Operational Validity Under Change

A Validity-Control Perspective on Robotaxi Scaling

INDEX

Introduction

Robotaxi deployment is increasingly moving from technical demonstration toward sustained operation and geographic expansion. As autonomous driving systems accumulate testing, operational experience and safety evidence, the challenge of scaling can appear primarily technical: vehicles can be reproduced, software can be deployed across fleets, and established systems can be introduced into additional operating areas.

Yet scaling raises a less visible problem. Evidence supporting operation is established under conditions. Those conditions can concern the system itself, its operating environment, relationships with external capabilities, and the institutional setting under which autonomous operation is permitted. When deployment changes, the technology may remain substantially the same while some of the conditions supporting its existing justification do not.

This creates a question that is not adequately captured by asking whether a robotaxi is technically capable of operating in a new environment. Nor can it be resolved by assuming that change itself determines the need for renewed validation. What matters is whether a change affects something on which the existing justification depends, whether the resulting state remains sufficiently covered by existing evidence, and what operation can still be justified under that state.

This paper examines that problem as one of **operational validity under change**.

The analysis first identifies the basis on which operational validity depends and distinguishes change from material change. It then examines how changes in validity can affect permissible operation and derives the resulting relation between validity and control. That relation is tested across materially different sources and mechanisms of change before its implications for robotaxi scaling are examined.

The purpose is not to propose a new autonomous-driving architecture or a universal theory of autonomous-system safety. The narrower objective is to determine whether materially different deployment changes can be understood through a common structure linking existing justification, changed conditions and permissible operation.

The paper therefore addresses the following research question:

Under what conditions can an existing claim of justified robotaxi operation remain valid when the conditions on which that claim depends change?

---

1. The Problem

A robotaxi can be proven to work without being proven to work everywhere.

This is not merely a geographic limitation. It follows from the way operational evidence is established. Testing, simulation, operational experience and safety assessment demonstrate performance under identifiable conditions; they do not establish an unlimited claim for every possible deployment environment. The resulting evidence is therefore bounded by the conditions under which the relevant capability has been demonstrated.

This matters for scaling. Once an autonomous driving system has accumulated extensive testing, operational experience and safety evidence, expansion may appear primarily to be a problem of replication. More vehicles can be produced, software can be deployed across a growing fleet, and a proven vehicle platform can be introduced into additional operating areas.

But the evidence supporting justified operation does not necessarily travel with the vehicle.

A robotaxi operates as part of an operational system. Its demonstrated capability is related to the conditions under which that capability was established: characteristics of the vehicle and ADS, the operating environment, relevant external relationships and, where applicable, the institutional conditions under which deployment is permitted.

Geographic expansion makes the transfer problem particularly visible. A move from one city to another can introduce different roads, traffic patterns, infrastructure, weather conditions, operating restrictions or authorization requirements. Yet a new geography does not itself establish that the existing justification has ceased to apply. Two geographically different areas can share relevant operating conditions, while two adjacent areas can differ in a condition on which an existing validity claim depends.

Nor does meaningful change require geographic movement.

Software can be updated while the vehicle remains on the same roads. A new vehicle platform can be introduced within an existing operating area. Supporting infrastructure or connectivity can change. An external dependency can become unavailable. Authorization conditions or operational-support arrangements can change while the physical deployment remains otherwise substantially unchanged. Conversely, a deployment can enter a new geographic area while many of the conditions relevant to its existing justification remain unchanged.

This prevents two premature conclusions.

A new location does not automatically require an entirely new justification merely because it is new. And a change does not automatically invalidate an existing justification merely because something has changed. Both conclusions would treat the occurrence of change as if it already determined the significance of that change.

The deeper issue is that operational validity is conditional rather than universal.

This changes the scaling problem.

The relevant question is no longer only whether the technology can be reproduced or whether the vehicle can technically operate in another city, on another vehicle platform or under another configuration. It is whether the justification supporting its operation remains applicable when the conditions surrounding that operation change.

A vehicle may be portable. Its software may be portable. Its components may be portable.

The validity established for their operation is not automatically portable with them.

This paper therefore examines the following question:

Under what conditions can an existing claim of justified robotaxi operation remain valid when the conditions on which that claim depends change?

Answering that question requires first identifying what such a validity claim actually depends on.

---

2. The Basis of Operational Validity

If operational validity is conditional, the next question is straightforward: what does it depend on?

For a robotaxi, the answer cannot be reduced to the vehicle or autonomous driving software alone. A robotaxi operates as part of a larger operational setting. Its ability to perform the dynamic driving task is essential, but justified operation also depends on the conditions under which that capability has been demonstrated, the relationships on which the particular system configuration relies, and the conditions under which its operation has been accepted.

This distinction matters because safety evidence always has a context. Testing, simulation, operational experience and safety assessment do not establish an unlimited claim that a system will perform appropriately under every conceivable condition. They support a claim within an identifiable range of system configurations, operating conditions, assumptions and, where relevant, external dependencies. For the purpose of this analysis, these dependencies can be distinguished across four dimensions rather than treating validity as a property residing exclusively in the vehicle.

Technical validity

The first dimension concerns the system itself.

Software configuration, sensing, perception, planning and control, vehicle characteristics, actuation, fallback capabilities and the configuration in which these elements operate can all contribute to the evidence supporting autonomous operation.

These elements may appear highly portable because software can be reproduced and a vehicle can physically be moved from one deployment environment to another. But physical sameness does not by itself establish validity equivalence. Software changes can affect the applicability of existing verification and validation evidence, and an unchanged component can encounter conditions under which the relevance of its previous evidence differs.

Technical validity is therefore necessary, but it does not describe the complete operational situation.

Operating-condition validity

The second dimension concerns the conditions in which the system operates.

Road characteristics, traffic conditions, weather, speed ranges, geographic restrictions and other elements of the operational design domain determine the conditions for which an ADS feature is designed and under which its performance has been demonstrated.

