Automating the For-Hire Fleet: Economic, Labour, Fiscal and Legal Scenarios for Autonomous Ride-Hail in the Denver-Boulder Region

A conditional scenario assessment, with policy options sorted by the authority required to act on them

Working paper. Scenario analysis, not a forecast, and not legal advice.
Evidence cut-off 2026-07-26 • Case geography: Denver-Boulder / DRCOG region
Prepared for: Office of the Governor; Colorado General Assembly; Public Utilities Commission; CDOT; CDLE; OEDIT; DRCOG and regional partners
Companion files: a reproducible model, an audit harness, and a calculation workbook
Every quantitative result is conditional on the parameter registry in Appendix A. Bracketed identifiers resolve to Appendix B.

ABSTRACT

Autonomous ride-hail arrived in Denver in the summer of 2026 in a form that is easy to misdescribe in either direction. It is not yet a public commercial market: the announced service was rider-only and limited to the operator's own employees, with public access expected later [S1]. Nor is it unregulated: the Public Utilities Commission granted the operator a luxury limousine permit in June, and incumbent taxi companies promptly intervened against a related vehicle waiver [S2]. The accurate description is stranger and more useful than either caricature, a pre-public deployment operating inside a statutory category built for a different business.

This paper builds a conditional model of that transition and reaches four conclusions. The for-hire market is small: 0.64% of regional person-trips and 0.88% of vehicle-miles, a finding that survives the full range of denominator assumptions. Automating it is consequential for the people and firms involved and is not, by itself, a metropolitan transformation. The economically decisive threshold lies elsewhere: when the cost of a robotaxi mile falls below the cost of owning a lightly-used household vehicle, which reorganises household transport spending, land use and the fiscal base simultaneously. Whether that transition raises or lowers regional vehicle travel turns on a single operational ratio: the share of fleet miles run empty, which is measurable, targetable, and currently unmeasured in Colorado. And the statutory exposure is not the hole it first appears to be: Colorado's per-ride fee attaches to services connecting riders to *drivers*, and a driverless fleet does not fit that language, but the observed consequence was not an absence of regulation. It was a quiet search for the nearest existing category, noticed first by the incumbents who bear the heavier obligations.

Because the useful responses are distributed across branches, Section 9 sorts options by the authority each requires. Several of the highest-value steps, establishing a data framework through the existing Autonomous Mobility Task Force, publishing the active-driver distribution already held by CDLE, seeking a joint agency classification determination, sit within executive authority and do not require a legislative session. Others do require statute. The sequencing logic is the same throughout: measure first, define functionally second, price externalities only once the first two exist.

1. Introduction

Most analysis of autonomous vehicles has been written in advance of them. The literature is consequently rich in projection and thin in observation, and it has developed a characteristic failure mode: it treats the arrival of the technology as the event, when the economically interesting events happen well before and well after arrival. Denver offers an unusual opportunity to correct for this, because the deployment there is real enough to make policy timely and immature enough to make confident forecasting obviously premature.

The region is also, for reasons that are partly accidental, a good place to observe the underlying economics. Colorado imposes no autonomous-vehicle-specific authorisation regime. In jurisdictions that do, deployment pace is confounded by the permitting queue, and an analyst cannot easily separate what the technology can do from what the regulator will allow. In Colorado the observed pace sits closer to the technology and cost constraint itself. The regulatory posture this paper questions on policy grounds is precisely what makes the region legible on empirical ones.

Three questions organise what follows. What does automated for-hire service actually cost, and which components of that cost are observable? How much of the regional transport system is genuinely in play, and on what timescale? And which instruments remain available to each level of Colorado government, given a statute that has already declared this a matter of statewide concern?

Three households recur through the argument, because the aggregate figures conceal the fact that this transition reaches very different people in very different ways. **Marcus** drives for a ride-hail platform in Denver. **The Delaneys** are a two-earner household in Arvada with two cars, one of which barely moves. **Ray** is an independent HVAC contractor working out of Commerce City. **Eleanor** is retired, lives in Lakewood, and owns a twelve-year-old sedan she paid off long ago. They are composites constructed for exposition, not survey respondents, and no number in this paper derives from them. Each is introduced where it clarifies a mechanism the arithmetic alone tends to obscure, and one of them, as Section 6 shows, is the reason the paper's most consequential estimate should be treated with suspicion.

A note on epistemic posture, because it shapes every number here. Colorado publishes no trip-level ride-hail dataset. No operator discloses a Denver cost stack. There is no public fare record for autonomous service in this market, because there is no public service. Every regional quantity is therefore reconstructed from statewide aggregates through named assumptions, each carrying an explicit range. What follows is a structured argument about magnitudes and mechanisms, not a measurement. Where the model would have to extrapolate beyond what its functional forms can bear, it declines to answer rather than producing a fragile number, and says so. That is a deliberate choice about what a policy reader should be able to trust.

2. The institutional setting

2.1 What is actually operating

As of the evidence cut-off, Waymo announced rider-only operations in Denver for its own EMPLOYEES with public access expected later [S1]. No public commercial market is established. Claims requiring observed Denver fares, utilisation or empty mileage cannot be validated.

This matters more than a pedantic point about launch dates. It means the observable quantities that would discipline a cost or demand model, fares, utilisation, wait times, empty mileage, customer mix, reliability in weather, do not yet exist for this market. Any analysis claiming to characterise Denver robotaxi economics from observation is, at this cut-off, characterising something else.

2.2 Oversight exists, inside a category built for something else

The PUC granted Waymo a luxury limousine permit effective 12 June 2026, and taxi operators intervened against a vehicle waiver [S2]. Luxury Limousine is a Limited Regulation Carrier permit: not rate-regulated, annually renewed, subject to operational and safety rules [S3]. Colorado oversight of for-hire passenger carriers therefore exists; what is contested is the fit of the category.

The choice of category repays attention. A luxury limousine permit is a Limited Regulation Carrier authorisation: it presumes prearranged, app-dispatched service, it is not rate-regulated, and it carries none of the minimum-fare or public-service obligations attached to taxi operation. Colorado never wrote a box for an autonomous fleet, so the operator selected the existing box that fit best, and the incumbents who bear the heavier obligations identified the competitive asymmetry immediately, which is why the proceeding is contested rather than routine.

> **Finding 1.** Colorado has no statutory category for an autonomous fleet, so the operator selected the existing category with the lightest obligations. Limited Regulation Carriers face no rate regulation and no minimum airport fares, which intervenors identified as the competitive issue [S2, S3]. The definitional gap is real; its consequence is category selection, not an absence of regulation.
>
> This is a more useful diagnosis than "regulatory gap," because it is observable and because it predicts where friction appears. A gap implies an absence that someone must eventually notice. Category substitution produces no absence at all: service operates, a permit exists, an agency has acted, and the statutory purpose has quietly stopped being served. The signal is not a complaint from the public. It is a complaint from the competitor.

2.3 The fee question

The prearranged ride fee attaches to transportation network companies, defined around connecting riders to drivers [S5, S6]. A Limited Regulation Carrier is not a TNC. The gap gap argument is therefore credible on the statutory text, but classification remains UNRESOLVED pending agency or judicial determination.

The practical significance is modest in dollars and large in structure. The revenue at stake in any scenario modelled here is in the single-digit millions annually. But the mechanism is instructive: no actor decides that the fee should stop applying, no proceeding is opened, and no dispute arises. The obligation simply ceases to attach, because the operative noun in the statute describes a person who is no longer present.

2.4 Preemption and the fee/tax boundary

SB 17-213 creates strong preemption risk for AV-specific operating rules [S4]. It does not by itself resolve technology-neutral curb management, general commercial-carrier rules, road-use charges, or proprietary action. Preemption analysis is instrument-specific.

A 2025 Court of Appeals decision upheld the challenged SB 21-260 enterprise structure [S16]. That is precedent, not blanket approval; it does not guarantee any future fee survives. Denver International Airport may hold additional leverage as a property owner rather than a regulator, since access terms negotiated by a proprietor are contractual instruments rather than ordinances. That distinction is real but litigable, and untested in Colorado on these facts. It appears in Section 9 as an option worth analysing, not as a conclusion.

3. Evidence base: prior literature, data, and method

3.1 Where this paper stands relative to the literature

The autonomous-vehicle literature is large, and it is worth stating plainly which of this paper's claims it supports, which it disputes, and which it has simply never addressed. The table below is organised by claim rather than by author, because the useful question for a policy reader is not what has been written but which conclusions can be leaned on.

Claim	Status in the literature	Refs	Basis
The for-hire market is a small share of regional travel (Finding 3)	**Corroborated**	`L1, L3`	Uncontroversial across the literature; for-hire modal share is small in every US metropolitan estimate.
Service and operating costs, not vehicle cost, dominate fleet economics (Finding 2)	**Corroborated**	`L1, L2`	Bösch et al. find service costs, not vehicle costs, dominate shared autonomous fleet economics. Consistent with this paper's finding that depot, supervision and insurance together exceed vehicle capital cost at the current state.
Fleet duty cycle near 40-60% of the service window (§4)	**Broadly consistent**	`L1, L12`	Bösch et al. report commercial-fleet active time of 57% peak, 59% off-peak and 30% at night; this model's 43% average duty cycle sits within that envelope. Cramer & Krueger establish occupied-time share as the operative productivity variable.
Empty miles are roughly 35-41% of for-hire vehicle-miles (§6.1)	**Corroborated**	`S10, L6, L10`	Denver 40.8% [S10] and California ~38.5% [L10] are independent estimates in close agreement. Erhardt et al. attribute measurable congestion increase to the mechanism.

Automation raises VMT through lower generalised travel cost (§5)	**Corroborated in direction, wide in magnitude**	L7, L8, L9	Automation lowers generalised travel cost and raises VMT; published ranges are very wide, which is why this paper caps rather than extrapolates its demand response.
A material share of vehicles is not substitutable at any price (§6, ceiling)	**Corroborated**	L1	Bösch et al. conclude from cost structure alone that private vehicle ownership will likely endure. Independent support for the functional substitutability ceiling.
Cheaper per-mile mobility causes households to shed vehicles (Section 6)	**CONTESTED, literature does not support a confident sign**	L3, L4, L5	Ward et al. (2019) find ~3% lower per-capita registrations at STATE level; Ward et al. (2021), substantially the same authors, find ~0.7% HIGHER registrations at CITY level and note that aggregation reverses the sign. Wang et al. find lower ownership among active users. The existing evidence tests ride-hailing at prices that never approached ownership parity, so it neither supports nor refutes this paper's mechanism, but it forecloses any confident claim that the effect is large.
Remote supervision is the binding cost constraint	**Literature silent**	none	No peer-reviewed estimate of remote-supervision ratios or their cost incidence was located. The academic gap mirrors the disclosure gap, which is itself a finding.
Statutory categories fail by substitution, not by visible gaps (Findings 1, 6)	**Literature silent**	none	Legal scholarship addresses AV permissibility and preemption, but no located work documents the mechanism whereby an operator selects an ill-fitting existing statutory category and the original obligation lapses without dispute. This is where the paper makes its most novel claim, and correspondingly where it is least corroborated.

Three observations follow. The **cost-structure result is well supported**: Bösch and colleagues, in the canonical bottom-up analysis of autonomous mobility services, conclude that service costs rather than vehicle costs are the main cost factor for shared autonomous fleets [L1], which is consistent with this paper's finding that depot, supervision and insurance together exceed vehicle capital cost at the current state. Their reported commercial-fleet active time of 57% at peak and 59% off-peak also brackets this model's 43% average duty cycle, which is reassuring for a parameter derived independently.

The **ownership channel is the paper's most exposed claim**, and honesty requires foregrounding that. The empirical record on whether ride-hailing reduces vehicle ownership is not merely uncertain, it is sign-ambiguous. Ward and colleagues found roughly 3% lower per-capita vehicle registrations following Uber and Lyft entry using state-level data [L3]; substantially the same authors, working at city level, found roughly 0.7% *higher* registrations and noted explicitly that aggregating to state level reverses the sign [L4]. Others find lower ownership among active users [L5], and Schaller has reported household vehicle counts rising in cities where ride-hailing use was heaviest.

There is a defensible reading of that record which is consistent with this paper, and it should be stated as a reading rather than a result. Every one of those studies observes ride-hailing at prices that never approached parity with vehicle ownership. A service at three to four dollars a mile is not a test of what happens at one dollar a mile, and finding no ownership effect at the former is weak evidence about the latter. But the converse also holds: nothing in the existing evidence *supports* a large ownership response either, and a policy reader should treat Section 6 as a conditional mechanism with an uncertain trigger rather than as a forecast. It is worth noting that Bösch and colleagues, reasoning from cost structure alone, concluded that private vehicle ownership would likely endure [L1], which independently supports the functional ceiling introduced in Section 6 even as it cautions against the channel's magnitude.

The **two silences are as informative as the findings**. No peer-reviewed estimate of remote-supervision ratios or their cost incidence was located, which mirrors the disclosure gap in §4.1, the number nobody publishes is also the number nobody has independently estimated. And no located work documents the category-substitution mechanism of Finding 1: legal scholarship addresses whether automated vehicles are permissible and how preemption operates, but not the process by which an operator selects an ill-fitting existing category and the original obligation lapses without anyone contesting it. That is where this paper claims most and is corroborated least, and it should be read accordingly.

3.2 The binding data constraint

> Colorado publishes no trip-level ride-hail dataset. Every regional quantity below is reconstructed from statewide aggregates through named assumptions. No result rests on observed Denver trip records, and none can be validated against a public robotaxi market, because none exists.

The reconstruction begins from 37227697 statewide transportation tasks reported by the Department of Labor and Employment [S7], of which 70% are allocated to the region. Regional denominators use a 3.3 million nine-county population. That figure is deliberately not the RTD service-area population of 3098079 [S8]; the two are commonly conflated, and using the smaller one shrinks the denominator and inflates every share reported against it.

