


HOW CUSTOMER EFFORT CHANGES RESOLUTION RATES.

What U.S. auto recall data reveals about whether a fix actually reaches the customer

Most CX programs measure whether a problem was identified or a case was closed. Far fewer measure whether the resolution actually reached the customer. We analyzed U.S. automotive recall data because it provides something most industries do not: a publicly reported outcome showing whether the customer actually received the fix.

EVERY YEAR MILLIONS OF AMERICAN CARS GET RECALLED, AND **THE DATA IS PUBLIC.**

A carmaker finds a fault. By law it must repair every affected vehicle free of charge, and it must report to the regulator how many owners actually got the repair. Almost no other industry is required to publish whether its fixes reached anyone at all.



Share of recalled vehicles eventually repaired, by the year the recall was issued, across roughly 4,000 campaigns, measured over the full life of a campaign. The series stops at 2024 on purpose: repair rates keep climbing for years after a recall is issued, so 2025 campaigns have not settled yet and plotting them would show how recent they are rather than how well they were repaired.

The rate has sat between 77% and 87% for a decade and none of it is contested. NHTSA's own published average across all recalls is 64%, rising to 83% for new vehicles. Our own figure is lower at 58.6%, because we cut every campaign at four quarters so that recent recalls are not compared against mature ones. Carfax counts 58.1 million US vehicles carrying an open recall today, roughly one in five on the road.

So we took every US light-vehicle recall old enough to measure and split it by a dimension not reported by NHTSA: how much the owner had to do to get the repair. NHTSA told GAO in 2024 that it holds no completion statistics at all for recalls fixed by a software update sent to the car.

THE COST OF ASKING THE CUSTOMER TO SHOW UP.

Every recall has one of two remedies. Some are software faults, fixed by an update sent to the car overnight, which asks the owner for nothing. The rest need the vehicle physically brought in.

	NOTHING ASKED OF THE OWNER software sent to the car	OWNER MUST ACT book and attend a service visit
ALL MANUFACTURERS	95.1% 59 campaigns · 17.1M vehicles	56.8% 3,156 campaigns · 345.4M vehicles
TESLA ALONE	98.4% 21 campaigns · 11.1M vehicles	68.7% 41 campaigns · 1.3M vehicles
EVERY OTHER MANUFACTURER	88.9% 38 campaigns · 6.0M vehicles	56.7% 3,115 campaigns · 344.1M vehicles

Share of covered vehicles remedied at four quarters, weighted by vehicles covered. Tesla is 65% of over-the-air volume, so the pooled row and both halves are shown rather than one number.

SOFTWARE REMEDY 0 actions required of the owner The car updates itself, usually overnight. Most owners never know it happened.	PHYSICAL REMEDY 4 actions required of the owner 1 Read the recall notice 2 Find a participating dealer 3 Book an appointment 4 Bring the car in
--	---

These two groups were never equivalent. Software remedies reach newer, connected vehicles and address a different class of fault, and which recalls can be fixed that way is decided by the defect rather than by chance. So this is a comparison of completion rates between unlike groups: an association, not a measured causal effect.

Every confound we could identify widens the gap rather than narrowing it, so the share attributable to effort alone is smaller than 38 points. We cannot say how much smaller and do not claim to. What survives is that the gap holds with the manufacturer held constant: inside Tesla alone it is 29.7 points.

Nothing here shows that owners are at fault. Every one of these vehicles had a working repair waiting for it, free of charge. The four steps above are the manufacturer's design, not the customer's failing.

TAKE OUT TESLA AND IT IS STILL 89% AGAINST 57%.

The previous page put over-the-air repairs at 95.1%. Tesla accounts for 65% of that volume, so the obvious objection is that the whole result is one unusual company. Take Tesla out completely and the remaining over-the-air repairs still reach 88.9%, against 56.8% for repairs that need a service visit. These are the ten manufacturers that make up that remainder.

MANUFACTURER	CAMPAIGNS	VEHICLES COVERED	SHARE OF THIS GROUP
Mercedes-Benz USA	7	1,949,206	32.5%
General Motors	10	1,938,021	32.3%
Stellantis (FCA US)	2	1,353,043	22.6%
Volvo Car USA	3	458,017	7.6%
Ford Motor Company	7	135,712	2.3%
Nissan North America	3	75,101	1.3%
Hyundai Motor America	3	58,126	1.0%
Kia America	1	14,163	0.2%
Rivian Automotive	1	7,873	0.1%
American Honda	1	4,174	0.1%
Ten manufacturers	38	5,993,436	100%

THOSE TEN COMBINED 88.9% 38 campaigns · 6.0M vehicles repaired over the air	TESLA ALONE 98.4% 21 campaigns · 11.1M vehicles repaired over the air	EVERY PHYSICAL REPAIR 56.8% 3,156 campaigns · 345.4M vehicles needing a service visit
--	--	--


Over-the-air campaigns at four quarters, excluding Tesla. We publish no completion rate for each manufacturer individually because most hold only one to three qualifying campaigns, which is too few to rate. The rates below are the pooled figures.

What survives every objection we could put to it is that the gap holds with the manufacturer held constant: inside Tesla alone it is 29.7 points, and across these ten it is 32.2. The reasons this group could still be flattering the result are set out in full on the method page.

Source: NHTSA recall API over-the-air field, every in-scope campaign, retrieved August 2026

THE PATH THAT WORKS COVERS ONE VEHICLE IN TWENTY.

Repairs that ask nothing of the owner reach 95.1%. They are also almost never the option available. Nearly every recalled vehicle is on the path that needs the owner to act, and that is where the repairs stop landing.



