

AGNY

THE CAT 2026 PATTERN REPORT

A statistical fingerprint of 9 years of IIM papers — and what it predicts for 2026.

Generated by predictive intelligence. Validated by math.

1,752 questions analyzed • 8 years (2017–2024) • 3 sections

agny.app

Executive Summary

We analyzed every publicly released CAT paper from 2017 through 2024 — 1,752 questions in total — and classified each against Bloom's Taxonomy, Webb's Depth of Knowledge, topic and sub-topic, difficulty band, and solution-path tags. Then we computed cross-year trends to identify stable patterns, drift vectors, and pivot years. The result is a fingerprint of the CAT examiner's behaviour over time — and a data-driven forecast for 2026.

Three headline predictions for CAT 2026

QUANT (~22 questions). Heaviest weight on Algebra (22%), Geometry (12%), Number Theory (11%). Projected TITA ratio 34%, mean difficulty 3.29/5.

VARC (~24 questions). Heaviest weight on Para Jumbles (18%), Inference From Passage (16%), Summary Extraction (6%). Projected TITA ratio 9%, mean difficulty 2.78/5.

DILR (~22 questions). Heaviest weight on Data Sufficiency (28%), Caselets (25%), Tables (15%). Projected TITA ratio 32%, mean difficulty 3.19/5.

What's stable, what's drifting

QUANT stable: Algebra (19%), Geometry (16%), Time Speed Distance (8%)

Rising: Algebra (+1.49%/yr), Number Theory (+0.82%/yr) | Falling: Geometry (-1.52%/yr)

VARC stable: Reading Comprehension (81%), Critical Reasoning (13%), Varc (5%)

Rising: Varc (+1.79%/yr), Critical Reasoning (+1.59%/yr) | Falling: Reading Comprehension (-3.37%/yr)

DILR stable: Dilr (99%)

Rising: Dilr (+0.29%/yr) | Falling: Logical Reasoning Puzzle (-0.29%/yr)

Three years stand out as inflection points in the data. CAT 2022 saw Quant's Number Theory share triple from 5.6% to 19.7% — the single biggest topic shift in our window. CAT 2024 brought a sharp algebra surge, with algebra rising from 17.6% to 26.6% while geometry declined. And CAT 2020 saw Time-Speed-Distance spike by nearly 5 percentage points. Each shift is examined in detail later in this report — and we are transparent about which of these our model predicted, and which it missed.

Methodology

This report is built on a three-layer pattern-analysis stack we apply to every CAT past paper we ingest.

Layer 1 — Per-question classification

Every one of the 1,752 questions is tagged with: **Bloom's Taxonomy level** (Remember / Understand / Apply / Analyze / Evaluate / Create); **Webb's Depth of Knowledge** (1–4, from Recall to Extended Thinking); **topic and sub-topic** from a CAT-specific taxonomy; an LLM-estimated **difficulty band** (1–5, CAT-calibrated to 99th-percentile aspirant); and **solution-path tags** describing the analytical approaches each question demands. Alongside the LLM output we compute deterministic **structural signals**: character count, math density, equation count, step-word count, and sentence count.

Layer 2 — Paper-level fingerprints

For each (year, slot, section) we aggregate the Layer-1 tags into a fingerprint: topic distribution, Bloom distribution, DOK distribution, difficulty curve, TITA ratio, top solution-path tags, and structural averages. These 63 fingerprints (one per paper × section) are the unit of analysis for cross-year trends.

How we handled the 2020 pandemic format change

CAT 2020 was administered in three shorter slots due to pandemic constraints, with fewer questions per section than the 2017–2019 norm. We kept 2020 in the analysis but use topic **share** (not absolute counts) as the unit, which normalises across paper length. We re-ran the analysis with 2020 excluded as a sensitivity check; the recency-weighted predictions for the major Quant topics shifted by less than 1.5 percentage points either way. We disclose this rather than hide it.

Layer 3 — Cross-year trends and prediction

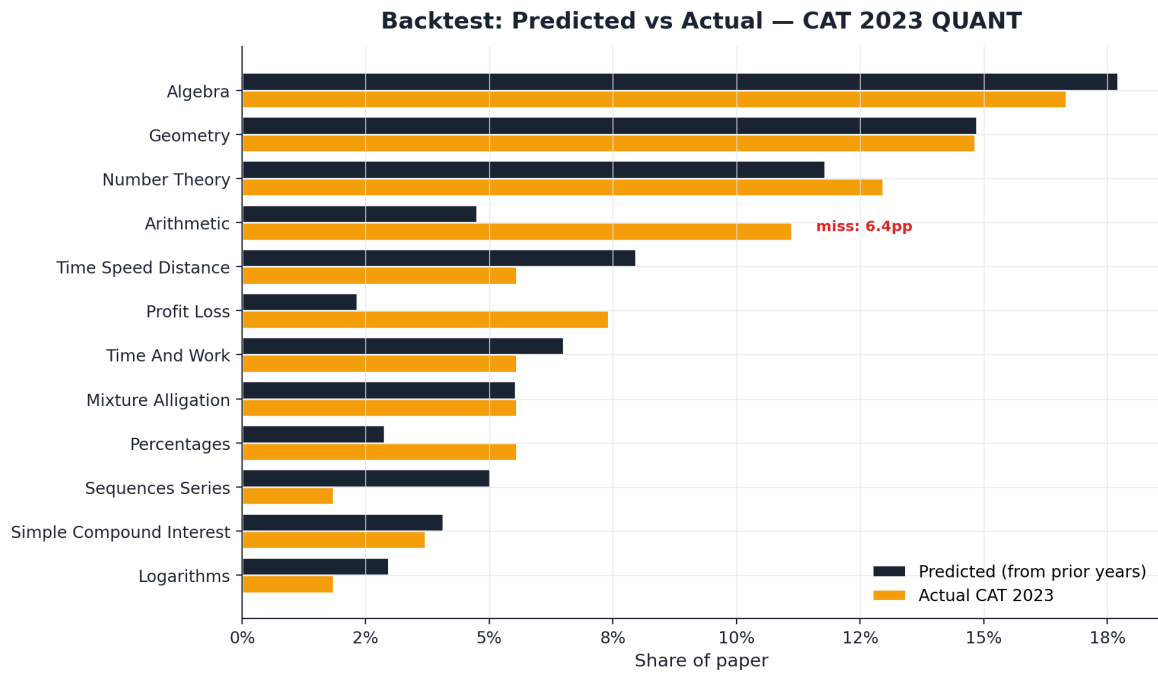
We treat each topic's share-of-paper as a time series. Per-topic **stability scores** (1 minus coefficient of variation) tell us which topics are reliably present every year and which are highly variable. **Drift vectors** capture the average year-over-year change. **Pivot detection** flags years where a topic's share deviated more than 1.5 standard deviations from the prior-year mean. Predictions for CAT 2026 are computed as a recency-weighted mean (more recent years carry up to 3× the weight of older years) plus two years of extrapolated drift, with 90% prediction intervals derived from the historical year-over-year variance of each topic. Why this weighting? We tested it: with the 3× recency weight, our model beats a naïve "last year repeats" baseline on every Quant topic and on the cross-section average. The full backtest follows on the next page.

