

Ten Theses After the 20260312 Response Cluster

Claude and MJC

March 12, 2026

1. The current MCP tail is stronger after normalization than before it.
2. The main recent empirical pressure lands on the combination rule, not on count structure as such.
3. If the claimed ablation is right, source selection is not enough to close the gap to softer combination rules.
4. If entropy-weighted blending really beats KN, then the interesting question is no longer “whether interpolation wins” but what structural principle confidence-weighted combination is exploiting.
5. The current MCP should be read as the exact theory of tropical max-min, not automatically as the final theory of count-based prediction.
6. A broader post-tropical theory is a legitimate target if it preserves count structure, quotient structure, and compression identities while relaxing the rigid combination rule.
7. Track U is only intellectually serious if it admits the tropical tax honestly rather than pretending max-min is already competitive.
8. Track C is only intellectually serious if its wins are reported with exact decoding, reproducible accounting, and explicit statements about what theoretical inheritance is being relinquished.
9. The project should alternate between one reviewed MCP unit and one verified benchmark unit rather than allowing either theory growth or benchmark rhetoric to dominate unchecked.
10. The right long-term ambition is not permanent bifurcation but a cleaner theory boundary: what belongs to tropical exactness, what belongs to count-based prediction in general, and what new combination theory is needed beyond both.