The Dead Mileage Playbook

A working document for fleet operators who want to measure and reduce empty running.

Before you start: what this is, and what it is not

This playbook is vendor-neutral. Every measurement method, cause and fix in it applies whatever you dispatch on - a modern cloud platform, a legacy on-premise system, or a controller with a radio and a whiteboard. If you read it, act on it and never speak to us again, it has done its job.

It also does not contain an industry benchmark table, and that is deliberate. There is no credible published figure for "the average fleet's dead-mileage ratio" that survives contact with reality, because the number depends almost entirely on things that vary between operators: how much of your work is pre-booked versus hailed, how far apart your zones sit, whether you run airport work, and whether drivers go home between jobs. A benchmark you cannot reproduce is worse than no benchmark, because it gives you a target you have not earned and a reason to stop measuring.

So the arithmetic in this document is arithmetic you perform on your own numbers. Section 2 gives you the formulas. Everything downstream depends on you actually doing it.

One honest note on where this comes from. TaxiDex is a young product. This playbook is drawn from how empty running is measured and reduced generally, not from a portfolio of fleets we have taken through it. When we have an operator who has run this and is willing to be named with their numbers, we will publish theirs instead of this paragraph.

Part 1 - What actually counts as dead mileage

Dead mileage is every kilometre driven with no fare on board. That sounds simple until you try to measure it, because it arrives in four distinct shapes and each one has a different fix. Separate them from day one or your baseline will be a single useless number.

- Approach miles. Driver accepts a job and drives to the pickup. Unavoidable in principle, but the distance is a direct consequence of allocation quality.
- Return-leg miles. Driver completes a drop-off and drives back toward their base, their home area, or a rank, with nothing on board. Usually the largest single block.
- Repositioning miles. Driver moves speculatively toward where they believe work will appear. Productive when informed, pure waste when it is a habit.
- Dwell and cruise miles. Circling near a rank, an airport holding area or a station, engine running, waiting for a job. Often invisible in mileage reports because the distance per event is small and the frequency is enormous.

A fifth category is worth tracking separately even though it is not strictly mileage: paid idle time. A driver sitting stationary on shift earns nothing and costs you the same driver-hour as a fare. If you are only counting kilometres you will miss it, and at airports it is frequently the bigger number.

Part 2 - Measuring it if you measure nothing today

The data you need

You need three things. Most operators already have all three and have never joined them up.

- Trip records from your dispatch system. Per job: timestamp of allocation, timestamp of arrival at pickup, timestamp of drop-off, pickup coordinates or postcode, drop-off coordinates or postcode, driver ID, vehicle ID, and metered or charged distance. Export as CSV or SQL. If your provider will only give you PDF reports, that is itself a finding - ask again in writing, citing the data-ownership clause in your contract.
- Total distance travelled per vehicle. From one of: telematics or tracker units, odometer readings captured at service intervals, or fuel-card litres converted using your measured consumption. In rough order of accuracy, telematics beats odometer beats fuel card.
- Cost inputs. Fuel or energy spend, tyres, servicing, and the distance-related portion of depreciation. All from invoices you already receive.

If you have no telematics and no odometer discipline, start with fuel cards. It is the least accurate route and it is still enough to establish a baseline you can improve against, which is the entire point.

The core ratio

- Dead-mileage ratio = (total distance driven ? revenue distance) ÷ total distance driven

Revenue distance is the sum of charged or metered distance across completed jobs, taken from your trip records. Total distance driven is from telematics, odometers or fuel. Both must cover the same vehicles over the same window. Use a full month, not a week - a week will be distorted by one bank holiday or one bad weather event.

Compute it three ways and keep all three:

- Fleet-wide, monthly. Your headline number.
- Per driver, monthly. This is where the money is. Empty running is almost never evenly distributed.
- Per zone or per job type, monthly. Airport work, school contracts, account work and street work behave completely differently and averaging them hides everything.

If you cannot get total distance at all

There is a fallback. Approach distance can be reconstructed from trip records alone: for consecutive jobs by the same driver, the straight-line distance from job N's drop-off point to job N+1's pickup point is a floor on the empty distance between them. Multiply by a road-network factor - 1.3 is a commonly used rule of thumb for urban road distance versus straight line, and you should calibrate it against a handful of real jobs where you know the actual distance rather than trusting it blindly.

This undercounts, because it captures nothing about the end of shift, nothing about cruising and nothing about the gap before the first job. Treat it as a lower bound, and label it as one so nobody quotes it later as if it were complete.

