

Estate Footfall Economics: Planned Traffic Does Not Follow the Footprint

Indivd Analytics

1 October 2026

1 At a glance

- **Tag:** ESTATE · GEOGRAPHY · DENSITY
- **Title:** A retail estate where footfall does not follow the footprint
- **Hook:** In one week across a 201-location, 16-market estate, a market holding 49 stores (24 % of the footprint) is forecast to contribute just ~0.7 % of the planned-trip total (a single-week point estimate), while a six-store market is forecast to contribute ~23 % of it (two of the top four forecast shares are single-venue markets). On this estate's forecast, store count does not by itself rank markets by demand — and that is exactly the signal a footprint map hides.
- **Key stat:** Store count does not rank markets by demand (forecast basis): one of the estate's largest store networks — 49 stores, 24 % of the footprint — carries ~0.7 % of the week's planned trips, i.e. ~0.03× the per-store average (its single-week band spans ~0.002–0.07× — see Uncertainty), while a six-store Iberian market is forecast at ~7.7×. (The single largest network, 76 stores, tracks the average at ~1.03×. Benchmarked against the multi-store markets only — the average otherwise includes three very high-volume single-venue markets — the same two markets read ~0.05× and ~11.7×.)
- **Chart:** per-store forecast index by market (bar chart; regenerated from the frozen capture above)

2 The pattern

Footfall planning is usually done on a map: count the stores, rank the markets, ship stock and staff accordingly. On a real multi-market estate the map is misleading, because **a store's footfall does not scale with how many stores its country has**.

Across an anonymized 201-location, 16-market retail estate (current forecast vintage, week of 1–7 October 2026), **traffic density per store is not proportional to market size**. One of the estate's largest store networks — **49 stores, 24 % of the footprint** — is forecast at only ~0.03× **the per-store average**, contributing just ~0.7 % of the planned-trip total, while a six-store Iberian market is forecast at ~7.7× the average, contributing ~23 %. The estate's single largest network — **76 stores, 38 % of the footprint** — sits almost exactly on the average (~1.03×). One caveat on the benchmark: the per-store average includes three single-venue markets that trade at very high levels, which lifts it. Benchmarked against an average of the multi-store markets only (~13,160 trips/store — 2,606,169 planned trips across 198 venues), the 49-store network reads ~0.05× and the six-store market ~11.7×. The spread between those two markets is unchanged by the benchmark choice (their per-store ratio is a property of those two markets — ~258× under either average); the direction of the finding does not depend on the benchmark. These figures are the estate's weekly demand **forecast** (planned visitor-trips), not observed counts — what the figures describe is the allocation pattern, not the absolute level.

The counter-intuitive consequence — a hypothesis for validation on observed data, not a prescription from forecast data alone — is that **in this estate, allocation decisions (stock depth, staffing, marketing weight) may belong on the density curve rather than on the map**. Between the

⁰Anonymized, estate-aggregated. Data vintage: forecast batch of 2026-10-01. Forecast, not observed; the frozen forecast capture is the single source for every figure below.

estate’s multi-store markets (single-venue markets excluded), per-store planned demand spans a more-than-250-fold range — from $\sim 0.03\times$ the estate per-store average ($33\times$ below it) to $\sim 7.7\times$ (nearly $8\times$ above it).

