ad-research.org

APPROVED FOR PUBLIC RELEASE

TECHNICAL PAPER

THE STANDING WAVE UNDER YOUR CAR

A practical diagnostic for porpoising onset from ordinary CFD force histories

REPORT DATE

May 31, 2026

PREPARED BY

Andrew Michael Brilliant

Director of Computational Physics

DISTRIBUTION

Public Release

VERSION

3.0

EXECUTIVE SUMMARY

This paper presents the practical screening method we use at ADR to catch porpoising onset from standard CFD force-time histories. The underbody shows downforce peaks across ride height; whether a peak causes porpoising depends on its force amplitude against the suspension's spring rate. We use negative kurtosis of the force time series to flag the ride heights worth a closer look, then confirm with the force trace, the power spectrum, and the 3D flow field.

Applied to 324 CFD runs on an LMP-class vehicle, the method flagged risk at two ride heights before the car turned a wheel; the actual wheel rate exceeded the critical rate by about a factor of two, predicting no porpoising, and a second vehicle showed no concern. The body lays out the diagnostic, the setup procedure it implies, and why passive damping cannot remove the underlying aero bistability. This is simply how we do it at ADR, shared in case other race engineers and aerodynamicists find it useful.

ADR-RPT-CF91B8EB-197F-47FF-AFD9-E7186BD7D1C4

THE SHORT VERSION

A car carries a whole heap of vortices, and they interfere, constructively and destructively, to make the force at the wheels. When their frequencies line up the right way, the standing wave of the car and of the individual parts suddenly goes constructive, and you get a downforce spike, or the opposite, a drop-out. So you either stop the suspension right there, or you fix the aero. And the thing to understand is this: you can move that peak around, but you cannot make it go away.

The Spoon in the Sink

Run the water in your kitchen sink. Angle the tap so the stream runs down the side of the basin, making a sheet of water along the wall rather than a jet into the middle. Now hold a spoon vertically, curved side toward the wall, and ease it toward the flow.

The water accelerates through the narrowing gap between the spoon and the wall. Low pressure forms, the same way a wing works. The spoon pulls itself toward the wall. Tap. The gap closes, the flow chokes, the pressure vanishes, and the spoon springs back. The gap reopens, the flow restarts, the low pressure rebuilds. Tap. Tap. Tap. That is porpoising.

The spoon is the car’s floor. The sink wall is the road. The water is the airflow. As the car gets closer to the ground, the air accelerates under the floor, low pressure forms, downforce increases, and the car is pulled down further, like the spoon pulling itself toward the wall. At some critical ride height, the flow breaks down. The downforce collapses. The springs push the car back up. The gap reopens, the flow reattaches, and the cycle starts again.

A fair objection: if porpoising is a dynamic motion, why read it off a static force–ride-height map at all? Because at the frequencies a heaving car actually runs, the dynamic underbody response collapses back onto its quasi-static shape. The static map is not a convenient approximation here; it is the right tool for the job in this regime. (The heaving-airfoil work behind that statement is in the reading list.)

WHY THE STANDARD EXPLANATION IS INCOMPLETE

The usual shorthand is “ride height sensitivity”: the aero load changes rapidly with ride height near the ground. That is true, but it does not by itself explain why one car starts oscillating while another with similar downforce does not, or why a setup change can make the problem appear or disappear without touching the underbody.

This note offers a practical framing that appears to fill part of that gap. Porpoising onset is not just an aero problem and not just a suspension problem. It is an accessibility problem. The underbody may contain a bistable ride-height band, but the car only porpoises if the suspension lets the chassis move through the full sequence.

One Peak Is Enough

The downforce curve of a ground-effect car has a peak: a ride height where downforce is maximum. Above the peak, downforce is building as the car gets closer to the ground. Below it, the flow has broken down and downforce is falling. As ride height falls, the diffuser flow passes from steady and symmetric, to symmetric but unsteady, to unsteady and asymmetric, with a hysteresis loop around maximum downforce.

That peak is the sink wall. On one side the spoon is pulled in; on the other it springs back. A single peak creates three zones: a stable region above the peak where downforce is building (a node), the peak itself where the flow is most active and can support two configurations (the antinode), and a stable region below where downforce has dropped (another node). Node–antinode–node. That is the minimum structure for porpoising. The node–antinode language is an engineering map, not a claim that the underbody contains a literal acoustic standing wave; it is just a way to mark where the system can rest and where it can flip state.

