Ruter#

The economics of a shared autonomous on-demand service

What are the potential commercial and societal implications of a shared autonomous on-demand service integrated into public transport?

Ruter has been exploring shared autonomous on-demand services as part of the future public transit mix, and how they can become a sustainable part of Ruter's mobility system.

December 2025

The purpose of this white paper is to explore potential economic and societal implications of shared autonomous vehicles. The material is illustrative and should not be regarded as definitive findings, formal policy positions, or exhaustive analysis. It reflects ongoing work and learning within Ruter's autonomous mobility initiatives and is offered solely to support further dialogue and learning.

Autonomous vehicles (AVs) are already on the roads in over 20 global cities—from San Francisco and Los Angeles to Oslo, Beijing, and Wuhan. As technology advances, particularly in level four autonomy, where vehicles can perform most driving tasks without human input or safety operators behind the wheel, new deployments are accelerating across regions. Leading players are securing early footholds in attractive markets, leveraging a land-and-expand approach to build scale and reinforce competitive advantage.

Industry leaders are concentrating their AV deployments on robotaxi services. These are geofenced ride-hailing operations that function within defined urban areas typically accessed through a proprietary app or in collaboration with existing ride-hailing platform, enabling a futuristic user experience.

What Ruter is striving for is a bit different—and closely aligned with the vision of *sustainable freedom of movement*. Traditional public transport (PT) takes you from *somewhere near* your location, at a time *close to* when you want to travel, to *somewhere near* your intended destination. This serves as a fundamental problem for PT in competition with other mobility modes. A robotaxi or private vehicle (PV) offers a different service experience: door-to-door, on-demand, and available 24/7. In terms of convenience and responsiveness, such modes will almost always outperform traditional PT. In terms of price, it's the other way around.

Ruter's ambition is to launch a Shared Autonomous On-demand Service as part of the public transport system. The service differs from a robotaxi in three important ways:

1. ***It is part of public transport.*** The service is integrated with Ruter's offering and booked through the existing app.
2. ***The rides are shared.*** Some call it *pooled*. Other passengers heading in the same direction may join your trip.
3. ***It is cost-efficient.*** The price per kilometer is expected to be significantly lower than ride-hailing, and close to or even lower than the cost of driving a private car today.

In small-scale deployments, the service will complement PT by extending its reach and flexibility. At scale, the service has the potential to *replace PV ownership and a large share of PV trips*, significantly reducing vehicle kilometers traveled (VKT) through sharing. The service will coexist with the PT ecosystem and enable service configurability to design a more dynamic and efficient holistic operation.

The form factor of the vehicles will resemble PVs or robotaxis, with a capacity of up to four people. Our simulations have consistently showed that vehicle occupancy rarely exceeds four passengers in an on-demand service.

For the service to be truly scalable, it must be a **financially self-sustaining service** without ongoing public subsidies. In this paper, we explore the economics behind a shared autonomous on-demand service integrated into PT, along with the strategic and societal opportunity.

Acting as market shaper, not just buyer

What happens when the service you want to offer simply doesn't exist in the market yet? Deploying a shared autonomous on-demand service at scale presents exactly this challenge.

Traditionally, Ruter issues tenders, receives bids from mature industry players, and selects the best offer based on cost, quality, and other criteria. These bids often represent joint efforts from multiple actors in the value chain, combining proven technologies and capabilities. In the case of AVs, the market is still immature. There are no established providers capable of delivering a complete fully operational, integrated, and scalable service today.

With this challenge, Ruter has invested time and resources into understanding how the value chain potentially comes together but also how value is created and distributed across it. This means breaking down the cost structure, identifying potential revenue models, and assessing how these change depending on fleet size, technology maturity, and service design. The economics of shared autonomous mobility are fundamentally different from traditional PT, but also from ride-hailing and PVs.

