Appendices for

"Rating competitors in games with strength-dependent tie probabilities"

.1

Appendix A Computation of Newton-Raphson updates

The one-step Newton-Raphson updates in (4.10) and (4.11) can be evaluated through a sequence of explicit calculations. For notational simplicity, we suppress the time index t throughout the remainder of this section. Additionally, conditioning on the system parameters is omitted except where their presence is relevant to the computations.

Let p_w, p_d and p_ℓ be the win, draw and loss probabilities for player i, as defined in (3.1), with the arguments θ_i and θ_j suppressed for brevity. In the expressions for these probabilities, the coefficients multiplying θ_i in the exponents can be written as

$$\begin{aligned} a_w &= 1 + x_{ij}\alpha_1/8, \\ a_d &= (1+\beta_1)/2, \\ a_\ell &= -x_{ij}\alpha_1/8. \end{aligned}$$

Define the following expectations over the outcome probabilities:

$$\begin{aligned} s_1 &= a_w p_w + a_d p_d + a_\ell p_\ell, \\ s_2 &= a_w^2 p_w + a_d^2 p_d + a_\ell^2 p_\ell. \end{aligned}$$

The quantity s_1 can be interpreted as the expected score for player i, treating a_w, a_d and a_ℓ as outcome-specific scores. Likewise, s_2 represents the expected square score. For example, when $\alpha_1 = 0$ and $\beta_1 = 0$, the scoring reduces to the conventional values of 1 for a win, $\frac{1}{2}$ for a draw, and 0 for a loss.

Under this framework, the first derivatives of the outcome probabilities with respect to θ_{it} are

$$\begin{aligned} \frac{\partial p_w}{\partial \theta_i} &= p_w(a_w - s_1), \\ \frac{\partial p_d}{\partial \theta_i} &= p_d(a_d - s_1), \\ \frac{\partial p_\ell}{\partial \theta_i} &= p_\ell(a_\ell - s_1), \end{aligned}$$

and the second derivatives are

$$\frac{\partial^2 p_w}{\partial \theta_i^2} = p_w(a_w^2 - s_2 - 2s_1(a_w - s_1)),$$

$$\frac{\partial^2 p_d}{\partial \theta_i^2} = p_d(a_d^2 - s_2 - 2s_1(a_d - s_1)),$$

$$\frac{\partial^2 p_\ell}{\partial \theta_i^2} = p_\ell(a_\ell^2 - s_2 - 2s_1(a_\ell - s_1)).$$

These expressions provide the components necessary to compute the first and second derivatives of the log-posterior in (4.9), evaluated at the prior mean μ_i, as required by (4.10) and (4.11).

For each opponent j, define the predicted outcome probabilities using the Gauss-Hermite quadrature nodes

$$\begin{aligned}
P_{wj}^- &= \mathrm{P}(Y_{ij} = 1 \mid \mu_i, \mu_j - \sigma_j),\\
P_{wj}^+ &= \mathrm{P}(Y_{ij} = 1 \mid \mu_i, \mu_j + \sigma_j),\\[1em]
P_{dj}^- &= \mathrm{P}(Y_{ij} = \frac{1}{2} \mid \mu_i, \mu_j - \sigma_j),\\
P_{dj}^+ &= \mathrm{P}(Y_{ij} = \frac{1}{2} \mid \mu_i, \mu_j + \sigma_j),\\[1em]
P_{\ell j}^- &= \mathrm{P}(Y_{ij} = 0 \mid \mu_i, \mu_j - \sigma_j),\\
P_{\ell j}^+ &= \mathrm{P}(Y_{ij} = 0 \mid \mu_i, \mu_j + \sigma_j).
\end{aligned}$$

Let P_j be the sum of the predicted probabilities for the observed game outcome:

$$P_j = \begin{cases} P_{wj}^- + P_{wj}^+ & \text{if } y_{ij} = 1,\\ P_{dj}^- + P_{dj}^+ & \text{if } y_{ij} = \frac{1}{2},\\ P_{\ell j}^- + P_{\ell j}^+ & \text{if } y_{ij} = 0. \end{cases}$$

