g)(2) (2012).

285. NHTSA, 2016 AUTOMATED VEHICLES POLICY, *supra* note 4, at 6; *see also id.* at 18 (“HAVs have great potential to use data sharing to enhance and extend safety benefits. Thus, each entity should develop a plan for sharing its event reconstruction and other relevant data with other entities. Such shared data would help to accelerate knowledge and understanding of HAV performance, and could be used to enhance the safety of HAV systems and to establish consumer confidence in HAV technologies.”).

286. *Id.* at 6.

sharing will make the insurance industry an integral component of the federal regulatory regime.

### C. Coordination of Federal and State Law

Concerned about the problem of uncertain and variable requirements of tort law, many have proposed to preempt these state requirements with federal law so that manufacturers will be uniformly regulated in the national market. “A patchwork of state laws governing the operation of self-driving cars threatens to stall their development, supporters told lawmakers as U.S. senators began consideration of a national standard for robotic vehicles” in March 2016.<sup>287</sup> Four industry insiders and an academic provided this testimony to a Senate committee; all urged Congress to enact laws that would uniformly regulate the safety of driverless vehicles.<sup>288</sup> Similarly, in a study sponsored by the California Department of Transportation, the RAND Corporation concluded that state tort liability “may lead to inefficient delays in the adoption of these technologies,” justifying a policy recommendation for the federal preemption of state law: “While federal preemption has important disadvantages, it might speed the development and utilization of this technology and should be considered, if accompanied by a comprehensive federal regulatory regime.”<sup>289</sup>

As we have found, NHTSA could regulate premarket testing and product warnings in a manner that would satisfy the tort obligations governing the manufacturer’s liability for the crash of a fully functioning autonomous vehicle. This liability question is the most important source of legal uncertainty manufacturers now face. By largely dissipating this uncertainty and its associated costs, these regulations would facilitate the safe deployment of autonomous vehicles in a manner that adequately coordinates federal regulatory law and state tort law.

#### 1. Overlap of Federal Regulatory Law and State Tort Law

For reasons of institutional comity, even if a court is not statutorily obligated to do so, it will defer to a legislative policy decision that is relevant to

---

287. Jeff Plungis & Keith Naughton, *Driverless Car Supporters Urge National Laws to Override State, Local, INS. J.* (Mar. 16, 2016), <http://www.insurancejournal.com/news/national/2016/03/16/402012.htm> [https://perma.cc/LW8W-6LP7].

288. *Hands Off: The Future of Self-Driving Cars: Hearing Before the S. Comm. On Commerce, Sci., & Transp.*, 114th Cong. (2016) (testimony of Dr. Chris Urmson, Director of Self-Driving Cars, Google X; Mr. Mike Ableson, Vice President, Strategy and Global Portfolio Planning, General Motors Co.; Mr. Glen DeVos, Vice President, Global Engineering and Services, Electronics and Safety, Delphi Automotive; Mr. Joseph Okpaku, Vice President of Government Relations, Lyft; and Dr. Mary (Missy) Louise Cummings, Director, Humans and Autonomy Lab and Duke Robotics, Duke University).

289. RAND REPORT, *supra* note 245, at 34, 37.

the resolution of any issue posed by a tort claim.<sup>290</sup> For example, if a motorist violates a traffic statute by speeding, courts will deem the violation to be negligence per se, even though the statute does not mandate this outcome.<sup>291</sup> Courts instead defer to the legislative policy decision about the reasonably safe speed for driving in this context. When rigorously applied across the full range of policy decisions implicated by a tort claim, this principle of common law deference largely coordinates state tort law with regulatory law.

