

Rapido: the minute after the accept

Rapido's weakest moment is the minute after a captain accepts, and the accept is not a commitment.

13 min read

backpocket.website/studies

IN SHORT

- Rapido's rider experience breaks in the minute after a captain accepts, when he rings, learns the destination, and decides again.
- Reads that failure as a contract problem: the accept behaves as a reservation, and the repricing happens by voice where the platform holds no record of it.
- Sets out a metric tree, one headline number, a five item build order, and the single unchecked assumption the order hangs on.
- Written from the rider app alone. Every claim about what a captain sees, earns or is penalised is marked as inference where it is made.

The trip that is awkward to buy

Rapido sells short urban trips, most of them pillion on a motorbike. Type a drop point and the app quotes a fare for bike, auto and cab before you commit to anything. The bike is why the product exists. A two kilometre hop in Bengaluru traffic is an awkward thing to buy, too far to walk and too short to feel good about cab money.

The rider I have in mind is not a tourist. Same three kilometres twice a day, to a metro gate or an office, and a ten rupee difference in the quote is visible. That is a persona I built, not something I found.

Supply is the half worth studying. Rapido calls its drivers captains, and from the rider side a captain is a name, a rating, a vehicle number and a dot on a map. Everything past that is inference, and I am working from a self-

directed fleet: no shift, no dispatcher who can force an assignment, and a bike he owns. Rented and fleet held bikes are common in two-wheeler gig supply here, so ownership is the shakiest of the three.

What is documented publicly, and what is not

Rapido books inside its app, so unlike a storefront there is no public flow to walk through. This piece is built from what the company publishes about how the service works, and from the mechanics of two sided dispatch that are common to every marketplace of this shape. That is a real limit, and it belongs at the top rather than in a footnote.

Nothing here is a report of a trip I took. Where the argument needs a fact about what a rider or a captain sees on screen, it is marked as inference, and that inference is the thing being argued about rather than evidence for it.

- Fares are quoted before booking, per vehicle type. Rapido publishes this as the model.
- Cancellation carries a fee once a captain is assigned, and the terms for it are published.
- Support runs through the app rather than a public channel, which is what the escalation argument assumes.
- Everything about the captain side is inference. No captain app seen, no ping seen, no captain interviewed.

Where the trip actually gets priced

A captain looking at a ping is working out his next two hours, and one trip is an input into that. A destination that puts him on the wrong side of a flyover at four in the afternoon carries a ride back with nobody on the seat. How much of a fare in the tens of rupees that empty return eats, I cannot compute from a rider's screen. Item three of the build order stands or falls on it.

