

On Killing Behavioural Conformity in 2030

A Thought Experiment on Recreating LMP1H-style Racing with LMDh-level Costs and Machinery

Abstract

Le Mans Daytona hybrid (LMDh) represents a remarkably successful approach in balancing manufacturer participation, cost control, and close competition as the current generation of top-level endurance prototypes. Allowing manufacturers to use a spec chassis from one of four approved suppliers (Dallara, Oreca, Ligier, and Multimatic), along with a common hybrid system, the LMDh regulations has more than succeeded in reducing technical and financial barriers to encourage manufacturer participation while enabling cars developed under different regulatory pathways and using different Internal Combustion Engine (ICE) architectures to compete within a common performance window.

However, the extensive convergence required to achieve this also raises a sporting question: whether increasingly similar powertrain and discernible vehicle behaviour leaves insufficient scope for meaningful manufacturer-dependent differences in on-track performance. When several cars operate inside closely controlled performance envelopes, on-track outcomes may consequently depend heavily upon driver execution, traffic and strategy. While that is also desirable and puts driver skill in its deserving spotlight, the role of the cars themselves are disproportionately distinguished.

This thought experiment outlines a power unit-based hypothetical regulatory framework heavily based on the current and what was announced for the unified 2030 FIA World Endurance Championship (WEC) Hypercar and IMSA GTP platform, intended to reintroduce controlled behavioural variance and this catalyse on-track action without recreating the technical freedom and development expenditure associated with late-generation LMP1H.

The proposed framework retains:

- a common rear-wheel-drive hybrid architecture
- common energy accounting
- a regulated ICE power envelope,

while introducing a single manufacturer-selected hybrid characteristic. Rather than using a time-, speed-, engine-speed-, or acceleration-event-dependent deployment waveform, the characteristic is indexed solely to gearbox state. Manufacturers select one of three homologated deployment philosophies: A, characterised by aggressive low-gear MGU torque availability; B, characterised by moderate deployment distributed across the gearbox; or C, characterised by lower low-gear deployment and aggressive

deployment in higher gears. The manufacturer-selected characteristic is fixed for the homologation period.

The ICE regulation is separated from hybrid deployment in this framework. The ICE remains subject to a common power-versus-normalised-engine-speed (kW vs N/N_{max}) ceiling, while MGU deployment is independently bounded by the amount of MGU torque permitted in each gear. The MGU therefore operates in addition to the ICE envelope rather than competing for capacity within a combined ICE-plus-MGU power ceiling. Here, an illustrative 400 kW ICE limit and 120 kW MGU limit produce a theoretical 520 kW peak combined output, aligning with what was jointly announced by IMSA, ACO, and FIA. A 550 kW sensitivity case using up to 150 kW MGU is also proposed for future investigation. These figures are not presented as validated regulatory specifications.

The central hypothesis is that gear-indexed hybrid characteristics can generate circuit- and situation-dependent vehicle behaviour without circuit-specific deployment programming. A slow corner leading onto a short straight can favour an A-type car, whereas a higher-speed corner leading onto a long straight can favour a C-type car. B exists as a “safe option” with minimal integration complexity. Consequently, the same vehicle characteristics can produce different competitive opportunities depending on circuit topology, restart location, traffic and tyre condition.

This paper is a conceptual thought experiment rather than a numerical vehicle-dynamics study. Its purpose is to demonstrate the regulatory logic and sporting mechanism. Exact deployment calibration, thermal sizing, battery behaviour, Balance of Performance authority and competitive equilibrium remain subjects for subsequent simulation and empirical validation.

1.) Introduction

The current generation of prototype endurance racing demonstrates the advantages of regulatory convergence. The LMDh regulations were explicitly constructed around cost control, a common hybrid system and common chassis architecture, while allowing manufacturers to retain branded bodywork and proprietary engines. ([IMSA](#)) The resulting convergence has subsequently enabled LMDh (and LMH) machinery to compete within a shared performance framework via Balance of Performance (BoP).

The next regulatory generation moves further toward a unified platform. On 12 June 2026, IMSA, the ACO, and FIA announced that the top prototype categories from 2030 would use a single rear-wheel-drive technical platform. Hybrid power will remain mandatory, manufacturers may either develop their hybrid system themselves or use a designated supplier, and manufacturers may pursue either bespoke chassis development or the existing constructor-supported model. Homologation is intended to remain stable for at least five years without routine performance evolution. ([IMSA](#))

This convergence solves a substantial economic and regulatory problem. It also creates the premise for the present thought experiment. The purpose of this paper is not to argue that technical diversity is inherently preferable to technical convergence, nor to claim that contemporary prototype racing lacks competitive action. Instead, it asks a narrower question:

Can a tightly bounded, manufacturer-controlled, and self-contained variable be introduced into a highly common hybrid prototype platform such that manufacturers produce measurably different vehicle behaviour, while retaining the cost and Balance of Performance advantages of the LMDh philosophy?

