LUMIS

CONTEXT-DRIVEN ESG ASSESSMENT, TAILORED FOR FINANCIAL PERFORMANCE INDICATION

Criteria and Methodology for LUMIS ESG Assessments

Indian Quick Service Restaurant Sector

EDITION	VERSION	REFERENCE	RECALCULATION
2026	1.1	MET-01	Daily

This document sets out the criteria and analytical framework applied in the preparation of LUMIS ESG assessments. It is issued for the information of users of those assessments and is subject to periodic revision.

Contents

Exhibits

1. Introduction and scope

1.1 Purpose

The LUMIS indicator assesses environmental, social and governance characteristics of listed companies for a single defined purpose: to capture attributes bearing on financial performance that are not represented in conventional financial statement analysis.

Financial metrics are widely available and rapidly incorporated into market prices. Attributes that affect operating performance but sit outside the audited financial statements are less consistently observed, and are therefore of greater analytical interest. LUMIS is constructed to measure that residual.

1.2 Financial materiality as the admission test

A factor is admitted to the framework only where a stated transmission mechanism connects it to the operating performance of the assessed entity. Factors of environmental or social consequence that carry no such mechanism fall outside the scope of this indicator.

Each admitted factor carries a written rationale describing the channel through which it is expected to affect margins, cost structure, revenue durability or regulatory exposure. These rationales are summarised in Section 4 and reproduced in full in Appendix A.

1.3 Materiality is not constant across time

Conventional ESG ratings fix the relative importance of each factor between reporting cycles. The resulting assessment is current on its publication date and progressively less informative thereafter, because the underlying assumption, that the relevance of a given factor is stable over time, does not hold in practice.

A quick service restaurant operator's dependence on liquefied petroleum gas is a minor consideration under ordinary conditions. Following disruption to supply, the same dependence becomes a leading determinant of gross margin. The company has not changed and its disclosures have not changed. What has changed is the weight a rational analyst would place on that particular attribute.

Framework proposition

LUMIS separates the assessment of what a company is from the assessment of what currently matters. Company attributes are scored periodically and change slowly. The weight attached to each attribute is recalibrated continuously in response to observed events. Two further signals, franchisor reputational exposure and consumer standing, are assessed independently and applied to the result. The composite is the product of all four, and therefore moves daily even where no new company disclosure has been made.

1.4 Scope limitations

LUMIS assesses ESG attributes selected for their bearing on financial performance. It is not a measure of an entity's aggregate environmental or social impact, and a favourable LUMIS assessment should not be interpreted as such a measure. LUMIS assessments are not investment advice, do not constitute recommendations with respect to any security, and do not address creditworthiness.

2. The assessment framework

The framework comprises four assessment layers and a computation stage. This section describes the relationship between them; the layers themselves are specified in Sections 4 to 10.

Figure 1: Assessment framework

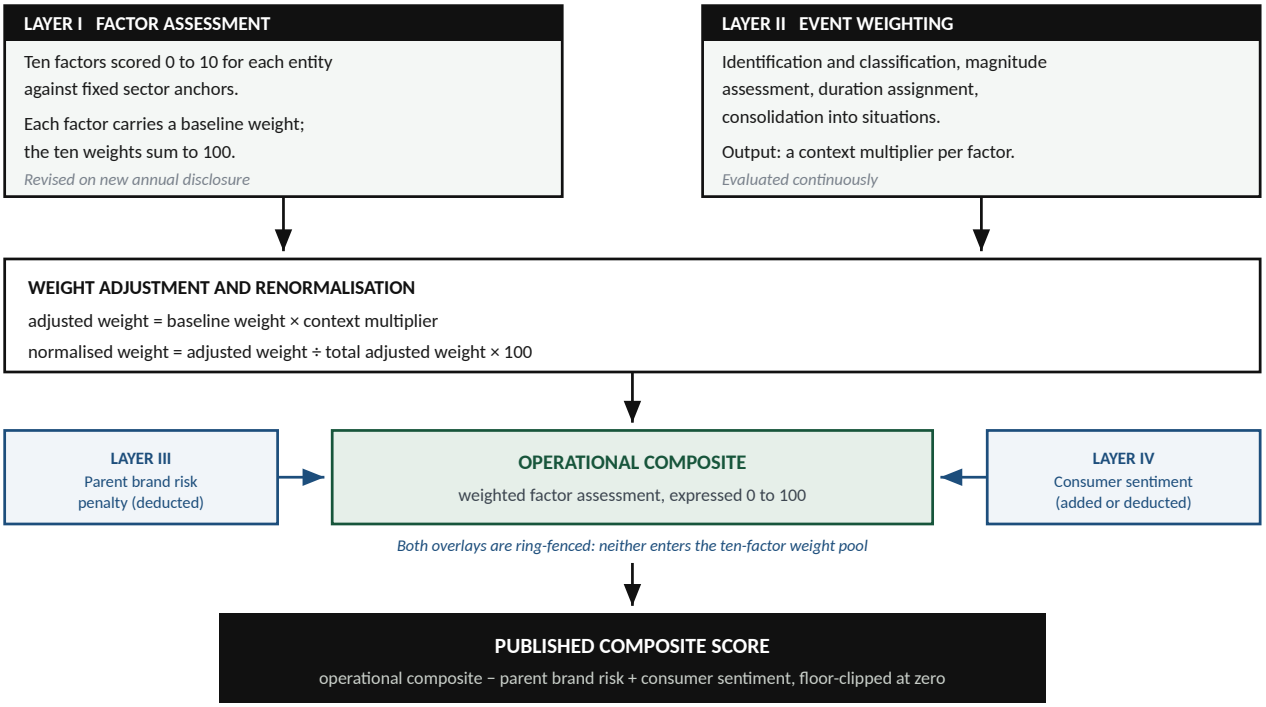

2.1 Rationale for ring-fencing the two overlays

Layers III and IV are computed independently of the ten factor framework and applied to the operational composite after it has been determined. Neither enters the weight pool, though the reasoning differs between them.

The entities within coverage operate as franchisees of international brands. A reputational event affecting a franchisor is material to the franchisee, but it is not evidence about the franchisee's own operational quality. Admitting such events into the weight pool would permit the conduct of an overseas parent to displace weight from factors measuring the domestic operator's own performance.

Consumer sentiment is excluded for a different reason. The ten factors are constructed from annual disclosures and revised annually, whereas sentiment is a continuously updating signal drawn from a live review corpus. Admitting it into a distribution calibrated on annual disclosure would subject every normalised weight to short term movement in a single series. The effect of both decisions is that a user can distinguish an entity's own operational quality from its exposure to franchisor conduct and from its standing with consumers, because the three are computed and reported separately.

2.2 Treatment of insufficient evidence

A common principle governs every threshold in the framework. Where the evidence does not support a confident determination, the framework declines to make one and applies a neutral default rather than a best estimate.

