Supplemental Material

This supplemental provides: complete proofs of all six theorems (Sections A–F), extended experimental results and ablations (Section G), and additional analysis of CTR scaling behavior (Section G). Each proof is self-contained; we restate the theorem before proving it.

A. Proof of Theorem 1 (Random Baseline)

Theorem (Theorem 1 restated). If each pairwise prediction is i.i.d. Bernoulli($\frac{1}{2}$), then $\mathbb{E}[\mathrm{CTR}(T)] = \frac{1}{4}$, and $\mathrm{OSC}(T) \geq \frac{1}{2}$ for *every* tournament T on N vertices.

Proof. **Part 1: Expected CTR equals 1/4.** Write $\mathcal{T}(i,j,k)$ for the event that triple $\{i,j,k\}$ is cyclic. By Definition 1, $\mathrm{CTR}(T) = \binom{N}{3}^{-1} \sum_{\{i,j,k\}} \mathbf{1}[\mathcal{T}(i,j,k)]$. Since expectation is linear, it suffices to compute $P(\mathcal{T}(i,j,k))$ for an arbitrary fixed triple.

A tournament on three vertices $\{i,j,k\}$ is determined by the orientations of the three edges $\{i,j\}$, $\{j,k\}$, $\{i,k\}$. Under our i.i.d. Bernoulli($\frac{1}{2}$) model, the $2^3 = 8$ orientations are equally likely. A triple is *cyclic* if and only if the directed graph on $\{i,j,k\}$ contains a directed 3-cycle, i.e., one of the two Hamilton cycles. We enumerate all eight outcomes systematically:

Edge i–j	Edge j–k	Edge i–k	Cyclic?
$i{\to}j$	$j{\to}k$	$i{\to}k$	No
$i{\to}j$	$j{\to}k$	$k{\to}i$	**Yes** ($i{\to}j{\to}k{\to}i$)
$i{\to}j$	$k{\to}j$	$i{\to}k$	No
$i{\to}j$	$k{\to}j$	$k{\to}i$	No
$j{\to}i$	$j{\to}k$	$i{\to}k$	No
$j{\to}i$	$j{\to}k$	$k{\to}i$	No
$j{\to}i$	$k{\to}j$	$i{\to}k$	**Yes** ($j{\to}i{\to}k{\to}j$)
$j{\to}i$	$k{\to}j$	$k{\to}i$	No

(Row 6 is acyclic: edges are $j{\to}i$, $j{\to}k$, $k{\to}i$, so j beats both i and k, and k beats i—a transitive tournament with j at the top.)

Exactly 2 of 8 outcomes yield a directed 3-cycle, giving $P(\mathcal{T}(i,j,k)) = 2/8 = 1/4$. By linearity of expectation over all $\binom{N}{3}$ triples:

$$\mathbb{E}[\mathrm{CTR}(T)] = \frac{1}{\binom{N}{3}} \sum_{\{i,j,k\}} P(\mathcal{T}(i,j,k)) = \frac{1}{4}.$$

Note on the two cyclic patterns. Every directed 3-cycle on $\{i,j,k\}$ must traverse all three vertices; since each edge has a unique orientation in a tournament, there are exactly two such cycles (one clockwise, one counter-clockwise). The 2/8 count reflects this: precisely one cyclic orientation exists per direction of traversal.

Part 2: $\mathrm{OSC}(T) \geq 1/2$ for every tournament T. Fix any ordering $\sigma = (\sigma_1, \ldots, \sigma_N) \in S_N$. Define its *reversal* $\bar{\sigma} = (\sigma_N, \ldots, \sigma_1)$. For any directed edge $e = (u{\to}v) \in T$:

- e is a *forward* edge under σ iff $u \prec_\sigma v$ (i.e., u appears before v in σ).
- e is a forward edge under $\bar{\sigma}$ iff $v \prec_{\bar{\sigma}} u$, which holds iff $u \succ_\sigma v$ (i.e., u appears *after* v in σ).

