



# Algorithmic Transparency Report 2025

Europe

Uber



# Table of Contents

	<b>Introduction</b> .....	1
	<b>A note on terminology</b> .....	3
	<b>The Role of Algorithms and AI at Uber</b> .....	4
	<b>How Uber Uses Algorithms and AI</b> .....	6
	Matching .....	6
	Pricing .....	8
	Use Cases specific to Uber Eats .....	12
	Algorithms Enhancing Safety .....	14
	<b>Responsible approach to AI: Governance and Privacy</b> .....	16
	Governance .....	16
	Privacy .....	16

# Introduction

**Uber's mission is to create opportunity through movement across Europe.** Operating across 29 countries and more than 3,000 cities, Uber connects millions of passengers and consumers, with hundreds of thousands of drivers, delivery people and merchants everyday.

Behind the scenes, algorithms and artificial intelligence (AI), including machine learning (ML) systems, work together to make those connections possible—matching people in real time, setting transparent prices, and helping keep trips safe and reliable. These systems are what allow Uber to operate efficiently at scale —powering hundreds of millions of trips and deliveries every year—while adapting to local conditions across Europe.

**Platforms like Uber fuel economic growth and innovation in Europe.** In 2024 Uber generated €24 billion of economic value to the European Union, and saved Europeans 48 million hours of unnecessary travel time every year, worth approximately €379 million to the European economy. Drivers and delivery people earned more than \$6 billion on our platform in the European Union 2024. It's not just on our app, we have thousands of engineers and staff across Europe focused on building and governing these algorithmic and AI systems responsibly, with three Centers of Excellence in Cracow, Limerick and Lisbon.

€24

Billion of economic value generated

48

Million hours of unnecessary travel time saved



**Designed with all platform users in mind:**

Algorithms and AI make journeys and delivery orders safer and reliable, economically attractive for all users of our platform and the in-app experience smoother. This benefits all of our users—from drivers and delivery people who can maximise their earnings; merchants who can increase their consumer base and generate extra revenues; passengers and consumers who enjoy daily convenient services to go anywhere and get anything.

Ultimately, **Uber users have the best experience when our platform works seamlessly, and all our users enjoy using it everyday.** This is why Uber is committed to ensuring our platform remains reliable and safe for all, through the responsible and transparent use of algorithms and AI.

**ALGORITHMS AND AI AT A GLANCE**

Every day, algorithms and AI power the heart of our platform by predicting traffic in real time, finding the best trip and delivery matches, and providing clear pricing on the spot. These systems support more than 1.5 million trips every hour, helping people get to work, make their flights, or simply skip cooking dinner. More importantly, they are a critical driver of safety, working in the background to protect passengers and drivers when it matters most. From guiding loved ones to doctor appointments to connecting people where public transit leaves off, Uber uses algorithms and AI to solve real problems for real people.

**29**  
Countries**3K+**  
Cities

**Uber's approach to algorithmic and AI governance is grounded in structure, accountability, and continuous improvement,** with a strong emphasis on transparency, fairness, and human oversight throughout the lifecycle of high-impact systems. We operationalise these principles through a combination of formal governance processes and cross-functional collaboration. Our governance approach is guided by industry standards, such as the National Institute of Standards and Technology (NIST) AI Risk Management Framework and International Organization for Standards (ISO) for AI Management System, and grounded in compliance **with European and national regulations.** And importantly, human judgment is embedded throughout the process—from model design and risk assessment to deployment and review—so that automation never replaces accountability.

As the role of algorithms, artificial intelligence (AI), and machine learning (ML) deepens, so too does **our responsibility to ensure these systems are transparent, fair, and aligned with European values.**

As part of our commitment to transparency, this report focuses specifically on how Uber's algorithmic and AI systems function within the European Union. It outlines the mechanisms behind decision-making on our platform, describes our robust governance structures, and shares our approach to fairness, safety, and accountability in the context of Europe's evolving digital landscape.



# A note on terminology

Algorithms, artificial intelligence (AI) and machine learning (ML) are closely related, but not identical, concepts. We want to start by highlighting some definitions and how they apply to our platform.

- **Algorithms:** Algorithms are step-by-step instructions for solving a problem. Some are deterministic, meaning they follow the same steps and produce the same outputs given the same inputs, such as calculating the shortest route between two points. Others incorporate the outputs of ML models, meaning they can adapt to new conditions or data over time. At Uber, algorithms often combine both types: fixed logic to structure decisions, and adaptive inputs from ML models to improve accuracy and relevance.