Converting distance into money

Do this in two layers, and keep them separate.

Layer one, running cost per empty kilometre:

- Fuel or energy cost per km = total fuel or energy spend for the period ÷ total distance driven for the period
- Wear cost per km = (tyres + servicing + distance-related depreciation for the period) ÷ total distance driven
- Running cost per empty km = fuel cost per km + wear cost per km
- Monthly running-cost loss = dead kilometres per month × running cost per empty km

Layer two, driver time:

- Empty hours per month = dead kilometres per month ÷ average operating speed in km/h
- Time cost = empty hours × your effective cost per driver-hour

For an employed-driver fleet, cost per driver-hour is a payroll figure you already have. For a self-employed or rental model you do not pay it directly - but the driver does, and it shows up as churn, as refused jobs and as pressure on your commission rate. Track it either way; just be clear which pocket it comes out of.

A note on our own calculator. The tool at taxidex.com/dead-mileage-calculator deliberately models fuel only - drivers × working days × dead km per driver per day × fuel cost per km. That makes it a deliberate floor, not a full picture. Once you add wear and driver time by the formulas above, your real figure will be higher than the calculator's. The calculator also pre-fills illustrative defaults (50 drivers, 22 working days, 30 dead km per driver per day, $0.18 per km) and applies a 23% recovery assumption. Those are the tool's assumptions for sizing the opportunity quickly, not measurements from your fleet. Replace every one of them with your own numbers as soon as you have them - that is what Part 2 is for.

Part 3 - Segment before you fix anything

Do not act on a fleet-wide average. Before choosing any intervention, split the same month of data four ways and look at the spread, not the mean:

- By driver, sorted by dead-mileage ratio, worst first. Look at the top decile specifically.
- By hour of day. Empty running at 03:00 is a different problem from empty running at 17:00.
- By zone pair - where jobs start and where they end. This exposes one-way corridors.
- By job source: pre-booked, account, app, rank, street.

The ranking in Part 4 is our expectation of what usually dominates for a mixed urban fleet. It is a starting hypothesis, not a measurement of your operation. If your segmentation disagrees with the order below, follow your data and not this document.

Part 4 - The causes, and what to change for each

4.1 Return-leg emptiness

The pattern. A car drops off eight kilometres out and drives eight kilometres back with nothing on board. Every job effectively costs double its distance. Look for it in your zone-pair analysis as corridors with heavy outbound flow and thin inbound flow.

What to change:

- Chain jobs before the current one completes. Allocate the next job while the driver is still on the drop-off leg, so they roll from fare to fare instead of turning around empty. This is the single highest-value change available to most fleets and it costs nothing but rule configuration.
- Build a back-fill habit for known one-way corridors. If your 08:00 flow is consistently suburb-to-centre, that is a standing opportunity for return work - school runs, hospital appointments, account collections - that is worth selling into rather than waiting for.
- Price the return leg into contract work. For regular account journeys on a one-way corridor, the empty return is a real cost of service. Either recover it in the rate or pair the job with return work; do not silently absorb it.
- Stop treating "return to base" as a status. If your system has a state that means the driver is heading home and unavailable, check how often it is used and how early in the shift. It is frequently a self-inflicted wound.

4.2 Poor zone design

The pattern. Zones drawn around administrative boundaries, historic radio coverage or a map that was accurate a decade ago, rather than around how vehicles actually move today. Symptoms: one zone permanently starved while its neighbour is permanently oversubscribed, or drivers refusing jobs at a zone edge because crossing it costs them queue position.

What to change:

- Redraw zones from your own drop-off density, not from council wards. Plot a month of drop-off points and let the clusters tell you where the boundaries belong.
- Kill the perverse incentives at the edges. If sitting in zone A and taking a job in zone B loses a driver their place in a queue, they will decline it, and someone further away will drive to it. Check what your zone rules reward.
- Size zones by travel time, not area. A zone should be roughly "the area a driver can serve without an unreasonable approach" - which in dense traffic is small and on a bypass is large. Uniform-sized zones on a map are almost always wrong.
- Review zone boundaries quarterly, not never. Roadworks, a new retail park or a changed one-way system can invalidate a boundary in a week.

4.3 Driver positioning habits

The pattern. Between jobs, a driver drifts to somewhere familiar - home, a favourite cafe, the rank they have used for fifteen years - rather than toward where the next job is likely to appear. This is not laziness. It is the rational response to having no better information.