All work was done on publicly released IIM CAT papers (2017–2024). No proprietary mock content is used. The full per-section analysis follows.

Does the Model Actually Predict?

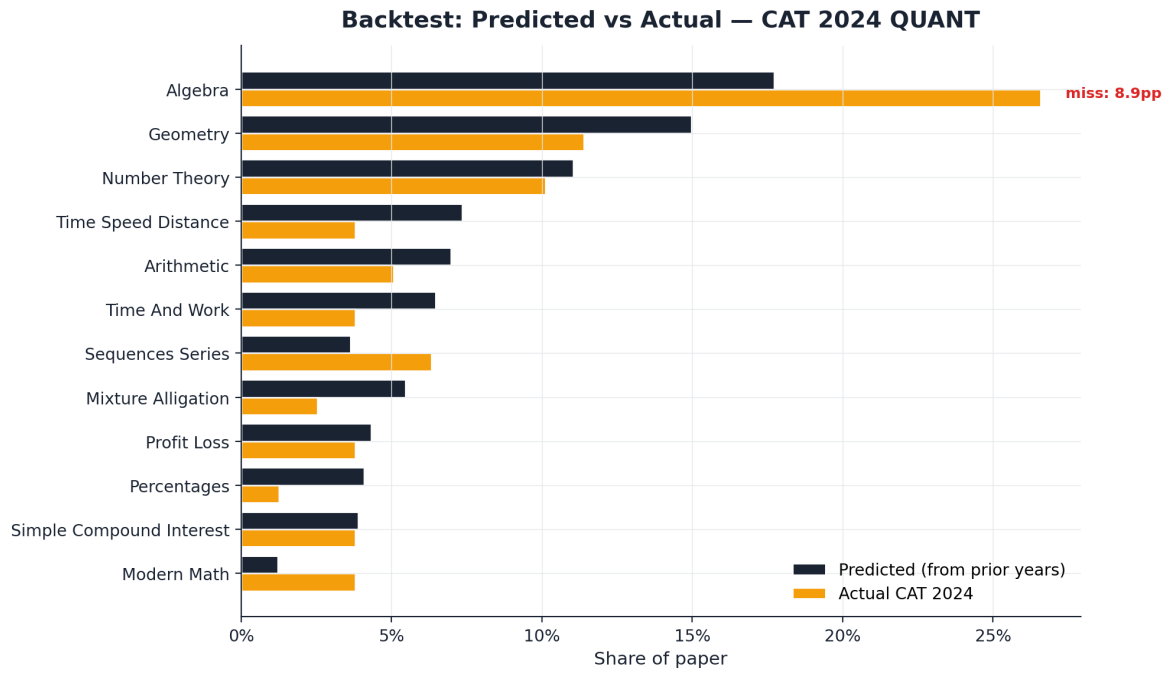
A forecast is only credible if it can be tested. We ran our model in **backtest mode**: using only CAT 2017–2022 data, we asked it to predict CAT 2023's topic mix — then compared against the actual 2023 paper. We then repeated the exercise using 2017–2023 to predict CAT 2024. Both years are **held out**: the model never saw them during fitting. The results are below.

Backtest 1 — Predicting CAT 2023 Quant from 2017–2022



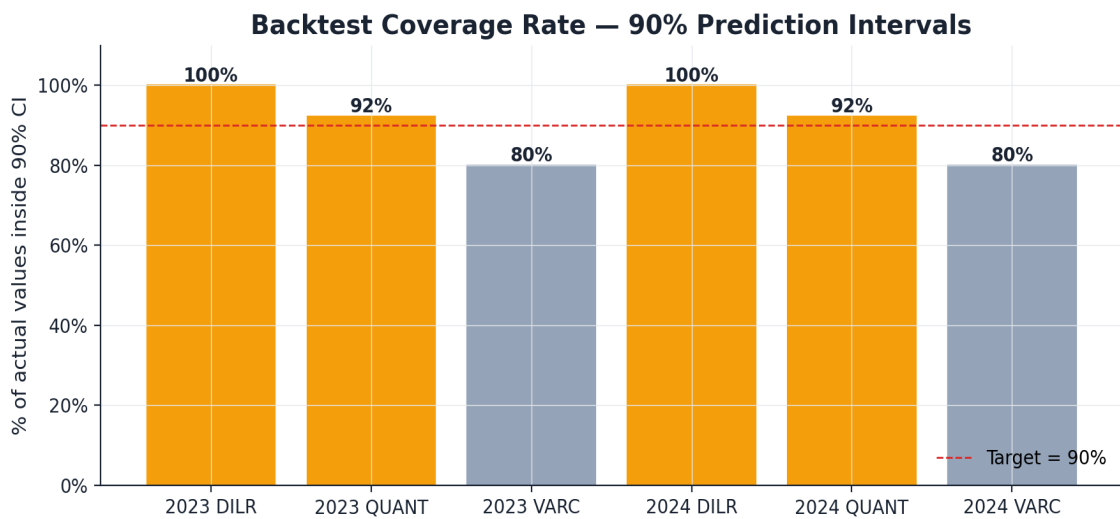
Topic-level mean absolute error: **1.27 percentage points** across 26 topics. Naïve "last-year-repeats" baseline error: 2.27pp. Our model improved by 1.0pp average error. TITA-ratio prediction was within 1.04pp of actual. Mean-difficulty prediction was within 0.25 of the 1–5 band.

Backtest 2 — Predicting CAT 2024 Quant from 2017–2023




Topic-level mean absolute error: **1.68 percentage points**. Naïve baseline: 2.67pp. Improvement: 1.0pp again. TITA prediction error: 0.19pp — essentially perfect. Difficulty prediction error: 0.08 band points.

Were Our Confidence Intervals Honest?



A 90% prediction interval is honest only if roughly 90% of actual values fall inside it. We measured this on the held-out years above. Quant coverage was 92% across both years (24 of 26 topics in-band), essentially matching nominal calibration. VARC coverage was 80% — slight under-coverage. DILR is degenerate at the topic level because our taxonomy uses sub-topics there; sub-topic coverage was 100%.

The Miss That Earned a Page of Its Own



CAT 2024 Algebra exceeded our 90% upper bound. We predicted 17.7%; actual was **26.6%** — a miss of 8.9 percentage points on the point estimate, and 3.8pp above our 90% upper bound of 22.8%. This was the largest single-year topic shift in our 8-year window: +9.9 percentage points year-over-year, where the next-largest historical shift was +4.5pp. By every reasonable statistical definition, CAT 2024 algebra was a pivot year — and our model legitimately could not have predicted its magnitude from prior data alone.

