ALEKSANDAR BAURANOV

A Crash Course for Robotaxis

After a plane crash, the entire aviation industry learns from the failure. Autonomous vehicle companies should adopt a similar standard.

On October 2, 2023, a hit-and-run driver struck a pedestrian in San Francisco, throwing her into the path of a Cruise robotaxi. The Cruise vehicle braked but could not avoid running her over. What happened next turned a traffic accident into a company crisis: The robotaxi's software, which didn't detect the pedestrian pinned beneath it, initiated a pullover maneuver and dragged her 20 feet across the pavement. Cruise then compounded the failure by submitting a crash report that omitted the dragging entirely. This set off a chain of events that led to the suspension of the company's operating permit, the CEO's resignation, and finally, General Motors shutting Cruise down.

This episode sparked a reckoning for Cruise, but it should have triggered an examination across the whole autonomous vehicle industry. Instead, eight months later, a Waymo vehicle hit a metal barrier while turning and tried to pull over. The barrier hooked onto the Waymo, which continued to push it, scraping five parked vehicles before finally coming to rest. After contact in these incidents, both the Cruise and Waymo vehicles chose to move rather than stop. The National Highway Traffic Safety Administration incident database shows that, in the two years since the Cruise disaster, robotaxis are repeating the same dangerous maneuvers: They keep moving after hitting cyclists, pedestrians, and pets.

This pattern of the same types of accidents recurring does not stop there. Robotaxi passengers have repeatedly opened doors into oncoming traffic. Waymo cars illegally passed stopped school buses more than 20 times just in the last school year. The industry's problem is not that accidents happen. It's that the same ones keep happening.

The United States is regulating robotaxis and autonomous trucks as if each company's safety failures are only a problem for that company to address—but that's not how mature safety systems work. Autonomy is now a public-policy and infrastructure issue. To better regulate this growing industry, regulators should look to aviation safety infrastructure, which encourages open investigation and cross-operator learning. In aviation safety, when one company finds a serious hazard, the discovery is not proprietary. It is information relevant to public safety.

Today, when passengers board commercial airplanes, they don't need to know the maker of the plane, never mind weigh the relative safety merits of one design over another. In aviation, a serious failure is treated as an input leveraged by the institutional structures to translate one operator's failure into mandatory changes for everyone else. A century of cross-learning in this fashion has improved US airlines' safety record to the point where in 10 of the last 15 years, they've recorded zero passenger fatalities. As robotaxis and other autonomous vehicles become more common on city streets and interstate highways, operators must commit to collectively learning from accidents and implementing far-reaching changes, which will help prevent repeated harms and improve public trust.

How aviation built its safety system

After Air France Flight 447 from Brazil to France crashed into the Atlantic in 2009, agencies spent two years searching for the fuselage at the bottom of the ocean, at an estimated cost of $40 million. They found that speed sensors had malfunctioned due to icing, the autopilot disengaged, and the pilots mistakenly held the plane in a stall. In addition to punishing Air France, international aviation agencies required airline pilots across the globe to undergo mandatory high-altitude stall training. Airbus replaced speed probes on all its A330 and A340 planes. Other manufacturers tested their own speed sensors in icy conditions. A network of international agreements allowed Brazilian and French safety bodies to share findings and coordinate changes across jurisdictions. This single crash reshaped training, hardware, and procedures of hundreds of airlines worldwide.

Since the birth of aviation a century ago, the safety of the industry has advanced one catastrophe at a time. For the first two decades after Orville and Wilbur Wright's 1903 flight, pilots were unlicensed, aircraft uninspected, and there was no air traffic control, training requirements, or federal regulation. The first death came in 1908 when a flight piloted by Orville Wright crashed due to a propeller issue, killing his passenger. The first formal accident investigation occurred in 1912, and although it established that accidents should be analyzed to prevent recurring failures, there was no institutional force behind it. At the time, aviation was largely dominated by barnstormers performing entertaining but wildly dangerous stunts at county fairs. In 1923 alone, there were 179 barnstorming accidents with 85 fatalities. By the end of the decade, the average accident rate was 100 fatal accidents per 100 million miles flown, a rate that would produce 20 fatal airliner crashes every day at today's traffic volumes.

Despite this record, opposition to federal intervention was fierce. Barnstormers feared they wouldn't be able to operate in a regulated environment and resisted the idea of federal licensing. However, airplane manufacturers, looking to expand their market, wanted federal safety standards to build public trust and enable commercial growth—people would not buy airline tickets if they associated flying with death. At their urging, Congress passed the Air Commerce Act of 1926, which became the foundation of US aviation regulation. The act charged the Department of Commerce with promoting the growth of airlines and the aviation industry, but also put it in charge of licensing pilots, certifying aircraft, and enforcing air traffic rules. In the wake of this, the accident rate during the 1930s fell to around a tenth of the previous decade's.