Baseline diagnostic	Central	Low - high	Basis
Allocated regional tasks	26059388	22336618 - 29782158	Official statewide count × allocation assumption
Regional person-trips	4.10 B	3.45 - 4.79 B	Population and trip-rate assumptions
Regional VMT	30.1 B	24.1 - 36.5 B	Population and VMT-rate assumptions
Share of person-trips	0.64%	0.65% - 0.62%	Modelled; below one percent throughout
Share of regional VMT	0.88%	0.72% - 1.06%	Modelled; includes human deadheading
Driver workload	5262 FTE	3759 - 7159	Tasks × hours ÷ 2,080
Hours per *licence* per week	8.6	8.4 - 8.9	Diagnostic only, see §3.2

3.3 The human benchmark is a band, not a number

Comparing automated to human service requires a human cost benchmark, and here the absence of Denver fare data bites hardest. It would be possible to calibrate against fares observed in another metropolitan area, and tempting, because it yields a single confident figure. It would also be wrong: fare levels vary substantially across markets with labour supply, regulation, congestion and airport structure, and a benchmark imported from a high-cost coastal market would systematically misstate the comparison here.

The benchmark is therefore built bottom-up from a driver's net labour compensation, vehicle cost, platform commission and gratuity, and reported as a **band** spanning the plausible labour rate: **$2.83 to $3.58 per occupied mile**, central $3.11. Every comparison against it inherits that width. This is the honest presentation when the most important input to a comparison is unobserved, and it has the useful property of directing the reader's attention to the assumption that actually matters.

3.4 Licences are not workers

The labour analysis carries a caveat that is easy to state and easy to forget. Colorado reports driver *licences*, not verified active workers.

> The denominator is LICENCES, not verified active drivers. A low hours-per-licence figure is consistent with many part-time workers OR with fewer full-time workers alongside dormant licences. It cannot establish the distribution of driving intensity.

> Marcus. He drives roughly thirty hours in a normal week, mostly airport runs and Friday evenings, and it is the larger share of what he lives on. In the state's records his licence is indistinguishable from that of a graduate student who drove four shifts in 2023 and has not logged in since. Both are one licence. Colorado has no way to tell them apart, and a transition programme sized to licence counts would treat them identically.

The temptation is to divide total driver-hours by licence counts, observe a small number, and conclude that ride-hail driving is overwhelmingly supplemental work. The arithmetic is right and the inference is not: the same average is produced by many part-time workers and by a smaller number of full-time workers alongside a large stock of dormant licences. These are different populations requiring different policy, and Colorado cannot presently distinguish them, though, as Section 9 notes, the data to do so is already collected.

3.5 Three modules, deliberately not combined

The model separates a conditional cost stack, a short-run for-hire market response, and a long-run household ownership channel. Short-run for-hire demand and long-run ownership substitution are reported separately and are never added, because the channels overlap and their interaction is unmeasured.

The separation is substantive rather than presentational. The two demand channels overlap: a household that sheds a vehicle because robotaxis are cheap is also, trivially, responding to price. Adding them double-counts. Taking a maximum or a weighted blend imposes an assumption about their relationship that no available evidence supports. Reporting them separately leaves the judgement with the reader, which is the only defensible option when the interaction is unmeasured. Epoch labels are conditional state names, not forecast dates.

4. What automation costs

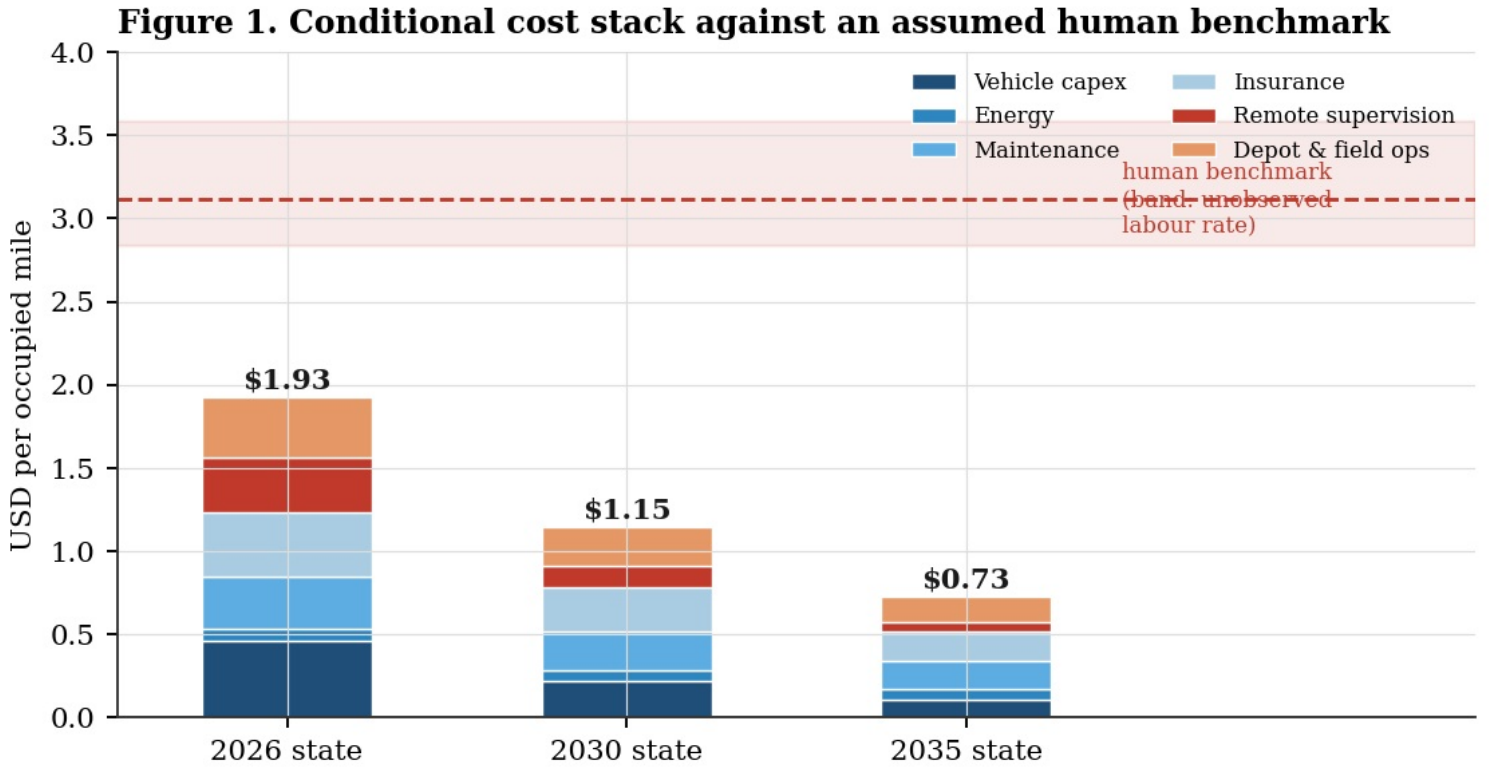

Figure 1. Component costs across three stipulated technology states, against the human benchmark band. These are conditional model states, not operator disclosures.

State	Operating cost per occupied mile	Cost-based retail per occupied mile	Empty-mile share	Supervision FTE per vehicle	Total FTE per vehicle
2026 state	$1.927	$2.312	35.5%	0.199	0.399
2030 state	$1.148	$1.377	32.4%	0.091	0.291

2035 state	$0.728	$0.873	29.6%	0.049	0.249

Removing the driver removes the largest single line in the human cost structure, but the saving is substantially consumed by two costs the human model does not carry: remote supervision, and depot and field operations. A robotaxi does not end its shift in a driveway its operator already pays for. It returns to leased property, is cleaned and charged by someone, and is watched in some fashion by someone else. Those obligations are not incidental; at the current state they are together comparable to the vehicle itself.

Supervision is charged here against *concurrently active* vehicles rather than fleet size, since an idle or charging vehicle generates no intervention requests. At current-state parameters the implied duty cycle is 43% of the service window. Demand is not flat, so staffing must cover a peak well above the average, and shift scheduling can chase that peak only partially. The resulting staffed fraction, not the duty cycle, and not the fleet count: is what determines supervision cost.

4.1 Which component actually drives the decline

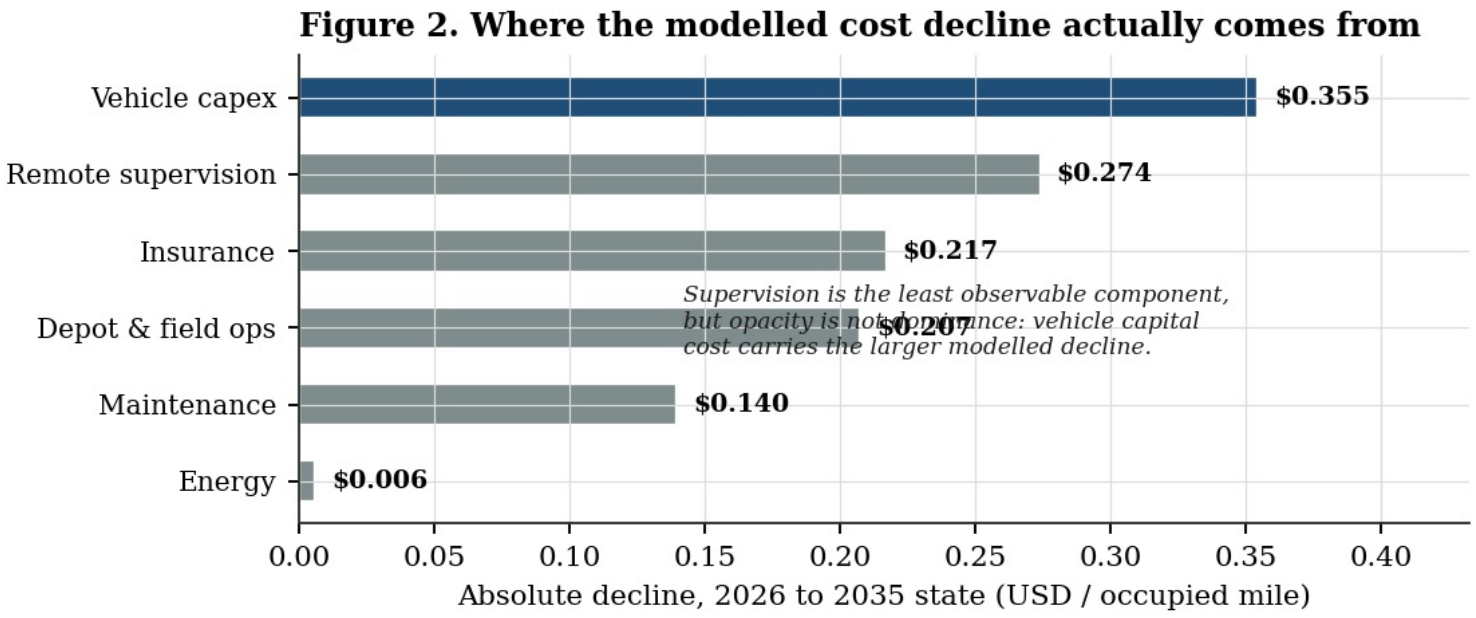

Figure 2. Absolute cost decline by component between the current and 2035 states. The ranking is computed from the model rather than assumed.

Component	Decline across the modelled states
vehicle capital cost	$0.3546
remote supervision	$0.2740
insurance	$0.2171
depot and field operations	$0.2074
maintenance	$0.1396
energy	$0.0059

Finding 2. The largest modelled decline is vehicle capital cost, not supervision. It is tempting to identify supervision as the uniquely binding constraint, because it is the component nobody discloses and the one that sounds most like the frontier of the technology. Opacity is not dominance, and these assumptions do not support the stronger claim.

What survives is a narrower observation about measurement. Supervision is the least observable component of the stack, and published ratios are ambiguous in a way rarely acknowledged: one operator to ten *fleet* vehicles and one operator to ten *concurrently active* vehicles are materially different claims about a business, and the distinction is almost never specified. Any external estimate of robotaxi economics, including this one, inherits that ambiguity.

4.2 Supervision cost and supervision employment are one quantity

Remote-operator hours per vehicle determine both supervision cost per mile and supervision full-time equivalents per vehicle. They are not two variables that happen to move together; they are the same quantity in different units. A scenario cannot deliver cheap supervision and abundant supervision employment at once, and any strategy promising both is internally inconsistent before it is examined empirically.

The identity does not extend to total fleet employment. Depot turnaround, field response, high-voltage maintenance, sensor calibration and charging operations are separate assumptions driven by vehicle count rather than supervision efficiency. A low-cost fleet can still support substantial local employment. What it cannot support is a large supervision workforce, a distinction that matters a great deal for where workforce policy should aim.