Exhibit 1

Stakeholder interaction and financial flows in the business model

Exemplary cashflows between stakeholders

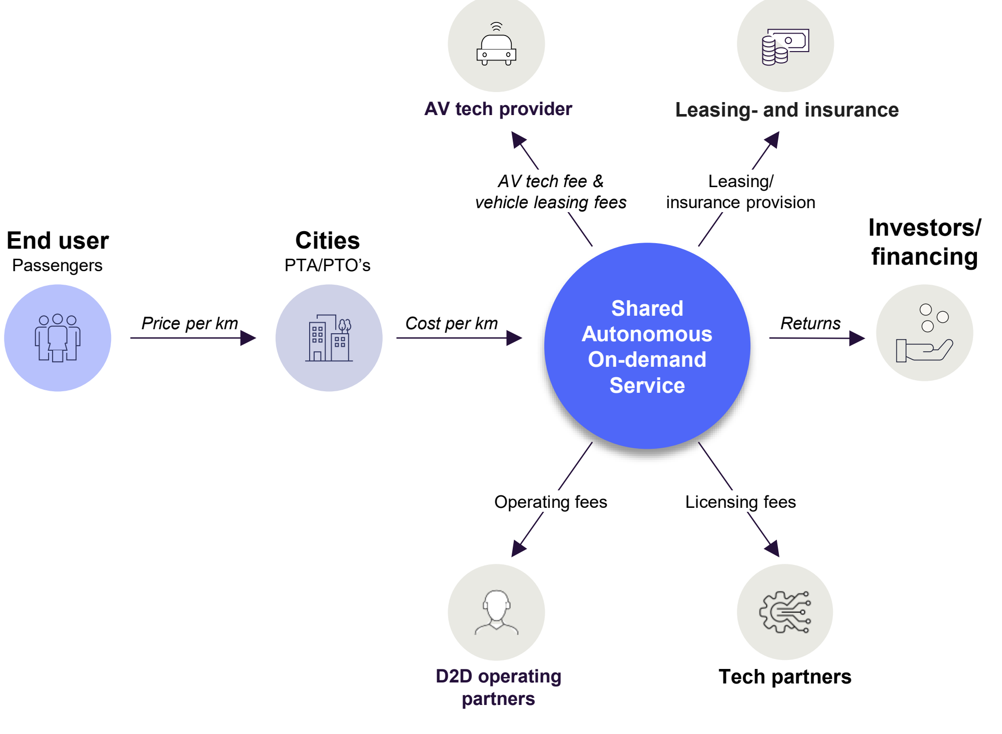

Exhibit 1 illustrates the high-level business model and cashflows between stakeholders in a shared autonomous on-demand service. End users pay for using the service to the PTA or PTO—in this case Ruter—which holds a contract with an entity responsible for delivering and operating the service. The orchestration of partnerships behind this entity is left to the players in the value chain to establish. **In this paper, we adopt the perspective of the contract entity**, and its underlying economics.

The first deployment: A balancing act

The service must either be more affordable or more convenient than owning and using a PV. Cost and utility serves as the balancing act for such a service. To determine pricing for a service we need to understand the full cost structure across physical assets, technology, operations, and both city-level and central organizational resources. For an initial deployment scenario in Oslo, we use an illustrative fleet size of 500 vehicles to examine potential implications.

Exhibit 2

What it costs to operate a fleet of 500 vehicles over 10 years across Ruter's service area

