

AREA 51 / WORKING PAPER / BETA v0.2

Poker 51

One Card Declared, Five Probabilities Preserved



DECLARED / OUT

Jeffery Lyn Huckstead · Cerebral Graphix
ORCID: 0009-0007-0234-2177

Unpublished working draft v0.2 · 28 September 2026. Adds Riemann, the published betting dealer, and the jbit rules (§4–§5). Prepared for the Project 51 collection [1]. This draft has not been deposited and has no assigned DOI.

Abstract

Poker 51 is a browser experiment in five-card draw with 2♣ declared outside the deck. Removing that card preserves the exact probabilities of one pair, two pair, three of a kind, full house and four of a kind. A transitive-incidence argument proves the invariance. Exhaustive censuses give the four changes: straight and straight flush probabilities increase by 5.10638% relative to the 52-card values; flush and high card decrease by 0.0399873%. The beta enumerates each possible draw under a stated information model. Version 0.2 adds jbit betting against Riemann, a dealer using published draw rules, sampled hand-strength estimates and fixed betting thresholds. We compare the policy with two conditional game-theoretic benchmarks and report simulations. No equilibrium or optimality is established; neither a player advantage nor a human learning effect is claimed. The name is a name: the dealer uses no zeta data and makes no claim about the Riemann hypothesis.

1. A frozen experiment

The active set is a standard deck without 2♣: 51 distinct identities. The outside marker is I. Poker uses this identity set, freshly shuffled for each hand; it does not retain the Invariant fixture’s card order or decode any stored zero label. Ace is high except in A-2-3-4-5. Suits do not break showdown ties.

Deal and decision order. The player receives five cards, then the House receives five. The player keeps any subset and replaces the rest; the House then draws under the rule below. Discarded cards never return during that hand. The House’s draw count is displayed. Betting then follows (§4). Both hands are revealed at the end of every hand, including after a fold.

House rule: first applicable row	Cards kept
Straight or better	All five
Three of a kind / two pair / one pair	All cards in the repeated ranks
Otherwise, four to a flush	Those four
Otherwise, four consecutive ranks, 2-3-4-5 through 10-J-Q-K	Those four; A-2-3-4 and J-Q-K-A do not qualify
Otherwise	The highest card; deterministic suit order ♠ ♥ ♦ ♣ if needed

The House uses its own initial hand only. It does not inspect the player’s cards or stamps. This is a declared deterministic policy, with no optimality claim.

2. Why five probabilities are preserved

Proposition (uniform deletion under transitive incidence). Let E be a class of five-card hands preserved by a permutation group acting transitively on the 52 card identities. If $T = |E|$ and I is the number of its hands containing one specified card, then every card has incidence I. Double-counting card-hand incidences gives $52I = 5T$. Deleting any specified card therefore retains $47T/52$ hands of E.

$$C(51,5) / C(52,5) = 47/52, \text{ so } (T - I) / C(51,5) = T / C(52,5).$$

For each of the five repeated-rank categories, arbitrary rank relabellings and independent suit permutations within each rank preserve the category and act transitively on identities. The proposition applies. It does not apply to the remaining four exclusive poker categories: their definitions depend on rank order, common suit, or exclusions of those events.

Exact census

Exclusive category	52 cards	51 cards	Relative probability change
Straight flush	40	38	+5.10638%
Four of a kind	624	564	0 (exact)
Full house	3,744	3,384	0 (exact)
Flush	5,108	4,615	-0.03999%
Straight	10,200	9,690	+5.10638%
Three of a kind	54,912	49,632	0 (exact)
Two pair	123,552	111,672	0 (exact)
One pair	1,098,240	992,640	0 (exact)
High card	1,302,540	1,176,825	-0.03999%

Denominators: $C(52,5) = 2,598,960$; $C(51,5) = 2,349,060$. Royal flush is included in straight flush. Straight excludes straight flush; flush excludes straight flush; high card excludes both straight and flush. Changes are relative percentages, not percentage points.

The four changed categories. The 2 belongs to two of the ten straight rank sets. For each such set, fixing $2\clubsuit$ leaves $4^4 - 1 = 255$ non-flush suit choices. Thus 510 ordinary straights and two straight flushes disappear: both retain 19/20 of their hands. Their probability ratio is $(19/20)/(47/52) = 247/235$.