## EXAMPLE:

When you request a trip, an algorithm helps determine which nearby driver to offer it to. The fixed logic ensures the system only considers available drivers, while ML-powered predictions, such as estimated travel times, influence which option is most efficient.

- **Machine learning (ML):** ML uses data to learn patterns and make predictions, like estimating the likelihood of an event, that can be used by algorithms to guide decisions.

- **AI:** AI is a broader umbrella that encompasses ML, referring to computer systems designed to perform tasks that typically require human intelligence, such as reasoning or decision-making. For the purposes of this paper and simplicity, we'll use AI and ML interchangeably.

## EXAMPLE:

Uber's ETA (Estimated Time of Arrival) predictions rely on machine learning. The system uses large amounts of past trip data (traffic patterns, time of day, weather, city events) to predict how long a driver will take to pick up a passenger or complete a trip. Unlike a simple algorithm, the model learns from past trip data and gets better over time as conditions change.

In practice, many of Uber's systems combine all three (Algorithms, ML, AI), using algorithms as building blocks and AI/ML to adapt complex patterns into real-world applications.



# The Role of Algorithms and AI at Uber

## **Our platform can only function at scale thanks to algorithms.**

Algorithms are essential to operating at the scale and complexity that Uber demands. With millions of passengers, consumers, drivers, merchants, and delivery people interacting on our platforms every day, manual or rules-based approaches are not sufficient to provide fast, accurate, and context-aware decisions. Algorithms allow us to manage this complexity in real time, turning vast volumes of data into actionable insights that improve user experience, system efficiency, and safety.

This technology may seem complex at first, but in this document we will do our best to break them down in an accessible way. We use some of the most advanced technologies in the world, and any discussion about it can rapidly become quite technical. All of this work is grounded in the need to operate efficiently at scale and ensure we can provide the best platform for customers, drivers and delivery people.

Rather than relying on static assumptions, ML models can learn from historical trends and real-time signals to improve outcomes over time.

### **EXAMPLE:**

The ML model used by Uber to estimate a trip's arrival time generates predictions based on factors such as distance, infrastructure, estimated traffic or estimated vehicle speed. This system learns from other similar trips, and can, for instance, adapt based on temporary or permanent road closures. This adaptability helps ensure the platform remains reliable and responsive, even as urban environments and customer expectations evolve.

## **Designed with all platform users in mind.**

In practice algorithms are essential tools that deliver meaningful benefits in people's daily lives.

Algorithms support a wide range of activities on our platform from determining which nearby drivers or delivery people to offer a trip or delivery to estimating arrival times so passengers and consumers have accurate waiting time information. We use algorithms to calculate prices dynamically to ensure that our services are reliable even in moments of high demand, meaning that trips are financially attractive

to drivers, and affordable for passengers. On Uber Eats, algorithms are also used to route orders between delivery people and merchants, factoring in preparation time, traffic, and location. Pricing is similarly designed to create attractive earnings for delivery people and merchants, and to make the offers appealing to consumers.

Algorithms can support systems that are more fair and consistent by applying trained models to data in a systematic way, rather than relying on individual judgment that may vary from case to case and is prone to bias. This reduces the likelihood of arbitrary outcomes and creates a stronger foundation for oversight. This does not eliminate the need for human governance—on the contrary, it makes governance even more important. Oversight is essential to ensure that models are applied responsibly, remain aligned with policy goals, and adapt appropriately as conditions change. And it offers a powerful tool for delivering consistent service across diverse geographies, while adapting to local conditions.

#### EXAMPLE:

**Fraudulent activity on Uber can sometimes involve a driver manipulating their GPS signal so it appears they are in a different location than they actually are. This tactic known as GPS spoofing can unfairly generate earnings or disrupt trip matching, harming both riders and drivers.**

**To address this type of fraud, if a driver's GPS location shows them suddenly jumping across the city in ways that don't match**



**real-world driving, the system may flag this as potential GPS spoofing, which may lead to the removal of the driver's access to the platform. Our internal review team is provided with clear explanations of what triggered the flag, if the team determines a driver's account should be restricted, the driver is informed. If the driver believes there is a mistake, they can submit evidence, which helps our team review the case. This combination of automated detection, human review, and transparent communication promotes safety on the platform while inserting accountability for our systems. The goal is to prevent fraud that impacts not just customers but other drivers who may have been matched with a trip.**

Finally, the use of algorithms empowers experimentation and innovation, which ultimately makes our platform better for all our users. Through techniques like A/B testing, reinforcement learning, and simulation, we can iterate on new features and policies at scale, measuring their impact and refining them in a structured way. This allows Uber to stay agile and data-informed as we build the next generation of mobility and delivery solutions.



