

Tesla's Autopilot Could Save the Lives of Millions, But It Will Kill Some People First

The complicated ethics of Elon Musk's grand autonomous vehicle experiment.

October 9, 2019, 6:00 PM GMT+9

On the last day of his life, Jeremy Banner woke before dawn for his morning commute. He climbed into his red Tesla Model 3 and headed south along the fringes of the Florida Everglades. Swamps and cropland whizzed past in a green blur.

Banner tapped a lever on the steering column, and a soft chime sounded. He'd activated the most complex and controversial auto-safety feature on the market: Tesla Autopilot. It's a computer system that performs all the functions of normal highway driving without any input from the driver. When the computer is in control, the car can speed up, change lanes, take exits, and—if it spots an obstacle ahead—hit the brakes.

Bloomberg Businessweek

October 14, 2019 • SPECIAL DOUBLE ISSUE

Tesla's Autopilot could save millions of lives. How many people will it kill first?



**The complicated
ethics of Elon Musk's
grand experiment**

Featured in *Bloomberg Businessweek*, Oct. 14, 2019. [Subscribe now.](#)

Photo illustration: 731: Photograph: AP Photo

[Tesla Inc.](#) aims to dominate the global auto market by building the world's first self-driving car, and it considers Autopilot to be the crucial first step. Customers adore it. They've logged more than 1.5 billion miles on Autopilot, often pushing the limits of the software. Although the owner's manual warns drivers to closely supervise the car at all times, that hasn't stopped some from reading books, [napping](#), [strumming a ukulele](#), or having sex. Most of the time, the car gets them where they're going.

But on that morning in March, Banner's sedan failed to spot a tractor-trailer crossing the four-lane highway ahead of him. So did Banner, whose attention had apparently strayed. He [struck the trailer broadside](#) at 68 mph, the top of his car shearing off like a sardine can. The 50-year-old [father of three](#) died instantly.

Computer mistakes don't look like human mistakes. Autopilot has lightning reflexes and its attention never flags, but it sometimes fails to spot hazards in its path. Such oversights appear to have played a role in four of five known fatalities since Autopilot was introduced in 2015. Banner's wreck, in fact, bore an uncanny resemblance to an [earlier one](#). In August, [Banner's estate sued Tesla](#) under Florida's Wrongful Death Act. The estate's argument is a straightforward [product-liability claim](#): Tesla promised a safe car and delivered a dangerously defective one.

But Autopilot is unlike almost any other consumer product in history, in ways that offer a preview of the uncomfortable questions we'll confront in the dawning robot age. Tesla's flamboyant chief executive officer, Elon Musk, says the technology saves lives, and legions of Tesla owners offer their own testimonies of hazards spotted and collisions avoided. (And they have [YouTube videos](#) to prove it.) It's possible that both sides are right, that the computers are killing a few drivers who otherwise would have lived, but that they're also saving the lives of many more. In the coming years, society—in particular, regulators and the courts—will have to decide whether that's an acceptable trade-off.

The question is no longer academic. Musk's decision to put Autopilot in the hands of as many people as possible amounts to an enormous experiment, playing out on freeways all over the world.

I was in the passenger seat, heading north on Interstate 405 in Los Angeles, when Omar Qazi took both of his hands off his steering wheel. We were going about 50 mph on the most [heavily traveled highway](#) in the country, and the wheel of Qazi's black Model 3 turned slightly to the left, keeping the car centered in the gently curving lane. "This is like L.A. rush-hour traffic, right?" said Qazi, a 26-year-old software engineer. "It's, like, flawless."

Tesla has legions of die-hard fans, many of them well-to-do, tech-obsessed, and male. Qazi is pretty close to the archetype. His Twitter handle, [@tesla_truth](#), is a bottomless font of Muskolatry. Before we met in August, he'd emailed Musk to give him a heads-up and encourage him to speak with me. The billionaire CEO, who declined to be interviewed for this story, replied to his fan the same day. "Your Twitter is awesome!" he said, before adding a warning: "Please be wary of journalists. They will sweet talk you and then wack you with a baseball bat." Musk cc'd me on the message. Tesla also declined to comment.