This is the shape of most customer experience problems. The mechanism that works is known, proven, measured, and deployed on a rounding error of the actual volume. The constraint is not evidence. It is that the remaining 95% of repairs still require the owner to do the work, and most manufacturers depend on an independent business to persuade them to do it.


A FIX THAT NEVER ARRIVES IS INDISTINGUISHABLE FROM A FIX YOU NEVER MADE.

A DIFFERENT QUESTION: WHOSE COMPLAINTS PILE INTO ONE COMPONENT?

Effort tells us whether the fix lands. Complaint mix tells us where to look next.

Vertical: how far a brand's worst component sits above the rate for all 25 brands.


Horizontal: whether complaints are rising or falling against that brand's own prior year. Every brand has a worst component; the question is how far above the industry it sits. Eight brands sit in the top right quadrant. Five of them are also rising more than 20%: Kia, Jeep, Mazda, Hyundai and Lincoln.



Source: NHTSA ODI complaint database, twelve months to 31 July 2026, deduplicated. Complaint counts are not normalised by fleet size, so this shows where a brand's own complaints concentrate, not how a brand compares on defect rate. Rivian omitted: no component cleared the 30-complaint floor.

HOW WE WORKED THIS OUT, AND WHERE IT COULD BE WRONG.

Every figure in this report is computed from public sources. Nothing here is licensed, modelled or estimated.



COMPLETION	NHTSA quarterly recall completion reports. Measured at a fixed campaign age of four quarters so recent campaigns are not compared against mature ones. Weighted by vehicles covered, not by campaign count. Campaigns issued October 2013 to June 2025, reported through 2026 Q2.
REMEDY PATH	NHTSA recall API over-the-air field, retrieved for every in-scope campaign. Complete coverage, no missing values.
COMPONENTS	NHTSA COMPDESC as assigned, with sibling categories merged on the root before the first comma or colon. Hydraulic brakes fold into service brakes. The unspecified catch-all is excluded.
SCOPE	Light vehicles only. Upfitters, van converters and wheelchair-accessible vehicle modifiers excluded. Manufacturer names normalised before aggregation.

EIGHT REASONS THIS COULD BE OVERSTATED.

Every one of these pushes in the direction of making the gap look bigger than it is. They are listed so a reader can discount the result themselves.

VEHICLE AGE

80% vs 56%

completion for cars under three years old against cars six to ten years old

Cars able to update themselves are newer, and newer cars complete more often whatever the path. Some of this gap is age, not effort.

REACHABILITY

Not measured

no public data on address quality by vehicle age

A connected car has a live account and a current address. A fifteen-year-old car on its third owner may have neither, and the manufacturer simply knowing where the customer is will carry some of the result.

PARTS DELAY

Not measured

interim notices are not flagged in the public file

Physical recalls are sometimes announced before the parts exist, which suppresses completion at four quarters for reasons that have nothing to do with the customer. Software remedies have no parts constraint.

FAULT TYPE

By definition

the two paths cannot carry the same repair

Software-remediable defects are a different class of fault from a part replacement. This compares two paths carrying different repairs, not two routes to an identical one.

NO DENOMINATOR

0 of 25

brands have a public US fleet-size figure

The brand complaint chart therefore has no exposure denominator. Read it as where a brand's own complaints concentrate, never as a defect rate.

SMALL SAMPLE

38 campaigns

and three makers hold 87% of the vehicles in them

Mercedes-Benz, General Motors and Stellantis dominate the group, so removing Tesla does not leave ten evenly weighted manufacturers behind it.

NOT RANDOM

11 makers

issued a software remedy at all in this period

They are the firms that had connected fleets to attempt one, which means newer fleets, which is the same bias as vehicle age above.

HOW WE COUNT

0.59

rank correlation between a four and eight quarter window

Manufacturer rankings are not stable, so none is published. Vehicle counts are also campaign-vehicle pairs: one car in three campaigns counts three times.

THREE QUESTIONS TO ASK ABOUT **YOUR OWN BUSINESS.**

A regulator counts this for carmakers. For everyone else it has to be measured deliberately. These are the questions this analysis is the automotive version of.

01

WHAT SHARE OF YOUR RESOLUTIONS ACTUALLY REACH THE CUSTOMER?

Not cases closed. Not responses sent. The share where the customer received the outcome and it took effect. Almost every CX programme measures whether a problem was found; very few measure whether the fix landed. Automotive has a legal mandate, a funded repair and a tracked outcome, and still reaches 56.8% when the customer has to act.

02

HOW DOES THAT RATE DIFFER BY DELIVERY ROUTE?

Tesla's own recalls run 98.4% when nothing is asked of the owner and 68.7% when a service visit is. Same manufacturer, same four-quarter measurement window. Because the company is held constant, that is the comparison worth running inside your own operation, on your own delivery routes.

03

HOW MANY ACTIONS DOES THE CUSTOMER HAVE TO TAKE?

Count them. Four required actions versus none coincided with a gap of nearly forty points. Step count is one variable you can control.

We measure this for organisations that do not have a federal regulator counting on their behalf. If you want to know what your own resolution rate looks like by delivery route, get in touch at intellimark.net.

THE DIFFERENCE WASN'T WHETHER A FIX EXISTED. IT WAS HOW THE CUSTOMER HAD TO GET IT.

Some owners had to do nothing. Some had to do four things. Four required actions versus none coincided with a gap of nearly forty points. The next question is what it looks like in your business, and public data cannot see that.