# How Uber Uses Algorithms and AI

## Matching Customers, Drivers, Delivery People and Merchants

At any given moment, millions of people around the world are logged on to our app, with new people logging on every second. That means millions of possible matches between passengers and drivers, delivery people, merchants and customers. Throw in real-world factors like traffic jams, and the task of matching passengers and drivers becomes more complex. That's where AI comes in.

## How matching works

**When a trip or delivery request is generated, the platform must quickly determine which driver or delivery person to offer it to.** Uber's matching system reviews who is online in the area, the type of trip requested, and conditions in the local city of operation. It looks at the distance to the passenger or customer, the product or vehicle type (for example, UberX vs. Comfort or bike vs. motorcycle for delivery), and any special requirements such as airport queues.

**The most important element in matching is location,** which depends on the layout of a city, traffic patterns, and geographic features. AI is essential in translating distance as the crow flies, into the real, lived experience of time and



distance. Closest doesn't always mean quickest. For example a one driver may appear closer but between the driver and a rider may be a series of one way streets or congestion that may cause the closest driver to have a much longer estimated time of arrival.

The matching algorithm is designed to optimise reliability and efficiency for all trips or deliveries at any one time, not just for an individual trip or individual delivery. It's about matching as many passengers who have requested a trip with as many drivers, and as many consumers with delivery people. Drivers or delivery people are then free to accept or reject offers.

For example, the mobility matching algorithm weighs whether a match is likely to reduce overall wait times, avoid unnecessary empty travel, and increase the chance of a driver receiving a

subsequent request after drop-off. For our Share products where multiple passengers are picked up by the same vehicle, the system considers whether routes overlap to minimise detours.

Sometimes, drivers ask why they're matched with certain passengers instead of others nearby. The reason is that the system is designed to reduce wait times for all users and make as many trips happen as possible. This can mean that a driver is matched with the passenger who isn't the closest, but whose trip leads to a better outcome for users as a whole.

By using algorithms in this way, Uber is able to balance millions of simultaneous potential matches across the platform—helping passengers get faster pickups, enabling drivers and delivery people to find more worthwhile trips, and keeping our platform dependable and affordable.

#### CASE STUDY:

In a [2024 paper](#), Uber Applied Scientist Dan Knoepfle, and economists Juan Camilo Castillo and Glen Weyl describe a problem in ride-hailing called a wild goose chase.

**This happens when there aren't enough available drivers compared to trip requests. Under a simple nearest-driver dispatch model, some passengers are matched with drivers close by while others are matched with drivers who are far away. That sets off a negative cycle for those matched with faraway drivers: passengers wait longer, drivers waste time driving without passengers, make less money, and some drivers log off, making the shortage worse.**

The study found that smarter dispatch systems and dynamic pricing technology—like the systems used by Uber—help minimise wild goose chase spirals by better balancing supply and demand. And without this technology, it would be harder for passengers to access reliable on demand transportation.



## Pricing

### Pricing Drives Reliability

Upfront pricing for rides shows Uber passengers the total fare before they request a trip, so they know exactly what they'll pay. It helps passengers make informed choices: whether to book immediately, wait for demand to ease, or consider another option and brings more clarity and predictability to travel.

#### Passengers see the fare upfront

When Uber introduced upfront pricing, it revolutionised passenger expectations by replacing the uncertainty of metered fares with clear, predictable costs. Passengers see a single fare before confirming a trip, calculated using multiple factors such as estimated time and distance, demand patterns, and anticipated traffic conditions. The system also accounts for tolls, taxes, surcharges, and fees, though wait-time fees are typically excluded from the upfront estimate. If dynamic pricing is in effect, that adjustment is reflected in the upfront fare as well.



#### RIDERS VALUE TRANSPARENCY:

In a 2024 Public First Poll, Uber riders said the most important reasons they chose to use the service were Uber's reliability and price transparency (90%). This was followed by the safety (88%) and convenience (86%) of using the service.



