Aalto-yliopisto
Perustieteiden korkeakoulu

Probabilistic Modelling of Chess Outcome Distributions and Estimation of the Fair Frontier

Topic Introduction

Ellis Saavalainen

30.4.2026

Ohjaaja: *Ahti Salo*
Valvoja: *Ahti Salo*

Motivation

Why probabilistic chess outcome modelling?

Predicting win / draw / loss outcomes from contextual variables sits at the intersection of chess analytics and machine learning. Rating differences alone have been studied extensively, but combining *rating*, *engine evaluation*, *clock*, and *move number* into one unified model has received less attention.

Why interpretability matters

A black-box predictor of WDL outcomes is of limited use to the chess community. By imposing structural priors (antisymmetric WDL logits, partial monotonicity, mask-aware learning), the model becomes both predictive *and* interpretable, exposing how each variable shapes winning chances.

The fair frontier

A central conceptual outcome is the **fair frontier**: the set of input combinations at which White and Black have equal win probabilities. It provides a domain-meaningful interpretation of balance between players.

Research Questions

Goal. Construct a probability model of chess outcomes that is both predictive and interpretable.

RQ1. How do rating, engine evaluation, clock, and move number *impact* outcome probabilities?

Concretely, how does each variable individually shift P(White win), P(Draw), P(Black win), and how do they interact?

RQ2. Which combinations of these parameters define the fair frontier, *of which* the White and Black win probabilities are equal?

What is the geometry of the level set $\{ x : P(White) = P(Black) \}$ *across feature pairs?*

Key Terms and Definitions

Rating (Elo). Scalar measure of player strength from the Elo system [Elo, 1978]. Higher rating ⇒ stronger player. We use Lichess blitz ratings.

Elo difference (ΔR). white_elo − black_elo. Standard predictor of expected score in chess.

Engine evaluation (cp). Centipawn score from White's perspective produced by Stockfish at fixed budget; clipped to [−1000, 1000]. Positive ⇒ White is better.

Clock fraction. Remaining time / initial budget (300 s for 5+0 blitz). 1.0 = full clock; 0.0 = flagged.

Clock-difference fraction. white_clock_fraction − black_clock_fraction. Positive ⇒ White has more time.

Full-move number. Number of completed full moves at the sampled position; coarse indicator of game phase.

Fair frontier ($\mathcal{F}$). Set of feature combinations at which P(White win) = P(Black win) under the model.

Background — Probabilistic WDL Modelling

Categorical outcomes via softmax. Predicted probabilities are obtained by mapping inputs to logits and applying softmax: $p(y|x) = \mathrm{softmax}(\eta(x))$

[Agresti, 2002; McCullagh & Nelder, 1989; Goodfellow et al., 2016].

Bradley–Terry / Davidson family. Bradley–Terry [1952] handles paired comparisons without draws. The Rao–Kupper [1967] and Davidson [1970] models extend this to ternary win–draw–loss outcomes through an explicit draw mechanism.

Elo and draw modelling. The classical Elo system [Elo, 1978] models expected score from rating differences but does not specify a draw probability. Recent work links Elo updates to Davidson-type draw models [Szczeciński, 2020].

Proper scoring rules. Probabilistic predictions are evaluated with log loss and the multiclass Brier score [Brier, 1950; Gneiting & Raftery, 2007].

Background — Structural Priors

Antisymmetric logits encode partial monotonicity. An advantage score a(x) and a draw score d(x) parameterise the logits as $(a, d, -a)$. The antisymmetric template makes a(x) the directional signal: $P(White) = P(Black)$ iff $a(x) = 0$, irrespective of d(x). This is the lever for *partial monotonicity* — domain knowledge fixes effect direction for some features (e.g. larger ΔElo favours White), and monotone subnetworks (nonnegative weights + monotone activations) make a(x) non-decreasing in those features [Davidson, 1970; Sill, 1998; You et al., 2017; Runje & Shankaranarayana, 2023; Szczeciński, 2020].

Gate constraint. Clock effects are injected through multiplicative gates in logit space: $\ell\_W = a + \log g\_W,\ \ell\_D = d,\ \ell\_B = -a + \log g\_B$, where $g\_W, g\_B \in [0,1]$ are side-specific clock gates. For raw clock fraction $t > 0, g(t;\, \tau) = \exp(-\tau \,/\, (t + \varepsilon))$ with a learned positive temperature $\tau(x, m)$; for t ≤ 0 the gate is 0 (hard flag boundary); if the clock is masked, the gate is 1 (neutral). The fair-balance condition becomes $a(x, m) + ½ \cdot (\log g\_W - \log g\_B) = 0$, reducing to a(x) = 0 only under neutral / symmetric gates.