The ordinary flush count loses $C(12,4) - 2 = 493$, giving 4,615. For any non-straight five-rank set, non-flush suit assignments are 255 times flush assignments in either deck: 1,020/4 in the full deck or when the 51-card rank set omits the 2, and 765/3 when the 51-card rank set includes the 2. Hence high card equals 255 × flush in both censuses. Their common relative change is exactly -24/60,019. The straight change depends on which rank is removed; the five invariances hold for deletion of any card.

3. What the exact draw panel means

Let H be the player’s initial five-card hand and K a specified subset to keep. Put $k = 5 - |K|$. Before using any stamp observation, the remaining 46 identities are exchangeable from the player’s seat. Although five are physically in the House’s initial hand, those identities are unknown and the House has not acted. Marginalizing that unknown hand makes every k -subset of the 46 equally likely as the player’s replacement set. All five initial player cards remain excluded, including discards. Write U for those 46 unseen identities.

For category j : $P(j \mid H, K) = \#\{D \subset U: |D| = k, \text{category}(K \cup D) = j\} / C(46,k).$

Cards drawn k	0	1	2	3	4	5
Possible draws	1	46	1,035	15,180	163,185	1,370,754

The engine visits every completion; it does not sample. The panel reports the category of the player’s final hand. “Higher category” excludes improvements within one category, such as a better kicker. It is neither a showdown win probability nor an expected-return calculation.

A replayable example

Start with $K\spadesuit K\heartsuit 7\diamondsuit 4\clubsuit 9\spadesuit$ and keep the kings. There are $C(46,3) = 15,180$ completions. The exact distribution is:

Final category	Completions	Probability
One pair	10,796	71.119895%
Two pair	2,412	15.889328%
Three of a kind	1,772	11.673254%
Full house	156	1.027668%
Four of a kind	44	0.289855%

All other categories have zero completions in this example. The counts sum to 15,180.

Stamps change the information, not the cards

A House card whose stamp was earned in Memory 51 displays that mark on its back. A player who recalls the mapping may infer the identity; the House ignores marks. The odds panel deliberately ignores all stamp observations, including whether a card is marked. It is therefore a baseline, not the full posterior distribution of a player using stamps.

The House’s later draw count is another observation of its fixed policy. A future win-probability model must condition jointly on that policy, the draw order, available marks and unseen discards. Simply maximizing the chance of reaching a higher hand category need not maximize the chance of beating the House. Quantifying the decision value of recall remains an open task for this beta.

4. Riemann: the published dealer

Both players ante 1 jbit. After the draw the player acts first: bet $B \in \{1, 2, 4, 8\}$ jbits, or all remaining jbits when fewer than 8, or fold the ante. Riemann folds, calls, or raises by B . The player may then call, re-raise by B once, go all in for less than the call, or fold; after a re-raise Riemann calls or folds. At showdown the better hand wins what both matched, $M = \min(P, R)$, where P and R are what each put in; any unmatched excess returns to its owner. A fold forfeits only what the folder put in. If the ante took the player's last jbit, the hand is shown for the antes.

What Riemann observes: its own final hand h , the player's draw count k , the bet B , and whether the player has jbits left to answer a raise. It does not read stamps or the player's cards.

4.1 The read

A belief table stores 1,000 quantiles per row: for $k \in \{0, \dots, 4\}$, final-hand strengths under the House draw rule; for $k = 5$, a separate uniformly random five-card baseline. The House never draws five. The table is built from 19,257,529 simulated deals with seed 51 (SHA-256 5858deb92a383c87..., rebuilt bit for bit by the verifier). Let $q_k(h)$ be the fraction of quantiles below h , with ties counted half. The House rule makes some rows sharp: a player who stood pat holds a straight or better, one who drew 2 holds trips or better, one who drew 3 holds a pair or better — each verified in every quantile.

A player need not follow the House rule: standing pat on nothing is legal. Riemann therefore blends the draw-conditioned row with the random-hand baseline:

$$r = 0.6 \cdot q_k(h) + 0.4 \cdot q_5(h).$$

This read is an estimated equity score, counting ties half, not an exact conditional win probability. The table is sampled and quantized; it does not remove cards known to Riemann or learn the player's strategy. The blend and subsequent range rescaling are modelling choices, not a Bayesian posterior.