90%

Reliability  
and price  
transparency



88%

Safety



86%

Convenience

#### Surge pricing helps maintain reliability

Surge pricing plays a crucial role in ensuring that passengers can secure a ride even when demand is high, by helping to restore balance between available drivers and passengers. It acts as a relief valve, automatically activating when there are more trip requests than available drivers to promote service reliability and reduce wait times. Prices adjust in real time within hyperlocal areas—leveraging AI for current and forecasted demand conditions to dynamically balance demand and supply. Because prices are shown upfront in real time, passengers who are able to can wait until surge conditions recede and prices decrease. In exceptional situations such as major emergencies that might affect public safety, Uber can intervene to cap surge pricing in affected areas.

**EXAMPLE:**

Thousands of attendees exit the venue and open the Uber app to request a ride.

- › **Without Surge Pricing:** Driver supply is quickly overwhelmed. Most passengers face long waits and many are never matched with a driver.
- › **With Surge Pricing:** Fares increase temporarily, signaling to drivers nearby that there are lots of passengers around waiting for a trip and that trips are now more attractive. This motivates drivers to move toward the concert area or stay online longer. As a result:
  - Passengers with flexibility can wait slightly longer for demand to ease, after which prices typically fall.
  - Overall, supply and demand rebalance faster, reducing the likelihood of gridlock in the system.

This isn't hypothetical. In 2020, surge was banned in Las Vegas as a result of the Covid state of emergency. Thousands of passengers were left stranded at the airport and local businesses, drivers lost earnings opportunities and the city had less reliable transportation options. The problem became so acute that the Governor had to create an exemption to allow surge pricing. Long term, surge is essential not just for drivers, but for passengers and cities alike.

## Pricing dynamically

Uber's goal is to keep trips and deliveries attractive for drivers and delivery people while helping them remain affordable for customers. To help maintain this balance, Uber uses dynamic pricing that adjusts fares using aggregated, anonymized trip patterns and real-time factors—such as time, distance, demand, and destination trends—in line with local regulation. The result is to help make trips more consistently appealing to drivers and reliable for riders.

With dynamic pricing, fares can vary depending on the specific route and the time of day, helping make transportation more accessible across different locations and times. This approach accounts for the unique, hyper-local dynamics of each city—factors like road layout, traffic, and fluctuating demand patterns—that can differ greatly even within the same urban area. By tailoring prices to these variable conditions, Uber's technology strives to optimise match rates and deliver pricing that is affordable for consumers and valuable for drivers and delivery people.

Once the trip is complete, passengers are generally charged the upfront price, unless the trip changes significantly, such as by adding stops, altering the destination, or taking a longer route.

**FACT CHECK:**

Upfront prices are not personalised. Our pricing algorithms do not use information about a passenger's personal characteristics (like their gender, race, or ethnicity), nor do we factor in details like a rider's phone model, operating system, or battery level.

**CASE STUDY:**

A natural experiment from Poland demonstrates the real world impact. In 2021, Uber began testing a new feature in Warsaw that showed drivers more upfront information with their offer. These trip summaries—known as offer cards—now displayed more upfront details, such as start and end zip codes and estimated earnings. Offer cards in other Polish cities, including Krakow, did not change.

At the time, both Warsaw and Krakow used a simpler version of Uber’s pricing algorithm based mainly on time, distance, and surge. Unlike our advanced dynamic pricing algorithm, this basic model did not account for other factors that make trips more or less appealing to drivers.

**The result was telling:** after upfront information was introduced in Warsaw, the percentage of completed trip requests fell significantly compared to Krakow. Passengers were less likely to be matched with a driver, and overall service reliability declined.

This experience highlighted that transparent offer cards give drivers valuable clarity, but they work best alongside advanced pricing. While upfront information helps drivers make informed decisions, those details need to be paired with pricing that adapts to real-world conditions, using aggregated, anonymized trip patterns and real-time factors like time, distance, and demand. Today, enhanced offer cards bring these elements together: drivers can see exactly how much they will earn before accepting a trip, while dynamic pricing keeps trips affordable for passengers and reliable on our platform.



## Pricing transparency for drivers and delivery people

In most cities, offers sent to a driver or delivery person include: pickup and drop-off locations, estimated travel time, and projected earnings. Drivers or delivery people can accept or decline, with full flexibility to make decisions based on what works best for them. Rejecting offers does not restrict future access to the platform.