As a matter of common law deference, courts consider whether a manufacturer's compliance with a federal safety regulation constitutes a complete defense to a tort claim. Unless the regulation is based on legislative policy decisions that fully resolve the tort claim, courts that fully defer to these policies must still make an independent tort judgment. That judgment could require the defendant manufacturer to take precautions beyond those required by the regulation. The failure to take any of those additional precautions would subject the manufacturer to liability, even though it complied with the regulation. The extent to which the regulatory policy decisions resolve the issues posed by the tort claim, therefore, determines whether regulatory compliance is a complete defense.<sup>292</sup>

Due to "the traditional view that the standards set by most product safety statutes or regulations generally are only minimum standards,"<sup>293</sup> regulatory compliance is usually not a complete defense. A safety regulation is a minimum standard for tort purposes when it does not account for all risks encompassed by the common law tort duty. By considering risks that the regulators did not account for, the court must make an independent tort judgment about whether the defendant was obligated to take care in excess of regulatory requirements. In such a case, the defendant's regulatory compliance is not a complete defense.<sup>294</sup>

But if a statute or regulation is based on safety decisions that fully resolve a tort claim, then regulatory compliance is a complete defense. By deferring to the legislative policy decisions embodied in such a regulation, the court can fully determine the defendant's tort obligations. No independent tort judgment

---

290. See Geistfeld, *Age of Statutes*, *supra* note 55, at 963–67 (discussing the principle of common law deference that supplies the basis for various tort doctrines).

291. See RESTATEMENT (THIRD) OF TORTS: LIABILITY FOR PHYSICAL AND EMOTIONAL HARM § 14 (AM. LAW INST. 2010) ("An actor is negligent if, without excuse, the actor violates a statute that is designed to protect against the type of accident the actor's conduct causes, and if the accident victim is within the class of persons the statute is designed to protect.").

292. See Geistfeld, *Age of Statutes*, *supra* note 55, at 991–1001.

293. RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 4 cmt. e (AM. LAW INST. 1998).

294. See Geistfeld, *Age of Statutes*, *supra* note 55, at 993–96; see also RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 4 cmt. e (AM. LAW INST. 1998) ("Thus, most product safety statutes or regulations establish a floor of safety below which product sellers fall only at their peril, but they leave open the question of whether a higher standard of product safety should be applied.").

is necessary. A defendant that complied with the regulation also necessarily satisfied the corresponding tort duty.<sup>295</sup>

The same conclusion would apply to a federal regulation that requires premarket testing showing that an autonomous vehicle at least halves the incidence of crashes relative to a conventional vehicle. The regulatory requirement would be a minimum standard in the sense that manufacturers could make their vehicles even safer. For tort purposes, however, the regulation would not be a minimum standard. The manufacturer would necessarily satisfy its tort obligation if the premarket testing showed that the vehicle at least halves the incidence of crashes relative to a conventional vehicle.<sup>296</sup> The manufacturer's compliance with such a regulation would be a complete defense to the associated tort claim.<sup>297</sup>

The same conclusion also applies to a federal regulation requiring the manufacturer to disclose the annual, risk-adjusted premium for insuring the autonomous vehicle. The regulation establishes a minimum standard only in the sense that the manufacturer could provide even more safety-related information. For tort purposes, however, compliance would also fully satisfy the manufacturer's associated tort obligation to warn about the inherent risk that the fully functioning operating system will cause the vehicle to crash.<sup>298</sup> Once again, regulatory compliance would be a complete defense to such a tort claim.

To be sure, not all states would necessarily apply the regulatory compliance defense in this manner.<sup>299</sup> Moreover, some states may formulate their liability rules differently from the majority, and so federal regulations based on the majority rule would not necessarily satisfy the associated tort obligations in those states. For either reason, some states could decide that a manufacturer's compliance with the foregoing regulations is not a complete defense. In that event, the federal constitutional doctrine of statutory preemption would instead coordinate these federal regulations with state tort law.

---

295. See Geistfeld, *Age of Statutes*, *supra* note 55, at 996–1001.

296. See *supra* Part II.B.2 (explaining why allegations of defective design for autonomous vehicles reduce to questions about the adequacy of premarket testing).

