

2026

Unlocking the Promise of Autonomy

Uber

Table of contents

3	Introduction
5	Executive summary
6	A hybrid network
7	Promise of autonomous technology
9	Responsible transition
10	How workers will be impacted
11	A responsible transition for workers
11	Serving all communities
12	Rethinking accessibility
13	Working with cities
14	Safety
15	Conclusion

Introduction



More than fifteen years ago, Uber set out to help cities move. Our ambition was simple: make it easier for people to get where they're going. Over time, that ambition expanded. What started as a new way to request a ride began to look like a new vision for transportation, one that is more efficiently shared, increasingly electric, and eventually autonomous.

Today, that future is beginning to take shape. Autonomous vehicles are moving out of research environments and onto public streets. The transition is happening faster than many people expected. Uber now has partnerships with more than 25 leading AV developers, driverless operations in multiple markets, and plans underway for dozens more cities. We are already facilitating millions of autonomous trips per year on our platform, and believe autonomy will play a central role in the future of cities.

The potential is significant. If deployed thoughtfully, autonomous vehicles will reduce crashes, accelerate electrification, and expand independence for people who cannot drive or should not have to. They will lower the cost of transportation and widen access to opportunity. That promise has drawn enormous investment and attention from companies, governments, and technical talent around the world.

But this technology does not arrive in controlled conditions. It arrives on real streets, in real weather, around real people. And while the momentum behind autonomy is undeniable, it raises important questions. How should safety be measured in a way the public can trust? What does responsible deployment look like when systems fail or when conditions are unpredictable? How do we ensure the benefits are broadly shared rather than concentrated in a small number of markets, and even in a small number of neighborhoods? And how do we manage the impact to millions of workers, including the drivers who help keep our cities moving today?

This paper aims to address these questions.

Progress in transportation has always depended on new technology, and autonomy will be the primary engine of that progress over the next decade. But innovation only succeeds if it earns public trust—that's a lesson Uber has learned over the years (sometimes the hard way). Building that trust requires transparency, clear standards, and accountability.



Operating in 70 countries and more than 15,000 cities, Uber sees first-hand how complex real-world transportation systems are. While autonomous technology is enabled by software, it's software that's deployed in and directly affects the real world. As a result, the deployment of autonomous technology will always be shaped by the realities of infrastructure, regulation, and the needs of the cities where it operates.

That requires partnership with regulators who set standards, with first responders who respond when things go wrong, with workers whose livelihoods are affected, and with communities that deserve a voice in how public roads are used.

Our goal is to be clear-eyed about what autonomy can deliver, and equally clear-eyed about what must be done to deploy it safely, responsibly, and at scale.

Executive Summary

Autonomous vehicles will reshape transportation in ways that are both significant and difficult to predict. At Uber, we believe three priorities should guide how the technology is introduced.

→ A transition for workers

We owe people the truth. Autonomous vehicles may expand mobility overall, but they could also mean less work for many drivers over the long run, including people who rely on platform work as a flexible safety net. We believe a phased transition to a hybrid model where drivers and autonomous vehicles work side-by-side offers the most promise for workers and greater reliability for consumers and cities. A too-rapid shift to driverless fleets risks exposing gaps in today's labor and benefit systems, pushing the burden onto the workers and communities to absorb the disruption. **If we want an autonomous future that holds together socially, we have to plan for any workforce disruption with the same seriousness we bring to the underlying hardware and software.**

📍 Serving all neighborhoods

Left to pure market dynamics, autonomy will concentrate in wealthier urban cores. That path risks recreating a two-tier transportation system, where access is determined by zip code and the neighborhoods that most need reliable mobility are the last to benefit. The move to driverless technology will also create a gap for riders who need physical assistance, including people using folding wheelchairs or managing mobility challenges that require a human hand, not just a vehicle at the curb. Accessibility and coverage cannot be treated as add-ons.

📋 Talking with cities, not at them

Autonomy must happen with cities, not to them. Poorly managed deployment could worsen congestion, cluster vehicles in already overcrowded areas, and put new strain on local infrastructure as cities contend with charging depots, noise, land use, and grid demand. And while AV technology has advanced rapidly, today's systems still struggle with complex edge cases that our response frameworks are not yet fully built to handle. Cities should be partners in shaping how this works, and accountable parties should be clear from day one. This is an area where we've learned quite a bit from experience, including our share of mistakes, during our early days.

A hybrid network

The future of mobility should not be a choice between humans and robots. The most resilient transportation systems have always combined both.

Autonomous vehicles are remarkable digital systems that must contend with the realities of the physical world: severe weather, emergencies, natural disasters, mapping needs, and infrastructure demands. A hybrid network allows AVs to scale thoughtfully, versus all at once, reducing the risk of flooding already congested streets with driverless fleets. It gives cities the ability to expand where the technology performs well, and ensures that communities outside areas where AVs operate continue to have access to reliable transportation.