THIS IS NOT A CFD ARTEFACT

That matters enough to be blunt about. In wind-tunnel tests the diffuser flow switches randomly between a high-downforce symmetric state and a low-downforce asymmetric state at the same ride height, with no change in model settings; the switch is the collapse of one of the two counter-rotating diffuser vortices, breaking the flow's symmetry. For the air, the two configurations are equally valid. It simply picks one.

And it is not a fixed-ground or Reynolds artefact: it shows up on a moving-ground belt with the floor boundary layer removed, it is near-independent of Reynolds number (which rules out boundary-layer merging), and when the same experiment is run on a fixed ground across a range of diffuser angles the behaviour is essentially unchanged, the curve drops in magnitude, the character does not. The same class of bistability turns up far outside diffusers: in the turbulent wake of a plain blunt body it is Reynolds-independent from Re ~ 10^2 to 10^6 and exists only above a critical ground clearance, bistability gated by ride height, in a completely different flow. When something this robust shows up in that many places, it is a feature of the physics, not the mesh.

HOW I HOLD IT IN MY HEAD

The cleanest way I have found to think about it is to forget the flow for a moment and watch only two forces:

the downforce pulling the car down, and the spring pushing it back up. The car always lives in downforce; it never makes lift. So the whole thing is a single inequality. While downforce exceeds the spring force, the car goes down. The moment the spring force exceeds the downforce, the car goes up. That is the entire dynamic.

The peak is nothing special in itself: it is just the ride height where downforce stops rising and starts to fall. The car passing through it is still going down, only now the downforce is dropping. Keep falling and the downforce drops below the spring force, the inequality flips, and the car is driven back up. It cannot stay; it is rejected from that ride height and returns higher. A node is a ride height where downforce holds above the spring force, so the car settles. An antinode is a ride height where downforce has fallen below it, so the spring sends the car away. Same curve seen from two sides, one spring always wanting to bring the car home. I am deliberately not saying what the flow is doing to make the downforce fall, whether a vortex bursts, a bubble opens, the diffuser stalls; the car neither knows nor cares. It feels only the two forces.

And it is never one source. The underbody is a system, diffuser pair, wing-tip vortices, edge and fence vortices, all coupled, sometimes reinforcing and sometimes cancelling. The standing wave, the way I use the term, is the interference pattern of that whole sum across ride height: nodes where the contributions cancel into a force that holds the car, antinodes where they line up so that the force backs off and the spring sends the car away. That is why two cars with “the same aero” porpoise differently, and why moving one device a few millimetres can shift the whole wave: you have re-timed one term in the sum. The full device-by-device decomposition is its own project; this is the working picture, and I am comfortable that it lines up with how the same kind of behaviour shows up wherever it has been measured carefully.

There is also a hysteresis loop in the force curve: the downforce at a given ride height depends on the direction of approach, and the high-downforce branch is unstable, so any disturbance can trip it down to the low-downforce branch. That much is measured, and quasi-static. The dynamic reading is mine: this path dependence is what sustains the cycle, with the car descending on one force branch and returning on the other.

Multiple peaks make it worse. A camel-back force curve creates several candidate antinodes, widening the ride-height range over which the car can be kicked. But one strong peak is enough.

THE CONDITION FOR OSCILLATION

The key question is whether the force change at the peak can move the car far enough to reach the adjacent node. If the spring is stiff enough, the chassis barely moves, the spoon is held too far from the wall; the water can pull, but not far enough:

$$\Delta F\_peak / k\_susp \geq \Delta h\_to\_node$$

If this is not met, the standing wave exists in the aerodynamics but the car is mechanically unable to traverse it. No complete traversal, no porpoising. The spoon twitches but never taps.

WHAT WE FOUND IN THE DATA

On the LMP-class vehicle, mapping the force histories with spectral analysis of standard CFD output separated the underbody flow into stable and bistable ride-height bands:

- Nodes (stable flow, single topology): FRH 15–16 mm, 19–20 mm, 24+ mm.
- Candidate antinodes (bistable signature): FRH 13–14 mm, 17–18 mm, 21 mm.

Three antinodes, separated by nodes. The primary antinode at FRH 17 mm had the strongest signature, with force oscillation equal to 4.1% of total vehicle load.

