

A Waymo autonomous taxi on Bush Street in San Francisco in December 2025. *Photographer: David Paul Morris/
Bloomberg*

CityLab | Perspective

We Still Don't Know if Robotaxis Are Safer Than Human Drivers

And even if self-driving technology proves to be less dangerous, there are many better ways to improve traffic safety and prevent fatal crashes.



By [David Zipper](#)

January 6, 2026 at 6:00 AM PST

Corrected January 20, 2026 at 10:46 AM PST



Save



Translate

If a chorus of wide-eyed boosters and enthralled journalists are to be believed, self-driving cars from companies like Waymo, Tesla, and Zoox can bring about road safety nirvana – if only US regulators would get out of their way.

Waymo has said that the for-hire autonomous vehicles it operates in several cities are “already making roads safer,” an assertion echoed by many media outlets. Since “robotaxis have fewer accidents than human drivers,” the *Economist* concluded, “they are almost certainly saving lives.” By implication, regulations that hinder AV deployments are effectively killing people. A neurosurgeon made a similar argument in a recent *New York Times* op-ed, writing that “there is a public health imperative” to expand robotaxis as quickly as possible.

A deus ex machina solution for crashes is a tantalizing prospect in the US, where residents are several times more likely to die in a collision than those in other rich countries.

But don’t pop the champagne just yet. In fact, don’t even take the bottle out of the fridge.

I’ve been writing and thinking about car safety since before Waymo was offering robotaxi service at all. In researching this piece, I spoke to many independent transportation and technology researchers about what we know and don’t know about the industry’s safety record thus far. Contrary to prevailing narratives, these experts are unconvinced that today’s self-driving cars are less crash-prone than those operated by humans. And even if robotaxis do end up being safer on individual trips, that doesn’t necessarily mean that America’s roadway body count will decline.

It’s time for a reality check about the safety of self-driving cars.

Human Error

Technologists have long imagined that vehicles outfitted with computers and sensors could reduce or even eliminate the traffic crashes that kill more than 40,000 people annually in the US. It’s been 16 years since Sebastian Thrun, the “godfather of the self-driving car industry,” floated that possibility in a talk at the University of Washington.

There are valid reasons to believe that computer-powered cars can be safer than those with a person behind the wheel. Unlike humans, AVs will not drive drunk, high or

fatigued. They will not become distracted while fiddling with their phones, trying to figure out the car's touchscreen controls or quieting a recalcitrant child. They won't become consumed with road rage and will reliably obey traffic laws (unless, of course, they are programmed otherwise).

On the other hand, self-driven cars make mistakes that humans would not, such as plowing into flood water or driving through an active crime scene where police have their guns drawn. During a major power outage in San Francisco on Dec. 20, many Waymos froze in place, jamming major intersections across the city when traffic lights went dark.

“In like 95% of situations where a disengagement or accident happens with autonomous vehicles, it's a very regular, routine situation for humans,” said Henry Liu, a professor of engineering at the University of Michigan who leads Mcity, the university's center for transportation technology and innovation. “These are not challenging situations whatsoever.”

One particularly gruesome self-driven mishap involved Cruise, a now-shuttered AV company that was a subsidiary of General Motors. On Oct. 2, 2023, a pedestrian in San Francisco was struck by a driver and landed beneath a Cruise robotaxi. Rather than halt – as a human driver likely would have – the Cruise vehicle rolled 20 feet to the curb, dragging the pedestrian pinned below it. (A firestorm followed, including a roughly \$10 million settlement paid to the surviving crash victim and a fine levied by California regulators who concluded that Cruise had been less than truthful about what had happened. GM closed Cruise in December 2024.)

Cruise self-driving vehicles in San Francisco in 2022. *Photographer: David Paul Morris/Bloomberg*

Waymo, the AV firm spun off from Google parent Alphabet Inc.'s pioneering self-driving unit, has not been linked to a crash nearly as calamitous as that one. (The company has been involved in collisions, including an incident that led to a lawsuit filed by a San Francisco cyclist who claimed they were doored by a Waymo blocking a bike lane and then struck by another of the company's vehicles.) Waymo has publicly emphasized its safety record, publishing data and studies to support the idea that its robotaxis are making streets less dangerous.