There are 8,000 cities in the US alone that Uber serves which likely won't be workable for AVs in the foreseeable future. If we rely on AVs only, we'll go back to a world where reliable transportation was reserved for a subset of the broader population.

A hybrid network also accounts for older adults or anyone who needs driver support getting into a car. And it ensures reliability when AVs have to shut down without warning during bad weather or blackouts.

Protecting access for riders is only half the equation. A hybrid network also protects drivers and allows for a phased transition.

A hybrid network can encourage more drivers to earn to meet peak demand, special events, and emergencies without requiring permanent overcapacity. That flexibility prevents a return to rigid, fixed-supply systems where service breaks down when demand spikes, and it ensures drivers remain a core part of the network. Without that flexibility, systems must build enough vehicle capacity to meet peak demand at all times, leading to larger, underutilized fleets and unnecessary strain on electric grids.

It also preserves meaningful earning opportunities on our platform over the long-term rather than triggering sudden displacement. That sequencing matters. In systems without drivers, reliability depends on deploying large-scale AV fleets to meet demand, while in a hybrid model, AVs can scale more gradually, giving drivers time to adapt while maintaining reliable service for riders.

The data bears that out. We've already seen that in Austin, Atlanta, and Phoenix, where AVs are already sharing the Uber app with drivers, trips are up and earnings have stayed consistent.

Promise of autonomous technology

Autonomous vehicles represent a structural shift in how transportation works. For more than a decade, we have argued that autonomy would become a defining feature of the future of mobility. Not because it was inevitable, but because the underlying benefits were so clear. Advances in computing power, sensing technology, and artificial intelligence have converged with growing demand for transportation that is safer, more efficient, and more sustainable.

We are beginning to see what autonomy can deliver in practice:



Safer streets

Road safety remains one of the most urgent public health challenges in the world. Human error contributes to the vast majority of crashes. Autonomous systems do not drive impaired, distracted, or fatigued. If deployed responsibly, they could help prevent tens of thousands of traffic deaths each year.



Broader access to mobility

For seniors, people with disabilities, and others who cannot drive, autonomous vehicles could expand independence and access to transportation.



Faster electrification

Paired with electric vehicles, autonomous fleets could accelerate the transition away from internal combustion engines and help reduce transportation emissions.



More affordable movement

Over time, lower operating costs could reduce the price of reliable transportation, making mobility more accessible for riders and more efficient for cities. While this is not yet true today, we expect to see these improvements in affordability as the technology (and cost to scale it) improves in the years ahead.

Realizing these benefits will depend not only on continued technical progress, but on how that technology is deployed. Autonomous systems are beginning to operate on public roads today, and the decisions made in this early phase will shape public trust in the technology for years to come.

The central challenge is not whether autonomous vehicles can work, but how they are introduced into the broader transportation system.

The central challenge is not whether autonomous vehicles can work, but how they are introduced into the broader transportation system. Safety must be measured in ways the public can trust. Deployment must happen with transparency and clear accountability. And the transition must be managed in a way that recognizes its impact on workers, cities, and communities.

Uber sits at the intersection of riders, cities, and technology developers. Through partnerships with more than 25 AV companies and operations across thousands of cities, we see both the opportunities and the challenges of autonomy in real-world transportation systems.



Responsible transition

Technological change creates both opportunity and risk. Uber has seen this first-hand. Over the past decade, we have helped billions of people gain access to reliable, on-demand transportation. And we have also seen how quickly innovation can reshape markets, institutions, and access to work.

As autonomy advances, we should not treat this moment as a clean slate. The lessons of the past decade matter. Growth without guardrails creates strain. If the next chapter of mobility is going to succeed, it must begin with a clear understanding of both the opportunities and the risks.



The potential of autonomous technology is significant, but there is no guarantee its benefits will be distributed evenly. Realizing the best outcomes will require honesty about the technology's current limitations, its edge cases, and the ways it could reshape work. The long-term benefits of autonomy will be stronger and more durable if workers, communities, and cities are part of the conversation from the start.

No single policy lever will resolve every tension created by this transition. But engagement matters. That means working with drivers, partnering with independent researchers to study the impact on drivers, and creating structured ways for drivers to share feedback as the technology evolves. It also means working with first responders including police, firefighters and EMS. Autonomous vehicles need to complement emergency services not complicate them. That requires both clear industry standards and thoughtful regulation.

The path forward should be built deliberately, informed by evidence, and grounded in shared responsibility. If we are transparent about what autonomy can do, and what it still can't, we can make sure the autonomous dividend is broadly shared. Regulating new technology requires nuance, which means there aren't a lot of easy answers, but we're committed to working together on smart solutions.

