

ODDSOFBLACKJACK

Technical Whitepaper

Model, conventions and evidence of the OOB engine

Version: V1 (final form, for publication approval)

Date: 7 October 2026

Engine build: 00B-ENG-b0681ede84a3

Methodology revision referenced: Methodology v1.0 (6 September 2026), with the additions described in this document

Status: Each piece of evidence in this document carries its own scope and status (Section 8). No global validation verdict is given for the model as a whole. Independent third-party validation has not been performed.

Contents

0. About this document	3
1. What OOB does, and what it does not	4
2. Game state and rules	5
2.1 The recorded composition	5
2.2 Rule parameters	5
2.3 Deck count and shoe depth	6
2.4 Which rule combinations have been checked	6
3. Expected Value conventions	7
3.1 Units	7
3.2 Available actions	7
3.3 Selection and precision	8
3.4 Exact ties	8
3.5 Player blackjack and surrender under ENHC	8
4. The evaluation model in production	9
4.1 Card probabilities	9
4.2 Conditioning and depletion conventions	9
4.3 Determinism	10
5. How each action is evaluated	11
6. Pre-bet EV and its calibration	12
6.1 What the pre-bet EV is	12
6.2 How it is formed	12
6.3 Why a calibration exists	12
6.4 What is public and what is not	12
6.5 Domain of the calibration	13
6.6 Evidence for the calibration	13
7. Initial bet sizing	15
7.1 Historical ramps V1 and K1 (simulation accounts)	15
7.2 MOV with a fixed unit (analysed retrospectively)	15
7.3 MOV sized by confirmed balance (the application)	15
8. Evidence: checks performed and their scope	17
8.1 Agreement with the source spreadsheet — porting check	17
8.2 Published house-edge values by deck count — historical calibration reference, incomplete provenance	17
8.3 Agreement with an internal reference program — internal independent-code check	18
8.4 Ten published reproducible states — reproducibility check	18
8.5 Controlled comparison with Basic Strategy, decision level — partial comparison	18
8.6 Audit of the production pre-bet EV — incomplete audit	19
8.7 Build identification	19
8.8 Independent validation	20
9. Evidence from simulation	21
9.1 The simulator	21
9.2 Earlier 300,000-card tracks	21
9.3 The first million — retrospective simulation	21
9.4 What the simulation evidence does not contain	22
10. Limitations	23
Appendix A. Glossary	24
Appendix B. Evidence labels	24
Appendix C. Provenance	25
External references	25

0. About this document

This whitepaper describes what the OddsOfBlackjack (OOB) engine calculates, under which conventions and approximations, and what evidence exists for those calculations. It also states what the evidence does not establish.

It does not disclose source code, operational pseudocode, caching, optimisation or infrastructure. Those remain proprietary and are kept server-side. The document allows the model and the published examples to be inspected. It does not demonstrate that every production result can be reconstructed in full from the public description (Section 6.4).

OOB is decision-support software. It does not operate a casino, accept wagers, or guarantee wins or financial results.

How to read the evidence labels. Every check in Section 8 is labelled by what it actually shows:

Label	Meaning
Porting check	Agreement with the spreadsheet model the engine was ported from. Shows faithful porting, not correctness.
Calibration-consistency check	Agreement with published values that were also used to set or check the calibration. Not validation on independent data.
Historical calibration reference — incomplete provenance	Published values recorded during calibration whose settings, result type or date of reading were not fully recorded. Kept as a record; not a fully verified direct comparison.
Internal independent-code check	Agreement with a separately written internal reference program on selected states. Internal, not third-party.
Reproducibility check	The same inputs produce the same outputs across runs or environments. Says nothing about correctness.
Partial comparison	A comparison against an external reference that has not been approved as a reference, or with an incomplete scope.
Incomplete audit	An audit whose design exists but whose execution has not reached a verdict.
Retrospective simulation	A result computed on a dataset that was also used to choose the rule being reported.

1. What OOB does, and what it does not

OOB takes the state of a blackjack game as entered by the user:

- the cards exposed so far,
- the player's hand or hands,
- the dealer's upcard,
- the active table rules.

From that state it returns:

1. the Expected Value (EV) of each action available in the current hand, and the action with the highest EV;
2. strategy tables (hard, soft, pair splitting) for the current composition;
3. a pre-bet EV for the next round, computed from the composition before the round is dealt;
4. a suggested initial bet, sized from the pre-bet EV and the balance the user confirms (Section 7.3).

All calculation runs on the server. The browser receives only results.

OOB does **not**:

- predict the next card;
- guarantee that a higher-EV action wins a given hand or session;
- guarantee profit from the suggested bet.

All outputs depend on the user entering the cards and rules correctly.

Some features described on the public landing page are not yet available in the application, for example Risk of Ruin and variance metrics. They are outside the scope of this document.