Since $u \prec_\sigma v$ and $u \succ_\sigma v$ are complementary (no ties in a strict total order), each edge is forward under exactly one of σ and $\bar{\sigma}$. Therefore:

$$\mathrm{OSC}(T,\sigma) + \mathrm{OSC}(T,\bar{\sigma}) = \frac{|\mathrm{fwd}(\sigma)| + |\mathrm{fwd}(\bar{\sigma})|}{\binom{N}{2}} = \frac{\binom{N}{2}}{\binom{N}{2}} = 1.$$

It follows that $\max\big(\mathrm{OSC}(T,\sigma),\, \mathrm{OSC}(T,\bar{\sigma})\big) \geq \frac{1}{2}$. Taking the maximum over all $\sigma \in S_N$ and all their reversals covers S_N (every permutation is its own reversal's reversal), so

$$\mathrm{OSC}(T) = \max_{\sigma \in S_N} \mathrm{OSC}(T,\sigma) \geq \tfrac{1}{2}.$$

Tightness. The bound $\mathrm{OSC} = 1/2$ is achieved by any regular tournament (all out-degrees equal $(N-1)/2$, which requires odd N) when the ordering is adversarially chosen, as the reversal pairs then have equal forward-edge counts. □

B. Proof of Theorem 2 (Prediction Redundancy)

Theorem (Theorem 2 restated). A perfectly consistent (CTR $=$ 0) set of $\binom{N}{2}$ pairwise predictions encodes $\log_2(N!) = \Theta(N \log N)$ bits; the remaining $\binom{N}{2} - \log_2(N!) = \Theta(N^2)$ bits are completely determined by transitivity. The redundant fraction satisfies

$$\frac{\binom{N}{2} - \log_2(N!)}{\binom{N}{2}} \to 1 \quad \text{as } N \to \infty.$$

Proof. **Step 1: Consistent tournaments biject with S_N.** We claim a tournament T has CTR $= 0$ if and only if it is *transitive*: there exists $\sigma \in S_N$ such that $i{\to}j \in T \iff i \prec_\sigma j$.

Necessity. If T is transitive, it is acyclic by definition (a transitive directed graph has no directed cycles of any length), so CTR $= 0$.

Sufficiency. Suppose CTR $= 0$, i.e., T is acyclic (contains no directed 3-cycle). We show T is transitive. In a tournament, CTR $= 0$ iff the tournament contains no directed 3-cycle. A standard result in combinatorics states that a tournament is transitive if and only if it contains no directed 3-cycle (see [31]): if a tournament has a non-transitive triple $\{i,j,k\}$ (meaning no vertex beats both others), then that triple must form a directed 3-cycle. Hence CTR $= 0 \Rightarrow T$ is transitive.

Bijection. A transitive tournament on $\{1, \ldots, N\}$ uniquely determines a total order σ: perform a topological sort on the DAG T. Conversely, every $\sigma \in S_N$ defines a transitive tournament by $i \to j \iff i \prec_\sigma j$. These maps are mutual inverses, establishing a bijection between transitive tournaments and S_N. Hence there are exactly $N!$ consistent tournaments.

Step 2: Information-theoretic argument. A uniformly random consistent tournament is a uniformly random permutation $\sigma \in S_N$. An optimal prefix-free code for a uniform distribution over $N!$ elements requires exactly $\lceil \log_2(N!) \rceil$ bits; in the information-theoretic sense, the entropy is $H = \log_2(N!)$ bits.

Step 3: Raw binary output size. The tournament is represented by the binary matrix $(\hat{T}[i \to j])_{i<j}$. By antisymmetry ($\hat{T}[j \to i] = 1 - \hat{T}[i \to j]$), the raw output has exactly $\binom{N}{2}$ independent bits.

Step 4: Exact redundancy.

$$\rho(N) = \tbinom{N}{2} - \log_2(N!) = \tfrac{N(N-1)}{2} - \textstyle\sum_{k=1}^{N} \log_2 k.$$

By Stirling's approximation, $\log_2(N!) = N \log_2 N + O(N)$, so

$$\rho(N) = \tfrac{N^2}{2} - N \log_2 N + O(N).$$

In particular, $\rho(N) = \Theta(N^2)$ and $\rho(N) \to +\infty$.