Operating costs, in million €

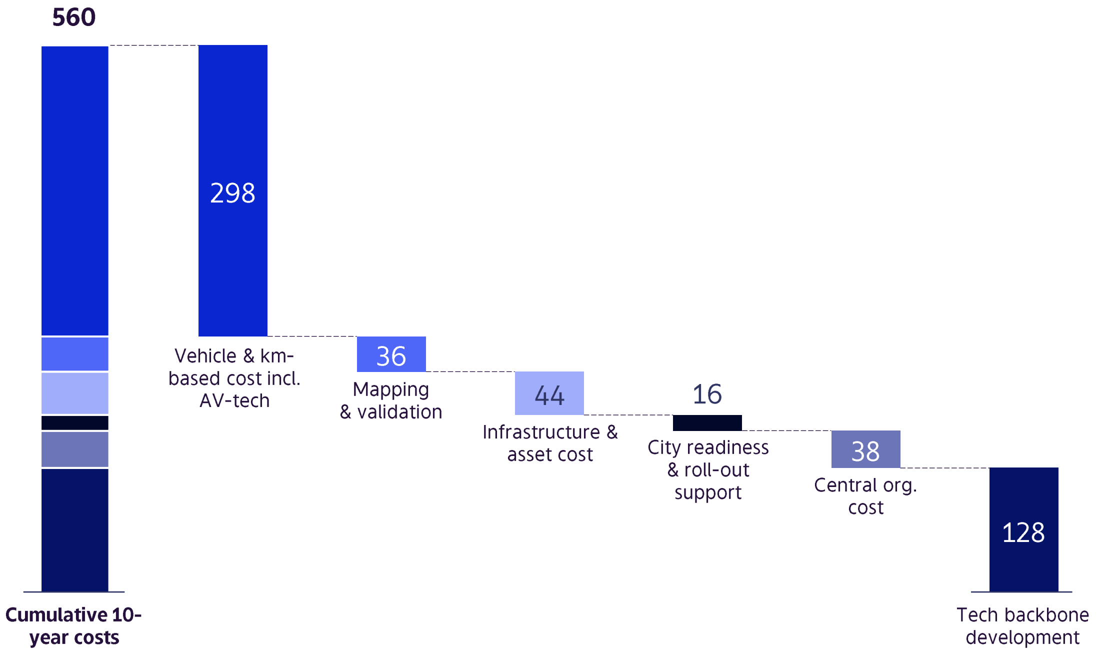

To assess the financial viability of a service with a 500-vehicle deployment in Oslo, it is essential to understand the full cost structure. Following is a description of the cost categories in Exhibit 2.

- ***Vehicle and kilometer-based cost.*** This category captures all recurring costs directly tied to the operation of the fleet. These include vehicle leasing, charging (electricity), regular maintenance, and system-level checks. Insurance, import duties, and on-road assistance also fall within this category. Operational resources such as vehicle control center operators, fleet and hub managers, and customer support are included, as are costs related to integration and replacement of AV systems. Cleaning and cybersecurity maintenance represent ongoing expenses. It also includes licensing fees paid to AV-stack and software providers, calculated on a per-vehicle basis.

- ***Mapping and validation cost.*** Safe AV-operations depends on high-precision, regularly updated maps of the vehicle's operating area (at least for now). This cost category includes leasing costs for pre-mapping and maintenance mapping vehicles, as well as labor for both drivers and engineers engaged in mapping and validation activities. These costs are particularly relevant during early deployments or when expanding into new geographic area.

- ***Infrastructure and asset cost.*** This category includes costs related to enabling both the physical and digital foundations of the service, with a focus on leased or rented assets. It covers the cost of AV testing and validation equipment, test vehicles used for verifying operational readiness within the defined Operational Design Domain (ODD), and the setup and operation of cloud infrastructure and data centers. It also includes leasing of roadside assistance vehicles and rented real estate such as depots, office space, and operational hubs. Office equipment and core tooling for operational and technical staff are also included, assuming a leasing or service-based model rather than capital investment.

- ***City readiness and roll-out support.*** This category combines the costs related to preparing locally for launch with the resources required to manage local operations during scale-up. It includes system integration work, local API configurations, and ensuring technical readiness in new geographies. Staffing costs span both internal roles and external contractors involved in roll-out planning and execution, as well as city-level leadership positions such as city managers and heads of operations and customer service. It also includes support roles such as data managers dedicated to both launch and extension phases. City-specific marketing, communication, and community engagement efforts are also captured here to ensure effective onboarding and adoption.

- ***Tech backbone.*** This category includes the internal software development required to support orchestration, system integration, and real-time operations across the service network. It covers in-house development of platforms and tools related to dispatching, monitoring, data flows, and service optimization. In our calculations, we assume that the service leverages Ruter's existing digital infrastructure, providing instant customer access through Ruter's app.