2. Game state and rules

2.1 The recorded composition

The engine works on ten rank counts: 2, 3, ..., 9, ten-valued cards (10/J/Q/K together) and ace. The **recorded composition** is the starting shoe minus every card the user has entered as exposed: the player's cards, the dealer's cards and any other cards removed from play. The dealer's hole card is not part of it until it is exposed.

The recorded composition is updated after each entered card and is the single input from which all probabilities are derived. How it is conditioned during the evaluation of a hand is described in Section 4.

2.2 Rule parameters

Parameter	Values supported
Decks	4, 5, 6, 8 (1-3 not offered)
Hole-card rule	US peek, or European No Hole Card (ENHC)
Dealer soft 17	stands (S17) or hits (H17)
Double	any two cards, 9-11 only, or 10-11 only
Double after split (DAS)	allowed / not allowed
Split	maximum total number of hands, 2 to 4
Resplit aces	allowed / not allowed
Hit split aces	allowed / not allowed
Late surrender	allowed / not allowed
Blackjack payout	3:2 or 6:5

The split parameter is a total, not a count of resplits. A value of 2 means a single split and no resplitting. A value of 4 allows up to three splits.

Default configuration (the game shown by default in the application and used by all simulations in this document):

- 8 decks;
- ENHC;
- S17;
- double on any two cards;
- no DAS;
- a single split (2 hands maximum);
- no resplitting of aces, no hitting split aces;
- no surrender;
- blackjack pays 3:2.

2.3 Deck count and shoe depth

The engine treats an N-deck shoe as an 8-deck shoe depleted to the same composition. The two give identical results (Section 8.2). Accordingly, the deck count enters only through the composition.

The practical limit is the number of **cards remaining**, not the deck count. The pre-bet EV calibration (Section 6) was measured down to 104 remaining cards. Below that depth, or outside a defined ten-density band, the uncalibrated value is shown instead. At a 50% cut, the depths reached are 208 (8 decks), 156 (6), 130 (5) and 104 (4) cards.

2.4 Which rule combinations have been checked

The interface allows many rule combinations. The checks in Section 8 cover:

- the default configuration (most checks);
- six full-shoe configurations against the source spreadsheet;
- three Wizard of Odds rule points recorded during calibration, with incomplete provenance;
- 4, 6 and 8 decks on a full shoe.

A complete public table of verified rule combinations, as opposed to combinations the interface merely allows, does not yet exist. (*Evidence gap L7.*)

3. Expected Value conventions

3.1 Units

Every action EV is the **net result per original bet** and includes any additional wager the action puts at risk:

Action	Wager at risk	Typical range
Stand	1 unit	-1 to +1
Hit	1 unit	-1 to +1
Double	2 units	-2 to +2
Split	2 or more units (all resulting hands)	-2 to +2 for one split
Surrender	½ to 1 unit	-1 to -½

A Double EV near -1.1 therefore reflects two units at risk. Split EV is the sum over all hands created by the split, expressed against the original bet.

Wizard of Odds uses the same normalisation per initial bet. A shared normalisation is necessary for comparing OOB with an external tool, but it is not sufficient. A direct comparison of two values is meaningful only when all of the following coincide:

- the rules;
- the shoe composition;
- the player's hand and the dealer's upcard;
- the conditioning conventions (Section 4.2) — for example, whether the player's own cards are removed before the dealer's outcomes are computed;
- the set of permitted actions, including the playing strategy assumed for later decisions in the hand.

When any of these differs, a difference between the two values does not by itself indicate an error in either.

3.2 Available actions

Only actions available in the current state are compared:

- Hit is available below 21.
- Double is available on two-card hands only, subject to the double restriction, and after a split only if DAS is allowed.
- Split is available on pairs, subject to the hand limit and the ace rules.
- Surrender is available only on the first two cards of the initial hand, never after a split, and only if the rules allow it.

An unavailable action has no value. This is distinct from an action whose computed value is zero.

3.3 Selection and precision

All values are held in IEEE 754 double precision. The highest-EV action is selected from the **unrounded** values. Display rounding to four decimal places happens only at presentation. As a result, two actions may look equal on screen while their internal values differ.

3.4 Exact ties

When two available actions have identical values in every significant digit, the current implementation selects deterministically, by comparing (value, action name) pairs. This resolves equal values by action name in reverse alphabetical order. **The ordering is deterministic but arbitrary.** It was not chosen for a mathematical reason and is documented because it is observable. Exact ties at full precision are rare.