The intended reference point is not the unrestricted technical freedom of LMP1H. Rather, the relevant characteristic is its capacity for **inter-manufacturer behavioural variance**: different combinations of engine, hybrid and vehicle philosophy could produce visibly different strengths and weaknesses.

The proposed answer is a fixed, gear-indexed hybrid deployment characteristic.

2.) Scope and Methodology

This work is a **conceptual regulatory thought experiment** aimed at the power unit (PU) aspect of the regulations. It does not attempt to establish an optimised A/B/C deployment table, calculate definitive lap-time deltas, establish engineering feasibility, or prove that a particular manufacturer would choose a particular profile. Such conclusions would require a complex, dedicated simulation model incorporating vehicle mass, tyre behaviour, aerodynamic characteristics, engine torque curves, gear ratios, MGU characteristics, battery state of charge, regeneration, thermal limits, fuel consumption, traction control and driver behaviour. In other words, I ain't got the time, effort, nor resource for dat big dawg

Instead, the present method is logic-based. The proposed system is evaluated by asking:

1. Does the regulatory variable have a clear and bounded definition?
2. Does it remain fixed throughout homologated operation?
3. Can it create different vehicle characteristics without circuit-specific programming?
4. Does it preserve the common resource and cost-control philosophy?
5. Does it generate plausible situations in which different characteristics are advantageous?
6. Does the regulatory mechanism avoid unnecessary state-machine complexity?

Numerical values presented in this paper are therefore **illustrative assumptions**, not outputs of an optimisation process. Future simulation is required to determine the magnitude of the resulting performance differences and to refine the precise gear-by-gear calibration.

3.) Design Objectives

The framework is constructed around seven requirements.

3.1 Behavioural variance

- Different manufacturers should be capable of producing recognisably different acceleration characteristics.

3.2 Resource parity

- Manufacturers should not gain access to different battery capacities, unrestricted electrical energy or fundamentally different hybrid architectures merely by selecting a behavioural profile.

3.3 Manufacturer commitment

- The chosen characteristic should be a genuine vehicle-design decision rather than a race-by-race software strategy.

3.4 Circuit neutrality

- The regulator should not prescribe circuit-specific hybrid deployment maps.

3.5 Regulatory simplicity

- The additional variable should be significantly simpler than an unrestricted manufacturer-specific hybrid control system.

3.6 Cost containment

- The framework should preferentially reuse the common components and control infrastructure already demonstrated by LMDh-type systems.

3.7 Endurance relevance

- The characteristic should remain meaningful across an entire stint rather than being merely a qualifying or single-lap performance trick.

These requirements motivate the central principle:

Resource parity without behavioural parity.

4.) The 2030 Baseline

The 2030 framework announced by IMSA, ACO, and FIA establishes a single rear-wheel-drive platform for the top prototype classes, with mandatory hybrid power and continued manufacturer-specific engine choice. Manufacturers retain the possibility of bespoke or supplied hybrid systems, but any such system must comply with common technical specifications. ([IMSA](#))

This paper therefore takes the following as its conceptual baseline:

- rear-wheel drive
- manufacturer-specific ICE (open engine formula)
- mandatory rear hybrid power
- highly constrained overall vehicle architecture
- common or tightly specified hybrid hardware where economically appropriate
- long homologation stability
- common sporting and energy-accounting mechanisms.

Current LMDh provides a useful technical reference, as current regulations already employ a common seven-speed gearbox, and the MGU is mechanically connected to the clutch shaft through homologated gearing rather than being a direct 1:1 engine-speed device. Current regulations also tightly control gearbox configuration and MGU integration. ([FIA](#)) The existing LMDh powertrain framework is likewise useful as a reference, because it already regulates power through an engine-speed-dependent wheel-power curve. The current nominal power curve is 500 kW, with defined lower and upper Balance of Performance margins. ([FIA](#))

This new framework deliberately departs from the current combined-power philosophy while retaining its regulatory discipline.

5.) Proposed Power Architecture

5.1 Separate regulation of ICE and MGU

The fundamental change from the earlier concept is that **ICE power and MGU deployment are not constrained by a single combined power envelope**. The ICE is subject to a common power ceiling:

$$P_{ICE} \leq P_{ICE,max}(N/N_{max})$$

where $N/N_{\max}$ represents normalised engine speed. This retains the conceptual advantages of an engine-speed-referenced power limit while allowing the detailed manufacturer ICE torque curve to remain comparatively open.