- ***Central organizational costs.*** This category includes the core support functions required to operate and scale the overall service. It includes leadership roles such as the CEO and functional heads (e.g., finance, technology, operations, people), as well as departments responsible for legal, procurement, HR, recruiting, and business development. It covers resources supporting city agreements and negotiations, and general administrative functions such as accounting, compliance, and internal training.

The demand for the shared autonomous on-demand service remains partially a black box. What is the demand for a service that does not yet exist, and how sensitive is it to price, assuming the service meets expected quality? To answer this, one can rely on benchmarking with comparable services such as robotaxis, ride-hailing, demand-responsive transport (DRT), and other similar offerings. In addition, simulation tools and strategic assumptions are used. Combined, we can build and estimate an adoption curve and pricing model with an associated price. Together, we can estimate revenue.

Exhibit 3

Estimated 10Y revenue development from preparation to steady state operations

Yearly revenue, in million €

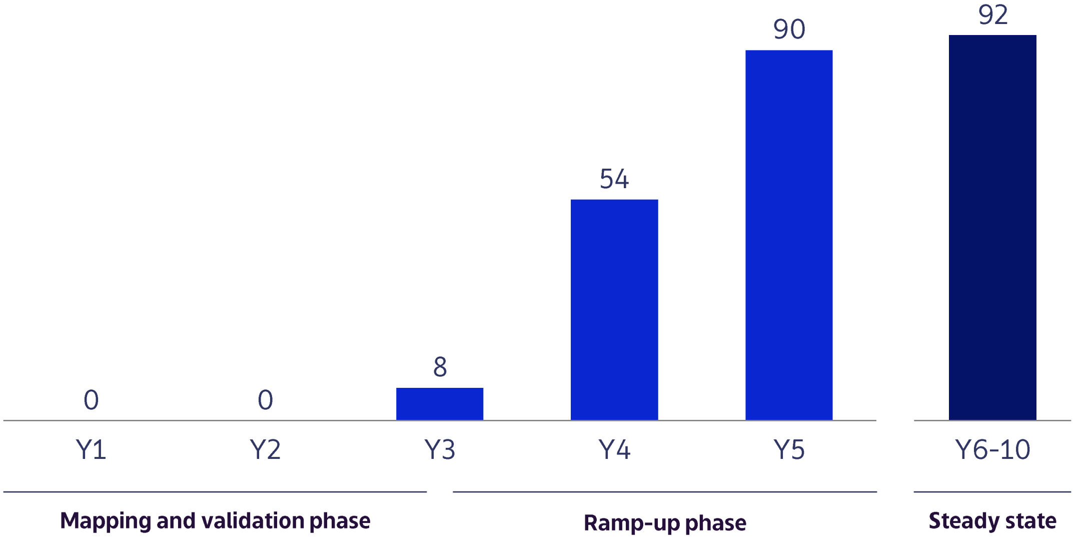

The first two and a half years are dedicated to mapping and validating the ODD, thus no revenue is expected in this period. Deployment begins gradually, with 100 vehicles added each quarter over the first five quarters. In steady state, utilization is modeled at 73 percent and occupancy at 1.9 passengers per trip. The vehicles travel at an average speed of 22.68 km/h. Each new deployment of vehicles uses five quarters to reach steady state. The price is set to 1.0 EUR per PKT, and yearly steady state revenue is estimated at approximately EUR 92 million yearly, or about EUR 183,000 per vehicle. On an hourly basis, this equates to roughly EUR 31, assuming 16 operating hours per vehicle per day.

The per-PKT price model reflects an average price rather than a detailed customer-facing pricing structure. The actual fare model, including potential variations by time of day, trip length, occupancy, or integration with PT tickets and revenue streams from configuring the vehicle such as ads, is not accounted for. Instead, the focus is on understanding the service's long-term economic viability at a representative revenue level. The price of EUR 1.0 per passenger-kilometer traveled (used in Exhibit 3) is derived from a cost-based approach, leaving a profit margin of 8% over the ten-year period.