297. Because the result follows as a matter of deference, a court is not statutorily obligated to make regulatory compliance a complete defense. See RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 4 cmt. e (AM. LAW INST. 1998) (“Such a conclusion *may* be appropriate when the safety statute or regulation was promulgated recently, thus supplying currency to the standard therein established; when the specific standard addresses the very issue of product design or warning presented in the case before the court; and when the court is confident that the deliberative process by which the safety standard was established was full, fair, and thorough and reflected substantial expertise.”) (emphasis added).

298. See *supra* Part II.D.2.

299. Cf. *Malcolm v. Evenflo Co., Inc.*, 217 P.3d 514, 522 (Mont. 2009) (rejecting the *Restatement (Third)* rule for the regulatory compliance defense because it “conflicts with the core principles of Montana’s strict products liability law”).

Statutory preemption is based on the constitutional supremacy of federal law over state law.<sup>300</sup> A federal regulation can preempt or displace state law either through the express statutory language that authorizes the regulatory scheme or by implication. Because the enabling legislation for NHTSA does not permit the express preemption of state common law,<sup>301</sup> any regulations NHTSA adopts can preempt state tort law only by implication.<sup>302</sup> Preemption by implication occurs when compliance with the state-imposed tort duty “may produce a result inconsistent with the objective of the federal statute.”<sup>303</sup>

Under this analysis, “[t]he identification of the relevant ‘federal purpose’ necessitates an answer to the ‘minimum standard’ versus ‘optimal balance’ question.”<sup>304</sup> When federal regulators have the purpose of establishing this type of minimum safety standard, they recognize that states could impose more demanding tort requirements. National uniformity is not a regulatory objective. A state court, therefore, can enforce its more demanding tort requirements without frustrating the federal regulatory purpose, eliminating any role for implied preemption. By contrast, when regulators adopt an optimal safety standard, they have made an all-things-considered safety decision that accounts for all of the regulatory costs and benefits, including those of national uniformity. “If the federal standard sets the optimal balance, then state laws that diverge from it—either to relax or tighten regulations—are in ‘conflict’ with the ‘federal purpose’ and therefore preempt[ed].”<sup>305</sup>

In this respect, the proposed federal regulations governing premarket testing and disclosure of the autonomous vehicle’s risk-adjusted insurance premium are optimal rather than minimum standards. According to NHTSA, its policy objective is to exclusively regulate the safe performance of autonomous vehicles in order to attain a nationally uniform body of regulation that will facilitate the reasonably safe deployment of this life-saving technology.<sup>306</sup> NHTSA also has the policy objective of retaining the traditional role of state tort law. Both policy objectives are furthered by the proposed regulations, which is why the regulations embody optimal standards for purposes of implied preemption.

The proposed regulations derive from the common law tort duty as formulated by most states, making regulatory compliance a complete defense

---

300. U.S. CONST. art. VI, cl. 2. “The Supreme Court has repeatedly identified the Supremacy Clause as the source of its authority to declare state law displaced (preempted).” Thomas W. Merrill, *Preemption and Institutional Choice*, 102 NW. U. L. REV. 727, 733 (2008).

301. *Geier v. Am. Honda Motor Co., Inc.*, 529 U.S. 861, 868 (2000).

302. *Id.* at 869–74.

303. *Rice v. Santa Fe Elevator Corp.*, 331 U.S. 218, 230 (1947) (citation omitted) (discussing other instances of preemption as well, none of which matter for present purposes).

304. 2 AM. LAW INST., REPORTERS’ STUDY: ENTERPRISE RESPONSIBILITY FOR PERSONAL INJURY 108 (1991).