The MGU is regulated separately. Rather than being assigned space beneath a combined ICE-plus-MGU power ceiling, it is defined by its permitted deployment characteristic:

$$T_{\text{MGU}} \leq T_{\text{MGU,max}}(\text{profile, gear})$$

The total power delivered by the vehicle is therefore simply:

$$P_{\text{total}} = P_{\text{ICE}} + P_{\text{MGU}}$$

This equation is an accounting identity, not a second regulatory envelope. There is consequently no requirement that:

$$P_{\text{ICE}} + P_{\text{MGU}} \leq P_{\text{total,max}}$$

under the proposed framework.

This separation is deliberate. If the ICE is operating at its full permitted output and the selected A/B/C characteristic permits maximum MGU deployment, the MGU contribution is not artificially clipped merely because the ICE is already near its ceiling. The consequence is that different A/B/C selections can naturally produce different instantaneous peak combined powers. This creates an important trade-off:

A car may sacrifice peak power in exchange for stronger acceleration in the region where its hybrid system is most heavily deployed.

That trade-off is central to the sporting concept.

5.2 Illustrative power levels

520 kW has been reported as the target output for the 2030 regulations, with the increase in power said to be from the ICE. ([IMSA](#)) Despite that, this paper adopts an **ICE peak power of 400 kW** and an **MGU maximum deployment of 120 kW** for the purposes of demonstrating the concept. The resulting theoretical maximum combined output is therefore:

$$400 + 120 = 520 \text{ kW}$$

Substantially greater MGU capability than current racing deployment is also technically conceivable without the need of a new motor, as the current motor is officially described as having a 50 kW regulated continuous output while possessing approximately 200 kW peak power and regenerative capability. ([Bosch Motorsport](#))

A second sensitivity case should also be considered in subsequent work where the ICE peak power is at 400-450 kW while having a maximum MGU deployment of up to 150 kW, yielding a theoretical 550-600 kW peak. In which the ICE produces power closer to the current LMDh regulations while pushing the existing MGU closer to its 200 kW limit.

Neither configuration should be interpreted as a proposed final regulatory specification. They merely establish plausible conceptual reference points around which the mechanism can be discussed.

6.) Manufacturer-Specific MGU Characteristic

The proposed additional manufacturer-controlled variable is a gear-indexed MGU torque characteristic. Manufacturers select one of three predefined profiles at homologation.

6.1 Profile A – Front-loaded

Profile A prioritises immediate acceleration. Its defining characteristic is high permitted MGU torque in low gears, approximately 1st through 2nd/3rd gear, followed by substantially reduced or zero deployment in higher gears. Conceptually:

Gear	Map A - Permitted MGU Deployment
1st	Maximum or high
2nd	High or medium
3rd	Medium or low
4th	Low or zero
5th	Zero
6th	Zero
7th	Zero

Table 1. Gear state versus permitted MGU deployment for profile A

The precise values are intentionally unresolved. The intended vehicle behaviour is strong acceleration immediately after slow corners, at the beginning of acceleration zones and during restarts originating from low-speed corners. A therefore exchanges high-speed hybrid assistance for stronger early acceleration.

6.2 Profile B — Distributed

Profile B provides a moderate MGU contribution across the gearbox, close to the current deployment characteristics found in LMDh cars. Conceptually:

Gear	Map B - Permitted MGU Deployment
1st	Moderate (constant)
2nd	
3rd	
4th	
5th	
6th	
7th	

Table 2. Gear state versus permitted MGU deployment for profile B

It represents the least specialised deployment philosophy and should therefore offer broad applicability across circuit types. Profile B also provides a natural engineering baseline against which the more specialised A and C profiles can be compared. The important point is not that B must reproduce the current LMDh map exactly, but that it represents the conceptual equivalent of a relatively continuous deployment strategy, and provides the lowest relative integration complexity.

6.3 Profile C — Rear-loaded

Profile C prioritises later acceleration and high speed torque. Its low-gear MGU contribution remains non-zero so that the car is not artificially disabled in slow corners, but deployment increases substantially through the higher gears. Conceptually:

Gear	Map C - Permitted MGU Deployment
1st	Low
2nd	Low or moderate
3rd	Medium

4th	Medium or high
5th	High
6th	High or maximum
7th	Maximum

Table 3. Gear state versus permitted MGU deployment for profile C

C therefore exchanges immediate low-speed acceleration for stronger acceleration once the car has reached the medium- and high-speed portion of the gearbox.