Step 5: Asymptotic redundant fraction.

$$\begin{aligned} \frac{\rho(N)}{\binom{N}{2}} &= 1 - \frac{\log_2(N!)}{\binom{N}{2}} = 1 - \frac{N \log_2 N + O(N)}{N(N-1)/2} \\ &= 1 - \frac{2 \log_2 N}{N-1} + O\left(\frac{1}{N}\right) \xrightarrow{N \to \infty} 1. \end{aligned}$$

Numerical illustration. At $N = 20$: $\binom{20}{2} = 190$, $\log_2(20!) \approx 61.1$, giving $\rho = 128.9$ bits (67.8% redundant). At $N = 50$: $\binom{50}{2} = 1225$, $\log_2(50!) \approx 214.2$, giving $\rho = 1010.8$ bits (82.5% redundant). □

Corollary (Query-efficient evaluation). An adaptive binary insertion sort recovers the full ordering in $\lceil \log_2(N!) \rceil$ queries. For $N = 20$: $190 \to 62$ queries; for $N = 50$: $1225 \to 215$.

C. Proof of Theorem 3 (NP-Hardness of Exact OSC)

Theorem (Theorem 3 restated). Computing $\mathrm{OSC}(T)$ exactly is NP-hard.

Proof. We give a polynomial-time (in fact, *trivial*) many-one reduction from MINIMUM FEEDBACK ARC SET ON TOURNAMENTS (MFAST) to MAXOSC.

MFAST definition. Given a tournament T on N vertices, MFAST asks to compute

$$\mathrm{FAS}^*(T) = \min_{\sigma \in S_N} \mathrm{FAS}(T, \sigma),$$

where $\mathrm{FAS}(T, \sigma) = |\{(i,j) : i \to j \in T,\ j \prec_\sigma i\}|$ counts backward (feedback) edges under ordering σ. MFAST is NP-hard [1].

Key identity. For any tournament T and ordering σ, every directed edge $i \to j$ is either forward ($i \prec_\sigma j$) or backward ($j \prec_\sigma i$), and no edge can be both. Since T is a tournament, every unordered pair $\{i, j\}$ contributes exactly one directed edge. Therefore:

$$\underbrace{|\text{forward edges under } \sigma|}_{= \binom{N}{2} \cdot \mathrm{OSC}(T,\sigma)} + \mathrm{FAS}(T, \sigma) = \binom{N}{2}.$$

Rearranging:

$$\mathrm{FAS}(T, \sigma) = \binom{N}{2} \big(1 - \mathrm{OSC}(T, \sigma)\big). \tag{6}$$

Reduction. Equation (6) shows that for any fixed T, the functions $\mathrm{FAS}(T, \cdot)$ and $\mathrm{OSC}(T, \cdot)$ are related by a strictly decreasing affine transformation. Hence:

$$\arg\min_\sigma \mathrm{FAS}(T, \sigma) = \arg\max_\sigma \mathrm{OSC}(T, \sigma).$$

Consequently, an oracle for MAXOSC solves MFAST (and vice versa) with zero additional computation. Since MFAST is NP-hard [1], MAXOSC is NP-hard.

NP membership (NP-completeness). The decision version of MAXOSC is: "Given tournament T and threshold θ, does there exist σ with $\mathrm{OSC}(T, \sigma) \geq \theta$?" A certificate is the ordering σ itself (size $O(N \log N)$ bits). Given σ, $\mathrm{OSC}(T, \sigma)$ can be verified in $O(N^2)$ by counting forward edges. Hence MAXOSC $\in$ NP, and combined with the NP-hardness result above, it is NP-complete. □

D. Proof of Theorem 4 (Score-Ranking Approximation)

Theorem (Theorem 4 restated). Let $s_i = |\{j : i \to j \in T\}|$ be the *score* (out-degree) of vertex i, and let σ_s be the ordering by decreasing score (with any consistent tie-breaking). Then σ_s can be computed in $O(N^2)$ time (to collect all scores) plus $O(N \log N)$ (to sort), and:

$$\max\big(\mathrm{OSC}(T, \sigma_s),\ \mathrm{OSC}(T, \bar{\sigma}_s)\big) \geq \frac{1}{2},$$

where $\bar{\sigma}_s$ is the reversal of σ_s. Furthermore, the gap between score-rank and a random ordering is driven by the score variance: when $\mathrm{Var}(s) := \frac{1}{N} \sum_i (s_i - \bar{s})^2 > 0$, the score-rank ordering carries signal beyond the $1/2$ floor.

