AVs: The Idea That Was A Long Time Coming

We have been spending a lot of time thinking about the robotaxi question over the past few months. It started as a focused piece of work on a couple of Chinese companies but it kept pulling us deeper, and what we found was interesting enough that I wanted to write it up properly.

This is not a stock pitch. It is an attempt to explain what is going on in this space and why commercialisation took so much longer than expected. Towards the end, we discuss the opportunities that exist in the space, but it is important to first walk through the thinking.

Executive Summary

- Autonomous vehicles took decades longer than expected to become a commercial reality because the earlier rule-based approach simply could not handle the unpredictability of the real world.
- China's domestic supply chain has driven robotaxi build costs down to US$35,000-40,000 per unit, compared to US$160,000 for Waymo. This gap is structural.
- Regulation has shifted meaningfully in both the US and China in 2025, finally allowing real commercial deployment to begin.
- Removing the driver does not mean all the economics flow to one entity. The margin gets shared across OEMs, sensor companies, chip makers, and distribution platforms.
- We think ride hailing platforms and compute chip makers are the most interesting parts of the value chain from an investment perspective.

Why Did Commercialisation Take So Long?

If you had asked anyone in 2016 where autonomous vehicles would be by 2026, almost nobody would have guessed that we would still be “figuring it out.”

GM paid a billion dollars for a startup called Cruise in 2016 and predicted commercial robotaxi services by 2019. Elon Musk had been saying full self-driving was one to two years away since roughly 2015. The Alphabet self-driving car project (Waymo) had logged 300,000 accident-free miles on California public roads by 2012, yet the driverless paid hailing service began only in 2020. Despite it being six years since then, they still have only about 2,500 cars on the road.

The reasonable assumption was that robotaxis would be rolled out much earlier than they actually have. The reason for the delay is worth understanding properly because it shapes a lot of what we are focusing on today as investors.

The early approach to autonomous driving was rule-based. Engineers wrote explicit instructions for every scenario they could think of – stop if the light is red, brake if the car ahead brakes, etc. This worked well under controlled conditions, but broke down constantly in the real world. This was something like the scientific experiments that work only “in a vacuum.”

In 2016, a Tesla on Autopilot drove into a truck crossing an intersection in Florida. The truck's white side blended against the bright sky and the forward obstacle detection saw open road instead of a barrier.

In 2018, an Uber test vehicle in Arizona failed to correctly classify a pedestrian crossing outside a crosswalk at night, switching between "unknown object," "bicycle," and "vehicle" classifications until it was too late. Both incidents were not about bad sensors, they were about systems that had been given specific rules for driving but had no clear way to apply them to the world they actually encountered.

Engineers had a name for this problem: the long tail of edge cases. For every thousand normal miles, the real world would produce one scenario that nobody had written a rule for. It could be absolutely anything: a construction worker waving you through a red light or a child's ball rolling into the road just ahead of the child.

At the kind of scale these companies were aiming for, the long tail was not occasional, it was actually a very frequent occurrence, and they could not write enough rules to cover them all.

What Changed?

A shift was made from rule-based systems to learned ones. Instead of giving vehicles instructions, engineers started feeding models with enormous amounts of driving data and letting them learn what good driving looks like.

The vehicles were not told what to do but were instead shown hundreds of millions of driving scenarios and trained to recognise patterns. This sounds straightforward, but the implications were significant. It meant that the technology that worked best was the technology with the most data.

This is where I would like to talk a little about China. Not only is the country quick to adopt and adapt to new technologies, but Chinese cities are also exceptionally valuable training environments with dense traffic, unpredictable behaviour, and complex interactions between vehicles, pedestrians, and cyclists. Guangzhou alone probably generates more genuine edge-case scenarios per mile than most cities in the world.