We know drivers and delivery people value having more information before they accept a trip or delivery, first and foremost precise information about how much they will earn for the trip or delivery. Reliability for customers has stayed high thanks to advanced pricing technology that leverages AI, adjusting for variables such as route efficiency, traffic, and local usage patterns.

Our pricing algorithms are dynamic and subject to regular internal reviews drawing from stakeholder insights. The reviews are iterative, incorporating feedback from operations teams who are most knowledgeable about local market conditions and who regularly engage with drivers and delivery people. Pricing reflects both dynamic conditions, such as temporary demand spikes, and structural considerations informed by that local insight.

For example our algorithms can identify the need for higher earnings through structural prices in popular summer destinations to meet expected demand.

Upfront Fares are not personalised—our fares algorithms do not use information about an individual driver's or delivery person's personal characteristics (like their gender, race, or ethnicity) or inputs like types of device, operating system or battery life to calculate fares.



### MACHINE LEARNING FOR PRICING

ML models and complementary rules-based techniques also improve the accuracy and efficiency of our pricing algorithms. For example, decision tree models—which break down complex decisions into a series of yes/no questions, branching at each step until reaching an outcome—can help predict whether variables like time of day, weather, and location will affect traffic conditions. Geofencing—a rules-based technique that uses GPS to draw virtual boundaries—complements this ML technique to ensure pricing accounts for the specific conditions in certain neighborhoods or venues. These tools continuously improve as more data becomes available, supporting more reliable predictions that inform pricing.

## Better Earnings Predictions

We use algorithms to help drivers and delivery people make the most of their time on the road. Probabilistic prediction is a mathematical model that learns patterns from data and then makes predictions about new unseen data. It is an essential part of many ML models. In some markets Probabilistic prediction power tools like the Earnings Heatmap which highlights for drivers when and where demand is likely to be higher. Real-time dispatch algorithms minimise downtime by efficiently matching drivers with nearby passengers and delivery people with optimised delivery routes—reducing empty miles and fuel consumption between trips and deliveries.

Drivers and delivery people also benefit from systems that adapt over time. For example, features like upfront fares and tailored promotions are all powered by algorithms that utilise historic data across the marketplace. These tools are designed to give drivers more control, flexibility, and transparency as they navigate their workday.



## Use cases specific to Uber Eats

On top of the algorithms and AI systems mentioned so far, Uber also uses specific systems to support tens of thousands of merchants on Uber Eats. From family-run restaurants to grocery chains, merchants rely on our digital infrastructure to bring their businesses online and widen their customer base. Algorithms assist them in setting up storefronts, managing menus, tracking orders, and gaining insights into customer behavior. For many businesses, this digital presence—powered by the Uber Eats Manager (our merchant facing application)—is a lifeline, especially during times when in-person foot traffic is unpredictable.

## Smarter Order Management and Delivery

### **Streamlining the flow of orders and deliveries.**

Order batching and delivery person matching algorithms optimise for efficient and high quality deliveries. By accounting for factors like order preparation time, the algorithms can help align timing between merchants and delivery people to minimise delays.

### **Helping merchants optimise promotions**

**selections**, gain critical insights into operational performance, and obtain insights into their customer bases through data-driven tools and features available on Uber Eats Manager. These insights help merchants make informed decisions and ultimately grow their businesses.

Larger businesses can utilise full API integration—a more advanced option that lets their existing point-of-sale connect directly with Uber’s algorithms for automation and scalability. Merchants can upload and update menus easily using batch import files. For those with large catalogs—such as supermarkets—our bulk upload tools use product IDs and barcodes to link images and names. In some cases, we use generative AI models to help classify and enhance menu content.

**EXAMPLE:**

**Uber Eats Manager includes an AI-powered tool that assists merchants in creating item descriptions. When a merchant adds a new item, key words are recommended as potential descriptions, which saves time in updating their menu.**

In every case, the goal of our algorithms is to automate tasks in ways that enhance user experiences, empower smarter choices, and help everyone on the platform thrive. We continue to refine these systems to help ensure that our technology delivers real value where it matters most.

**Discovery on Uber Eats:** When a customer opens the Uber Eats app, the first screen they see—our Discovery page—is shaped by a machine learning model that determines which merchants, products, and services are surfaced and prioritised for the customer.