“Waymo is already improving road safety in the cities where we operate, achieving more than a ten-fold reduction in serious injury or worse crashes,” said Trent Victor, Waymo's director of safety research and best practices. He also pointed to the dozens of peer-reviewed articles coauthored by company employees.

Tesla and Zoox, Waymo's most prominent robotaxi rivals in the US, have had only limited deployments in far fewer cities. (Neither company responded to requests for comment.) As a result, Waymo's record dominates the AV safety discourse.

“The reason we're only talking about Waymo is because they have enough deployment to study and because they voluntarily release data like vehicle miles traveled that enable such studies,” said Joseph Young, spokesperson for the Insurance Institute of Highway Safety. “There isn't sufficient data available on Zoox' or Tesla's autonomous ridehailing services to draw any conclusions.”

Although Waymo has shared more data than its rivals, some researchers say they aren't ready to judge its safety performance. Liu told me that he “doesn't have enough information” to know whether the company's vehicles are currently safer than humans, particularly because the company maintains control over the information it chooses to share. “We have seen many reports from autonomous vehicle developers, and it looks like the numbers are very good and promising,” said Liu, who has extensively studied self-driving technology. “But I haven't seen any unbiased, transparent analysis on autonomous vehicle safety. We don't have the raw data,” he said. He added that he would like to see “designed experiments where we can control the situation,” taking into account variables like weather and road conditions.

Waymo disputed that view, noting the crash and locational datasets that the company collects and publishes online.

Other researchers question Waymo's use of human “teleoperators” – remote backup drivers (some based in call centers located as far away as the Philippines) who monitor and occasionally take the controls. Teleoperators provide assistance across the self-driving industry, but they are largely unregulated and the full extent of their

interventions isn't known. Missy Cummings, the director of George Mason University's Autonomy and Robotics Center, has noted that their input could affect safety estimates of driverless systems, and that they may have played a role in the vehicle stoppages that Waymo suffered during San Francisco's blackout last month. "It is not enough to claim your car is safe without including information on the nature of human babysitting," she wrote recently on LinkedIn.

In response, Waymo pointed out that its remote assistance program had been audited by TÜV SÜD, a German company specializing in regulatory and safety compliance. TÜV SÜD concluded that the company's use of teleoperators adhered to industry standards for best practices.

But the automotive journalist Junko Yoshida maintains that unanswered questions about remote operations remain – particularly after the San Francisco blackout. "No one seems to know," she wrote in a recent newsletter, "how many teleoperators are deployed per how many robotaxis or where 'call centers' for teleoperators are located. Nor is there information on the criteria for hiring 'teleoperators,' or the minimum latency for self-driving remote operations."

Several self-driving Waymo vehicles stopped at intersections during a power outage in San Francisco in December. *Photo by Tayfun Coskun/Anadolu via Getty Images*

Premature Victory Lap

Liu cited a further reason for caution. So far, Waymo vehicles have mostly driven on

Liu cited a further reason for caution. So far, Waymo vehicles have mostly driven on urban streets with a speed limit of 35 miles per hour or less – much slower than the highways and arterials that human drivers regularly traverse. “It’s not really fair to compare that with human driving,” Liu said, suggesting that the 127 million miles Waymo has logged doesn’t necessarily reflect conditions in which people typically drive: “It’s not apples to apples.”

He also hypothesized that robotaxi crashes could be more deadly than equivalent collisions involving human-driven vehicles because fewer robotaxi occupants will use seat belts, mistakenly assuming that they will be fine without one. (Previous studies have found that ridehail passengers are significantly less likely to wear seat belts than those inside personally driven vehicles.)

Phil Koopman, an emeritus engineering professor at Carnegie Mellon, pointed out another data complication. In California, around half of Waymo’s total mileage involves deadheading, with no passengers inside the vehicle. “I would expect robotaxis to have only a third as much harm to occupants as regular taxis,” said Koopman. “They spend half their time empty and half the time with one person instead of half the time with one person and half the time with two people.”

Furthermore, companies like Waymo frequently update their self-driving software, raising the possibility that their vehicles become significantly safer or more dangerous afterward. Liu questioned “whether it makes sense to compare the autonomous vehicle safety record from version 1.0 to 3.0 to the human safety record. I don’t really think that’s a fair comparison.”

Koopman agreed. “You have no idea if next Tuesday all of a sudden it’ll be a death machine because of bad software,” he said, suggesting that the companies should “reset the odometer” for incident data following major updates.