Qazi met me at the charging station outside Tesla's L.A.-area offices, with one of Musk's SpaceX booster rockets looming nearby like an industrial obelisk. Qazi wore a day's worth of stubble and blue Nike Airs. He immediately showed me the experimental Smart Summon feature, at the time available only to a select group of Tesla beta testers. (Qazi got it after begging Musk on Twitter; the feature rolled out to regular customers in September.) He pressed a button on his phone, and his car pulled out of its spot. Qazi watched it cross the parking lot and roll toward him. "It's not useful—yet," he said, grinning. But he loves showing off this trick so much he's been known to linger in a parking lot, waiting for an audience.

Cumulative Tesla Vehicle Deliveries

Data: "Tesla Vehicle Deliveries and Autopilot Mileage Statistics," Lex Fridman

Smart Summon offers a tiny glimpse of the driverless future Musk is promising, but for road driving, Autopilot is as close as it currently gets. Tesla says the technology isn't reliable enough yet for humans to turn their attention away, even for a second, so it requires them to keep their hands on the wheel. Because most U.S. states are still figuring out how they'll handle driverless cars, this also serves a legal purpose. To state regulators, Autopilot is just an advanced driver-assistance program—a souped-up cruise control, basically. Autopilot can't yet tackle off-highway features such as traffic lights and stop signs. But during its four years on the road, it has gradually shouldered more complex tasks: merging smoothly, avoiding cars that cut in, and [navigating from one highway to another](#).

"It can't drive itself perfectly, but the rate of advancement of the software is like—every couple of weeks you get an update, and the car's driving a little more humanlike. It's very eerie," Qazi said. A few minutes later, a silver sedan cut into our lane, and the car smoothly braked to let it in. "See that?" he asked.

It's not as if human drivers set the bar very high. In Los Angeles, on the day I met Qazi, an [illegal drag racer died](#) when his Mazda hit a parked truck; a motorcyclist [fatally struck](#) a broken-down van in a carpool lane; and a high school junior on a bicycle was critically injured after a [car dragged him](#) 1,500 feet and then sped away.

In fact, driving is one of the most dangerous things most adults do. [It killed](#) 40,000 Americans last year and 1.4 million people [globally](#). And yet we're all pretty complacent about it. In 1974, in the name of fuel savings, the U.S. capped highway speed limits at 55 mph. [One study found the change](#) cut highway deaths by at least 3,000 in its first year. But people like driving fast, and Congress later removed the cap. A few years ago, traffic deaths began inching up, a development experts attribute to distractions from smartphones. Still we drive, and text, and drive.

Whatever their flaws, computers don't get drunk, or tired, or angry, or feel an irrepressible urge to check Instagram while driving on interstates. Autonomy promises to preserve our car-centered lifestyle but eliminate the estimated 94% of crashes caused by human error. Viewed from that angle, the self-driving car could be a lifesaver in the same class as penicillin and the smallpox vaccine.

Qazi has done the math: He says autonomous cars will someday save 3,000 lives a day. By his logic, anyone standing in the way of that progress has blood on their hands. "Imagine someone delaying the software by one day," he says. "You are really going to end up killing a lot of people."

Less than two months after Banner's fatal crash, Musk invited about 100 investors and analysts to Tesla's headquarters in Palo Alto, greeting them in a cavernous meeting hall. Born and raised in South Africa, he made a fortune in Silicon Valley and then undertook a series of audacious projects: commercial rockets, high-speed tunnels, brain implants, electric cars. His many admirers consider him a world-changing visionary; his foes, a bloviating phony. On that April morning, Musk occasionally interrupted the Tesla scientists who shared the stage with him and [mused freely](#) about, among other things, whether life might be a computer simulation.

Tesla's stock had been sinking for months. Despite having delivered the [world's bestselling electric car](#), the Model 3, the company was still far from profitable, and Musk would soon be forced to raise more cash from investors. Over the course of the 2 1/2-hour presentation, Musk pointed investors toward a new focus: building the first truly driverless car. Cars on the road today, he said, would be able to use Autopilot on local roads within months. By sometime in 2020 they'd no longer need human oversight and could begin earning money as drone taxis in their downtime.

