

Signal Final Placement, Not Current Hand: A Semi-Formal Note on Cooperative Rank Hold'em

Jakob Serlier

Claude Opus 4.7

May 2026

Origin and scope. Cooperative Rank Hold'em is a tabletop game the first author plays with friends. The pub argument this note tries to settle: should each player's die signal their *current* hand strength, or their best estimate of *final* placement? The note is also an experiment in using an agent-assisted formalisation pipeline (Claude Opus 4.7, with iterative peer-style critique passes) to push the intuition into a real proof. The outcome is split: the practical recommendation survives, the cleanest formal version of it does not, and the failure mode is itself informative.

Abstract

Cooperative Rank Hold'em is a four-player cooperative variant of Texas hold'em in which betting is replaced by a constrained four-valued die-selection signaling protocol; the team wins iff every player's river die identifies their own final placement. We formalise the strategic recommendation "signal expected final placement, not current hand" as a comparison between two convention classes: encoders factoring through the joint posterior $\pi_i^t = \mathbb{P}(Y \mid X_i^t)$ over placement, and encoders factoring through current rank R_i^t . The main positive result is a Blackwell-style sufficiency lemma for π_i^t together with a team-game dominance theorem: *under* conditional independence of player information states given the final placement vector, posterior signaling matches the unrestricted optimum and dominates current-rank signaling. We then show, by exhaustive enumeration on three reduced decks, that the natural attempt to lift this dominance to hold'em via set inclusion fails: $\Sigma_{\text{curr}} \not\subseteq \Sigma_{\text{post}}$, hence the two classes are Blackwell-incomparable as abstract channels. The practical claim is therefore preserved as a conjecture for hold'em, and the abstract claim is downgraded to a cautionary lemma about which formalisations are available.

1 The Game and the Pub Argument

Components. Four players, a standard 52-card deck, and four sets of four four-sided dice (one set per street). Each set contains exactly one die showing each value in $\{1, 2, 3, 4\}$. Die faces represent predicted final placement: 1 = best, \dots , 4 = worst.

Streets and rounds. A hand has four streets (pre-flop, flop, turn, river); each street has one round of die selection.

Per-round protocol. Four dice are placed in the centre. Players act in any order, taking a die from the centre or stealing a die from another player; players may not swap or directly give dice. The round ends when the fourth centre die is taken. Each set contains one die of each value and each player ends a round with exactly one die, so the team's signal vector at any street is a permutation of $\{1, 2, 3, 4\}$.

Win condition. The team wins iff every player’s river die matches their final placement, with ties resolved softly (any assignment within a tied placement block counts as correct).

Communication. Only indirect: die choices, timing, hesitation, steals, and cross-round consistency. Pre-game discussion is permitted; mid-hand verbal communication is not.

The pub argument. Should a player’s die at street t encode their *current* hand strength R_i^t (Option 1), or their best estimate of their *final* placement at showdown (Option 2)? The first author’s intuition: Option 2, because the die’s job is to predict final placement, and current rank discards information (e.g. the four-flush draw on a heart-heavy flop) that bears on that prediction. The rest of this note tries to make that intuition precise.

2 Modelling

We model each encoder as a measurable Markov kernel $\sigma_i^t : X_i^t \rightarrow \Delta(\mathcal{A})$ with $\mathcal{A} = \{1, 2, 3, 4\}$, where $X_i^t = (h_i, B^t)$ is player i ’s information at street t , abstracting over take/steal sequencing and treating each player’s signal as an independent draw from \mathcal{A} rather than a coordinate of a permutation.¹

Let $N = \{1, \dots, 4\}$ index the players and $\mathcal{T} = \{0, 1, 2, 3\}$ the streets. All random variables are defined on a common probability space with standard Borel state spaces, so regular conditional distributions exist. Let $Y_i \in \mathcal{A}$ be player i ’s final placement, $Y = (Y_1, \dots, Y_4)$, R_i^t the standard hold’em rank computed from X_i^t , and $\pi_i^t : \omega \mapsto \mathbb{P}(Y \in \cdot \mid X_i^t)(\omega)$ a regular conditional distribution.