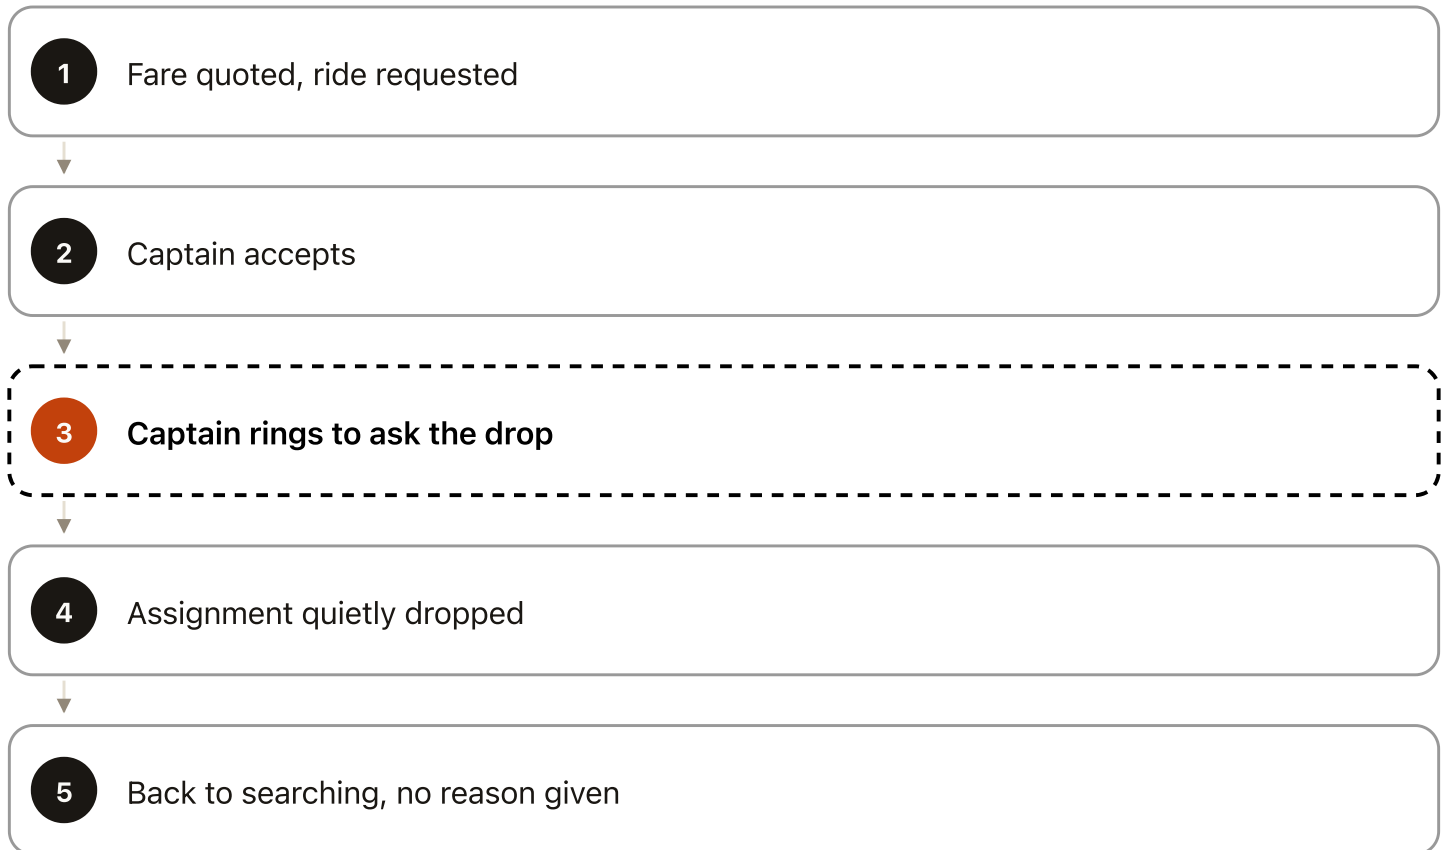
So the sequence is accept, ring, ask where you are going. Whatever the ping shows him, the call is where the trip gets priced. That is the defect I would fix first. The accept has been allowed to become a reservation the captain can reconsider off-platform, by voice, where the company cannot see it or log it.

There is a second explanation and it points the other way. If acceptance rate is scored and declines are penalised, the rational captain accepts everything and sorts it out afterwards. That would make the call a consequence of the

scoring rather than of the information, and I cannot separate the two from outside. Both explanations survive the evidence I have, which is a reason to state them both rather than pick.

Under that reading, a binding accept with the decline penalty left in place gives the captain two penalties and one exit. Which of the two explanations dominates is answerable from inside the company in a day. It decides whether the first item in the build order works or backfires.

A captain ringing to confirm a landmark is doing something reasonable, and a captain ringing to decide whether to keep the trip is not. The published cancellation reason list includes an option covering the captain having asked the rider to cancel, which I read as the company having noticed the second pattern and given it a home in the interface. That is an inference from a design choice. It is the only evidence I have for any of this, and everything below rests on it.



Where the trip actually gets priced. The accept is not the commitment, and the call is where the real negotiation happens.

Who absorbs a failed match

When a match fails after assignment, somebody pays. The rider pays in time every time, and in money depending on which side gets recorded as having cancelled.

The captain's side of that is inference and I want it labelled here, because the second item in the build order is built on it. I am assuming his own cancel button costs him something and that asking the rider to press it costs him close to nothing. If captain cancellation is already free, he has no reason to route the decision through the rider, and this part of the argument collapses.

