SNATCH rummy

Testing a Smarter Bob

Previous versus updated strategy | A matched-deal simulation article

A legal move is only the beginning

Watching the bots play raised a useful question: did they understand what a negative joker was doing to the score, or were they simply finding somewhere legal to put it? A wild card can complete a combination while still costing its team points. The destination matters.

What changed

All four profiles received a shared baseline for evaluating negative-joker plays. Bob also received new Secret Pass and discard planning based on his own hand and public table information. Those changes gave us a testable question: would updated bots outperform their previous versions under the same game rules?

A fairer comparison

We played updated and previous partnerships against each other, swapped their seats on matched deals, and tested every current combination of Joker and final-turn rules. The completed study contains 384 four-hand matches, or 1,536 hands. It uses 12 initial deal-seed blocks, reused across profiles and rules; the large hand count should not be confused with 1,536 independent experiments.

Bob provided the clearest signal

Updated Bob partnerships recorded 59 wins, 35 losses, and 2 ties in 96 matches, with a mean advantage of 62.4 points per match. The estimated score interval was +24.8 to +100.3 points. This is a promising, exploratory result for Bob, not a claim that he has mastered the game.

The whole lineup tells a more cautious story

Across all profiles, updated bots averaged 20.0 more points per match, but the uncertainty interval included zero. Tactician, Runner, and Setter did not show a clear individual improvement in this sample. Updated bots made fewer direct negative-joker placements into their own secured cards, but also fewer attacks on opponents. Better judgment is not simply more aggression.

How to read the figures

The following pages separate match performance from card-placement patterns. Interval bars show uncertainty from resampling deal-seed blocks. The method and limitations are included so the findings can be checked. The results support continued testing, and a measured story about Bob, rather than a blanket claim that every bot is stronger.

Did the updated bots play better?

384 four-hand matches | 1,536 hands | 192 mirrored pairs | 12 deal-seed blocks

Overall: updated win share 53.4% (95% interval 47.8-58.3%). Mean score advantage +20.0 points per match (95% interval -0.1 to +40.6).

Bob: 59 wins, 35 losses, 2 ties; +62.4 points per match.

Score margin by bot profile

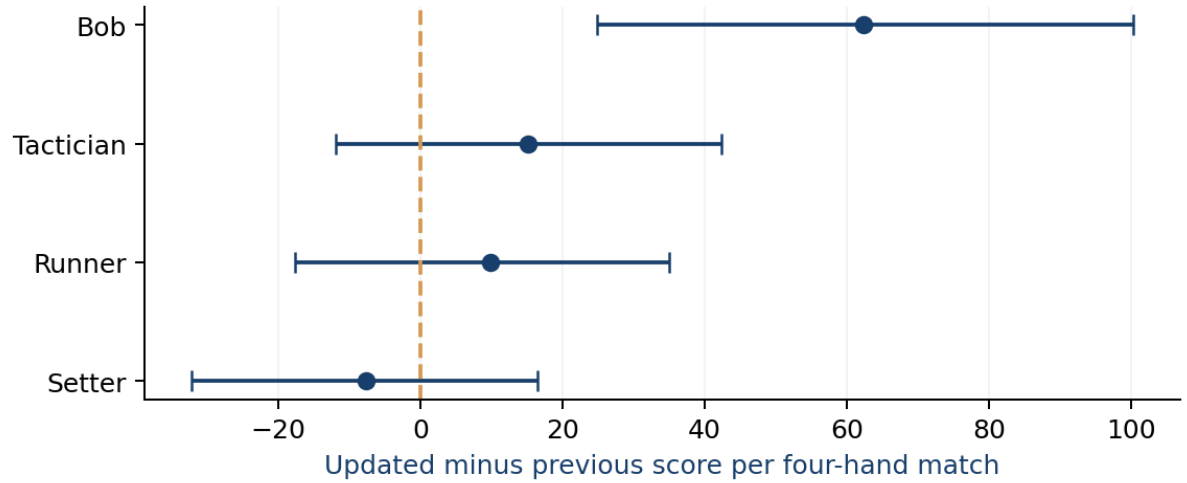

Dots show the mean. Lines show 95% intervals from resampling deal-seed blocks.