The same technical system can therefore encounter operating conditions for which its existing evidence has different relevance.

This does not mean that every environmental difference invalidates previous evidence. A changed condition may still lie within the range already represented by testing, validation or operational experience. The important point at this stage is narrower: operating conditions form part of what the evidence must be understood against.

Operational validity therefore cannot be inferred from technical capability alone.

Relational validity

Some dependencies cannot be assigned exclusively to either the robotaxi or its environment.

A particular autonomous-driving configuration may rely on communication, maps, road or localization infrastructure, backend systems, remote assistance, operational support or other external capabilities. This forms a distinct dependency class because neither side of such a relationship necessarily determines its validity on its own. A vehicle can remain functional and an external system can remain functional while the relationship required between them no longer satisfies the conditions on which operation depends.

The existence of external infrastructure or support, however, does not automatically make it part of the validity basis.

External support becomes validity-relevant where the particular autonomous-driving configuration relies on the function or relationship it provides. Regulatory treatment likewise recognizes that external infrastructure can form part of an ADS configuration without constituting a universal requirement for every configuration.

This produces an important distinction:

A component may be portable while the relationship that made its operation valid is not.

The relevant question is therefore sometimes not whether an individual component still works, but whether a relationship on which justified operation depends still holds.

Institutional validity

Robotaxi operation also takes place within an authorization framework.

Permits, operating conditions, accepted safety evidence, assurance requirements and jurisdiction-specific authorization conditions can determine whether a technically capable system is permitted to operate in a particular deployment setting.

This creates a fourth dependency that cannot simply be collapsed into technical validity:

Technical validity is not the same as authorization to operate.

A system can therefore remain technically unchanged while the institutional conditions governing its permissible deployment differ.

At the same time, this dimension must remain bounded. Not every legal or regulatory requirement affecting a robotaxi operator belongs to the operational-validity problem examined here. Institutional conditions become relevant where they directly affect the basis under which autonomous operation itself is justified or permitted. This distinction prevents regulation from being treated either as identical to technical validity or as the universal explanation for robotaxi deployment constraints.

The validity basis

Taken together, these four dimensions explain why operational validity cannot be treated simply as an inherent property of a robotaxi.

It is better understood as a claim supported by an identifiable validity basis: the relevant system properties, operating conditions, relationships and, where applicable, institutional conditions on which justified operation depends.

The four dimensions are analytically distinct, but a particular validity claim can depend on more than one of them at the same time. Operational validity can therefore involve system properties, operating conditions, external relationships and institutional conditions in combination.

This does not require every conceivable dependency to be catalogued in advance. Nor does the existence of a technical component, infrastructure element, service or regulatory condition automatically make it validity-relevant. What is required is that the dependencies supporting the particular claim of justified operation can be identified.

A different software version, vehicle, road environment, interface, external service or authorization condition therefore has no determinate validity consequence merely because it is different. **Its significance depends on its relationship to the validity basis it may affect.**

---

3. When Does Change Matter?

Robotaxi deployment is not static. Software changes, vehicle platforms evolve, operating areas expand, infrastructure relationships change, and regulatory or operational conditions can develop over time.

But the occurrence of change alone says little about its significance.

A large technical modification may leave the relevant basis of an existing validity claim substantially unaffected. A seemingly minor change may alter a condition, relation or assumption on which that claim directly depends.

Change is not the same as material change.

California autonomous-vehicle rules provide an external regulatory comparison because they distinguish safety-relevant changes across heterogeneous categories such as roadway type, operating speed, geographic area and operating restrictions. What these changes have in common is not their magnitude or novelty, but their potential relevance to the conditions under which operation has been authorized. **[California autonomous-vehicle regulations]**

Materiality therefore cannot be inferred simply from how large, visible or novel a change appears.

It is relational.

A change matters to an existing validity claim insofar as it affects the basis on which that claim depends.

Dependency

The first question is therefore:

Does the changed element affect something on which the existing validity claim depends?

If not, the change has not been shown to be material to that claim.

Consider a software update. The existence of a new software version establishes that the system has changed. It does not establish that a validity-relevant capability or assumption has changed. A modification affecting perception, planning, fallback behavior or another safety-relevant function may require a different assessment from a modification unrelated to the basis supporting autonomous operation. This distinction also appears in deployment practice: Waymo explicitly identifies safety-relevant software updates, rather than software change as such, as one of the circumstances in which deployment readiness is reassessed. **[Waymo, "Safe to Deploy," 2025]**

The same principle applies to geography.

Two geographically different areas can share many validity-relevant operating conditions, while adjacent areas can differ in road characteristics, traffic conditions, environmental constraints or authorization requirements that matter directly to the existing validity basis. The relevant test is therefore not simply whether the place is new, but whether the new deployment introduces a validity-relevant condition not sufficiently represented by the existing basis.

The category of change is therefore not enough.

Its dependency relationship matters.

That dependency can also become validity-relevant without disappearing.

A communication path, information source, external service or other supporting capability may remain technically available while no longer providing a sufficiently trustworthy basis for the role assigned to it. Technical availability must therefore be distinguished from sufficient trustworthiness: a dependency can remain present while its output can no longer be relied upon for autonomous operation.

A manipulation case makes the distinction particularly clear:

failure
→ a required dependency may become unavailable or degraded;

manipulation or rejection
→ the dependency may remain technically available but can no longer be accepted as a sufficiently trustworthy operational basis.

Both can therefore become material to the same validity basis without requiring a separate materiality logic.

At the same time, the state of one dependency must not automatically be confused with the state of the function it supports.

Loss, degradation or rejection of a supporting dependency does not by itself establish that the required function has disappeared. The function may remain sufficiently fulfilled through another valid path. A required function is therefore not necessarily identical to any single component or dependency supporting it.

But the existence of another path is not enough to infer resilience.

Where continued support is attributed to an alternative path, independence must be assessed relative to the dependency actually under challenge. Two paths may appear structurally different while sharing the same communication layer, information source, service, infrastructure or other underlying condition.