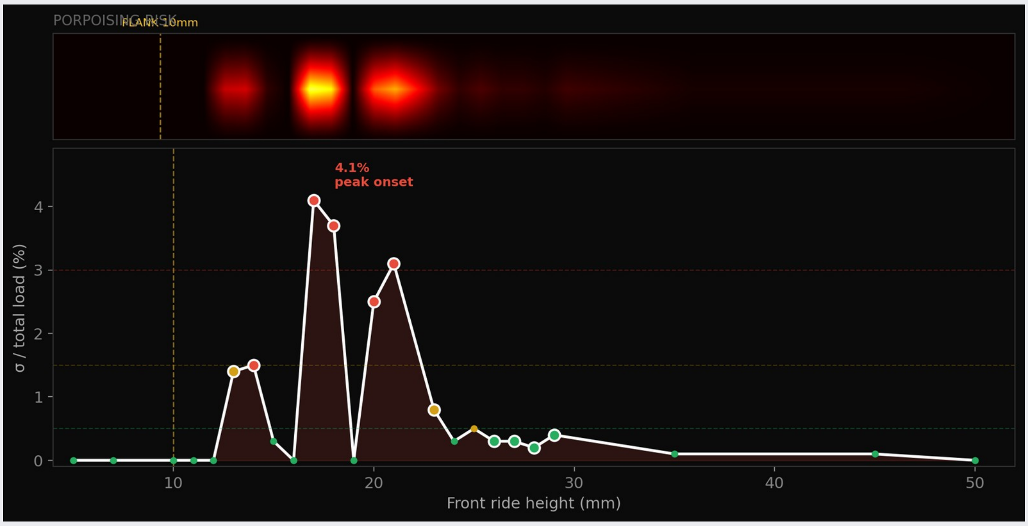

Figure 1. Porpoising risk map, LMP-based vehicle. Three candidate antinodes at FRH 13–14, 17–18, and 21 mm, separated by stable nodes. The map does not prove porpoising by itself; it identifies the ride-height bands where the suspension-accessibility calculation matters.

Why a Pinned Car Doesn’t Porpoise

Put your finger between the spoon and the sink wall. The spoon still feels the pull, but it cannot reach the wall. No tap. No cycle. That is the bump-stop idea in its cleanest form.

If the car could be magically pinned at FRH 17 mm with infinite stiffness, it would not porpoise. It might sit on peak downforce. The CFD force trace might still show an unsteady or bistable underbody. But if the chassis cannot move, the force fluctuation cannot become a ride-height oscillation. Porpoising requires range: the car has to traverse from one node, through the antinode, to the other node. No range, no complete oscillation.

The cycle looks like this:

- The car sits above the peak. Flow is attached enough to build downforce, but the chassis is not yet in the most sensitive part of the map.
- Speed increases, downforce builds, and the chassis is pulled down against the spring.

- The car reaches the antinode. Downforce is high but the flow is marginal. A snap in the underbody force pushes the chassis past the peak.
- Below the peak, the high-downforce state collapses. The springs now carry more stored energy than the reduced aero load needs, so they drive the chassis back up.
- The car passes back through the antinode and overshoots toward the upper node.
- The cycle repeats if the aero, spring, and ride-height range keep feeding it.

The descent is driven by increasing aero load against the spring; the upward leg by stored spring energy after the downforce collapses. The antinode is not where the car rests, it is where the car changes state.

How to Map the Standing Wave

You do not need special instrumentation to start. The first signal is already in ordinary CFD output: force versus iteration. Most workflows already monitor the standard deviation of the force history for convergence, that gives the amplitude of the unsteadiness, but little about its shape.

WHAT TO ADD: KURTOSIS

Kurtosis measures the shape of the force distribution rather than its amplitude. In Python it is one line: k = kurtosis(Fz_timeseries, fisher=True). The sign is a first-pass diagnostic:

- Near zero: the force distribution is roughly single-state. Usually safe, assuming small amplitude.
- Positive: rare spikes on an otherwise compact signal. Worth watching as possible precursor behaviour.
- Negative: a flattened or two-level distribution, consistent with the force spending time in more than one state. A two-level (bimodal) distribution is the recognised fingerprint of a system flickering between two attractors near a bistable threshold, and it is exactly how bistability is identified in turbulent bluff-body wakes. This is where porpoising risk can live.

Kurtosis is a cheap flag, not a proof. A negative value tells you which ride heights deserve a closer look in the force trace, PSD, and 3D flow field.