Crashes continued to drive further improvements. When beloved University of Notre Dame football coach Knute Rockne died in a 1931 crash, intense public grief forced the Department of Commerce to release all accident reports publicly for the first time, helping make aviation safety a matter of public accountability. In 1935, when another crash killed US senator Bronson Cutting, the investigation became a political firestorm, ultimately driving the creation of an independent Air Safety Board, which had a mandate to investigate all aviation accidents across the country. In 1956, when two planes collided over the Grand Canyon, killing all 128 people aboard both aircraft, public shock catalyzed the Federal Aviation Act of 1958, which created the Federal Aviation Agency (FAA). The National Transportation Safety Board (NTSB) began operating in 1967 and has since investigated more than 153,000 aviation accidents. Cross-industry attention has cumulatively improved the safety record from 10 fatal accidents per 100 million miles in the 1930s, to 1 accident per 100 million miles in the 1950s, to 0.1 accidents per 100 million miles in the 1970s. The trend continued: Aviation today is another 1,000 times safer than in 1970s—many years see no fatalities at all.

Public transparency in the aviation industry extends beyond crashes. The Aviation Safety Reporting System (ASRS), established in 1976, allows pilots, controllers, and other aviation professionals to confidentially report near-misses and safety concerns without fear of retribution. The system, administered by NASA, has grown and became one of the most important data sources for regulators, researchers, and industry professionals who analyze the failures and propose solutions to prevent their recurrence.

Aviation's century-long arc demonstrates that if safety systems are not built proactively, they will be willed into existence one catastrophe at the time. The autonomous vehicle (AV) industry should consider whether it can afford to treat safety as a fragmented problem unique to each company when the repercussions of individual failures reach far beyond. As the Cruise incident showed, a single major accident involving a robotaxi can end a company. How safe does the AV industry need to be to gain public trust?

The United States is regulating robotaxis and autonomous trucks as if each company's safety failures are only a problem for that company to address—but that's not how mature safety systems work.

How safe is safe enough?

Autonomous vehicles are not, of course, commercial airliners; most people would group them more closely with cars and trucks. And society has made implicit judgements about what sort of safety record is acceptable for personal mobility. On the one hand, traffic safety has improved over the last 60 years from 5 fatalities per 100 million miles driven in 1966 to 1.1 in 2025. On the other, 40,000 people continue to die on US roads each year, suggesting that society accepts a different level of risk on the roads than in the sky.

When people evaluate acceptable risk, they don't only consider the likelihood of death—they also weigh their sense of agency and sense of randomness. When we drive, we feel agency and thus take a certain amount of responsibility for the outcome. Researchers such as engineer Chauncey Starr in the late 1960s and psychologist Paul Slovic in the late '80s found that people tolerate voluntary risk at rates roughly 1,000 times higher than involuntary risk. If I drive my Ford into a telephone pole, for example, I wouldn't think of suing Ford any more than I would think to sue Nike if I tripped over my shoelaces.

What would such a mechanism look like in practice?

The potential benefits of AV use can materialize only if the technology is safe at scale. So far, that scaling has been slow. Robotaxis have logged about 0.2 billion miles—a rounding error compared to total annual vehicle miles on US roads of about 3.3 trillion. In the current fragmented regulatory landscape, each locality is creating its own rules. For companies, the cost of compliance multiplies with each new market. Yet without overarching federal regulation, states compete to attract AV companies by making their rules as permissive as possible, and the weakest rules effectively set the national floor.

Today, the National Highway Traffic Safety Administration (NHTSA) occupies the same position the Department of Commerce held for the aviation industry in 1935, but it lacks the tools FAA has today. It is tasked with both fostering automotive innovation and investigating failures. But it needs new authorities, more like FAA, to intervene effectively in governing autonomous vehicles. NHTSA has the legal authority to identify defects in vehicles of the same design and manufacture and to compel recalls, which works well for mass-produced hardware. A faulty ignition switch in one Chevy Cobalt is identical to that in every other Chevy Cobalt, because they all came off the same production lines with the same parts. But FAA, beyond identifying faulty hardware, can identify unsafe behaviors or a broad accident category and issue orders that require airlines and manufacturers to rectify the situation.

Empowering NHTSA to be more FAA-like requires several steps. First, regulators should establish a common taxonomy for autonomous vehicle incidents and hazardous behaviors. The industry cannot learn systematically if every company categorizes edge cases differently. A shared structure for reporting would make it easier to identify patterns across firms and cities. NHTSA should develop, maintain, and—as needed as the industry matures—update an incident taxonomy based on different levels of autonomy in consultation with operators, academic researchers, and accident investigators. FAA used such a collaborative model in developing its Aviation Safety Information Analysis and Sharing system. Second, certain classes of incidents should automatically trigger cross-operator review. Congress should grant NHTSA the authority to issue cross-operator safety directives for autonomous vehicles, modeled on FAA's Airworthiness Directive process. FAA does not prescribe how

Aviation's century-long arc demonstrates that if safety systems are not built proactively, they will be willed into existence one catastrophe at the time.

