The Price of the Tropical Semiring

Claude and MJC

March 12, 2026

Abstract

Recent ablation claims, if they hold up, sharpen the status of the Universal Model rather than merely refuting it. A gap of roughly 5.347 bpc to 2.418 bpc between rigid flat max-min and interpolation means that source selection alone is not the main bottleneck; rigid max-min itself is expensive. But the newer gap-weighted tropical blend result (2.244 bpc at 1M, 2.559 bpc at 100K) also shows that the semiring is richer than one flat aggregation law, and that the tropical confidence signal is sample-efficient enough to beat KN even before entropy-weighted blending does. A still newer normalized-conviction hybrid (H3) now appears to improve further at higher orders, which suggests the frontier has already moved beyond gap alone versus entropy alone. But this does not imply that the MCP is obsolete. It implies that the MCP is the theory of a specific constitutive choice—the tropical max-min rule—whose algebraic clarity comes at a measurable compression cost. The right response is not to dissolve everything into "Track C" pragmatism, but to ask what survives, what fails, and what generalizes if a softer confidence-weighted combination rule wins empirically.

1 What The New Numbers Would Mean

Suppose the reported ablation is approximately:

Policy	bpc
max-min	5.347
sharpest-LPP	5.100
KN-interp	2.418
entropy-weighted blend	2.337

Proposition 1. *On these numbers, source selection is not the main bottleneck. Sharpest-LPP improves max-min, but does not recover more than a small fraction of the full gap to interpolation. But gap-weighted tropical blending also shows that most of the old "price of tropical was really the price of using one rigid competition law.*

That matters because it changes the burden of proof. The old hope was that richer competition over fixed tropical scoring might close most of the gap. If the reported result is real, that hope is false. The forward rule is doing real compression damage.

2 What This Does Not Mean

Remark 1. *It does not follow that the UM or the MCP were mistakes. It follows that rigid flat max-min is expensive.*

The MCP remains the theory of the object defined by its axioms. If grafs 30 through 37 derive from count structure plus the tropical forward rule, then they still describe that object correctly even if a different scoring rule compresses better.

So the new result would not say "the algebra was wrong." It would say:

> the algebra is exact for a system whose compression objective is not globally optimal.

3 Why The Two-Track Split Is Still Too Neat

The current Track U / Track C framing is useful but incomplete.

Proposition 2. *If entropy-weighted blending beats KN interpolation, then the problem is no longer "algebra versus engineering." It becomes: what structural principle is that confidence-weighted rule exploiting?*

That is a mathematical question, not just an engineering one. A confidence-weighted combination rule may still admit:

- a higher-order event-space interpretation,
- a partition-function interpretation,
- or a soft quotient over competing predictor spaces.

If so, the right conclusion is not "Track C wins, stop thinking." The right conclusion is that the current MCP may be the rigid limit of a larger theory.

4 What Changes In The MCP

The recent MCP tail should now be read more sharply:

1. some grafs are about count structure in general;
2. some are about the specific tropical max-min rule;
3. some may survive under softer confidence-weighted combination.

Remark 2. *This is clarification, not collapse. The MCP becomes more honest when it states which theorems are constitutive of max-min and which belong to a broader counting-based prediction theory.*

In particular, the recent tail suggests a clean distinction:

- grafs 30, 35, and 36 are most tightly tied to the tropical rule;
- grafs 31 and 34 have a broader life as compression/counting statements;
- graf 33 is epistemic and may generalize beyond the current rule.

5 The Real Frontier

The strongest new empirical claim is no longer merely that interpolation wins. It is that two different confidence-weighted rules improve on KN in different ways: entropy-weighting eventually beats KN, gap-weighting appears to beat KN earlier and stay better, and a normalized-conviction hybrid now appears to beat both at higher orders.

If that survives scrutiny, then the real frontier is:

> characterize confidence-weighted combination sharply enough that we can say exactly which MCP claims fail, which survive, and which generalize.

That is better than the old A/B debate. It is a content question.

6 Conclusion

The price of the tropical semiring is not a fixed 2.9-bit tax. The newer results suggest that most of that tax belongs to rigid flat max-min, not to tropical confidence-weighted combination as such. If the gap-weighted blend result survives, then max-min is a principled but expensive corner of a richer tropical family.

The intellectually serious response is neither to retreat into pure editorial theory nor to dissolve everything into compressor pragmatism. It is to use the new empirical pressure to separate:

1. what is true of counting-based prediction in general,
2. what is true only of tropical max-min,
3. and what a broader post-tropical MCP would need to explain.

References

[1] Claude and MJC. *MCP.* Working manuscript, March 2026.

[2] Claude and MJC. *The Combination Problem: Three Layers, One Benchmark.* Hutter archive, 12 March 2026.

[3] Universal Model Project. *Normalized Conviction: The* $^{g/H}$ *Rule.* $Hutterarchive, 12March2026$.

[4] Graf. *Conviction Depth: Fewer Wins, Bigger Wins.* Hutter archive, 12 March 2026.

[5] Graf. *The Conviction–Accuracy Tradeoff: Why Support-Gap Blending Beats Entropy Weighting.* Hutter archive, 12 March 2026.

[6] Claude and MJC. *Reviewing the MCP Tail: Grafs 30 Through 37 After Normalization.* Hutter archive, 12 March 2026.