Compared to the cost of driving a PV, a service with a fleet of 500, under a financially sustainable model over a 10-year period—would initially be *slightly* more expensive under these assumptions. As illustrated in Exhibit 4, the price of PVs cost EUR ~0.7 per PKT. This assumes the PV is purchased with 100% equity financing and assigns no monetary value to the time saved or the strain and responsibilities avoided by not owning a PV.

Exhibit 4

How pricing compares to other urban transport modes

Price per PKT, in €/PKT (Oslo)

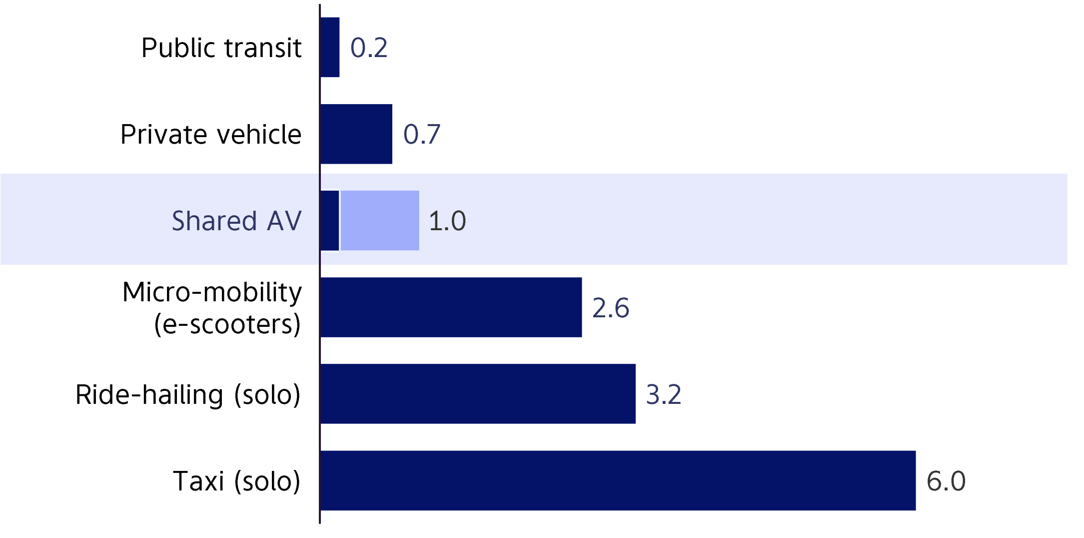

Long term, the price per PKT of the service has the potential to fall below that of PVs. The 10-year average cost per PKT in a 500-vehicle case is EUR 0.92, inflated by front-loaded investments such as mapping and building the tech backbone. By year 10, the isolated cost stabilizes at EUR 0.63 underscoring the model's long-term efficiency, with an attractive margin of 37%. **By 2040, costs could even fall to as low as EUR 0.2 per PKT**.

The key uncertainty is the pace of this decline—whether it unfolds gradually through technology maturity, operational efficiency, and competition, or accelerates through technological breakthroughs. We expect **costs to decrease in three principal areas**:

- ***Efficiency in operations***. Substantial efficiency gains have already been realized through automation of sensor calibration, in-depot driving, and the utilization of depot staff. Additional improvements are expected from the automation of cleaning, docking, and charging processes. Continued advancements in AV systems will further reduce the need for human intervention, lowering requirements for control center operators, fleet managers, roadside assistance, customer support, and system-level checks, while also diminishing the need for maintenance mapping and validation.

- ***Vehicles and sensors.*** Current fleets are largely retrofitted, a practice that is both costly and suboptimal. The introduction of purpose-built vehicles produced on dedicated production lines is expected to lower costs and extend vehicle lifetimes. Hybrid production models, where existing PV lines are adapted for sensor integration, have also emerged and are expected to launch shortly. Meanwhile, sensor suite costs are declining, with improved durability and performance, and ongoing enhancements in AV systems are expected to reduce the overall number of sensors required.