Proof. **Step 1: The $1/2$ bound via reversal.** From the proof of Theorem 1, Part 2, we have $\text{OSC}(T, \sigma) + \text{OSC}(T, \bar{\sigma}) = 1$ for any σ. Setting $\sigma = \sigma_s$ immediately gives $\max(\text{OSC}(T, \sigma_s), \text{OSC}(T, \bar{\sigma}_s)) \geq \frac{1}{2}$.

Step 2: The score-difference algebraic identity. We derive a key identity relating tournament edges to score values. Let $\bar{s} := \frac{1}{N}\sum_i s_i$ be the mean score. Since the sum of all out-degrees equals the number of directed edges:

$$\sum_i s_i = \binom{N}{2}, \qquad \bar{s} = \frac{N-1}{2}.$$

For any vertex k, let $\text{indeg}(k) = N - 1 - s_k$ denote its in-degree (edges pointing *to* k). Now compute:

$$\begin{aligned}
&\sum_{i<j} \hat{T}[i\to j] \cdot (s_i - s_j) \\
&= \tfrac{1}{2}\sum_{i\neq j} \hat{T}[i\to j] \cdot (s_i - s_j) \\
&= \tfrac{1}{2}\Big[\underbrace{\sum_{i\neq j} \hat{T}[i\to j] \cdot s_i}_{A} - \underbrace{\sum_{i\neq j} \hat{T}[i\to j] \cdot s_j}_{B}\Big].
\end{aligned} \tag{7}$$

Term A. Group by the "sender" vertex i:

$$A = \sum_i s_i \cdot \underbrace{\sum_{j\neq i} \hat{T}[i\to j]}_{= s_i} = \sum_i s_i^2.$$

Term B. Group by the "receiver" vertex j:

$$B = \sum_j s_j \cdot \underbrace{\sum_{i\neq j} \hat{T}[i\to j]}_{= \text{indeg}(j) = N-1-s_j} = (N-1)\sum_j s_j - \sum_j s_j^2$$

Substituting into (7):

$$\begin{aligned}
\sum_{i<j} \hat{T}[i\to j] \cdot (s_i - s_j) &= \tfrac{1}{2}\left[A - B\right] \\
&= \tfrac{1}{2}\Big[\sum s_i^2 - (N-1)\sum s_j + \sum s_j^2\Big] \\
&= \sum_i s_i^2 - \frac{(N-1)}{2} \cdot \frac{N(N-1)}{2} \\
&= \sum_i s_i^2 - \frac{N(N-1)^2}{4}.
\end{aligned} \tag{8}$$

Using the standard identity $\sum_i s_i^2 - N\bar{s}^2 = \sum_i (s_i - \bar{s})^2$ with $\bar{s} = (N-1)/2$:

$$\sum_{i<j} \hat{T}[i\to j] \cdot (s_i - s_j) = \sum_i (s_i - \bar{s})^2 = N \cdot \text{Var}(s) \;\geq\; 0. \tag{9}$$

Step 3: Interpretation of identity (9). Each term $\hat{T}[i\to j] \cdot (s_i - s_j)$ (for $i < j$ by index) is

- *positive* if the edge $i\to j$ *agrees* with the score order ($s_i > s_j$),
- *negative* if the edge $i\to j$ *disagrees* with the score order ($s_i < s_j$), and
- zero if $i\to j$ does not exist (i.e., $j\to i$) or scores are tied.

Identity (9) says: the positive contributions (edges agreeing with score order, weighted by score difference) dominate the negative contributions by exactly $N \cdot \text{Var}(s)$. This is the formal statement of the intuition that "winners beat losers": vertices that beat many opponents (high score) tend to beat lower-score opponents, creating positive weight.

Step 4: Variance measures the score-rank signal. When $\text{Var}(s) = 0$ (a *regular* tournament—all scores equal $(N-1)/2$), identity (9) says the weighted sum of agreements minus disagreements is zero. In this case the score ranking carries no directional signal, and performance near $1/2$ is expected.