THE PROCEDURE

- Run the aero map sweep as normal (ride-height variations).
- For each run, compute the Fz standard deviation and kurtosis.
- Plot kurtosis against ride height. Negative regions become candidate antinodes; inspect the force trace and PSD there.
- Identify the adjacent stable nodes.
- Compute the mechanical accessibility: can the force impulse move the chassis from the antinode to the nearest node?
- Set the spring and bump-stop package so the car cannot traverse a complete node–antinode–node sequence.

Kurtosis flags candidate antinodes cheaply, but it is a heuristic, not a test. To put a number on it, Hartigan's dip test provides a formal test of non-unimodality with a significance level: run on a downsampled (decorrelated) force series, the dip test assumes independent samples, which a raw CFD time series is not, it turns a kurtosis-flagged ride height into a statistically defensible one, with a p-value rather than just an eyeballed shape.

THE SPECTRAL FINGERPRINT

The visual check is the power spectral density. Stack the PSDs by ride height. The antinodes show broadband excitation; the nodes show cleaner, narrower features. The transition is usually obvious once you know where to look.

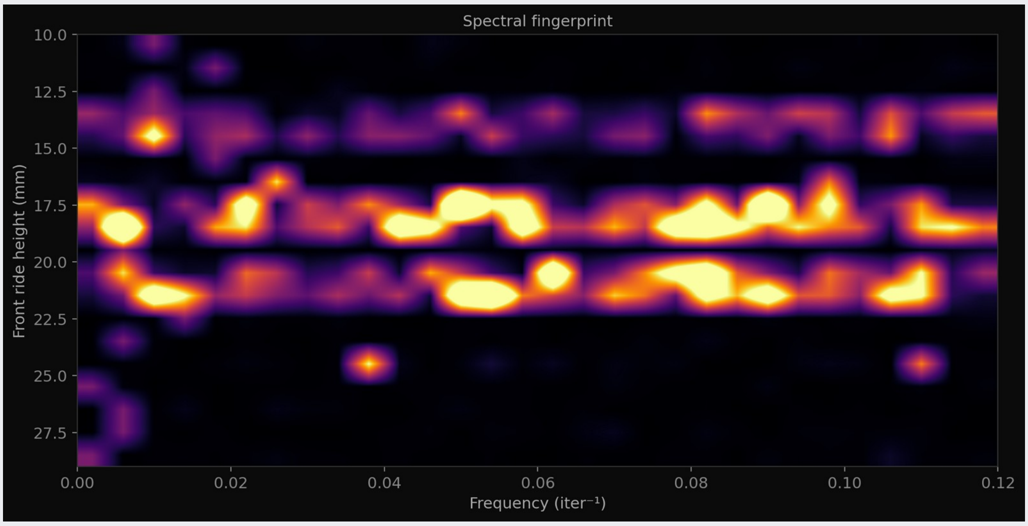

Figure 2. Spectral fingerprint of the same LMP vehicle. Each horizontal band is one CFD run. The candidate antinodes at FRH 13–14, 17–18, and 21 mm show broadband excitation; the nodes show cleaner narrow-band structure.

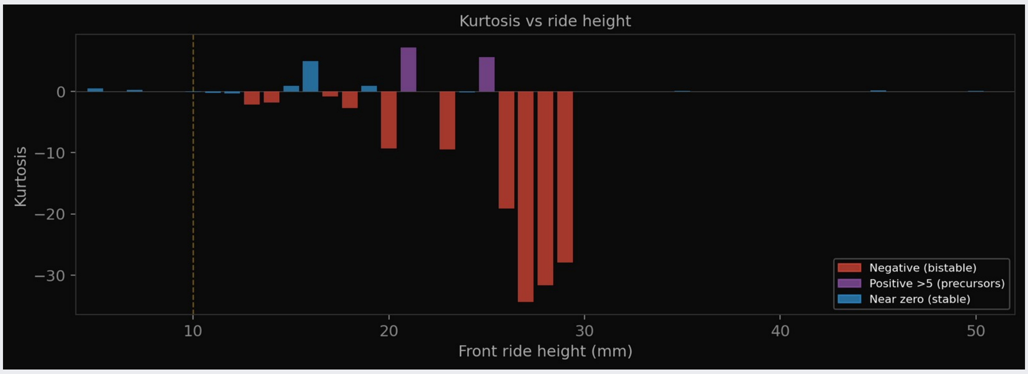

Figure 3. Kurtosis versus ride height. Negative values flag candidate bistable ride heights for closer inspection; positive values mark spike-like behaviour; near-zero is usually more stable. The plot's practical value is telling you which neighbouring CFD cases to compare in the 3D flow field.