5. The short-run market: small, and likely to stay small

Figure 3. Short-run for-hire market only, ownership substitution excluded

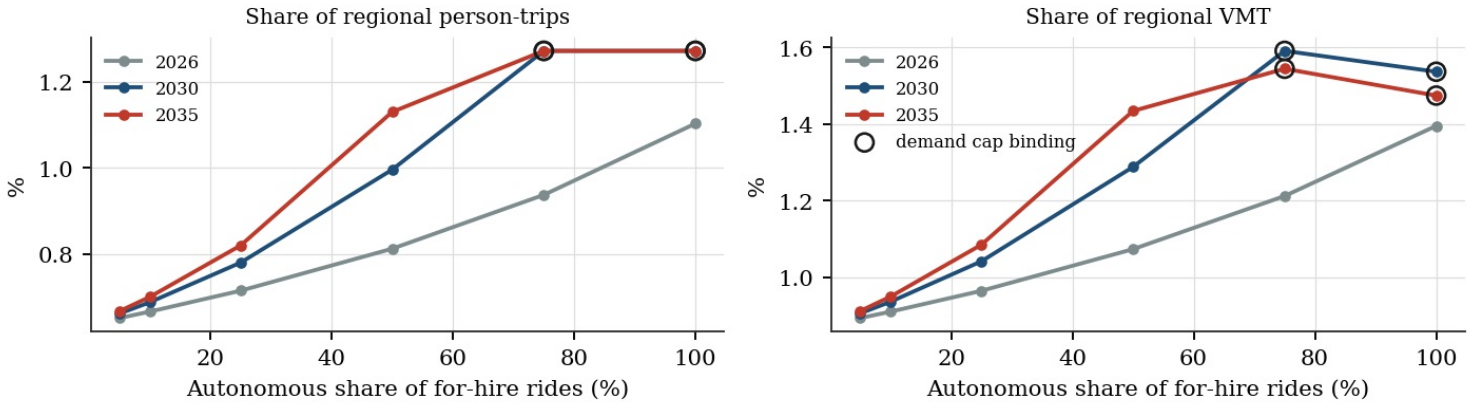

Figure 3. Short-run for-hire channel only; household vehicle substitution is excluded. Circled markers indicate scenarios where the demand guardrail binds and the plotted value is the cap rather than a model result.

State	AV share	Price index	Raw response	Bounded	% person-trips	% regional VMT	Driver FTE lost	New fleet FTE
2026	5%	0.982	1.02×	1.02×	0.65%	0.89%	150	79
	10%	0.963	1.05×	1.05×	0.67%	0.91%	308	163
	25%	0.908	1.12×	1.12×	0.71%	0.96%	831	436
	50%	0.816	1.28×	1.28×	0.81%	1.07%	1904	991
	75%	0.724	1.47×	1.47×	0.94%	1.21%	3324	1717
	100%	0.632	1.73×	1.73×	1.10%	1.40%	5262	2694
2030	5%	0.969	1.04×	1.04×	0.66%	0.91%	69	48
	10%	0.938	1.08×	1.08×	0.69%	0.94%	146	100
	25%	0.844	1.23×	1.23×	0.78%	1.04%	426	283
	50%	0.688	1.57×	1.57×	1.00%	1.29%	1143	723
	75%	0.532	2.13×	2.00× capped	1.27%	1.59%	2631	1386
	100%	0.377	3.23×	2.00× capped	1.27%	1.54%	5262	1848
2035	5%	0.962	1.05×	1.05×	0.67%	0.91%	25	35
	10%	0.924	1.10×	1.10×	0.70%	0.95%	54	73
	25%	0.810	1.29×	1.29×	0.82%	1.09%	178	215
	50%	0.619	1.78×	1.78×	1.13%	1.43%	587	592
	75%	0.429	2.76×	2.00× capped	1.27%	1.54%	2631	1000
	100%	0.239	5.58×	2.00× capped	1.27%	1.47%	5262	1333

Finding 3. Ride-hail accounts for 0.64% of regional person-trips, and stays under 1.3% even where the entire market is automated and demand expands to the model's ceiling. Automating this market matters intensely to the drivers, carriers, riders and regulators directly involved. It is not, by itself, a metropolitan transformation, and analyses presenting it as one are generally importing the ownership channel of Section 6 without saying so.

> Marcus, competing. What reaches him first is not a price war but a change in the shape of his day. For every hundred miles he carries a passenger, he drives roughly sixty-nine empty [S10], repositioning after a drop-off, circling before a pickup, returning from the airport. He is paid for the first hundred and absorbs the cost of the rest. An automated competitor carries the same empty-mile burden, but as a line on a fleet operator's balance sheet rather than as unpaid time. The first thing Marcus is likely to notice is not a fare cut. It is the best hours getting thinner while the empty miles stay exactly where they were.

This is the single most useful thing to internalise early, because it determines whether the appropriate posture is crisis response or preparation. It also cuts against the instinct that a small present share implies a small future one, the ownership channel is a different mechanism operating on a much larger base, and the two should not be reasoned about together.

5.1 Full conversion is a bound, not a scenario

The tables above run the autonomous share to 100% because a bounding case is analytically useful, not because it is achievable. Several categories of for-hire trip resist automated service independently of cost: wheelchair-accessible vehicle trips requiring equipment and often assistance; travel outside any plausible service area, which at present is a handful of neighbourhoods; travel during weather-related suspension, which in this market is not a rare event; large parties and cargo; and riders without a smartphone or a payment

instrument. The 100% column should be read as an upper envelope on the channel, and the realistic ceiling is materially lower, a point Section 6 makes quantitatively for the ownership channel, where the constraint can actually be parameterised.

5.2 Where the model declines to answer

> A constant-elasticity demand curve extrapolated across a price change far outside any calibration range produces a raw response of 5.58× baseline rides at full autonomous share in the 2035 state. That is not a finding; it is a functional form failing quietly. The model reports a 2.0× guardrail instead and marks those cells. They should be read as *unknown*, not as an estimate.

The alternative, publishing the raw figure with a caveat, would put a large, memorable and essentially fictional number into circulation, where it would outlive the caveat. Refusing to answer is less satisfying and more defensible. What the model cannot establish is how many trips would be newly generated, diverted from transit or active modes, shifted from private cars, or induced among people who cannot currently drive at all. Each requires trip-level data and a mode-choice model that Colorado is not presently in a position to build.

6. The ownership frontier: where the transition actually happens

The economically decisive threshold is not substitution for Uber. It is parity with owning a car. The distinction is easy to state and consistently underweighted, because the ride-hail comparison is the visible one and the ownership comparison is the consequential one.

The mechanism is a fixed-cost problem. Ownership costs are dominated by depreciation, insurance, registration and finance, all of which accrue whether or not the vehicle moves. A household driving 6500 miles a year pays roughly $1.40 per mile on an average-new-vehicle basis, most of it for the privilege of the vehicle existing. Per-mile mobility priced below that figure does not merely compete with ride-hail; it makes the second car in a two-car household economically indefensible. That is a far larger base to act on, and it moves household spending, parking demand, land use and the fuel-tax base at the same time.

> **The Delaneys.** Two cars, two commutes, and a second vehicle that turns over about 6500 miles a year: school runs, the hardware store, the occasional trip to the airport when the schedules collide. On an average-new-vehicle basis it costs them roughly $1.40 a mile, and almost none of that is fuel. It is depreciation, insurance, registration and a finance payment that arrives whether the car moves or not. They are not paying to drive. They are paying for the option to drive, and the option is priced as though they exercised it far more often than they do. At $2.31 a mile the alternative is not close. At $0.87 it is not close in the other direction, and the second car becomes difficult to justify to anyone who looks at the number.

> **Ray.** He runs an HVAC business out of Commerce City with two vehicles. The van is fitted out, shelving, parts inventory, a ladder rack, a compressor bolted to the floor, and does about 11,000 miles a year. The other is a fifteen-year-old pickup that covers maybe 4,500 miles: hauling a water heater, towing a small trailer, getting up an unplowed county road in March to a service call the van cannot reach, starting at 5:40 in the morning somewhere no fleet has ever sent a vehicle.
>
> The model, looking only at annual mileage, classifies that pickup as one of the most attractive shed candidates in the region. Ray would sooner sell the house. The truck is not transport; it is a tool that happens to drive, and no price per mile makes a robotaxi carry a water heater up an unplowed road at dawn.

Ray is the reason the response curve needs a ceiling rather than only a discount. Low annual mileage is not evidence of substitutability, and for occasional-use work and access vehicles it is actively misleading, the vehicles that move least are disproportionately the ones kept for the specific jobs nothing else can do. The model therefore separates two limits that are easy to conflate. A *behavioural* discount, applied at 45%, covers households that could substitute but will not: habit, garaging, trip-chaining, uncertainty. A *functional* ceiling, set at 80%, covers vehicles that cannot be substituted at any price, purpose-built work vehicles, towing and hauling capability, equipment carriage, unplowed and unpaved access, and travel outside any plausible service window. The ceiling applies first, and the behavioural discount applies to what remains.

The 80% figure is the weakest number in this paper, and it should not stay that way. It is estimable from records Colorado already holds, and the estimate requires nothing whatever from the fleet operators.

Registration records give functional class directly: VIN decoding yields body style, gross vehicle weight rating, drivetrain and fuel type, and the registration itself carries plate class (passenger, commercial, farm, apportioned) along with geography [S18]. What registration lacks is use. That is supplied by the emissions programme, which since it covers gasoline vehicles more than seven model years old across nine Front Range counties on a biennial cycle, touches almost exactly the population of older vehicles this channel cares about [S17]. The two systems are jointly administered by CDPHE and the Department of Revenue, and emissions results already upload electronically to the county registration record. The linkage exists operationally. It has simply never been asked this question.

Joining them on VIN yields the functional-class by annual-mileage distribution that this parameter currently approximates, and a defensible bracket falls out immediately. A *floor* of clearly non-substitutable vehicles (commercial and farm plates, cargo vans, Class 2b and 3 trucks) is unambiguous. A *ceiling* of all pickups, vans and four-wheel-drive vehicles registered outside plausible service geography is the maximum defensible claim. The truth lies between, and a modest household survey closes the gap by establishing what share of lightly-driven pickups are Ray's tool rather than someone's lifestyle vehicle. Two caveats belong with this: body style is a proxy rather than a fact, and electric vehicles are exempt from emissions testing, so the coverage of this method degrades as the fleet electrifies. The window for using it is open now and narrowing.

On central assumptions that ceiling sets aside roughly 96320 of the 481600 low-mileage household vehicles in the region as permanently outside this channel. It is worth being precise about what the ceiling is doing: at the modelled prices the price-response heuristic yields

75%, which sits just below the ceiling, so the ceiling *bounds* the mature scenario without currently *binding* it. Under more aggressive cost assumptions it binds directly, and it is the difference between a model that reports a plausible ceiling and one that reports full conversion of every lightly-driven vehicle in the metropolitan area, which no evidence supports and which Ray's pickup alone falsifies.

Figure 4. Long-run ownership channel, reported separately from short-run demand

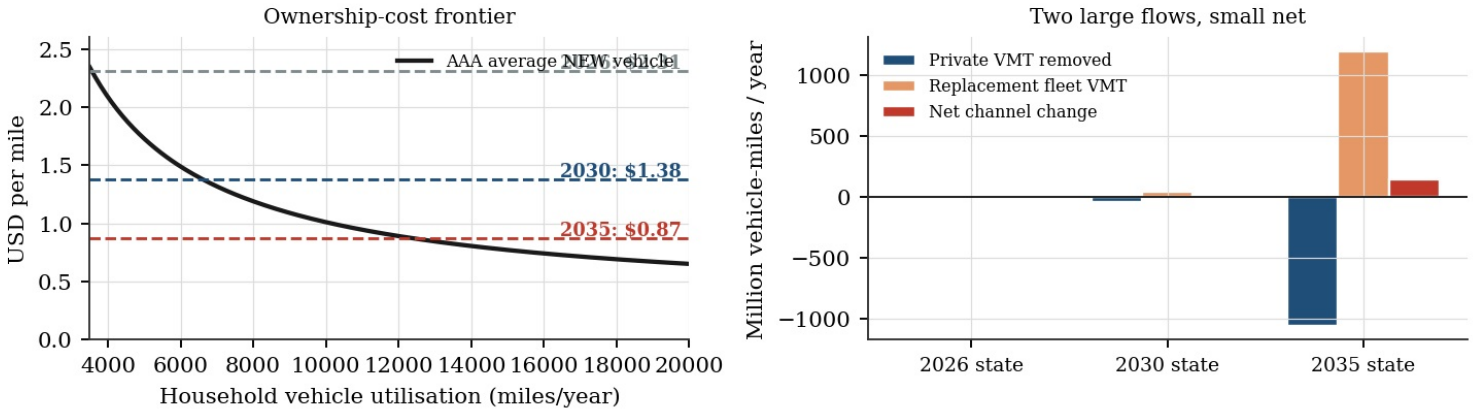

Figure 4. Ownership-cost frontier and the two offsetting VMT flows. Crossing points are conditional on an average-new-vehicle benchmark and the model's retail markup.

State	AV retail	Ownership cost at 6500 mi/yr	Price response	After functional ceiling	Vehicles shed	Net channel VMT as % of regional
2026 state	$2.31	$1.40	0.0%	0.0%	0	0.00%
2030 state	$1.38	$1.40	2.9%	2.9%	6249	0.02%
2035 state	$0.87	$1.40	75.0%	75.0%	162532	0.48%

Three cautions belong with these figures. The first is that the response curve is a HEURISTIC: percentage cost advantage divided by 50%, capped at 100%, times a realisation rate. Shed counts are outputs of that rule, not forecasts or observed preferences. The second is that the empirical literature on ride-hailing and vehicle ownership is sign-ambiguous (§3.1): the best-identified studies disagree with each other, and in one case the same research team reaches opposite signs at different levels of geographic aggregation [L3, L4]. Those studies observe prices far above ownership parity and therefore do not test this mechanism, but they equally do not support it. The third caution is more concrete, and is best stated through the household the model handles worst.

> **Eleanor.** She is retired, drives about 4,000 miles a year, and owns a twelve-year-old sedan she finished paying for a decade ago. On the average-new-vehicle benchmark used above she is the *most* attractive candidate in the region for shedding a vehicle: her utilisation is the lowest, so her modelled cost per mile is the highest, roughly $2.09. Her actual cost is nothing like that. There is no finance payment, the depreciation that remains is trivial, and her insurance is a fraction of a new vehicle's. Registration, insurance and repairs come to something on the order of $1400 a year, and with fuel and maintenance she is driving for closer to $0.70 a mile. The benchmark overstates her cost by roughly 3.0 times.

