

# Supplementary Materials for

## DeepStack: Expert-Level AI in No-Limit Poker

### **Game of Heads-Up No-Limit Texas Hold'em**

Heads-up no-limit Texas hold'em (HUNL) is a two-player poker game. It is a repeated game, in which the two players play a match of individual games, usually called hands, while alternating who is the dealer. In each of the individual games, one player will win some number of chips from the other player, and the goal is to win as many chips as possible over the course of the match.

Each individual game begins with both players placing a number of chips in the pot: the player in the dealer position puts in the small blind, and the other player puts in the big blind, which is twice the small blind amount. During a game, a player can only wager and win up to a fixed amount known as their stack. In the particular format of HUNL used in the Annual Computer Poker Competition (50) and this article, the big blind is 100 chips and the stack is 20,000 chips or 200 big blinds. Resetting the stacks after each game is called “Doyle’s Game”, named for the professional poker player Doyle Brunson who publicized this variant (25). It is used in the Annual Computer Poker Competitions because it allows for each game to be an independent sample of the same game.

A game of HUNL progresses through four rounds: the pre-flop, flop, turn, and river. Each round consists of cards being dealt followed by player actions in the form of wagers as to who will hold the strongest hand at the end of the game. In the pre-flop, each player is given two private cards, unobserved by their opponent. In the later rounds, cards are dealt face-up in the center of the table, called public cards. A total of five public cards are revealed over the four rounds: three on the flop, one on the turn, and one on the river.

After the cards for the round are dealt, players alternate taking actions of three types: fold, call, or raise. A player folds by declining to match the last opponent wager, thus forfeiting to the opponent all chips in the pot and ending the game with no player revealing their private cards. A player calls by adding chips into the pot to match the last opponent wager, which causes the next round to begin. A player raises by adding chips into the pot to match the last wager followed by adding additional chips to make a wager of their own. At the beginning of a round when there is no opponent wager yet to match, the raise action is called bet, and the call action is called check, which only ends the round if both players check. An all-in wager is one involving all of the chips remaining the player’s stack. If the wager is called, there is no further wagering in later rounds. The size of any other wager can be any whole number of chips remaining in the player’s stack, as long as it is not smaller than the last wager in the current round or the big blind.

The dealer acts first in the pre-flop round and must decide whether to fold, call, or raise the opponent's big blind bet. In all subsequent rounds, the non-dealer acts first. If the river round ends with no player previously folding to end the game, the outcome is determined by a showdown. Each player reveals their two private cards and the player that can form the strongest five-card poker hand (see "List of poker hand categories" on Wikipedia; accessed January 1, 2017) wins all the chips in the pot. To form their hand each player may use any cards from their two private cards and the five public cards. At the end of the game, whether ended by fold or showdown, the players will swap who is the dealer and begin the next game.

Since the game can be played for different stakes, such as a big blind being worth \$0.01 or \$1 or \$1000, players commonly measure their performance over a match as their average number of big blinds won per game. Researchers have standardized on the unit milli-big-blinds per game, or mbb/g, where one milli-big-blind is one thousandth of one big blind. A player that always folds will lose 750 mbb/g (by losing 1000 mbb as the big blind and 500 as the small blind). A human rule-of-thumb is that a professional should aim to win at least 50 mbb/g from their opponents. Milli-big-blinds per game is also used as a unit of exploitability, when it is computed as the expected loss per game against a worst-case opponent. In the poker community, it is common to use big blinds per one hundred games (bb/100) to measure win rates, where 10 mbb/g equals 1 bb/100.

## Poker Glossary

**all-in** A wager of the remainder of a player's stack. The opponent's only response can be call or fold.

**bet** The first wager in a round; putting more chips into the pot.

**big blind** Initial wager made by the non-dealer before any cards are dealt. The big blind is twice the size of the small blind.

**call** Putting enough chips into the pot to match the current wager; ends the round.

**check** Declining to wager any chips when not facing a bet.

**chip** Marker representing value used for wagers; all wagers must be a whole numbers of chips.

**dealer** The player who puts the small blind into the pot. Acts first on round 1, and second on the later rounds. Traditionally, they would distribute public and private cards from the deck.

**flop** The second round; can refer to either the 3 revealed public cards, or the betting round after these cards are revealed.