How workers will be impacted

Autonomous vehicles are AI in the physical world, and like AI, they will reshape how people work. That shift will create new opportunities, but it will also displace workers who rely on platform work. The question is not whether work will change. It is how we manage that change.

Autonomous vehicles are AI in the physical world, and like AI, they will reshape how people work.

As autonomous vehicles begin operating alongside human drivers, they will influence work and earnings dynamics across markets. We believe overall demand for mobility will continue to grow. But we should be honest: in many cities, drivers will face fewer, or different, trip opportunities over time. Acknowledging that reality is a necessary starting point for handling the transition responsibly.

The impact will not be uniform. Most drivers engage in platform work part-time, relying on it as flexible, low-barrier income. Globally, 61% of drivers and couriers use Uber fewer than 20 hours per week, and 85% do so fewer than 40 hours per week. For many — including women, caregivers, and workers from lower-income communities — platform work functions as a safety net and an entry point into the labor market.

In California, one AV does the work of about four drivers.

We have seen early signals that this impact is not just theoretical. In San Francisco and Los Angeles, where drivers compete against AV-only networks, driver utilization and hourly earnings declined last year. Given that the vast majority of drivers work part time, and an AV is online most of the day, we see that in California,

one AV does the work of about four drivers. Over time, the drivers who rely on the flexibility of Uber the most will be disproportionately affected in these first-party-only markets, while others will have to work longer or at different times and places to earn the same amount.

A phased transition to a hybrid network substantially reduces these risks. AVs will not need to scale as fast on a hybrid network, giving drivers more time to adapt to changes. Hybrid networks may not need as many autonomous vehicles either, preserving work for existing drivers. The alternative of a rapid transition, with thousands of directly-competitive vehicles appearing overnight, risks uprooting an untold number of drivers.

But a hybrid network alone will not be enough to mitigate the labor impacts. Industry and governments will need other solutions. That is why local research, transparent data, and sustained engagement with workers and experts are essential. Policies should reflect observed outcomes, not assumptions. We plan to support high-quality, market-specific research and hope others will as well.

A responsible transition for workers

Worker protections will also need to evolve alongside AV deployment. As autonomous services expand, changes in demand patterns and operating models may strain driver benefits and protections designed for a context when AVs were not yet a reality.

In cities where protections and benefits are tied directly to fares, this creates a risk of regulatory arbitrage. For example, AV operators may be incentivized to prioritize markets with higher labor costs, not because those markets need autonomy most, but because replacing human drivers yields greater economic returns in those places. Policies designed to protect workers should not unintentionally accelerate their displacement. To remain effective, regulatory frameworks will need to recognize that AVs and human drivers operate within a single ecosystem, and ensure protections continue to support workers through the transition.

Autonomy may create new roles in fleet management, maintenance, remote operations, and oversight. But those positions are unlikely to absorb the majority of affected drivers. A responsible transition will therefore require rethinking protections for workers, alongside policies that support alternative pathways.

We are engaging in efforts to support this transition, but we don't have all the answers. Durable solutions will require broader collaboration among governments, industry, and worker representatives.

Serving all communities

Autonomous vehicles will not expand opportunity for everyone, if they are only accessible to a few. For residents of underserved neighborhoods, where reliable transportation can determine access to jobs, healthcare, and education, autonomy offers independence. But that outcome is not automatic.

Autonomous vehicles will not expand opportunity for everyone, if they are only accessible to a few.

On-demand transportation must work for more than the easiest trips and the most predictable streets. Companies and policymakers share responsibility for ensuring innovation reaches entire communities, rather than reinforcing the gaps that already exist.

We have seen what happens when access is left to convenience and profitability alone. Before ridesharing platforms expanded coverage, taxis concentrated in airports, and dense downtown corridors, leaving many neighborhoods unserved. As AVs scale, we cannot allow a quiet return to a system where service depends on zip code.

The risk of a two-tier model is already real; AV services tend to launch and concentrate in well-resourced areas with predictable traffic patterns and strong demand, while



deprioritizing transportation deserts or communities with more complex operating conditions.

In San Francisco, for example, AV operations have expanded into some of the wealthiest neighborhoods while remaining absent from nearby cities like Oakland. Some of these choices reflect technical constraints. Some reflect business decisions aimed at maximizing vehicle utilization.

Without guardrails, these factors will continue to produce uneven coverage. A two-tier system will narrow the very benefits autonomy promises to expand.

Rethinking accessibility

If we are redesigning vehicles from the ground up, we cannot replicate old exclusions in new technology.

As AV platforms scale, manufacturers and operators have an opportunity to address historical wheelchair accessibility issues head on and build accessible vehicles from the ground up.