7.) Why Gear State As the Independent Variable

Earlier formulations of the concept considered time-dependent, speed-dependent, engine-speed-dependent, and acceleration-event-dependent deployment profiles. These approaches created unnecessary regulatory states. A time-dependent system requires an answer to the question of when an acceleration phase begins, how long it lasts and whether a brief lift resets the deployment sequence. A speed-dependent system embeds road speed into the regulation and can create circuit-specific effects. An RPM-dependent system can produce repeated deployment windows after upshifts because the engine repeatedly passes through the same speed region. An acceleration-event system requires the control system to identify when acceleration begins and ends.

Gear state avoids these problems, as the defining relationship simply becomes:

gear → homologated MGU characteristic

The car does not need to know where it is on the circuit, how long it has been accelerating or how long it has been deploying. For example, a lift in 5th gear to avoid traffic does not reset anything, nor does it need to. Reapplying the throttle in 5th gear simply makes the 5th-gear characteristic available again. An upshift naturally transitions the vehicle into the next homologated characteristic. This is substantially closer to an ordinary drivetrain characteristic than a complex time-domain waveform.

8.) Torque Rather Than kW for the MGU Characteristic

The proposed A/B/C variable is most naturally expressed through MGU torque, rather than requiring a separate kW waveform.

The physical relationship is:

$$P = T \times \omega$$

A constant permitted MGU torque therefore produces increasing electrical/mechanical power as rotational speed rises, until another limit is encountered. This is conceptually useful because A/B/C is intended to describe **how much hybrid tractive effort is available in a given gear**, rather than simply how many kilowatts may be delivered for an arbitrary time interval.

The distinction between torque and power also allows the regulatory architecture to remain layered:

- **ICE:** regulated primarily through a common power-vs-engine-speed curve, seen in the current regulations
- **MGU:** regulated through gear-dependent MGU torque availability.
- **Combined vehicle:** naturally produces the sum of these two contributions.

The proposed system therefore does not require a combined ICE-plus-MGU power ceiling. The precise reference point for the torque limit (MGU shaft, gearbox input, or another prescribed measurement location) would need to be specified in a full technical regulation. This is an implementation question rather than a conceptual requirement, and for now that is assumed as the same torque sensor system used currently.

9.) ICE Character and the A/B/C Choice

The separation between an ICE power envelope and an open manufacturer torque characteristic is particularly important. A common kW vs N/Nmax ceiling does not require identical torque curves. A naturally aspirated engine and a turbocharged engine can therefore approach the same power ceiling through different torque and RPM behaviour. This therefore creates another interaction with the A/B/C profiles.

A manufacturer using a highly responsive, high-revving engine may find the front-loaded A characteristic attractive because the ICE and MGU can jointly prioritise immediate acceleration. Alternatively, a turbocharged engine with slower transient response may also find A attractive to compensate for its limitations. Likewise, a C-type vehicle can place greater emphasis on the higher-speed region where the powertrain is further into its optimal operating range.

The proposal therefore preserves engine architecture as a source of manufacturer identity rather than using the hybrid characteristic to erase it. This is consistent with the existing LMDh philosophy, in which engine design is free within specified restrictions and manufacturer-specific engines coexist with a common regulated performance framework. ([IMSA](#))

10.) A, B, and C as Vehicle Philosophies

The objective is not merely to create three software maps as “gimmicks.” Each profile is intended to encourage a corresponding vehicle-development philosophy:

A: acceleration-optimised

The manufacturer prioritises strong immediate deployment. This creates incentives to develop:

- strong traction
- tyre-friendly low-speed behaviour
- sufficient thermal capacity for concentrated deployment
- packaging capable of dealing with the resulting electrical and thermal demand
- etc.

B: generalist

The manufacturer prioritises broad applicability and comparatively straightforward integration.

C: high-speed acceleration specialist

The manufacturer prioritises later acceleration. This can interact naturally with:

- higher-downforce vehicle concepts
- higher-drag configurations
- strong medium/high-speed cornering
- sufficient thermal capacity for concentrated deployment
- packaging capable of dealing with the resulting electrical and thermal demand
- etc.

These are not mandatory design prescriptions, but are potential consequences of the chosen hybrid characteristic. The resulting differentiation is therefore not purely software-level, and the manufacturer must build a car capable of exploiting the characteristic it chooses.

11.) Sporting Mechanism

The central sporting hypothesis is that the same fixed A/B/C characteristic will become advantageous in different situations depending on circuit topology. Consider two circuits:

11.1 Slow corner feeding a short main straight

Suppose a circuit has a slow hairpin immediately before the start/finish line. The critical acceleration begins in 1st or 2nd gear. An A car has access to its strongest MGU contribution at precisely this point. A C car has comparatively limited MGU contribution available. Therefore the A car can obtain a relative acceleration advantage immediately after the restart or while attacking/defending on corner exit.

