Final Recommended CAT Predicted Papers

Disclaimer: Pattern-based prediction only. This is not a leaked or official paper.

Methodology

PYQ analysis is loaded into SQLite, converted into prediction specs, used to create internal candidate prompts/imports, verified, scored for CAT-likeness, filtered for weakness/duplicates, assembled into variants, backtested, and weight-adjusted.

Selected Papers

Variant	Type	Expected Overlap	Risk	Diversity	Reason
Arithmetic-Heavy / QA Expected Paper	arithmetic_heavy	0.865	0.35	0.636	Selected as one of the best-scoring non-wildcard variants.
Reasoning-Heavy Paper	reasoning_heavy	0.858	0.35	0.727	Selected as one of the best-scoring non-wildcard variants.
Balanced High-Probability Paper	balanced	0.849	0.35	0.864	Selected as one of the best-scoring non-wildcard variants.
Wildcard Paper	wildcard	0.817	0.65	0.636	Included as a volatility hedge.

Candidate Content Status

- Selected candidate pool: 65

- Selected paper-eligible candidates: 37
- Selected LaTeX candidates: 26
- Specs represented: 19
- Final assembly readiness override: ALLOWED WITH RISK FLAGS.

Coverage Risk

CAT_QA_SPEC_14 Circles remains under-covered; final paper variants should not force a weak circle candidate unless required by portfolio diversity.

No rejected Circles candidate was promoted during assembly.

Limitations

- Exact CAT question prediction is not realistic.
- VARC/DILR extraction and tags are rule-based.
- Candidate quality depends on manual LLM/local-model generation and verification.
- Final output should stay limited and high-precision, not become a bulk question bank.