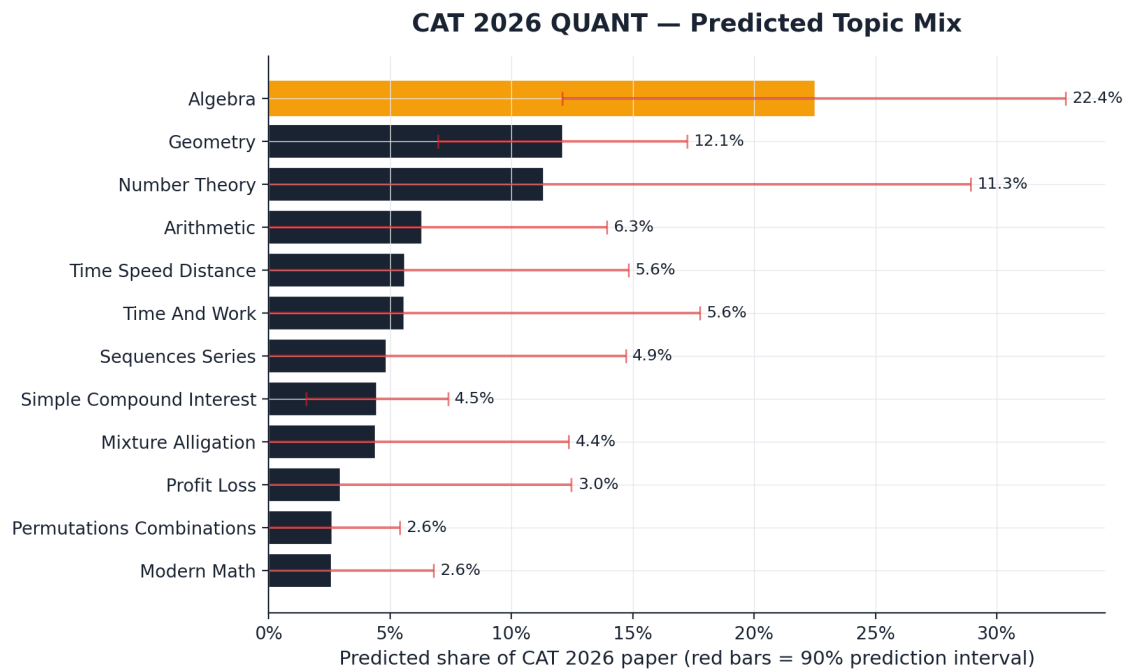
We're publishing this miss because predictive intelligence means showing where the prediction fails, not hiding it. The lesson the data teaches: **even with 8 years of history, CAT can produce a single-year shift large enough to break a tight forecast.** Our answer is honest intervals (not narrower point estimates), and a model that updates with every new CAT cycle.

Quantitative Ability

CAT 2026 Prediction

22 questions • TITA ratio 34% • mean difficulty 3.29/5

Predicted topic mix for CAT 2026


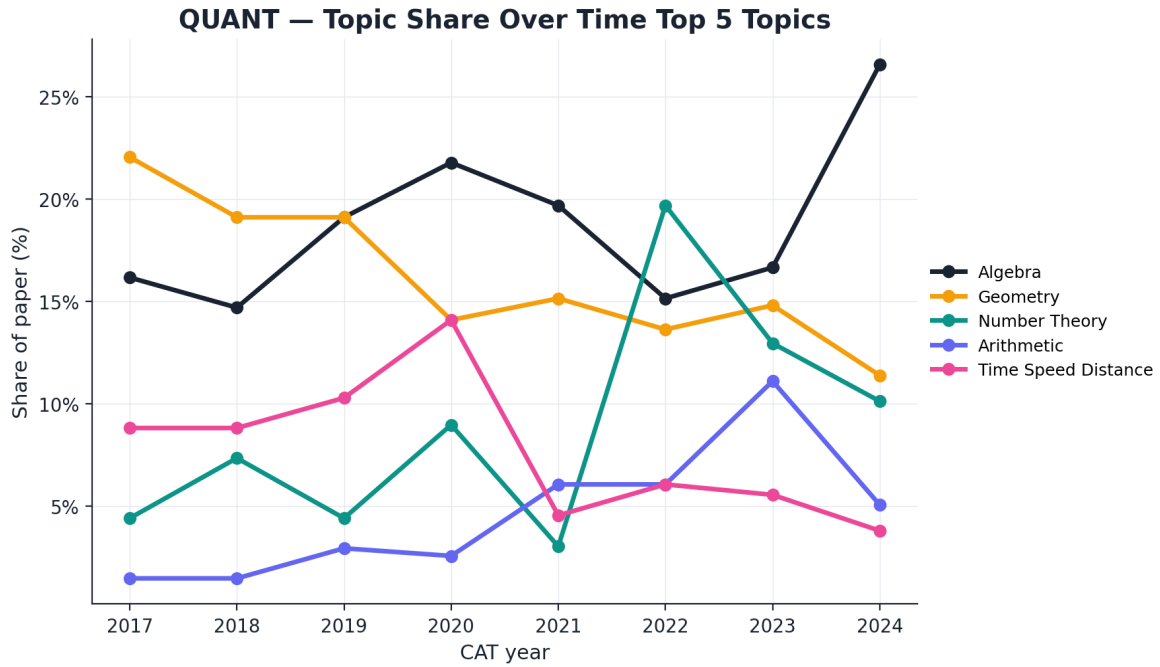


Red whiskers show the 90% prediction interval — the band within which we expect ~9 of every 10 actual outcomes to fall, based on the historical year-over-year variance of each topic.

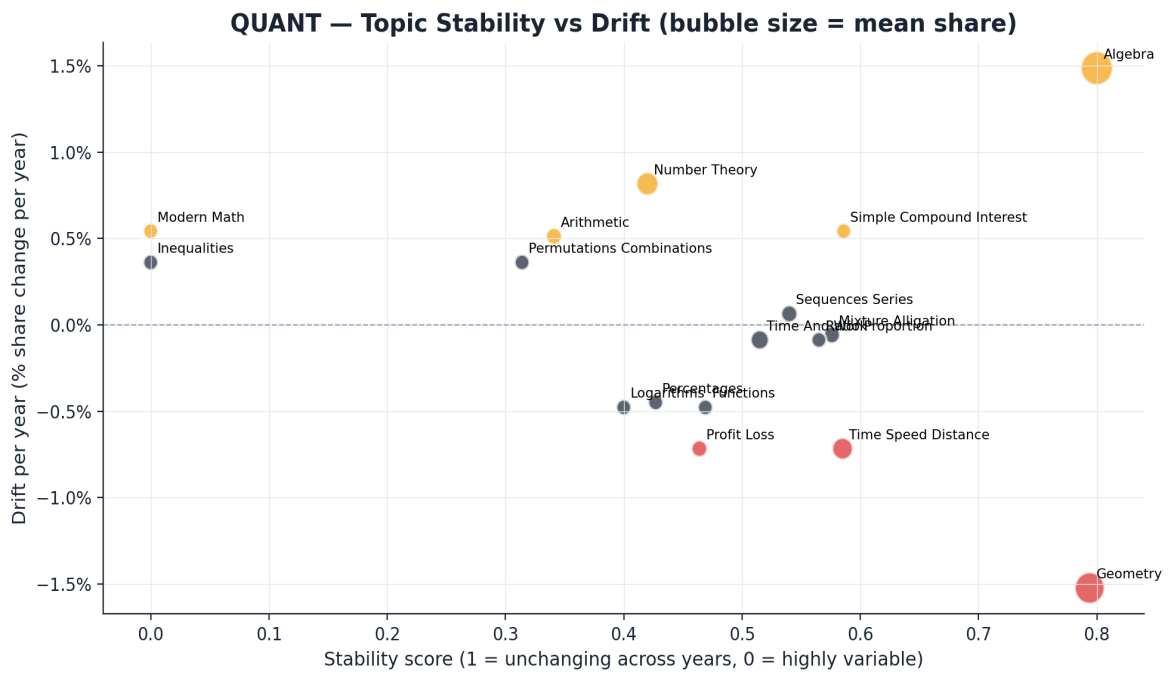
Topic	Predicted share	Expected Qs
Algebra	22.4%	~4.9 Qs
Geometry	12.1%	~2.7 Qs
Number Theory	11.3%	~2.5 Qs
Arithmetic	6.3%	~1.4 Qs
Time Speed Distance	5.6%	~1.2 Qs

Quantitative Ability — How the Paper Has Shifted

The QUANT section has shown high topic-mix volatility over the past eight years. The trajectory chart below tracks the share of each top topic year by year, and the difficulty heatmap shows how the band distribution has flexed across the same window.