What the rider is left holding is a dropdown. He picks the row that best describes what happened and hopes the fee comes off. The rider has the least information about what went wrong and he is the one recording it.

If that reason list carries a row for the captain having asked the rider to cancel, it is the strongest signal available from outside the company: it would mean the pattern was noticed internally and given a home in the UI. I have not seen the list, so this is the first thing I would check rather than something I am reporting, and either way it would be an inference from a design choice and not a measurement of how often it happens.

The second cost is the restart. A failed match puts the rider back at the beginning, and the supply around him has moved by then. No place is held for him, and the app says nothing about what happened.

One asymmetry sits on top of all of it, and it is an assumption rather than an observation. I am taking captain supply here as heavily male and the rider side as not. An unexplained call from an assigned stranger is not the same event for everybody who receives one.


That cuts at my own recommendation. A binding accept settles who honours the trip, and it also settles who is stuck with whom. Any commitment regime needs a rider side release that carries no fee and no reason list, on the rider's word alone.

Split liquidity from satisfaction, and keep captain pay outside both

Completed rides sits at the top, being the line the captain is paid on and the thing the rider came for. The cancellation fee complicates that, because it is income arriving out of a failure. If that fee is material, the internal case for fixing attribution is weaker than the rider side argument makes it look. Somebody should size it before this tree gets adopted.

Split hard underneath. The two branches move in opposite directions often enough that averaging them hides the thing you are looking for. Liquidity can be bought with rider pain: tolerate the negotiation, keep pushing pings out until somebody says yes, and assignment rate holds up while riders stop opening the app.

When clean fulfilment drops, the first line I would open is the share of assigned rides that die before pickup, cut by which side the failure is attributed to. Everything else in the tree is context for that one.



Liquidity and satisfaction kept apart, with captain pay outs held outside both so they cannot be traded against rider outcomes.

Clean fulfilment, and the five things it hides

Clean fulfilment: of all ride requests, the share that end in a completed ride at the quoted fare, with the rider waiting no longer than the quoted pickup time plus two minutes, and with no abandonment attributed to the captain on the way. Two minutes is an example threshold. Naming it is the point, since an unnamed tolerance widens quietly.

Two things make me choose it. The denominator is requests rather than completed rides, so failures cannot leave the room. The promise sits inside the definition, so liquidity bought by breaking the promise does not score.

It is blunt and lagging. Weekly, by city and by vehicle type, is the fastest cut I would trust, and only where the cell carries enough requests for one week to move the number. In a small city that means fortnightly.

- The rider who never requests. Quote him a fifteen minute pickup, watch him close the app, and clean fulfilment improves. It only ever gets read next to requests per active rider and search abandonment by quoted ETA.
- The captain. This number can be pushed up by sending captains on long unpaid pickups. It rises for a month and then supply thins out, which is why the guardrails sit outside the tree.

- Quiet coercion. A ride where the captain talked the rider into paying cash, or into a pickup point two lanes away, still completes at the quoted fare and still counts as clean. Fare charged against fare quoted is already in the satisfaction branch, and it travels with the headline number as a companion, the same way quoted ETA error does.
- Its own threshold. Widen the quoted ETA and the tolerance test gets easier to pass. Quoted ETA error and absolute median pickup wait get reported beside it, never folded into it.
- The attribution it depends on. The last clause of the definition needs the attribution system that is item two of the build order. So today this number either cannot be computed, or has to be computed from the dropdown I have just called unreliable. Item two's own false positives then land in this metric's numerator, which means the metric and the attribution engine ship together or neither is worth reading.