Mask-aware learning. Real tabular data have missingness. Concatenating an explicit mask m ∈ $\{0,1\}^6$ to a zero-filled input lets one shared model produce subset-conditioned predictions [Little & Rubin, 2002; Che et al., 2018; Van Ness et al., 2024]. Stochastic masking acts as input-level dropout [Srivastava et al., 2014].

Aalto-yliopisto
Perustieteiden
korkeakoulu

Background — Chess-Specific Covariates

Engine evaluation as positional summary. Stockfish reports a centipawn score; the project also provides fitted WDL models that map evaluations to win rates [Stockfish FAQ; Stockfish WDL]. Lichess publishes a logistic-style centipawn-to-win-percentage mapping used in its analysis tools [Lichess Accuracy].

Time pressure and clock state. Reduced thinking time can systematically alter move quality and decision-making in online blitz [Carow & Polania, 2025]. Clock fractions therefore carry independent predictive signal beyond rating and evaluation.

Move number / game phase. Used as a coarse proxy for opening / middlegame / endgame. Large-scale studies of chess advantage dynamics and first-move effects support its predictive value in population data [Ribeiro et al., 2013; Cook et al., 2024].

Data Source and Sampling

Data source. Public Lichess monthly dump.

Filter. Rated, standard variant, 5+0 blitz, decisive or drawn result.

Source month. February 2026.

Unit of observation. One sampled in-game position; sampling rate $\lceil (M-5)/10 \rceil$ per game.

Target. Outcome class ∈ {White win, Draw, Black win}.

Split. By unique game_id (no leakage across positions of the same game).

Split	Rows
Train	3,999,340
Validation	500,470
Test	499,365

Sampled Position: Example from Lichess

Each row in the dataset corresponds to one such position: *(white_elo, elo_diff, eval_cp_white_pov, full-move number, white & clock-diff fractions)* plus the eventual game result.

Stockfish evaluation is computed at a fixed 20 ms per position (capped to ±1000 cp). Clock values come directly from the PGN clock annotations.

Input Features

Feature vector $x \in \mathbb{R}^6$

$x_1 =$ full-move number *(game phase)*

$x_2 =$ white_elo *(absolute rating level)*

$x_3 =$ elo_diff $=$ white_elo $-$ black_elo

$x_4 =$ eval_cp_white_pov *(Stockfish, capped ±1000)*

$x_5 =$ white_clock_fraction $\in [0, 1]$

$x_6 =$ clock_diff_fraction = *x_5 − b_clock_fraction*

Standardised by training-split mean / std before being fed to the network.

Observation mask $m \in \{0,1\}^6$

$m_j = 1$ feature x_j is observed

$m_j = 0$ feature x_j is hidden (zero-filled)

Each m_j indexes the same coordinate as x_j, so e.g. m_3 = 1 means elo_diff is observed for this row.

Why a mask?

Lets one model produce subset-conditioned predictions p(y | observed subset, m), so it can also be queried in scenarios where some features are unavailable. Masks are sampled stochastically during training (40% full mask, 60% over the 62 other non-empty masks).

Formal Prediction Goal

Symbol cheat sheet

θ collection of all trainable model parameters (network weights & biases).

$p\theta(\cdot \mid x, m)$ predicted conditional distribution over outcome classes given features x and mask m.

Δ^{W,D,B} the 2-simplex of valid probabilities over (White, Draw, Black): the set of feasible 3-D probability vectors.

Model output

$$p\theta(\cdot \mid x, m) = (p\_W, p\_D, p\_B) \in \Delta\text{^}\{W,D,B\},\ \ p\_W + p\_D + p\_B = 1,\ \ p^r \geq 0.$$

Fair frontier

$$\mathcal{F} = \{\ x \in \mathbb{R}^6 : p\theta(W \mid x, m) = p\theta(B \mid x, m)\ \}$$

i.e. the set of feature configurations at which the predicted White-win and Black-win probabilities are equal. RQ2 concerns the geometry of $\mathcal{F}$ across feature pairs.

From Topic to Results

What this presentation covered (Topic Introduction)

- Motivation, research questions, and key terms
- Background: probabilistic WDL modelling, structured logits, monotonicity, mask-aware learning, chess covariates
- Data source, sampling unit, train/val/test split
- Input features (x, m) and the formal prediction goal in terms of $p\theta(\cdot \mid x, m)$ and the fair frontier $\mathcal{F}$

What comes next (Results presentation)

- Training & evaluation setup (objective, optimiser, learning-rate sweep)
- Predictive performance and calibration
- Fair frontier geometry across feature pairs
- Time-trouble effects, absolute-rating effect, and conclusions