- ***Transitioning to end-to-end systems.*** End-to-end architectures has the potential to integrate perception, prediction, and planning into a single learned controller, thereby reducing engineering, integration, and maintenance needs. They also decrease reliance on expensive mapping and validation processes, shifting progress to a data-driven paradigm where improvements scale with data availability. This makes development cycles faster, more adaptive, and increasingly cost-efficient.

Short term, the most **effective lever for reducing cost per PKT at launch is scale**. In this way, one can leverage one-time investments, improve the economics of licensing AV technology, vehicles and the sensor stack. As illustrated in Exhibit 5, expanding the fleet size fundamentally changes the economics of the service.

Exhibit 5

How scale could affect the 10Y PKT cost in an initial launch

Cost per PKT for a fleet of 500 vs. 3,000 vehicles, in €/PKT

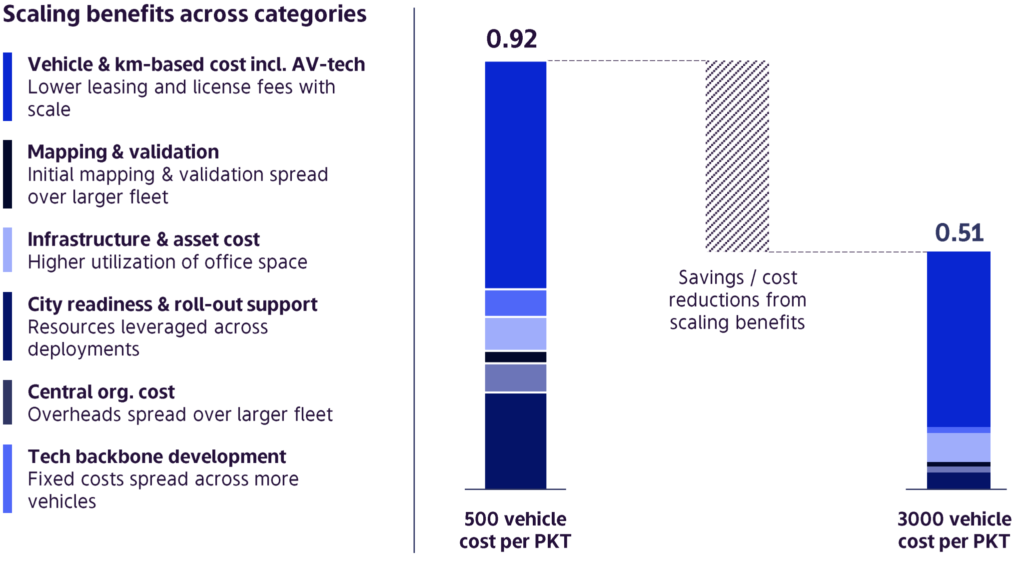

In Exhibit 5, vehicles are deployed gradually, adding 300 vehicles to the fleet over 10 quarters.

Scale has the potential to lower the overall cost per PKT, but more importantly reshapes what is possible in terms of end-customer pricing. Over time, sufficient scale allows the service to compete with PVs directly on price. In the scenario of 3,000 vehicles, the vehicles are deployed at a rate of 300 per quarter. The 10-year average cost per PKT is 0.51, in year ten the isolated cost per PKT is 0.41, leaving potential for offering an average end-customer price below that of PVs cost today (EUR 0.7).

While Exhibit 5. provides useful insights, there are multiple important caveats to consider;

- **AV technology availability and pricing:** This analysis assumes that AV-technology suppliers will license their systems at a fixed fee. In practice, today's dominant models revolve around vertically integrated robotaxi operations, where suppliers retain full control or partner with ride-hailing platforms. The pricing models of AV technology remains immature, with few competitors in the market. Limited competition could drive prices higher than assumed.