Additionally, the Chinese supply chain makes collecting that data at scale economically viable in a way it is not in the US or in any other country. Hesai, a Shanghai-based LiDAR manufacturer, had brought sensor costs down to a few hundred dollars per unit by 2024. Waymo, Alphabet’s robotaxi company in the US, had been buying the same component for US$75,000 per unit a decade earlier.

Cheap LiDAR sensors and other components mean that you could install them on more vehicles, which translates into more data and, thus, better models.

Today, in 2026, looking carefully at both ecosystems, one can conclude that there is no meaningful gap in software quality between Waymo and the leading Chinese robotaxi players. Chinese companies like WeRide, Pony.ai, and Baidu's Apollo have all converged on the same approach.

However, while the software landscape has become roughly equal, the hardware landscape is still worlds apart.

The Cost Gap

In April 2025, Pony.ai launched its seventh-generation robotaxi with four LiDAR sensors, fourteen cameras and a domain controller, i.e. a full autonomous driving kit. This cost US$38,000. Baidu's equivalent came in at US$35,000. WeRide's at US$40,000.

Waymo's sixth-generation vehicle cost approximately US$160,000 that same year.

There is a four-to-one cost gap between the most competitive Chinese vehicles and the leading US vehicle in the same year with comparable technology, certainly not a marginal difference. In any business where one deploys capital at scale and waits years for payback, a four-to-one difference in unit cost is extremely meaningful.

This gap comes from the supply chain. China has a domestic ecosystem for every component that matters in a robotaxi: LiDAR from Hesai and RoboSense, compute chips from Horizon Robotics, cameras from Sunny Optical, and the rest. These are companies manufacturing at volumes that western suppliers are not, and with volume come cost efficiencies.

Pony.ai's progression illustrates this well. Its fifth-generation vehicle in 2020 cost the equivalent of US$140,000, sixth generation in 2022 cost US$115,000, seventh generation in 2025 costs US$38,000.

The projections going forward do not show this gap closing. By 2030, estimates put Chinese vehicle costs somewhere between US$22,500 and US$30,000 while Waymo is expected to reach around US$120,000. Still three to four times more expensive. The gap is structural, built on industrial capacity that takes years to develop and cannot be replicated quickly.

The Unit Economics

All this brings us to the question: why are we so keen on having robotaxis in the world? Yes, they look extremely cool, but there is another reason.

To understand why the robotaxi business is so compelling, start with how ride-hailing works today. When you book a US$20 trip through Uber in the US, US$14 goes to the driver and US$6 goes to Uber.

That 30% “take rate” is essentially Uber's entire business. DiDi and Grab run similar models, keeping around 20% in China and Southeast Asia respectively. The driver captures the vast majority of every fare, which is exactly why removing the driver changes everything.

Human drivers work 8 to 10 hours a day. They sleep, eat, have families and mortgages, and they choose which trips to take. A robotaxi does none of that. It runs continuously, through the night, on public holidays, and it does not have a family to look after.

Uber reported that the Waymo robotaxis on its platform in Austin were 99% busier than human drivers in the same area. This difference in utilisation is not just about revenue. In addition to zero labour costs, robotaxis allow fixed costs like depreciation and insurance to be spread across far more trips, making each individual ride cheaper to deliver.

Industry estimates suggest that taking the driver out of the picture allows for gross margins of 63% per ride in China and 70-75% in the US and Western Europe where labour costs are higher. As mentioned above, DiDi and Uber today operate at around 20%-30% gross margin per ride. This gap in gross margin, with robotaxis in the picture, is significant.