3.5 Player blackjack and surrender under ENHC

- **Player blackjack.** Under ENHC, a player blackjack is a push against a dealer blackjack. Its value is therefore $\text{payout} \times (1 - P(\text{dealer blackjack}))$.
- **Surrender.** Under ENHC, surrender forfeits half the bet only if the dealer does not then draw blackjack. If the dealer does, the full bet is lost. Hence:
$$\text{Surrender EV} = -0.5 - 0.5 \times P(\text{dealer blackjack}).$$
- **The denominator.** $P(\text{dealer blackjack})$ is computed with the number of remaining cards n as denominator. The hole card is one of the n remaining cards; the visible cards were already removed when the composition was recorded.

A previous version used $n - 1$, removing the upcard a second time. This was corrected on 6 September 2026 (Section 8.4).

4. The evaluation model in production

This section describes the calculation path the application actually uses. It separates three things that are easy to confuse:

- **(A)** the recorded composition, described in Section 2.1;
- **(B)** the conditioning conventions used while evaluating the actions of a hand, described below;
- **(C)** the calibration applied only to the pre-bet EV, described in Section 6.




Figure 1. Conceptual overview. The recorded composition feeds two separate calculations: the evaluation of the actions in a hand, and the pre-bet EV of the next round. Only the pre-bet EV is calibrated, and only inside its domain. The initial bet is derived from the pre-bet EV and, in the application, the confirmed balance. The figure shows concepts, not the internal architecture.

4.1 Card probabilities

The probability of drawing rank r is its remaining count divided by the total remaining:

$$P(r \mid S) = n_r / N$$

where n_r is the count of rank r and N is the total number of cards in the composition S used for that step. What matters for accuracy is **which** composition is used at each step. That is the subject of 4.2.

4.2 Conditioning and depletion conventions

The production path evaluates each decision from the recorded composition, under fixed conventions. Within the evaluation of a single hand, the composition is **not re-derived after every hypothetical card**:

1. **The dealer's outcome distribution.** For each possible upcard, the engine computes the probabilities of the dealer finishing on 17, 18, 19, 20, 21 or busting. These are computed from the recorded composition, and the dealer's own successive draws deplete the composition between them. The distribution is **not** updated for the hypothetical cards the player draws while the hand is being evaluated (hits, double cards, cards dealt to split hands).
2. **The player's draws.** When the engine evaluates a hit or a double, each possible next card is weighted from a fixed draw composition derived from the recorded composition. That composition is **not** updated for the player's earlier hypothetical draws within the same hand. Hit values depend on the hand total and soft/hard status, not on which specific cards produced that total.
3. **Split hands.** Each hand produced by a split is evaluated under the same conventions. The cards hypothetically dealt to one split hand are not removed before the other hand or the dealer is evaluated.
4. **Dealer blackjack.** Under ENHC, the probability of a dealer blackjack uses the recorded composition with the full remaining count as denominator (Section 3.5).

Why these conventions. They are the conventions of the spreadsheet model from which the engine was ported. They make a full set of decision tables computable fast enough for live use. Two slower evaluations exist for measurement only; neither is the production path. One is a **composition-dependent mode of the same engine**, which removes the player's own cards and conditions the dealer on them; it is the target of the calibration (Section 6). The other is a **separate internal reference program**, written without the engine's code, which enumerates with full depletion at every draw (Section 8.3). The two are not identical.

What has been measured. The recorded composition reflects every card actually entered as exposed. The production values have been compared on selected states with an internal reference program that enumerates with full depletion at every draw (Section 8.3). The differences below are observed totals. **They have not been fully decomposed by cause** and are not attributed to any single convention:

- **Stand.** On recorded compositions, the production Stand values matched the reference exactly in the measured cases. Stand involves no player draws, so only the dealer distribution matters.
- **Hit and Double.** On shoes of about 290 cards, similar to those in the simulations, the differences were up to about 0.002 (Hit) and 0.004 (Double) EV units in the measured states. On very small shoes (around 20–40 cards) they reached several hundredths.
- **Split.** One measured split state on a small shoe differed by about 0.016 EV units.

These are case measurements, not a statistical bound. Whether an approximation of this size changes the highest-EV action is checked case by case. It has not been quantified over a large sample. (*Evidence gap; see Section 10.*)

4.3 Determinism

Given the same recorded composition, hand, upcard and rules, the production path returns the same action EVs and the same selected action. Core decision values do not use random sampling.

5. How each action is evaluated

Every action uses the conventions of Section 4.2.

- **Stand.** The player's total is compared against the dealer's outcome distribution for the upcard: win if the dealer busts or finishes lower, push on an equal total, lose otherwise. A player 21 does not push against a dealer blackjack.
- **Hit.** Each possible next card is weighted by its probability. After the card, the hand either busts, or continues with the better of Stand and further Hit, evaluated recursively.
- **Double.** Exactly one more card, then Stand, with the result counted on two units.
- **Split.** The pair is divided and each hand receives a card. The rules then determine what each hand can do:
 - resplitting is allowed if the same rank appears and the hand limit is not reached (for aces, only with resplit-aces);
 - split aces receive one card and stand, unless hitting split aces is allowed;
 - otherwise, the best of Stand, Hit and (with DAS) Double.