Also, define the outcome-weighted expectations:

$$\begin{aligned}
s_{1j}^- &= a_w P_{wj}^- + a_d P_{dj}^- + a_\ell P_{\ell j}^-,\\
s_{1j}^+ &= a_w P_{wj}^+ + a_d P_{dj}^+ + a_\ell P_{\ell j}^+,\\
s_{2j}^- &= a_w^2 P_{wj}^- + a_d^2 P_{dj}^- + a_\ell^2 P_{\ell j}^-,\\
s_{2j}^+ &= a_w^2 P_{wj}^+ + a_d^2 P_{dj}^+ + a_\ell^2 P_{\ell j}^+.
\end{aligned}$$

Then the j-th term in the first derivative sum from (4.9) is

$$\delta_{1j} = \begin{cases} \left(P_{wj}^-(a_w - s_{1j}^-) + P_{wj}^+(a_w - s_{1j}^+))\right)/P_j & \text{if } y_{ij} = 1,\\ \left(P_{dj}^-(a_d - s_{1j}^-) + P_{dj}^+(a_d - s_{1j}^+))\right)/P_j & \text{if } y_{ij} = \frac{1}{2},\\ \left(P_{\ell j}^-(a_\ell - s_{1j}^-) + P_{\ell j}^+(a_\ell - s_{1j}^+))\right)/P_j & \text{if } y_{ij} = 0. \end{cases}$$

Note that the factor of $\frac{1}{2}$ in the definition of u_{ijt} appears in both the numerator and denominator and cancels out, so it is omitted here.

The corresponding second derivative term is:

$$\delta_{2j} = \begin{cases} \begin{gathered}\Big(P^-_{wj}(a_w - s^-_{2j} - 2s^-_{1j}(a_w - s^-_{1j})) + \\ P^+_{wj}(a_w - s^+_{2j} - 2s^+_{1j}(a_w - s^+_{1j}))\Big)/P_j - \delta^2_{1j}\end{gathered} & \text{if } y_{ij} = 1, \\ \begin{gathered}\Big(P^-_{dj}(a_d - s^-_{2j} - 2s^-_{1j}(a_d - s^-_{1j})) + \\ P^+_{dj}(a_d - s^+_{2j} - 2s^+_{1j}(a_d - s^+_{1j}))\Big)/P_j - \delta^2_{1j}\end{gathered} & \text{if } y_{ij} = \frac{1}{2}, \\ \begin{gathered}\Big(P^-_{\ell j}(a_\ell - s^-_{2j} - 2s^-_{1j}(a_\ell - s^-_{1j})) + \\ P^+_{\ell j}(a_\ell - s^+_{2j} - 2s^+_{1j}(a_\ell - s^+_{1j}))\Big)/P_j - \delta^2_{1j}\end{gathered} & \text{if } y_{ij} = 0. \end{cases}$$

As before, the factor of $\frac{1}{2}$ cancels and is omitted in the expressions above.

With the quantities δ_{1j} and δ_{2j} computed as above, the posterior mean and variance updates in (4.10) and (4.11) can be rewritten as

$$\mu^*_{it} = \mu_{it} + \frac{\sum_{j\in\Omega_{it}} \delta_{1j}}{\sigma^{-2}_{it} - \sum_{j\in\Omega_{it}} \delta_{2j}}, \tag{A.1}$$

$$\sigma^{*2}_{it} = \frac{1}{\sigma^{-2}_{it} - \sum_{j\in\Omega_{it}} \delta_{2j}}.$$

These updates are computed independently and in parallel for each player i, enabling efficient batch processing of all players' posterior distributions.