Eleanor is not an edge case. The vehicles most likely to be lightly driven are disproportionately older and paid off, which is precisely the population for which an average-new-vehicle benchmark is least appropriate [S11]. The model's response curve therefore systematically overstates how attractive shedding looks, and the overstatement is largest exactly where the curve reports the biggest opportunity. The shape of the frontier is robust; its position is not, and it is likely too far to the right.

There is a second thing Eleanor illustrates, running the other way. Reliable per-mile mobility without a driving licence is worth considerably more to her than the fare difference suggests, and it becomes worth more still on the day she stops driving. That gain is real, it is concentrated among older and disabled residents, and this model captures it only as a price response, which means the welfare accounting in Section 8 understates it. The two biases do not cancel; they operate on different populations and in different directions.

6.1 Why the empty-mile share decides the sign

Displacing a private vehicle removes its miles from the network but returns most of those trips as fleet miles, which carry an empty-mile penalty the private vehicle did not. Whether the exchange is favourable turns on a single ratio, and the algebra is unusually clean:

> neutral deadhead ratio = (1 ÷ trip retention) − 1 ⟹ at 80% retention, 0.25 empty miles per occupied mile, or 20% of fleet vehicle-miles

Finding 4. Replacement fleet travel exactly offsets shed private travel when the empty-mile share is 20%. Modelled empty-mile shares are 35.5%, 32.4%, 29.6% across the three states, so the ownership channel raises vehicle travel throughout. This is an identity within one channel, it says nothing about induced demand, transit substitution, or total regional VMT, but within that channel it is exact, and it identifies the operational variable policy can actually reach.

The practical implication is that the interesting policy target is not the vehicle, the technology, or the trip. It is the empty mile: measurable in principle, unmeasured in Colorado at present, and directly priceable once measured.

7. Labour

Figure 5. Labour structure, levels are assumption-sensitive, directions are not

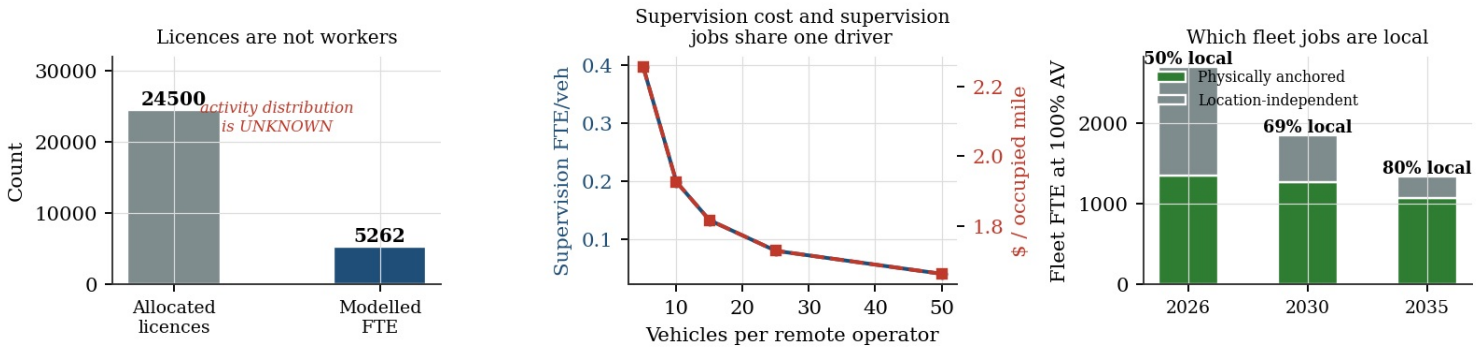

Figure 5. Licence counts against modelled full-time equivalents; the supervision cost and employment identity; and the split between physically anchored and location-independent fleet roles.

Finding 5. Two directions are robust: at a fixed ride count, removing the driver removes human driver-hours, and improving remote-assistance efficiency reduces supervision jobs per vehicle. Neither the level nor the net sign is robust, because ride growth, utilisation, active-driver hours, physical staffing, outsourcing and service geography are all weakly observed. This paper advances no net-employment claim, and readers should be sceptical of anyone who does.

What can be said about the *location* of the work is considerably more durable, and it is the part most relevant to regional economic policy. Remote assistance is not real-time control. The vehicle encounters an ambiguous situation, requests guidance, and a human supplies a high-level input; the vehicle continues to drive itself throughout. Latency tolerance is therefore high and physical presence unnecessary, which makes the function location-independent and, in principle, offshorable. Depot turnaround, field response and recovery, high-voltage maintenance, sensor calibration and charging operations must happen where the vehicles are.

The locally captured share of fleet employment consequently rises across the technology states, from 50% to 80%. This is not the good news it appears to be. The offshorable share shrinks not because it is retained but because it is eliminated: the jobs remaining in the region are the ones automation cannot remove, not the ones policy successfully anchored. Workforce strategy should therefore target the physically anchored functions, which are durable, rather than teleoperations, which is simultaneously the largest near-term category and the one most exposed to both automation and relocation. Training displaced drivers into the job the cost curve exists to eliminate would be an expensive and predictable error.

7.1 Small fleets face a penalty, not a prohibition

Fleet	Operators required	Effective ratio	Cost per occupied mile
7	1	1:7.0	$2.428
15	1	1:15.0	$1.985
30	2	1:15.0	$1.985
50	3	1:16.7	$1.946
100	6	1:16.7	$1.946
500	30	1:16.7	$1.946
2,000	117	1:17.1	$1.937
12,000	688	1:17.4	$1.930

Distributed or cooperative fleet ownership is a recurring proposal, and the supervision layer is where it runs into difficulty. A standalone operation offering continuous service must staff an integer number of supervisors and cover stochastic demand peaks, so a seven-vehicle operator achieves an effective one-to-seven ratio regardless of what ratio the technology supports. The penalty is substantial at that scale and largely gone by roughly fifty vehicles.

This is a real constraint but not a theorem of impossibility, and the distinction matters for anyone designing around it. Pooled supervision across operators, purchased supervision as a service, restricted operating windows, or lower intervention rates each move the threshold. The layer where distributed ownership remains straightforwardly viable is the physically anchored one, depot services, field response, maintenance, which is capital-light relative to vehicles, carries no intervention liability, and requires a licence to no one's autonomy stack.

8. Fiscal exposure and economic interpretation

Exposure or instrument	Illustration	Reading
Prearranged ride fee gap	$9.4M/yr	Only in the bounded full-AV scenario, only if autonomous rides fall outside the operative category, and at a rate last verified through June 2026 [S6]. **Rate requires refresh.**

Gasoline excise	$9.3M/yr	Conditional ownership scenario; excludes other fuel-related fees and behavioural offsets.
Flat $0.50 per trip	$26.1M/yr	Revenue comparator; behaviourally inert.
Uniform $0.10 per vehicle-mile	$44.4M/yr	Prices all for-hire road use equally.
$0.05 occupied + $0.35 empty	$61.6M/yr	Targets the variable identified in Finding 4, but requires audited occupancy data and instrument-specific legal review.

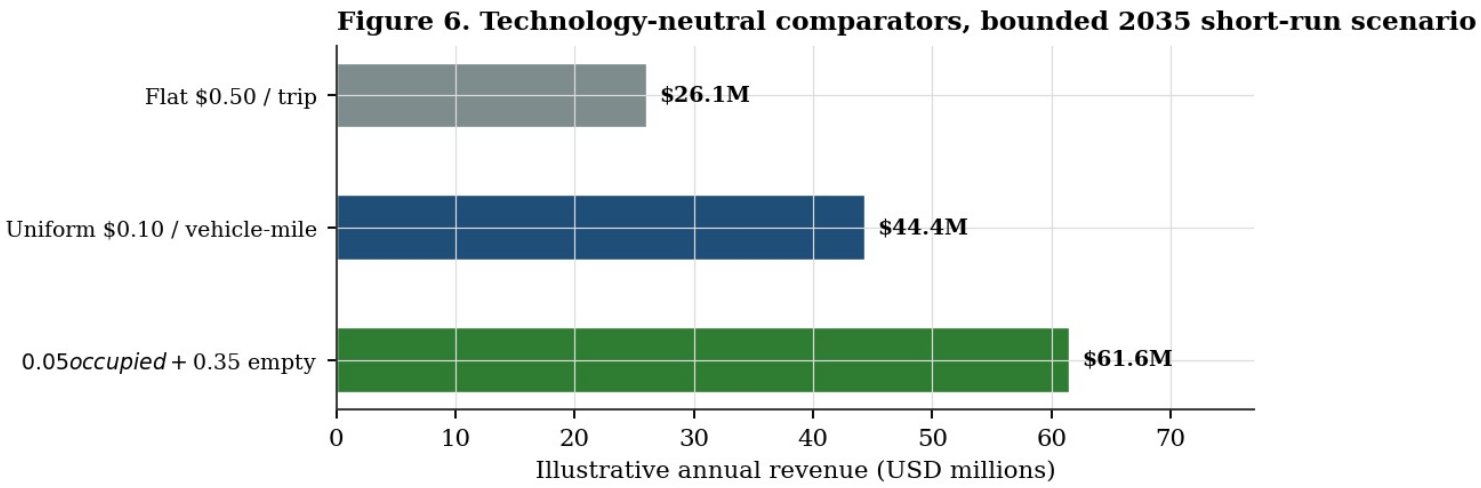

Figure 6. Comparators drawn from the same bounded scenario and the same rate registry as the text. Revenue is illustrative and does not establish legal validity.

The magnitudes are worth stating plainly: tens of millions annually against a state transport budget in the billions. This is not a fiscal crisis, and presenting it as one would be counterproductive. What distinguishes the exposure is that it is structural rather than cyclical, it does not recover, and that it is concentrated in the revenue streams that fund the road network the fleet uses.

8.1 Transit exposure is not a farebox question

RTD's fare revenue is a small fraction of its budget beside sales and use tax, so direct farebox loss is the least of the transit channels and the one most often cited. The more consequential exposures are indirect: ridership drives federal formula funding, affects network productivity, sustains the political constituency for a sales-tax levy, and determines service for riders without alternatives. This paper does not estimate those effects and should not be read as having bounded them.

8.2 On the temptation to estimate economic impact

Any policy audience will reasonably ask how large this transition is economically. The direct answer is that it is small: no channel modelled here approaches macroeconomic significance against a metropolitan economy of roughly $312 billion [S13]. The more useful answer is that gross product is a poor instrument for this particular question, because several welfare-relevant effects enter it with the wrong sign or not at all.

> **The Delaneys, counted.** Suppose they sell the second car. They stop buying its fuel, its insurance, its oil changes, its tyres, its registration, and eventually its replacement. Every one of those was output, and a good deal of it was income to someone in the region, a mechanic in Arvada, an insurance agency, a tyre shop. They now buy mobility instead, but less of it in dollar terms. Their position improves and measured regional output falls. This is not a paradox to be resolved; it is a statement about what the measure was built to count.

A household that sheds a vehicle stops purchasing fuel, insurance, repair, parking and eventually the vehicle, all of which are output. Netted against the replacement mobility it must now buy, that household releases roughly $4542 annually ($9084 avoided against $4542 substituted), which is an improvement in its position and a reduction in measured output. Consumer surplus from lower prices does not appear in output statistics at all; on the model's demand assumptions it lies between $712M and $1822M annually, a range wide enough to be read as an order of magnitude rather than an estimate. And gross product is entirely indifferent to the composition change, driver labour income becoming fleet capital income, that this paper argues is the substance of the transition.

> Reported as a BAND, never a point. The estimate extrapolates a constant-elasticity curve far outside any calibration range and is not a welfare conclusion: producer surplus, taxes, externalities and distribution are excluded. It is retained rather than deleted because a legislative audience will otherwise supply its own number.

9. Policy options, sorted by the authority required

The useful responses to this transition are distributed across branches and agencies, and several of the highest-value steps do not require legislation. This section therefore sorts options by who can act rather than by topic. Two organising principles run through all of them.

The first is a dependency logic. Measurement precedes functional definition, which precedes any attempt to price an externality. Each step is largely unadministrable until the one above it exists, so the ordering is not a ranking by ambition but by feasibility.

The second is technology neutrality. Colorado's preemption clause bars applying standards to an automated driving system that differ from those applied to a human driver [S4]. Any instrument whose incidence depends on a vehicle being autonomous carries serious risk. Every option below is drafted to attach to compensated passenger service or to road use, which is both more legally durable and, as it happens, better policy: the externalities at issue are not caused by autonomy.

9.1 Available to the executive without new legislation

These are the near-term options, and collectively they would resolve most of this paper's evidentiary limitations within a single budget cycle.