305. *Id.*

306. *See supra* notes 247–51 and accompanying text.

to the associated tort claims. In these states, there is no conflict between the federal interest in the uniform regulation of national markets and the traditional state interest in tort law. Permitting a few states to impose more demanding safety requirements on manufacturers would not only frustrate the regulatory objective of uniformity, it would also elevate the interests of these states over the substantially larger number that favor uniformity, thereby frustrating the regulatory purpose of optimally solving the federalism problem. For these and other reasons, the two federal regulations would impliedly preempt any state tort claims that seek to impose more demanding premarket testing or warning requirements on the manufacturers of autonomous vehicles.<sup>307</sup>

The same reasoning applies to other regulations with these same attributes. For example, NHTSA has adopted regulations concerning airbags and other passive restraint systems (such as automatic seat belts) that impliedly preempt any state-imposed tort requirements that are inconsistent or otherwise incompatible with the federal regulatory purpose.<sup>308</sup> So, too, new federal safety regulations that NHTSA adopts for automated driving technologies will preempt inconsistent or incompatible tort claims, further ensuring the coordination of federal regulatory law and state tort law when they have overlapping safety requirements.

## 2. *Federal Regulatory Law and State Tort Law as Supplements*

NHTSA will not be able to comprehensively regulate all safety aspects of automated driving technologies.<sup>309</sup> NHTSA has limited resources that it should expend on technologies with the greatest safety potential, much like it has done for the regulation of seat belts, airbags, antilock brakes, and rear-view cameras.<sup>310</sup> The resultant gaps in the federal regulatory regime will be filled by

---

307. The U.S. Supreme Court in a series of recent decisions has effectively conducted the implied preemption inquiry in tort cases by asking whether the regulation entails a safety decision (defined in cost-benefit or risk-utility terms) that is inconsistent with the one required by the tort claim. *See Geistfeld, Age of Statutes, supra* note 55, at 1004–17. Consider in this regard the preemptive effect of the federal regulation requiring premarket testing. The associated tort claim alleges that the vehicle is defective because a reasonable design of the operating system embodies more extensive testing (or machine learning) that would have prevented the vehicle from crashing in the case at hand. To recover, the plaintiff must prove that the cost of more extensive testing is *less* than the associated safety benefit, thereby rendering the existing design defective. *See supra* Part II.B.2 (explaining why allegations of defective design for autonomous vehicles reduce to questions about the adequacy of premarket testing). The federal regulation fully resolves the identical safety decision in a contrary manner. The premarket testing requirement is based on the policy conclusion that the safety cost of more extensive testing (due to delayed deployment of the life-saving technology) is *greater* than the safety benefits of more extensive testing. The regulation embodies the policy decision that further pursuit of safety via more extensive testing would be self-defeating. Permitting the plaintiff to recover for this claim, therefore, would be inconsistent with regulatory policy decision and is preempted for this reason alone.

308. *Geier v. Am. Honda Motor Co., Inc.*, 529 U.S. 861, 874–75 (2000).

309. *See supra* notes 241–51 and accompanying text.

310. *See* NAT'L HIGHWAY TRAFFIC SAFETY ADMIN., DOT-HS-812-069, LIVES SAVED BY VEHICLE SAFETY TECHNOLOGIES AND ASSOCIATED FEDERAL MOTOR VEHICLE SAFETY STANDARDS, 1960 TO 2012 (2015) (evaluating the effectiveness of safety technologies addressed by

state tort law, which imposes various obligations on manufacturers to supply reasonably safe or non-defective products. In a world of scarce regulatory resources, state tort law is a necessary supplement to federal regulation. Together they yield a regime that comprehensively regulates the safe performance of automated driving technologies.