The relevant questions are therefore:

Does the alternative path avoid the challenged dependency?

and, if it does:

Does it remain sufficiently available, trustworthy and capable for the resulting operational state?

Structural duplication alone does not establish independent resilience.

Coverage

Dependency, however, is only the first part of the assessment.

A validity-relevant condition can change while the resulting state remains sufficiently supported by existing evidence.

This creates a second distinction:

Being affected is not the same as being insufficiently covered.

Three states can therefore be distinguished:

Unaffected.
The change does not materially affect a condition, relation or assumption on which the existing validity claim depends.

Affected but covered.
A relevant part of the validity basis changes, but the resulting state remains within demonstrated or accepted coverage.

Affected beyond existing coverage.
The change reaches a validity-relevant condition for which the existing evidence is insufficient to carry the previous validity inference forward without additional justification.

The third case requires careful interpretation.

Insufficient coverage does not establish that the robotaxi has been demonstrated to be unsafe.

It establishes something narrower: the existing evidence is no longer sufficient, by itself, to support the same validity claim for the resulting state.

This distinction is important because it separates lack of demonstrated support from demonstrated failure.

A new state may therefore require further evidence, adaptation, restriction or another justified response without the analysis asserting that unsafe behavior has already been proven.

A practical materiality test

The reasoning can now be reduced to three questions:

Step	Question
Dependency	Does the change affect something the existing validity claim depends on?
Coverage	Is the resulting state still sufficiently supported by existing evidence and accepted conditions?
Consequence	Can the existing validity claim still be carried forward without additional evidence, adaptation or restriction?

This sequence is deliberately narrower than the shortcut:

change → revalidation

because change alone does not determine the result.

A software update, a geographic expansion or a changed external relation can each produce the same three outcomes: unaffected, affected but covered, or affected beyond coverage.

The common question is therefore:

What changed relative to the basis of the existing validity claim, and is the resulting state still sufficiently covered?

Geography is a source of change, not the measure

This distinction is particularly important for robotaxi scaling because geographic expansion is highly visible.

A new operating area can introduce different road types, traffic patterns, environmental conditions, infrastructure relationships, speed ranges or authorization requirements. Any of these may affect the existing validity basis.

But geographic novelty itself does not determine materiality.

Waymo's deployment-readiness methodology provides a useful operational reality check. Expansion into a new geographic area is treated as a reason to reassess deployment readiness because the new ODD may expose the system to different risks and uncertainty—not because geographic novelty is itself equivalent to invalidity.
[Waymo, “Safe to Deploy,” 2025]

4. From Validity to Permissible Operation

Determining that an existing validity claim can no longer be carried forward unchanged does not, by itself, determine what a robotaxi should do.

A state that is no longer sufficiently supported for normal operation is not automatically a state in which the entire system has been shown to be unsafe. Nor does every such state necessarily require the same response.

The operational question is more specific:

Given the current validity state, which operations remain justified?

Existing ADS regulation provides concrete examples of this distinction through operational responses such as fallback and minimal-risk manoeuvres when continued normal operation is no longer appropriate. **[Commission Implementing Regulation (EU) 2022/1426]**

Different validity states can justify different responses

Under conditions sufficiently covered by the existing validity basis, normal operation may remain justified.

When those conditions change, a different response may become appropriate. Depending on the system and the condition encountered, that response can involve adapting behavior, restricting operation, entering a fallback or recovery state, moving toward a minimal-risk condition, or withholding normal operation until sufficient conditions have been restored.

The important point is not any one of these responses.

It is the relation between the **current validity state** and the **set of operations that remain justified**.

Operational validity should therefore not be interpreted as a simple binary distinction between “operate” and “stop.”

Validity situation	Operational consequence
Existing conditions remain sufficiently covered	Normal operation may continue
Changed conditions remain covered through an established response	Adapted or restricted operation may be justified
Existing support is insufficient for normal operation	Normal operation cannot simply be carried forward
Conditions required for normal operation are sufficiently restored	Normal operation may become justified again

The distinction matters because behavior appropriate to one validity state may no longer be appropriate to another.

Functional continuity does not require identical operation

A change affecting a supporting dependency does not by itself determine the permissible response.

Loss, degradation or rejection of one dependency does not necessarily imply loss of the function it supported. A required function may remain sufficiently fulfilled through another valid path. But functional continuity alone does not establish that the previous operation remains permissible. The resulting configuration must still support justified operation under the intended conditions.

This produces an important distinction:

Preserving a required function is not the same as preserving the previous scope of operation.

An alternative path may preserve the function while providing less capability than the original configuration. The resulting operational state may therefore support only a narrower speed range, smaller operating area, reduced duration, restricted functionality or another limited operating envelope.

Likewise, where the originally required function can no longer be sufficiently fulfilled, another justified response may still remain available.

A fallback or minimal-risk response should not automatically be described as replacing the function that was lost. It may instead perform a different required function: handling uncertainty or failure under the resulting condition.

Changing behavior can preserve validity

A change in system behavior does not necessarily indicate failure of operational validity. In some circumstances, changing behavior is precisely what preserves justified operation.

Consider a state in which continued normal driving is no longer sufficiently supported, but a previously validated fallback or recovery response is.

The system may transition away from normal operation toward a restricted or recovery-oriented state. If that response is itself supported for the condition encountered, the behavioral change does not mean that operational validity has disappeared. It means that a different response is justified by the resulting validity state.

This gives rise to an important distinction:

Loss of normal operation does not necessarily mean loss of operational validity.

Operational validity can therefore require preservation of **appropriate behavior**, not preservation of **identical behavior**.

For robotaxi operation, this distinction is fundamental. A system that continued behaving identically after a material change would not necessarily demonstrate robustness. It could instead be failing to respond to information that should alter its permissible operation.

Constraint is not substitution

Continued operation may become possible only by narrowing the conditions under which the system operates.