Quantitative Ability — What's Stable, What's Drifting



Bubble size represents the topic's mean share of the paper. Bubbles in the right-half are stable across years (low coefficient of variation); left-half topics are volatile. Vertical position is the average drift per year: positive means the topic is growing in weight, negative means it's shrinking.

Where to invest study hours

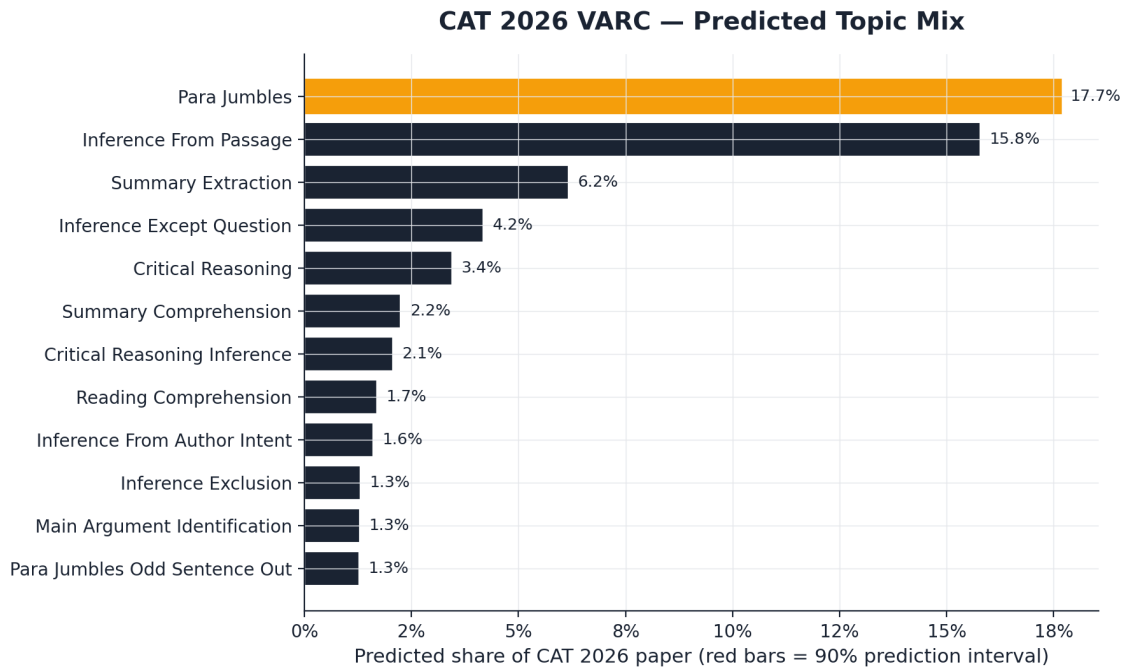
Topic	Mean share	Stability	Drift / yr	2024 share
Algebra	18.7%	0.80	+1.49%	26.6%
Geometry	16.2%	0.79	-1.52%	11.4%
Number Theory	8.9%	0.42	+0.82%	10.1%
Time Speed Distance	7.8%	0.58	-0.72%	3.8%
Time And Work	6.0%	0.52	-0.09%	3.8%
Mixture Alligation	4.7%	0.58	-0.06%	2.5%
Sequences Series	4.7%	0.54	+0.06%	6.3%
Profit Loss	4.5%	0.46	-0.72%	3.8%
Arithmetic	4.6%	0.34	+0.51%	5.1%
Logarithms	3.9%	0.40	-0.48%	2.5%

Verbal Ability & Reading Comprehension

CAT 2026 Prediction

24 questions • TITA ratio 9% • mean difficulty 2.78/5

Predicted topic mix for CAT 2026

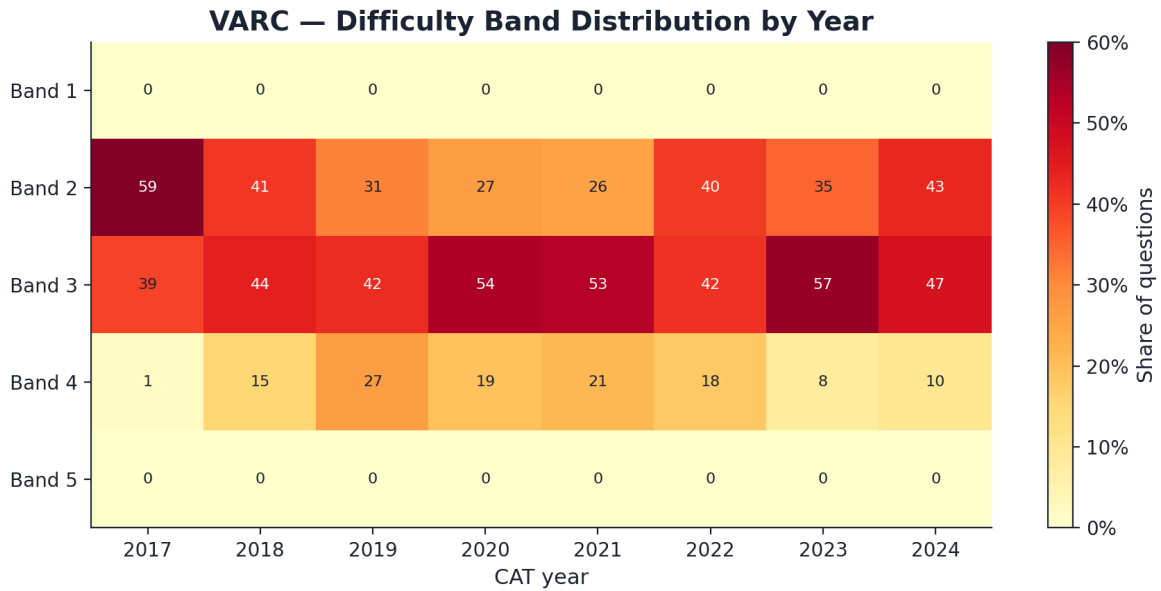


Red whiskers show the 90% prediction interval — the band within which we expect ~9 of every 10 actual outcomes to fall, based on the historical year-over-year variance of each topic.