This follows from the structure of the computation. Because weights are renormalised across all ten factors, a single erroneously elevated multiplier does not distort only its own factor; it suppresses the remaining nine. An event missed leaves one factor correctly at baseline, whereas an event falsely detected corrupts the entire distribution. The two failure modes are not equivalent, and thresholds are set accordingly. The same principle governs both overlays, each abstaining explicitly rather than contributing an unsupported figure.

3. Coverage

The 2026 edition of the criteria applies to three listed Indian quick service restaurant operators.

Table 1: Entities assessed

Entity	Brands operated in India	Listing	Outlets
Devyani International DIL	KFC, Pizza Hut, Costa Coffee, Vaango, New York Fries, Tealive	NSE: DEVYANI	1,664
Jubilant FoodWorks JFL	Domino's Pizza, Popeyes, Dunkin'	NSE: JUBLFOOD	2,304
Westlife Foodworld WFL	McDonald's (West and South India)	NSE: WESTLIFE	438

3.1 Basis for sector selection

The quick service restaurant sector was selected for initial coverage on four grounds.

1. **Limited external interdependence.** Operations, customer base and regulatory environment are substantially domestic. Geopolitical developments transmit through identifiable channels, principally cooking fuel and agricultural inputs, rather than through diffuse pathways that cannot be modelled.
2. **Direct transmission to operating performance.** Energy intensity, packaging compliance and food safety arrangements affect operating cost and revenue durability through mechanisms that can be stated explicitly.
3. **Structural comparability.** The entities operate a common format, address a comparable customer demographic, and compete for the same urban demand, so cross-entity comparison is defensible.
4. **Consistency of available information.** All three entities publish annual disclosures on a common regulatory schema, providing a consistent basis for factor construction.

Structural feature of the sector

All three entities operate under international franchise agreements and are therefore exposed to franchisor reputational risk. This is not treated as an impediment to comparability but as a modelled dimension in its own right, isolated within the Layer III overlay specified in Section 9.

4. Layer I: factor assessment

Ten factors are scored 0 to 10 for each entity from published annual disclosures. This layer establishes the attributes of each entity. It does not determine how much each attribute currently matters.

4.1 Criteria for admission

A candidate factor is admitted only where three conditions hold simultaneously: a stated transmission mechanism to financial performance; consistent availability across all entities within coverage on a comparable basis; and sufficient variation between entities to be informative. Candidates failing the second condition are excluded even where conceptually material, as set out in Section 4.6.

4.2 Scoring against sector anchors

Factor scores are computed against fixed sector anchors rather than by normalising across the entities within coverage.

This is a material methodological choice. Where coverage comprises three entities, relative normalisation mechanically assigns 0 to the weakest and 10 to the strongest on every factor, irrespective of whether the underlying difference is substantial or negligible. The resulting distribution conveys rank order while destroying information about magnitude.

Under the anchor method each factor is bounded by two externally reasoned values: a plausible worst case corresponding to 0, and a plausible best case corresponding to 10. Anchors are set at levels achievable within the sector as it presently operates, not at aspirational or theoretical limits.

```
score = (value - worst anchor) / (best anchor - worst anchor) x 10
the numerator is inverted where a lower observed value represents better performance
```

One consequence is that where all entities within coverage cluster near the weak end of a plausible range, all of them receive low scores. Renewable energy share illustrates this. The strongest observed position within the sector remains a fraction of one percent, which is negligible against any defensible sector benchmark, and the anchor method reports it as such rather than presenting the least weak position as though it were strong.

Provenance of anchors

Of the anchors applied, one is derived from a binding external standard: the lower anchor for board independence is set at 33.3 percent, being the minimum proportion of independent directors required under Regulation 17 of the SEBI Listing Obligations and Disclosure Requirements. The remaining anchors are reasoned sector estimates. The distinction is disclosed rather than elided, and reasoned anchors are recorded as provisional and subject to periodic recalibration.

4.3 Ordinal factors

Three factors are ordinal rather than continuous, admitting a small number of defined states rather than a measured quantity. These are scored against the fixed bounds of their own scale, not against the observed range across entities.

- **Food safety system depth.** A composite scored 0 to 3, counting supplier certification, accredited third party laboratory testing, and a proprietary programme exceeding the regulatory minimum. Scored as (composite ÷ 3) × 10.
- **Supply chain certification coverage.** Graded 1.0 to 3.0, where 3.0 denotes full systematic certification across all suppliers. The lower bound is set below any observed entity, on the basis that a materially weaker programme is realistic within the sector.
- **Packaging compliance stage.** Three defined stages, from undisclosed progress through to a substantially completed transition.

Ordinal scoring assumes equal spacing between adjacent states. This assumption is not independently validated, and movements driven by ordinal factors accordingly warrant marginally greater interpretive caution than movements in continuously measured factors.

4.4 Modelled factors

One factor, last mile delivery carbon liability, is assessed on modelled rather than uniformly disclosed information, and is identified as such wherever it appears.

The factor measures emissions intensity arising from delivery logistics, expressed in tonnes of carbon dioxide equivalent per outlet. Only one entity within coverage discloses the relevant emissions category. The disclosed figure of 4,070 tonnes across 1,664 outlets establishes an observed rate of 2.446 tonnes per outlet for a fully outsourced delivery model.

The second entity operating an outsourced model is assigned the same observed rate, on the basis of a structurally identical delivery arrangement. The entity operating an owned electric delivery fleet is modelled by applying a 40 percent emissions reduction, being the midpoint of published lifecycle comparisons between electric and petrol two wheelers in the Indian context, to the electrically covered portion of deliveries under a 50 percent fleet electrification assumption.

Modelling assumption

The 50 percent fleet electrification figure is an assumption. The existence of the electric fleet is verified; its proportion of total delivery capacity is not publicly available and could not be independently confirmed. A neutral midpoint was adopted in preference to a figure favourable to any entity. The upper anchor for this factor is set at 75 percent fleet electrification rather than 100 percent, reflecting the realistic sector frontier.

4.5 Disclosed override

Override: packaging compliance stage

One factor score departs from its formula output. On packaging compliance stage, one entity within coverage is assigned a score above the floor of the scale, where the formula output for Stage 1 is the floor itself.

That entity's disclosures were examined directly for a quantified packaging transition metric and none was located, correctly placing it at Stage 1. The absence of a quantified metric is not equivalent to confirmed absence of progress, and both comparator entities disclose quantified figures. Assigning the absolute floor of the scale would assert a stronger negative finding than the available evidence supports. The override raises the score above the floor while preserving the entity's rank position.

No other manual override is applied to any factor for any entity.

4.6 Factors considered and excluded

Four candidate factors satisfying the materiality test were excluded for failing the comparability condition. They are recorded because their absence is itself methodologically relevant.