An operational restriction can reduce the demands placed on the remaining capabilities. A lower speed, smaller operating area, restricted environmental condition or another limitation may establish a narrower state in which those capabilities are sufficient for the operation being considered.

But the restriction does not create a capability that is absent.

Constraint can restore alignment between available capability and permissible operation without substituting for the missing capability itself.

Otherwise, a constrained deployment could be described as functionally equivalent to the configuration from which it was derived.

It is not necessarily equivalent.

What may have been preserved is **justified operation within a different boundary**.

Minimal-risk behavior makes the relation visible

Minimal-risk behavior provides an observable operational example.

EU ADS regulation provides for defined responses under conditions in which continued normal operation cannot be maintained as intended. A minimal-risk manoeuvre can move the vehicle toward a minimal-risk condition rather than simply continuing normal operation, while departure from the resulting minimal-risk condition is conditional on confirmation that the cause of the manoeuvre is no longer present. **[Commission Implementing Regulation (EU) 2022/1426]**

The resulting operational pattern can be represented as:

normal operation
→ condition no longer sufficiently supports that operation
→ **fallback / minimal-risk / recovery response**
→ triggering condition assessed
→ **possible return to normal operation after sufficient restoration**

This should not be interpreted as a universal state machine. The form and ordering of responses depend on the system and on what has actually been justified for the condition encountered.

Different states can support different permissible behavior.

The system is therefore not simply switching between “valid” and “invalid.”

Its operation changes according to what the current validity state can support.

Permission is broader than prohibition

Permissible operation is not merely a prohibition mechanism deciding whether the robotaxi may move.

Depending on the validity state, justified responses can include:

continue · adapt · restrict · fallback · recover · resume

The operational problem is one of selecting or constraining behavior according to what remains sufficiently justified.

This is more precise than a simple stop/go model. Once different validity states justify different actions, validity participates directly in the regulation of system behavior.

The resulting boundary is therefore one of **selective permission**, not simply prohibition.

A functional boundary appears

On one side are operations sufficiently justified under the current validity state.

On the other are operations for which that justification is absent or insufficient.

This boundary is not necessarily permanent. Conditions can change, evidence can be extended, restrictions can be introduced, dependencies can be restored, and normal operation can become justified again.

What remains constant is the underlying relation:

The current validity state determines which operations can still be treated as justified.

For that boundary to remain operationally meaningful, the authority responsible for enforcing permissible operation must remain intact.

Information, interpretation or external support can contribute to operation without possessing authority to determine or execute vehicle action. If an unauthorized component or actor can override the relevant decision or execution path, the permissible-operation boundary can no longer reliably constrain behavior according to the validity state.

The integrity of operational authority is therefore a condition for the effective enforcement of the permissible-operation boundary.

---

5. The Validity-Control Relation

Operational validity depends on an identifiable basis. Change matters when it affects that basis beyond what existing evidence sufficiently covers. A resulting change in validity state can then alter which operations remain justified.

At that point, the problem is no longer only one of validation.

It becomes one of control.

From validity to control

Control, in this context, does not mean that an external authority continuously directs the robotaxi. Nor does it refer only to stopping the vehicle when something goes wrong.

It describes a more general function:

Control uses information about the current state to regulate system behavior so that operation remains within justified conditions.

This function follows from the preceding analysis rather than being assumed in advance.

If normal operation remains sufficiently supported, it may continue.

If changed conditions remain covered through an established response, behavior may adapt or operation may be restricted.

If normal operation is no longer sufficiently justified, it cannot simply be carried forward. A fallback, recovery or minimal-risk response may instead remain justified where that response is supported for the condition encountered.

If the conditions required for normal operation are sufficiently restored, normal operation may again become permissible.

The significance of these differences is not simply that the system can behave in different ways. It is that behavior justified in one validity state may no longer be justified in another.

Control therefore does not determine whether change itself is acceptable.

It regulates which actions remain justified after the consequences of change have been assessed.

The Validity-Control Relation

The relationship can now be stated:

Operational validity is conditional on a validity basis. Assessing whether change is material requires that the relevant basis be identifiable. Change becomes material when it affects that basis beyond the coverage of the existing validity claim. Where the resulting validity state changes which operations remain justified, permissible operation must be regulated accordingly. This regulation performs a control function.

Why the intermediate steps matter

A shorter statement would be tempting:

Change requires control.

But that is not what the analysis has established.

Some changes do not affect the existing validity basis. Others affect it while remaining sufficiently covered. Still others affect the basis beyond existing coverage without thereby demonstrating that the system has become unsafe.

The control relation becomes operationally significant only where the resulting validity state has consequences for justified operation.

The same problem arises with:

A new deployment requires revalidation.

This again makes the occurrence of change the criterion rather than its relationship to the existing validity basis.

The Validity-Control Relation therefore preserves three distinctions:

- change is not automatically material;
- material change does not automatically imply demonstrated failure;
- a change in justified operation does not automatically require complete cessation of operation.

These distinctions allow continuation, adaptation, restriction, fallback and restoration to be considered without reducing operational validity to a binary safe/unsafe decision.

A boundary defined by justified operation

The control function regulates the boundary between operations that remain sufficiently justified under the current validity state and operations that do not.

That boundary can shift as conditions change. Normal, restricted, fallback or recovery operation may be justified under different states, and normal operation may become justified again when its required conditions have been sufficiently restored.

The boundary is therefore dynamic without being arbitrary. Its reference remains the validity basis and the evidence supporting the resulting state.

The relation can be expressed as a continuous chain:

change → validity basis → materiality → coverage → validity state → permissible operation

The control function arises where the resulting validity state constrains the operations that can still be treated as justified.

Scope of the relation

The Validity-Control Relation has been derived here from the operational problem of robotaxi deployment.

It should not, on that basis alone, be treated as a universal architecture for autonomous systems. Such a claim is neither necessary nor established here.

Within the robotaxi operational context examined here, conditional validity, material change, evidence coverage and permissible operation form a coherent relation that performs a control function.