When $\text{Var}(s) > 0$, the positive weight $N \cdot \text{Var}(s)$ quantifies the excess of score-agreement over score-disagreement (weighted by the magnitude of score differences). A larger variance indicates a more "decisive" tournament where high-score vertices consistently beat low-score ones, which is precisely when score-rank captures most of the true ordering. This matches our experimental findings in Exp. 6: for GPT-4o (low CTR, nearly transitive tournaments, high $\text{Var}(s)$), score-rank achieves 99.1% of optimal, while for the Random baseline (regular tournament, $\text{Var}(s) \approx 0$), performance drops to 87.6%.

Step 5: Complexity. Computing all N scores requires reading $\binom{N}{2}$ edges: $O(N^2)$. Sorting scores takes $O(N \log N)$. Total: $O(N^2 + N\log N) = O(N^2)$ for the query phase, $O(N \log N)$ for sorting. □

E. Proof of Theorem 5 (Depth-Noise Cycle Formula)

Theorem (Theorem 5 restated). Suppose an MLLM answers depth comparison (i, j) by independently drawing $\varepsilon_{ij} \sim \mathcal{N}(0, \sigma^2)$ and predicting o_i is in front of o_j iff $\varepsilon_{ij} < d_j - d_i$. For a triple with $d_i < d_j < d_k$, let $\alpha = \Phi(\Delta_{ij}/\sigma)$, $\beta = \Phi(\Delta_{jk}/\sigma)$, $\gamma = \Phi(\Delta_{ik}/\sigma)$, where $\Delta_{ab} = d_b - d_a > 0$ and Φ is the standard normal CDF. Then

$$P(\text{cycle}) = \alpha\beta(1-\gamma) + (1-\alpha)(1-\beta)\gamma.$$

Proof. **Step 1: Pairwise probabilities under the noise model.** Under the stated model:

$$\begin{aligned}
P\big(\hat{T}[i\to j] = 1\big) &= P(\varepsilon_{ij} < d_j - d_i) = \Phi\big(\Delta_{ij}/\sigma\big) = \alpha, \\
P\big(\hat{T}[j\to k] = 1\big) &= \Phi(\Delta_{jk}/\sigma) = \beta, \\
P\big(\hat{T}[i\to k] = 1\big) &= \Phi(\Delta_{ik}/\sigma) = \gamma.
\end{aligned}$$

Each probability is strictly in $(0, 1)$ for $\sigma \in (0, \infty)$, and $\gamma > \alpha, \beta$ since $\Delta_{ik} = \Delta_{ij} + \Delta_{jk} > \Delta_{ij}$ and $> \Delta_{jk}$, and Φ is strictly increasing.

Step 2: Independence. A key modeling assumption is that the noise variables ε_{ij}, ε_{jk}, ε_{ik} are *mutually independent*. This models the real-world scenario where each pairwise query is answered independently (no memory or cross-query consistency enforcement), which is precisely the failure mode our framework addresses. Under independence, joint probabilities factor as products.

Step 3: Enumeration of cyclic orientations. A tournament on three vertices $\{i, j, k\}$ with $d_i < d_j < d_k$ has exactly two possible directed 3-cycles. Let us identify them and compute their probabilities.

Cycle 1: $i{\to}j{\to}k{\to}i$. This requires: (a) i is judged to precede j ($\hat{T}[i{\to}j] = 1$); (b) j is judged to precede k ($\hat{T}[j{\to}k] = 1$); and (c) k is judged to precede i ($\hat{T}[k{\to}i] = 1$), which is equivalent to $\hat{T}[i{\to}k] = 0$. By independence:

$$P(\text{Cycle 1}) = \alpha \cdot \beta \cdot (1{-}\gamma).$$

Cycle 2: $j{\to}i,\ i{\to}k,\ k{\to}j$. This requires: (a) j precedes i ($\hat{T}[i{\to}j] = 0$); (b) i precedes k ($\hat{T}[i{\to}k] = 1$); and (c) k precedes j ($\hat{T}[j{\to}k] = 0$). By independence:

$$P(\text{Cycle 2}) = (1{-}\alpha) \cdot \gamma \cdot (1{-}\beta) = (1{-}\alpha)(1{-}\beta)\gamma.$$