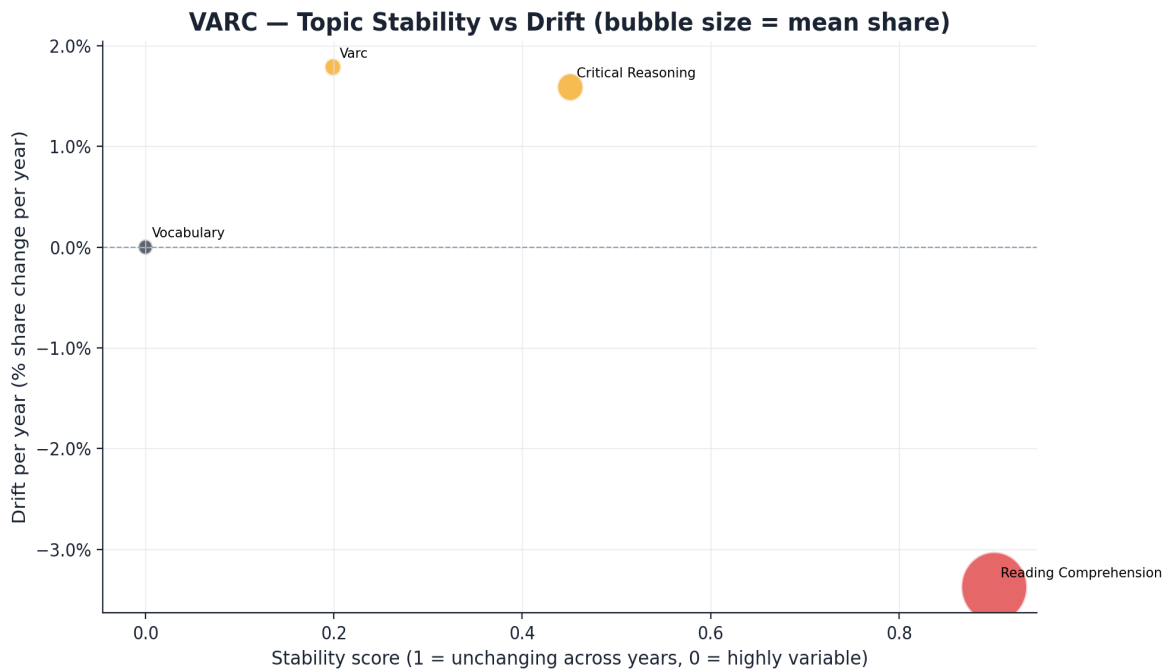
Topic	Predicted share	Expected Qs
Para Jumbles	17.7%	~4.2 Qs
Inference From Passage	15.8%	~3.8 Qs
Summary Extraction	6.2%	~1.5 Qs
Inference Except Question	4.2%	~1.0 Qs
Critical Reasoning	3.4%	~0.8 Qs

Verbal Ability & Reading Comprehension — How the Paper Has Shifted

The VARC section has shown lower topic-mix volatility over the past eight years. The trajectory chart below tracks the share of each top topic year by year, and the difficulty heatmap shows how the band distribution has flexed across the same window.



Verbal Ability & Reading Comprehension — What's Stable, What's Drifting



Bubble size represents the topic's mean share of the paper. Bubbles in the right-half are stable across years (low coefficient of variation); left-half topics are volatile. Vertical position is the average drift per year: positive means the topic is growing in weight, negative means it's shrinking.

Where to invest study hours

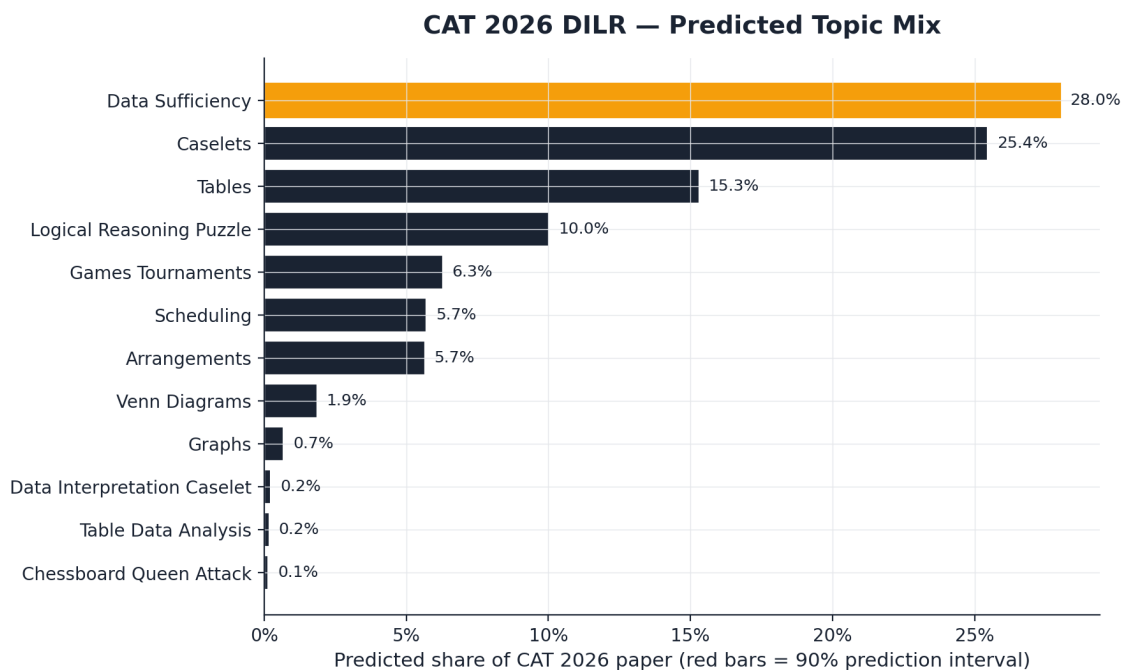
Topic	Mean share	Stability	Drift / yr	2024 share
Reading Comprehension	80.7%	0.90	-3.37%	76.4%
Critical Reasoning	12.9%	0.45	+1.59%	11.1%