---

6. Testing the Relation Across Change

The Validity-Control Relation was derived from a particular problem: an existing basis for justified robotaxi operation encounters change.

A relation derived through one path, however, may reflect characteristics specific to that path.

The relation can therefore be tested by varying the **location and mechanism of change** while holding its underlying logic constant. Six materially different situations are considered, moving the point of change from the system itself to the operating context, external relationships, institutional conditions and operational support.

The question remains the same throughout:

Does the change require a different validity-control logic, or can it still be assessed through the relation already established?

Six different locations of change

Case	What changes?	Central question	Result
Software update	System	Can validity be affected without changing the deployment environment?	Relation holds
Geographic / ODD expansion	Operating context	Is a new location inherently material, or only the conditions it changes?	Relation holds
Infrastructure dependency	System–environment relation	Can validity depend on a relationship rather than an individual component?	Relation holds
Vehicle platform	System + interfaces	Does the relation remain sufficient when several dependencies change together?	Relation holds
Authorization	Institutional conditions	Can permissible deployment change while technical capability remains unchanged?	Relation holds, with scope distinction
Operational support	Operational relation	Can normal operation cease while justified fallback operation remains possible?	Relation holds

The cases are deliberately different. Their purpose is to determine whether materially different kinds of change require materially different validity-control structures.

System change and context change

Software updates and geographic expansion provide the simplest contrast.

A robotaxi can remain in the same operating area while its software changes. If the update affects a capability, assumption or dependency on which existing validity depends, the change may require assessment even though nothing has changed geographically.

The opposite can occur as well. The software and vehicle can remain unchanged while the robotaxi enters a new operating area. The new geography matters only to the extent that it introduces conditions relevant to the existing validity basis.

Neither software change nor geographic change is therefore the fundamental variable.

In both situations, the same questions remain sufficient:

Does the change affect the validity basis? If so, does the resulting state remain sufficiently covered?

When the dependency lies between things

Infrastructure introduces a different problem.

Some conditions required for justified operation cannot be assigned exclusively to the robotaxi or exclusively to its environment. They exist in the relationship between them.

A vehicle may function correctly. An external system may also function correctly. Yet the communication, compatibility, information exchange or another relationship required between them may no longer satisfy the conditions on which operation depends.

The relevance of such external support is architecture-relative. Infrastructure does not become part of the validity basis merely because it exists or could support autonomous operation. It becomes validity-relevant where the particular autonomous-driving configuration relies on the function or relationship it provides.

This produces an important result:

Component portability does not imply relational portability.

Moving or reproducing individual components does not necessarily reproduce the relationship that supported their previous operation.

Loss of one infrastructure dependency does not necessarily imply loss of the function it supported if another sufficiently independent and capable path remains available. Apparent redundancy is insufficient where the alternative path shares the dependency under challenge.

When several dependencies change together

A vehicle-platform change tests the relation differently because several elements may change together.

Sensor configuration, vehicle dynamics, interfaces, actuation or fallback capabilities can be affected simultaneously.

The number of changed elements does not require a different validity logic. What matters is whether their combined effect reaches the validity basis beyond existing coverage.

The relation therefore remains applicable to coupled as well as isolated change.

Technical validity and institutional permission

Authorization provides a deliberately different test.

A robotaxi can remain technically unchanged while the conditions governing its permissible deployment change.

This means:

Technical validity and institutional authorization are not the same thing.

A system may remain technically capable of performing as before without automatically retaining the same permission to operate under changed authorization conditions.

But the distinction also sets a limit.

Not every regulatory change is a change in operational validity. A legal or administrative requirement matters to the relation examined here only where it affects the basis under which autonomous operation itself is justified or permitted.

This case also demonstrates why not every validity-relevant change can be treated as a problem of technical substitution.

A changed authorization condition cannot simply be compensated for by adding another sensor, communication path or technical capability.

Validity without identical behavior

Operational support provides a different test because loss of a dependency required for normal operation need not eliminate every justified response.

If another valid path preserves the required function, operation may continue within the state that path supports.

If the function cannot be preserved, an already justified fallback or recovery response may still remain available.

Loss of normal operation does not necessarily mean loss of operational validity.

Failure and manipulation

The six cases distinguish **where** a validity-relevant change occurs. A separate variable is **how** the relevant condition becomes challenged.

This distinction becomes important when a dependency remains available but can no longer be relied upon in the same way.

A communication path, information source, external service or support capability may continue to function technically while becoming insufficiently trustworthy for the role assigned to it. Manipulation can therefore affect the validity basis without producing an ordinary availability failure.

Consider two different disturbances to the same dependency.

In the first, the dependency fails and its required output becomes unavailable.

In the second, the dependency continues to provide output, but there is insufficient reason to trust that output for the operational role assigned to it.

The observable technical states differ:

unavailable

and

available but insufficiently trustworthy

The validity questions remain structurally the same:

Does the affected dependency belong to the validity basis?

Does the resulting state remain sufficiently covered?

Which operations remain justified under that state?

A robotaxi therefore cannot infer operational sufficiency from availability alone. A dependency can remain present while the conditions under which its output can legitimately support operation have changed.

Nor does an alternative path settle the issue automatically. If both the primary and nominally independent path rely on the same compromised information source, communication layer or support service, the apparent redundancy does not remove the challenged dependency.

Failure and manipulation therefore test the relation differently without requiring a different underlying validity logic.

What remains constant across the cases

The location and mechanism of change vary, but the assessment remains:

What changed?
→ Does it affect the existing validity basis?
→ Does the resulting state remain sufficiently covered?
→ Which operations remain justified?

Across the cases examined, no separate validity-control logic is required for system change, geographic change, relational change, coupled change, authorization change or operational-support change.

The manipulation test extends this result: the relation also remains applicable where a dependency is technically available but no longer sufficiently trustworthy for its assigned role.

What this test does not establish

This result does not establish that every possible change affecting every autonomous system can be reduced to the cases examined here.