Action	Vehicle	Why it matters	Constraint
Task the Autonomous Mobility Task Force with a data framework. Direct it to negotiate voluntary, privacy-protected trip-level data sharing with for-hire operators, on the Chicago schema [S14].	Executive direction The Task Force already exists in statute under SB 17-213, no new body required.	Every disputed quantity in this paper is reconstructed rather than measured. A voluntary framework can be standing before any bill is drafted, and produces the evidence a bill would need.	Voluntary participation is not guaranteed; the framework should be designed so that a later statutory mandate can adopt it unchanged.
Join registration to emissions records. Query DMV registration [S18] against Air Care emissions records [S17] on VIN to produce a functional-class by annual-mileage distribution of the regional household fleet.	Agency direction Both systems are already jointly administered by CDPHE and the Department of Revenue, and already linked by VIN.	Replaces the weakest parameter in this analysis: the share of lightly-driven vehicles that are functionally substitutable, with an estimate. Uniquely among the actions here, it requires **nothing at all from the fleet operators**: no mandate, no negotiation, no voluntary framework. It is the state querying its own records.	Confirm whether odometer is captured at inspection; if not, adding the field is a near-zero-cost change to an existing process. Release at ZIP or tract level by body class and mileage band to manage re-identification. EV exemption means coverage declines over time, run it soon.
Publish the active-driver distribution. Direct CDLE to report hours and earnings distribution, not only licence counts and aggregate task totals.	Agency direction Data already collected under the TNC Transparency Act.	Resolves §3.2 directly. Without it, workforce programmes are sized against licences and will be mis-scaled in a direction the state cannot currently determine.	Requires aggregation thresholds to protect individual drivers; genuinely low cost otherwise.
Seek a joint classification determination. Convene PUC, Department of Revenue and the Attorney General to state, prospectively, how driverless for-hire service is treated under Title 40 and the prearranged-ride fee.	Convening Interpretive, not legislative.	Creates a record and removes ambiguity for operators and incumbents alike, whichever way it lands. Also identifies precisely what statutory change is needed.	**The PUC is exercising quasi-judicial authority in a contested proceeding.** Prospective guidance on classification generally is appropriate; direction on the pending vehicle waiver is not, and the distinction should be observed carefully.
Order a labour-defined category audit. Direct each principal department to identify statutes and rules whose operative category presumes a human worker, and report which are load-bearing.	Executive order Departments know their own statutes.	Operationalises Finding 6. Cheap, requires no position on autonomous vehicles, and produces a prioritised list for the legislature rather than an open-ended request.	Findings are only as good as departmental engagement; a reporting deadline and a named coordinating office matter.
Refresh fee rates and model exposure. Direct the Department of Revenue to publish current prearranged-ride fee rates and OSPB to carry the exposure in fiscal outlooks.	Agency direction	Rates in this paper are last verified through June 2026 and are marked accordingly [S6]. No legislative scoring should proceed on stale figures.	Routine.

9.2 Requiring the General Assembly

Action	Vehicle	Why it matters	Constraint
Mandate trip-level reporting. Require privacy-protected reporting from all for-hire passenger operators: zones, times, occupied and empty mileage, fare, service availability.	Statute	Converts the voluntary framework of §9.1 into a durable obligation, and is a precondition for any mileage instrument.	Privacy design is the substantive work: aggregation thresholds, suppression in low-density zones, audit rights and release rules. Far easier to impose on an entrant than on an established industry.
Define for-hire service functionally. Amend Title 40 and the prearranged-ride fee so coverage turns on compensated passenger service rather than on the presence of a human driver.	Statute	Addresses the category-substitution dynamic of Finding 1 at its source, and forecloses the quiet-lapse mechanism rather than patching one instance of it.	Must be drafted so it does not bite only on autonomous services, which would invite a preemption challenge. Functional definition is the safeguard.
Authorise an empty-mile pilot.	Statute	The only instrument modelled here that prices the	Requires the reporting above to be administrable at

Permit a provider-neutral commercial for-hire mileage charge with a higher rate on verified empty miles, structured as a dedicated enterprise fee.		variable determining the sign of the ownership channel, and it leaves the operator free to choose how to comply.	all. Preemption, TABOR nexus, privacy, equity and administrative cost each need instrument-specific review. Incumbent human operators bear it first and will say so.
Fund sliding-scale transition support. Base assistance on verified earnings and driving intensity rather than a universal payment or a hard hours threshold.	Appropriation	The active-driver distribution is unknown, so both uniform payments and cliff-edge tests are poorly targeted. A hard twenty-five-hour threshold would decide Marcus's case on a technicality about which weeks were counted.	Should follow §9.1's CDLE publication, not precede it. Direct training at demonstrated local hiring demand, not at supervision roles the cost curve exists to eliminate.

9.3 Federal and interstate positioning

Colorado holds an unusual position that is worth using deliberately rather than incidentally. Most accumulated autonomous-vehicle operating experience comes from warm, dry, grid-planned metropolitan areas. A cold-weather, polycentric, mountain-adjacent market generates evidence that those places structurally cannot, and it does so first.

- **Federal engagement.** Colorado has standing to argue for winter-performance and weather-suspension reporting in federal crash and performance frameworks, on the straightforward basis that it will hold data no Sun Belt state can produce. This is a Governor's-office function rather than a legislative one, and it is materially strengthened if §9.1's data framework exists first, advocacy backed by a dataset is a different proposition from advocacy backed by an argument.
- **Cold-weather state coordination.** Colorado, Michigan, Minnesota, Illinois and comparable states share an interest in winter operating standards, snow-condition service suspension, and cold-weather sensor performance that early-deployment Sun Belt states do not. A working group among those states is a governors' vehicle, needs no legislation, and would give Colorado a leadership position it cannot obtain acting alone.
- **Schema interoperability.** Adopting a reporting schema substantially compatible with Chicago's [S14] costs little and makes Colorado's data immediately comparable to the largest open corpus in the country, which raises its research value and its usefulness in federal proceedings. Inventing a bespoke schema forfeits both for no gain.

9.4 Regional and municipal

Local authority here is narrower than commonly assumed, and the boundary is worth stating precisely so that effort is not wasted.

- **Proprietary terms at airport property.** Denver International Airport can set commercial access terms for dwell, staging, curb use and empty circulation as a property owner rather than a regulator, and provider-neutral terms would reach one of the region's densest trip generators. Municipal, proprietary The proprietary/regulatory boundary is litigable and untested in Colorado on these facts; obtain instrument-specific review and monitor diversion to off-airport pickup.
- **Regional data participation.** DRCOG's travel modelling and CDOT's planning functions are the natural consumers of any trip-level dataset, and involving them in schema design raises the probability the data answers questions planners actually have. Regional
- **Curb management in the public right-of-way** is regulatory territory and carries materially higher preemption risk than airport property. It is not foreclosed, but it should be approached technology-neutrally and with counsel, not by analogy to the airport case.

9.5 Not recommended, and why

- **An autonomy-specific tax or permit.** High preemption risk under a statute barring different standards for automated systems, and unnecessary when the policy target can be defined by service or road use instead.
- **A blanket municipal conclusion in either direction.** Some tools are preempted; others may rest on general police power, commercial-carrier law, contracting or proprietary capacity. Neither "cities can act" nor "cities cannot act" is a defensible summary.
- **Mandated human operators.** Preserves a cost category by blocking the technology rather than addressing safety, distribution or transition, and does not survive contact with the labour structure in Section 7.
- **A firm minimum fleet size.** The threshold is assumption-dependent (§7.1) and would ossify a number that legitimate pooling arrangements can change.
- **Pricing before measuring.** An empty-mile fee without audited occupancy data is unadministrable and would discredit the instrument for a decade.
- **A large welfare headline in either direction.** Demand form, mode substitution, safety and replacement costs are individually too uncertain to support one, and jointly far too uncertain.

9.6 What the sequence looks like in practice

If the analysis here is right, the next twelve months are unusually cheap and unusually consequential. The executive can stand up a voluntary data framework through a body that already exists, publish labour data the state already collects, and obtain a prospective classification determination from agencies that already have the authority to give one. None of that requires a bill, an appropriation, or a position on whether autonomous vehicles are desirable. All of it produces exactly the evidence a legislature would need in a subsequent session to define categories functionally and, if the evidence supports it, to pilot a road-use instrument.

The alternative sequence, legislating first and measuring afterwards: is available and has been attempted elsewhere. It tends to produce instruments that cannot be administered, obligations calibrated to the wrong population, and a subsequent decade of arguing about numbers nobody collected.

10. What this case says about deploying automation generally

Autonomous ride-hail is one of the few contemporary applications of machine autonomy that is simultaneously at commercial scale, safety-critical, and gated by physical capital. That combination makes it unusually legible. A fleet cannot outrun vehicle manufacturing, depot construction or charging infrastructure, so the transition proceeds slowly enough to be watched, which is not true of automation mediated entirely by software, where the same economic logic can play out in a quarter rather than a decade.

Four structural observations from this case appear to us to generalise. They are offered as hypotheses for testing elsewhere, not as findings established here.

1. **Residual human oversight, not autonomous capability, tends to set the economics.** Wherever an automated system performs work for which a human retains accountability, the ratio of instances to supervisors enters unit cost directly. That ratio is typically undisclosed and, as §4.1 shows, frequently ambiguous even when disclosed, because the denominator goes unspecified. An industry can discuss its own economics for years without that number ever being pinned down.
2. **Oversight cost and oversight employment are the same quantity.** Where the dominant residual human role is also the dominant residual cost, no strategy delivers both a low unit cost and a large oversight workforce. This is an identity rather than a tendency, and it disposes of a common proposition, that automation will be fine because it creates supervisory jobs, without requiring any empirical dispute.
3. **Physical capital sets the clock.** This case is a *slow bound* on automation-driven labour transition, not a typical one. It should be read as a rehearsal rather than a reassurance: the institutional adaptations it demands, measurement, functional definitions, targeted transition support: are precisely the ones that will be needed faster, and with less warning, wherever no capital constraint applies.
4. **Statutory categories built around a human worker do not fail loudly.** They fail by substitution. An operator finds the nearest existing category and proceeds; no gap is declared, no proceeding opened, and the statutory purpose lapses silently. Here the signal came from a competitor with heavier obligations rather than from the regulator or the public, which is a fragile and accidental detection mechanism.

These four observations are developed into a scoring instrument in Appendix D, so that a researcher working in another sector can locate their own case relative to this one rather than reasoning by analogy. The exercise is worth doing before assuming this case is representative: on that instrument, autonomous ride-hail turns out to be unusually *observable* and unusually *unprotected*, a combination which makes it a good venue for building measurement practice and a poor guide to how automation will proceed where neither condition holds.

> **Finding 6.** The fourth observation carries an action inside existing executive authority, requiring no position on autonomous vehicles: audit the statute book for operative categories defined by the presence of a human worker, and establish which are load-bearing (§9.1). The legislative repair then becomes a template rather than a single fix. It requires no forecast about when, or whether, automation reaches any particular sector, only the observation that where it does, these categories will not announce their own obsolescence.

11. Robustness and limitations

Claim	Status	Basis
For-hire activity is under one percent of regional person-trips	Relatively robust	0.65%-0.62% across the full denominator range
Public commercial service was not open at the cut-off	High-confidence as-of statement	Operator and local reporting describe staged employee operations [S1, S2]
20% is the ownership-channel neutral empty share at 80% retention	Analytic	Closed-form identity, within that channel only
Supervision cost and supervision FTE move as one	Identity	Both derive from operator-hours per vehicle
Cost falls across the three states	Arithmetic within assumptions	The states stipulate lower capital cost, insurance, supervision, maintenance, depot cost and deadheading
The largest decline is vehicle capital cost	Computed, not asserted	The companion audit fails if the stated conclusion and the computed ranking diverge
Any state label is an adoption date	Not supported	Labels are conditional state names, retained only for readability
Net employment sign in full-automation cases	Not robust	Depends on ride growth, utilisation, staffing and service geography
Autonomous rides fall outside the prearranged-ride fee	Not settled	Textual ambiguity is real; no reviewed agency or court decision resolves it
Any mileage instrument is lawful	Not established	Technology neutrality helps; preemption, TABOR, nexus, privacy and administration each require review

11.1 Principal limitations

1. **No Denver trip-level data.** The binding constraint on everything above, and the reason measurement leads Section 9.
2. **No public operating record.** Cost, retail markup and utilisation are conditional assumptions rather than observations, and cannot be validated in this market at this time.
3. **No fare calibration.** The human benchmark is a band (§3.3), and it is the single largest source of uncertainty in every comparison the paper makes.

4. **Fee rates are stale** for the current annual period and must be refreshed before any fiscal scoring [S6].
5. **Demand form is weak at large price changes** which is why the guardrail exists. The guardrail itself is transparent but arbitrary.
6. **The ownership response is a heuristic** and the AAA benchmark describes an average new vehicle rather than the older vehicles most likely to be shed.
7. **The functional substitutability ceiling is an assumption where it could be an estimate.** The 80% figure is the weakest number in this paper. Unlike most of the others it is not limited by missing data, the registration and emissions records that would establish it already exist and are already linked (§6, §9.1). This limitation is the one most cheaply removed, and the paper's ownership-channel results should be treated as provisional until it is.
8. **The Denver empirical anchor is thin.** The deadheading and mode-replacement evidence [S10] rests on a small single-driver sample whose authors present their figure as a lower bound.
9. **Winter reliability is unmodelled.** Snow, cleaning, sensor performance and weather-related suspension would raise cost and reduce utilisation, biasing these estimates optimistically in a cold-weather market, and it is the gap §9.3 argues Colorado is best placed to close.
10. **Safety is excluded entirely.** Exposure-adjusted effects could exceed every fiscal magnitude here, in either direction, and their omission is the largest single gap in the welfare accounting.
11. **This is not a legal opinion.** Classification and instrument design require Colorado counsel and agency interpretation.

12. Conclusion

Denver-Boulder is a useful case precisely because it is early. The deployment is real enough to make policy timely and immature enough to make confident market forecasting obviously premature, a combination rarer than it sounds, since most jurisdictions get one condition or the other.

Scale the near-term claim correctly. The for-hire market is under one percent of regional person-trips. Its automation is consequential for those directly involved and is not, on its own, a metropolitan transformation. Getting this wrong in either direction produces bad policy: overstatement invites crisis measures against a small base, and understatement misses the ownership channel entirely.

Treat the ownership frontier as the consequential scenario. If per-mile cost falls below the cost of owning a lightly-used vehicle and service becomes broadly reliable, household vehicle shedding reorganises transport spending, land use and the fiscal base together. Timing and magnitude remain genuinely uncertain. The sign of the vehicle-travel effect within that channel does not: it turns on the empty-mile share, a measurable operational variable that Colorado currently does not measure.