Definition 1 (Convention classes). Σ_{curr} : encoders factoring through R_i^t . Σ_{post} : encoders factoring through π_i^t . Σ_{full} : encoders that are arbitrary measurable kernels of X_i^t .

Definition 2 (Team value). $V(C) = \mathbb{E}[\prod_{i \in N} \mathbf{1}[\text{decode}_i(H) \approx Y_i]]$, where $a \approx b$ holds iff a and b lie in the same tied placement block of Y and H is the public die history.

The pub argument, formally, is the inequality $\sup_{\Sigma_{\text{post}}} V \geq \sup_{\Sigma_{\text{curr}}} V$.

3 Posterior Sufficiency and the CI Dominance Theorem

Lemma 3 (π_i^t is Blackwell-sufficient for Y given X_i^t). *For any measurable Markov kernel $f : X_i^t \rightarrow \Delta(\mathcal{A})$, there exists a measurable Markov kernel $g : \Delta(\mathcal{A}^N) \rightarrow \Delta(\mathcal{A})$, realised with auxiliary randomness independent of (X_i^t, Y) , such that $(g(\pi_i^t), Y) \stackrel{d}{=} (f(X_i^t), Y)$.*

Proof. π_i^t is by construction a regular conditional distribution of Y given X_i^t , so $Y \perp X_i^t \mid \pi_i^t$. Take $g(\pi)$ to be the conditional law of $f(X_i^t)$ given $\pi_i^t = \pi$, sampled with fresh auxiliary randomness. Since f ’s internal randomness is independent of (X_i^t, Y) , $f(X_i^t) \perp Y \mid \pi_i^t$ as well, and the joint $(g(\pi_i^t), Y)$ has the same finite-dimensional law as $(f(X_i^t), Y)$ by direct computation. \square

Remark 4 (Same R , different π : why π is at least sometimes richer). Player i holding $A \heartsuit K \heartsuit$ on a flop $7 \heartsuit 2 \heartsuit 9 \spadesuit$ versus the same player holding $A \spadesuit K \diamond$ on the same flop: both states have the same R_i^1 (A-high, K kicker, no pair) but the suited holding has a four-flush draw, so π_i^1 differs

¹The permutation constraint is genuine; the per-player kernel replacements used in the proofs use independent auxiliary randomness and break it. A version compatible with the constraint requires a joint replacement argument we do not develop here.

against any non-degenerate opponent range. So $\Sigma_{\text{post}} \not\subseteq \Sigma_{\text{curr}}$, and there exist states the posterior distinguishes that current rank cannot. This is the formal ghost of the pub argument.

Definition 5 (CI at street t). X_1^t, \dots, X_n^t are conditionally independent given Y at street t if $\mathbb{P}(X_1^t, \dots, X_n^t | Y) = \prod_{i \in N} \mathbb{P}(X_i^t | Y)$.

Theorem 6 (Dominance under CI). *Suppose Definition 5 holds at street t , and let $C^\circ \in \Sigma_{\text{full}}$ have encoders $\sigma_i^\circ = f_i(X_i^t)$. Let C^* have encoders $\sigma_i^* = g_i(\pi_i^t)$ with g_i from Lemma 3 and mutually independent auxiliary randomisations. Then $(\sigma_1^*, \dots, \sigma_n^*, Y) \stackrel{d}{=} (\sigma_1^\circ, \dots, \sigma_n^\circ, Y)$, hence $V(C^*) = V(C^\circ)$, and*

$$\sup_{\Sigma_{\text{post}}} V = \sup_{\Sigma_{\text{full}}} V \geq \sup_{\Sigma_{\text{curr}}} V.$$

Proof. Per-player conditionals given Y agree by Lemma 3. Both joints factor over Y : σ_i° are functions of X_i^t alone (with independent internal randomness), conditionally independent given Y by hypothesis; σ_i^* likewise factor through X_i^t via π_i^t with mutually independent U_i . The two product factorisations agree term-by-term, so the joint laws of $(\sigma_1, \dots, \sigma_n, Y)$ agree, and any decoder yields the same team value. \square

Theorem 6 is the cleanest formalisation of the pub argument. *Under CI*, signaling final placement is provably at least as good as signaling current rank, and in fact saturates the unrestricted optimum.