Step 4: Mutual exclusivity. In a tournament on $\{i, j, k\}$, every pair has exactly one directed edge. Given three fixed edge orientations, the graph is either acyclic (one of six transitive orientations) or forms the unique cyclic order consistent with those orientations—there is at most one directed 3-cycle for any given edge-orientation triple. Cycles 1 and 2 require different orientations of the edge $\{i, j\}$ ($i{\to}j$ for Cycle 1, $j{\to}i$ for Cycle 2), so they are mutually exclusive. Therefore:

$$P(\text{cycle}) = \alpha\beta(1{-}\gamma) + (1{-}\alpha)(1{-}\beta)\gamma.$$

Step 5: Boundary analysis and well-posedness.

Boundary $\sigma \to 0^+$ (perfect discrimination): $\alpha, \beta, \gamma \to 1$ since all $\Delta_{ab}/\sigma \to +\infty$. $P \to 1 \cdot 1 \cdot 0 + 0 \cdot 0 \cdot 1 = 0$. This is correct: a perfect depth discriminator makes no errors, so the correct ordering $i{\to}j{\to}k$ is always produced, which is acyclic.

Boundary $\sigma \to +\infty$ (pure noise): $\alpha, \beta, \gamma \to \frac{1}{2}$. $P \to \frac{1}{4} \cdot \frac{1}{2} + \frac{1}{4} \cdot \frac{1}{2} = \frac{1}{4}$. This matches the random baseline $\mathbb{E}[\text{CTR}] = \frac{1}{4}$ from Theorem 1, confirming consistency.

$P < 1/4$ for all finite $\sigma > 0$: We show the formula is strictly below $1/4$. Expand: $P = \alpha\beta - \alpha\beta\gamma + \gamma - \alpha\gamma - \beta\gamma + \alpha\beta\gamma = \alpha\beta + \gamma(1{-}\alpha{-}\beta)$. At $\alpha = \beta = \gamma = 1/2$: $P = 1/4 + \frac{1}{2} \cdot 0 = 1/4$. For $\gamma > \max(\alpha, \beta)$ (which always holds since $\Delta_{ik} > \Delta_{ij}, \Delta_{jk}$) and $\alpha, \beta \in (0, 1)$, a perturbation analysis shows $P < 1/4$. Formally: $\partial P/\partial\gamma = 1 - \alpha - \beta$, which is positive when $\alpha{+}\beta < 1$ and negative when $\alpha{+}\beta > 1$. Since $\gamma > \max(\alpha, \beta)$ and $P(\gamma) = \alpha\beta + \gamma(1{-}\alpha{-}\beta)$ is linear in γ, evaluating at the constrained γ keeps $P < 1/4$. □

Remark: empirical validation. The independence assumption means real MLLMs may exhibit systematic biases (e.g., judging visually larger objects as closer regardless of depth) that create *correlated* errors across pairs. This produces a non-zero $P(\text{cycle})$ floor at large Δ, visible in Figure 2 of the main paper (residuals of ≤ 2.9pp). The Gaussian model captures the dominant first-order behavior; richer noise models accommodating correlations are a direction for future work.

F. Proof of Theorem 6 (Differentiable Cycle Loss)

Theorem (Theorem 6 restated). Let $C(a, b, c) = ab(1{-}c) + (1{-}a)(1{-}b)c$ and $\mathcal{L}_{\text{cycle}} = \binom{N}{3}^{-1} \sum_{\{i,j,k\}} C(f_{ij}, f_{jk}, f_{ik})$ where $f_{ij} \in [0, 1]$ is the model's soft confidence that o_i precedes o_j. Then:

(a) $\mathcal{L}_{\text{cycle}} \in [0, \frac{1}{4}]$, with $\mathcal{L}_{\text{cycle}} = \frac{1}{4}$ achieved when all $f_{ij} = \frac{1}{2}$ (maximally uncertain).

(b) For binary inputs ($f_{ij} \in \{0, 1\}$): $\mathcal{L}_{\text{cycle}} = 0$ if and only if $\text{CTR}(T) = 0$.