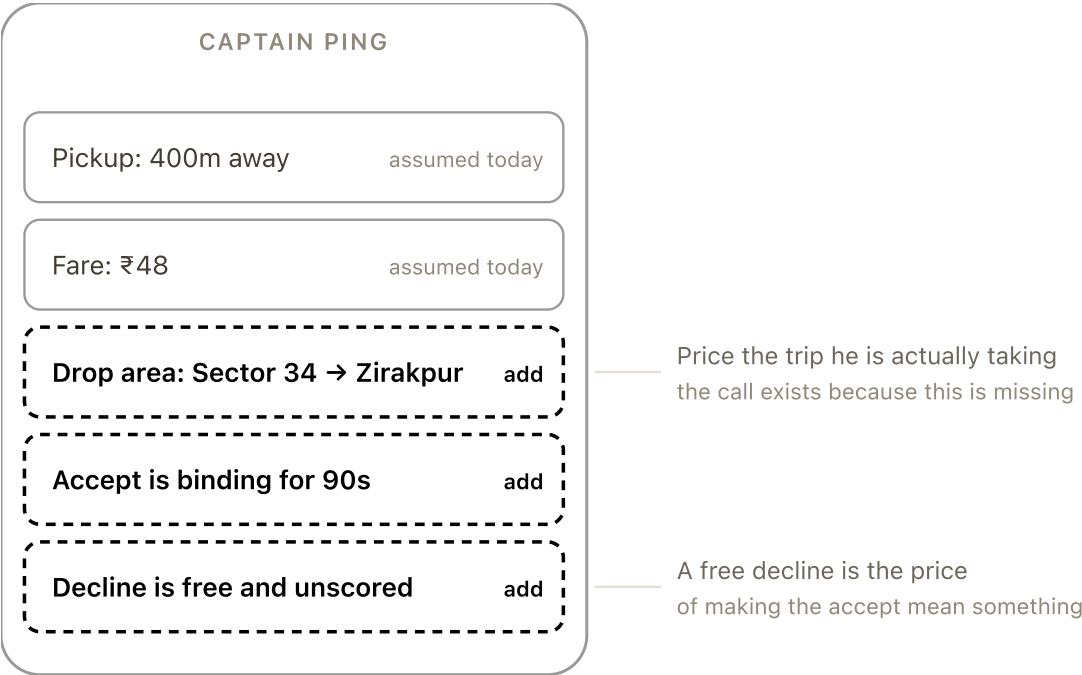
The build order, and the assumption holding it up

The ranking is judgement. The rule I want it judged against is how much of the failure each item removes per unit of engineering. From outside I can put a magnitude on neither side of that, so what follows is an argument for an order with no arithmetic under it.

I am ranking on the assumption that the ping is thin. If it already carries the destination and the fare, this table starts at item three. That assumption is checkable in twenty minutes by asking a captain at pickup to turn his screen round, and I did not do it.

A commitment needs a mechanism, and the word on its own is not one. What I mean by it: an accepted trip that ends without a pickup counts against the captain whoever presses cancel, the count escalates from a warning to reduced ping flow for the rest of that session, and declining a ping costs him nothing at all. Those sizes are illustrative.

The pairing is the part that matters. A binding accept laid on top of an existing decline penalty leaves the captain no free way to refuse a trip, and the phone call is still sitting there.



The ping as I assume it is, against the ping I would send. I have never seen one, which is the whole reason this is drawn as an assumption and set beside the alternative rather than presented as a finding.

#	WHAT	WHY IT RANKS HERE	WHAT IT COSTS
1	Put the drop area, the fare and the payment mode into the ping, and make the accept binding on the terms set out above, with declining made free	The call is where trips get repriced. Removing the reason for the call stops the cascade behind it, and removing the decline penalty is what keeps the captain on the platform while it happens	Acceptance per ping falls and time to assignment rises. That is the trade I am accepting, with the floor described in the next section
2	Attribute cancellations from platform signals rather than the rider's dropdown, and waive the fee automatically when the signals point at the captain	The signals are cheap if the call is on-platform: a call placed, a captain whose dot never moves toward the pickup, a rider cancelling soon after that call ends. I have not confirmed the number is masked, and without masking this row loses its cheapest input	False positives penalise some captains wrongly, so it needs an appeal route and a person reading the edge cases. Those same false positives land in the headline metric
3	Price the direction: pay a dead leg allowance on trips that pull a captain out of a demand pocket, and reflect part of it in the rider's quote	If the fare does not cover the empty ride back, no rule makes a rational captain accept. Whether it covers it is the number I could not get from the rider side	Fares become less predictable, which works against the metric above. Needs a cap on how far a quote can move
4	Rebuild the failure moment: hold the rider's place, widen the search, say plainly that the captain dropped off, and cover the fare difference on a rebook inside a few minutes	This is the moment the habit breaks. Ranked here on two observed vanishings, which establish that the failure exists and say nothing about its rate	A real subsidy line, and it can hide a supply shortage that ought to be visible
5	Not doing: a rider facing captain reliability score, or any further loyalty currency	A score the rider cannot act on in a thin supply area is anxiety before a ride, and a coin balance does not repair a match	Some riders will ask for the score. What would move it onto the list is a version they can act on, which needs enough supply density for a decline to be free, so dense cells first if at all