But accessibility is not limited to wheelchair access. Even a technically accessible vehicle may not meet every rider's needs. Many customers require some form of human assistance, particularly in complex or unexpected situations.

Older adults with mobility challenges, riders using folding wheelchairs, and individuals with limited use of their hands or arms often rely on driver support. Others may need assistance navigating pickup points or adjusting to changing curb conditions. These needs do not disappear simply because the vehicle drives itself – in many cases they are made more acute.

A transportation system that serves everyone must account for the full range of physical, cognitive, and situational challenges riders face. That will require deliberate policy choices, intentional platform design, and sustained collaboration with cities and disability advocates. Autonomy offers a chance to improve access meaningfully – but only if inclusion is treated as a core design principle, and not just a compliance requirement.

Working with cities

As autonomous vehicles expand, the question is not simply how well they perform, but how they will shape the cities around them. Autonomy could strengthen urban transportation systems, but that's not a given. Whether AVs reduce congestion or deepen it will depend on deployment choices and on how cities plan ahead.

Whether AVs reduce congestion or deepen it will depend on deployment choices and on how cities plan ahead.

There is a credible risk that poorly managed AV deployment could increase congestion in dense city centers. Human-driven rideshare networks contain a natural constraint: drivers weigh their time against potential earnings. When demand drops, many log off. This is a fundamental tenet of how Uber matches supply and demand.

AVs operate under different constraints. An AV can remain in circulation so long as operating costs are covered. As the costs of autonomous vehicles decrease, the incentive to drive around empty will increase. In economic terms, it may be rational to keep vehicles on the road until just past the point that likely revenue is exceeded by operating costs. In practical terms, that could mean more vehicles cruising without passengers. At the extreme, one could imagine privately owned AVs circling blocks for hours to avoid parking costs. Cities will need strategies to discourage AVs from driving around empty.

Recognizing this challenge upfront gives us the opportunity to be deliberate about how to address it. Hybrid networks will naturally need fewer AVs and have higher utilization. Integrated into a highly utilized hybrid network, AVs could reduce the total number of vehicles required to move people through a city. But that outcome depends on policy decisions around congestion pricing, curb management, fleet utilization standards, and integration with public transit.



Mass rollouts also raise land use questions. Large AV fleets require charging infrastructure, storage facilities, and maintenance depots. Decisions about where those sites are located will affect neighborhoods differently.

Autonomy will not operate in isolation. It will interact with buses, bikes, pedestrians, delivery vehicles, and emergency services. The opportunity to make cities greener, safer, and more vibrant is significant. So is the responsibility to ensure that in solving one set of urban challenges, we do not create another.

Safety

Nowhere is the case for autonomy stronger than safety. Human error remains the leading cause of serious crashes. Autonomous vehicles do not drink, text, or lose focus. In a country where more than 40,000 people die each year in traffic incidents, even incremental improvements in safety translate into thousands of lives saved.

The long-term safety potential is real. But real-world operation is not a laboratory environment. Public trust will depend not only on safety outcomes, but on transparency around how systems perform when things inevitably go wrong.

Autonomous systems still struggle in the kinds of ambiguous, fast-moving scenarios where humans rely on instinct and contextual judgment. Edge cases, like erratic human behavior or infrastructure failures, test the limits of current technology. Recent fleet-wide outages in San Francisco and Wuhan highlighted how AVs can be challenged by sudden disruptions that alter the streetscape in ways systems were not trained to anticipate.

Edge cases may sound rare. But at scale, rare events are constant. When tens of millions of trips occur each day, multiple one-in-a-million scenarios happen daily. Companies must plan not only for optimal conditions, but for edge cases and failure modes with clear protocols, remote support, and accountability built in from the outset.

Responsibility also extends beyond manufacturers. Cities will need infrastructure capable of supporting new vehicle systems. That is why early and sustained collaboration between cities and industry is essential. Questions around depot noise, coordination with first responders, navigation through construction zones, and response during natural disasters cannot be solved unilaterally. They require joint planning and continuous refinement as deployment expands.



Conclusion

Autonomy will improve how people move: fewer crashes, broader access, faster electrification, and lower costs over time. But those outcomes are not guaranteed. If we roll out technology without guardrails, we risk eroding trust and concentrating the benefits. The transition is already underway, and now is the time for governments and industry to act.

The right path is phased and hybrid, moving fast but not all at once. This reduces real risks: worker displacement, uneven access, and pressure on city infrastructure and safety systems. We should stay focused on three priorities. First, a dignified transition for workers that avoids abrupt disruption and preserves earning opportunities. Second, broader accessibility as a requirement, ensuring autonomy expands access and doesn't limit it. Third, partnership with cities to manage congestion, infrastructure, and safety with clear accountability. If we get this right, we build trust, preserve flexibility, and put autonomy on a path to scale that holds.