Table 2: Factors considered and excluded

Candidate factor	Grounds for exclusion
Employee training hours	Not available for one entity within coverage in any form. No comparable three-way metric can be constructed.
Water intensity per outlet	Reporting boundaries are irreconcilable across entities. Converted figures diverge by a factor of approximately nineteen, exceeding any plausible operational difference and indicating differing definitions of included consumption.
Used cooking oil recycling rate	Not available for one entity. Expected differentiation is minimal in any case, given mandatory compliance with the national recovery framework.
Data privacy breach incidence	All three entities report nil for the assessment period, so the factor is non-differentiating at baseline. Retained within the event framework as a dormant trigger that activates on occurrence.

4.7 Baseline weights

Each factor carries a baseline weight reflecting its assessed importance to operating performance under ordinary conditions, absent any active event. Weights sum to 100 and are held constant between revisions of these criteria. The complete factor set, with anchors, observed values, scores and weights, is reproduced in Appendix A.

5. Layer II: event identification and classification

Layer II determines which developments are materially relevant to the entities within coverage, and what form that relevance takes. It operates continuously against news and commodity price sources.

5.1 Event categories

Seven operational categories are defined, each mapped to the factors whose materiality it affects. An eighth category, parent brand risk, is defined as an overlay and specified separately in Section 9. Category definitions and affected factors are set out in Appendix B.

5.2 Text based identification

Textual sources are processed through four sequential stages. The stages are separated because each addresses a different question against a different form of evidence, and each therefore warrants an independent threshold.

Figure 2: Event identification pipeline

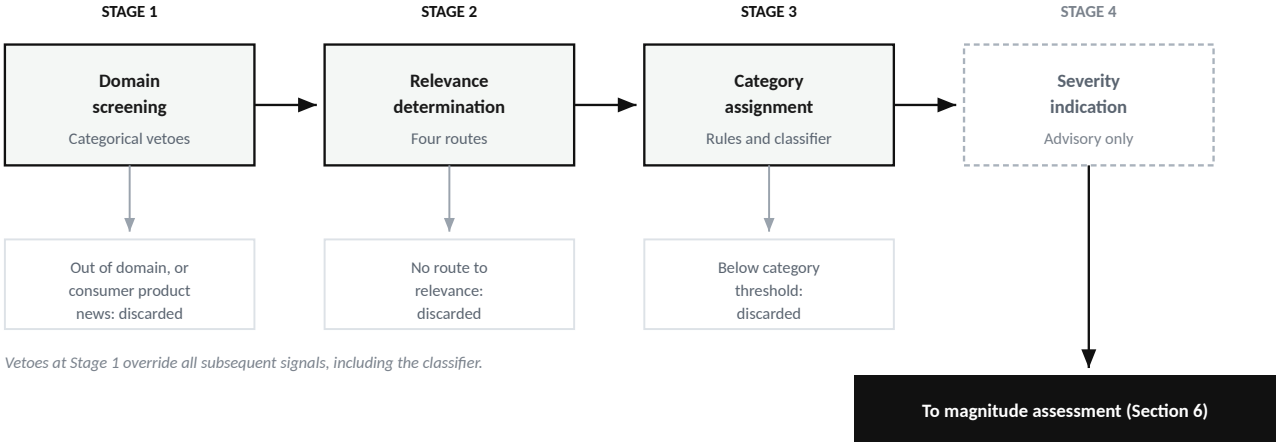

Stage 1: domain screening

Two categorical vetoes execute before any other assessment, and override all subsequent signals including the classifier.

The first excludes material outside the business domain, such as entertainment, sport, weather and elections. This veto exists because the classification model is trained exclusively on Indian business copy. Presented with text drawn from outside that distribution, the model produces confident but meaningless output. A keyword veto is more reliable than a model in this specific circumstance, because the model holds no representation of the material at all.

The second excludes consumer product news, such as menu launches, new items and promotional pricing. Items of this kind name brands within coverage and contain extensive food vocabulary, but concern consumer experience rather than ESG materiality. The veto is narrow by construction: it fires only where launch or menu language appears in the absence of any regulatory or incident term, so that a product withdrawal following contamination is routed correctly to the food safety category. Consumer experience is not discarded by the framework; it is assessed separately within the Layer IV overlay specified in Section 10.

Stage 2: relevance determination

Four independent routes may establish relevance to the entities within coverage.

1. **Named entity.** The item names the sector, a brand within coverage, or a listed peer.
2. **Named input.** The item names a material input such as cooking fuel, edible oils or packaging resin.
3. **Cross-cutting legal instrument.** The item concerns an instrument binding by legal class rather than by sector. Extended producer responsibility rules bind every brand owner, wage codes bind every employer, and data protection and board

composition requirements bind every listed company. An entity within coverage is all of these by definition, so such an instrument establishes relevance without any sector wording.

4. **Model confidence.** Where the classifier assigns high confidence to a category in the absence of rule support, that confidence is itself treated as evidence of relevance.

The fourth route was introduced following holdout measurement which established that a purely rule gated design discarded a substantial majority of genuine events. Items of the form "two outlets of a listed chain were closed following inspection" are plainly material yet contain no rule detectable term. Approximately half of all genuine events carry no rule signal.

Stage 3: category assignment

Two independent determinations are combined. A rule layer maintains, for each category, a high precision list of named regulators and instruments, together with a lower precision list of subject vocabulary. A category fires on any match against the former, or on two or more distinct matches against the latter. A multi-label classifier trained on approximately 2,700 headlines, incorporating 780 hand labelled examples oversampled to prevent dilution by weak labels, captures paraphrasing that fixed vocabulary cannot.

Where the rule layer has anchored a category, a competing classifier assignment to a different category must clear a substantially elevated confidence threshold. This prevents incidental vocabulary overlap from producing spurious secondary classifications. Each category carries an independently tuned decision threshold, reflecting that the seven classes differ markedly in separability.

Stage 4: severity indication

An advisory severity indication is computed and surfaced for analyst review. It does not feed the multiplier. An earlier design admitted textual severity into magnitude sizing directly. Measurement established accuracy of approximately 50 percent with errors in both directions, and the pathway was removed. Given renormalisation across ten factors, a confidently incorrect magnitude is more damaging than the absence of a magnitude estimate.

5.3 Price based identification

Two categories, fossil fuel and energy shock and supply and commodity shock, are additionally sized from observed commodity price movement rather than from text. This pathway is isolated from the natural language pipeline, because a report that prices have risen sharply conveys nothing about whether the movement cleared a defined threshold.

Percentage change is computed across a 30 day rolling window. Movements below 5 percent generate no event. Movements at or above 5 percent generate a first tier event, and movements at or above 15 percent generate a second tier event. Within each band the multiplier is interpolated proportionally, so that movements at the lower and upper extremes of a band are not treated as equivalent.