State tort law would further supplement the federal regulatory regime by providing a robust enforcement mechanism. A good example involves the exploding airbags manufactured by Takata, which have recently led to a flurry of litigation and regulatory action.<sup>311</sup> Pursuant to federal regulation, every passenger vehicle manufactured since September 1, 1997, must be equipped with airbags that satisfy minimum performance standards.<sup>312</sup> In the Takata cases, the tort claims seek to hold the manufacturer responsible for its airbags that exploded (malfunctioned) and physically harmed consumers.<sup>313</sup> By enforcing the mandated performance standard for airbags, tort law (the doctrine of negligence per se) gives all manufacturers a strong financial incentive to fully comply with the federal regulation.<sup>314</sup>

Tort law could also incentivize manufacturers to follow regulatory proposals regarding the safety performance of autonomous vehicles. For example, to reduce the risk of crash from malfunctions of the operating system, NHTSA proposes that manufacturers adopt a “fall back minimal risk condition” for the vehicle that “should encompass designing the intended functions such that the vehicle will be placed in a safe state even when there are electrical, electronic, or mechanical malfunctions or software errors.”<sup>315</sup> If a malfunctioning vehicle is not placed in a safe state and crashes as a result, the malfunction would subject the manufacturer to strict tort liability, giving all manufacturers a financial incentive to adopt an effective “fall back minimal risk condition” as NHTSA urges.<sup>316</sup>

---

Federal Motor Vehicle Safety Standards), <https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/812069> [<https://perma.cc/HG3X-XG5E>].

311. See Danielle Ivory & Hiroko Tabuchi, *Takata Says No to Fund for Victims of Defective Airbag*, N.Y. TIMES (July 9, 2015), <https://www.nytimes.com/2015/07/10/business/takata-says-no-to-fund-for-victims.html>? [<https://perma.cc/GK22-D77E>].

312. Occupant Crash Protection, 58 Fed. Reg. 46,551, 46,563 (Sept. 2, 1993) (codified at 49 C.F.R. §§ 571, 585 (2017)).

313. See Jonathan Soble, *Takata Expects Return to Profit Despite Facing Airbag Lawsuits*, N.Y. TIMES (May 8, 2015), <https://www.nytimes.com/2015/05/09/business/international/takata-says-airbag-defect-fallout-behind-it.html> [<https://perma.cc/XRU8-V6LM>].

314. RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 4(a) (AM. LAW INST. 1998) (“[A] product’s noncompliance with an applicable product safety statute or administrative regulation renders the product defective with respect to the risks sought to be reduced by the statute or regulation.”); *id.* § 4 reph. n. cmt. d (observing that this rule “finds its origin in a common-law rule holding that the unexcused omission of a statutory safety requirement is negligence per se”).

315. NHTSA, 2016 AUTOMATED VEHICLES POLICY, *supra* note 4, at 14, 20.

316. See *supra* Part II.A (explaining why malfunctions attributable to computer bugs are subject to strict products liability); see also *supra* Part III.A (explaining why malfunctions attributable to third-party hacking are subject to strict products liability).

A similar incentive would be created by the rule of strict manufacturer liability for malfunctions caused by unauthorized third-party hacking.<sup>317</sup> Under the regulatory approach proposed by NHTSA, “[m]anufacturers and other entities should follow a robust product development process based on a systems-engineering approach to minimize risks to safety, including those due to cybersecurity threats and vulnerabilities.”<sup>318</sup> This regulatory objective could be attained without NHTSA adopting regulations that specify detailed performance standards, as some have suggested.<sup>319</sup> Instead, the prospect of incurring strict tort liability for product malfunctions will give each manufacturer a financial incentive to figure out the most cost-effective way for minimizing these risks of malfunction, obviating the need for detailed regulations.

Federal law, in turn, can supplement state tort law by helping to clarify issues otherwise posed by a tort claim. A good example involves the manufacturer’s tort obligation to design the vehicle so that it equally treats the interests of consumers (the owner and users of the autonomous vehicle) and bystanders (those in other vehicles and so on).<sup>320</sup> According to NHTSA, “[a]lgorithms for resolving these conflict situations should be developed transparently using input from Federal and State regulators, drivers, passengers and vulnerable road users, and taking into account the consequences of an [autonomous vehicle’s] actions on others.”<sup>321</sup> A federal regulation requiring such transparency would considerably simplify a tort inquiry asking whether the design of an autonomous vehicle unreasonably risks danger to bystanders.