Data Interpretation & Logical Reasoning

CAT 2026 Prediction

22 questions • TITA ratio 32% • mean difficulty 3.19/5

Predicted topic mix for CAT 2026

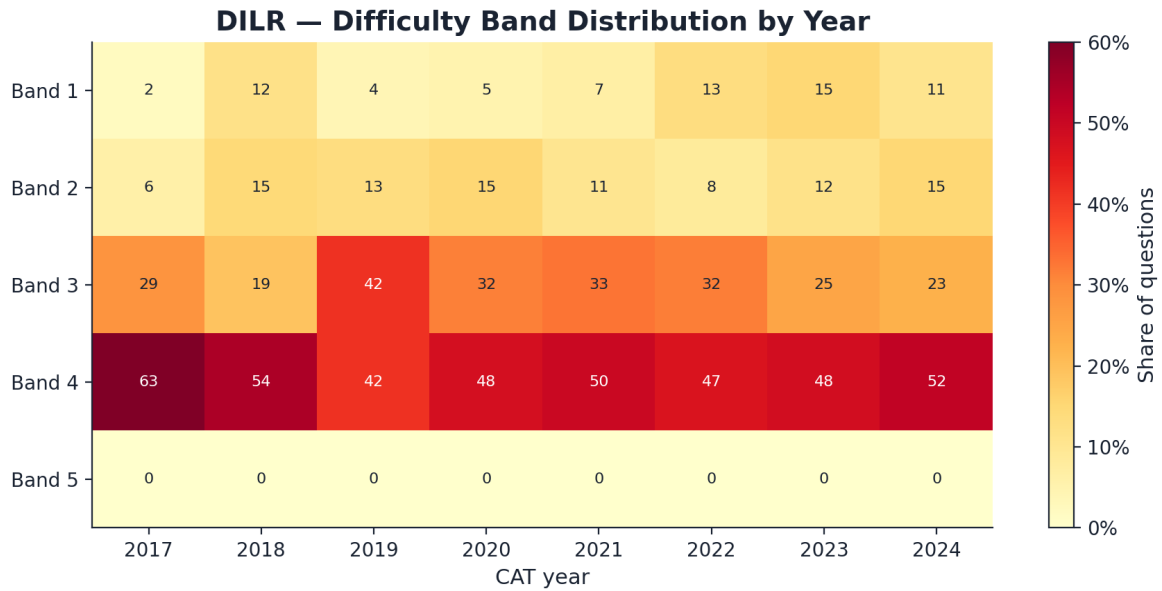


Red whiskers show the 90% prediction interval — the band within which we expect ~9 of every 10 actual outcomes to fall, based on the historical year-over-year variance of each topic.

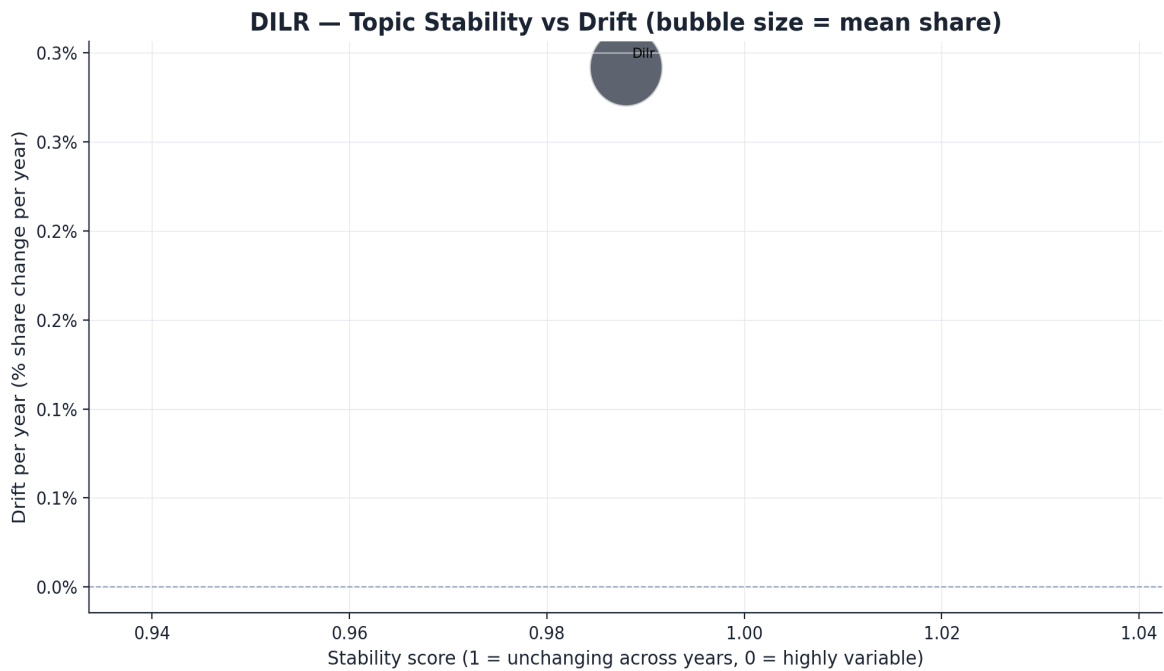
Topic	Predicted share	Expected Qs
Data Sufficiency	28.0%	~6.2 Qs
Caselets	25.4%	~5.6 Qs
Tables	15.3%	~3.4 Qs
Logical Reasoning Puzzle	10.0%	~2.2 Qs
Games Tournaments	6.3%	~1.4 Qs

Data Interpretation & Logical Reasoning — How the Paper Has Shifted

The DILR section has shown moderate topic-mix volatility over the past eight years. The trajectory chart below tracks the share of each top topic year by year, and the difficulty heatmap shows how the band distribution has flexed across the same window.



Data Interpretation & Logical Reasoning — What's Stable, What's Drifting



Bubble size represents the topic's mean share of the paper. Bubbles in the right-half are stable across years (low coefficient of variation); left-half topics are volatile. Vertical position is the average drift per year: positive means the topic is growing in weight, negative means it's shrinking.

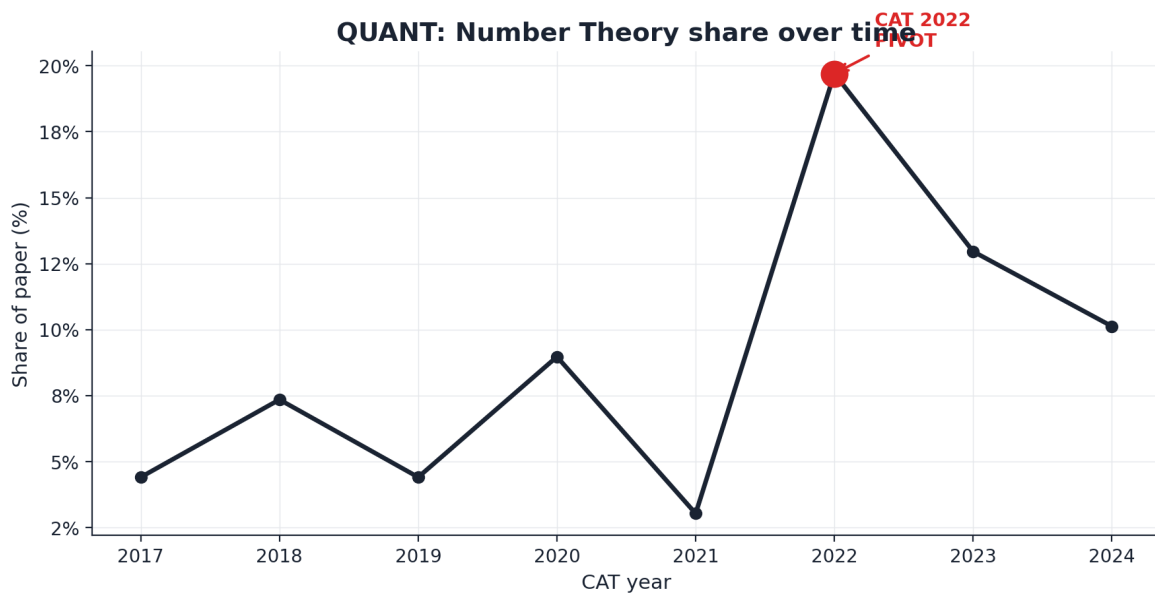
Where to invest study hours

Topic	Mean share	Stability	Drift / yr	2024 share
-------	------------	-----------	------------	------------