Nor does it establish that technical, operational, relational and institutional conditions are interchangeable. Their evidence requirements and practical consequences can differ substantially.

The existence of alternative paths also does not establish independent resilience. Independence alone does not establish that the resulting capability is sufficient for the previous operating envelope.

The result is narrower:

Within the robotaxi contexts examined here, differences in where change occurs and in how a validity-relevant condition becomes challenged did not require a different underlying validity-control relation.

---

7. Implications for Robotaxi Scaling and Operational Safety

Robotaxi scaling is often visible as physical expansion: more vehicles, larger operating areas, additional cities or new vehicle generations. Each expansion brings concrete engineering, regulatory and operational work.

But physical expansion is only part of the problem.

The deeper question is whether the basis supporting justified operation remains applicable as deployment changes.

Scaling is not primarily a geography problem

Geographic expansion can change many conditions at once: roads, traffic behavior, infrastructure, environmental conditions and authorization requirements.

But a new geography is not itself the relevant measure of validity. Conditions encountered in a new operating area may already be sufficiently covered, while software, vehicle, infrastructure or operational changes within an existing area may materially affect the basis supporting operation.

The relevant scaling question is therefore not simply:

Can the robotaxi operate in another city?

It is:

Which functions are required for the intended operation, how are they currently provided, and which parts of the existing justification remain applicable when those conditions change?

This moves the problem from geographic replication to **validity transition**.

Scaling as a sequence of validity transitions

A deployed robotaxi does not remain static while it scales.

Software changes. Operating areas expand. Vehicle platforms evolve. Infrastructure relationships change. Authorization conditions differ. Operational-support arrangements develop.

These changes remain technically and institutionally distinct. A software update is not the same problem as a new authorization regime, and a vehicle-platform change is not the same problem as entering a new operating area.

Yet each can create the same underlying challenge:

An established justification for operation encounters a changed state.

Scaling can therefore be viewed as a sequence of validity transitions.

The practical task is not to eliminate those transitions. Change is inherent to deployment and development. The task is to determine what existing evidence can legitimately continue to support as those transitions occur and, where required functions are affected, whether those functions can remain sufficiently available under a changed configuration.

Successful operation at scale cannot be inferred from replication alone.

The ability to reproduce a vehicle or software stack does not necessarily reproduce the complete basis on which its operation was justified.

From infrastructure gaps to functional requirements

This becomes particularly important when the conditions of one deployment cannot simply be reproduced in another.

The first question should not be whether the target environment contains the same infrastructure as the existing deployment.

It should be:

What function does the autonomous-driving configuration require, and where is that function currently provided?

A function may be provided onboard the vehicle, through infrastructure, through operational support, or through a combination of these.

The relevant scaling problem is therefore not infrastructure equivalence by itself. It is **functional sufficiency**.

An environment with different infrastructure may still provide the functions required by a particular autonomous-driving configuration. Conversely, an environment with extensive and highly developed infrastructure may still lack a function or relationship on which another configuration depends.

Infrastructure readiness is meaningful only relative to the functions and dependencies of the autonomous-driving configuration being scaled.

This distinction becomes particularly visible when comparing highly developed deployment environments.

Germany and Japan illustrate the problem. Both possess highly developed transport and communications infrastructure, yet that general characterization says little by itself about their suitability for a particular robotaxi configuration. Differences in the conditions under which infrastructure, communication and operational support are provided can affect autonomous-driving configurations differently.

The comparison does not support a ranking in which one country is simply more "robotaxi-ready" than the other.

It demonstrates something narrower:

General infrastructure quality is not a sufficient proxy for functional suitability.

The relevant question is whether the particular functions and relationships required by the autonomous-driving configuration are sufficiently supported in the target deployment environment.

Functional allocation creates a scaling option

Where a required function is not sufficiently supported in the target configuration, deployment is not necessarily impossible.

The function may be capable of being provided differently.

A capability previously supported through infrastructure may be provided onboard. A function performed onboard may be supplemented through external support. Other valid combinations may also be possible.

The relevant question is not whether the original implementation can be copied unchanged, but whether the required function can be preserved through a valid alternative allocation.

This creates an important scaling option:

A deployment condition does not necessarily have to be reproduced in its original form if the function it provided can be fulfilled sufficiently through another valid configuration.

Dependency relocation

Changing where a function is provided changes what the system depends on.

Moving a function from infrastructure to the vehicle may reduce one external dependency while increasing requirements for onboard sensing, computation, energy, software or validation.

Adding external support may reduce demands on an onboard capability while introducing dependencies on communication, interfaces, information sources, service availability or operational actors.

Functional redistribution therefore does not establish equivalence.

An alternative allocation has to be treated as a different dependency configuration. The fact that the same high-level function remains available does not establish that the conditions supporting its operation remain unchanged.

Where scaling relies on an alternative path, that path must be independent relative to the dependency under challenge and sufficiently capable for the resulting operating state. Structural duplication alone does not establish resilience, and independence alone does not establish functional sufficiency.

Functional continuity and the operating envelope

If an alternative allocation preserves the required function sufficiently, operation may remain justified.

But functional continuity does not automatically preserve the previous operating envelope.

An alternative configuration may provide the function with reduced range, performance, precision, duration or availability. In that case, the function remains present while the operations it can support become narrower.

A target environment therefore does not necessarily present a binary choice between full replication and no deployment.

A configuration may remain valid for a more limited form of operation.

Constraint as a second scaling response

Not every functional gap has to be closed by adding or reallocating capability.

Another possibility is to restrict operating conditions until the available capabilities become sufficient for the intended operation.

A narrower operating area, speed range, environmental condition or another operational constraint can reduce the demands placed on the system.

But:

Constraint is not substitution.

Restricting the operating envelope does not create the capability that is absent. It changes the conditions under which the available capability is required to operate.

Two different forms of adaptation therefore become visible:

Functional adaptation changes how a required function is provided.