An artifact of performing the rating updates in parallel is that the draw score value, $a_d = (1 + \beta_1)/2$, can introduce unintended effects when $\beta_1 \neq 0$. For example, if $\beta_1 > 0$, implying that stronger players are more likely to draw, then two players with equal prior means who draw a game will both experience an increase in their posterior means, according to Equation (A.1). This behavior is undesirable, as a draw between equally rated players should not systematically increase their estimated strengths. To address this, the implemented system overrides the default value of a_d, setting it to exactly $\frac{1}{2}$, regardless of the value of β_1. This adjustment ensures that when two players with equal prior means draw a game, their posterior means remain unchanged, preserving the symmetry and interpretability of the rating update.

Appendix B Validation of the normal approximation using Gauss-Hermite quadrature

Several numerical approximations are employed to replace the posterior density for θ_{it} in (4.6) with the normal approximation defined by the updates in (4.10) and (4.11). In this section, we assess the adequacy of this approximation.

Although the distribution of θ_{it} described in (4.6) does not admit a closed-form representation, it can be accurately approximated using Gauss-Hermite quadrature. Here we focus on the calculation of the posterior mean and variance. The posterior mean of θ_{it} is given by the ratio of two nested integrals:

$$
\begin{aligned}
\mathrm{E}(\theta_{it} \mid y_{ijt}) &= \frac{\displaystyle\int_{-\infty}^{\infty} \theta_{it} \cdot \mathrm{N}(\theta_{it} \mid \mu_{it}, \sigma_{it}^2) \left(\int_{-\infty}^{\infty} p(y_{ijt} \mid \theta_{it}, \theta_{jt}) \cdot \mathrm{N}(\theta_{jt} \mid \mu_{jt}, \sigma_{jt}^2)\, d\theta_{jt} \right) d\theta_{it}}{\displaystyle\int_{-\infty}^{\infty} \mathrm{N}(\theta_{it} \mid \mu_{it}, \sigma_{it}^2) \left(\int_{-\infty}^{\infty} p(y_{ijt} \mid \theta_{it}, \theta_{jt}) \cdot \mathrm{N}(\theta_{jt} \mid \mu_{jt}, \sigma_{jt}^2)\, d\theta_{jt} \right) d\theta_{it}} \\
&= \frac{\displaystyle\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \theta_{it} \cdot p(y_{ijt} \mid \theta_{it}, \theta_{jt}) \cdot \mathrm{N}(\theta_{it} \mid \mu_{it}, \sigma_{it}^2)\mathrm{N}(\theta_{jt} \mid \mu_{jt}, \sigma_{jt}^2)\, d\theta_{jt}\, d\theta_{it}}{\displaystyle\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} p(y_{ijt} \mid \theta_{it}, \theta_{jt}) \cdot \mathrm{N}(\theta_{it} \mid \mu_{it}, \sigma_{it}^2)\mathrm{N}(\theta_{jt} \mid \mu_{jt}, \sigma_{jt}^2)\, d\theta_{jt}\, d\theta_{it}}.
\end{aligned}
\tag{B.1}
$$

Similarly, the posterior second moment of θ_{it} is given by

$$
\mathrm{E}(\theta_{it}^2 \mid y_{ijt}) = \frac{\displaystyle\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \theta_{it}^2 \cdot p(y_{ijt} \mid \theta_{it}, \theta_{jt}) \cdot \mathrm{N}(\theta_{it} \mid \mu_{it}, \sigma_{it}^2)\mathrm{N}(\theta_{jt} \mid \mu_{jt}, \sigma_{jt}^2)\, d\theta_{jt}\, d\theta_{it}}{\displaystyle\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} p(y_{ijt} \mid \theta_{it}, \theta_{jt}) \cdot \mathrm{N}(\theta_{it} \mid \mu_{it}, \sigma_{it}^2)\mathrm{N}(\theta_{jt} \mid \mu_{jt}, \sigma_{jt}^2)\, d\theta_{jt}\, d\theta_{it}},
$$

so that

$$
\mathrm{Var}(\theta_{it} \mid y_{ijt}) = \mathrm{E}(\theta_{it}^2) - \left(\mathrm{E}(\theta_{it})\right)^2. \tag{B.2}
$$