Win share by bot profile

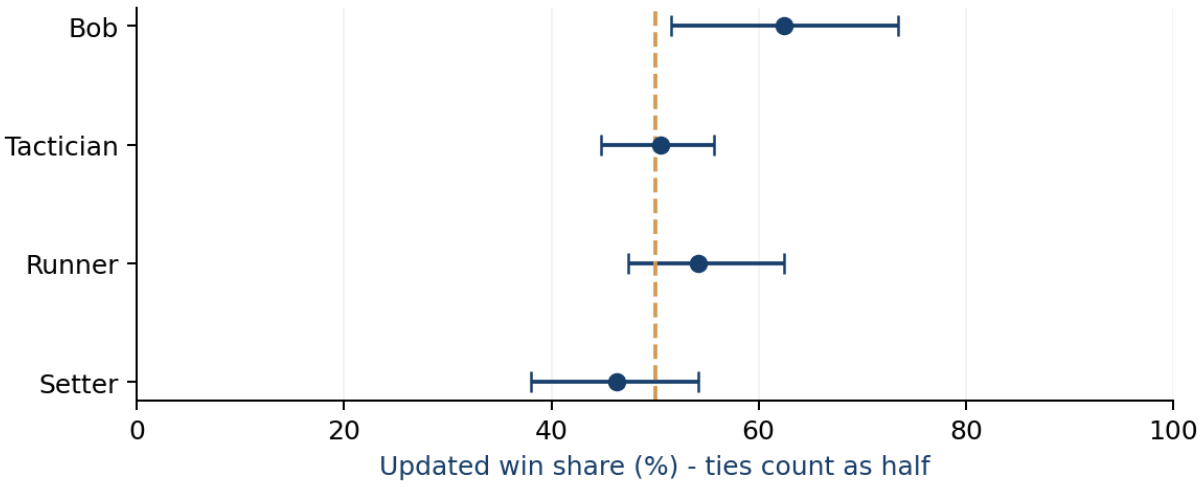

Dashed reference: 50%. Each profile contributes 96 matches; ties count as half a win.

Where did the negative jokers go?

Same deals and rules. Previous strategy in gold; updated strategy in blue.

Direct placements

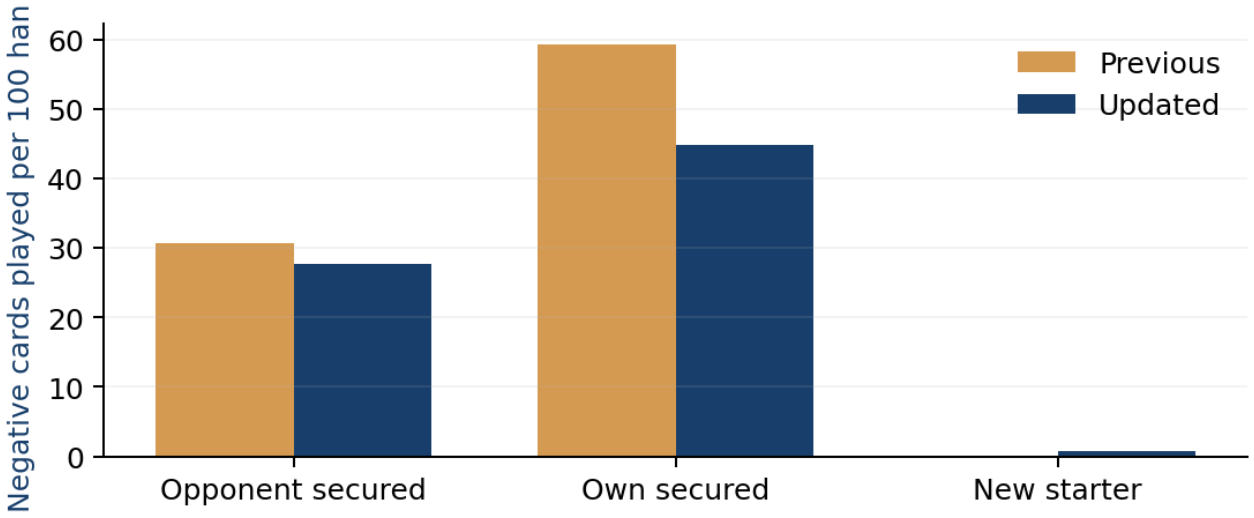

Placements count negative cards added from hand. Own-secured placements include completions and extensions; swaps are counted separately. These are action frequencies, not a measure of whether each play was best.