Operational adaptation changes the conditions under which the available functions are expected to support operation.

Both can contribute to a justified scaling configuration, but they do so differently.

A functional-validity scaling sequence

The resulting scaling logic can be expressed as an assessment sequence:

Step	Question
1 — Function	What function is required for the intended operation?
2 — Allocation and dependency	Where is that function provided, and what does that allocation depend on?
3 — Target condition	Which relevant dependencies remain available, trustworthy and valid in the target configuration?
4 — Alternative configuration	If necessary, can the function be sufficiently preserved through another allocation or independent path?
5 — Operational constraint	If full functional equivalence cannot be preserved, can a narrower operating state align available capability with operational demand?
6 — Resulting validity	What operation can the resulting configuration sufficiently support, and which operations remain justified?

This sequence does not prescribe an autonomous-driving architecture.

It identifies the questions that need to be resolved before a changed deployment configuration can be treated as functionally equivalent, sufficiently adapted or validly constrained.

Trustworthiness and authority remain design conditions

Functional reallocation must preserve more than technical availability.

Dependencies introduced through a new allocation must remain sufficiently trustworthy for their assigned role, and functional contribution must not be confused with operational authority.

A configuration cannot be treated as functionally sufficient merely because the required output remains technically obtainable.

Availability, functional sufficiency and trustworthiness are related but distinct conditions.

Operational safety as preservation of justified operation

Operational safety cannot mean preventing change. Robotaxi systems and their operating environments change over time.

Nor can it mean maintaining normal operation under all circumstances. Conditions can arise under which continuing normal operation is no longer justified.

Within the validity problem examined here, operational safety therefore includes:

the preservation of justified operation under changing conditions.

The relevant requirement is not identical behavior, but that the operation performed remains compatible with what the current validity state can support.

The underlying scaling constraint

Robotaxi scaling is not constrained only by whether autonomous-driving capability can be reproduced. It is also constrained by whether the functional conditions and justification supporting that capability can survive or be validly reconfigured as deployment changes.

The underlying scaling constraint can therefore be stated more precisely:

Within the validity problem examined here, the underlying scaling constraint is not change itself or the inability to reproduce an existing deployment environment exactly. It is the ability to preserve or validly reconfigure the functional conditions on which justified operation depends when change materially affects the existing validity basis.

What this does not solve

This approach does not solve every constraint on robotaxi scaling.

It does not determine perception performance, eliminate rare-event uncertainty, harmonize regulation, provide missing infrastructure, resolve liability, establish commercial viability or determine public acceptance.

Nor does it prescribe a universal sensor architecture, communication system, infrastructure model, operational-support structure or evidence package.

Functional reallocation cannot simply be presumed equivalent to the original configuration. It can create new dependencies, validation requirements and limits on permissible operation.

The contribution is narrower:

It provides a structured way to determine which functions an existing deployment relies on, how those functions are allocated, what dependencies that allocation

creates, and whether justified operation can be preserved through existing capability, alternative allocation or a validly constrained operating state.

---

8. Discussion and Conclusion

The analysis began with a practical problem: evidence supporting robotaxi operation under one set of conditions cannot automatically be assumed to support operation after those conditions change.

Following that problem produces a more specific result. Operational validity depends on an identifiable basis. Change becomes material when it affects that basis beyond what existing evidence sufficiently covers. Where the resulting validity state changes which operations remain justified, permissible operation must be regulated accordingly.

Testing this relation across different locations and mechanisms of change showed that substantially different changes can confront the same underlying problem: an established justification encounters a changed state. Their technical, operational and institutional consequences differ, but the validity question remains concerned with what changed, whether the change affects the existing validity basis, whether the resulting state remains sufficiently covered, and what operation can still be justified.

The scaling analysis adds a further implication:

Portability of technology is not equivalent to portability of operational validity. But non-portability of an existing validity basis does not necessarily imply non-portability of the required function.

A function on which justified operation depends may remain sufficiently available through the existing configuration, may be provided through another valid allocation, or may support operation only within a more constrained operating state.

This creates a connection between validity assessment and scaling design. The question is no longer only whether an existing justification survives change. It is also whether the functions underlying justified operation can remain sufficiently fulfilled under the resulting configuration without introducing dependencies that undermine the intended operational state.

Scope and limits

The Validity-Control Relation is not a complete theory of robotaxi safety.

It does not describe every mechanism required to design, validate or operate an autonomous-driving system. It does not determine which technical evidence is sufficient for a particular deployment, nor does it replace safety engineering, regulatory assessment or operational decision-making.

The cross-case test does not establish that every possible change can be represented by the situations examined here. Stability across those cases supports the relation within the investigated robotaxi context; it does not establish a universal law.

Several distinctions must therefore remain intact.

Technical validity is not the same as institutional authorization. A technically supported system may not have permission for equivalent operation under a different authorization regime.

Insufficient coverage is likewise not evidence that a system has been demonstrated to be unsafe. It means that the existing justification is insufficient, by itself, to support the same claim for the resulting state.

Operational safety also cannot be reduced to the Validity-Control Relation. The narrower result is that preserving justified operation under changing conditions forms one relevant part of operational safety.

Boundaries of functional reasoning

The functional analysis developed here is an analytical method, not a reference architecture.

Different autonomous-driving configurations can allocate required functions differently across the vehicle, infrastructure, external services or operational support. The purpose of functional analysis is therefore to separate **what function is required** from **how and where that function is provided**.

This distinction creates several boundaries.

Functional redistribution does not establish equivalence.

Changing the allocation of a function changes the dependencies through which that function is provided. An alternative allocation must therefore be assessed as a new dependency configuration rather than presumed to reproduce the validity of the previous one.

Functional redistribution is a design possibility, not a presumption of equivalence.

Infrastructure sufficiency is architecture-relative.

There is no universal infrastructure threshold established here for robotaxi deployment. Infrastructure becomes validity-relevant through the functions and relationships on which a particular autonomous-driving configuration depends.