The values of all resulting hands are summed.

- **Surrender.** The fixed value of Section 3.5, compared against the other available actions.

Split restrictions in the application. The restrictions on hands created by a split are applied uniformly to every split hand: DAS, double restriction, split-ace rules. Before a correction on 23 August 2026, split hands were advised using the main-hand table, without these restrictions. The correction is covered by 14 regression tests.

6. Pre-bet EV and its calibration

6.1 What the pre-bet EV is

The pre-bet EV is the expected net result per unit of initial bet for the **next** round, given the recorded composition before that round is dealt. It is the input to bet sizing (Section 7).

6.2 How it is formed

1. **The initial deal.** The engine enumerates every initial deal: the player's two cards and the dealer's upcard. It weights each deal by its probability of being dealt from the recorded composition, without replacement.
2. **The value of each deal.** Each deal is valued with the conventions of Section 4.2:
 - player blackjack is handled separately;
 - under peek, dealer blackjack is checked;
 - under ENHC, it is accounted for in the push of a player blackjack.

The probability of a dealer blackjack given a particular deal takes that deal's cards into account.

3. **Calibration.** A calibration correction is added to the result (6.3), within its domain (6.5).

6.3 Why a calibration exists

The conventions of Section 4.2 introduce a systematic difference between the uncalibrated pre-bet EV and the engine's slower composition-dependent mode on the same composition (Section 4.2). The difference was measured on a full 8-deck shoe and at three deeper penetrations. It is approximately **inversely proportional to the number of cards remaining**, and it grows as the shoe is depleted. The calibration adds a correction of that form, plus small rule-dependent terms. Its purpose is to bring the pre-bet EV to the level of that composition-dependent mode.

The calibration target is itself a model, not an exact answer. On the states measured in Section 8.3, the composition-dependent mode differed from the independent full-enumeration program by up to about 0.004 EV units for Stand and Hit, and about 0.008 for Double, on shoes of about 290 cards. The effect of that difference on the pre-bet EV has not been measured.

6.4 What is public and what is not

Public:

- that a calibration exists;
- its purpose;
- its functional form (a depth-dependent term inversely proportional to cards remaining, plus rule-dependent terms);
- its domain (6.5);
- the evidence for it (6.6).

Not public: the numerical coefficients and implementation details. They are recorded in a separate internal document.

Consequence. This document allows the model to be inspected: what is calculated, under which conventions, within which domain, and with which evidence. It also allows the published reproducible states (Section 8.4) to be examined. It does **not** demonstrate that all production results can be reconstructed in full from the public description. In particular, the calibrated pre-bet EV cannot be reproduced without the unpublished coefficients. No proprietary detail is added to this document to support a stronger claim.

6.5 Domain of the calibration

The calibrated pre-bet EV is used only when both conditions hold:

- **at least 104 cards remain**, the deepest level at which the depth law was measured;
- **the proportion of ten-valued cards lies between 26% and 34%**. A neutral shoe has 30.77%. Outside this band the correction is not supported by the measurements and can overcorrect.

Outside the domain, the uncalibrated pre-bet EV is used. A less refined value is preferred to an extrapolated one.

6.6 Evidence for the calibration

Depth law. The depth law was measured against the engine's composition-dependent mode at 416, 312, 208 and 104 cards remaining on 8-deck compositions.

Rule-dependent terms.

- One composite rule term was set from **a single published rule point**. Its accuracy for individual rule changes is therefore not established.
- For the no-DAS case, no term is applied. The source spreadsheet over-penalised the absence of DAS. The engine does not reproduce that error, so a compensating term would double-correct. This was arbitrated with the engine's composition-dependent mode.

Historical comparison with published values — *historical calibration reference, incomplete provenance*. During calibration, the calibrated result was compared with three rule points read from the Wizard of Odds Blackjack House Edge Calculator (reference W1, Appendix C). The values are house edges per initial bet. The rules below are given exactly as recorded at the time; settings that were not recorded are not filled in here.

Point	Rules as recorded	Result type recorded	Wizard of Odds (W1)	OOB, calibrated	Difference
1	base configuration, US peek	not recorded	0.43000%	0.43368%	+0.0037 pp
2	US peek, H17, double 10-11 only, up to 3 hands, RSA, HSA, late surrender, blackjack 3:2, full shoe	basic strategy with continuous shuffler	0.53038%	0.52644%	−0.0039 pp
3	ENHC, H17, no DAS, up to 2 hands, HSA, late surrender	not recorded	0.66306%	0.67242%	+0.0094 pp

All three points have incomplete provenance: for points 1 and 3 the result type is not recorded, and for none of the points are all rule settings or the date of reading recorded. The calculator offers three result types that differ in strategy and in shuffling (Appendix C). Because the type is unknown for points 1 and 3, no statement is made about whether the

strategy behind those values matches the strategy OOB uses. For these reasons the values are kept as a historical record of the calibration, **not** as a fully verified direct comparison. If the original record is recovered, the provenance will be completed in a later revision.