The third tier, reserved for geopolitical escalation and interruption of imports, is unreachable from price data alone, because no price series can express that condition. An earlier design permitted textual severity to escalate a price reading to the third tier. This was removed after a 5.89 percent price decline was escalated to a shock multiplier, the accompanying headline having been framed in terms of a crisis averted, which the severity rules could not parse for negation.

A staleness guard governs both series. Where the interval between the most recent observation and the evaluation date exceeds a defined threshold, no reading is issued, and the affected factors remain at baseline rather than being scored against outdated information.

5.4 Classification performance

Classification performance is measured on a held out evaluation set across five folds, and is reported as a frozen figure: precision of 0.727, recall of 0.754, and an F1 measure of 0.741 across the seven operational categories.

The parent brand risk overlay is architecturally distinct, assessing reputational spillover rather than enforcement action, and is evaluated separately. Folding it into this measure would conflate the performance of two dissimilar systems.

6. Magnitude assessment

Magnitude assessment determines how far an event moves the weight of the factors it affects. The framework does not infer severity from tone, adjectives or narrative framing.

6.1 Evidentiary signals

Four signals are evaluated independently for each event. Each is narrow, enumerable and high precision by construction.

Table 4: Evidentiary signals

Signal	Basis of assessment
Instrument rank	A four level ontology of enforcement instruments: advisory or show cause notice; formal notice or failed inspection; penalty imposed or licence suspended; prosecution or ordered closure. Assessed from named instruments only, not from characterisation.
Source tier	Publisher classification across three tiers. National wire services, major business dailies and regulators' own notifications constitute the first tier.
Corroboration	Count of distinct independent publishers within a 72 hour window. Syndicated reproductions of a single wire report are collapsed to one source.
Escalation trigger	A restricted vocabulary of unambiguous severity terms: recall, fatality, hospitalisation, ordered closure, prosecution, nationwide prohibition. The list is not expanded with softer synonyms, as doing so would reintroduce the tone inference failure the signal exists to avoid.

6.2 Concurrence requirement

An event is escalated above the neutral floor only where at least two signals concur. Where no signal, or a single signal, fires, the event registers at the neutral multiplier and is recorded as an abstention.

Table 5: Multiplier by signal concurrence

Concurring signals	Multiplier	Classification
0 or 1	1.50	Neutral floor. Event registered, magnitude not asserted.
2	1.80	Weak concurrence
3	2.20	Moderate concurrence
4	2.60	Strong concurrence

6.3 Calibration against market evidence

The food safety severity gradient has been tested against published event study literature. Recalls involving a confirmed health hazard are associated with approximately 1.15 percent negative abnormal return over a five day window, against approximately 0.63 percent where no health hazard is present, a ratio of approximately 1.8. The framework's corresponding span from neutral floor to strong concurrence is 1.73. The calibration is therefore consistent with observed market behaviour rather than assumed.

Comparable literature informs the governance and reputational categories. Data breach events exhibit short horizon effects of approximately 2.5 percent, with severe cases involving financial data compounding substantially over longer horizons. Boycott events exhibit a mean effect of approximately 2.7 percent but with pronounced geographic asymmetry, an observation that supports the transmission tier construction described in Section 9 rather than any single uniform multiplier.

7. Duration and decay

The influence of an event is bounded in time. Four regimes govern the behaviour of an assigned multiplier following identification. The regime is determined by the nature of the event, rather than by a uniform decay applied to all categories.

Figure 3: Duration regimes

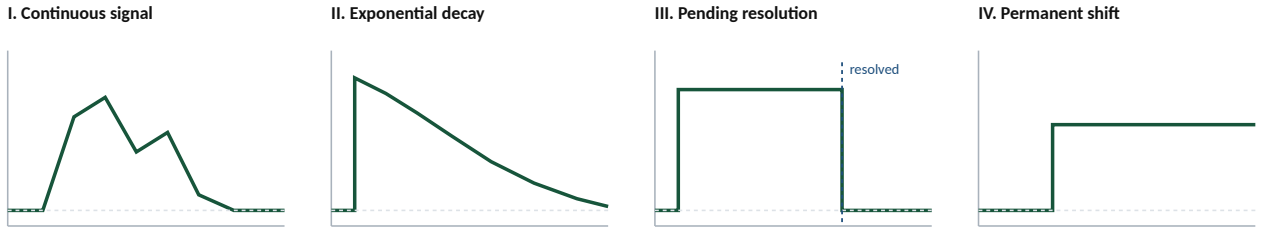

Vertical axis: effective multiplier. Horizontal axis: days since identification. Dotted line indicates the neutral floor.

```
M(t) = 1 + (M0 - 1) x 2^(-t / half life)
t is measured in days from identification; M0 is the assigned peak multiplier
```

7.1 Proposed and enacted regulation

Regulatory events are distinguished by stage. A consultation paper, draft rule or introduced bill signals prospective materiality, and is assigned exponential decay. A gazette notification, enacted statute or circular carrying a compliance deadline signals realised structural change, and is routed for permanent revision of the baseline weight.

Where the stage cannot be determined with confidence, the framework applies the consultation treatment. An unconfirmed permanent revision requires manual reversal, and is materially more damaging than an unnecessary decaying multiplier.

7.2 Direct and peer exposure

Food safety events are further distinguished by whether an entity within coverage is the subject of the incident, or whether the incident concerns a sector peer. Direct exposure carries a longer minimum active window, reflecting that resolution of an entity's own regulatory position takes materially longer than dissipation of sector wide attention. Complete duration parameters are set out in Appendix C.

8. Consolidation and combination

Related items are consolidated into single situations before any multiplier arithmetic is applied. This consolidation, rather than the combination formula, is the principal safeguard against distortion driven by volume of coverage.

8.1 Situation consolidation

Twenty items reporting a single supply disruption over three days constitute one event reported repeatedly, not twenty independent shocks. Absent consolidation, any combination rule, however conservative, would respond to the volume of coverage rather than the severity of the underlying development.

Items are clustered on four joint conditions: shared entity or common sector wide designation; shared event category; temporal proximity within a defined window, extended where a common instrument is referenced; and textual similarity above a calibrated threshold. Clustering is single link, permitting developing accounts to chain across successive reports.

A situation is not static. Items arriving subsequently that match an open cluster update its aggregate signals, and may revise its magnitude or trigger its resolution. They do not create an additional situation for combination purposes.

8.2 Combination of concurrent situations

Genuinely distinct situations may affect a common factor concurrently, as where an energy price shock and an unrelated commodity disruption both bear on delivery cost exposure. Such situations are combined by decomposing each multiplier into its increment above neutral, applying diminishing weights in descending order of magnitude, and reconstituting.

```
combined = 1 + d1 + (d2 x 0.5) + (d3 x 0.25) + ...
where d(i) = M(i) - 1, sorted in descending order of magnitude
```

Multiplicative combination is not applied, as it compounds without bound and implies that independent cost pressures amplify one another. Simple selection of the maximum is also rejected, as it discards the additional exposure created by concurrent stressors.