The Suspension Fix

If porpoising requires a complete node–antinode–node sequence inside the accessible range, the first fix is straightforward: do not let the car access a complete mode.

CRITICAL SPRING RATE

A force impulse ΔF at the antinode produces a ride-height excursion:

$$\Delta h = \Delta F / k$$

The car can traverse the wave only if that excursion reaches the nearest node. Per corner:

$$k\_min = \Delta F\_per\text{-}corner / \Delta h\_to\text{-}nearest\text{-}node$$

For the LMP-class vehicle, the primary antinode produced a force oscillation of about 2% of total downforce on the relevant basis. With a 60% front distribution, the per-corner impulse is about 111 N. The nearest node is 1 mm below the antinode, giving k_min ≈ 111 N/mm. The actual front wheel rate at top speed was about 200 N/mm, so the same impulse moves the chassis only Δh ≈ 111/200 = 0.56 mm, short of the 1 mm needed to reach the node. The wave is present, but the suspension package prevents the car from traversing it.

THE PORPOISING NUMBER

For setup work, collapse the same idea into a dimensionless number. Use one basis throughout; this note uses per-corner quantities. With ε the relative force-oscillation amplitude, F_0 the per-corner peak downforce at the antinode, k the per-corner wheel rate, and Δh_node the distance to the nearest node:

$$\Pi = \varepsilon F_0 / (k \cdot \Delta h\_node)$$

Porpoising becomes mechanically accessible when Π > 1. For the LMP-class vehicle, Π ≈ 0.5 in the analysed condition, useful margin, not a universal safety law. As a practical design target, aim below Π = 0.5 when you can. Between 0.5 and 1.0, real-world perturbations matter. Above 1.0, the car is mechanically able to traverse the mode.

A SPIKE IS NOT A DEFECT

A downforce spike is not automatically bad. What matters is the spike against the spring, and there is no universal danger threshold. On a softly sprung car, a small kick, say 50 kgf, can be enough to move the chassis across the window and start the cycle. On a stiffer car that same 50 kgf barely registers; the chassis hardly moves and nothing happens. In practice the two tend to scale together: a car making enough downforce to throw a high-severity, high-frequency event has usually had to run springs stiff enough to carry that load. Speed matters as much as size. A 3 mm drift over ten seconds is imperceptible to the driver and barely shows up to the car; the same 3 mm in a tenth of a second is porpoising. The map does not ask you to flatten every peak. It tells you which peaks matter for this car, with these springs, at this speed.

BUMP STOPS

A progressive bump stop is the second defence. If it engages above the antinode and reaches high stiffness by the antinode, the car can be pinned near peak downforce without reaching the lower node. The stop must carry the peak load, the bistable snap, and a margin for bumps and curbs:

$$k\_bump \geq (\Delta F\_antinode + F\_margin) / (h\_engage - h\_antinode)$$

For the LMP-class vehicle the required order of magnitude was a few hundred N/mm over about 3 mm of travel: firm, but not exotic.

THE RACE ENGINEER'S PLAYBOOK

Two defences, one procedure. First, make sure the wheel rate is above k_min. Second, set the bump stop so the car cannot access the full node–antinode–node range. Shake down from the safe side: hard springs and conservative stops first, then soften one step at a time while watching for onset. The map gives the boundaries; the shakedown walks toward the edge without crossing it.

WHY DAMPING DOES NOT REMOVE THE ROOT PROBLEM

Damping can reduce severity and slow the motion, and in a specific setup it may even prevent traversal. But it does not remove the underlying aero bistability. The bistable snap at the antinode is not a smooth sinusoid; it is a fast force transition relative to the suspension response time. A damper resists velocity ($F_d = c \cdot v$), so at the start of the event it has almost no velocity to work with; by the time velocity builds, the chassis may already have moved through the critical region. An inerter resists acceleration ($F_i = b \cdot a$), so it helps earlier in the event, but practical inertance is finite and the effect is bounded.