"It's financially insane to buy anything other than a Tesla," Musk said, throwing up his hands. "It will be like owning a horse in three years."

Musk's timetable sounded particularly bold to anyone following the self-driving car business. Some three dozen companies, including General Motors, Daimler, and Uber, are racing to develop the technology. Many observers consider the strongest contender to be Waymo LLC, the Google spinoff that's been working on the problem for more than a decade. None of them is anywhere near selling a driverless car to the public.

Tesla will overtake them all, Musk told the assembled investors, thanks to the more than 500,000 Autopilot-enabled Teslas already on the road. Although he didn't use these words, Musk described Autopilot as a kind of rough draft, one that would gradually grow more versatile and reliable until true autonomy was achieved.

Releasing still-incomplete software to customers now, and hoping to work out bugs and add capabilities along the way, is, of course, how Silicon Valley often introduces smartphone apps and video games. But those products can't kill people. Waymo, GM, and the others have rough drafts, too, but they're installed in only a few hundred test models, deployed in a handful of carefully chosen neighborhoods around the country, and almost always supervised by professional safety drivers. Safety is an obsession, especially after an Uber test car mowed down a pedestrian last year. GM's prototypes crawl San Francisco's hilly streets at a maximum speed of 35 mph.

Miles Driven

Data: "Tesla Vehicle Deliveries and Autopilot Mileage Statistics," Lex Fridman

Musk, on the other hand, is putting his rough draft into consumers' hands as fast as he can. This allows Tesla engineers to collect terabytes of data from customers and use the information to refine the Autopilot software based on real-world conditions. Even Teslas that aren't on Autopilot pitch in: They silently compare the human driver's choices with what the computer would have done. Every few weeks, Tesla completes a new and improved version of Autopilot and uploads it to the cars, to the delight of Qazi and other fans.

"Everyone's training the network, all of the time," Musk said in Palo Alto. He called this virtuous cycle "fleet learning," comparing it to the way Google's search engine improves with each of the 1.2 trillion queries a year it fields. Someday soon, he declared, the software will be so good, drivers will start unbolting the steering wheels from their cars and throwing them away.

When a Morgan Stanley analyst pressed Musk about Autopilot's safety record, he quickly changed the subject to the dangers of human driving and the potential for technology to fix it. He compared cars to old-fashioned elevators controlled by human operators. "Periodically, they would get tired, or drunk or something, and then they'd turn the lever at the wrong time and sever somebody in half," he said. "So now you do not have elevator operators."

Considering the life-and-death stakes, it isn't surprising that Musk sometimes talks about driverless cars as a kind of righteous crusade. He once said it would be "[morally reprehensible](#)" to keep Autopilot off the market. But he and his acolytes aren't the only ones to talk this way. The first U.S. driver to die on Autopilot was Joshua Brown, a Navy veteran from Ohio who, like Banner, also rammed into a crossing semi. After his crash in 2016, his family issued a statement that basically endorsed Tesla's moral calculus. "Change always comes with risks," they wrote. "Our family takes solace and pride in the fact that our son is making such a positive impact on future highway safety." Brown had become, in effect, a martyr to Musk's cause.