4.2 Bets narrow the range; pot odds decide

A bet is read as strength. Riemann keeps the top fraction $f(B)$ of the player's range — $f = 0.55$ for $B \geq 8$, 0.70 for $B \geq 4$, 0.85 for $B \geq 2$, and 1 otherwise — and rescales $n = \text{clamp}((r - (1 - f)) / f, 0, 1)$. If a raise is not selected, calling B produces a final pot of $2 + 2B$, giving the call threshold

$$\text{call if } n \geq B / (2 + 2B) \quad (25\% \text{ at } B = 1; 44.44\% \text{ at } B = 8).$$

Riemann first value-raises when $n \geq 0.70$, or bluff-raises when $n < 0.10$ and the FNV-1a hash of its sorted card codes is divisible by 8 — one of eight hash classes, publicly replayable. It never raises a player with no jbits left. Facing a re-raise it halves f and calls if $n \geq B / (2 + 6B)$, the break-even for a further B into a pot of $2 + 5B$.

Opening decision: first applicable row	Rule
Raise (value)	$n \geq 0.70$ and the player still holds jbits
Raise (bluff)	$n < 0.10$, hash class $0 \bmod 8$, and the player still holds jbits
Call	n at or above $B / (2 + 2B)$, when neither raise applies
Fold	Otherwise
After a re-raise	call if $n \geq B / (2 + 6B)$ with f halved, else fold

5. Game theory: what the thresholds are, and are not

The **call thresholds** are pot-odds break-evens applied to an estimated equity score. The value-raise threshold, bluff cutoff, hash class, random-hand weight and range fractions are design choices. The bet sizes are the Blackjack 51 stakes. No Nash equilibrium or optimality is established. Two elementary benchmarks help frame the tests, under explicit assumptions.

Bluffing frequency. After Riemann raises by B, the pot is $2 + 3B$ and the player’s call costs B. For a polarized range in which a caller beats every bluff and loses to every value hand, with no ties, indifference requires a bluff share $B / (2 + 4B)$: 16.7% at $B = 1$ and 23.5% at $B = 8$. Against a House-rule player who bets every hand, the policy labelled 1.6% and 16.3% of its raises as bluffs. These labels describe the dealer’s score and hash test; they are not the caller’s actual equity. Comparing those frequencies alone does not establish under-bluffing or a profitable response.

Minimum defence. For a one-shot, zero-equity bluff into the two-ante pot, a continuation frequency of at least $2 / (2 + B)$ removes automatic bluff profit: 66.7% against 1 and 20% against 8. Against the same every-hand bettor Riemann continued 81.6% and 44.3%. Those aggregate frequencies exceed the benchmarks, but do not prove that arbitrary betting strategies lose in this game. The following simulations measure particular strategies directly.

5.1 Simulation and search (status C)

Player strategy (200,000 hands each, seed 7)	jbits / hand
House draw; always bet 1; always call	-0.418 ± 0.005
House draw; always bet 8; always call	-2.903 ± 0.021
House draw; always bet 8; always re-raise	-4.330 ± 0.029
House draw; bet 8 with two pair or better, 1 with a pair, fold less; call raises with trips or better; re-raise with a straight or better	-0.068 ± 0.009
House draw; fold unless two pair or better, then bet 8 and always call	-0.269 ± 0.010
Draw-count bluff: stand pat on anything; bet 1; fold to a raise	-0.555 ± 0.004
Draw-count bluff: stand pat on anything; bet 8; fold to a raise	-0.433 ± 0.008
Draw-count bluff: stand pat on high card, else House draw; bet 1; call raises with two pair or better	-0.329 ± 0.004

Results are mean ± one Monte Carlo standard error. Each hand starts with an unlimited player bank; these are per-hand returns, not finite-bank session forecasts. The simulations use no stamps.

A search of 480 parameter combinations (including duplicates) — six draw rules including standing pat or drawing one on weak hands, four betting floors, two sizes, and call and re-raise floors — screened each on 5,000 hands (seed 21) and re-tested the best 12 on 150,000 (seed 99). Every reported retest mean was negative; the best was -0.088 ± 0.005 jbits a hand (draw: honest; bet 1 with a pair or better, else fold; call a raise with two pair or better; re-raise with a straight or better). This is evidence about the tested configurations only, not a best-response bound. A player who reads stamps holds information the search did not model.