This model considers various inputs to surface options that are relevant and optimal for the customers based on a range of contextual signals. After all, a customer who wants to order breakfast at 9 a.m. likely wants to see different options than when browsing to order dinner in the evening. To make sure customers see options that are relevant and likely of interest to them, our machine learning model considers inputs like their location, time of day, merchant availability and proximity to the customer, and their historical order preferences. Estimated delivery time is also a factor designed to prioritise options with faster delivery times. Overall, we aim to ensure that our models balance fairness, efficiency, and availability to provide customers with options that they like.

In addition to organic recommendations, Uber Eats offers restaurants and brands the opportunity to run paid promotions. It’s important that promotions reflect consumer demand and connect merchants to the right consumers. Our ad ranking algorithm considers several factors including: how closely the restaurant matches a customer’s current needs, how relevant their menu is to the customer’s past orders, and the merchant’s bid price. For example, if a customer often orders vegetarian meals, the algorithm may rank a sponsored plant-based restaurant listing over a higher-spending burger restaurant during recommendation calculations.

Customers may opt out of personal advertising, in which case sponsored placements will still appear but will not be tailored based on account details or past order history.



## Algorithms Enhancing Safety and Support

At Uber, safety is fundamental to how we design and deploy technology. Our safety algorithms are a combination of preventative tools, real-time monitoring, and post-incident interventions. While some of these tools operate behind the scenes, they are critical to maintaining trust in the platform.

**AI-powered driver and delivery person identity verification.** When a new driver or delivery person signs up to use Uber, they go through a multi-step screening process that includes document identity verification checks. But verification doesn't end at sign-up—we also require drivers and delivery people to periodically reconfirm their identity by submitting a selfie to compare against their account profile photo using facial

verification technology. This helps confirm that the person using the app is the verified account holder. This not only deters fraud but also ensures that passengers, consumers, and merchants know who they're interacting with. The facial verification technologies we use undergo an extensive review process that promotes fairness and incorporates privacy principles into the product design.

**Proactive monitoring to address service issues in real time.** For example, if our algorithms detect an unusually high number of delayed deliveries in a certain area or if a trip takes unusually long compared to our estimates, we may automatically trigger communications to passengers, customers and/or merchants, explaining the delay and offering assistance. By using algorithms to detect these types of issues, and pairing automation with human support, we are able to manage millions of daily interactions without losing sight of individual experience.

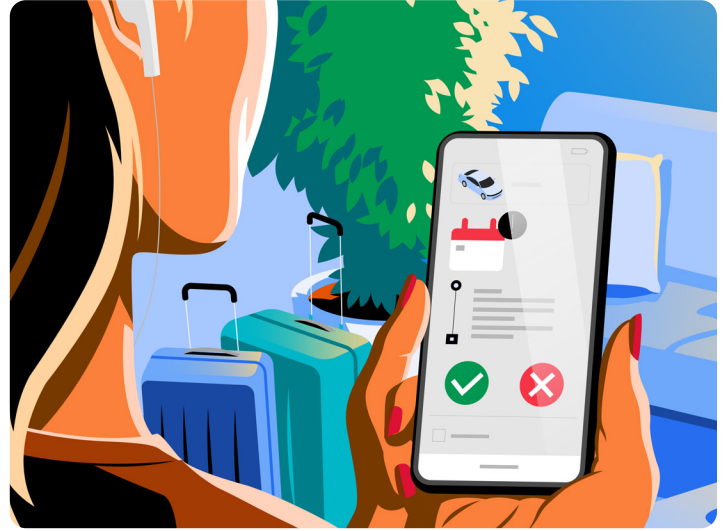
## Deactivation Transparency

When issues arise—such as reports of unsafe behavior, fraudulent activity, or other violations of our terms—drivers and delivery people can temporarily lose access to accounts or face permanent account deactivation. We recognise that deactivation can have an impact on an individual’s ability to earn money—and it’s not a decision we take lightly. Our approach centers on clarity, fairness, and the opportunity to appeal.

**Fraud detection.** We continuously monitor the platform for behaviors that suggest suspicious activity—such as repeated voucher abuse, stolen accounts or unauthorised account takeovers, or the creation of fake accounts. These signals are flagged with behavioral anomaly detection models, and depending on severity, may result in temporary account holds, additional verification requests, or escalation to human reviewers.