This is the key point about bistability: at the antinode the flow is going to fall one way or the other. A device that resists acceleration buys time, it slows the traversal and makes a given excursion less likely to complete, but it cannot remove the fall itself, and that effect is bounded. It is like surfing a wave and dropping too fast: you drag a hand to slow the descent, but you never stop dropping, and if you drag too long you go over the front and end up right back where you started. Slowing the traversal is not a partial path to stopping it; past a point, fighting it harder just routes the energy somewhere worse.

And it does go somewhere worse. If you over-damp to suppress chassis motion, the energy can reappear at the tire sidewall, the contact patch, or in chassis vibration. On a high-downforce vehicle at 241 km/h we saw exactly that: the visible porpoising was reduced, but the tire unloaded during rebound. The problem moved from ride quality to grip and safety. The durable fixes are geometric exclusion or aero modification, springs, ride height, bump stops, and floor stays keep the car out of the dangerous range; underbody development moves the dangerous range itself.

WHAT "RAISE THE RIDE HEIGHT" ACTUALLY MEANS

Be precise here, because the static number itself does nothing. Several thousand kilograms of downforce will blow straight through any static ride-height setting; the car goes where the aero load and the spring let it go. What locks the car out of a dangerous ride height is the stiffness: spring rate, and bump stops once engaged. "Raise the ride height" and "stiffer springs" are not two fixes; they are one fix described two ways. Stiffness is the mechanism, ride height is the result. The vehicle-dynamics work on this finds the same thing: simulate the full menu of suspension responses: damping, spring stiffness, bump-stop gap and stiffness. They give little or no benefit on their own; only a significant increase in static ride height (achieved through that stiffness) delays or removes the porpoising, and only at a measured cost in downforce and lap time. Suspension cannot fix the aero. It can only keep the car out of the part of the range where the aero goes bistable.

What the Map Gives You Before Track Testing

The LMP-class configuration here has not been on track. The map identified candidate antinodes at FRH 17 mm and 21 mm before a wheel turned, and the spring-rate comparison then predicted that the chosen setup should not porpoise in the analysed condition. That is the workflow in one pass: map the force histories, flag candidate antinodes, identify the adjacent nodes, compute the spring-rate threshold or Π, compare with the actual setup, and act only if the car can mechanically traverse the mode.

The map is not just a warning light. It tells the aerodynamicist which CFD cases contain the physics. In a normal ride-height sweep, adjacent force values can look ordinary; but comparing the node at FRH 16 mm with the adjacent antinode at FRH 17 mm revealed a global reorganisation of the flow topology, more than 85% of surface nodes showed a significant eddy-frequency change, and 39% showed substantial TKE change. Without the map, that comparison would not be obvious; with it, it becomes the first place to look.

FROM MAP TO ROOT CAUSE

Once the map identifies an antinode, the aero development process is direct:

- Identify the antinode and its adjacent node.
- Compare the two CFD solutions directly.
- Use blink comparison, frequency-labelled structures, and surface/volume diagnostics to find where the topology changes.
- Trace those structures back to the geometry that sheds or feeds them.
- Modify the geometry to reduce the bistability, move it below the plank, or broaden the peak enough that the spring rate can contain it.
- Re-run the map.

In practice the structures this isolates are not anonymous: the kurtosis map flags which ride-height window to inspect, and the power spectrum identifies the responsible structures by frequency in the 3D field, so each can be labelled and tracked across ride height (this vortex in state A at one height, state B one millimetre away). For a diffuser underbody these are the longitudinal counter-rotating vortices that roll up along the diffuser edges, whose breakdown is the established downforce-loss mechanism, so the frequency-labelled structures correspond to known physics, not unnamed features. The goal is not a floor with no peaks; it is a floor whose peaks do not line up with an accessible mechanical mode.

Why You Need Margin

The porpoising number Π is not a cliff edge at exactly 1.0. Real cars do not sit perfectly still at a single ride height. Track undulations, yaw changes, tire deformation, wake from other cars, gusts, and surface-temperature changes all perturb the operating point. Most of those perturbations are transient, they do not need to repeat continuously. If the car is close to threshold, one bump or wake encounter can kick it across the antinode, and the standing wave then keeps the cycle going.

Clean CFD and wind-tunnel environments remove most of that noise. If the clean environment says the car is marginal, the track will be worse. If it says $\Pi \approx 0.5$, there is useful margin. Speed matters because aerodynamic load scales roughly with v^2: the same geometry might be harmless at 200 km/h and marginal at 300 km/h. Track surface, weather, traffic, and setup then decide whether marginal becomes visible.