Abstentions do not aggregate

Only evidenced situations, being those clearing the two signal concurrence requirement or carrying a quantified price reading, contribute increments to the combination. A situation registering at the neutral floor is an explicit abstention: it records that an event is live while declining to assert its size. Aggregating abstentions would manufacture apparent significance from volume of coverage alone. The strongest active situation establishes the floor, and only evidenced situations add to it.

8.3 Ceiling

The effective weight of a factor may not exceed three times its baseline weight under any combination of concurrent situations. Where the combination formula would exceed this bound, the result is clipped and the clipping is recorded in the audit log.

8.4 Audit trail

Every combined multiplier written to the weighting model retains a complete record of its constituent situations, the increment contributed by each, the diminishing weight applied, the item count within each cluster, and whether the ceiling was engaged. Any published multiplier can therefore be reconstructed to the specific items that produced it.

9. Layer III: parent brand risk overlay

Reputational exposure arising from an international franchisor is assessed independently of the ten factor framework, and is expressed as a deduction in composite points. It does not enter the weight pool.

9.1 Structural separation

The overlay carries no factor mapping. Its output is directed to a distinct penalty term rather than to the multiplier of any factor. The ten baseline weights, and their summation to 100, are structurally unaffected by any parent brand risk event.

An implementation directing overlay output to the multiplier of an existing factor would satisfy the arithmetic constraint while defeating the purpose, permitting franchisor conduct to displace weight from factors measuring the domestic operator's own performance. The separation is enforced structurally rather than by convention.

9.2 Franchisor mapping

Table 6: Entity to franchisor mapping

Entity	Franchisor	Weighting basis
DIL	Yum! Brands (KFC, Pizza Hut); The Coca-Cola Company (Costa Coffee)	Majority weight to flagship
JFL	Domino's Pizza Inc; Restaurant Brands International (Popeyes); Inspire Brands (Dunkin')	Majority weight to flagship
WFL	McDonald's Corporation	Full weight

Franchisors are weighted by revenue and outlet share within each entity rather than equally. Two franchisors are excluded from quantitative model training while remaining active for event identification. One is privately held and has no observable security. The other is a diversified enterprise of sufficient scale that a brand specific event would be statistically indistinguishable from unrelated movement.

9.3 Transmission tiers

A franchisor event does not transmit uniformly to domestic operations. A recall confined to a single overseas market and a coordinated multi market boycott are not equivalent exposures, and the framework does not treat them as such.

Table 7: Transmission tiers

Tier	Determination	Treatment
Tier I `HIGH`	Explicit regional relevance, such as named South Asian or Middle Eastern markets, or boycott mobilisation observable in domestic coverage	Full transmission
Tier II `MODERATE`	Global brand equity or conduct event without explicit regional anchoring, but with plausible spillover through shared brand identity	Substantially reduced
Tier III `LOW`	Event affirmatively confined to a single overseas market with no identifiable transmission channel	Minimal

Tier III requires positive evidence of confinement. An event naming no market is assigned Tier II rather than Tier III, on the basis that the absence of geographic specification is not evidence of geographic containment. The default posture of the framework under ambiguity therefore assumes moderate rather than negligible transmission.

Tier assignment is evaluated per item rather than against concatenated cluster text, so that a single regionally relevant report within a larger cluster is not diluted by unrelated coverage.

9.4 Role within a shared incident

Within a single situation, entities within coverage may occupy materially different positions. Where contamination is traced to a shared supplier, one franchisor may be the subject of confirmed harm while another withdraws the affected input as a precaution, with no illness attributed to it.

These are not equivalent reputational events. The framework classifies the role of each attributed entity, per item, as originator, being confirmed harm arising from its own product, or responder, being precautionary action tied to a third party incident. Responder classification attracts a discount grounded in the same published severity gradient applied elsewhere in the framework.

Where the role cannot be determined with confidence, the framework applies originator treatment and no discount. An unverified precautionary framing does not reduce an assessed penalty in the absence of supporting language. Aggregation across a situation is asymmetric: a single item establishing originator status prevails over any number of items framed as precautionary, because precautionary framing is the account a culpable party has the strongest incentive to place in circulation.

9.5 Attribution requirement

An entity is attributed to a multi entity situation only where at least one item within that situation names one of its own franchisors. Incidental co-mention within an item concerning a different brand does not establish attribution. This requirement is enforced by regression test across the full live feed.

9.6 Penalty computation

Situation severity is converted to a penalty expressed in composite points. The conversion is uncapped but concave. Severity carries no fixed maximum, so that a genuinely severe event produces a correspondingly severe penalty, while the concave form prevents any single input from scaling the output without bound.

Concurrent situations are combined under the same diminishing weight construction applied within the operational layer, preventing repeated coverage of one crisis from compounding a penalty. The final composite is floor clipped at zero. This clipping is a presentation constraint, and does not limit the penalty computation itself.

9.7 Validation

The overlay is validated against a paired historical test constructed to detect over transmission as well as under detection. A coordinated multi market boycott and a single market contamination outbreak, both concerning the same franchisor, are assessed against the same entity within coverage. The framework separates the two by a factor exceeding the required threshold, with the boycott producing materially the larger penalty. A framework producing comparable penalties for both cases would have identified two events without having learned the transmission distinction.

10. Layer IV: consumer sentiment overlay

Consumer standing is assessed independently of the ten factor framework and expressed as a bidirectional adjustment in composite points. Like Layer III it does not enter the weight pool. Unlike Layer III it may raise a composite as well as reduce it.

10.1 Structural separation and direction

The rationale for excluding this layer from the weight pool is set out in Section 2.1. In summary, the ten factors are constructed from annual disclosures and revised annually, whereas sentiment is a continuously updating signal drawn from a live corpus of consumer reviews. Admitting a continuously updating input into a distribution calibrated on annual disclosure would subject every normalised weight to short term movement in a single series.

The layer is bidirectional, which distinguishes it from Layer III. A reputational event originating with a franchisor cannot constitute favourable evidence about the assessed entity, so Layer III is a deduction only. Consumer standing is different in kind: a genuine improvement in how an entity is regarded by its customers is a real and creditable development bearing on revenue durability, and the framework reports it as such rather than merely declining to penalise it.

10.2 Basis of measurement

Sentiment is assessed from a corpus of consumer reviews. Each review is first tested for relevance and then assigned a polarity, and the resulting values are aggregated into two figures for each brand: a long run baseline rating establishing that brand's established position, and a live rating reflecting its current position. Both are expressed on a scale of 0 to 100.