2.1 Data-grounding block

- **Source:** Indivd Analytics estate forecaster, `location_visits` metric — the estate’s venue-visit forecast metric, counted only during each venue’s opening hours (“planned visitor-trips” below is that metric aggregated per market over the forecast week). Coherent-vintage production pull captured 2026-10-01 (all rows from the single forecast generation batch — the forecaster’s weekly run — of 1 October 2026); that frozen capture is the single source for every figure below.
- **Method:** estate total $\div$ locations per market = per-store planned trips; index vs estate mean (read alongside the companion methodology paper, *Methodology Paper §6* — the weather-compensated estate trend — on the Papers page: <https://analytics.indivd.com/papers/indivd-insights-methodology.pdf>). Figures follow the served forecast’s own store-hours gating (open hours only); every venue in the current vintage carries data-derived opening hours (201 of 201 venues), so the basis is uniformly gated to opening hours. Per-market intensity is reported on this single gated basis.
- **Date range:** forecast window 2026-10-01 $\rightarrow$ 2026-10-07 (headline), with the second forecast week 2026-10-08 $\rightarrow$ 2026-10-14 recorded in the artifact as a second-week agreement check (the served horizon extends to 10-15 22:00 UTC; the partial final day is disclosed and unused — see the figures note).
- **Figures (week 2026-10-01 $\rightarrow$ 10-07, current estate):** 4,008,802 planned visitor-trips across **201 locations**, 19,072 hourly data points; estate mean $\sim 19,944$ trips/store. Top-4 markets $\sim 89\%$ of planned trips — SE 38.8% and ES 23.0% are multi-store markets, while SK 16.0% and QA 11.3% are single venues, so this concentration headline mixes multi-store and single-venue markets. Per-store indices vs the estate mean: the six-store Iberian market $\sim 7.7\times$; the 49-store network $\sim 0.03\times$ (24.4% of stores, 0.73% of trips — the market’s planned trips are spread across its venues, none above $\sim 5\%$ of the market total); the 39-store Low-Countries network $\sim 0.09\times$; the 76-store Nordic anchor market $\sim 1.03\times$. The single-venue markets SK, QA and PL carry 16.0% , 11.3% and 7.7% of planned trips (one venue each) and read $\sim 32\times$, $\sim 23\times$ and $\sim 15\times$ — single-venue figures, not market-density estimates. Density *index* per market, not identifying store-level plans.
 - *Gating and point count:* the 19,072 hourly points reflect per-location open-hours gating — every venue contributes only its open hours, ~ 13.6 open hours per venue-day on average across the estate ($19,072 \text{ points} \div 1,407 \text{ venue-days} \sim 13.6$).
 - *Second-week agreement:* the second forecast week (10-08 $\rightarrow$ 10-14) agrees on rank order for every market above $\sim 10,000$ weekly trips (7 of 16 markets; the nine tail markets together carry under 1% of planned trips) — an agreement check between the two weeks of the same forecast run, not a stability claim; the smallest tail markets swap positions.
 - *Horizon edge:* the served horizon runs from 10-01 00:00 to late evening on 10-15 — 14 full days plus a partial final day — because the generation batch wrote its full horizon in one run; forecast rows for 10-01’s early hours were part of that same run, so week 1 has full hourly coverage. The partial final day is disclosed and unused.
- **Data-quality confessions (carried, not hidden):** (1) **coherent-vintage basis:** every figure comes from the single forecast generation batch of 2026-10-01 — the pipeline’s actual current vintage, with complete hourly coverage of both weeks and zero duplicate rows (confirmed by quality checks at capture). A naive pull filtered only on the calendar window would assemble leftover rows from superseded batches: at capture, the week-1 calendar window additionally held $\sim 3,050$ rows ($\sim 25.5\text{M}$ trips) belonging to venues that have since **left** the forecast estate — their last batch’s rows linger until those venues are removed from the forecast store and are excluded here (confirmed: none of those rows belong to current-estate venues); (2) **the estate changed between vintages:** the archived 20 September capture covered 237 forecasted locations including twelve stale flat-model venues ($\sim 1.9\text{M}$ forecast trips/week each, then $\sim 62\%$ of the estate total) and a five-venue Middle-East market; the current vintage forecasts **201** locations and none of those seventeen venues receive forecast rows any more — the stale-model and ungated-basis caveats of the previous version are retired with them, and the estate-size change ($237 \rightarrow 201$) is part of the reported basis. The 16 markets of the current vintage all existed in the archived capture’s forecast too, but the two vintages’ market orderings are not comparable: the archived capture’s second-largest market by

trips (the Low-Countries network, $\sim 25.7\%$ then) reads $\sim 1.8\%$ now — its total was dominated by venues that no longer receive forecasts; (3) **absolute forecast levels are not comparable across vintages**: the current-vintage week-1 total ($\sim 4.0\text{M}$) is materially lower than the archived September capture ($\sim 36.7\text{M}$) — forecast levels are planning numbers re-read each weekly run and this estate’s served forecast level changed between vintages. What is claimed here is the **relative cross-market pattern on the current vintage**; the current vintage is the only one evaluated in this piece, and one week of one vintage is a single point estimate, re-read at every forecaster run; (4) all figures are **forecasts (planned visitor-trips), not observed counts**; (5) an archived copy of the September capture exists for audit and is **not** republished as current.

- **Uncertainty**: prospective forecast; per-location confidence bands apply; the density indices are single-week point estimates, re-read each weekly run. The headline week’s estate total carries a band of $\sim 3.84\text{M}$ – 4.25M trips ($-4\%/+6\%$ around the point estimate; the forecaster’s per-venue confidence bands summed over the estate), and the market figures below sum each venue’s own confidence band. The headline conclusions — which markets index far above or far below the average — are unchanged across the bands: the six-store Iberian market reads ~ 7.2 – $8.1\times$ the per-store average band-inclusive; the 49-store network stays far below average across its much wider band (~ 0.002 – $0.07\times$ — so its headline “ $\sim 0.03\times$ ” is an order-of-magnitude statement, not a precise ratio); the single-venue markets’ bands are tight ($\leq 1.2\%$). The two low-index networks’ relative order is not band-robust — both are far below the estate mean.
- **Accuracy/trust**: the demand forecast behind these figures is a planning number with per-location confidence bands — not a verified outcome, and no forecast-accuracy claim is made here. The 1.17% KPMG verification covers the underlying **counting methodology for observed data only**; it says nothing about the forecast figures above.
- **Privacy**: everything here is estate- and market-aggregated, country-level only; no store, brand, region or customer data appears. Country codes plus per-market store counts are the intended level of disclosure; within-market distribution notes (such as a largest-venue share) are aggregate statistics that identify no individual venue. Venue-level figures — including any table of individual stores — are not published.

Originally published 2026-09-03; recomputed on the forecast’s own opening-hours basis 2026-09-20; recomputed again on the current coherent forecast vintage 2026-10-01 — the figures here supersede the earlier versions. Anonymization: estate- and market-aggregated, country-level only; no venue-level or store-level figures are published.