The same environment may sufficiently support one configuration while leaving a validity-relevant functional gap for another.

General infrastructure development therefore does not determine functional suitability independently of the architecture being deployed.

Independence is dependency-relative.

Multiple paths, providers, modalities or services do not establish resilience merely because they are structurally distinct. Independence must be assessed relative to the dependency under challenge.

But independence itself is not sufficient. An independent alternative may provide less capability than the original path. The capability that remains must still be sufficient for the operation being considered.

Constraint does not create capability.

Reducing speed, geographic scope, environmental exposure or another element of the operating envelope can reduce the demands placed on available capabilities and may establish a narrower state in which those capabilities are sufficient.

But the restriction does not create the capability that was absent.

What is preserved in such a case is justified operation within a different boundary, not functional equivalence with the original configuration.

Functional sufficiency is not sufficient trustworthiness.

Technical availability and functional output do not exhaust the validity question. A dependency may remain operational while becoming insufficiently trustworthy for its assigned role.

Functional reallocation must therefore preserve not only capability but also the conditions under which that capability can legitimately support operation.

The same boundary applies to operational authority. External information, infrastructure or operational support can contribute to autonomous operation without thereby possessing authority to determine or execute vehicle behavior.

Functional contribution must not be confused with authority to determine or execute vehicle behavior.

A functionally sufficient configuration is therefore not necessarily a sufficiently trustworthy configuration.

Theoretical correspondence

The Validity-Control Relation was derived from the robotaxi problem rather than from a general theory of information, representation or autonomous systems.

That order matters.

The empirical and functional reasoning developed here stands independently of any broader theoretical correspondence.

A conceptual correspondence can nevertheless be observed.

A separate theoretical line of work distinguishes a sequence in which a distinction is formed, represented, interpreted and subsequently produces behavioral consequences:

Distinction → Representation → Interpretation → Behavioral Consequence

The examination of manipulation in the present analysis produces an independently derived observation that is structurally compatible with this sequence.

Suppose multiple downstream systems interpret the same upstream representation. The existence of several interpretive paths can appear to provide diversity. But if all of them inherit the same compromised upstream distinction, information source or representation, their diversity occurs only after the common dependency has already been affected.

Multiple interpretations therefore do not necessarily constitute independent evidence.

The relevant independence question has to be asked at the level of the dependency under challenge.

This correspondence exposes a broader structural point behind the manipulation analysis:

Diversity at one level cannot automatically compensate for a shared dependency at an earlier level.

The correspondence is not required for the Validity-Control Relation, nor does it establish that the two structures form a unified theory.

It is an analytical connection between independently developed lines of reasoning, not evidence for the validity of either one.

Conclusion

Robotaxi scaling is often described through visible expansion: more vehicles, larger operating areas, additional cities and new vehicle generations.

The analysis developed here suggests that the deeper problem lies elsewhere.

A robotaxi operates on the basis of conditions under which its operation has been justified. Those conditions can change even when the vehicle does not move, and they may fail to transfer even when the underlying technology does.

Technology may travel while its operational validity does not automatically travel with it.

But this does not mean that every changed deployment requires reproduction of the original environment.

The more precise task is to identify the functions on which justified operation depends, determine how those functions are currently provided, identify the dependencies created by that allocation, and assess whether those functions remain sufficiently fulfilled under the changed configuration.

Where they do not, another valid allocation may preserve the required function.

Where allocation changes, dependencies can change with it.

Where alternative paths are relied upon, their independence must be assessed relative to the dependency under challenge.

Where sufficient capability cannot be preserved for the previous operating envelope, permissible operation may need to become narrower.

And where the resulting validity state changes what remains justified, operation must be regulated accordingly.

The resulting conclusion is narrower than a universal theory of autonomous-system scaling, but more useful than treating each deployment transition as an unrelated problem:

Robotaxi scaling is better understood not as replication of a deployment environment, but as the preservation or justified reconfiguration of the functional conditions on which operational validity depends.

This does not remove the engineering, regulatory, institutional or operational complexity of scaling.

It provides a common structure for determining where that complexity becomes relevant to justified operation.

References

California Department of Motor Vehicles. (n.d.). *California Autonomous Vehicle Regulations*. California DMV. https://www.dmv.ca.gov/portal/vehicle-industry-services/autonomous-vehicles/california-autonomous-vehicle-regulations/

European Commission. (2022). *Commission Implementing Regulation (EU) 2022/1426 of 5 August 2022 laying down rules for the application of Regulation (EU) 2019/2144 as*

regards uniform procedures and technical specifications for the type-approval of the automated driving system (ADS) of fully automated vehicles. Official Journal of the European Union, L 221, 1–64.
https://eur-lex.europa.eu/eli/reg_impl/2022/1426/oj

Federal Ministry for Digital and Transport. (2024, December 4). *Die Zukunft fährt autonom: Strategie der Bundesregierung für autonomes Fahren im Straßenverkehr*.
https://www.bmv.de/SharedDocs/DE/Publikationen/StV/die-zukunft-faehrt-autonom.html

Ministry of Land, Infrastructure, Transport and Tourism of Japan. (n.d.). *Autonomous Driving Infrastructure Committee (自動運転インフラ検討会).*
https://www.mlit.go.jp/road/ir/ir-council/jido-infra/index.html

Waymo. (2025, June 16). *Safe to Deploy: How We Know the Waymo Driver Is Ready for the Road*.
https://waymo.com/blog/2025/06/safe-to-deploy/

Zoox. (2024, March 13). *The Next Steps for Our Robotaxi in Las Vegas, Foster City, and Beyond*.
https://zoox.com/journal/zoox-update-03-14

rightsofrobots.com. (2026, July 5). *Information, Meaning and Behavior.* Research area comprising *The Semantic Scaling Problem*, *Meaning Under Conditions of Scale*, *From Structure to Action*, and the extension *Distinction and Representation*.
https://www.rightsofrobots.com/research/information-meaning-behavior/