What to change:

- Give drivers information, not instructions. A driver who is told where demand is expected will usually go there, because it pays them. A driver who is ordered there without explanation will resent it.
- Show each driver their own number, never a fleet average. "The fleet runs 34%" changes nothing. "You ran 41% last month, the fleet ran 28%, and here are the three trips that account for most of the gap" starts a conversation.
- Target the top decile only. Empty running is concentrated. Coaching everybody dilutes your effort and irritates the drivers who are already doing it right.
- Re-measure the same drivers two weeks later. Coaching without a follow-up measurement is a conversation, not an intervention. Part 7 gives you the scripts.

4.4 Allocating on proximity alone

The pattern. This is the subtle one, and it catches fleets that think they have already solved dispatch. Assigning "the nearest car" by straight-line radius looks efficient and often is not, because straight-line distance ignores three things that matter enormously:

- Direction of travel. A car two kilometres away heading away from the pickup is worse than a car three kilometres away heading toward it.
- The road network. Two kilometres across a river with one bridge is not two kilometres.
- What the assignment does next. Taking the nearest car may strand it in a low-demand area, creating a longer empty leg on the following job. The locally optimal choice is not always the globally optimal one.

What to change:

- Score by estimated time of arrival on the road network, not by radius. This is the single most common upgrade available and most modern systems support it if it is configured.
- Include heading and current job destination in the score, so a car already travelling toward the pickup area wins over a car that has to turn around.

- Look one job ahead where you can. Prefer the allocation that leaves the fleet better positioned for known upcoming bookings, not just the one that minimises this approach.
- Keep a manual override and watch how often controllers use it. A high override rate is not controllers being difficult - it is your rules being wrong, and the override log is a free list of the rules to fix.

4.5 Airport and station dwell

The pattern. Airports and stations concentrate two costs at once: long empty approach legs to reach them, and long unpaid waiting once there. A car in a two-hour holding queue is generating no revenue and no mileage, so it disappears from a distance-based dead-mileage report entirely while quietly being one of your most expensive assets.

What to change:

- Measure dwell in hours, not kilometres, and report it alongside the mileage ratio. Otherwise your worst-utilised vehicles look clean.
- Match arrivals to departures deliberately. If you have both drop-offs at the terminal and pickups from it, pairing them is free money; letting them happen independently means paying for the approach twice.
- Set a queue-length threshold above which drivers are better off leaving. There is a crossover point where the expected wait exceeds the value of returning to normal work. Work it out and tell drivers what it is.
- Use flight and train arrival times, not booked times, to release cars toward the terminal. Sending a car for a delayed flight converts one job into a long approach plus a long dwell.
- Check the rank rules you do not control. Airport concessions often mandate the holding area. Where you cannot change the rule, at least stop pretending the wait is free.

Part 5 - The 90-day plan

Each block has a named owner. Migrations and improvement programmes run by committee slip.

Weeks 1-2: baseline. Owner: whoever owns reporting

- Pull one full month of trip records and one full month of total distance.
- Compute the fleet ratio, the per-driver ratio and the per-zone ratio.
- Compute running cost per empty km and the monthly money figure.
- Write the baseline down, dated, and circulate it. A baseline nobody has seen is a baseline that will be disputed in week ten.

Weeks 3-6: allocation. Owner: dispatch lead

- Switch allocation scoring from radius to road-network ETA if it is not already.
- Turn on job chaining so the next job allocates before drop-off completes.
- Review the override log weekly and fix the rule behind the top three overrides each week.
- Re-measure the fleet ratio at the end of week 6. Allocation changes show up fast - within two weeks if they are working.

Weeks 7-10: zones and positioning. Owner: operations manager

- Redraw the two worst zone boundaries identified in Part 3. Two, not all of them.
- Publish expected-demand guidance to drivers for the three busiest windows.
- Run coaching conversations with the top decile only, using the scripts in Part 7.
- Re-measure those specific drivers at the end of week 10.

Weeks 11-13: lock it in. Owner: whoever owns the board

- Stand up the KPI board in Part 6 on a screen the control room actually sees.
- Set a weekly review slot, fifteen minutes, same time every week.
- Re-baseline in full using the same method as weeks 1-2, so the comparison is honest.
- Document what worked and what did not, including the interventions that changed nothing.

Do not run these in parallel. If you change allocation, zones and coaching in the same fortnight you will get a result and have no idea which change produced it, which means you cannot repeat it.

Part 6 - The KPI board

Six numbers, one screen, updated at least daily. More than six and nobody reads it.