The overlay keys on the deviation between them, being the live rating less the baseline rating, rather than on the raw live rating. This is consistent with the treatment of context throughout the framework. What is measured is departure from an entity's own established position, not its absolute standing, so that a brand with a structurally lower rating is not perpetually penalised for a position already reflected in its baseline.

10.3 Brand coverage and aggregation

Sentiment is measured for each brand. Entities operating several brands therefore require aggregation to a single figure. Brands are weighted by their share within the entity, on the same revenue and outlet basis applied to franchisor weighting in Section 9.2.

Table 8: Sentiment brand coverage by entity

Entity	Brands within sentiment coverage	Brands outside coverage	Covered share
DIL	KFC, Pizza Hut	Costa Coffee	0.90
JFL	Domino's Pizza	Popeyes, Dunkin'	0.90
WFL	McDonald's	None	1.00

Brands outside coverage are excluded from the aggregation and the covered share is reported alongside the result, rather than an uncovered brand being treated as neutral. Treating an unmeasured brand as neutral would understate the deviation of the brands actually measured.

Disclosed assumption: brand split within a shared franchisor block

One entity within coverage operates two brands under a single franchisor block, and the outlet split between those two brands is not separately available. In the absence of that figure the two brands are weighted equally within the block.

This is an assumption standing inside a published figure, and it is recorded as such in every output artefact rather than presented as measured. It is not a calibration parameter but a substitute for a datum that does not presently exist, and a single disclosed outlet split would retire it. Where that entity's two brands diverge materially in sentiment, the aggregated deviation is correspondingly sensitive to this assumption.

10.4 Gating

Two gates govern whether the overlay contributes at all. Each produces an explicit abstention rather than a silent zero, so that a recorded absence of contribution is distinguishable from genuinely neutral sentiment.

- **Minimum sample.** Where the aggregated review count for an entity falls below a defined threshold, the overlay contributes nothing and records the reason. A deviation computed from a small number of reviews is not a reliable measurement of that entity's standing.
- **Data currency.** The overlay reads the age of the most recent scored observation. Where that age exceeds fourteen days, the overlay contributes nothing and records the reason. An overlay reading stale sentiment while presenting a current figure would be materially worse than no overlay at all, and this gate is load bearing rather than precautionary.

10.5 Conversion to composite points

The aggregated deviation is converted to an adjustment in composite points under a signed concave transformation.

```
adjustment = sign(d) x scale x ln(1 + |d|)
where d is the aggregated deviation, being live rating less baseline rating
```

The construction mirrors the penalty conversion in Section 9.6 in being uncapped but concave. A genuinely large deviation in either direction produces a correspondingly large adjustment, while the concave form prevents any single input from scaling the output without bound. In practice a threefold widening of the input range produces substantially less than a threefold widening of the resulting adjustment.

The concave function applied in Section 9.6 could not be reused directly. That function clamps negative inputs to zero, which is correct for a penalty that cannot be favourable, but would silently discard the entire signal for any entity whose sentiment had deteriorated. A signed variant preserving the identical curve is applied here, and the distinction is enforced by test.

The scale factor is provisional. It is set so that an adjustment remains a visible but subordinate contribution relative to the operational composite and the parent brand risk penalty, and is recorded for recalibration as observation accumulates.

10.6 Duration

No separate decay mechanism is applied to this layer. The underlying live rating is itself a moving average which attenuates older observations as new reviews arrive, and the deviation therefore returns toward zero without intervention once an episode has passed. The framework does not impose a second decay on top of it, which would attenuate the same signal twice.

The currency gate specified in Section 10.4 governs the case the moving average cannot address, namely a corpus that has stopped updating altogether, in which the most recent deviation would otherwise persist indefinitely as though current.

11. Composite score computation

The composite is recalculated daily. Company attributes enter through Layer I and are revised annually. Weights enter through Layer II and are revised continuously. Both overlays are evaluated on each recalculation.

```
01  adjusted weight      = baseline weight x active multiplier
02  normalised weight    = adjusted weight / total adjusted weight x 100
03  operational (0-10)   = SUM (normalised weight / 100 x factor score)
04  operational (0-100)  = operational (0-10) x 10
05  final                = operational (0-100) - parent brand risk penalty
                                 + consumer sentiment adjustment
06  published            = max(final, 0)
```

Renormalisation at step 02 is the mechanism through which context alters the assessment. Elevating the weight of one factor necessarily reduces the proportional influence of the remainder. An entity performing strongly on a factor that has become materially relevant will see its composite rise, and an entity performing poorly on that same factor will see its composite fall, without any change to the underlying attributes of either.

The two overlay terms at step 05 are applied after the operational composite has been determined in full. Neither participates in the renormalisation at step 02.

11.1 Illustrative computation for a single assessment day

The computation below is an illustration. It reproduces, in full, how the framework arrived at a published composite for one entity on one particular day. It is presented to demonstrate the mechanics of the framework, namely how context multipliers alter the effective weight of each factor, how those weights are renormalised, and how the two overlay terms are applied to reach the final figure.

The multipliers and overlay values shown were those active on that day. They are not fixed values and will differ on any other day, as will every figure derived from them.

Table 9: Illustrative computation for one entity on a single day

Factor	Baseline weight	Multiplier	Adjusted	Normalised %	Score	Contribution
Energy intensity per outlet	11.0	1.9135	21.049	10.123	9.22	0.9333
GHG emissions intensity	10.0	1.5000	15.000	7.214	9.06	0.6536
Renewable energy share	6.0	1.9135	11.481	5.522	0.67	0.0370
Fossil fuel share of energy	13.0	1.9135	24.875	11.963	5.58	0.6676
Packaging compliance stage	8.0	1.5000	12.000	5.771	5.00	0.2886
Female workforce representation	9.0	1.8000	16.200	7.791	4.24	0.3303
Supply chain certification	12.0	3.0000	36.000	17.314	10.00	1.7314
Board independence ratio	9.0	1.5000	13.500	6.493	4.00	0.2597
Delivery carbon liability	11.0	2.2568	24.825	11.939	4.30	0.5134
Food safety system depth	11.0	3.0000	33.000	15.871	10.00	1.5871
Total	**100.0**		**207.930**	**100.000**		**7.0019**

The operational composite of 7.0019 rescales to 70.0188. A parent brand risk penalty of 12.3783, comprising three evidenced situations combined under diminishing weights, is deducted. A consumer sentiment adjustment of 2.9507 is then applied, negative on this day because the entity's aggregated live rating stood below its own established baseline. The published composite for that day is 54.6898.

```
70.0188 - 12.3783 + (-2.9507) = 54.6898
```

11.2 Illustrative assessment across the coverage universe

Extending the same single day illustration to all three entities shows how the operational composite and the two overlay terms combine, and how the three components move independently of one another. On this day the consumer sentiment term is the sole positive contribution recorded for any entity, and it is sufficient to alter the rank ordering relative to the operational composite alone.