To evaluate the expressions in (B.1) and (B.2), the numerator and denominator can be computed separately through Gauss-Hermite quadrature. This method is specifically designed for approximating integrals of the form:

$$
\int_{-\infty}^{\infty} f(z) e^{-z^2}\, dz \approx \sum_{r=1}^{R} w_r f(z_r),
$$

where $\{z_r\}_{r=1}^R$ are the roots of the degree-R Hermite polynomial $H_R(z)$, and $\{w_r\}_{r=1}^R$ are the corresponding quadrature weights, given by

$$
w_r = \frac{2^{R-1} R! \sqrt{\pi}}{R^2 (H_{R-1}(z_r))^2}.
$$

To adapt this technique to integrals involving normal densities of the form $\mathrm{N}(\theta \mid \mu, \sigma^2)$, we apply the change of variables $\theta = \mu + \sqrt{2}\sigma z$, which yields the quadrature approximation

$$
\int_{-\infty}^{\infty} f(\theta) \cdot \mathrm{N}(\theta \mid \mu, \sigma^2)\, d\theta \approx \frac{1}{\sqrt{\pi}} \sum_{r=1}^{R} w_r \cdot f(\mu + \sqrt{2}\sigma z_r).
$$

Applying this to both the inner and outer integrals in (B.1), we define for each quadrature node $\theta_{it}^{(s)} = \mu_{it} + \sqrt{2}\sigma_{it} z_s$, and approximate the inner integrals over θ_{jt} for opponent j as

$$
I_j(\theta_{it}^{(s)}) \approx \frac{1}{\sqrt{\pi}} \sum_{r=1}^{R} w_r \cdot p(y_{ijt} \mid \theta_{it}^{(s)}, \mu_{jt} + \sqrt{2}\sigma_{jt} z_r).
$$

The posterior mean of θ_{it} can then be approximated as

$$\mathrm{E}[\theta_{it} \mid y_{ijt}] \approx \frac{\sum_{s=1}^{R} w_s \cdot \theta_{it}^{(s)} \cdot I_j(\theta_{it}^{(s)})}{\sum_{s=1}^{R} w_s \cdot I_j(\theta_{it}^{(s)})},$$

and a similar approximation can be used to compute $\mathrm{E}(\theta_{it}^2 \mid y_{ijt})$, which in turn yields the posterior variance via (B.2). These approximations serve as the ground truth for evaluating the accuracy of the normal approximation method described in the main text. In our implementation of the outcome probability model $p(y_{ijt} \mid \cdot)$, we found that using $R = 9$ quadrature points provides sufficient accuracy for evaluating these integrals.

We assessed the accuracy of the normal approximation used in the rating update algorithm by comparing it to the more precise Gauss-Hermite quadrature evaluation of the posterior density in (4.6). To do so, we computed prior means and variances for all players during the final period of the ICCF dataset (Q1 2022), using the ICCF-implemented system parameters from Table 2. For each game, we performed single-game posterior updates of the strength parameter for the player with white, once using the approximate method described in this paper and once using Gauss-Hermite quadrature as a more accurate benchmark.

In both approaches, the same outcome probability model $p(y_{ijt} \mid \theta_{it}, \theta_{jt})$ and the same system parameters from Table 2 were used. These comparisons were conducted separately for all 17,414 games from the validation period. Table 3 summarizes the results of the comparisons. Each row summarizes results for a specific subset of the 17,414 validation games. Subsets were