Use the cheap window. The most valuable steps available in the next year sit inside existing executive authority, a data framework through a body already constituted in statute, labour data already collected, a classification determination from agencies that already have the power to give one, and a join of registration to emissions records that would replace this paper's weakest parameter with an estimate while requiring nothing at all from the industry. They cost little, require no position on whether autonomous ride-hail is desirable, and generate precisely the evidence that later statutory choices will require.

It is worth holding the three households together at the end, because they are not on the same timetable and no single instrument reaches all of them. Marcus is exposed now, in a market that is under one percent of regional trips and therefore easy to overlook, and what he needs is a transition test that can distinguish him from a dormant licence. The Delaneys are exposed later, at a threshold this paper can locate only approximately, and what they change when they cross it is much larger than the market Marcus works in. Ray is never exposed in this channel at all: his truck is a tool, and the model's own ceiling now says so rather than counting him among the converts. Eleanor may never be exposed in the way the model predicts, her car is cheap in a way the benchmark cannot see, while being the person for whom the service, if it arrives reliably, is worth the most. A policy framework that gets the aggregate right and these three wrong has not got very much right.

The most transferable lesson is not about vehicles. A fee that attaches to services connecting riders to *drivers* does not generate a controversy when there is no driver. It generates a quiet search for the nearest available category, and in this instance the search was noticed not by the regulator or the public but by a competitor who still bore the heavier obligations. Legal and economic frameworks that define a service by the presence of a human worker should be audited before that search begins, not after, and the audit costs almost nothing compared to discovering, some years later, which of them stopped functioning and when.

Appendix A. Parameter registry

Evidence tiers: **A** official or administrative · **B** peer-reviewed · **C** industry or reputable reporting · **D** author scenario assumption. Tier D values are not observations, and each carries an explicit range.

Parameter	Central	Unit	Range	Tier	Source	Note
pop_drcog_region	3.3e+06	persons	3.15e+06 - 3.45e+06	C	S9	DRCOG NINE-COUNTY region. Distinct from the RTD service area of 3,098,079 (S8); conflating the two shrinks the denominator and inflates every share.
pop_rtd_service_area	3.098e+06	persons		A	S8	Reported for transit analysis only. Not the regional denominator.
person_trips_per_capita_day	3.4	trips/person/day	3 - 3.8	D	SA	
vmt_per_capita_day	25	vehicle-miles/person/day	21 - 29	D	SA	Largest single uncertainty in the regional denominator.
tnc_tasks_colorado	3.723e+07	transportation tasks/yr		A	S7	Statewide official count, most recent reporting year.
tnc_driver_licences_colorado	3.5e+04	licensed drivers	3e+04 - 4e+04	A	S7	LICENCES, not verified active workers. See activity caveat.
metro_share_of_state_tnc	0.7	share	0.6 - 0.8	D	SA	
tnc_trip_occupied_miles	6	occupied miles/ride	5 - 7	D	SA	
tnc_trip_driver_hours	0.42	driver-hours/ride	0.35 - 0.5	D	SA	
deadhead_ratio_human	0.69	empty miles/occupied mile	0.55 - 0.85	B	S10	
transit_substitution_share	0.34	share of net-new rides	0.25 - 0.5	B	S10	
driver_net_labour_hourly	15	USD/hr net of vehicle cost	12 - 20	D	SA	No Denver fare observation exists, and fares from other metropolitan areas (S15) are not transferable. This is the model's single largest source of uncertainty, and the benchmark is therefore reported as a BAND rather than a point.
driver_vehicle_cost_per_total_mile	0.76	USD/vehicle-mile	0.6 - 0.9	A	S12	IRS business rate as a proxy. A tax convention, not an observed cost.
platform_take_rate	0.25	share of pre-tip fare	0.2 - 0.32	D	SA	
tip_rate	0.175	share of pre-tip fare	0.1 - 0.2	D	SA	
av_vehicle_capex	9e+04	USD/equipped vehicle	7e+04 - 1.5e+05	D	SA	
av_vehicle_life_miles	3e+05	total vehicle-miles	2e+05 - 4e+05	D	SA	
av_paid_miles_per_day	110	occupied miles/vehicle/day	70 - 160	D	SA	
av_energy_kwh_per_total_mile	0.35	kWh/vehicle-mile	0.28 - 0.45	C	SA	
av_electricity_cost	0.13	USD/kWh	0.1 - 0.18	D	SA	
av_maintenance_per_total_mile	0.2	USD/vehicle-mile	0.12 - 0.35	D	SA	
av_insurance_per_total_mile	0.25	USD/vehicle-mile	0.1 - 0.5	D	SA	
av_vehicles_per_remote_operator	10	vehicles/operator	3 - 25	D	SA	No operator discloses this, and the DENOMINATOR is ambiguous: fleet size or concurrently active vehicles are different claims.
av_remote_operator_cost_hourly	32	USD/loaded hour	26 - 42	D	SA	
av_service_hours_per_day	20	hours/day	16 - 22	D	SA	
av_depot_cost_per_vehicle_day	40	USD/vehicle/day	25 - 70	D	SA	
deadhead_ratio_av	0.55	empty miles/occupied mile	0.35 - 0.8	D	SA	
av_retail_markup	0.2	share above operating cost	0.1 - 0.4	D	SA	
av_urban_speed_mph	20	mph	16 - 24	D	SA	
demand_peak_ratio	2.1	peak/average concurrent	1.7 - 2.5	D	SA	
supervision_shift_flex	0.7	share of peak-average gap schedulable	0.5 - 0.85	D	SA	
supervision_z	1.645	std deviations (P95)		A	SA	
local_physical_fte_per_vehicle	0.2	FTE/vehicle	0.13 - 0.28	D	SA	Depot turnaround, field response, high-voltage maintenance, calibration, charging O&M.

fte_hours_per_year	2080	hours/FTE/yr		A	SA	
demand_elasticity	-1.2	own-price elasticity	-2 - -0.8	D	SA	
short_run_expansion_cap	2	multiple of baseline rides	1.5 - 3	D	SA	GUARDRAIL, not an estimate. When it binds, the scenario is uninformative and is labelled as such rather than reported as a result.
aaa_ownership_cost_annual	1.158e+04	USD/yr	1e+04 - 1.35e+04	C	S11	Average NEW vehicle. Overstates the cost of older paid-off vehicles.
aaa_assumed_annual_miles	1.5e+04	miles/yr		C	S11	
aaa_fixed_cost_share	0.62	share of annual cost	0.55 - 0.7	D	S11	
hh_vehicles_drcog	1.72e+06	household vehicles	1.5e+06 - 1.95e+06	D	SA	
low_mileage_hh_vehicle_share	0.28	share	0.2 - 0.36	D	SA	
low_mileage_annual_miles	6500	miles/yr	5000 - 8000	D	SA	
functional_substitutability_ceiling	0.8	share of candidate vehicles a fleet can replace	0.65 - 0.9	D	SA	HARD functional limit, distinct from behavioural friction. Purpose-built work vehicles, towing and hauling capability, equipment and cargo carriage, unplowed and unpaved access, and trips outside any plausible service area or service window cannot be served by a standard fleet vehicle at ANY price. Low mileage does not imply substitutability. MEASURABLE, NOT INHERENTLY ASSUMED: a join of DMV registration records [S18] to Air Care emissions records [S17] on VIN would yield the functional-class by annual-mileage distribution directly. This central value should be replaced by an estimate as soon as that join is run. See policy recommendation on registration data.
shed_realization_rate	0.45	share of attractive sheds realised	0.2 - 0.65	D	SA	BEHAVIOURAL friction among households that could substitute: habit, garaging, trip-chaining, uncertainty. Applied after the functional ceiling, never instead of it.
replacement_trip_retention	0.8	share of shed VMT returning as occupied	0.65 - 0.95	D	SA	
prearranged_ride_fee_standard	0.34	USD/ride		A	S6	Last verified through 30 June 2026. Rates publish annually. REFRESH BEFORE USE. **REFRESH**
prearranged_ride_fee_zev	0.18	USD/ride		A	S6	Last verified through 30 June 2026. REFRESH BEFORE USE. **REFRESH**
co_gasoline_excise_per_gallon	0.22	USD/gallon		A	S6	
fleet_mpg	25	miles/gallon	22 - 28	D	SA	
metro_gdp	3.119e+11	USD/yr		A	S13	Scale denominator only. Explicitly NOT a welfare measure.
policy_flat_trip_fee	0.5	USD/ride	0.25 - 0.75	D	SA	
policy_uniform_vmt_fee	0.1	USD/vehicle-mile	0.05 - 0.2	D	SA	
policy_occupied_mile_fee	0.05	USD/occupied mile	0 - 0.15	D	SA	
policy_empty_mile_fee	0.35	USD/empty mile	0.15 - 0.75	D	SA	

Conditional technology states

State labels are names, not dates. They define the conditions under which the consequences described above follow, whenever those conditions obtain.

Parameter	2026 state	2030 state	2035 state
av_vehicle_capex	9e+04	5.5e+04	3.5e+04
av_vehicle_life_miles	3e+05	3.75e+05	4.5e+05
av_paid_miles_per_day	110	135	160
av_vehicles_per_remote_operator	10	25	50
av_maintenance_per_total_mile	0.2	0.16	0.12
av_insurance_per_total_mile	0.25	0.18	0.12
av_depot_cost_per_vehicle_day	40	32	25
deadhead_ratio_av	0.55	0.48	0.42

Appendix B. Academic literature cited

Positioned claim by claim in §3.1. Findings are summarised as used here; readers should consult the originals rather than relying on these summaries.

ID	Reference	Finding as used, and its limits
L1	Bösch, Becker, Becker & Axhausen (2018), Cost-based analysis of autonomous mobility services, Transport Policy 64:76-91	Canonical bottom-up AV cost analysis. Finds service costs, not vehicle costs, are the main cost factor for shared autonomous fleets, and that private vehicle ownership will likely endure. Reports commercial-fleet active time of 57% peak, 59% off-peak, 30% night. Swiss cost base; not directly transferable in levels.
L2	Becker, Becker, Abe, et al. (2020), Impact of vehicle automation and electric propulsion on production costs for mobility services worldwide, Transportation Research Part A 138:105-126	Multi-country extension of the cost-based approach. Establishes that automated fleet cost structures vary substantially by market, which is why cost figures should not be transferred across jurisdictions.
L3	Ward, Michalek, Azevedo, Samaras & Ferreira (2019), Effects of on-demand ridesourcing on vehicle ownership, fuel consumption, VMT and emissions per capita in U.S. states, Transp. Res. Part C 108:289-301	State-level difference-in-differences: ~3% average decline in per-capita vehicle registrations following Uber and Lyft entry, 2005-2015. Found no significant state-level TNC-VMT relationship.
L4	Ward, Michalek, Samaras, Azevedo, Henao, Rames & Wenzel (2021), The impact of Uber and Lyft on vehicle ownership, fuel economy, and transit across U.S. cities, iScience 24(1)	CITY-level analysis by substantially the same authors reaching the OPPOSITE sign: ~0.7% increase in vehicle registrations, larger in cities with prior high motorisation. The authors note state-level aggregation reverses the sign. Central caution for any ownership-channel estimate.
L5	Wang, Shi & Chen (2021), Impact of ride-hailing usage on vehicle ownership in the United States, Transp. Res. Part D 101:103085	Finds regular and active ride-hailing users own fewer vehicles. Usage-based rather than availability-based identification, with attendant endogeneity.
L6	Erhardt, Roy, Cooper, Sana, Chen & Castiglione (2019), Do transportation network companies decrease or increase congestion? Science Advances 5(5)	San Francisco: TNCs materially increased congestion. Supports the deadhead mechanism central to Section 6.
L7	Wadud, MacKenzie & Leiby (2016), Help or hindrance? The travel, energy and carbon impacts of highly automated vehicles, Transportation Research Part A 86:1-18	Establishes the rebound framework: automation lowers the generalised cost of travel and may raise VMT even as it improves vehicle efficiency.
L8	Taiebat, Stolper & Xu (2019), Forecasting the impact of connected and automated vehicles on energy use, Applied Energy 247	Income-stratified rebound estimates for household VMT under automation. Wide range; upper estimates are large.
L9	Harper, Hendrickson, Mangones & Samaras (2016), Estimating potential increases in travel with autonomous vehicles for the non-driving, elderly and people with travel-restrictive medical conditions, Transportation Research Part C 72	Access-driven demand from populations that cannot currently drive. The welfare channel this paper's elasticity framework understates.
L10	California Air Resources Board (2019), SB 1014 Clean Miles Standard, base-year emissions inventory	Administrative estimate that roughly 38.5% of TNC vehicle-miles in California are deadhead. Independent corroboration of the Denver figure at [S10].
L11	Choi, Guhathakurta et al. (2022), Ride-hailing and regional VMT growth, Atlanta	Empirical-Bayes counterfactual estimate of ~0.6% average annual regional VMT growth attributable to ride-hailing. Materially smaller than trip-level deadhead studies imply, and a useful upper-bound check.
L12	Cramer & Krueger (2016), Disruptive change in the taxi business: the case of Uber, American Economic Review 106(5)	Capacity-utilisation comparison between taxis and TNCs. Establishes that occupied-time share is the operative productivity variable in for-hire service, the human analogue of this paper's duty-cycle treatment.