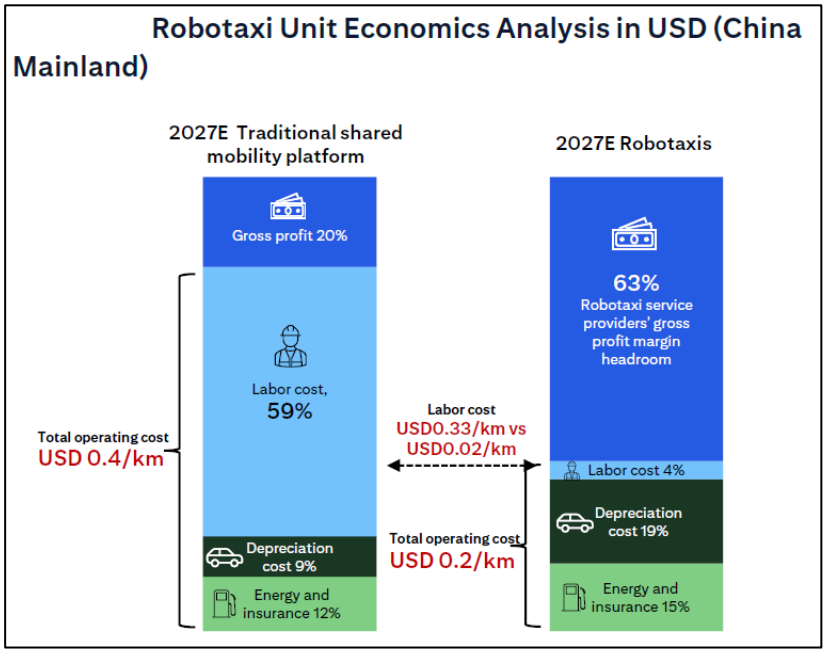

Source: Citibank

But if these numbers sound too good to be true, that is because the full picture is more complicated. The 63% gross margin does not all flow to the company running the app. Think about everything that goes into a single robotaxi ride: the vehicle could be built by Toyota, the sensors came from a company like Hesai, the chips from Horizon or Nvidia, another company is managing remote supervision, and the vehicles themselves are owned by another company or a combination of the players mentioned above.

Each layer of that stack captures a share of the economics, and the question of who captures the most value over time is still open.

That said, our conversations with people in the industry have led us to believe that ride hailing platforms will have a real opportunity to raise their take rates on autonomous rides, maybe not from 20% to 60%, but comfortably to 30% or beyond.

The combination of higher take rates and higher volumes can be quite disruptive (in a good way) to the current business models of companies like DiDi and Uber.

Regulation has progressed

Another thing that has genuinely changed is the regulatory environment. For a long time, regulation was the invisible ceiling on this whole industry. The technology was improving but regulators were not letting the cars onto the road.

What changed is that both the US and China moved toward a more enabling posture at around the same time. Around 40 US states had completed or were drafting robotaxi legislation by end-2025. Waymo ended the year doing over 400,000 rides per week, with a target of 1 million by end-2026.

In China, the pace was equally striking. In the first half of 2025 alone, China issued 9 national-level, 21 provincial-level, and 33 municipal-level policy documents on autonomous driving. Pony.ai holds fully driverless permits across all four tier-1 cities. WeRide is operating commercially in China, the UAE, Singapore, Saudi Arabia, and Dubai.

But here is the thing. Regulation, an enabler today, could become the biggest obstacle over time.

The Regulatory Hurdle

Right now, robotaxis are additive in the sense that they are not replacing humans. They may cover hours that human drivers avoid, or take on the additional supply to meet demand and regulators can look at that and feel comfortable.

That will change as fleets scale and robotaxis start replacing human drivers rather than supplementing them. There are approximately 3 million for-hire drivers in the US, roughly 2% of the working population. In China, the number is around 8 million. These are people with mortgages and families and no obvious employment alternative.

The usual counter-argument is that new technology creates new jobs. But the industry's own projections expect one just teleoperator managing 50 to 100 robotaxis at maturity.

Regulators, who are today competing to attract and support AV companies, will eventually have constituents asking why their livelihoods were handed to a technology company. That political pressure is not here yet but it is bound to come, and be a major issue when it does.