Claim	Status
Sampled decisions match the published rule; sampled table paths conserve jbits without overdraw	A-CAT (verifier; COMPUTED in Area 51 words)
The belief table regenerates from seed 51	A-CAT (bit-for-bit rebuild)
Reported player means are negative in the simulations and search retest	C (seeded simulation and search)
Riemann is optimal, or an equilibrium strategy	not claimed

6. Verification, limits and next questions

Evidence in this package	Scope
Supplied verifier: 25/25 checks pass	Two evaluators agree on all 2,598,960 standard-deck hands and all 2,349,060 active-deck hands; rule, draw and sampled deal checks pass.
Independent C++ census	Enumerates standard-deck hands and filters out 2♣; reproduces every count above without importing the game engine.
Independent draw and showdown replay: 35/35 checks pass	Every category for all 32 hold masks of the example hand agrees; 200,000 deterministic sampled showdowns agree with a separate full ranking key.
v0.2 betting verifier: 20 checks pass	Belief-table hash and structure; 100,000 Riemann decisions recomputed from the published rule; 300,000 random hands conserve jbits exactly with no overdraw; all-in, re-raise and check paths exercised.

The all-mask draw fixtures contain 2,349,060 completions in total, a partition check: each final five-card set uniquely determines which original cards it retains. This does not exhaust all starting hands. The supplied showdown test checks consistency with the engine’s own key; the added replay independently constructs that key. Sampled showdowns are not an exhaustive strategy proof.

Reproduce from the site root

```
node area-51/poker-51/verify-poker51-v0-1.mjs
node area-51/poker-51/verify-independent-v0-1.mjs
node area-51/poker-51/verify-poker51-v0-2.mjs --regenerate
node area-51/poker-51/simulate-riemann-v0-2.mjs
node area-51/poker-51/search-riemann-v0-2.mjs
```

The separate C++ source includes its compile/run command and regenerates INDEPENDENT_CENSUS_v0_1.json. Both verifiers write JSON receipts next to the game. The evidence manifest records source hashes. PP288 [3] supplied a corroborating census report; this draft’s numbers are backed by the new local replay, not solely by that report.

Implementation and interpretation limits

The beta uses browser cryptographic random words, rejection sampling for bounded integers and Fisher-Yates shuffling. Uniformity follows under the stated uniform-word model; the censuses do not certify browser entropy. The House hand exists in client memory and can be inspected. This is a research exhibit, not a secure adversarial game.

Hands are played for jbits in the shared browser-local Area 51 bank. Each ante, bet, call, raise and settlement is posted with its own receipt id in the same atomic step that saves the hand, so a reload resumes the hand and a repeated click cannot post twice. The table keeps the last 100 hands for its receipt. jbits have no monetary value. Poker does not alter earned stamps. The game is not a controlled study of cognition. No result here concerns zeta zeros or RH.

The animated Riemann avatar, thinking bubble and short table comments are presentation. A decision pause lasts 1–2.4 seconds, increasing with the pot. Original dialogue and attributed historical excerpts are distinguished on the game’s Rule Card. Comments depend only on public actions and revealed outcomes; neither comments nor delay alter the betting policy, random deal or jbit settlement. Riemann is a rule-based simulated opponent, not a language model.

Open research questions. What is the best response to Riemann, and how far is Riemann from an equilibrium of this betting game? What policy maximizes showdown success against this exact House rule? When does a decoded stamp change the best hold decision? How much does that decision improvement depend on the House’s observable draw count? Answering these requires a stated objective and a conditioned opponent model; neither an optimal policy nor a quantified memory advantage is claimed.

References and draft citation

[1] Huckstead, Jeffery. Project 51: Reading the Record Across Rules. Zenodo record, 28 September 2026.

doi:10.5281/zenodo.23004789. Intended collection destination; this existing DOI does not identify the Poker draft.

[2] Huckstead, Jeffery Lyn / Cerebral Graphix. Poker 51 beta v0.1. Supplied site delta, engine and verifier, 28 September 2026; v0.2 adds poker51-riemann-v0-2.js, poker51-table-v0-2.js and their verifier. Game route: cerebralgraphix.com/area-51/poker-51/.

[3] PP288: daily crawler return, 28 September 2026, §C-4. Author-supplied constructive-seat research report, retained under the game’s provenance/ directory. Other PP288 numerical claims are outside this paper’s scope.

Cite this draft: Huckstead, Jeffery Lyn. (2026). Poker 51: One Card Declared, Five Probabilities Preserved. v0.2, unpublished working draft, 28 September 2026. Cerebral Graphix. © Jeffery Lyn Huckstead. CC BY 4.0.