11.2 Fast corner feeding a long main straight

Now consider a circuit where the main straight is preceded by a high-speed corner. The car begins its decisive acceleration in 4th, 5th or higher. The C car is now operating inside its preferred deployment region. A, having largely exhausted its hybrid advantage in the lower gears, relies much more heavily on ICE power. The relative advantage has therefore reversed.

Variability has been induced, with no circuit-specific regulation has been introduced nor are they necessary. The circuit topology determines which characteristic becomes valuable, and the regulation merely defines the cars. This is the key sporting mechanism of the proposal.

12.) Restarts and Race Situations

Restarts provide an especially clear illustration because the acceleration opportunity immediately after a restart can have direct positional consequences. A restart from a slow-speed corner naturally rewards an A car's low-gear advantage. Inversely, a restart from a high-speed corner naturally exposes a C car's higher-gear advantage.

Consequently, the relative advantage of the profiles can vary not only between circuits but between **different restart locations on the same circuit**. The same principle applies to:

- overtaking attempts
- defensive acceleration
- traffic rejoining
- corner-exit battles
- acceleration following a brief lift
- tyre degradation
- changing grip conditions

This framework therefore creates the possibility for repeatedly changing strategic circumstances without requiring the manufacturer or regulator to alter the homologated hybrid map.

13.) Peak Power as an Intended Trade-Off

Because MGU deployment is permitted in addition to the ICE envelope, the three profiles can produce different peak combined outputs in different operating regions. For a simplified example:

A: potentially reaches approximately 520 kW in the lower gears, but provides little or no MGU assistance at high speed.

B: can approach an intermediate combined output over a wider proportion of the gearbox.

C: potentially reaches approximately 520 kW in the higher gears, while producing a lower combined peak at low speed.

Thus, peak power is no longer a universal attribute of the car. A manufacturer can effectively exchange **where peak power occurs** for a different acceleration characteristic. This is desirable because it makes the A/B/C choice a genuine design trade rather than a free performance multiplier.

14.) Stint Energy and the Meaning of “Self-Contained”

The proposed MGU characteristic remains “self-contained” over a stint because its definition contains no temporal or circuit-dependent state. The manufacturer’s characteristic is fixed:

$T_{\text{MGU}} = f(\text{profile, gear, permitted driver demand})$

It is not a function of:

- elapsed time
- distance
- circuit position
- detected acceleration phase
- previous acceleration history
- manufacturer-selected strat mode

The physical amount of energy actually consumed by A, B and C can differ, but it is not a defect. A car that chooses to use more of its finite hybrid resource in low gears will naturally have a different state-of-charge (SoC) history from a car that reserves more of its hybrid contribution for high gears. The invariant is instead the **common resource framework**. The cars remain subject to common limits on battery capacity, allowable state of charge, permissible stint energy and system capability. The essential distinction is therefore:

The resource is common; its spatial distribution within the mechanical acceleration sequence is not.

The current WEC/IMSA energy-accounting framework provides an existing example of this philosophy. A virtual energy tank is defined, with cumulative deployed energy per stint calculated from the integral of driveshaft torque sensor measurements. ([FLA](#)) This provides a natural foundation for retaining common energy accounting while allowing the hybrid contribution to be distributed differently.

15.) Regeneration and Energy Management

Regeneration is intentionally not part of the A/B/C variable. A/B/C governs positive MGU deployment. Regeneration remains a common system function governed by the ordinary braking, lift and energy-management rules. This separation prevents the framework from creating a second manufacturer-controlled energy variable. The different profiles will nevertheless produce different SoC histories. A may spend energy rapidly during low-gear deployment and therefore place greater importance on energy conservation later in the stint. C may retain more energy initially and use it more heavily later. This is an intended consequence of the concept. Specific functions such as lift-and-coast or other common energy management strategies may remain available, but they are not required to define the A/B/C characteristic. Likewise, additional ICE-based charging strategies should not be made part of the core framework. They should only be considered if subsequent modelling and simulation demonstrates a genuine systemic need.

16.) Safety Car, VSC, and FCY

Safety Car, VSC and FCY conditions initially appeared to create a major problem for the earlier temporal waveform concept. A time-dependent waveform would require the regulations to determine whether its sequence should pause, continue, restart or otherwise compensate for a neutralisation.

The gear-indexed system substantially reduces this problem. During neutralisation, **A remains A. B remains B. C remains C.** There is no temporal waveform to pause nor any acceleration-event state to reconstruct. There is no requirement to restart a counter after the field goes green. The existing virtual energy framework can continue to provide the common energy accounting.

