

# Supplementary Material

## Adaptive Noise in Human Game Play: An Information-Theoretic Signature that Separates Learning Models on Matching-Pennies Data

[Author name removed for double-blind review] [Affiliation removed for double-blind review]

---

This document supplements the main text with (i) a guide to the reproducibility package, (ii) the full per-game, per-model transfer-entropy table that underlies the summary in the main text, and (iii) the complete cross-game-class results, including entropy rates and significance tests, that expand the in-text table. All numbers are produced by the accompanying code from fixed random seeds; no quantity here is hand-entered. Section, table, and equation references of the form “Section 4.5” refer to the main manuscript; supplementary exhibits are numbered S1, S2, ...

## S1. Reproducibility package

All results in the manuscript and this supplement are generated by a single self-contained Python codebase that accompanies the submission. The code depends only on `numpy`, `scipy`, `pandas`, `statsmodels`, and `matplotlib`, runs on a standard workstation in a few minutes per analysis, and is deterministic given the random seeds fixed within it. No external datasets are required beyond the period-by-period Selten–Chmura records, which were obtained from the original authors and are not redistributed here.

### S1.1 Repository structure

The package is organized into self-contained modules:

- `estimators/` — discrete entropy estimators (plug-in, Miller–Madow, Grassberger), the transfer-entropy estimator with shuffle correction, the generalized transfer entropy for arbitrary action support and history order, and the order-1 Markov entropy rate.
- `games/` — payoff matrices for the twelve Selten–Chmura games and the seven canonical comparison games (Section S3 below).
- `agents/` — the five learning models (Q-learning, reinforcement, stochastic fictitious play, self-tuning EWA, impulse-matching) behind a common `select()/observe()` in-

terface.

- **pipeline/** — the analysis drivers: signature computation, the five-model comparison, the temperature sweep, the robustness grid, the surrogate test, the subject-level mixed-effects analysis, the cross-game extension, and the figure scripts.
- **results/** — stored numerical outputs (JSON and CSV) and the publication figures in vector and raster form.

## S1.2 Mapping of results to exhibits

Each manuscript exhibit is reproduced by a named driver, reading from or writing to a stored result file:

**Table S1.** Manuscript exhibits and the code paths that produce them.

| Exhibit            | Driver                                                                  | Stored result                           |
|--------------------|-------------------------------------------------------------------------|-----------------------------------------|
| Figs. 1–3, Table 2 | <code>pipeline/figures.py</code> , <code>signature_comparison.py</code> | <code>signature_comparison.json</code>  |
| Table 3, Figs. 4–5 | <code>pipeline/model_comparison.py</code> , <code>figures.py</code>     | <code>model_comparison.json</code>      |
| Tables 4–5         | <code>pipeline/robustness.py</code>                                     | (computed at run time)                  |
| Table C1 (mixed)   | <code>pipeline/subject_level.py</code>                                  | <code>subject_level.csv</code>          |
| Table 6, Fig. 7    | <code>pipeline/cross_game.py</code>                                     | <code>cross_game_signatures.json</code> |

To regenerate the cross-game analysis, for instance, one runs the cross-game driver, which simulates both models on every game and writes the signature table consumed by Figure 7 and Table 6.

## S2. Full per-game, per-model transfer entropy

Table 3 of the main text reports transfer entropy and entropy rate averaged over all 108 groups. Table S2 gives the underlying per-game breakdown for all six sources, so that the human gradient across games (Section 4.3 of the main text) and its absence in every model can be inspected game by game. Values are mean transfer entropy in bits across the groups of each game (12 groups for the constant-sum games G1–G6, 6 for the non-constant-sum games G7–G12).

The negative entries for the reward-only and cumulative models are shuffle-correction residuals scattered around zero; none differs meaningfully from the transfer-entropy floor. Each source is computed over 864 subjects or agents in total (the per-source totals in the subject-level analysis), so the per-game means are each averages over 48 (G1–G6) or 24 (G7–G12)

**Table S2.** Per-game mean transfer entropy (bits) by source. Human transfer entropy rises with the centrality of the mixed equilibrium (G1, G7 most extreme; G6, G12 most central); self-tuning EWA tracks and exceeds it in every game; the four reward-only and cumulative models remain at the floor throughout.