- **Feasibility of large-scale deployment:** No operator has yet reached a fleet of 3,000 AVs. A sustained deployment of 300 vehicles per quarter has never been tested at this scale. Supply chain challenges, technical bottlenecks, and organizational limits could all slow the process, making the assumed pace difficult to achieve.

- **Capital intensity:** Achieving a 3,000-vehicle fleet requires heavy investment not only in the vehicles themselves but also in supporting infrastructure such as charging networks, depots, and maintenance facilities. Securing this level of financing may prove difficult, particularly in the early stages before commercial viability is demonstrated.

- **Public acceptance and demand:** Scaling also depends on customer willingness to adopt autonomous technology. Even if the technology is available and vehicles are deployed, uptake could be slower than expected if trust, safety perceptions, or user experience fall short. Without strong ridership demand, the economics might be weakened.

Still, the potential is clear; A financially sustainable and scalable service that offers much of the same utility as PVs at competitive prices is highly attractive. Beyond the economics, there is also a societal dimension.

Shared autonomy: A potential driver of societal impact

By enabling shared rides, a shared autonomous on-demand service have the potential to reduce VKT. This serves as a driver for societal impact. The impact of small deployments will be incremental, but at scale the effect could become transformative and deliver societal benefits across multiple dimensions.

By gradually *replacing private vehicle ownership* and a larger share of PV trips, one can start to reduce the number of vehicles on the road. When we consider how much of today's PV PKT could be served across different fleet sizes, the impact potential becomes clearer.

Exhibit 6

The service has the potential to replace PKTs provided by PVs today

Share of total PV PKT within Ruter's service area that the service can serve, in percent

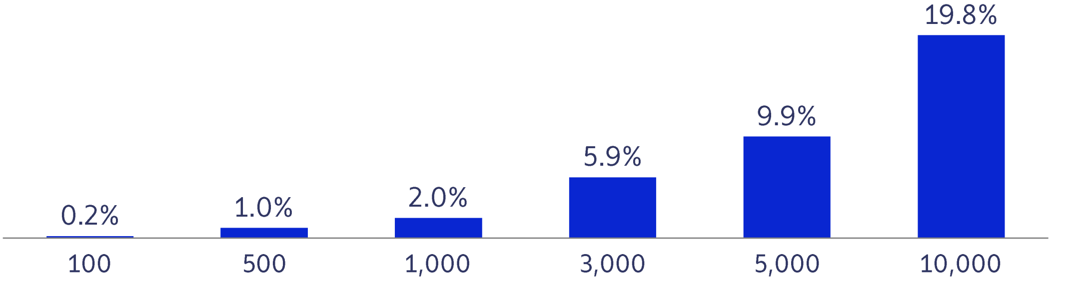

In Exhibit 6 we have used the same assumptions as in our cost and revenue estimations; steady state occupancy of 1.9, utilization of 73% and 16 operational hours a day per vehicle traveling an average speed of 22.68 km/h. In Exhibit 7 we illustrate the potential avoided VKT reduction and PV replacement. In between trips we assume that the vehicle stands still 50% of the time.