A neutralisation may still influence the number of times each profile operates in particular gears, and therefore may produce different realised SoC histories. This is a sporting consequence that should be acknowledged rather than artificially removed. A separate neutralisation-specific A/B/C map is therefore **not part of the core proposal**. If future modelling identifies a systematic sporting or energy-accounting issue, a common neutralisation operating mode could be considered. Such a mechanism would be an external safety/sporting state rather than a fourth manufacturer profile.

17.) Traction Control and the Driver

For profile A specifically, the presence of traction control does not inherently negate its advantages. A simply provides a larger amount of permitted low-gear MGU torque. Whether that torque can actually be exploited depends on:

- available tyre grip
- vehicle traction
- throttle application
- surface conditions
- tyre wear
- traction-control intervention.

The driver therefore remains responsible for deciding how aggressively the available torque can be applied. On a fresh tyre, an A car may exploit its high low-gear torque more aggressively. On worn tyres, excessive throttle application may increase traction-control intervention or compromise tyre performance.

The manufacturer must consequently engineer the vehicle so that its chosen profile remains usable. This creates an additional interaction among manufacturer design, driver skill and tyre state rather than eliminating the distinction. No separate driver MGU override is proposed. The driver operates the characteristic through the ordinary accelerator and gearbox controls; the manufacturer-defined A/B/C characteristic itself remains fixed.

18.) Gear Selection as an Emergent Strategic Variable

Gear-indexed deployment creates one new strategic interaction that is worth acknowledging. An A driver may have an incentive to remain in a lower gear to preserve access to strong MGU deployment. A C driver may have an incentive to reach a higher gear earlier. This does not represent an additional manufacturer-controlled variable. Gear selection is already a normal driver input.

Of course, such manipulations are not free. The A driver, by virtue of sustaining high ICE RPM to retain maximum MGU allowance, will consume both fuel and SoC much quicker. Lift-and-coast and/or similar energy management strategies will have to follow in order for them to not prematurely exhaust their available stint energy. Conversely, if a C driver attempts to short shift too rapidly to reach the high gears quicker for maximum MGU allowance, the maneuver is self-defeating by construction. As the MGU is coupled to the clutch shaft rather than the transmission output, its rotational speed collapses in lockstep with engine RPM under an aggressive short-shift. Delivered MGU power remains low despite the available torque ceiling for that gear. Hence, availability of torque allowance is not sufficient for it to be usable. This effect holds independent of engine architecture. A secondary penalty may compound it: lugging the ICE at low RPM under sudden load can produce a transient torque deficit, though the severity is architecture-dependent. This is one of the biggest advantages of the framework: natural, passive deterrence against gear-based MGU exploitations.

Nevertheless, the interaction should be investigated in later vehicle modelling to ensure that manufacturers cannot exploit gearbox strategy in a manner that overwhelms the intended A/B/C distinction. The existence of this interaction may also be beneficial: the hybrid characteristic becomes part of the driver's understanding of how best to operate the car.

19.) Gear-Ratio Homologation

Current LMDh regulations provide a useful precedent for tightly controlled gearbox architectures. The gearbox uses seven forward ratios, and current regulations allow limited homologated ratio-set options. The MGU gearing is also homologated. ([FIA](#))

For the proposed framework, unrestricted event-by-event gearset selection would be undesirable because gear-indexed hybrid deployment means that the gearset determines the road-speed region in which each profile operates. The preferred conceptual approach is therefore:

A manufacturer selects its permitted gearset at homologation and retains that choice.

The same A/B/C deployment characteristic then applies in gear-space throughout the homologation period.

An alternative implementation could permit two predefined homologated gearset variants provided that:

1. the manufacturer must select one before homologation;
2. it cannot switch between them opportunistically by event;
3. the A/B/C deployment characteristic is identical in gear-space for both.

The latter arrangement would preserve limited existing gearbox freedom without converting gear selection into a circuit-specific hybrid map. A single common gearset would provide an even purer expression of the concept, but is not strictly necessary for the underlying logic.

20.) Thermal and Vehicle Integration Consequences

The new framework deliberately does not attempt to eliminate all consequences of manufacturer choice. Doing so would undermine the concept. A manufacturer selecting Profile A knowingly concentrates significant hybrid deployment into low gears. It may therefore need to engineer its vehicle around the resulting operating load through:

- cooling capacity
- sidepod packaging
- vents and ducting
- thermal management
- component margins
- airflow management

A manufacturer selecting C may face a different operating profile and may optimise the vehicle accordingly.

This is a deliberate feature. The regulations provide common component capability and hard limits, but manufacturers remain responsible for designing vehicles capable of exploiting the profile they choose. The resulting engineering freedom is bounded rather than unrestricted.

