Supplementary Figures

Companion to: *The Dictionary Collision Effect in Computational Decipherment*

Matthew Ruckman

April 2026

This supplement contains the four figures referenced as Supplementary Figures S1 through S4 in the main paper. Each figure is reproduced at full size with its full caption. The corresponding source files are in the paper repository at `paper/v1/figures/`.

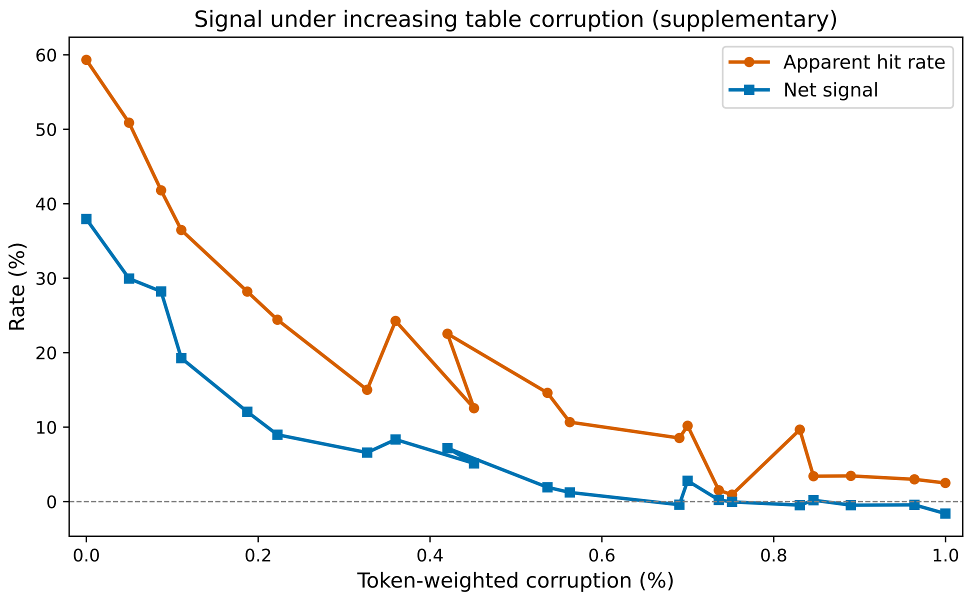

Supplementary Figure S1: Continuous corruption sweep. Apparent hit rate and four-category net signal as decode-table corruption sweeps from 0 to 100 percent (Latin plaintext, Latin 10K dictionary, 5-percent steps). The two curves diverge: apparent hit rate decays slowly, with non-zero residuals even at 100 percent corruption from chance dictionary collisions on wrong characters; net signal collapses rapidly, crossing zero around 30 percent corruption and staying near or slightly below zero thereafter. The gap between the curves at any single corruption value is the dictionary-collision overstatement. Source: `corruption_sweep.pdf`. Referenced in main-paper Section 6.3.

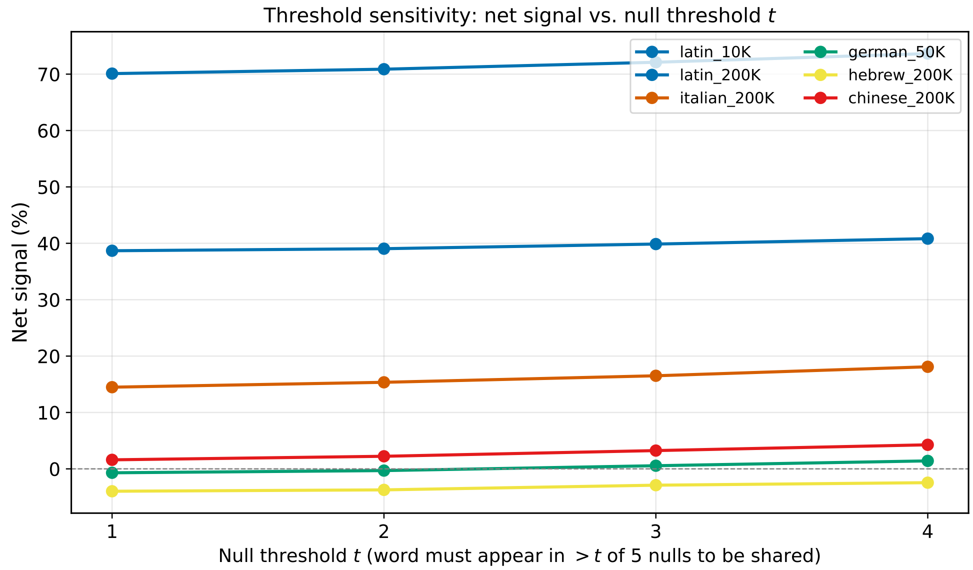

Supplementary Figure S2: Threshold sensitivity. Net signal as a function of the null threshold t for six representative cells spanning correct, wrong, and cognate languages. The direction and relative ordering of every result is preserved across $t \in \{1, 2, 3, 4\}$: correct-language cells are large and positive; wrong-language cells (Hebrew, Chinese) are near zero or slightly negative; cognate and partial cells (Italian, German) are intermediate. Absolute magnitudes shift by 1 to 2 percentage points as t varies, confirming that the specific threshold is not a load-bearing choice. Source: `fig_threshold_sensitivity.pdf`. Referenced in main-paper Section 7.