Under this approach, the manufacturers of autonomous vehicles would be subject to uniform regulation across the national market with respect to the most important aspects of safe product performance. The sole regulatory authority on these matters would reside with NHTSA, but the federal regime would nevertheless retain state tort law by deriving the mandated performance standards from the associated tort obligations and otherwise relying on tort law as a necessary supplement. Federal and state law can work together to comprehensively regulate automated driving technologies.

To be sure, NHTSA might ultimately leave these important regulatory issues for the states to determine. In its recently released 2017 policy statement that outlines a “path forward for the safe deployment of automated vehicles,” NHTSA “offers a nonregulatory approach to automated vehicle technology

---

317. See *supra* Part III.A (explaining why the operation of the vehicle by an unauthorized third-party hacker would be a product malfunction subject to strict liability).

318. NHTSA, 2016 AUTOMATED VEHICLES POLICY, *supra* note 4, at 21.

319. See LIPSON & KURMAN, *supra* note 30, at 105 (“Driverless cars will need a redundant real-time operating system that contains built-in independent, self-testing systems that are required by law.”); *id.* (“Similar to those airplanes, the car’s wiring and on-board computers should be physically walled off, safe from the tinkering hands of innocent passengers or malevolent hijackers.”).

320. See *supra* Part II.C.

321. NHTSA, 2016 AUTOMATED VEHICLES POLICY, *supra* note 4, at 26–27.



safety” based on voluntary industry guidelines and recommended best practices for state legislatures.<sup>322</sup> NHTSA continues to envision that the states will retain their traditional responsibilities to “regulat[e] motor vehicle insurance and liability.”<sup>323</sup> Absent the requisite federal regulations, however, manufacturers will soon face the considerable uncertainties that inhere in the state tort regimes of products liability. A combination of federal and state law is the best method for enabling manufacturers to confidently assess their potential liabilities for autonomous vehicles.

#### CONCLUSION

In addition to its other impacts, the emerging technology of autonomous vehicles will disrupt the practice of tort law. The majority of tort cases in the state courts now involve automobile accidents allegedly caused by a driver’s negligence.<sup>324</sup> By eliminating the human driver, autonomous vehicles will eliminate these tort claims. The manufacturer will instead be responsible for the driving performance of the autonomous vehicle, potentially making it liable for a crash. Autonomous vehicles will alter the mix and number of tort cases, causing a massive shift from ordinary negligence claims to those based on products liability.

This dynamic will inevitably put pressure on the doctrines of products liability. Disagreement about the potential scope of manufacturer liability for the crash of an autonomous vehicle is compounded by the potential for variations among the different state tort systems across the country. Unable to assess their potential liabilities and other tort obligations within the national market, manufacturers face an overly uncertain legal environment, which generates costs that could impede the emergence of this life-saving technology.

In an effort to address this problem, NHTSA has announced a strategy of promulgating nationally uniform safety regulations that function alongside of the state tort and insurance systems.<sup>325</sup> The strategy is not yet fully specified because it is still in development, but it already presents an obvious set of problems. What is the role of state tort law? What ensures that it will be

---

322. NAT’L HIGHWAY TRAFFIC SAFETY ADMIN., AUTOMATED DRIVING SYSTEMS 2.0: A VISION FOR SAFETY, at i–ii (2017), [https://www.nhtsa.gov/sites/nhtsa.dot.gov/files/documents/13069a-ads2.0\\_090617\\_v9a\\_tag.pdf](https://www.nhtsa.gov/sites/nhtsa.dot.gov/files/documents/13069a-ads2.0_090617_v9a_tag.pdf) [<https://perma.cc/5BC5-SDS6>].

323. *Id.* at 20.