Last month, Waymo issued a voluntary recall to address a pattern of its vehicles driving past stopped school buses. In a post, Cummings suggested that a flawed software update could be responsible. (Problematic software updates have previously caused safety issues for automakers like Tesla.)

Waymo declined to address that possibility directly, but Victor said that “every major software update undergoes rigorous testing and readiness review,” including “a combination of real-world testing and advanced simulations.”

Even putting all these caveats aside, Waymo’s current safety record is not anything to celebrate, according to Matthew Raifman, a senior researcher at the University of California, Berkeley’s Safe Transportation Research and Education Center. With around 127 million self-driven miles under Waymo’s belt, the company has been involved in two fatal crashes (although not directly responsible for either). A simple calculation dividing

total miles by total fatalities – which Raifman called “the most accurate measure of safety that we have” – would give the company a higher per mile death rate than average American drivers, who drive around 80 million miles for every fatality.

But Waymo doesn't evaluate its data that way. Instead, it lumps fatalities together with serious injuries, a category of crashes where data is more likely to be missing – and which, unlike road deaths, is not included in the federal government's quarterly road safety updates.

For now, at least, Raifman and Koopman agreed that Waymo simply has not driven enough miles to draw clear conclusions about its safety record, because fatal crashes are relatively rare. In an influential 2016 study, the RAND Corporation concluded that “autonomous vehicles would have to be driven hundreds of millions of miles and sometimes hundreds of billions of miles to demonstrate their safety.” More recently, Koopman estimated that clear comparisons between humans and robotaxis could be drawn once Waymo had around 2 billion miles under its belt.

“The jury is still out on Waymo fatality rates, and will be for quite some time,” Koopman wrote in a recent blog post.

Waymo's Victor acknowledged that “there is not yet sufficient mileage to make statistical conclusions about fatal crashes alone,” adding that “as we accumulate more mileage, it will become possible to make statistically significant conclusions on other subsets of data, including fatal crashes as its own category.”

To summarize: We don't yet know whether a robotaxi trip is more or less likely to result in a crash than an equivalent one driven by a human. The answer likely depends on the trip and the autonomous vehicle company, and it might change in the future.

The Hidden Risk

Let's give robotaxi companies the benefit of the doubt and assume that independent researchers will eventually agree that self-driven journeys are significantly safer than human-driven ones. Many transportation and robotics experts do believe this will ultimately happen; Liu said he was “quite confident” about it.

Can we then assume that crash deaths will fall as AVs scale? No, we cannot.

The reason is that self-driving cars are poised to increase driving – and thus create more opportunities for collisions. As Koopman put it: “Civilized countries look at car deaths per capita instead of car deaths per mile.” Reducing deaths per capita presents a much higher hurdle for AVs to clear.

In San Francisco, many residents appear willing to pay the premium that Waymo charges over human-driven ridehail. That is unsurprising, given the comfort, privacy

and general pleasantness of self-driven trips. Since robotaxis offer a more appealing experience for many riders, allowing them to do other tasks and play their favorite music en route, it is a reasonable bet that people will take more and longer car rides if and when they become widely available at a more affordable price, much like the process through which added highway lanes induce new driving.

A Zoox robotaxi drives in San Francisco in November 2025. Photographer: Justin Sullivan/Getty Images North America

Robotaxis will also spend many miles deadheading, posing at least some risk to other road users even as they drive around empty (and potentially causing gridlock that slows everyone else on the street). If AVs become privately owned, they could be used on longer journeys that would have otherwise occurred on different modes – or not happen at all: Think daily multi-hour commutes to distant workplaces. (Or these cars might be empty, too, dispatched to perform ridehailing duty during their off hours.)

All this extra motion brings greater risks. If AVs are 20% safer than human drivers and expand total driving by 30%, total crashes would *increase* by 4%. (The effect would be higher still if people shift to AVs from transit, which is exceptionally safe.)

This is the standard that self-driven vehicles must meet to produce an overall reduction in crashes. Simply outperforming human-driven cars is not enough; they must do so to such an extent that they compensate for any added miles driven. But current AV safety analyses seldom take such network effects into account.

