

# Toward a Post-Tropical MCP

Claude and MJC

March 12, 2026

## Abstract

If entropy-weighted combination survives empirical scrutiny, then the current MCP cannot be the final theory of count-based prediction. That does not make the MCP useless. It means the present MCP should be read as the exact theory of a rigid constitutive choice: tropical max-min combination. A broader post-tropical MCP would need to separate what is true of counting-based prediction in general from what is true only of max-min, and then explain why confidence-weighted combination improves compression without discarding the count-structural insights that made the MCP valuable in the first place.

## 1 The Pressure

The recent empirical pressure is simple:

1. max-min appears much worse than interpolation on the same count tables;
2. source selection improves max-min but does not close most of the gap;
3. entropy-weighted blending may outperform plain KN interpolation;
4. gap-weighted tropical blending may outperform both while staying native to support-space competition.

If that picture is right, then the main question is no longer whether the tropical rule is elegant. The main question is what part of the MCP is really about count structure, and what part is only about the tropical choice of combination operator.

## 2 Three Layers Of The MCP

A useful way to restate the theory is:

1. **Count layer.** Counts, marginals, conditionals, quotienting, factorization, and code-length identities.
2. **Combination layer.** The rule that combines multiple pieces of evidence into one predictive distribution.
3. **Architecture layer.** Event spaces, factor towers, hidden bottlenecks, and recurrent transport.

**Proposition 1.** *Most of the current empirical challenge lands on the combination layer, not on the count layer.*

That is why the right response is not to throw away the MCP, but to refactor it.

### 3 What Likely Survives

The following claims plausibly survive into a broader post-tropical MCP:

- counts induce conditionals;
- code length is expected negative log-probability;
- quotienting splits coarse and residual information;
- coarsening loses mutual information;
- hidden state and recurrent transport can still be analyzed as bottlenecks and path sums.

These are not inherently tropical claims. They are structural claims about count-based prediction and compression.

### 4 What Likely Becomes Conditional

The following claims should probably be restated as conditional on the specific max-min rule:

- lattice/residuation derivations of the forward pass;
- the exact tropical-integer bridge in its current form;
- monotonicity and fixed-point claims tied to the rigid operator;
- any theorem whose proof uses max-min as a constitutive axiom rather than as one admissible rule among others.

**Remark 1.** *This is not failure. It is a cleaner statement of scope.*

### 5 What The New Theory Must Explain

If entropy-weighted combination really wins, a broader theory must explain at least three things:

1. **Confidence as a structural quantity.** Why should entropy or sharpness determine evidence weight?
2. **Soft competition among sources.** Why is a distributed weighting over predictors better than either winner-take-all selection or fixed recursive discount?
3. **Compatibility with count semantics.** How can the new combination rule still be understood as operating over count-derived evidence rather than as an unrelated numerical hack?

## 6 A Minimal Program

The smallest intellectually serious next program is:

1. keep the current MCP as the exact max-min theory;
2. mark clearly which grafes are constitutive of that theory;
3. develop a second layer of statements that hold for any count-derived predictive operator;
4. characterize entropy-weighted combination inside that broader layer.

**Proposition 2.** *If this program succeeds, the project does not split into “pure UM” versus “mere compressor.” It becomes a nested theory: the tropical MCP as one exact corner of a broader count-combination framework.*

## 7 Formal Boundary

Gap-blend sharpens the distinction between *tropical-native ingredients* and *the old rigid max-min object*. The following steps remain tropical-native:

1. extract per-order support tables from counts;
2. compute the top-two support gap  $g_k = s_{1,k} - s_{2,k}$ ;
3. convert that gap into a power-of-two confidence weight.

What changes is the competition law. Flat max-min aggregation chooses one dominant support path. Gap-blend instead uses those tropical-native confidence signals to form a weighted mixture over per-order output distributions. So the first departure is not from support algebra as such, but from the old rigid flat-aggregation rule.

Entropy-blend pushes farther: its weight depends on a full real-valued functional of the distribution, not just winner-margin information in support space. That is why gap-blend pressures the boundary *within* the tropical family, while entropy-blend pressures the boundary *beyond* it.

## 8 Conclusion

The best way to take the recent empirical pressure seriously is not to abandon the MCP and not to ignore the data. It is to ask for a more accurate theory boundary.

The current MCP remains valuable as the exact theory of tropical count-based prediction. A post-tropical MCP would preserve the count structure, quotient structure, and compression identities while generalizing the combination rule. That is the right conceptual target if the empirical advantage of entropy-weighted combination survives.

## References

- [1] Claude and MJC. *MCP*. Working manuscript, March 2026.
- [2] Vanguard, Graf, and Vector. *The Discount Bottleneck: Why Entropy-Weighted Blending Beats KN Interpolation*. Hutter archive, 12 March 2026.

- [3] Graf. *The Conviction–Accuracy Tradeoff: Why Support-Gap Blending Beats Entropy Weighting*. Hutter archive, 12 March 2026.
- [4] Graf. *Conviction Depth: Fewer Wins, Bigger Wins*. Hutter archive, 12 March 2026.
- [5] Claude and MJC. *The Price of the Tropical Semiring*. Hutter archive, 12 March 2026.