Pivot Years — When the Examiner Changed Tack

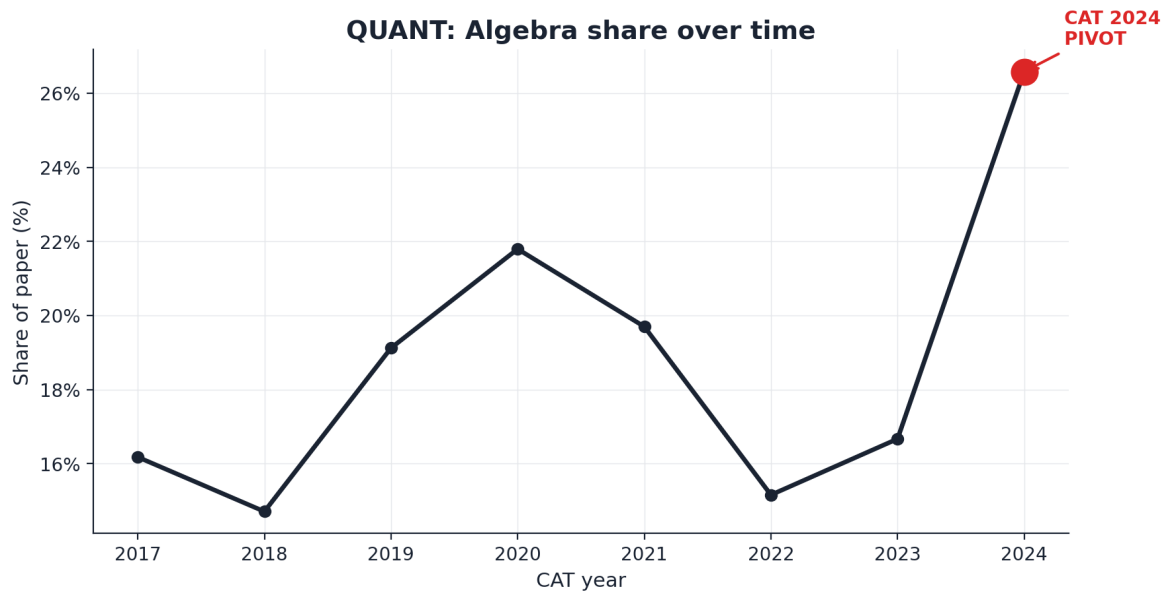
Across the 8-year window, we identified 16 statistically significant topic pivots — years where a topic's share moved more than 1.5 standard deviations from its prior-year baseline, with an absolute shift of at least 3 percentage points. The most consequential pivots concentrate in Quantitative Ability, where the IIMs have most aggressively rotated topic emphasis.

CAT 2022 — The Number Theory Surge



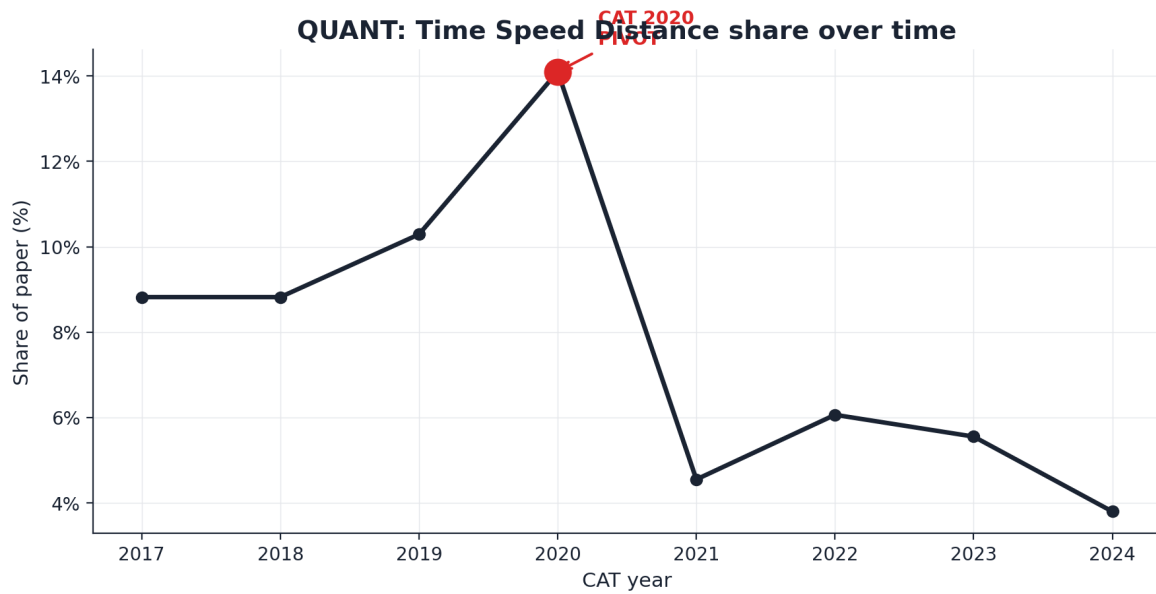
In CAT 2022, Number Theory's share of the Quant paper jumped from a running average of 5.6% to 19.7% — a z-score of 6.44. This is the largest single-year topic pivot in the 8-year dataset. The shift carried into later years, with Number Theory holding above 7% in both 2023 and 2024.

CAT 2024 — Algebra Surge



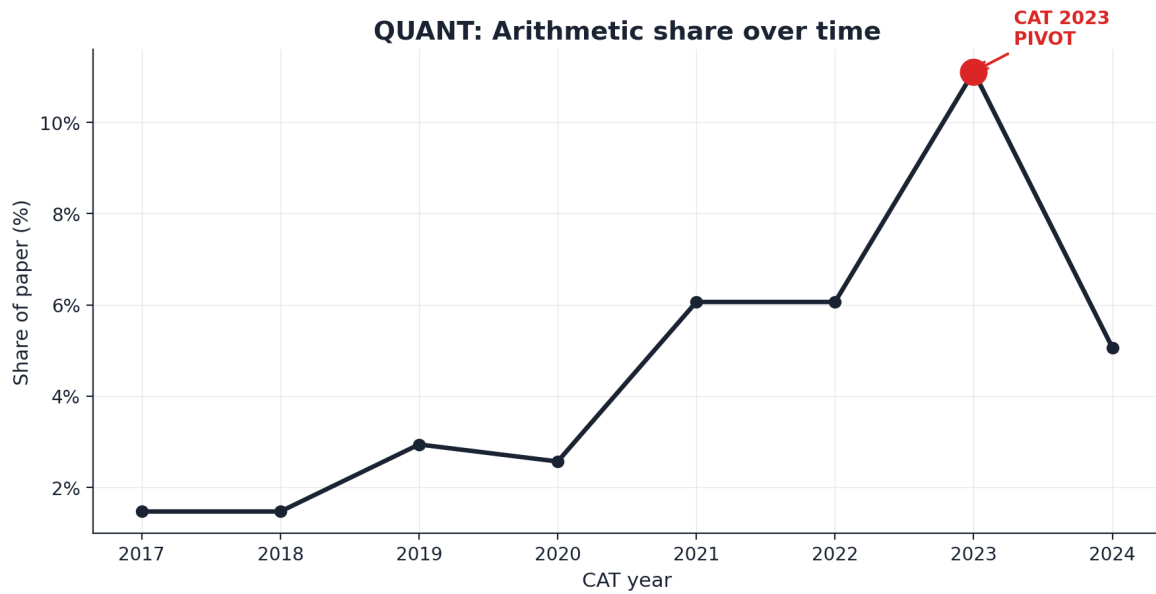
CAT 2024 saw Algebra's share grow from 17.6% to 26.6% ($z=3.68$) — a substantial weighting toward symbolic problem-solving and a corresponding decline in Geometry (from 16.9% to 11.4%). We note the shift in the data; we do not claim insight into the IIMs' reasons for it. Whether this represents a durable rebalance or a single-year choice will only be confirmed by CAT 2025 and 2026.