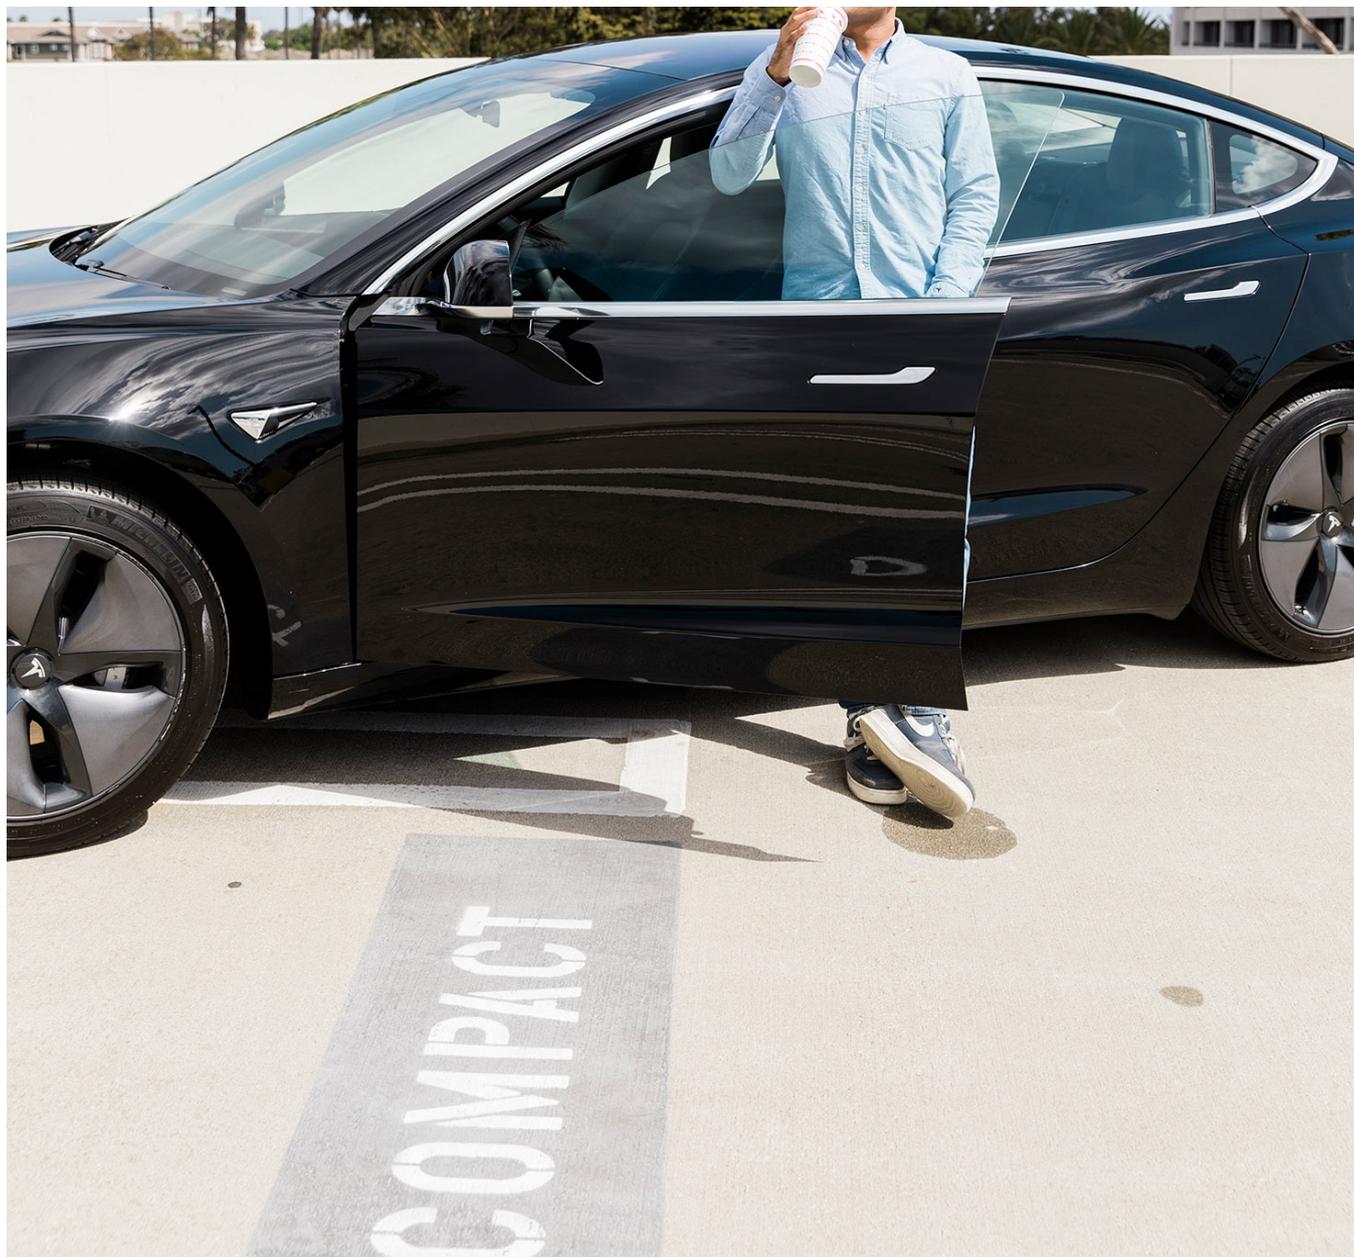
Until drivers go the way of elevator attendants, Musk says, Autopilot is the next best thing: all the safety of a human driver, plus an added layer of computer assistance. But automation can cut both ways. When we cede most—but not all—responsibility to a computer, our minds wander. We lose track of what the computer is supposed to be doing. Our skills get rusty. The annals of aviation are full of screw-ups caused by humans' overreliance on lowercase "a" autopilot. Two Northwest Airlines pilots once zoned out so completely they overshot Minneapolis by 100 miles.