Appendix C. Primary sources and use limitations

ID	Source	Use and limitation
S1	Waymo, From the road, 8 July 2026	Denver status: rider-only EMPLOYEE operations announced as upcoming; public access later. Does not establish a public commercial market.
S2	Denverite, Taxi companies want to block Waymo's plans in Denver, 9 July 2026	PUC granted Waymo a luxury limousine permit effective 12 June 2026; Metro Capital Group and Ski Town Transportation intervened against a vehicle waiver for the Ojai.
S3	Colorado PUC, Passenger Carrier FAQ / Transportation	Luxury Limousine is a Limited Regulation Carrier permit: not rate-regulated, annual renewal, subject to operational and safety rules at 4 CCR 723-6.
S4	Colorado General Assembly, SB 17-213 (2017)	Statewide-concern declaration and the different-standards preemption clause.
S5	C.R.S., 40-10.1-602 Definitions	TNC, driver, personal vehicle, prearranged ride definitions.
S6	Colorado Dept of Revenue, Prearranged Ride Fee rate table	Last-verified rates through 30 June 2026. REQUIRES REFRESH.
S7	Colorado Dept of Labor and Employment, TNC annual transportation tasks	Official statewide task counts.
S8	FTA National Transit Database, RTD agency profiles, 2017 and 2024	Audited transit service and financial data. Service-area population 3,098,079 is the RTD SERVICE AREA, which is NOT the DRCOG nine-county region.
S9	DRCOG, Regional planning and mobility pages	Nine-county regional population scale and daily VMT benchmark.
S10	Henao and Marshall, The impact of ride-hailing on VMT, Transportation 46 (2019)	Denver-specific deadheading and modal replacement. Small single-driver sample; authors present the deadhead figure as a lower bound.
S11	AAA, Your Driving Costs, 2025	New-vehicle ownership benchmark at 15,000 miles/year. Represents an average NEW vehicle, not the existing fleet.
S12	IRS, Optional standard mileage rates, 2026	Business-mile rate used as a vehicle cost proxy. It is a tax convention, not an observed driver cost.
S13	BEA via FRED, GDP, Denver-Aurora-Lakewood MSA	Scale denominator only. Not used as a welfare measure.
S14	City of Chicago, Transportation Network Providers Trips	Trip-level reporting schema precedent. Not used to generate Colorado values.
S15	TechCrunch / Obi, Waymo-Uber price gap, Jan 2026	EXTERNAL BENCHMARK ONLY. Figures derive from simulated San Francisco Bay Area requests. NOT Denver observations and NOT used for calibration in r3.
S16	Colorado Court of Appeals, Americans for Prosperity v. Colorado, 2025COA46	Upheld the challenged SB 21-260 enterprise structure. Precedent, not blanket approval of any future fee.
S17	CDPHE / Air Care Colorado, Motor vehicle emissions inspection programme	Biennial inspection for model year 1982+ gasoline vehicles after a seven-model-year exemption, across nine Front Range counties. Jointly administered with the Department of Revenue, Division of Motor Vehicles; results upload electronically to the county registration record. EVs are exempt, so coverage degrades as the fleet electrifies.
S18	Colorado DMV, Vehicle registration and title records	VIN, body class, GVWR, drivetrain, fuel type, model year, plate class (passenger, commercial, farm, apportioned) and registration geography. Odometer disclosure is federally required at transfer under ten years.
SA	Author, Scenario assumption	Tier D input. Not an observation. Carries an explicit range.

Appendix D. A transferability framework for automation case studies

This appendix is intended for researchers applying this paper's approach elsewhere. Its purpose is to make the reading-across explicit rather than intuitive: a sector is scored on ten dimensions, the scores resolve to three indices, and the indices imply what kind of case it is and what kind of attention it warrants. Nothing here is specific to vehicles.

The instrument exists because the most common error in reading a case study across is to treat it as representative. This case is not. It sits at an unusual corner of the space, heavily gated by physical capital, and simultaneously among the least institutionally protected workforces in the economy. Read naively as "the slow case", it produces the wrong conclusion for a licensed profession, where capital is light but incumbents can raise barriers that ride-hail drivers never could.

D.1 The dimensions

Each dimension is scored 0 to 4 against the stated anchors. Dimensions are worded so that a higher score always means more of the named quantity; the column marked *sign* records the direction in which the dimension enters its family index, so a researcher can re-weight without re-scoring.

Dimension	Scoring question	Sign	Anchors
Family A, RATE: how fast does displacement proceed once technically feasible?			
`capital_gating`	How much physical capital must be built per unit of displaced labour?	−	**0** None; software deploys on existing infrastructure **2** Moderate; some equipment or facility investment **4** Heavy; vehicles, depots, energy infrastructure gate the rollout
`oversight_requirement`	How much residual human oversight does each automated instance require at maturity?	−	**0** None; output is accepted without human review **2** Sampled or exception-based review **4** Per-instance human accountability is required and cannot be delegated
`labour_share`	What share of pre-automation unit cost is the labour being displaced?	+	**0** Minor component **2** Roughly a third **4** Dominant; the largest single line in unit cost
Family B, FRICTION: what institutionally slows it?			
`occupational_closure`	Is the work protected by licensure, credentialing or professional monopoly?	+	**0** No credential required to perform the work **2** A licence exists but confers no monopoly over the task **4** Statutory scope-of-practice reserves the task to credentialed persons
`collective_representation`	Are the workers collectively organised, and does that bargaining reach deployment?	+	**0** None; workers are atomised or classified as contractors **2** Partial density or sectoral bargaining **4** High density with contractual reach over technology deployment
`liability_concentration`	Does an identifiable party bear accountability severe enough to constrain deployment?	+	**0** Diffuse; errors carry little concentrated legal consequence **2** Commercial liability, insurable and priced **4** Personal, professional or criminal exposure attaches to a named party
Family C, LEGIBILITY: can anyone observe it in time to act?			
`jurisdictional_boundedness`	Is there a government with clear standing over where the work occurs?	+	**0** Borderless; work can be performed anywhere **2** Nominal jurisdiction, weak practical reach **4** Physically situated; a single jurisdiction plainly has authority
`incumbent_asymmetry`	Is there an incumbent bearing obligations the automated entrant escapes?	+	**0** No incumbent, or incumbents face identical obligations **2** Incumbents exist but the asymmetry is minor **4** Incumbents carry material obligations the entrant avoids, and will say so
`transaction_observability`	Are units of output separately priced and administratively recorded?	+	**0** Output is internal to a firm; no per-unit price or record exists **2** Priced but not administratively captured **4** Discrete priced transactions already recorded by a public body
`statutory_anchoring`	Do operative legal categories name the human role being automated?	+	**0** Rules attach to outcomes or entities, not to a person performing a task **2** Mixed; some obligations reference a human role **4** Core obligations are defined by the presence of a named human worker

D.2 Indices and what they mean

> family index = mean of its dimensions, sign-adjusted, normalised to [0, 1]
> surveillance priority = rate × (1 − friction) × (1 − legibility)

The **rate index** asks how quickly displacement proceeds once it is technically possible. The **friction index** asks what institutional structures slow it. The **legibility index** asks whether anyone can observe it happening, and it is the primary axis, because a transition nobody can see cannot be governed regardless of its speed or the resistance it meets.

The priority formula follows directly. The combination that should concern a policymaker is fast, unresisted, and invisible. A sector that is slow and visible may still be consequential, but it will be noticed; a sector that is fast and invisible will not be, and the displacement will be substantially complete before the measurement framework that would have detected it is proposed.

D.3 Reference scorings

The ride-hail row is calibrated against the body of this paper. The remainder are the authors' judgements, offered as anchors to make a researcher's own scoring comparable, and explicitly not as findings about those sectors.

Sector	Rate	Friction	Legibility	Priority	Reading
Customer support operations	0.83	0.17	0.06	**0.651**	HIGHEST PRIORITY, fast, unresisted, and unobservable. Displacement is likely to be substantially complete before any measurement framework exists, and statutory exposure will surface late or not at all. There is no incumbent positioned to raise the alarm.
Contract and document review	0.92	0.33	0.19	**0.497**	HIGHEST PRIORITY, fast, unresisted, and unobservable. Displacement is likely to be substantially complete before any measurement framework exists, and statutory exposure

					will surface late or not at all. There is no incumbent positioned to raise the alarm.
Warehouse picking	0.58	0.17	0.38	**0.304**	SUBSTANTIAL PRIORITY, unresisted and unobservable, with rate uncertain. Establish measurement on the strength of the friction and legibility scores alone; waiting for evidence of displacement means waiting for the thing measurement was supposed to detect.
Long-haul freight	0.42	0.58	0.75	**0.043**	VISIBLE, MODERATELY RESISTED, the transition will be recorded as it proceeds. Timing depends on institutional outcomes rather than on detection capability.
Autonomous ride-hail (this paper)	0.50	0.33	0.94	**0.021**	OBSERVABLE CASE, visible and institutionally unprotected. Displacement can be measured as it happens, which makes this a suitable venue for building measurement practice and statutory technique that later transfers to less legible sectors. This is the profile of the case in this paper.
Diagnostic radiology	0.58	0.83	0.94	**0.006**	SLOW AND VISIBLE, resistance and administrative visibility both high. Policy attention is likely to be adequate and possibly excessive relative to the displacement actually occurring. The risk here is misallocated attention, not missed detection.

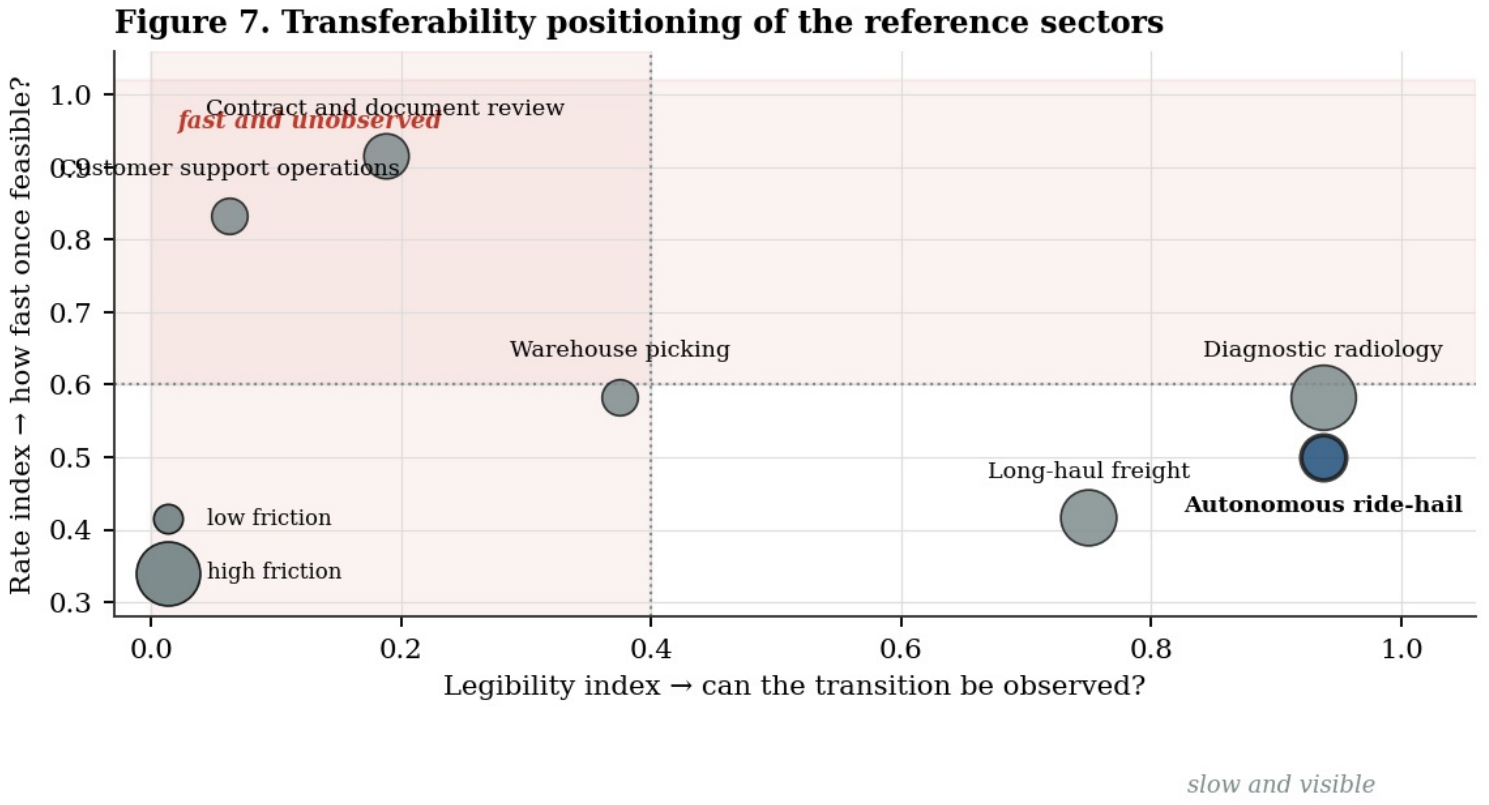

Figure 7. Reference sectors positioned on legibility and rate, with marker size proportional to institutional friction. The shaded region is the fast-and-unobserved quadrant. The case examined in this paper sits at the opposite corner.

Three features of that ordering are worth drawing out. The sector studied here ranks *lowest but one* on surveillance priority, which is the point: it was tractable to study precisely because it is observable, and its tractability is not evidence that other sectors will be similarly tractable. Diagnostic radiology inverts the ride-hail profile almost exactly, light capital, heavy licensure, and complete administrative visibility through billing records, and the risk there is misallocated attention rather than missed detection. And the two highest-priority sectors share a characteristic that ride-hail lacks entirely: no incumbent bearing obligations the entrant escapes, and therefore no actor with both the information and the motive to raise an objection. In this case the alarm was raised by a taxi company. In an unbounded software deployment there may be no equivalent party.

D.4 What transfers regardless of profile

Five results from this paper are not contingent on the scoring above. They should hold in any sector, and a researcher can treat them as starting assumptions rather than hypotheses to re-establish.