However, people perceive injury as unjust if they feel they had no agency, were the victims of others' decisions, or suffered a random breakdown in a system. In this way, robotaxis more closely resemble aviation, because most passengers have no agency in how the vehicle or aircraft operates. Likewise, random glitches or malfunctions and the repetition of similar failures signal not just technical immaturity, but institutional weakness. They suggest that no one is fully in charge.

The American public's acceptance of rather significant risks while driving their own cars, in other words, doesn't map directly onto AVs. Obviously, applying aviation's safety system to ordinary driving—to the 240 million licensed drivers and nearly 300 million registered motor vehicles in the United States—was never realistic. But autonomous vehicles change that calculus. A fleet of robotaxis can be trained, tested, updated, and monitored with the kind of institutional discipline that human driving never could absorb. What's more, doing so builds a structure to earn public trust, because a new technology can survive isolated failures and thus overcome randomness if the public believes the system learns from those failures.

each airline implements a fix. It defines the unsafe condition and requires each operator to demonstrate resolution in its own fleet. NHTSA could operate the same way: at the level of what the vehicle does, not how the software is written. When one operator demonstrates a hazardous failure mode, NHTSA should be empowered to compel all others to test for the same vulnerability, implement corrective action, and demonstrate compliance on a mandatory timeline. The new directive should describe the unsafe behavior, set a deadline for compliance, and apply to every company operating autonomous vehicles on public roads, not just the one that happened to fail first.

Third, NHTSA's Standing General Order on Crash Reporting captures crashes but excludes near-misses, system shutdowns, sensor failures, and close calls. Critics have called for expanded reporting beyond crashes. NHTSA should broaden the order to capture events such as unplanned disengagements of the autonomous system, remote human-operator takeovers, sensor degradation episodes, emergency braking maneuvers, and others, which will enable better identification of root causes and prevention of hazards.

Finally, a subset of significant incidents should receive independent, public-facing analysis. This does not require replicating the entire machinery of aviation accident investigation. But it does require moving beyond opaque company summaries and fragmented incident narratives. Currently, AV operators publish limited information about safety. As of February 2026, for example, Waymo had logged 200 million miles with two fatalities, a somewhat better rate than human drivers in the United States. However, companies publish limited narratives about their accidents and face no requirements to disclose more. Tesla does not even go that far; their incident summaries in NHTSA's database are redacted. However, even if companies wanted to share more, there is no confidential, nonpunitive mechanism for reporting near-misses or other unexpected behaviors. A well-designed voluntary reporting system could capture far more safety-relevant information than crash reports alone, enabling the industry to improve proactively instead of reacting to serious incidents.

The time is right

The AV industry has an opportunity that aviation never had—to learn from history rather than relive it. Institutional conditions right now are nearly identical to aviation's early stages: fragmented oversight, self-certification, no independent investigation, and no cross-operator learning.

But the autonomous vehicle industry is better positioned for cross-operator safety learning than aviation ever was. The number of operators is small. Waymo alone accounts for more than 95% of all robotaxi rides in the United States. Even as Zoox, Tesla, Aurora, and others expand, the number of companies will remain in the single digits for the foreseeable future, and coordination among a handful of companies is simpler than among hundreds of airlines. More importantly, autonomous vehicles generate replayable digital evidence through cameras, lidar, and detailed logs. Cause-finding is faster and more precise than recovering a black box from the bottom of the ocean. And fixes propagate at the speed of software—orders of magnitude faster than grounding hundreds of Boeing 737 MAXs.

Unfortunately, right now, regulation is moving in the opposite direction. NHTSA lost 25% of its staff in 2025. The agency's AV framework from April 2025 has three stated principles, including "unleash innovation by removing unnecessary regulatory barriers." Meanwhile, the industry is also split. Some established operators like Aurora and Waymo, whose CEO recently expressed support for government regulation, recognize that public trust can be a bottleneck for commercial expansion. Others, like Tesla, resist standardized reporting and external oversight. But safety is a necessary precondition for the sustainability of the AV industry, and the federal government needs to provide a critical input: safety regulation.

The AV industry has an opportunity that aviation never had—to learn from history rather than relive it.

Robotaxis, over time, could lead to safer roads. International markets may be more open to operators who have proven safety records, especially for autonomous logistics players and trucking companies. But we might never get there if every operator is left to learn on its own. Aviation history suggests a better path. The fix is not to ban robotaxis. The standard should not be "zero crashes," which is neither realistic nor especially useful as a governance principle. Nor should it be simply "fewer crashes than human drivers," a benchmark that can obscure machine-specific failure patterns. In the fight against 40,000 traffic deaths, the goal isn't zero accidents. It's zero repeats.

Aleksandar Bauranov *is a principal at Boston Consulting Group. He previously worked as an aviation safety researcher at the University of California, Berkeley, with the Federal Aviation Administration and NASA.*