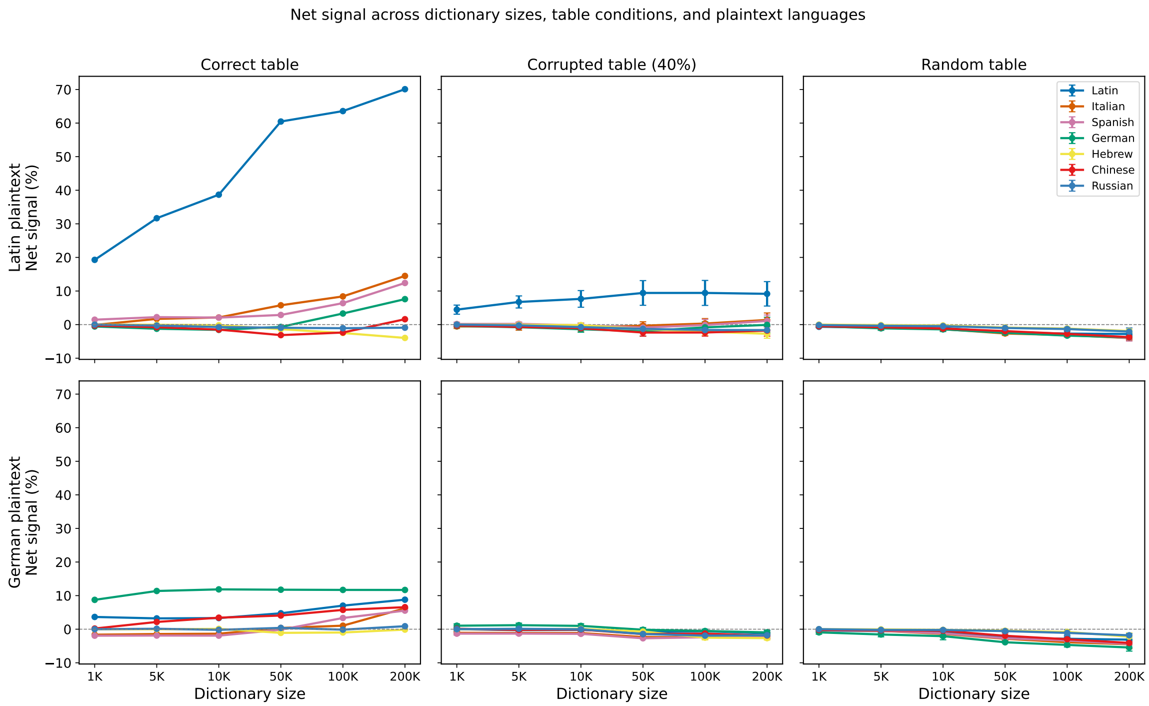

Supplementary Figure S3: Full language grid (2×3). Net signal across the experiment grid, including the 40-percent-corrupted table condition omitted from the compact 2×2 main-paper figure. Rows: plaintext language (Latin, German). Columns: decode-table condition (correct, 40-percent corrupted, random). Within each panel: one line per evaluation language, with error bars (standard deviation over 5 seeds for corrupted and random; correct condition is deterministic). The correct language dominates in the correct-table column; decoder corruption flattens discrimination while preserving the direction of the result; random tables collapse everything to near zero. The effect is symmetric across plaintexts. Source: `fig_language_grid.pdf`. Referenced in main-paper Section 5.1.

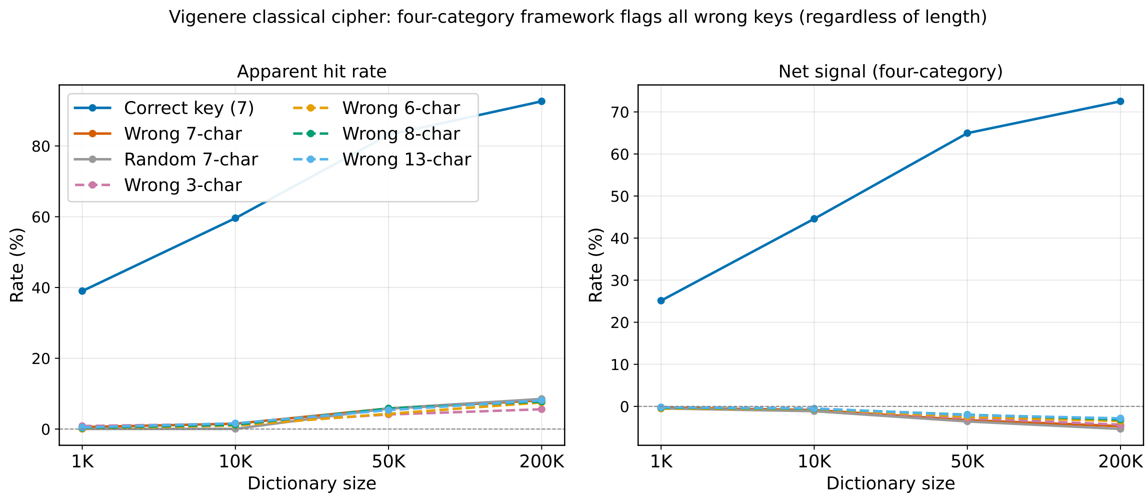

Supplementary Figure S4: Classical Vigenère cipher results. Left: apparent hit rate. Right: four-category net signal. Seven key conditions: correct 7-character key, two wrong same-length keys, and four wrong mismatched-length keys (3, 6, 8, 13 characters). For every wrong key, regardless of length, apparent hit rate at 200K climbs to 5 to 8 percent, but four-category net signal is correctly negative (-2.8 to -5.4 percent). The correct key reaches 72.5 percent net signal at 200K. The framework is robust to key-length mismatch on a real historical cipher, not only the synthetic syllabic design. Source: `fig_vigenere_results.pdf`. Referenced in main-paper Section 9.2.