**Uber explains deactivations and next steps.** Uber provides in-app or email notifications that explain to drivers and delivery people the reason for their deactivation decision and the steps they can take next. This often includes the option of immediately appealing their deactivation through an in-app process. Drivers and delivery people are also able to resolve many issues on their own, like updating their licensing or insurance information or completing other steps directly and quickly.

**Importantly, humans are in the loop for these decisions.** They carefully review consumer reports, notify and collect information from drivers



and delivery people, and determine whether a deactivation is warranted. If a driver or delivery person appeals a deactivation, our team closely examines the appeal along with any further information they’ve submitted, including media.

**We work to improve our deactivation processes.** In collaboration with driver associations, unions and independent experts, we continually explore ways to make our deactivation policies even more transparent and predictable. This includes efforts to publish more detailed guidelines about what may trigger a deactivation, what evidence is reviewed, and how decisions are made. Our goal is to ensure that drivers and delivery people not only have visibility into the process, but a genuine voice in shaping it.

We know that transparency is not a one-time promise but a continuous commitment. As our systems evolve, so must our efforts to explain them clearly, invite feedback, and ensure that technology serves people—not the other way around.



# Responsible approach to AI: Governance and Privacy

## Governance

**Uber's approach to AI governance is grounded in structure, accountability, and continuous improvement**, with a strong emphasis on transparency, fairness, and human oversight throughout the lifecycle of high-impact systems. We operationalise these principles through a combination of formal governance processes and cross-functional collaboration.

Senior leadership groups spanning Tech, Legal, and Compliance teams—including the AI Law and Ethics Council—set the high-level direction, policies, and principles that guide Uber's governance approach, such as responsibilities for high risk systems.

The **Office of Responsible AI**, housed within Uber's Legal Department, along with other teams focused on AI governance, help translate these principles into practice by supporting model oversight. This includes maintaining inventories of high risk models, ensuring that model documentation including model cards are created and maintained, and coordinating model impact assessments. They also monitor global regulatory trends and update internal processes accordingly.

**Expert teams across Uber contribute to governance in specific areas.** For example, Uber's Marketplace Fairness Team evaluates fairness considerations as part of Uber's assessment

process while privacy and security specialists help ensure that appropriate safeguards are in place before deployment. Uber's AI ARK supports consistency in how models are categorised, deployed, and maintained—for instance, ensuring critical models affecting safety or earnings are subject to stricter monitoring and review.

This layered governance structure enables Uber to support innovation while embedding accountability and adaptability into our use of AI. This technology is constantly evolving, and we know our governance approach will have to evolve with them. What works today may not work forever and we're constantly iterating on ways to improve the process through the lens of transparency, fairness, and human oversight.

## Privacy, Security, and Data Protection

Privacy is critical to earning and maintaining customers' trust. Uber has implemented an industry-leading privacy program designed to protect users' data and empower them with control and transparency, including as relates to use of data for AI systems.

The program is driven by a cross-functional team of privacy and cybersecurity experts who work to embed privacy across our operations and implement appropriate data governance policies.

Through its privacy-by-design framework, we integrate privacy reviews into product development to identify and mitigate risks before launch. This framework guided our development of key algorithms and AI systems, including those relating to safety and fraud detection and prevention.

We have stringent controls on data collection, retention, and deletion, as well as clear communication explanations of our data practices. For example, our privacy notices provide clear and comprehensive descriptions of our use of data, including within key automated processes like matching pricing, and safety/fraud.

Our cybersecurity risk-management program is anchored in the ISO 27001/2 framework. The key features of this program include annual independent audits and certifications, regular vulnerability assessments and internal audits, a robust cyber incident management process that includes proactive monitoring, investigation, and remediation, and a third-party risk management that impose strict requirements and limitations on Uber's suppliers.

Finally, Uber's Public Safety Response team handles law enforcement data requests through a secure portal, ensuring legal compliance and protecting user privacy by limiting disclosures to legitimate, process-backed cases.



## Conclusion

**Transparency in AI isn't just about publishing a report.** We believe that every algorithm that affects people's lives should be understandable, fair, and accountable. This report represents just one part of our effort to meet that standard.

**As the regulatory environment in the European Union evolves, so will our systems and governance practices.** We welcome input from researchers, regulators, and the communities we serve. Because building trustworthy AI isn't just about avoiding harm—it's about earning the confidence of everyone who uses our platform.

For comments, questions, or collaboration,  
please reach us at: [eu-policy@uber.com](mailto:eu-policy@uber.com)