Invariant	Claim	Scope	Consequence for analysis
The oversight identity Section 4.2 of this paper	Where a human retains accountability for automated work, oversight cost per unit and oversight employment per unit are the same quantity, not two variables.	Sector-independent. Holds wherever residual human review is required.	No strategy delivers both a low automated unit cost and a large oversight workforce. Any proposal promising both is internally inconsistent before it is examined empirically.
Definitional exit Findings 1 and 6 of this paper	An obligation attaching to a person who performs a task ceases to attach when the task is automated, without any actor deciding that it should.	Highly general. Applies to revenue, licensing, labour, disclosure and safety obligations wherever the operative category names a human role.	Failure is silent. It produces no dispute, no proceeding and no gap that anyone is obliged to report. Detection depends on an incumbent noticing an asymmetry, which is an accidental mechanism.
Measurement lag Sections 3.2 and 9 of this paper	Administrative categories are built around the prior production function and stop describing the activity before anyone notices they have.	Sector-independent.	The window to establish measurement is before displacement, not during. Obligations are far easier to impose on an entrant than on an established industry.
Headcount is not exposure Section 3.4 of this paper	Where work is fragmented across many partial participants, headcount overstates aggregate labour displaced and understates the number of	Any sector with contractor, gig, or multi-employer participation.	Programmes sized to participant counts are mis-scaled in a direction that cannot be determined without activity-distribution data.

	households affected.		
Functional non-substitutability Section 6 of this paper	A share of apparent demand is not addressable by the automated substitute at any price, for reasons of physical function rather than preference.	Any sector where the incumbent artefact or worker performs bundled functions.	Full-conversion scenarios are bounds, not scenarios. The non-substitutable share should be estimated from administrative records rather than assumed.

D.5 Using the instrument

1. **Score the sector on all ten dimensions** before computing anything. Partial scoring invites the researcher to fill gaps in the direction of their prior. The companion model raises an error on missing or out-of-range dimensions rather than defaulting them.
2. **Read the legibility family first.** If it scores low, the analytical priority is establishing measurement, and every other question is downstream of that. Waiting for evidence of displacement in a low-legibility sector means waiting for exactly the thing measurement was supposed to detect.
3. **Identify the lowest-scoring legibility dimension** and treat closing it as the first-order intervention. In this paper's case that dimension was transaction observability, the trips exist and are priced, but no public body records them, which is why measurement leads Section 9 and every other recommendation is sequenced behind it.
4. **Do not aggregate across sectors.** The indices are ordinal comparisons within a common frame, not cardinal quantities. A priority score of 0.65 is not ten times worse than 0.065; it means the sector belongs in a different attention class.
5. **Score sensitivity on the family with the widest internal spread.** A family whose dimensions disagree is doing less work than its mean suggests, and the disagreement usually points at a real structural feature worth examining directly.

> **Limits of the instrument.** The dimensions are the authors' selection, chosen because they explained variation in this case and in the reference sectors; they have not been validated against outcomes, because the outcomes have not yet occurred. The scoring is ordinal and judgemental. The reference scorings other than ride-hail are unvalidated. Anyone using this should treat it as a structured way of organising an argument about a sector, not as a measurement of it, and should report their scores so that others can disagree with them specifically rather than generally.

Appendix E. Three horizons

Extending this analysis over time invites a specific error, and it is worth naming before committing it. Attaching dates to the conditional technology states would convert a conditional structure into a forecast, falsifiable in a way that discredits the structure when the dates prove wrong, and they will. The horizons below are therefore kept separate from the technology states, and a reader may reject any horizon without disturbing anything else in this paper.

The more consequential design point is that three years, five years and ten years are not the same question asked three times. They are governed by different variables, admit different methods, and warrant different kinds of answer.

Horizon	Governed by	Method	Epistemic status	Appropriate output
H1 ~3 years Commitment horizon	Decisions already taken: capital in the ground, permits filed, vehicles ordered, contracts signed, service areas announced.	Inventory what is already committed and let it play out. Very little of what happens in this window is still undecided.	**LARGELY KNOWABLE. Falsifiable within the horizon. Errors here are failures of research, not of forecasting.**	Point estimates with stated ranges.
H2 ~5 years Threshold horizon	Whether unit economics cross behavioural thresholds, most importantly the ownership-parity frontier of Section 6.	Threshold analysis. Identify the crossing conditions and report what follows on each side, rather than predicting which side obtains.	**PARTLY MODELLABLE. The consequences of crossing are more reliable than the timing of the crossing.**	Conditional statements: if the threshold is crossed, then X.
H3 ~10 years Institutional horizon	How institutions respond: whether measurement exists, whether statutory categories were repaired, whether professions closed ranks, whether transition support was targeted.	Scenario branching. Two or more institutional paths, developed in parallel and presented as alternatives.	**NOT FORECASTABLE. The dominant variables are political choices that have not been made. A point estimate at this horizon is a category error, however carefully it is hedged.**	Branches, never a central case. Report what each branch requires.

> **A point estimate at the ten-year horizon is a category error.** The dominant variables there are political choices that have not been made: whether measurement exists, whether statutory categories were repaired, whether professions closed ranks. Those are not uncertain in the way a cost curve is uncertain. They are undetermined, and partly determined by whether analysis of this kind is done and acted on. The honest output at that horizon is branches.

E.1 What moves when nobody intervenes

The dimensions of Appendix D are not fixed. Each has an expected direction of travel in an untreated sector, and they do not all move the same way.

Dimension	Untreated drift	Reasoning
`capital_gating`	falls	Hardware cost falls; shared infrastructure is built once.
`oversight_requirement`	falls	Reliability improves; supervision ratios widen.
`labour_share`	ambiguous	Roughly stable until the displaced work is gone.
`occupational_closure`	rises	Professions respond to threat by tightening scope of practice. Friction is partly ENDOGENOUS to rate.
`collective_representation`	ambiguous	Ambiguous: threat drives organising, fragmentation defeats it.
`liability_concentration`	ambiguous	Ambiguous: courts may concentrate liability, legislatures may create safe harbours.
`jurisdictional_boundedness`	falls	Functions relocate and dematerialise; remote operation crosses borders.
`incumbent_asymmetry`	falls	INCUMBENTS EXIT. The party that raised the alarm stops existing.
`transaction_observability`	falls	Work moves in-house; priced external transactions become internal costs.
`statutory_anchoring`	ambiguous	Stays high until deliberately repaired; repair is the intervention.

Two of these deserve emphasis. **Occupational closure rises**, because professions respond to automation by tightening scope of practice, which means friction is partly endogenous to rate, and a static instrument that treats the two as independent will understate resistance in credentialed sectors and only there. And **incumbent asymmetry falls**, because incumbents exit. The party positioned to notice that an obligation had lapsed stops existing.

Figure 8. The observation window closes because observers leave, not because the transition accelerates

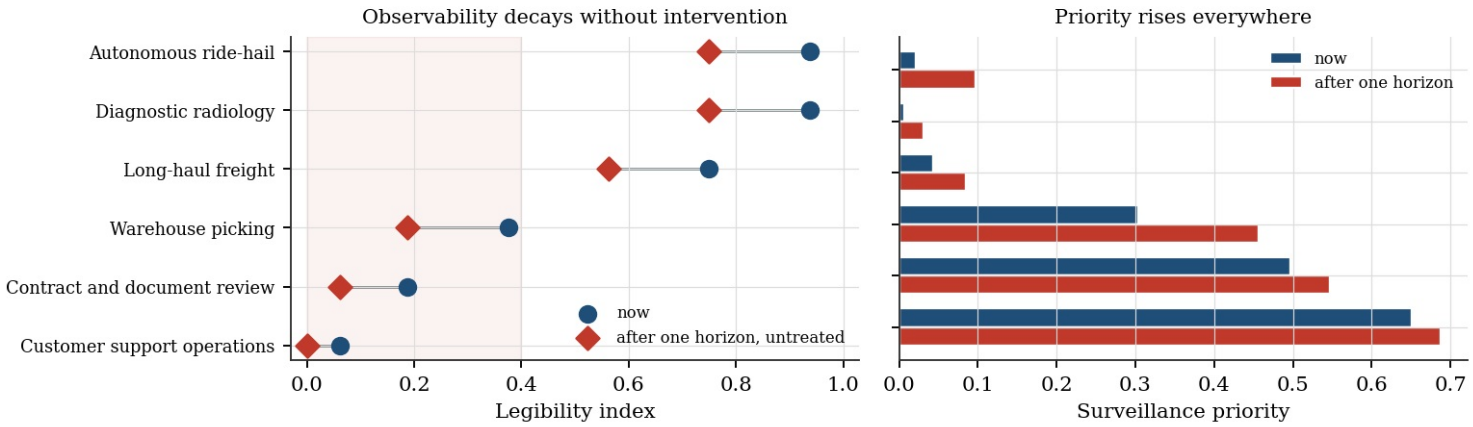

Figure 8. One horizon-step of untreated drift applied to the reference sectors. Legibility falls in every case and surveillance priority rises in every case.

Sector	Legibility now	after	change	Priority now	after	change
Autonomous ride-hail	0.94	0.75	-0.19	0.021	0.097	+0.076
Diagnostic radiology	0.94	0.75	-0.19	0.006	0.031	+0.025
Long-haul freight	0.75	0.56	-0.19	0.043	0.085	+0.042
Warehouse picking	0.38	0.19	-0.19	0.304	0.457	+0.153
Contract and document review	0.19	0.06	-0.12	0.497	0.547	+0.050
Customer support operations	0.06	0.00	-0.06	0.651	0.688	+0.036

Finding 7. Observability decays without intervention, and it decays fastest where it is currently highest. The sector examined in this paper loses more legibility per horizon-step than any of the low-legibility comparators, because those have little left to lose. The mechanism has nothing to do with the technology: incumbents exit, functions relocate, and priced external transactions become internal costs. *The window for observation closes because the observers leave, not because the transition accelerates.*

This is the strongest available argument for the sequencing in Section 9. Measurement established now is cheap and effective; the same measurement attempted two horizons later is expensive and partial, and by then the party who would have complained has gone out of business.

E.2 The three horizons, told through four households

What follows is narrative rather than model output, and it inherits the epistemic status of its horizon: the first is largely knowable, the second conditional, the third explicitly branched. The households are the composites introduced in Section 1.

Horizon 1, the commitment horizon, roughly three years

Marcus. Public service exists now, across a handful of neighbourhoods and the airport run. He notices it on Friday nights downtown: the surge that used to hold until two in the morning softens by midnight. His average hour is much as it was. His *best* hours are worse, and best hours were the reason the work was viable. He is still driving. Nothing about this appears in any dataset.

The Delaneys. Nothing changes. They use the service twice all year, once when both cars were committed, once after a dinner downtown. It is a convenience, not a substitute. The second car sits in the driveway exactly as before.

Ray. Nothing changes. Nothing will.

Eleanor. Service does not reach Lakewood. She reads about it, and has formed an opinion about it.

What is actually established at this horizon: the marginal hour of one occupation got worse, and no record of it exists.

Horizon 2, the threshold horizon, roughly five years

Marcus. This is the decision point, and which way it goes turns on a quantity nobody measured. If lower prices expanded the market, there is more work at worse rates. If they did not, there is less work. He is deciding whether to replace his car, a five-year commitment against a market he cannot read, and neither can the state.

The Delaneys. The second car needs replacing, and for the first time they do the arithmetic properly. It is close. What decides it is not the model's threshold but whether the transmission fails before or after the lease on the first car ends. The ownership frontier is crossed household by household, on replacement timing, not on price.

Ray. Still nothing, though his van is now serviced at a depot that also turns around fleet vehicles, and the technician who works on it trained on a fleet contract. The transition reaches him as a supplier, not as a competitor.

Eleanor. Service reaches her. She starts using it for the medical appointments she had begun to dread driving to. She keeps the car.

What is conditional here: everything. Each household's outcome depends on a threshold crossing whose timing this paper cannot fix.

Horizon 3, branch A, Colorado measured and acted

Marcus. Transition support was sized on published activity data, so he was identified as a full-time driver rather than as one of twenty-four thousand licences. Bridge income, then depot operations, one of the physically anchored roles that stayed in the region.

The Delaneys. They shed the second car. Because an empty-mile charge held fleet deadhead near the twenty percent threshold, the miles the fleet added were fewer than the miles their car removed. Regional vehicle travel fell.

Ray. Unaffected, and the registration and emissions join meant his truck was correctly excluded from every state projection for a decade.

Eleanor. She gave up driving at eighty-four and has mobility she would not otherwise have had. The gain was documented, so when service contracted elsewhere, hers was defended on the record.

Horizon 3, branch B, nothing was measured

Marcus. Displaced. Support was sized to licence counts, so it was either far too generous or far too thin. Nobody can say which, because the distribution was never published.

The Delaneys. They shed the second car anyway. Deadhead ran near thirty-five percent, so the fleet added more miles than their shedding removed. Regional vehicle travel rose, and no one attributed it to anything.

Ray. Still fine. But every state projection for ten years counted his truck among the conversion candidates, so every travel and revenue forecast in that period was wrong in the same direction, and nobody could say by how much.

Eleanor. The same gain, still uncounted. Because it had never entered a cost-benefit analysis, the route to Lakewood was cut in a later contraction on the grounds that it was not carrying its cost.

The two branches differ in almost no technical particular. The vehicles are the same, the costs are the same, and the households make substantially the same decisions. What differs is whether anyone knew, and therefore whether the empty-mile externality was priced, whether transition support reached the right people, whether a decade of forecasts were biased in a knowable direction, and whether an uncounted benefit survived the first budget contraction. The ten-year outcome is not a prediction to be refined. It is a choice being made now, by default, in the absence of measurement.

---

27 July 2026. Scenario analysis, not a forecast, not legal advice. All quantitative results are conditional on Appendix A. The companion model and audit harness permit full replication and re-parameterisation; the audit fails if any stated conclusion diverges from the computed result.