Remark 7 (CI fails substantively in hold'em). Conditioning on Y is structurally restrictive: $Y_i = 1$ requires i 's best five-card hand to beat every other player's best five-card hand on B^3 , which restricts each losing player's hand to those that lose to h_i on B^3 in a way that depends on h_i and creates first-order cross-player dependence. CI does not hold in hold'em and the failure is not a finite-deck perturbation.

4 The Natural Shortcut, and Why It Fails

Theorem 6 as stated requires CI. The natural attempt to bypass that hypothesis and recover an unconditional version of the pub argument is structural rather than distributional: ask whether $\Sigma_{\text{curr}} \subseteq \Sigma_{\text{post}}$ as a set of channels, i.e. whether every current-rank encoder can be replicated as a posterior encoder. This holds iff R_i^t is a.s. a $\sigma(\pi_i^t)$ -measurable function of π_i^t .

Definition 8 (Refinement hypothesis). R_i^t is $\sigma(\pi_i^t)$ -measurable, i.e. for almost every information state, the value of R_i^t is determined by π_i^t .

If Definition 8 holds, $\Sigma_{\text{curr}} \subseteq \Sigma_{\text{post}}$ trivially and the pub-argument inequality $\sup_{\Sigma_{\text{post}}} V \geq \sup_{\Sigma_{\text{curr}}} V$ holds without any CI assumption. The hypothesis is concrete and computationally checkable on reduced decks.

4.1 Computational falsification

We instantiate Cooperative Rank Hold'em on three reduced decks preserving the rule structure (4 players, 2 hole cards each, 5-card board, river only) and the relevant combinatorial features (multiple ranks, multiple suits, full hold'em hand evaluator).

For each information state (h_1, B) we compute π_1 exactly by enumerating all opponent hole-card assignments from the residual deck and tallying the placement vector Y . Two states with identical

Deck	Ranks \times Suits	Canonical states	Distinct π	π -classes with $> 1 R$	Rate
14 cards	7×2	72,072 (raw)	1,141	788	69%
16 cards	4×4	11,120	212	101	48%
20 cards	5×4	73,550	1,252	691	55%

Table 1: Failure rate of the refinement hypothesis on three toy decks.

π_1 but different R_1 falsify the refinement hypothesis. To remove obvious symmetry artifacts, we work modulo the suit-permutation action on the deck, identifying suit-isomorphic states.

A representative collision pair on the 14-card deck, hand-verified by independent enumeration:

- Board $B = (0\spadesuit, 0\heartsuit, 1\spadesuit, 1\heartsuit, 2\spadesuit)$.
- State A : $h_1 = (2\heartsuit, 3\spadesuit)$. Best hand: two pair (2s and 1s) with kicker 3. $R_1^A = (2, 2, 1, 3)$.
- State B : $h_1 = (2\heartsuit, 4\spadesuit)$. Best hand: two pair (2s and 1s) with kicker 4. $R_1^B = (2, 2, 1, 4)$.
- Joint posteriors over $Y \in \mathcal{A}^4$ are equal as 27-entry distributions: $\pi_1^A = \pi_1^B$.

The mechanism is a deck automorphism (the rank transposition $3 \leftrightarrow 4$) that fixes the board, maps h_1^A to h_1^B , and induces a measure-preserving bijection on opponent assignments preserving placement order, but does not preserve the absolute kicker value used by R_1 . A symmetry-quotient diagnostic (`symmetry_diagnostic.py`) shows that $\sim 39\%$ of same-board collision pairs on the 14-card deck are *not* explained by any deck automorphism: structurally distinct hand types (e.g. made two-pair-with-kicker vs. pocket-pair-improving-the-board) can produce identical π_1 .

