Z408 Final-18 Bayesian / Likelihood-Ratio Comparison

This report attempts to quantify the earlier qualitative phrase that the final 18 characters are more naturally explained by deliberate or semi-deliberate filler construction than by ordinary continuation text. The analysis compares explicit generative models and estimates likelihood ratios using the same drift-0, same-column pull-down score used in the preceding analysis.

1. What is being quantified

The observed evidence E is defined as the final 18 characters achieving a drift-0 pull-down score at least as high as the observed score. The observed final-18 string is:

```
6VEXr9WI6qEHM)=UIk
```

The observed score is 26. The score rewards contiguous same-column copied fragments by length squared. Thus a four-symbol aligned repeat is much more influential than four isolated one-symbol matches.

2. Likelihood-ratio logic

For two hypotheses H1 and H2, the likelihood ratio is:

```
LR = P(E | H1) / P(E | H2)
```

This is not the same as the probability that H1 is true. It says how much more expected this evidence is under one specified model than another specified model. The result is therefore model-dependent.

3. Models tested

Model	Description
H_cont_freq	Ordinary/non-copy continuation approximated by independent symbols sampled from the prior Z408 symbol distr
H_cont_markov	Ordinary/non-copy continuation approximated by a first-order symbol Markov chain trained on prior Z408 cipherte
H_lazy_col_equal	Non-block filler: each target position independently reuses a symbol from the same column above, equally weigh
H_lazy_col_recent	Non-block filler: same as above, but recent rows are favoured.
H_copy_p0.25	Pull-down filler: same-column fragment copying; probability of continuing a fragment is 0.25.
H_copy_p0.40	Pull-down filler: same-column fragment copying; probability of continuing a fragment is 0.40.
H_copy_p0.55	Pull-down filler: same-column fragment copying; probability of continuing a fragment is 0.55.
H_copy_p0.70	Pull-down filler: same-column fragment copying; probability of continuing a fragment is 0.70.

4. Simulation results

Each model was run for 5,000 trials. Add-one probability estimates were used, so that a zero observed count is treated as less than 1/(trials+1) rather than as a literal zero probability.

5. Interpretation

Against the simplest non-copy continuation model, H_cont_freq, all copy or same-column reuse models produce much higher probabilities of reaching the observed score. The pull-down copying models are thousands of times more likely than H_cont_freq to produce a score at least as high as the observed final 18.

The stricter comparison is against H_cont_markov, because that model permits some non-random symbol sequencing. Against this stronger continuation model, the mild pull-down model H_copy_p0.25 is about 55 times more likely to produce the observed-or-higher score. Stronger copy-continuation settings produce likelihood ratios of about 64 to 65.

The non-block same-column reuse models are also more favourable than continuation, but much less so than explicit fragment copying. This is useful because it shows that the final 18 need not be explained by fully intentional block copying; even semi-deliberate column-local filler is a better fit than simple continuation, though not as strong as fragment copying.

6. Cautions

These likelihood ratios are not universal probabilities. They depend on how the competing models are defined. In particular, the Markov continuation model is only a rough proxy for ordinary homophonic continuation; it is not a full simulation of English plaintext encoded by the Z408 key. Similarly, the filler models are stylised approximations of possible human behaviour.

Accordingly, the most precise defensible claim is model-relative: under the models tested here, the observed final-18 pull-down score is tens to thousands of times more likely under same-column filler/reuse or pull-down copying models than under the tested ordinary continuation models.

7. Revised wording suggested

A better wording than the earlier qualitative phrase would be:

Under explicit same-column filler models, the observed final-18 pull-down structure is substantially more likely than under the tested ordinary continuation models. Depending on the continuation baseline, the estimated likelihood ratio ranges from roughly tens-to-one to thousands-to-one. This supports deliberate or semi-deliberate filler construction over ordinary continuation text, while remaining dependent on the assumptions of the generative models and without proving authorial intent.