- Dead-mileage ratio, fleet-wide, trended weekly against your dated baseline.
- Dead-mileage ratio for the worst decile of drivers, trended weekly.
- Average approach distance and approach time per job, by zone.
- Paid idle hours per driver per shift, including airport and rank dwell.
- Jobs per driver per hour, trended weekly.
- Empty-mile cost per month in currency, using your own cost per km from Part 2.

Two rules about the board. Never show a metric you are not willing to act on - an unactioned red number teaches everyone to ignore the board. And always show the baseline alongside the current figure, because a ratio without a reference point is decoration.

If you are reporting emissions to corporate or public-sector clients, CO2e per completed trip belongs here too. It moves with the same lever: fewer empty kilometres is less fuel is lower emissions, and having it already on the board is easier than reconstructing it at tender time.

Part 7 - Coaching scripts

Use these close to verbatim. The design principle is that the driver's numbers are shown, not described, and the ask is specific and small.

Opening, one-to-one, with the driver's own report on the screen between you:

"I want to show you something and get your view on it. Last month across the fleet, about a third of the distance we drove had nobody in the car. Your share was higher than most - here it is. I am not here to tell you off, because half of this is our allocation and our zones and that is on me. I want to work out which half is which."

When the pattern is end-of-shift return legs:

"Most of your empty distance is the last hour of your shift. You finish out west and drive back across town. Would it help if we tried to hold you a job heading back that way for the last part of the shift? If we can get you one, you get paid for a drive you are doing anyway."

When the pattern is speculative repositioning:

"You spend a fair bit of time moving between jobs looking for work. That is on us for not telling you where the work is. From next week you will get guidance on where the next hour's jobs are expected. Try it for two weeks and tell me straight if it is wrong, because if it is wrong I need to know."

When the pattern is rank or airport dwell:

"You waited two hours at the terminal eleven times last month. On a normal shift you would have cleared three or four jobs in that time. If the queue is longer than forty minutes, come back into town - I would rather

you were earning."

Closing, always:

"I will pull the same report in two weeks and we will look at it together. If it has not moved, that probably means the problem is ours, not yours, and I will fix it."

That last line matters more than the rest of the script. It puts a measurement date on the conversation and it makes the operator accountable too, which is what makes drivers engage instead of nodding.

Part 8 - Four ways this goes wrong

- Chasing zero. Some empty running is structurally necessary; a fleet with a zero dead-mileage ratio is a fleet turning down work. The target is not zero, it is the point where the next reduction costs more in service quality than it saves.
- Optimising mileage while destroying service. Longer waits and refused jobs are an easy way to cut empty kilometres and a fast way to lose regulars. Watch time-to-pickup and refusal rate alongside the ratio, always.
- Measuring the metric into meaninglessness. If drivers work out that the ratio affects their standing, the ratio will improve without the behaviour changing. Cross-check against fuel spend, which is much harder to game.
- Changing everything at once. Covered above and worth repeating, because it is the most common failure. One change, measured, then the next.

The one-page checklist

- ☐ Exported one full month of trip records with allocation, arrival and drop-off timestamps
- ☐ Obtained total distance driven for the same vehicles over the same month
- ☐ Computed fleet-wide dead-mileage ratio and written it down with the date
- ☐ Computed per-driver ratio and identified the worst decile
- ☐ Computed per-zone and per-job-type ratios
- ☐ Calculated running cost per empty km from own fuel, tyre, servicing and depreciation invoices
- ☐ Calculated empty driver-hours and the cost of that time
- ☐ Named an owner for each block of the 90-day plan
- ☐ Switched allocation scoring from straight-line radius to road-network ETA
- ☐ Enabled job chaining so the next job allocates before drop-off completes
- ☐ Started reviewing the dispatcher override log weekly
- ☐ Identified the two worst zone boundaries from drop-off density
- ☐ Published expected-demand guidance for the three busiest windows
- ☐ Held coaching conversations with the top decile, using their own numbers
- ☐ Diarised the two-week re-measure for those specific drivers
- ☐ Measured airport and rank dwell in hours, separately from mileage
- ☐ Set a queue-length threshold at which drivers should leave the terminal
- ☐ Stood up the six-metric KPI board where the control room can see it
- ☐ Booked a fifteen-minute weekly review in the calendar
- ☐ Re-baselined at day 90 using the identical method

Size the opportunity in about a minute at taxidex.com/dead-mileage-calculator, and if you want to walk through your own trip data with someone, taxidex.com/book-demo.