Skill has an edge.

Under your baseline, current Tactician teams beat a coherent simpler strategy more often than chance in both hands and full matches. The advantage is meaningful, but it does not remove deal-driven swings.

Measure	Result	95% interval
Four-hand match wins	**60.5%**	54.5% to 67.0%
Single-hand win score (ties half)	**56.2%**	54.5% to 57.9%
Mean score edge per hand	**+13.9 points**	+10.7 to +17.1
Mean score edge per match	**+53.1 points**	+31.8 to +73.9

A CONSISTENT EDGE ABOVE THE 50% REFERENCE

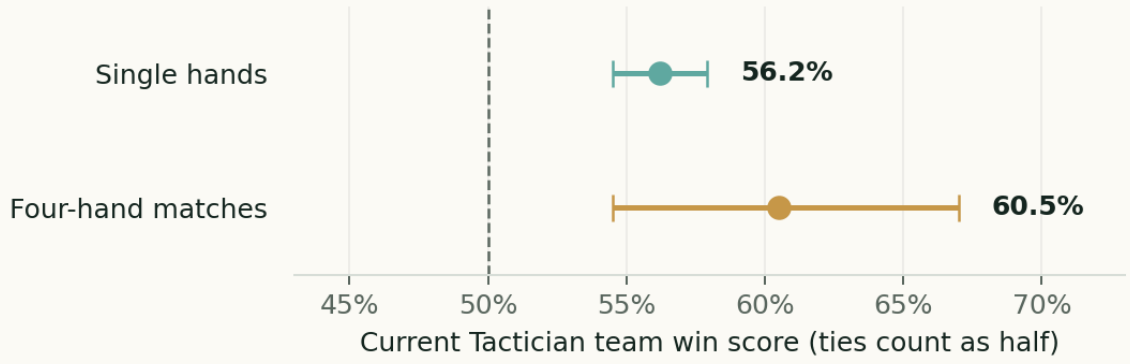

Intervals resample entire deal or paired-match groups. The 9,600 hand replays are not treated as 9,600 independent deals. Strict single-hand wins were 55.0%, with 2.4% ties; all 200 full matches had a winner (121 current-strategy wins, 79 losses).

WHAT THIS SUPPORTS

Decisions can improve results on matched cards. This is evidence that the baseline supports a measurable strategy advantage in this model. It does not tell us which part of that advantage comes from passing, discard choices, Joker decisions or future planning.

44.4% of reference losses became wins when the same team retained its Tactician policy and the opponents switched to the simpler policy (95% interval: 42.0% to 46.9%). Initial cards, seats and stock were held fixed; later play and draws could diverge.

The deal still has leverage.

With Tactician on all four seats, the median hand gap was 55 points and the 90th-percentile gap was 130. A 100-point-or-more gap occurred in 22.7% of hands. Close scores alone would not prove skill.

FINAL HAND GAPS / ALL TACTICIAN

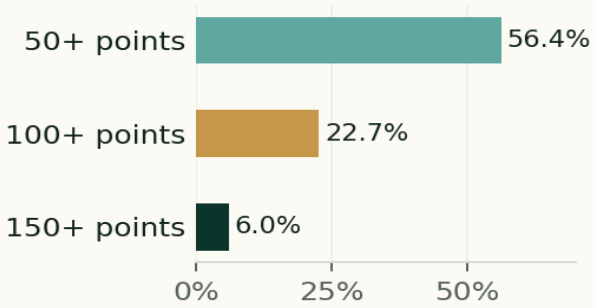

OTHER CURRENT PROFILES

Profile	Win score
Tactician	**57.2% (53.9-60.5)**
Bob	**58.0% (55.0-61.0)**
Runner	**58.7% (54.5-62.6)**
Setter	**58.7% (55.3-61.9)**

Same 24 deals and 24 seat assignments per profile. Ties count as half; parentheses show 95% intervals. All four profiles had a positive edge. These overlapping intervals do not establish a best bot.

OPENING ADVANTAGES MATTER, WITHOUT FIXING THE WINNER

Initial Joker advantage and legal starter-pair advantage were moderately associated with final margins (r = 0.38 and 0.37). These are incomplete, overlapping proxies for hand quality, not a percentage of the game attributable to luck. Every one of the 400 opening hands won and lost at least once across changed partners and seats; that also changes turn order and later draws.

METHOD AND COMPLETION CHECK

Main hands: 100 independent shuffled deals x 24 seat assignments x four policy setups = 9,600 runs. Setups: all Tactician, all basic, and each color taking the Tactician team against basic opponents. First player balanced across deal seeds.

Matches: 100 independent paired four-hand sequences, swapping policy colors, with dealer rotation and cumulative totals fed back into decisions. 200 matches / 800 additional hand runs. Other-profile checks add 3,456 runs on the first 24 main deals.

Basic policy: legal immediate-point/card-shedding plays; takes discards that improve its immediate play; retains useful own-hand pairs and Jokers on discard. No secret passes, opponent denial or future-turn planning. It is not a validated human novice model.

Validation: all 13,856 evaluated hand runs finished. Moves were checked by the live rule engine; 54 unique cards conserved; one Joker per combination; scores reconciled; all three Last Lap seats completed final turns. A 304-hand pilot reproduced exactly. Nine-card starts; negative Jokers off.

Practical limit: this tests your baseline, not whether it is better than other rule setups. Bot results support a skill signal; they do not establish the balance of human skill and luck. The useful next question is which specific partnership decisions produce the advantage.