How this forecast was built

A short overview of the analysis behind the Paper 2 forecast — what was read, what was tested, what was found, and how far it can be trusted.

The corpus

Every IB Mathematics: Analysis & Approaches **SL** past paper available was read **question by question** through its mark scheme — not skimmed, mapped:

- **15 Paper 2 papers** and **14 Paper 1 papers**, spanning May 2022 to November 2025, across all time zones supplied.
- For each, every question was tagged with its topic(s), genre, command term, marks and section, and the paper was rolled up into a per-topic mark total.
- Paper identities were verified against the internal **IB file code** printed in each PDF — the filenames in the source folder were unreliable and had to be decoded and corrected.

What was done with it

1 · Structural analysis. Ten analysis passes built the picture of what Paper 2 *is*: the recurring Section A genres and how often each appears, the Section B "chain templates", the topic-mark distribution, and how Paper 2 differs from Paper 1 (calculator-driven, modelling-wrapped, Statistics- and Calculus-heavy).

2 · Five independent prediction models. Five different forecasting methods were run separately — within-sitting balancing, a time-zone-2-specific model, a recency/trend projection, a slot-by-slot conditional model, and a topic-rotation model — then compared for where they agreed and where they genuinely diverged.

3 · The pairing study. The intuitive idea that "a topic heavy on your Paper 1 will be light on your Paper 2" was tested directly: 14 matched same-sitting, same-time-zone paper pairs, with a per-topic correlation between Paper 1 and Paper 2 topic splits.

4 · A methodology audit — then a full re-audit. The first forecast was audited and found to over-reach; it was rebuilt base-rate-led. The pairing study was then re-audited from the ground up: every roll-up re-derived from question-level data so each paper sums to exactly 80 marks, under two explicit splitting conventions, and the correlation re-run and verified three independent ways (Pearson, Spearman, a 20,000-run permutation test).

5 · IB's own design philosophy. IB's official documentation was researched to check whether the IB itself claims to balance topic coverage across papers.

What was found

Finding	Status
Statistics and Calculus **co-dominate** Paper 2 (~23 marks each); Geometry & Trig ~16, Number & Algebra ~10, Functions ~8	High confidence — from clean, fully-reconciled data
A normal-distribution question (~85%) and a regression question (~73%) in Section A are near-certain	Calibrated from corpus frequency
The "Paper 1 predicts Paper 2" compensation effect	**Not statistically detectable** — every correlation's confidence interval straddles zero
IB's stated philosophy	Confirms it: IB balances coverage only loosely, *within* a paper — there is **no** published cross-paper rule
Who owns the hardest Section B question (Q9)	Genuinely unresolved — Statistics or Calculus

The re-audit also corrected a real error: the first version **under-counted Calculus** (the source roll-ups had dropped marks). Fixing it moved Calculus from a distant second to level with Statistics.

What it produced

- A **slot-by-slot forecast** of the most likely paper, with honest confidence levels.
- A **prioritised revision plan** — what to drill tonight, in order.
- A **practice-paper ranking** — all 15 past Paper 2s scored and tiered by ROI for this exam.
- A **60-page revision guide**, and the depth sections behind the forecast.

How far to trust it

This is a **well-grounded prior, not a leaked paper.** The corpus is 15 Paper 2s — enough to read the base rate reliably, not enough to pin down subtle effects (the pairing study's confidence intervals are wide; n is small and honestly flagged). The base-rate findings — Statistics and Calculus co-dominant, the near-certain Section A genres, the Section B templates — are solid and survive every robustness check. The specific slot-by-slot guesses are judgement on top of that. Its real job is telling you where to spend revision time, and that answer is robust.