Exhibit 7

The service can potentially replace PVs and avoid VKT

Avoided yearly VKT across fleet sizes, in million VKT

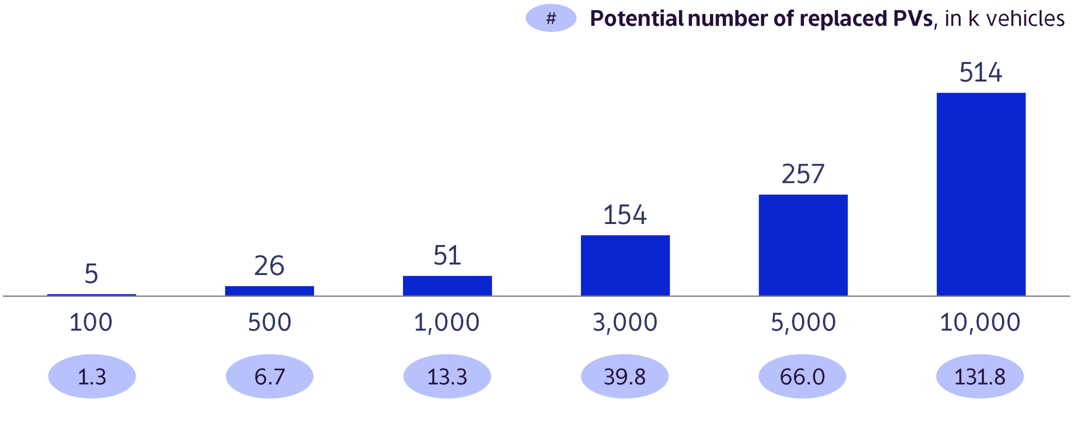

The long-term societal potential is clear if sufficient adoption and scale are achieved. A fleet of 10,000 vehicles could potentially serve 19.8% of all PV trips in Ruter's service area and potentially replacing 131.8k PVs, while driving 514 million kilometers less (on a yearly basis).

Such a service could deliver a range of societal benefits:

- **Lower emissions:** Less VKT and less congestion equals lower emissions. A shared fleet also enables higher vehicle utilization, and purpose-built vehicles support longer lifetimes, all improving air quality in dense urban environments.

- **More livable cities:** Fewer cars on the road free up valuable space that can be repurposed for green areas and active mobility, contributing to more livable urban spaces.

- **Access and affordability:** The service can increase access to affordable transport for groups underserved by traditional PV ownership, helping reduce inequality in urban transport systems.

- **Improved road safety:** Advanced AV systems today cause up to 90% fewer serious crashes than human drivers, greatly reducing fatalities and injuries. By removing human error, AVs can make urban transport systems far safer.

- **Economic resilience:** Widespread service adoption can strengthen economic resilience. By reducing congestion and enabling more efficient mobility, the service can improve labor market accessibility, shorten travel times, and support productivity gains in cities.

Thus, shared autonomy represent more than a technological breakthrough—they mark a potential turning point in how cities can organize mobility. As shown throughout this paper, the economics of scaling such systems are challenging but achievable. Early phases will carry high costs, but scaling effects can bring the price per PKT to a competitive level. Realistic pricing and coordinated efforts across stakeholders are essential to move from pilot operations to deployable services.

The broader benefits also strengthen the case for deployment. A future where a shared autonomous on-demand service reach scale promises cleaner air, safer streets, more efficient land use, and expanded mobility access. The societal dividends of lower emissions, better affordability, enhanced inclusion, and economic resilience, all highlight the transformative potential of this technology. The challenge now is not whether the service can deliver value , but how quickly it could be deployed.

About the analysis and sources

This paper builds on Ruter's internal analyses, proprietary financial models, and data from ongoing autonomous vehicle pilots and research collaborations. Market figures and cost assumptions have been validated against international benchmarks, industry reports, and input from leading AV technology providers. While the report does not include in-text references, all underlying sources have been cross-checked for relevance, accuracy, and consistency.

The field of shared autonomous mobility is developing rapidly. New technologies, regulations, and market dynamics may shift the assumptions and findings presented in this paper. The analysis reflects the best available data and understanding at the time of publication, and Ruter reserves the right to update its perspectives as the industry evolves.

Ruter also welcomes researchers, analysts, experts, and suppliers within the industry to share and discuss their perspectives with us to further advance the understanding of shared autonomous mobility.

This paper is a collaborative effort by Endre Angelvik, Vibeke Harlem, Bjørn Oscar Unander, Eivind Graff Christiansen, and Carl Isac Jerner, based on Ruter's continued work in autonomous mobility.

Endre Angelvik is Executive Vice President of Radical Innovation at Ruter, **Vibeke Harlem** is Director of Radical Innovation at Ruter, **Eivind Graff Christiansen** is Chief Marketing and Business Development Officer at Ruter, and **Bjørn Oscar Unander** is Head of Business Development at Ruter. **Carl Isac Jerner** is a consultant and advisor from Knowit Management Consulting.