Point 2 is the point from which the composite rule term was set, so it is not an independent check of that term. Because published values were used to set or check the calibration, this is at most a **calibration-consistency check**, not validation on independent data.

Not available:

- validation of the calibration outside the fitted points, for other rule sets and compositions (*gap L2*);
- a measurement of the pre-bet EV's error on a large sample of real compositions (*see the incomplete audit, Section 8.6*).

7. Initial bet sizing

Three different bet-sizing rules appear in OOB's history and evidence. They are kept separate throughout this document.

7.1 Historical ramps V1 and K1 (simulation accounts)

All rules here share participation and cover:

- no bet below a pre-bet EV of -1.0% ;
- one unit for $-1.0\% \leq \text{EV} < +0.1\%$.

Above $+0.1\%$, the stake rises with EV in steps of 1, 2, 3, 4, 6, 8, 10 and 12 units of \$10. V1 uses an edge of 0.20% per unit, and K1 0.10% per unit, so K1 reaches each step at half the EV of V1. The maximum stake is 12 units (\$120) for both.

V1 and K1 were simulated with full accounting on the first million (Section 9.3).

7.2 MOV with a fixed unit (analysed retrospectively)

Pre-bet EV	Stake
below -1.0%	no bet
-1.0% to below $+0.1\%$	1 × unit
$+0.1\%$ to below $+1.1\%$	2 × unit
$+1.1\%$ and above	13 × unit

Thresholds are inclusive. In the experiment of Section 9.3 **the unit was a constant \$10**, so the stakes were \$10, \$20 and \$130 throughout, independent of the balance.

This rule was **selected after examining the first-million dataset**: 13× is the first top multiple, in steps of one unit, for which the model contribution on that dataset is positive. The rule was applied retrospectively to that track (Section 9.3).

7.3 MOV sized by confirmed balance (the application)

The application uses the same thresholds and multipliers, with a **dynamic** unit:

- unit = confirmed balance / 100, rounded to the nearest whole number, with halves rounded up (for example, $9.5 \rightarrow 10$ and $10.5 \rightarrow 11$);
- initial bet = unit × MOV multiplier, where the pre-bet EV thresholds select the multiplier 0, 1, 2 or 13.

A confirmed balance below 50 therefore gives a unit of 0 and a suggested bet of 0.

How the application handles the balance.

- Only a positive number is accepted as a confirmed balance. A confirmation that is zero, negative or not a number is refused.
- When a confirmation is refused, the last confirmed balance stays unchanged.
- Without any confirmed balance, no amount is suggested; the application shows only the multiplier and asks for a balance.

- A bet already fixed for the current round is not changed by any confirmation, valid or refused.

The pre-bet EV used for the thresholds is the **estimate given by the production model** (Section 6), calibrated within its domain. It is not a demonstrated expected value of the whole betting policy as played.

The bet is **fixed for the round**: it is computed once, when the round's first card is entered, from the pre-bet EV and the balance confirmed at that moment. Doubles and splits add the same initial bet again. A balance confirmed during a round applies from the next round.

This rule has not been tested on any simulated track. The fixed-unit MOV of 7.2 and the balance-sized MOV of 7.3 coincide only while the confirmed balance lies in the interval [950, 1050), where the unit is 10. Outside that range they are different rules. No result reported in this document describes 7.3.

8. Evidence: checks performed and their scope

Every item states what was checked, its scope, and its label (Section 0).

8.1 Agreement with the source spreadsheet — *porting check*

The engine was ported from a spreadsheet model. Six full-shoe 8-deck configurations were compared; five agree to within about 0.005 percentage points. The sixth (peek, no DAS) differs by about 0.06 pp, by design: the spreadsheet over-penalises the absence of DAS, and the engine does not reproduce that error. This was confirmed with the engine's composition-dependent mode.

A manual validation plan covers seven groups: full-shoe EV, hard and soft boundary cells, splitting, depleted shoes, deck counts, integrity. Two parts were completed and recorded on 23 August 2026:

- **the integrity group**, all checks passed: rank limits, count reset, split seeding, automatic double, undo;
- **five composition scenarios**, with a mean difference of 0.0005 pp and no opposite signs.

The remaining groups are not recorded as completed.

Scope: shows that the port reproduces the spreadsheet where intended. It does not show that the spreadsheet, or the engine, is correct.