"It's just human nature that your attention is going to drift," says Missy Cummings, a former Navy fighter pilot and a professor at Duke University's Pratt School of Engineering who wants Autopilot taken off the market. Waymo, the Google spinoff, developed an Autopilot-like system but abandoned it six years ago. Too many drivers, it said, were texting, applying makeup, and falling asleep.

Computers, meanwhile, can mess up when a driver least expects it, because some of the tasks they find most challenging are a piece of cake for a human. Any sentient adult can tell the difference between a benign road feature (highway overpass, overhead sign, car stopped on the shoulder) and a dangerous threat (a tractor-trailer blocking the travel lane). This is surprisingly hard for some of the world's most sophisticated machine-vision software.

Tesla has resisted placing limits on Autopilot that would make it safer but less convenient. The company allows motorists to set Autopilot's cruising speed above local speed limits, and it lets them turn on Autopilot anywhere the car detects lane markings, even though the manual says its use should be restricted to limited-access highways.





Photographer: Liz Kuball for Bloomberg Businessweek

To those who'd test the car's limits, Musk himself offers winking encouragement. When he showed off a Model 3 to Lesley Stahl on *60 Minutes* in December, he did precisely what the manual warns against, turning on Autopilot and [taking his hands off the wheel](#). Then in May, after an Autopilot porn video went viral, Musk responded with a [joking tweet](#): "Turns out there's more ways to use Autopilot than we imagined." Qazi says he and an ex-girlfriend used to make out while it was on. He did his best to keep one eye on the road.

Given that Autopilot now has [more than 1.5 billion miles](#) under its belt,

determining its safety record ought to be easy. Musk has claimed driving with Autopilot is about twice as safe as without it, but so far he hasn't published data to prove that assertion, nor has he provided it to third-party researchers. Tesla discloses [quarterly Autopilot crash-rate figures](#), but without more context about the conditions in which those accidents occurred, safety experts say they're useless. [An insurance-industry study](#) of Tesla accident claims data was mostly inconclusive.

After Brown's 2016 crash, the National Highway Traffic Safety Administration investigated Autopilot and found no grounds for a recall. It based its conclusion, in part, on a finding that Teslas with Autopilot installed were crashing 40% less than those without. But that was based on a series of dubious calculations. While Tesla had handed over mileage and collision data on 44,000 cars, key data was missing or contradictory for all but 5,700 of them. Within that modest group, the crash rate with Autopilot was actually higher. The faults came to light only when Randy Whitfield, an independent statistics consultant in Maryland, pointed them out this year. The NHTSA has said it stands by the finding.