Penalty remaining at the end of a hand

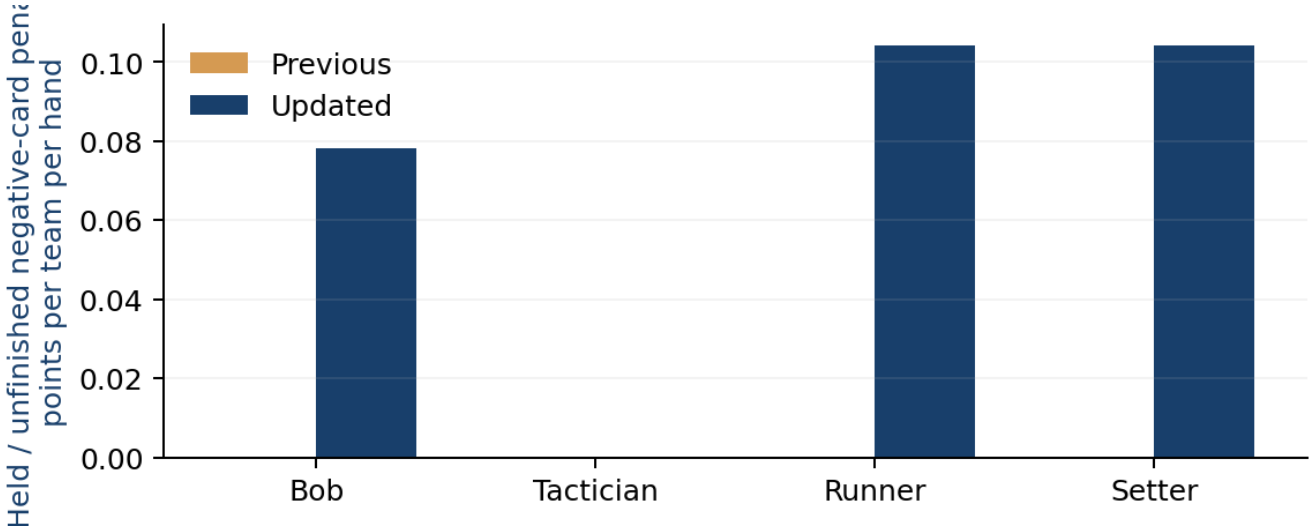

Each held or unfinished negative joker contributes a 10-point penalty. This chart excludes the value of other cards and negative jokers already in secured scoring.

What the comparison can support

Rule results, opportunity counts, and practical limits.

Updated results by rule combination

Fixed / Last Lap: 52.1% win share; +0.2 points per match.

Fixed / Snatcharama Last Lap: 52.1% win share; +10.5 points per match.

Permanent / Last Lap: 53.6% win share; +27.4 points per match.

Permanent / Snatcharama Last Lap: 55.7% win share; +41.7 points per match.

Available attacks are not necessarily the best play

Previous bots chose an immediate negative-joker attack at 471 of 727 eligible play decisions; updated bots did so at 425 of 1,454. Opportunities can repeat within a turn. A skipped attack can be a sound decision, so these counts are not an error rate.

Method

Each table used one bot profile in all four seats. Updated bots occupied one partnership and previous bots the other, then swapped partnerships on the same deals. Four hands rotated the dealer. All four profiles were tested under Fixed and Permanent Jokers, each with Last Lap and Snatcharama Last Lap. Negative jokers were on, exposure was at the end of the next turn, and each player started with nine cards.

The previous source predates the latest shared negative-joker baseline and Bob’s pass/discard planning. Both versions use identical legality and scoring. The isolated harness exposes the existing move selector and suppresses cosmetic comments; it does not change bot choices or game rules. Initial deals were verified by card-identity hashes. Strategies can cause subsequent stock recycling to diverge.

Limits and interpretation

The overall interval includes no score advantage; this sample does not establish a general improvement. There are only 12 independent deal-seed blocks, reused across profiles and rules. The 95% intervals use 10,000 seed-block bootstrap resamples. Profile intervals are exploratory and not adjusted for multiple comparisons. These are bot-only, homogeneous-profile partnerships; results do not establish human-play strength or the effect of a single bot in a mixed lineup.