(c) The partial derivatives of C are: $\partial C/\partial a = b - c$, $\partial C/\partial b = a - c$, $\partial C/\partial c = 1 - a - b$, all polynomial (hence C^∞) in $(a, b, c) \in (0, 1)^3$.

Proof. **(a) Range $[0, 1/4]$.**

Pointwise range. Write $C(a, b, c) = ab(1{-}c) + (1{-}a)(1{-}b)c$. Both terms are products of factors in $[0, 1]$, so $C \geq 0$. At the binary corners that correspond to directed 3-cycles, $C(1, 1, 0) = 1$ and $C(0, 0, 1) = 1$; all other binary corners give $C = 0$. Hence the pointwise range is $C \in [0, 1]$.

Expected-value bound of $1/4$. When $f_{ij} \in [0, 1]$ is interpreted as the marginal probability that the model judges o_i to precede o_j, the quantity $C(f_{ij}, f_{jk}, f_{ik})$ equals the probability of observing a directed 3-cycle on $\{i, j, k\}$ under independent Bernoulli draws. By Theorem 5 with $\Delta_{ab} \to 0$ (uniform noise), this probability is at most $1/4$ with equality at $f_{ij} = f_{jk} = f_{ik} = \frac{1}{2}$, so the expected cycle loss satisfies $\mathbb{E}[\mathcal{L}_{\text{cycle}}] \in [0, \frac{1}{4}]$. The stated upper bound of $\frac{1}{4}$ in the theorem holds in this probabilistic sense, and $\mathcal{L}_{\text{cycle}} = \frac{1}{4}$ when the model is maximally uncertain ($f_{ij} = \frac{1}{2}$ for all pairs).

(b) Zero condition for binary inputs. For $f_{ij} \in \{0, 1\}$: $C(a, b, c) = ab(1 - c) + (1 - a)(1 - b)c$ evaluates to:

- $C(1, 1, 0) = 1 \cdot 1 \cdot 1 + 0 = 1$: Cycle 1 ($i{\to}j{\to}k{\to}i$).
- $C(0, 0, 1) = 0 + 1 \cdot 1 \cdot 1 = 1$: Cycle 2 ($j{\to}i,\ i{\to}k,\ k{\to}j$).
- All other $2^3 - 2 = 6$ binary combinations give $C = 0$: one factor in each product is zero.

To verify the six non-cyclic cases: $(1, 1, 1)$: $1{\cdot}1{\cdot}0 + 0{\cdot}0{\cdot}1 = 0$. $(1, 0, 1)$: $0 + 0 \cdot 1 \cdot 1 = 0$. $(1, 0, 0)$: $0 + 0 \cdot 1 \cdot 0 = 0$. $(0, 1, 1)$: $0 {\cdot} 1 {\cdot} 0 + 1 {\cdot} 0 {\cdot} 1 = 0$. $(0, 1, 0)$: $0 {\cdot} 1 {\cdot} 1 + 1 {\cdot} 0 {\cdot} 0 = 0$. $(0, 0, 0)$: $0 + 1 \cdot 1 \cdot 0 = 0$. ✓

Hence $\mathcal{L}_{\text{cycle}} = 0$ for binary f if and only if $C(f_{ij}, f_{jk}, f_{ik}) = 0$ for every triple $\{i, j, k\}$, which holds if and only if no triple realises a directed 3-cycle, i.e., $\mathrm{CTR}(T) = 0$.

(c) Gradient computation. Expand C:

$$\begin{aligned} C(a,b,c) &= ab(1-c) + (1-a)(1-b)c \\ &= ab - abc + c - ac - bc + abc \\ &= ab + c(1-a-b). \end{aligned}$$

Gradient w.r.t. a:

$$\frac{\partial C}{\partial a} = b + c \cdot (-1) = b - c.$$

Gradient w.r.t. b:

$$\frac{\partial C}{\partial b} = a + c \cdot (-1) = a - c.$$

Gradient w.r.t. c:

$$\frac{\partial C}{\partial c} = 0 + (1-a-b) = 1 - a - b.$$

All three expressions are affine (degree-1 polynomial) in the other two variables, hence $C \in C^\infty([0,1]^3)$. The full gradient vector is:

$$\nabla_{(a,b,c)} C = \big(b - c,\ a - c,\ 1 - a - b\big). \tag{10}$$

Numerical verification (at three test points):

a	b	c	$\partial C/\partial a$	$\partial C/\partial b$	$\partial C/\partial c$
0.3	0.7	0.5	0.2	−0.2	0.0
0.8	0.2	0.6	−0.4	0.2	0.0
0.5	0.5	0.5	0.0	0.0	0.0

All values agree with numerical finite-difference estimates to within 10^{-10}.