324. According to court data collected from seventeen states, auto tort cases in 2015 ranged from a low of 32 percent of all tort cases (Missouri) to 75 percent (Texas), with auto tort cases exceeding 50 percent of all tort cases in the majority of these states. *See* 2015 Statewide Auto Tort Cases per 100,000 Population, NAT’L CTR. FOR STATE COURTS, COURT STATISTICS PROJECT, <http://www.courtstatistics.org/Civil.aspx> (follow “Launch DataViewer” hyperlink; then select “Statewide Auto Tort Rates (Chart)”). “Auto tort cases arise from auto accidents and typically allege personal injury, property damage, or wrongful death resulting from negligent operation of a motor vehicle.” *Id.*

325. *See supra* notes 246–51 and accompanying text.

adequately coordinated with federal regulatory law? What is the role of insurance? Rather than clearly resolving matters, NHTSA's proposed regulatory approach at this point seems only to deepen the uncertainty.

These problems, perhaps ironically, can be largely solved by the technology itself for reasons that become apparent once we consider why the technology can cause an autonomous vehicle to crash.

As with conventional vehicles, hardware malfunctions can cause an autonomous vehicle to crash. As with conventional vehicles, these defects will subject the manufacturers of autonomous vehicles to strict liability, giving them a financial incentive to adopt reasonably safe methods of quality control. All crashes caused by defective hardware in the vehicle clearly fit within the existing liability regime.<sup>326</sup>

The remaining cases involve crashes caused by the software components of the operating system. Just as hardware malfunctions, an autonomous vehicle's software can malfunction. Programming errors or bugs can cause the operating system to freeze, which in turn can cause the vehicle to crash. As with hardware malfunctions, these defects will subject the manufacturers of autonomous vehicles to strict liability, giving them a financial incentive to subject the vehicle's programming to reasonably safe methods of quality control.<sup>327</sup>

The most pressing problem—and potential source of substantial legal uncertainty—involves manufacturer liability for the crash of a fully functioning autonomous vehicle. When the fully functioning operating system causes a crash, the vehicle was engaged in systemized driving performance that can be evaluated with aggregate driving data for the fleet. Under widely adopted rules of products liability, the programming or design of the fully functioning operating system would necessarily satisfy the tort obligation if the data show that the autonomous vehicle performs at least twice as safely as conventional vehicles, eliminating defective design as a potential source of manufacturer liability in these cases.<sup>328</sup>

To avoid liability for crashes proximately caused by a fully functioning autonomous vehicle, the manufacturer must also adequately warn consumers about this inherent product risk. Once again, the crash involves systemized driving performance, and so aggregate driving data provide a determinate method for the manufacturer to satisfy the tort obligation. These data will enable insurers to establish the premium for insuring the vehicle based on the underlying systemic risk of crash. By disclosing this risk-adjusted premium to consumers, the manufacturer would satisfy its obligation to warn about the

---

326. See *supra* notes 70–71 and accompanying text.

327. See *supra* Parts II.A, III.A.

328. See *supra* Part II.D.1.

inherent risk of crash, eliminating this final source of manufacturer liability for crashes caused by a fully functioning autonomous vehicle.<sup>329</sup>

These liabilities are based on tort rules that have been widely adopted across the country. States that do not follow the majority approach might reach different conclusions. To ensure that manufacturers face uniform obligations across the national market, NHTSA could adopt two regulations that clearly fit within its proposed regulatory approach, with each one respectively derived from the manufacturer's associated tort obligations to adequately test the vehicle and warn about the inherent risk of crash.

By complying with both of these federal regulations, the manufacturer would avoid tort liability for the crash of a fully functioning autonomous vehicle in most states under the regulatory compliance defense, and in the remaining states under the doctrine of implied preemption.<sup>330</sup> The regulations would attain national uniformity while retaining tort law in the vast majority of states, an optimal solution to the federalism problem.

Although the manufacturer would not be liable in these cases, those who are injured by the crash—the occupants of the vehicle, its owner, and third parties—might still be able to recover from others involved in the crash. The legal regime is not a system of no-fault liability that eliminates tort liability for motor vehicle crashes.