Part of the problem with assessing Autopilot, or fully autonomous technology for that matter, is that it isn't clear what level of safety society will tolerate. Should robots be flawless before they're allowed on the road, or simply better than the average human driver? "Humans have shown nearly zero tolerance for injury or death caused by flaws in a machine," said [Gill Pratt](#), who heads autonomous research for Toyota Motor Corp., in a 2017 speech. "It will take many years of machine learning, and many more miles than anyone has logged of both simulated and real-world testing, to achieve the perfection required."

But such a high standard could paradoxically lead to more deaths than a lower one. In a [2017 study for Rand Corp.](#), researchers Nidhi Kalra and David Groves assessed 500 different what-if scenarios for the development of the technology. In most, the cost of waiting for almost-perfect driverless

cars, compared with accepting ones that are only slightly safer than humans, was measured in tens of thousands of lives. "People who are waiting for this to be nearly perfect should appreciate that that's not without costs," says Kalra, a robotics expert who's [testified before Congress](#) on driverless-car policy.

Key to her argument is an insight about how cars learn. We're accustomed to thinking of code as a series of instructions written by a human programmer. That's how most computers work, but not the ones that Tesla and other driverless-car developers are using. Recognizing a bicycle and then anticipating which way it's going to go is just too complicated to boil down to a series of instructions. Instead, programmers use machine learning to train their software. They might show it thousands of photographs of different bikes, from various angles and in many contexts. They might also show it some motorcycles or unicycles, so it learns the difference. Over time, the machine works out its own rules for interpreting what it sees.

The more experiences they have, the smarter these machines get. That's part of the problem, Kalra argues, with keeping autonomous cars in a lab until they're perfect. If we really wanted to maximize total lives saved, she says, we might even put autonomous cars on the road while they're still more dangerous than humans, to speed up their education.



Videos and photos of Tesla drivers who appear to be napping are a fixture of social media.

Even if we build a perfect driverless car, how will we know it? The only way to be certain would be to put it on the road. But since fatal accidents are [statistically rare](#)—in the U.S., about one for every 86 million miles traveled—the amount of necessary testing would be mind-boggling. In another Rand paper, [Kalra estimates](#) an autonomous car would have to travel 275 million failure-free miles to prove itself no more deadly than a human driver, a distance that would take 100 test cars more than 12 years of nonstop driving to cover.

Considering all that, Musk's plan to simultaneously refine and test his rough draft, using regular customers on real roads as volunteer test pilots, doesn't sound so crazy. In fact, there may be no way to achieve the safety gains of autonomy without exposing large numbers of motorists to the risk of death by robot. His decision to allow Autopilot to speed and to let it work on

unapproved roads has a kind of logic, too. Every time a driver wrests control from the computer to avoid an accident, it's a potential teachable moment—a chance for the software to learn what not to do. It's a calculated risk, and it's one that federal regulators, used to monitoring for mechanical defects, may be ill-prepared to assess.

The U.S. already has a model for testing potentially lifesaving products that might also have deadly side effects: phased clinical drug trials. Alex London, a philosophy professor at Carnegie Mellon University, is among those [calling for auto regulators](#) to try something similar, allowing new technology onto the road in stages while closely monitoring its safety record. "Even if my proposal is not the best proposal, I can tell you what the worst proposal is," he says. "The worst proposal is to take the word of the person who designed the system, especially when they are trying to sell it to you."

On my last ride with Qazi, we drove to Rancho Palos Verdes, snaking through rolling brown hills and bluffs overlooking the Pacific. We weren't on a limited-access highway, but Qazi, defying the manual again, turned on Autopilot. The car did great, mostly.

As the road skirted a steep cliff, we approached a cyclist headed the same way. The Tesla correctly identified him as a biker and moved to overtake him. Just before it pulled alongside, Qazi braked, allowing the man to advance to a wider part of the road before we passed. He said he hoped the computer would have done the same, but he wasn't willing to find out.

But Qazi seemed resigned to the statistical certainty that as Teslas proliferate on the world's roads, there will be more Autopilot fatalities. "The biggest PR nightmares are ahead," he told me before we parted. "There's only one way to the goal. Through the minefield." —*With Dana Hull and Ryan Beene*