On the other hand, the changes that people like you and me can expect are:

- Taxi ride costs come down significantly over time due to increased supply.
- Possibly a future where our personal cars operate as robotaxis while we are at work or travelling.
- And maybe a future where we do not own personal cars altogether. Or perhaps one in which we own one car, instead of two or three.

That said, if regulation does not completely derail the pace of robotaxi deployment, we can expect the incoming volumes to be incredible.

The Volumes

To put some numbers around this, there are currently 3 to 3.5 million for-hire vehicles in the US and approximately 270 million personal passenger vehicles.

Estimates put the US robotaxi fleet at around 250,000 vehicles by 2030. That is less than 10% of the existing for-hire market, let alone any penetration into personal vehicle market. We do not think that is a particularly conservative estimate, but it does illustrate how early we are.

The more interesting long-term question is what happens beyond 2030. Robotaxis do not just compete with for-hire drivers. Over time, they compete with car ownership itself. As reliable, cheap, always-available rides become the norm, the calculus around owning a personal vehicle, especially in cities, will start to shift.

The real TAM here is not 3 million for-hire vehicles. It is a slice of 270 million personal cars. Even modest penetration into that market produces numbers that are drastically larger than anything the near-term models are pricing in.

That is what makes us genuinely excited about this space over a longer horizon. The near-term story is interesting. But the long-term story, if even a fraction of it plays out, is enormously exciting.

And Finally, Where Do the Opportunities Lie?

We have spent a lot of time thinking about this and we keep coming back to the same conclusions. Let us walk through them.

1. The software companies are genuinely impressive. These are the teams that made autonomous driving actually work, that solved the edge case problem, that built the models capable of handling real-world complexity. There is real craft in what they do.

   But when we look at the investment case, a few things give us pause. The space is getting crowded. A few years ago, there were a handful of serious players. Now there are a lot of them, across China and the US, and the differentiation between them is not always obvious from the outside.

   We are not saying the deep autonomous driving expertise these companies have built becomes worthless, but we are saying that the moat around pure software capability feels narrower than it did.

2. We are quite optimistic about the legacy ride hailing platforms: Uber in the US and DiDi in China. Our reasoning is pretty simple: most of the robotaxi and software companies do not want to handle distribution themselves.

   Building a consumer-facing ride hailing platform from scratch is expensive and slow, so these new entrants are going to the platforms. Waymo, Pony.ai, and WeRide have partnered with Uber in multiple markets.

   The platforms are effectively becoming the storefront for the whole industry, with almost none of the capital intensity of actually owning and operating the vehicles.

3. The chip layer sits below everything else in the autonomous driving stack. Every software company and every OEM needs capable automotive-grade AI chips, and the barriers to entry for this business are meaningful high.

   Horizon Robotics is the name we keep coming back to in China. Even in a world where Nvidia and Qualcomm dominate globally, there is significant room for a domestic Chinese player here.

   There is a genuine and sustained push for Chinese companies to use Chinese-made chips. That is a structural tailwind that exists independent of whether or not Horizon makes the best chip in the world.

   Sometimes being the right chip in the right place at the right time matters as much as being the best chip.

Closing

The reason autonomous vehicles took so much longer to be commercialised than anyone expected was that the initial approach was wrong, the regulatory environment was hostile for a long time, and the economics only started making sense once hardware costs fell to the level that the Chinese supply chain reached.

Now that this has all changed and the industry has moved forward, the question is: How would the economics actually get divided across the value chain as it scales, and what are the opportunities that arise for savvy investors? That question does not have a clean answer yet, but it is the right question for all of us to be working on.

Disclaimer

This material is not intended as an offer or solicitation for the purchase or sale of any financial instrument. Information has been obtained from sources believed to be reliable. However, neither its accuracy and completeness, nor the opinions based thereon are guaranteed. Opinions and estimates constitute our judgment as of the date of this material and are subject to change without notice. Past performance is not indicative of future results. This information is directed at accredited investors and institutional investors only