“If AVs are leading to fewer trips by train or subway and more miles on the road, that also could affect the overall safety benefit,” said IIHS’ Young. “It’s really important that

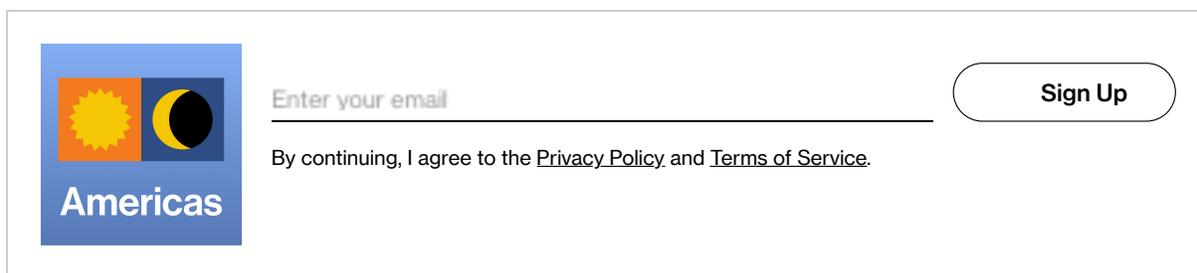
researchers examining AV safety report their findings in this context.”

A Much Cheaper Fix

Let's give AV companies yet another benefit of doubt and assume that their technology proves so powerful that it produces a net reduction in crash deaths even with an increase in total car use.

Still, that is not enough to justify government leaders prioritizing AVs as a road safety solution.

The reason is quite simple. Good policymaking entails choosing the most cost-effective ways to address a public problem, in this case traffic deaths. Self-driving technology is only one of many tactics available to reduce crashes, and it is not at all clear that it offers the highest return on investment.



The image shows a sign-up form for 'Americas'. On the left is a logo with a blue square containing a yellow sun and a black moon, with the word 'Americas' in white below it. To the right is a text input field with the placeholder 'Enter your email'. Below the input field is a line of text: 'By continuing, I agree to the [Privacy Policy](#) and [Terms of Service](#).' To the right of the input field is a rounded rectangular button labeled 'Sign Up'.

To offer just a few alternative approaches: Cars could be outfitted with Intelligent Speed Assist – a far simpler technology that automatically limits the driver's ability to exceed posted limits. Regulators could restrict the size of oversized SUVs and pickups that endanger everyone else on the road. Cities could build streets with features that are proven to reduce crashes, such as protected bike lanes, wider sidewalks, roundabout intersections and narrower travel lanes. States could legalize the installation of automatic traffic cameras that deter illegal driving. Bus and rail service could be expanded.

Unlike autonomous vehicles, these strategies have been reliably shown to reduce crashes – and without AVs' mitigating risk of expanding total car use. Many would cost a pittance and be available at lightning speed compared to the astronomically expensive, decades-long, and highly uncertain vision of replacing human-driven cars with self-driven ones.

Of course, road safety countermeasures are not mutually exclusive; many can (and should) be pursued concurrently. But if AVs are presented as a future panacea, wary policymakers will gain a new excuse to dither on controversial decisions (like implementing road diets) and ask, “Do we really need this? Let's just wait and see if technology makes the problem go away.”

Overblown safety claims could also justify industry-requested federal preemption

blocking states and cities from managing a technology that remains very much a work in progress.

No one should claim that self-driving cars cannot, at some point in the future, potentially reduce crashes across American roadways. But it would be equally wrong to suggest that they offer the only pathway to do so.

For proof, just look outside the US, where plenty of countries have managed to bring traffic fatalities down dramatically. Most recently, the Nordic cities of [Helsinki](#) and [Oslo](#) have gone an entire year without any cyclist or pedestrian deaths at all. How did they accomplish this feat? Local leaders credited a range of infrastructural and policy changes, from slower speed limits and wider sidewalks to higher car fees and stiffer enforcement.

Neither city has a single robotaxi plying its streets.

– David Zipper is a senior fellow at the MIT Mobility Initiative, where he examines the interplay between transportation policy and technology.

(Corrects fatality rate for US drivers in paragraph 28.)

Bloomberg

This story was featured in Bloomberg's Weekend section.

[**Read More Here**](#)

Contact us:

[**Provide news feedback or report an error**](#)

Confidential tip?

[**Send a tip to our reporters**](#)

Site feedback:

[**Take our Survey**](#)

Before it's here, it's on the [Bloomberg Terminal](#)