# The Ride-Hailing Startup Launch Roadmap

A working document for anyone standing up a local ride-hailing, private-hire or minicab operation from zero. It is written for the person who will actually be in the control room at 23:00 on a Friday, not for an investor deck.

## Before you read: what this guide does not contain

It contains no benchmark numbers. No driver acquisition cost, no churn rate, no supply-to-demand ratio, no "what the first 1,000 rides cost".

That is deliberate. TaxiDex is a young product with essentially one operator running on it. We do not have a body of launch data across many cities, and we are not going to print numbers we cannot stand behind - an invented benchmark is worse than no benchmark, because you would plan against it.

What you get instead is the arithmetic. Every number that matters here is one you calculate from your own city, your own fares and your own costs, and we show you exactly how to work it out. That is more useful anyway: a driver acquisition cost from Nairobi tells you nothing about Bradford, and a supply ratio from a dense city centre tells you nothing about a market town.

The same applies to regulation. Licensing rules vary by country, by city and sometimes by council, and anyone who gives you one answer for all of them is guessing. Where local specifics matter, this guide tells you which question to ask and who to ask it of, rather than pretending to know your city.

One more thing worth saying at the top, because it is the single most common mistake: almost all of the hard part is operational, not software. The technology to dispatch a small fleet is a solved, commodity problem. Getting licensed, getting drivers, and getting enough repeatable demand to keep those drivers earning is not solved, and no platform solves it for you.

## Phase 0 - Licensing and regulatory approval

Do this first. Not first among equals - first. It is the long pole, it is largely outside your control, and it is the only phase that can stop your launch outright rather than just make it worse.

Most people build the app, recruit a few drivers, then discover the licence takes months. Reverse that order and everything else gets easier, because licensing runs in the background while you do the rest.

In England and Wales, private hire is licensed in three separate parts by your local licensing authority: an operator licence for the business that takes bookings, a vehicle licence for each car, and a driver licence for each driver. All three are separate applications with separate timelines. Other countries and cities structure it differently - some license the operator only, some run a permit cap, some treat app-based work as a distinct category - so treat the three-part shape as an illustration of the problem, not as your answer.

Two structural points hold almost everywhere and are worth internalising:

- You generally cannot accept bookings before the operator licence is granted. Not "soft launch", not "friends and family for testing". Taking a booking is the regulated act. This is what turns a slipped licence into a slipped launch.
- Driver licensing is the true long pole, because it is per-person, sequential, and depends on other people's calendars - background checks, medicals, and in some authorities a local knowledge test. You cannot parallelise your way out of it. This is why the fastest route to supply is recruiting drivers who are already licensed.

Ask your licensing authority these, in writing, and get the answers before you commit money:

- ☐ What licences does my business need, and can I accept a single booking before they are granted?
- ☐ What is the current realistic processing time for operator, vehicle and driver licences - not the target, the actual queue right now?
- ☐ Is there a cap or moratorium on new vehicle or driver licences?
- ☐ What are the vehicle requirements - age limit, emissions standard, plating, signage, testing frequency, accessibility obligations?
- ☐ What are the driver requirements - background check, medical, knowledge or safeguarding test, language requirement, right to work?
- ☐ What insurance class is required, and do I need a policy in place before or after licensing?
- ☐ What are the record-keeping obligations - booking records, driver and vehicle records, retention periods, and what an inspection looks like?
- ☐ Can drivers licensed by a neighbouring authority work for me, and can I take bookings that start outside my area?

That last one catches people constantly. Cross-border rules decide whether your coverage map is legal, and they are rarely intuitive.

Alongside the licence, get the boring compliance in place: hire-and-reward insurance (ordinary motor insurance does not cover carrying passengers for payment, and a driver who assumes it does is an uninsured crash waiting to happen), data protection registration and a privacy policy, and whatever tax registration applies to you. Also settle now, in writing, whether your drivers are contractors or employees under your local law. That classification affects your pricing, your costs and your legal exposure, and it has been litigated hard in several markets. Take actual advice on it - it is not a question to answer from a blog post.

## Phase 1 - Supply before demand

A marketplace dies on the supply side first. A rider who opens your app and sees no cars does not try again; a rider who books and waits twenty-five minutes tells other people. You need drivers live before you promise anyone a pickup time.

## How many drivers you actually need

Do not guess, and do not use someone else's ratio. Do this arithmetic with your own numbers:

- Decide the pickup promise you want to make. Ten minutes is a very different business from three.
- Estimate the full job cycle in minutes: drive to pickup, plus waiting, plus the trip, plus repositioning afterwards. Time three real journeys in your corridor at your target hour and use that.
- Estimate jobs per hour in the busiest hour you intend to serve.
- Drivers needed live in that hour is roughly jobs per hour multiplied by job cycle in hours, divided by the utilisation you are willing to accept.