Table 10: Illustrative assessment, all entities, single day

Rank	Entity	Operational composite	Parent brand risk	Consumer sentiment	Published composite
1	**Jubilant FoodWorks**	60.79	−4.99	+2.41	58.21
2	**Devyani International**	70.02	−12.38	−2.95	54.69
3	**Westlife Foodworld**	56.00	−10.67	−2.90	42.43

Illustration only, for a single assessment day. Assessments are recalculated daily and will differ from the figures shown. Current assessments are published on the ratings page.

The illustration demonstrates a property of the framework worth stating explicitly. The entity with the strongest operational composite is not the entity with the highest published composite. Devyani International leads on the weighted factor assessment by more than nine points, and finishes second once franchisor exposure and consumer standing are applied. An assessment reporting only the operational layer would not convey this.

12. Materiality of score movement

Not every movement in a daily assessment warrants attention. The framework distinguishes routine fluctuation from material change through two independent tests.

12.1 Magnitude test

The magnitude test is graduated. Until an entity has accumulated sufficient observed history, a fixed threshold of four composite points applies. This is a provisional value and is disclosed as such.

Once twenty or more observations have accumulated, the framework transitions to a volatility based test. The rolling standard deviation of day over day movement is computed for that entity across its most recent observations, and a movement exceeding two standard deviations is designated material. The rolling figure is recomputed on each assessment and refines as the observation series lengthens.

The transition occurs per entity and without intervention. The framework reports which test governed each entity on each assessment, so that the basis of any designation is unambiguous.

Basis for a graduated test

A volatility based threshold requires an observed distribution of normal movement. Constructing one from simulated or backfilled values would produce a statistically framed threshold resting on observations that were never made. The framework therefore operates a disclosed fixed threshold until genuine history exists, rather than asserting statistical rigour it has not yet established.

12.2 Rank order test

Independently of magnitude, any change in the relative ordering of entities within coverage constitutes a material event and is always reported. For a comparative indicator, a modest movement that reorders the assessed universe may carry greater analytical significance than a larger movement that preserves it.

13. Self-training quantitative layer

Alongside the rule based magnitude assessment, LUMIS operates a quantitative event study model which learns the relationship between event characteristics and realised market outcomes. It is trained and operational, and is presently authorised to override nothing.

13.1 Construction

The model is trained on 261 historical events with verified dates, drawn from statutory material event filings for internationally listed comparators and from verified domestic regulatory actions. Filing dates carry legal precision, eliminating the date ambiguity that would otherwise corrupt outcome measurement.

For each event, abnormal return is computed against a market model. The security is regressed on a sector benchmark across a clean pre-event window, isolating the return attributable to the event from broad market and sector movement. Cumulative abnormal returns are computed at five, thirty and ninety day horizons, and weak fits are down weighted rather than treated as equivalent evidence.

Published event study findings enter as informative priors, so that the model begins from established knowledge of comparable event types rather than from zero. Validation is walk forward exclusively, training on events preceding a rolling cutoff and testing on those following it, because random partitioning would leak subsequent information into training.

13.2 Authorisation gates

The model may supersede the rule based assessment for a given category only on satisfying three conditions simultaneously.

1. **Sufficient observation.** At least forty verified historical events within the category.
2. **Demonstrated calibration.** Predicted confidence intervals must contain realised outcomes at approximately their stated frequency on held out data.
3. **Informative precision.** The predicted interval must be narrow enough to constitute a more informative estimate than the rule based assessment it would replace.

Table 11: Authorisation gate status by category

Category	Observations	Calibration	Status
Food safety incident	97	0.69	Deferred: interval precision
Labour and social regulation	76	0.84	Deferred: interval precision
Governance and data event	61	0.77	Deferred: interval precision
Supply and commodity shock	19	n/a	Deferred: observation count
Parent brand risk	8	n/a	Deferred: observation count

Every category presently defers to the rule based assessment. Three satisfy both the observation and calibration requirements, and all are held at the precision gate, because an honestly calibrated interval on short horizon abnormal returns spans too wide a range to improve on the rule based estimate.

Self-training model, shortly to enter operation

This outcome represents the gate performing its function. The model has assessed its own predictive precision, found it insufficient to displace a validated rule based estimate, and declined to override without discretionary intervention.

The framework closes this gap on its own. Events identified in live operation enter a review queue, mature for ninety days until a realised market outcome exists, and following verification join the training corpus at the next scheduled retraining. Each cycle re-evaluates all three gates. When the evidence base for a category becomes sufficient, quantitative assessment for that category activates automatically, and each subsequent assessment within it is sized from observed market behaviour rather than from rules.

Verification is conducted by analysts. Live classification operates at approximately 73 percent precision, and admitting unverified classifications directly into training would compound identification error into the learning process. The framework therefore accumulates evidence at the pace at which it can be confirmed.

14. Information sources and controls

Table 12: Information sources

Source	Application
Published annual disclosures	Construction of all Layer I factors. Filed by each entity within coverage on a common regulatory schema.
Annual reports and investor communications	Supplementary operational data, outlet counts and governance composition.
News aggregation	Textual event identification across the seven operational categories and the parent brand risk overlay.
Consumer review corpus	Construction of the Layer IV sentiment overlay. Reviews are scored for relevance and polarity and aggregated per brand.
Petroleum Planning and Analysis Cell	Indian crude basket price series for energy shock assessment.
Commercial cylinder price notifications	Cooking fuel cost exposure. Commercial grade pricing is tracked, being the grade used in commercial kitchens.
Statutory material event filings	Historical event corpus for the quantitative learning layer, providing legally precise event dating.

14.1 Provenance controls

Every price observation carries a provenance record identifying its source. Any series lacking a verified source is flagged and excluded from scored output by an explicit gate. Such series may be used in testing but cannot reach a published assessment.

Series are additionally subject to the staleness guard described in Section 5.3, and the sentiment corpus to the currency gate described in Section 10.4. Where an observation window cannot be constructed from sufficiently current information, no reading is issued and the affected component remains at baseline.

14.2 Audit trail

Each published assessment retains the baseline weights applied, each active multiplier together with its constituent situations, the item level provenance of each situation, the normalised weight distribution, both scale expressions of the operational composite, the parent brand risk penalty with its constituent contributions, the consumer sentiment adjustment with its covered brand share and gate status, and the resulting composite. Any published figure can be reconstructed from its inputs without re-execution.

A. Factor reference

The table below sets out the factors assessed, their baseline weights, the anchors against which each is scored, and the mechanism by which each is expected to affect financial performance. Entity level scores are not published in these criteria; current assessments are published on the ratings page.