8.2 Published house-edge values by deck count — *historical calibration reference, incomplete provenance*

On a full shoe, the internal record lists the following values, measured against the Wizard of Odds calculator (W1):

Decks	OOB house edge	Wizard of Odds (W1)
8	0.434%	0.43%
6	0.408%	0.40%
4	0.357%	0.35%

Both are house edges per initial bet. The record keeps the published values to two decimals. It does not record the rule settings, the calculator's result type or the date of reading, and none of these is filled in here. A value for 5 decks was not recorded and is not given.

Separately, an N-deck shoe gave results identical (0.00000 pp) to an 8-deck shoe depleted to the same composition; this is an internal consistency check of the engine, not a comparison with an external value. The rule points used in the calibration are discussed in 6.6.

Scope: historical values with incomplete provenance; not a fully verified direct comparison and not validation on independent data. If the original record is recovered, the provenance will be completed in a later revision.

8.3 Agreement with an internal reference program — *internal independent-code check*

A separate internal program computes EVs by complete enumeration, with full depletion of the shoe at every player and dealer draw. It does not use the engine's code. The results below come from the verification of 13 September 2026 and belong to the version of the reference program identified there (a79594ccbc59). Later versions of the program do not inherit these results. The comparison used selected states: very small shoes, which magnify any difference, and shoes of about 290 cards. Results:

- Stand on recorded compositions: no difference in the measured cases;
- Hit and Double: differences up to about 0.002 and 0.004 EV units on ~290-card shoes, not decomposed by cause;
- ENHC dealer-blackjack handling: no difference in the measured Stand cases;
- splitting: one state measured under the production conventions (a pair of 8s against a 6 on a small shoe, difference about 0.016 EV units); a second split state (aces against a 6) was measured only for the engine's composition-dependent mode;
- the engine's composition-dependent mode, the calibration target of Section 6: on ~290-card shoes it differed from the reference by up to about 0.004 EV units for Stand and Hit and about 0.008 for Double.

Scope: a small number of selected states. Internal code, not third-party. No large-sample measurement (*gap*).

8.4 Ten published reproducible states — *reproducibility check*

Ten fully specified states are published:

- the exact remaining composition, hand, upcard, rules and action EVs;
- the engine build and methodology revision.

On 6 September 2026 all ten were recomputed on the same engine build in two environments (Python 3.11.4 and 3.14.4), and all ten matched (10/10).

Reviewing these states found the ENHC denominator error described in 3.5. Its correction changed four of the ten states: three by 0.0001, as predicted. It changed no decision and no Stand/Hit margin.

Open discrepancy (N3). In state OOB-RS-0005, the Double EV changed by 0.0011 between the earlier and the corrected run. The dealer upcard in that state is a 4, where the denominator correction has no effect. **The cause of the earlier value has not been established.** It is not attributed to the correction and remains open.

Scope: shows that the published values are reproduced by the same build. It does not validate their correctness.

8.5 Controlled comparison with Basic Strategy, decision level — *partial comparison*

A controlled run (9 September 2026, identifier bs_vs_oob_full_20260909_182319) compared the action recommended by a published Basic Strategy chart with the action OOB selects from the exact composition.

The sample:

- 1,000 generated states under the default rules;
- 813 within the domain: initial two-card, non-pair, non-terminal hands;
- 131 pairs and 56 naturals were excluded by design.

Results:

- 782 agreements (96.19%) and 31 deviations (3.81%);
- every deviation was a strict EV difference;
- the EV given up by following the chart in deviating states averaged 0.0281 per unit (maximum 0.1415);
- both actions in each state were evaluated by the same engine.

Scope:

- the reference chart (BlackjackInfo Basic Strategy Engine, reference B1 in Appendix C) has the status *candidate, not approved as reference*; a second chart (Wizard of Odds Basic Strategy Calculator, W2) is also a candidate;
- pairs and post-split hands are outside the domain;
- this is a decision-level comparison of EV, **not** a comparison of bankroll results;
- **no Basic Strategy comparison exists for the first-million simulation.**

8.6 Audit of the production pre-bet EV — *incomplete audit*

An audit protocol (v1.1) was designed to estimate, round by round, the difference between the production pre-bet EV, calibration included, and a full-play reference. The design covers a sample of 113 rounds drawn from an earlier 300,000-card track.

Status:

- the initial batch evaluated 4 rounds and, by design, produced no aggregate estimate and no verdict;
- a Monte Carlo pilot for intractable cases passed five concordance controls (largest $|z| = 2.00$);
- one hard case was measured for cost and dispersion only.

Scope: no verdict exists on the sign or size of the error in the production pre-bet EV (*gap L3*). The audit population is the earlier 300,000-card track, not the first million.

8.7 Build identification

Two identifiers accompany published results:

- **The build identifier** (00B-ENG-...): a SHA-256 of the engine source files, the same on any machine. This is the published one.
- **The runtime identifier** (00B-RT-...): a fingerprint of the code held in memory, which depends on the Python version.