Utilisation is where the trap is. If you want a driver busy 70% of the time, divide by 0.7. But a driver at 70% utilisation cannot also be idle within ten minutes of every new request, so a tight pickup promise forces low utilisation, and low utilisation means a driver earns less per hour logged in. Pickup speed and driver earnings are in direct tension. Pick your side of that trade explicitly, because if you do not choose it, the drivers will choose for you by logging off.

Then account for the gap between drivers signed up and drivers actually logged in at the hour you need them. It is always a fraction, and at launch it is a small one. Recruit against drivers live at peak, never against total signups.

## Where drivers actually come from

- Already-licensed drivers, first. They can start now. Everyone else is on the licensing queue from Phase 0, which means months.
- Driver referral, which in most markets is the main channel. Drivers trust other drivers about earnings and about how an operator behaves in a dispute.
- Physical places drivers already are: rank queues, airport holding areas, the licensing office itself, vehicle-hire and plating garages, driver WhatsApp and Facebook groups.
- Vehicle rental partners. In many markets, the constraint on supply is not people, it is access to a compliant, plated car. An arrangement with a rental garage can unlock drivers who want to work but have no vehicle.

Assume from day one that your drivers are multi-apping. They will run you alongside every other source of work they have. You are not competing for their loyalty; you are competing for the moment they accept your job instead of somebody else's. That means job quality, not app features.

## The cold-start problem, stated honestly

At launch you have almost no jobs. That means terrible utilisation, which means poor earnings per hour logged in, which means drivers leave - right at the moment you need them present so that riders see cars.

There is no clever way around this. Every operator pays for it somehow. The options, in rough order of how well they usually work for a small independent:

- A guaranteed earnings floor for a defined window: a fixed amount per hour logged in during specific hours, for a specific number of weeks, in your one launch corridor. It is expensive but it is the only lever that reliably puts cars on the road on night one. Cap it in writing, cap it by hours, and put an end date on it before you offer it.
- Pre-booked and contract work, which is schedulable. If you know the jobs exist at 07:00 tomorrow, you can staff exactly that and the driver is never idle. This is the cheapest form of liquidity there is and it is chronically underused by new operators.
- Only running the hours you can actually fill. A service that is excellent from 06:00 to 10:00 and closed the rest of the day beats one that is unreliable all day. Reliability is the product.

Whatever you choose: model the total cost of it before you switch it on, and know the date it ends.

## Phase 2 - Own one corridor before you own a city

The instinct is to cover the whole city so that any request can be served. It is the wrong instinct. City-wide coverage with thin supply produces long pickups everywhere, which is a worse product than fast pickups in one place.

Pick one corridor and saturate it. A corridor is a repeatable flow with a predictable time pattern - not a neighbourhood, a flow. Candidates worth looking at:

- Station or transport hub to a business park or industrial estate at commuter hours.
- Hospital shift changes, which are punctual, high-volume and poorly served late at night.
- Hotels to the airport, which is pre-booked, higher value, and sold through concierges rather than through consumers.
- The night-time economy strip between 23:00 and 03:00, which is high demand and high hassle - good volume, but plan for the hassle before you take it.
- School, care and dialysis contracts, which are scheduled weeks ahead and pay on invoice.

Choose on four criteria: the flow is repeatable, the timing is predictable, you can reach the decision-maker, and the incumbent service there is visibly poor.

Two things follow from picking a corridor. First, you can staff it precisely, so drivers earn and stay. Second, you can win it on reliability rather than on price, because in a defined corridor people notice reliability.

The underrated strategy here is account and contract work. It is unglamorous, it is sold by walking into a hotel or a care provider or a factory and asking who handles their transport, and it is how a large share of private-hire operators actually stay solvent. Scheduled work lets you promise a pickup time honestly, which is the only promise that matters. Consumer app demand is the expensive way in - you buy it with discounts, and it leaves when the discount does.

Only open a second corridor when the first one stands up without subsidy.

## Phase 3 - Pricing and payment rails

## Pricing: fares and the driver's share

Decide first between a fixed quote at booking and a metered fare. Fixed quotes are easier to sell, easier to dispatch and remove the argument at the end of the journey, but they put route and traffic risk on you. Metered protects you and unsettles the passenger. Most new operators are better off with fixed quotes in a corridor they know well, precisely because they know it well.

Price against the local alternative, not against your costs - but know your cost floor per job before you set a price, so you know which jobs you are losing money on. Check whether your regulator caps fares for any class of work; some do for hackney carriage and not for private hire.

Then decide the driver's share, and understand it as a pricing decision rather than an afterthought, because it determines supply behaviour:

- Percentage commission scales with fare and keeps drivers aligned on longer trips, but drivers feel it on every job.
- A weekly or daily flat rent gives you predictable revenue and drivers keep the upside, which attracts high-volume drivers and punishes part-timers.
- Zero commission with a subscription is a strong recruiting message, but your revenue no longer grows with volume, so model that carefully.