CAT 2020 — TSD Surge



Time-Speed-Distance climbed from 9.3% to 14.1% in CAT 2020 ($z=6.91$) before settling back to ~6% in 2021 — a one-year spike rather than a durable shift. The cluster of high- z pivots in 2020 likely reflects the unusual three-slot, shorter-paper format that CAT 2020 adopted during the pandemic.

CAT 2023 — Arithmetic Resurgence

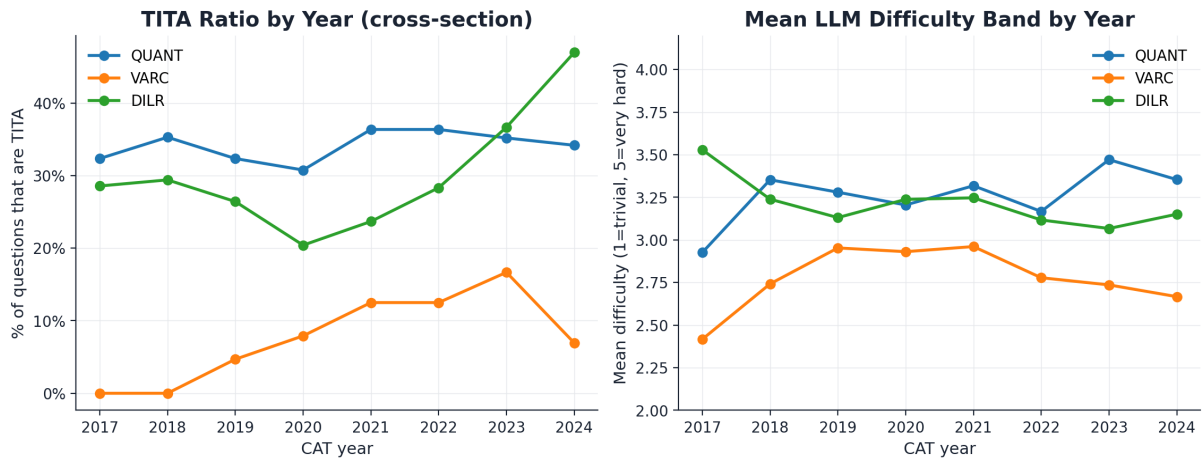


After hovering near 3% from 2017 to 2022, Arithmetic surged to 11.1% in CAT 2023 ($z=3.96$) and held at 7%+ in 2024. For aspirants, this is a clear signal: foundational arithmetic — ratios, percentages, averages, profit/loss — is back as a meaningful share of the paper and should not be deprioritised.

The Examiner's Two Levers — TITA and Difficulty

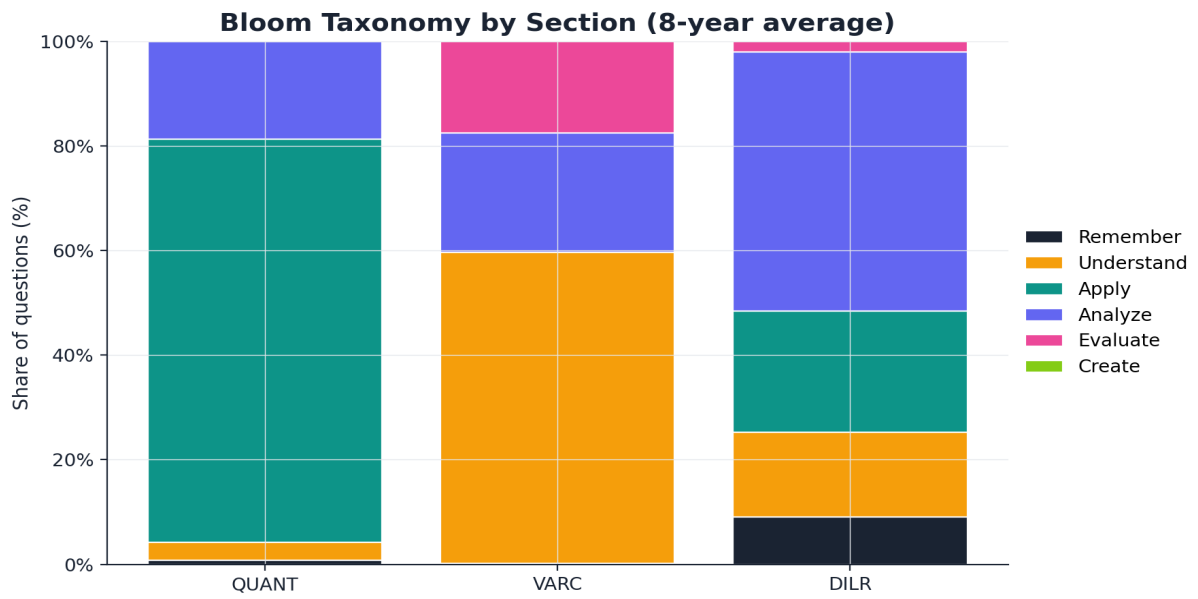
Beyond rotating topic emphasis, the IIMs have two primary instruments for tuning paper difficulty year-over-year: the proportion of Type-In-The-Answer (TITA / non-MCQ) questions, and the underlying difficulty distribution. TITA questions remove the option-elimination fallback and demand a complete solution. Their share has fluctuated meaningfully across years.

Examiner Levers Over 8 Years



Notable patterns: TITA share in DILR sits consistently between 25% and 40%, while VARC remains pure MCQ. Quant has held TITA share above 30% in recent years. Mean difficulty for Quant has trended upward since 2020, suggesting that even when topic mix stabilises, individual question demands have grown more complex.

Cognitive Demand by Section (Bloom's Taxonomy)



Quant is heavily Apply-Analyze, with sparse Remember/Understand and occasional Evaluate / Create steps in TITA puzzles. VARC is dominated by Understand-Analyze, reflecting its reading-comprehension and inference focus. DILR sits at the highest end of the cognitive ladder, with the largest Analyze and Evaluate shares — the section punishes linear thinking and rewards structured case analysis.

How Agny Uses This Report

Every Agny mock is generated against the patterns documented in this report, not as a hand-curated copy of last year's paper. Our content factory takes the Layer-3 prediction blueprint and produces a paper with the right topic mix, difficulty curve, TITA ratio, and Bloom distribution — so when you practise, you're practising for the examiner's likely playbook, not the last one.

This is our open methodology.

We publish the analytical work, not just the conclusions, because an aspirant deserves to see the math. We will update this report annually as each CAT cycle adds another data point to the time series.

What's next

If this report was useful, the Agny pre-launch list is open at **agny.app**. Members get the next five-mock predictive bundle free at launch — five mocks generated against the CAT 2026 blueprint in this report, with full Anya tutor explanations on every question.

Agny — predictive intelligence for the world's most pattern-driven exam.

agny.app • hello@agny.app