21.) Balance of Performance

The central sporting principle is to **keep A/B/C outside the Balance of Performance toolbox**. The profile is a manufacturer homologation decision and should not be changed in response to event-by-event performance. BoP should instead act on aggregate vehicle performance.

Unlike the current WEC and IMSA approaches, which can apply separate low- and high-speed power adjustments, the proposed framework would discard this second power dimension. In this framework, such a mechanism would partially counteract the deliberately speed-dependent behavioural variance created by A/B/C. A conceptual hierarchy is:

Primary actuator: ICE power

- Apply a single scalar percentage to the manufacturer-specific ICE power envelope, scaling the curve without altering its shape, while leaving its A/B/C profile unchanged.

Secondary actuator: vehicle mass

- Apply weight adjustments where power adjustment alone is insufficient.

Tertiary actuator: stint energy

- Use energy allocation only where persistent discrepancies remain.

This retains the conceptual separation:

Manufacturer chooses the behavioural characteristic, and BoP equalises overall competitive resource.

Importantly, because the MGU is no longer constrained by a combined power envelope, BoP power reduction does not require rewriting the A/B/C deployment characteristic. A car can therefore remain an A car even when its ICE power is adjusted. Its characteristic remains the same, while the magnitude of its ICE contribution is scaled proportionally across the permitted power curve. The exact controllability of such a BoP hierarchy is not claimed here and remains a subject for future work.

22.) Cost and Manufacturer Incentives

A major requirement of the proposal is that behavioural diversity should not require the return of unrestricted hybrid development. The proposed mechanism intentionally uses a simple regulatory distinction:

common hardware constraints + fixed gear-dependent control characteristic.

A manufacturer choosing a profile therefore gains a reason to optimise its car around that profile without necessarily requiring an entirely different hybrid architecture.

The financial incentive also works in both directions. A manufacturer that is highly committed to extracting maximum value from its chosen characteristic can invest in corresponding vehicle integration.

A manufacturer that is cost-sensitive can accept the prescribed profile with less extensive optimisation. This means the framework does not require every manufacturer to pursue the same level of bespoke development. That asymmetry is itself compatible with the cost-control objective.

23.) Enforcement

The new architecture is considerably simpler to police than a free-form deployment waveform. A compliance system would need to verify that:

1. the manufacturer is operating the declared A, B or C profile;
2. the permitted MGU torque/output is consistent with gearbox state;
3. the driver-demand and common system limits are respected;
4. the common stint-energy allocation is not exceeded.

The existing LMDh architecture already integrates gearbox state, powertrain control and MGU arbitration. The current regulations tightly prescribe software, control systems and homologated parameters, providing a technical precedent for regulatory monitoring of such behaviour. ([FIA](#)) The proposal therefore requires additional compliance logic but not a fundamentally new regulatory infrastructure. Exact scrutineering algorithms are outside the scope of this conceptual paper.

24.) What the Framework Deliberately Does Not Regulate

The proposal deliberately avoids attempting to prescribe:

- exact engine torque curves
- exact aero philosophy
- exact thermal architecture
- exact battery thermal strategy
- exact driver technique
- circuit-specific deployment maps
- manufacturer-specific safety-car maps
- dynamic A/B/C switching
- arbitrary manufacturer-selected gear-by-gear maps
- a time-based deployment counter
- an acceleration-event state machine

- a separate combined ICE-plus-MGU power ceiling

These exclusions are important, as the value of the framework depends partly on what it **does not** regulate.

25.) Principal Risks and Limitations

Several questions remain unresolved and should not be presented as solved by the conceptual model.

25.1 Profile calibration

- The precise MGU torque levels for A, B and C remain undetermined. The profiles must be sufficiently distinct to matter while remaining viable over an endurance race.

25.2 Competitive equilibrium

- A, B and C must not collapse into:
 - ◆ one universally superior profile
 - ◆ a profile that is categorically non-competitive
 - ◆ a situation in which B provides every useful property of A and C without their trade-off

This is a calibration problem rather than a logical defect in the framework.

25.3 Gear-selection effects

- Driver gear choice may become strategically linked to hybrid deployment. The magnitude of this effect requires investigation.

25.4 Thermal feasibility

- Different profiles may impose different thermal loads. This is a deliberate engineering consequence, but component capability must remain sufficient for all permitted profiles.

25.5 Neutralisation effects

- SC/VSC/FCY conditions may alter the distribution of opportunities between profiles. The gear-based architecture removes the principal temporal-control problem, but not every possible sporting consequence.