For example, an inattentive human driver or an oncoming defective autonomous vehicle could have caused the crash, in which cases the injured parties could seek recovery from the negligent driver or the responsible manufacturer. Insurance can also cover these injuries and associated liabilities, with the states retaining their traditional role of establishing the required amounts of minimum insurance coverage for vehicle owners and commercial operators.

For the foreseeable future, autonomous vehicles will share the road with human drivers. After autonomous vehicles are first commercially distributed, there will be a transitional period when human drivers sometimes have difficulty anticipating the driving behavior of autonomous vehicles, resulting in crashes.<sup>331</sup> A regulatory-compliant, fully functioning autonomous vehicle foreseeably creates the risk of these crashes. Consequently, the autonomous vehicle would be a proximate cause of such a crash. This inherent risk of crash, however, factors into the relative safety performance of the autonomous vehicle and is accordingly addressed by the proposed regulatory requirements of adequate premarket testing and disclosure of the risk-adjusted premium for insuring against the inherent risk of crash. If the vehicle has otherwise been designed in a manner that equally accounts for the interests of other motorists

---

329. *See id.*

330. *See supra* Part IV.C.1.

331. *See generally* Surden & Williams, *supra* note 28 (discussing problem of humans predicting the behavior of autonomous vehicles).

(bystanders), then a regulatory-compliant, fully functioning autonomous vehicle that proximately causes a crash with a conventional vehicle would not subject the manufacturer to tort liability.<sup>332</sup>

Within the realm of crashes only involving autonomous vehicles of this type, there will be no tort liability. The absence of manufacturer liability will channel injury costs from the relatively inefficient tort system, where they are now largely located, into the more cost-effective insurance system. The same outcome is ideally attained by no-fault automobile insurance.<sup>333</sup>

The only remaining significant potential for manufacturer liability involves cybersecurity. An autonomous vehicle that crashes because a third-party hacked into the operating system did not function as expected or intended. The malfunction would subject the manufacturer to strict liability, although the matter is not clear at this point.<sup>334</sup> The complexity of cybersecurity often makes it difficult to establish negligence liability, requiring strict liability to effectively enforce the requirement of reasonable safety. However, the extent of liability could be vast, particularly for hacks motivated by terrorism. Based on the concern that strict liability would unduly threaten the stability of the market, courts could limit manufacturer liability to negligence claims despite the attendant difficulties of proof. That outcome is hardly certain, so for now manufacturers should expect to incur strict liability for crashes caused by an operating system that malfunctions because of hacking.

Within this legal framework, a regulatory-compliant autonomous vehicle would subject the manufacturer to tort liability only for crashes caused by malfunctioning physical hardware (strict products liability); malfunctions of the operating system caused by either programming error (same) or third-party hacking (strict liability again, with an important caveat); the manufacturer's failure to adopt a reasonably safe design or to provide adequate warnings for ensuring safe deployment of the vehicle (an ordinary products liability claim); or the manufacturer's failure to treat consumers and bystanders equally when designing the vehicle (an ordinary negligence claim). A manufacturer would also be subject to tort liability for not complying with the federal regulations (negligence per se). In addition to giving the manufacturer the requisite financial incentive to ensure that the driving performance of the autonomous vehicle is reasonably safe and well understood by consumers, these liabilities are not overly uncertain. The road ahead is clear.

---

332. See *supra* Part II.C.

333. In a tort suit, both parties typically need legal representation, and so by eliminating tort liability, no-fault auto insurance can reduce total costs for consumers. See generally Nora Freeman Engstrom, *An Alternative Explanation for No-Fault's "Demise,"* 61 DEPAUL L. REV. 303 (2012) (explaining the rationales for no-fault auto insurance and identifying the reasons why the plans adopted by states were designed in a manner that fatally undermined these objectives); see also *supra* notes 202–05 and accompanying text (explaining why tort compensation is more expensive for consumers than other forms of insurance).

334. See *supra* Part III.