On 9 September 2026 the runtime mechanism was found to be unstable: it serialised code objects, and the serialisation depends on reference counts during execution. It was replaced with a canonical representation, verified both stable and sensitive to real code changes. Earlier records keep their original identifiers; none was reconciled retroactively.

Scope: *reproducibility check* of identity, not of correctness.

8.8 Independent validation

Not yet performed. No third party has reproduced OOB's calculations (*gap L4*).

9. Evidence from simulation

9.1 The simulator

The simulator deals a shuffled 8-deck shoe to a cut card at 208 of 416 cards, with a fixed seed. Every playing decision is taken by the engine from the recorded composition, and the pre-bet EV is recorded before each round. Accounts with different bet-sizing rules can play the same track.

For every account and round, the simulator reconciles:

- wins, losses and pushes against settled hands;
- money flows against balances;
- non-negative balances after settlement.

A mismatch stops the run. Funding rule: \$1,000 at the start, and \$1,000 added only when a stake cannot be covered. Net results exclude deposits.

9.2 Earlier 300,000-card tracks

Two earlier tracks of about 300,000 cards each were used to compare candidate ramps and to choose K1 for further testing. They are archived. They are context for the choice of rules, not evidence of profitability.

9.3 The first million — *retrospective simulation*

The track. One track of 1,000,005 cards (seed 20261001), dealt between 29 September and 4 October 2026, with 187,406 rounds: 135,411 wagered and 51,995 observed without a bet.

The accounts:

- V1 and K1 were simulated with full accounting. The reconstruction from the round journals matched the simulator's balances on every wagered round.
- VERDE, ROZ and fixed-unit MOV were applied retrospectively. Because settlement is linear in the initial stake, each round's stakes and payouts were rescaled exactly; cards, decisions and participation are unchanged.

Definitions. R is the realised net result. $C = \Sigma (\text{initial stake} \times \text{pre-bet EV})$ is the model contribution.

Variant	R (USD)	C (USD)	Initial stakes	C / stakes	Max. drawdown	Lowest net	Deposits
V1	+3,685.00	−636.29	1,823,680	−0.035%	7,315	−3,290	3
K1	+7,300.00	+3,379.67	2,337,040	+0.145%	11,090	−7,005	7
VERDE	+4,645.00	−1,658.66	1,718,680	−0.097%	6,440	−1,050	1
ROZ	+7,085.00	+195.39	1,830,000	+0.011%	8,040	−2,290	2
MOV (fixed \$10 unit)	+7,570.00	+204.25	1,845,630	+0.011%	7,925	−810	0




Figure 2. Fixed-unit MOV on the first million: cumulative realised net result R and cumulative model contribution C , against cards dealt. Retrospective simulation; the rule was selected on this dataset. Same data and geometry as the public chart at oddsOfBlackjack.com/proof/1-million-cards.

The fixed-unit MOV track:

- finished +\$7,570 from a \$1,000 start, without deposits;
- realised an average of +\$5.59 per 100 wagered rounds. This is a realised average, not an expected rate.

Reading the result:

- **The difference $R - C$.** It is the realised outcome minus the model contribution; it is not attributed entirely to variance.
- **Dispersion.** An indicative dispersion scale for R on this track (empirical round standard deviation $\times \sqrt{n}$) is about \$8,850. No statistical-significance conclusion is drawn from it.
- **Stability.** C for MOV was negative in 3 of 10 blocks of about 100,000 cards. The blocks are parts of one track, not independent tests.
- **Liquidity.** The lowest cash after placing stakes, before settlement, was \$10. The absence of deposits is a property of this track, not evidence of safe capital.

9.4 What the simulation evidence does not contain

- A bankroll comparison between OOB and Basic Strategy on the same track. **There is none:** the first million contains no Basic Strategy track.
- An out-of-sample test of the MOV rule. The rule was selected on the data it is reported on.
- Any result for the balance-sized MOV used by the application.

10. Limitations

1. **Production conventions.** Action EVs are computed under the conditioning conventions of Section 4.2, not with full depletion at every hypothetical card. Measured differences are small on shoes of a few hundred cards and larger on small shoes. They have not been quantified on a large sample. Their effect on decision selection is not established in general.
2. **Calibration.**
 - The pre-bet EV includes an empirical calibration, measured at four depths, with a rule term set from a single published point.
 - It applies only within its domain: at least 104 cards remaining and ten-density between 26% and 34%.
 - Its coefficients are not public, so the calibrated pre-bet EV cannot be fully reproduced from this document.
3. **Pre-bet EV accuracy.** The production pre-bet EV has not been audited to a verdict (8.6). It is the basis of the model contribution C and of bet sizing. A systematic error of the order of the MOV margin (about 0.01 percentage points of stakes) would remove that margin.
4. **Rule coverage.** Checks concentrate on the default configuration and selected published points. Other combinations are supported by the same model but not separately verified. Decks 1–3 are excluded.
5. **Simulation evidence.** One track, with the MOV rule selected in-sample.
6. **Basic Strategy.** No bankroll comparison with Basic Strategy exists. The decision-level comparison uses a reference not yet approved, and excludes pairs and post-split hands.
7. **The application's bet sizing.** The balance-sized MOV is untested on any track.
8. **Exact ties** are resolved by an arbitrary deterministic order.
9. **Independent validation** has not been performed.
10. **Input accuracy.** All outputs assume the cards and rules were entered correctly.
11. **Open discrepancy N3** (OOB-RS-0005) is unresolved.
12. **External reference values.** The Wizard of Odds values used during calibration have incomplete provenance (6.6, 8.2). They are not a fully verified direct comparison.