Conclusions

This note addresses two questions. The race engineer asks: how do I set up the suspension so this car does not porpoise? The aerodynamicist asks: how do I change the underbody so the problem is not there in the first place? The standing-wave map serves both.

- Underbody force histories can reveal alternating stable and bistable ride-height bands.
- Porpoising requires a complete node–antinode–node sequence inside the accessible ride-height range.
- Negative kurtosis is a practical flag for candidate bistability, not a proof by itself. PSD, force traces, and flow-field comparison complete the diagnosis.
- The standing wave existing is necessary but not sufficient: the car must also be mechanically able to traverse it.
- The useful setup outputs are k_min and Π.
- Springs and bump stops control access to the mode. Aero development moves or removes the mode.
- The method flagged porpoising risk on an LMP-class vehicle before track testing, then predicted the chosen wheel rate should prevent traversal in the analysed condition.
- The same method predicted no porpoising concern on a GT-class vehicle, because the only candidate bistable point sat below plank contact.

THE TRADE-OFF, AT A GLANCE

- Overdamping: reduces motion severity; cost: lose mechanical grip, may move energy elsewhere.

- Bump stops: truncate the ride-height range; cost: constrain the setup window.
- Raise / stiffen the platform: keep the car above the mode; cost: lose peak downforce and compliance.
- Fix the aero: moves or removes the mode; cost: development, packaging, validation.

There Is No Spoon

The standing-wave map tells you where the antinodes are. The next question is why they are there: which underbody structures create the bistability, and whether geometry can move the problem below the operating window or remove it entirely.

A stopper prevents the spoon from reaching the wall. A stiffer handle prevents the water from pulling it in. Those work, but they constrain the spoon. The permanent fix is to redirect the water so the pull never appears in the first place. That is what fixing the aero means: not merely holding the car away from the dangerous ride height, but reshaping the underbody so the dangerous ride height is no longer in the usable range, or the bistability is no longer strong enough to overcome the suspension.

Redirecting the water, not constraining the spoon, is what we do.

Technical Basis & Further Reading

References are kept out of the body on purpose, this is a note, not a literature review. None of the background physics here is novel, and if any specific claim makes you want to argue, the work below is where to start. It covers diffuser ground effect and the hysteresis around peak downforce, the bistability and symmetry-breaking seen in both diffuser and bluff-body flows, the statistics of bimodal distributions and the dip test, the vehicle-dynamics setup studies, and inerter theory.

A. E. Senior and X. Zhang, "The force and pressure of a diffuser-equipped bluff body in ground effect," ASME J. Fluids Eng. 123(1), 105–111 (2001). doi:10.1115/1.1340637

A. Ruhrmann and X. Zhang, "Influence of diffuser angle on a bluff body in ground effect," ASME J. Fluids Eng. 125(2), 332–338 (2003). doi:10.1115/1.1537252

O. H. Ehirim, K. Knowles and A. J. Saddington, "A review of ground-effect diffuser aerodynamics," ASME J. Fluids Eng. 141(2), 020801 (2019). doi:10.1115/1.4040501

M. Grandemange, M. Gohlke and O. Cadot, "Turbulent wake past a three-dimensional blunt body. Part 1. Global modes and bi-stability," J. Fluid Mech. 722, 51–84 (2013). doi:10.1017/jfm.2013.83

J. Molina and X. Zhang, "Aerodynamics of a heaving airfoil in ground effect," AIAA J. 49(6), 1168–1179 (2011). doi:10.2514/1.J050369

M. Scheffer, J. Bascompte, W. A. Brock, et al., "Early-warning signals for critical transitions," Nature 461, 53–59 (2009). doi:10.1038/nature08227

J. A. Hartigan and P. M. Hartigan, "The dip test of unimodality," Ann. Statist. 13(1), 70–84 (1985). doi:10.1214/aos/1176346577

M. Gadola, D. Chindamo, P. Magri and G. Sandrini, "Analyzing porpoising on high downforce race cars: causes and possible setup adjustments to avoid it," Energies 15(18), 6677 (2022). doi:10.3390/en15186677

M. C. Smith, "Synthesis of mechanical networks: the inerter," IEEE Trans. Autom. Control 47(10), 1648–1662 (2002). doi:10.1109/TAC.2002.803532

Applied Dynamics Research · ad-research.org · amb-aero.com, Computational aerodynamics for motorsport, aerospace, and research. Sapporo & Berkeley.