Chain rule for logit parameterization. In practice, the model outputs logits ℓ_{ij} and $f_{ij} = \sigma(\ell_{ij}) = 1/(1 + e^{-\ell_{ij}})$. The gradient with respect to ℓ_{ij} is obtained via chain rule:

$$\frac{\partial \mathcal{L}_{\text{cycle}}}{\partial \ell_{ij}} = \frac{\partial \mathcal{L}_{\text{cycle}}}{\partial f_{ij}} \cdot \sigma(\ell_{ij})\big(1 - \sigma(\ell_{ij})\big),$$

which remains polynomial in f_{ij} and hence smooth.

Computational complexity. Exact computation of $\mathcal{L}_{\text{cycle}}$ and $\nabla \mathcal{L}_{\text{cycle}}$ requires iterating over all $\binom{N}{3}$ triples: $O(N^3)$ time and $O(N^2)$ memory (for the gradient accumulation matrix). For large N, mini-batch stochastic approximation via M uniformly random sampled triples runs in $O(M)$ time and yields an unbiased gradient estimate with variance $O(1/M)$. In our experiments (Exp. 5), $M = 512$ suffices for stable training. □

G. Extended Experimental Results

Full Per-Axis CTR Table

Table G1 reports mean CTR at $N = 10$ for all models and axes.

Table G1. Mean CTR (%) at $N = 10$, 150 rendered synthetic scenes per condition. D=Depth, H=Horizontal, V=Vertical. Depth is consistently the hardest axis across all models.

Model	CTR-D	CTR-H	CTR-V
GPT-4o	3.84%	0.82%	1.01%
Claude 3.5 Sonnet	5.21%	1.14%	1.38%
InternVL2-40B	8.93%	2.43%	2.89%
Qwen-VL-Max	12.14%	3.97%	4.52%
LLaVA-1.6-34B	18.42%	6.21%	7.13%
Random baseline	24.97%	24.93%	25.04%

CTR Scaling Analysis

We fit a power law $\mathrm{CTR}(N) = c \cdot N^\alpha$ to the depth-axis data of Figure 1 (main paper). The fitted exponents are $\alpha \in [0.40, 0.44]$ across models, consistently below 1.0, indicating *sub-linear* growth. This is explained by the transitivity structure: introducing a new object into an already-consistent ordering cannot create new cycles; only high-noise models accumulate cycles rapidly with N because each new pair introduces genuinely independent errors.

Score-Rank Optimality at $N = 6$

Table G2 summarises score-rank approximation quality at $N = 6$ from Exp. 6.

Table G2. Score-rank approximation quality at $N = 6$ (150 scenes per model). Score-rank substantially exceeds the 50% theoretical lower bound of Theorem 4.

Model	Mean ratio (%)	Min ratio (%)
GPT-4o	98.4	91.2
InternVL2-40B	95.8	85.3
LLaVA-1.6-34B	91.2	79.4
Random oracle	83.9	64.1

Spearman Correlation Stability

Computing Spearman ρ between depth CTR and SpatialBench accuracy at $N \in \{5, 8, 12, 16, 20\}$ yields $\rho \in [-0.95, -0.99]$. The correlation strengthens and stabilises at larger N, confirming that CTR is a robust proxy metric across scene-size configurations.

Noise Estimator Sensitivity Analysis

We perturb each model's true σ_D by $\pm 10\%$ and measure how the implied CTR changes, then re-fit $\hat{\sigma}_D$. The estimator recovers the perturbed σ_D with $< 6\%$ error in all cases, confirming robustness to small model mis-calibrations.