Appendix A. Glossary

- **A (in simulation tables).** $R - C$.
- **C, model contribution.** Σ (initial stake $\times$ pre-bet EV) over wagered rounds: what the model attributes to the bets.
- **Calibration.** A correction added to the uncalibrated pre-bet EV within a defined domain (Section 6).
- **Cut card / penetration.** The point at which the shoe is reshuffled. Here 208 of 416 cards, i.e. 50%.
- **DAS, RSA, HSA.** Double after split; resplit aces; hit split aces.
- **ENHC.** European No Hole Card: the dealer takes the second card after the player acts. A dealer blackjack then beats all player wagers except a player blackjack, which pushes.
- **EV.** Expected net result per original bet (Section 3).
- **Fixed unit / dynamic unit.** A constant stake unit (\$10 in the experiment), versus a unit derived from the confirmed balance (the application).
- **Pre-bet EV.** The EV of the next round, before it is dealt (Section 6).
- **R.** Realised net result: Σ (amount returned – amount staked), excluding deposits.
- **Recorded composition.** The starting shoe minus all cards entered as exposed (Section 2.1).
- **S17 / H17.** Dealer stands / hits on soft 17.
- **Unit.** The base stake to which multipliers apply.

Appendix B. Evidence labels

See the table in Section 0.

Appendix C. Provenance

Item	Identifier
Engine build	00B-ENG-b0681ede84a3
Methodology revision	v1.0, 6 September 2026
Reproducible states	10 states, verified 6 September 2026, build above, Python 3.11.4 and 3.14.4
BS decision-level comparison	run <code>bs_vs_oob_full_20260909_182319</code> , seed 20260908; reference status: candidate
First-million track	seed 20261001; 1,000,005 cards; 187,406 rounds; 167 batches
First-million journals	SHA-256 aggregate 1366b4a33008e9e1b695c7990e9700f1a1d5a8cb9ac361e4fdb00173f8c43d79
First-million simulator	00B-K1-b564264c55fd
Public result page	oddsOfBlackjack.com/proof/1-million-cards
Internal reference program (8.3)	version <code>a79594ccbc59</code> , verification of 13 September 2026; later versions do not inherit these results

Internal records — verification reports, protocols, audit files and journals — are held privately. They are identified in the internal companion document.

External references

The values attributed to these sources are those recorded internally at the time of reading; they were not re-read for this document, and published external values may have changed since.

Ref.	Source	Configuration and normalisation	Used in
W1	Michael Shackleford, <i>Blackjack House Edge Calculator</i> , Wizard of Odds, https://wizardofodds.com/games/blackjack/calculator/ (page last updated 3 August 2026; results are rendered in the browser)	House edge in % of the initial bet. Three result types: "optimal results" (perfect composition-dependent strategy, shuffle after every hand); "basic strategy with cut card" (total-dependent basic strategy, cut card); "basic strategy with continuous shuffler" (total-dependent basic strategy, shuffle after every hand). Rule configurations as listed in 6.6 and 8.2. Result type recorded for one point only; rule settings incomplete and date of reading not recorded (6.6, 8.2). Status: historical calibration reference, incomplete provenance.	6.6, 8.2
W2	Michael Shackleford, <i>Blackjack Basic Strategy Calculator</i> , Wizard of Odds, https://wizardofodds.com/games/blackjack/strategy/calculator/ (read 6 September 2026)	Settings: 4 or more decks, dealer stands on soft 17, double after split not allowed, no surrender, no peek / no hole card (European). Output: action chart, no EV values.	8.5 (candidate)

Ref.	Source	Configuration and normalisation	Used in
B1	Ken Smith, <i>Basic Strategy Engine</i> , BlackjackInfo, https://www.blackjackinfo.com/blackjack-basic-strategy-engine/?numdecks=8&soft17=s17&dbl=all&das=no&surr=ns&peek=no (read 7 September 2026)	Settings: 8 decks, S17, double any two cards, no DAS, no surrender, no peek (European). Output: action chart, no EV values; read from the generated tables.	8.5 (candidate)