The rate does not vanish as the deck grows toward hold'em proportions: going from 4 ranks to 5 ranks at fixed suit count increases it from 48% to 55%, so the failure is not a small-deck artifact. We conclude:

Proposition 9 (Incomparability on toy decks). *On the 14-card, 16-card, and 20-card four-player toys defined above, the refinement hypothesis (Definition 8) is false: there exist information states with identical π_1 and distinct R_1 . Combined with Remark 4, neither of the convention classes $\Sigma_{\text{curr}}, \Sigma_{\text{post}}$ is contained in the other as a set of channels.*

Conjecture 10 (Incomparability on hold'em). *Proposition 9 extends to Cooperative Rank Hold'em on a 52-card deck: the refinement hypothesis fails, and Σ_{curr} and Σ_{post} are Blackwell-incomparable as abstract channels.*

The toy mechanisms (rank-relabeling near board, made-vs-pocket structural coincidences) are generic, not artifacts of small rank counts; the empirical rate does not decrease with deck size in the range we can fully enumerate.

5 What Survives, and What Doesn't

The pub argument has two readings, and they fare differently.

Practical reading. *A human at the table should usually signal their best estimate of final placement rather than literal current rank.* This survives. Theorem 6 gives it under CI; Remark 4 gives a constructive same-board example where current rank throws away information that bears on placement (the four-flush draw); the conjectured extension to hold'em (Conjecture 10) does not undermine the recommendation, it only blocks a clean unconditional proof of dominance.

Formal reading. Σ_{post} Blackwell-dominates Σ_{curr} as a channel class. This does not survive in the abstract setting. Section 4 exhibits states with identical π and distinct R , so $\Sigma_{\text{curr}} \not\subseteq \Sigma_{\text{post}}$. Each class throws away information the other keeps:

- π keeps correlations between own placement and opponents' placements that R discards (Remark 4).
- R keeps absolute hand-feature distinctions (e.g. kicker rank in two-pair situations) that π collapses when the residual deck has structural symmetries (Section 4.1).

Why the practical claim still beats the formal one. The kicker-rank distinctions R preserves and π collapses are precisely the distinctions that, by construction, do not affect the placement vector Y and therefore do not affect the team's win condition. The symmetry-induced collisions are real Blackwell incomparabilities and they break the abstract dominance, but they are decision-irrelevant for this game's payoff structure. So the formal incomparability is a genuine cautionary lemma about which formalisations are available, not a counterexample to the strategic recommendation.

Practical caveats. (i) Posterior estimation at the table is cognitively expensive; a noisy posterior implementation may underperform a clean rank implementation. (ii) Current-rank trajectories preserve causal ties between belief shifts and public board events (a sudden rank drop after a third heart implies a flush threat) that a posterior trajectory can collapse into a scalar update.

6 Summary

The pub argument splits into a practical claim (signal final placement, not current rank) and a formal claim (Σ_{post} dominates Σ_{curr}). The joint posterior π_i^t is sufficient for Y given a player's information (Lemma 3), and under CI of information states given Y , posterior signaling matches the unrestricted optimum and dominates current-rank signaling (Theorem 6). CI fails substantively in hold'em (Remark 7). The natural shortcut to a CI-free version is the refinement hypothesis, which is empirically false on three toy versions of the game (Proposition 9, Table 1); we conjecture it is false on real hold'em as well (Conjecture 10). The practical claim therefore survives as a CI-conditional theorem and an unconditional conjecture; the abstract channel-dominance claim does not. Code and verification scripts accompany this note.

Acknowledgments

The variant was encountered in tabletop play; the strategic question is a long-standing pub argument among the first author's group. The formalisation, the conditional-independence theorem, the identification of the structural obstruction in hold'em, and the computational verification establishing incomparability were developed jointly with Claude Opus 4.7 (Anthropic) across multiple iteration rounds, including peer-style critique passes by additional agents. Code and verification scripts (`cooperative_rank_check.py`, `symmetry_diagnostic.py`, `scale_up_fast.py`) accompany this note.

References

- [1] David Blackwell. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24(2):265–272, 1953.

- [2] Persi Diaconis and David Freedman. Finite exchangeable sequences. *The Annals of Probability*, 8(4):745–764, 1980.