25.6 Balance of Performance authority

- It remains necessary to establish whether ICE power adjustment provides sufficient control authority to balance different vehicle concepts without excessive use of ballast or energy adjustment.

25.7 Exact numerical output

- The 400 kW ICE and 120 kW MGU values used here are illustrative.

26.) Why the Framework Remains a Thought Experiment

The proposal should not be mistaken for a completed technical regulation. No dedicated vehicle simulation has yet established:

- lap-time deltas
- tyre degradation differences
- optimal gear-by-gear torque values
- SoC trajectories
- thermal margins
- fuel consequences
- driver strategy
- BoP controllability
- or the magnitude of racing variance

That is acceptable within the intended scope, as the present work demonstrates a **logical regulatory architecture**. Its central claim is not that a particular A/B/C table has already been proven optimal. Instead, it is that a simple and physically interpretable mechanism exists by which a future prototype regulation could introduce bounded behavioural variance without requiring circuit-specific software, unrestricted hybrid development or a fundamentally different energy-accounting system.

Future work should therefore proceed in stages:

Stage 1: conceptual regulatory formulation.

Stage 2: simplified logic-based modelling.

Stage 3: dedicated vehicle-dynamics and endurance simulation.

Stage 4: engineering and regulatory feasibility assessment.

Stage 5: empirical sporting validation.

The current paper occupies Stage 1.

27.) Discussion

The principal conceptual advantage of the gear-indexed approach is that it transforms the proposed manufacturer freedom from an abstract mathematical waveform into a recognisable powertrain characteristic.

A time-based waveform needs to answer the question:

“When should the hybrid deploy?”

The proposed framework asks:

“In which part of the gearbox should the car’s hybrid power be strongest?”

That distinction is significant, as the latter can be homologated once and applied everywhere. It survives brief lifts, traffic, changing circuit layout and restarts without needing a temporal state machine.

The resulting sporting behaviour is therefore emergent rather than explicitly programmed. A manufacturer does not select a “slow-corner map” and a “fast-corner map.” It selects A, B or C. The circuit then determines which characteristic becomes valuable. This creates the desired relationship:

Regulation → vehicle characteristic → circuit interaction → competitive opportunity.

rather than:

Regulation → circuit-specific deployment logic → competitive opportunity.

The former is much closer to the concept of inter-manufacturer behavioural variance.

28.) Conclusion

The proposed framework seeks to introduce a narrow form of manufacturer freedom into a highly unified prototype racing platform. Its central mechanism is a **fixed, gear-indexed MGU torque characteristic**, selected by each manufacturer at homologation as one of three broad philosophies:

A: aggressive low-gear deployment

B: distributed deployment

C: aggressive high-gear deployment

The ICE remains regulated through a common power-vs-normalised-engine-speed envelope, while the MGU is regulated independently through gear-dependent deployment limits. MGU is therefore permitted to operate in addition to the ICE envelope rather than being clipped by an artificial combined ICE-plus-MGU ceiling.

An illustrative 400 kW ICE limit combined with a 120 kW MGU produces a theoretical 520 kW peak output, while a 550 kW sensitivity case with up to 150 kW of MGU could be considered in future study. These values are conceptual rather than validated specifications.

The resulting peak-power differences are themselves part of the intended trade-off. A manufacturer choosing A may accept reduced high-speed hybrid assistance in exchange for stronger low-gear acceleration; C may accept weaker initial acceleration in exchange for stronger later acceleration; B provides a comparatively distributed compromise.

The sporting mechanism is particularly relevant to endurance racing, because circuit topology can determine which profile is advantageous without the regulations ever becoming circuit-specific. A slow corner leading onto a short straight can favour A, while a high-speed corner feeding a long straight can favour C. The same principle can affect restarts, overtaking opportunities, defensive acceleration and traffic interactions.

The concept therefore attempts to recreate one narrow aspect of LMP1H: **meaningful behavioural variance between manufacturers**, while deliberately retaining the commonality, resource constraints and homologation stability associated with LMDh-style regulation.

Most importantly, the proposed variable is self-contained. A/B/C is fixed by homologation, indexed only to gearbox state, and cannot create additional battery capacity, stint energy or total ICE entitlement. The actual amount of energy consumed by each profile may differ, but the underlying resource constraints remain common. The framework is consequently best understood not as a completed 2030 regulation, but as a hypothesis:

A tightly bounded, gear-indexed redistribution of a common hybrid resource may be sufficient to reintroduce meaningful manufacturer-dependent vehicle behaviour without abandoning the economic and regulatory architecture that made LMDh successful.

Whether that behavioural variance is sufficiently large, whether all three profiles remain viable, and whether the resulting differences translate into measurable sporting value are questions for subsequent modelling and empirical study.