| Game | Human | EWA   | Q-learn | Reinf. | Fict. play | Impulse |
|------|-------|-------|---------|--------|------------|---------|
| G1   | 0.023 | 0.077 | 0.001   | −0.000 | 0.002      | 0.001   |
| G2   | 0.052 | 0.115 | 0.000   | 0.001  | 0.000      | 0.002   |
| G3   | 0.027 | 0.096 | 0.003   | 0.001  | −0.004     | 0.002   |
| G4   | 0.059 | 0.127 | 0.005   | 0.001  | 0.000      | 0.004   |
| G5   | 0.084 | 0.141 | 0.003   | 0.000  | 0.002      | 0.003   |
| G6   | 0.103 | 0.160 | 0.002   | −0.001 | 0.004      | 0.001   |
| G7   | 0.017 | 0.065 | −0.000  | 0.001  | 0.004      | 0.003   |
| G8   | 0.082 | 0.100 | −0.002  | −0.000 | 0.001      | 0.005   |
| G9   | 0.038 | 0.089 | 0.001   | −0.001 | −0.002     | −0.001  |
| G10  | 0.067 | 0.105 | 0.002   | 0.001  | 0.001      | 0.003   |
| G11  | 0.077 | 0.127 | 0.003   | 0.000  | 0.003      | 0.001   |
| G12  | 0.093 | 0.146 | 0.003   | −0.001 | 0.005      | 0.002   |

row-and-column players.

### S3. Complete cross-game-class results

Table 6 and Figure 7 of the main text summarize the model-side contrast across game classes. Table S3 gives the complete numbers, including both entropy rates with their gap and the one-sided Mann–Whitney test that self-tuning EWA transfer entropy exceeds Q-learning’s. Each cell is computed from 24 independent groups per model per game, simulated under the identical protocol used for the matching-pennies analysis (four row and four column agents, 200 periods, random rematching), with transfer entropy and entropy rate estimated on the game’s own action support (two or three actions).

**Table S3.** Complete cross-game-class signatures (24 groups per model per game). TE and ER are mean transfer entropy and entropy rate in bits;  $p$  is the one-sided Mann–Whitney test that EWA transfer entropy exceeds Q-learning’s. The entropy-rate gap (EWA – Q-learning) is positive in every game; the transfer-entropy gap is significant wherever opponent-responsiveness is strategically relevant, and absent in the dominance-solvable Prisoner’s Dilemma and the degenerate Shapley cyler.

| Game                      | Class           | EWA TE | QL TE | $p$                  | EWA ER | QL ER | ER gap |
|---------------------------|-----------------|--------|-------|----------------------|--------|-------|--------|
| SC matching pennies       | zero-sum (data) | 0.137  | 0.003 | $1.5 \times 10^{-9}$ | 0.829  | 0.075 | +0.75  |
| Matching Pennies          | zero-sum        | 0.024  | 0.006 | $3.1 \times 10^{-8}$ | 1.001  | 0.908 | +0.09  |
| Rock–Paper–Scissors       | zero-sum        | 0.027  | 0.003 | $2.2 \times 10^{-6}$ | 1.587  | 1.471 | +0.12  |
| Shapley                   | zero-sum        | 0.018  | 0.025 | 0.89                 | 1.586  | 1.214 | +0.37  |
| Stag Hunt                 | potential       | 0.088  | 0.010 | $6.2 \times 10^{-8}$ | 0.564  | 0.133 | +0.43  |
| Coordination $3 \times 3$ | potential       | 0.079  | 0.048 | $2.7 \times 10^{-6}$ | 1.280  | 0.416 | +0.86  |
| Battle of the Sexes       | general-sum     | 0.084  | 0.019 | $1.5 \times 10^{-9}$ | 0.941  | 0.236 | +0.71  |
| Prisoner’s Dilemma        | general-sum     | 0.007  | 0.010 | 0.72                 | 0.634  | 0.189 | +0.45  |

Two points reinforce the reading in the main text. First, the entropy-rate gap is positive in all eight games, including the two where the transfer-entropy gap is absent: even in the Prisoner’s Dilemma and Shapley’s game, self-tuning EWA sustains more policy entropy than Q-learning, so the two failures are specific to the *responsiveness* signature and not to stochasticity. Second, the two non-significant games are exactly the cases where opponent-responsiveness should not register: the Prisoner’s Dilemma is dominance-solvable, so the payoff-maximizing action is independent of the opponent and both models sit at the transfer-entropy floor; and Shapley’s game is a classical example in which learning dynamics fail to converge, leaving both models near the floor with overlapping distributions. The contrast thus generalizes across game classes precisely where opponent-responsiveness is a meaningful strategic quantity.

We reiterate the scope stated in the main text: this analysis extends the *model-side* contrast only. No human data exist for the seven comparison games, so these results do not speak to whether human subjects display the maintained-entropy, opponent-responsive signatures outside the matching-pennies setting. Establishing that would require new experiments, which the protocol and estimators here are designed to accommodate without modification.