Factor	Baseline weight	Anchors 0 to 10	Direction	Transmission mechanism
Energy intensity GJ per outlet	11	1200 to 350	Lower is better	Direct exposure of operating margin to electricity tariff and fuel price movement.
GHG intensity tCO_2e per outlet	10	220 to 50	Lower is better	Forward exposure to carbon pricing and mandatory emissions disclosure.
Renewable share % of energy	6	0% to 3%	Higher is better	Insulation against grid tariff escalation across the transition horizon.
Fossil fuel share % of energy	13	75% to 15%	Lower is better	Principal channel through which commodity and geopolitical developments reach operating margin.
Packaging compliance stage 1 to 3	8	Stage 1 to 3	Higher is better	Exposure to producer responsibility penalties and transition cost.
Female workforce % of headcount	9	15% to 48%	Higher is better	Proxy for hiring pool depth and attrition exposure in constrained urban labour markets.
Supply chain certification ordinal 1 to 3	12	1.0 to 3.0	Higher is better	Supply resilience during disruption, and exposure to recall and regulatory penalty.
Board independence % independent	9	33.3% to 75%	Higher is better	Oversight quality, and exposure to governance related market penalty.
Delivery carbon liability tCO_2e per outlet	11	3.0 to 1.712	Lower is better	Delivery fuel cost exposure and forward electrification mandate risk.
Food safety depth composite 0 to 3	11	0 to 3	Higher is better	Leading indicator of incident probability. Incidents materially affect same outlet sales.
Total	**100**			

The lower anchor for board independence is set at the minimum proportion of independent directors required under Regulation 17 of the SEBI Listing Obligations and Disclosure Requirements. All other anchors are reasoned sector estimates, recorded as provisional and subject to periodic recalibration.

B. Event categories and affected factors

Table 3: Event categories

Category	Affected factors	Principal sources	Illustrative trigger
Fossil fuel and energy shock	Fossil fuel share; energy intensity; renewable share; delivery carbon liability	Crude basket; commercial cylinder notifications	Supply disruption affecting cooking fuel availability or price
Carbon and emissions regulation	GHG intensity; renewable share; packaging compliance	Environment ministry notifications; disclosure circulars	Carbon pricing mechanism; mandatory emissions disclosure
Labour and social regulation	Female workforce; supply chain certification	Labour ministry notifications; provident fund circulars	Wage floor revision; workforce classification change
Food safety incident	Food safety depth; supply chain certification	Food regulator notices; state enforcement; sector reporting	Licence suspension; contamination recall; standards revision
Plastic and packaging regulation	Packaging compliance stage	Pollution control notifications; gazette	Category prohibition; producer responsibility penalty revision
Governance and data event	Board independence; dormant privacy trigger	Securities regulator orders; incident notifications	Enforcement action; disclosed data breach
Supply and commodity shock	Supply chain certification; delivery carbon liability; fossil fuel share	Edible oil indices; import notifications	Input price spike; import restriction on a material input
Parent brand risk `Layer III overlay`	None. Directed to composite penalty.	International news across franchisor entities	Boycott; franchisor conduct event; multi market withdrawal
Consumer sentiment `Layer IV overlay`	None. Directed to composite adjustment.	Consumer review corpus	Sustained movement in a brand's live rating away from its established baseline

C. Duration parameters

Category	Condition	Regime	Parameters
Fossil fuel and energy	Price derived	Continuous signal	Tracks series; resolves on return within bounds
	Geopolitical text	Exponential decay	45 day half life; 7 day minimum
Carbon and emissions	Consultation	Exponential decay	60 day half life; 14 day minimum
	Enacted	Permanent shift	Baseline revision on confirmation
Labour and social	Consultation	Exponential decay	45 day half life; 10 day minimum
	Enacted	Permanent shift	Baseline revision on confirmation
Food safety	Direct exposure	Pending resolution	30 day minimum; 90 day maximum
	Peer or sector	Pending resolution	14 day minimum; 45 day maximum
Plastic and packaging	Consultation	Exponential decay	60 day half life; 14 day minimum
	Enacted	Permanent shift	Baseline revision on confirmation
Governance and data	Data incident	Pending resolution	21 day minimum; 120 day maximum
	Enforcement	Pending resolution	30 day minimum; 90 day maximum
Supply and commodity	Price derived	Continuous signal	Tracks series; resolves on return within bounds
	Restriction text	Pending resolution	14 day minimum; 60 day maximum
Parent brand risk	Flash conduct	Exponential decay	10 to 14 day half life
	Sustained boycott	Pending resolution	60 to 90 day minimum; 270 to 365 day maximum
	Franchisor distress	Decay or permanent	Permanent on insolvency or restructuring
Consumer sentiment	All	Moving average attenuation	No imposed decay. Currency gate at 14 days.

D. Glossary

Term	Definition
Abstention	A registered event or overlay input that does not clear its governing threshold. Its presence is recorded, its magnitude is not asserted, and it contributes nothing to the composite.
Anchor	A fixed reference value defining the 0 or 10 point of a factor's scale, set at a level plausibly achievable within the sector.
Baseline rating	A brand's long run established position within the consumer review corpus, against which its live rating is compared to produce a deviation.
Baseline weight	The assessed importance of a factor under ordinary conditions, absent any active event. The ten baseline weights sum to 100.
Context multiplier	The coefficient applied to a baseline weight in response to an active event. A value of 1.0 denotes no active adjustment.
Covered share	The proportion of an entity's brand portfolio, by revenue and outlet weight, for which a sentiment measurement exists. Brands outside coverage are excluded from aggregation rather than treated as neutral.
Deviation	A brand's live rating less its baseline rating. The input to the Layer IV overlay, expressing departure from an established position rather than absolute standing.
Evidenced situation	A situation clearing the two signal concurrence requirement, or carrying a quantified price reading, and therefore eligible to contribute to combination.
Live rating	A brand's current position within the consumer review corpus, computed as a moving average which attenuates older observations as new reviews arrive.
Operational composite	The weighted factor assessment before application of either overlay, expressed on a scale of 0 to 100.
Originator and responder	Within a shared incident, the entity whose own product caused confirmed harm, as distinguished from an entity taking precautionary action in response.
Ring-fenced	Computed independently of the ten factor weight pool and applied to the operational composite after it has been determined. Both Layer III and Layer IV are ring-fenced.
Situation	A cluster of related items treated as a single underlying development for magnitude and combination purposes, irrespective of the volume of coverage.
Transmission tier	The assessed degree to which a franchisor event reaches domestic operations, graded across three levels by geographic relevance.

Revision

These criteria constitute the 2026 edition, version 1.1. Version 1.1 introduces the Layer IV consumer sentiment overlay specified in Section 10, together with consequential revision to Sections 1, 2, 5, 11 and 14 and to Appendices B, C and D. LUMIS reviews its criteria periodically and on the occurrence of developments that materially affect the analytical framework. Parameters identified in these criteria as provisional are subject to recalibration as operational data accumulates.