Set your cancellation, no-show and waiting-time rules at the same time, and publish them to drivers and riders. These are the disputes that consume your day if the policy is unwritten.

## Payment rails: who holds the money

The operational question is not "can I take cards" - you can. It is who holds the money and when the driver gets paid.

- Work out honestly how much of your market is cash. In many places it is still most of it, and cash changes everything downstream: the driver already has the money, so settlement runs the other way, and you are now invoicing drivers rather than paying them. Decide whether you will accept it before launch, not after.
- If you take card payments, use a mainstream processor and check three things specifically: the total percentage plus fixed fee per transaction, the payout delay, and how chargebacks and disputes are handled. On low-value fares, a fixed per-transaction fee is a much bigger deal than the percentage.
- Driver payout frequency is a recruiting lever, not just an accounts task. Faster and more predictable payouts win drivers from operators who pay slowly. Whatever cadence you choose, never miss it - a late driver payment is the fastest way to lose a driver cohort.
- Account and corporate work needs invoicing, credit terms and a purchase-order reference on the booking. If you are chasing contract work in Phase 2, you need this at launch, not later.

- Reconcile daily from day one: jobs completed, fares collected by method, driver pay owed, fees taken. Doing this daily for a small number of jobs is trivial; retrofitting it at a thousand jobs a month is miserable.

## Phase 4 - What to launch with, and what to defer

Launch with the shortest path that lets a passenger get a ride and you get paid:

- [ ] A way to take a booking - phone, web form and app, in that order of reliability
- [ ] A way to allocate it to a driver and know they accepted
- [ ] A way for the passenger to know the car is coming
- [ ] A way to take the money and record it
- [ ] Records that satisfy your licensing authority
- [ ] A phone number a human answers, during every hour you operate
- [ ] Written procedures for the six things that will go wrong on night one

Defer all of this until the corridor is stable: surge pricing, loyalty schemes, referral engines, corporate self-service portals, demand forecasting, anything described as AI, automated driver deactivation by rating, a second city, and a second vehicle class. None of these create demand you do not have. Several will actively hurt you if switched on before you have the volume to interpret what they are doing.

Write the six procedures before night one, because they are what actually consumes your first month:

- A driver does not turn up for an allocated job.
- A card is declined at the end of the journey.
- A passenger leaves property in a vehicle.
- A passenger is too intoxicated, or a driver refuses a passenger.
- A collision or a breakdown with a passenger on board.
- A safeguarding concern involving a vulnerable passenger or a child.

Each needs a named owner, a script, and a record. Agree them with your drivers in advance. Handling these well is most of what "a good operator" means to a driver, and drivers talk.

## Phase 5 - Measure liquidity, then prove the economics

The first thousand rides are a liquidity problem, not a marketing problem. Track five numbers on one sheet, daily, from your first day of live jobs:

- Requests you could not fulfil at all, as a share of requests. This is your single most important number and the one most often not measured, because unfulfilled requests leave no trace unless you deliberately record them.
- Time from request to a driver accepting.
- Promised pickup time versus actual pickup time. Not the average - the share of jobs where you beat your promise.
- Driver idle share during the hours you are paying for or guaranteeing.
- Share of riders who book a second time within 30 days.

Four of those five are supply diagnostics. That is the correct emphasis at this stage.

Then the economics, which is arithmetic on your own numbers rather than anything anyone can benchmark for you:

- Contribution per ride equals the fare, minus driver pay, minus payment processing, minus any incentive or guarantee attributed to that ride, minus refunds and goodwill.

- Monthly fixed costs are licences, insurance, vehicles if you hold any, software, phone and office, and a realistic value for your own time.
- Breakeven rides per day equals monthly fixed costs, divided by contribution per ride, divided by the days you operate.

Calculate breakeven rides per day in week one, write it on the wall, and compare it to what you actually did today. Most operators discover their contribution per ride is negative while incentives are running, which is fine and expected - but only if you know the date the incentives stop and what contribution becomes on that date.

Do not expand to a second corridor until the first one clears its own contribution without subsidy. Expansion on top of a corridor that only works with incentives multiplies the loss.

## The honest summary

Licensing sets your real timeline and you do not control it. Supply is harder than demand and costs more than you planned. One corridor served reliably beats a city served badly. Contract and account work is dull and it is how you survive your first year. The software matters least - a whiteboard and a phone can dispatch forty jobs a day, and nothing at all can dispatch drivers who are not there.

Get the licences moving this week. Everything else can start in parallel; that cannot.

If it would help to walk through your own supply arithmetic and corridor plan with someone, taxidex.com/book-demo.