What a binding accept costs, and who pays it

A commitment grade accept means captains decline more. Some requests sit longer. In thin areas at thin hours some will not fill at all, and the rider gets a slower fill in place of a faster false start.

I will take that. A wait a rider can see is a cost he can plan around. An abandonment at minute seven arrives after he has already told the office he is on his way. The floor goes underneath it: if the share of requests assigned inside three minutes falls below eight in ten for a given city and hour band, the commitment relaxes there rather than everywhere.

Eight in ten is a placeholder for whatever a city team can defend, and the number matters less than having one written down before the rule ships.

The uncomfortable part is that this bites hardest on the trips that are hardest to serve, and those are often the ones with the fewest alternatives. Somebody in an outer ward at eleven at night gets a cleaner failure and still no ride. The only answer I have is paying for the direction, which is item three.

What would change my mind

All but one of these is a query against data the company already holds. The exception needs somebody to open the captain app. One of them takes two quarters and cannot be hurried.

- If the ping already shows the drop area and the fare, and captains still ring, then the constraint is pricing and item one has nothing left to remove.
- If captains with high decline rates then receive fewer pings, the scoring story is live and the emphasis in item one inverts: make declining free first, and watch whether the calls stop on their own before any commitment rule ships. Acceptance rate against subsequent ping volume, per captain, settles it.
- If most post assignment failures sit with a tail of captains smaller than one in twenty of the active fleet, this is enforcement and a much cheaper job. I would stop redesigning matching and start suspending accounts.
- If cash share on abandoned assignments runs above cash share on completed ones by more than a few points, captains are sorting on payment mode and the work belongs in payments.
- If riders turn out to be more sensitive to fill time than to abandonment, my trade-off is backwards and reliability has to be bought with incentives rather than rules. The test is request abandonment plotted against quoted ETA, alongside repeat rate after a slow fill against repeat rate after an abandonment.
- If assignment rate tracks captain hours per demand cell and nothing much else, supply is simply short in the cells where this happens and the honest answer is captain acquisition. That voids most of the above, and both series are already in the tree.

- If clean fulfilment climbs across two quarters with repeat rate flat behind it, the metric is measuring the promise and not the thing the promise was for. I would drop it and go back to the tree.

What I take from it

- The weak point is the accept. A captain can reprice a trip by phone, off the platform, where nothing about the decision is recorded.
- The cheap fix is attribution from signals the platform already holds, and it depends on the captain's call running through the platform's own masking layer.
- A binding accept only works with the decline penalty removed at the same time, and that pairing is the part most likely to get dropped in delivery.
- The headline metric needs the attribution engine that does not exist yet, so the two ship together and get read together.

What this is based on

Written from the outside, with no access to anything internal. Specifically:

- Rapido's own published material about how booking, fares and cancellation work, read in August 2026.
- No trip taken and no app session observed. Rapido books in-app and there is no public flow to walk through, so nothing here reports something I watched happen.
- No view of the captain app, no ping ever seen, and no captain interviewed. Every claim about what a captain sees, earns or is penalised by is marked as inference where it appears.
- The general mechanics of two sided dispatch, accept then assign then abandon, are treated as common to marketplaces of this shape rather than as facts about Rapido specifically.
- No internal data, and no figures on market size, revenue or volumes. Thresholds in the piece are examples I would set and are named as such.
- Independent analysis written from the outside, with no relationship to the company.