		Mean comparisons				Log Std Dev comparions
	N	Δ_{approx}	Δ_{GH}	$R^2_{y=x}$	Δ_{approx} - Δ_{GH}	$R^2_{y=x}$
All games	17414	0.0402	0.0405	0.9855	0.0076	0.9644
Decisive games only	5149	0.1023	0.0997	0.9912	0.0115	0.9536
Drawn games only	12265	0.0141	0.0157	0.9169	0.0059	0.9765
All games, $\mu_1 > 5.06$	5842	0.0148	0.0133	0.9623	0.0048	0.9155
Decisive games, $\mu_1 > 5.06$	785	0.0654	0.0572	0.9580	0.0094	0.8144
Drawn games, $\mu_1 > 5.06$	5057	0.0070	0.0065	0.8107	0.0041	0.9403
All games, $3.43 < \mu_1 \leqslant 5.06$	5774	0.0261	0.0256	0.9733	0.0063	0.9273
Decisive games, $3.43 < \mu_1 \leqslant 5.06$	1278	0.0807	0.0741	0.9797	0.0107	0.8912
Drawn games, $3.43 < \mu_1 \leqslant 5.06$	4496	0.0105	0.0118	0.8739	0.0050	0.9445
Games, $\mu_1 \leqslant 3.43$	5798	0.0798	0.0828	0.9874	0.0117	0.9634
Decisive games, $\mu_1 \leqslant 3.43$	3086	0.1205	0.1210	0.9924	0.0124	0.9535
Drawn games, $\mu_1 \leqslant 3.43$	2712	0.0334	0.0393	0.9191	0.0109	0.9784

Table 3: Summaries of single-game updates for 17,414 ICCF games played during Q1 2022. Updates were computed using both the approximate normal method and Gauss-Hermite quadrature, under the ICCF-implemented system parameters from Table 2.

defined based on two criteria: whether the game was decisive or drawn, and the strength of the player with white, measured by their prior mean μ_1. To capture differences across player strength, we divided the games into terciles based on μ_1: players with $\mu_1 > 5.06$ (strongest), $3.43 < \mu_1 \leqslant 5.06$ (intermediate), and $\mu_1 \leqslant 3.43$ (weakest). This categorization yielded 12 distinct evaluation subsets.

Each summary in the table reports results for a subset of games, with the following columns:

- N: the number of games in the subset.
- Δ_{approx}: the average absolute change from prior to posterior mean using the approximation method developed in this paper.
- Δ_{GH}: the average absolute change from prior to posterior mean using Gauss-Hermite quadrature.
- $R^2_{y=x}$: the coefficient of determination from regressing the approximate posterior mean changes against those from Gauss-Hermite quadrature, relative to the identity line $y = x$.
- $\Delta_{\text{approx}} - \Delta_{\text{GH}}$: the average absolute difference between the posterior mean changes from the approximate method and those from Gauss-Hermite quadrature.
- Log Std Dev $R^2_{y=x}$: the R^2 statistic comparing the changes in log posterior standard deviations from the approximate method and Gauss-Hermite quadrature, again relative to the line $y = x$.

As shown in Table 3, the first two columns, Δ_{approx} and Δ_{GH}, are consistently close, indicating strong agreement between the approximate method and Gauss-Hermite quadrature in estimating the magnitude of the prior-to-posterior mean updates. The accompanying R^2 values for the mean changes (column 4) further support this, showing that the updates from both methods lie tightly along the identity line $y = x$. Notably, the few cases where $R^2 < 0.95$ correspond to subsets with very small average posterior mean changes, as seen in the first two columns, differences that are unlikely to be practically significant.

The fifth column, representing the average absolute difference in the prior-to-posterior mean changes between the two methods, confirms the numerical closeness of the two approaches across all subsets. Moreover, the final column demonstrates that the updates to the posterior standard deviations (on the log scale) also agree closely between methods.

The only subset where the agreement appears weaker, based on a lower R^2 in the final column, is for decisive games involving high-strength players. This discrepancy likely arises from the relatively small changes in posterior standard deviation in such cases, making even minor numerical differences more visible in the R^2 statistic. Nonetheless, the correlation between posterior standard deviations for this subset (row 5 of the table) is exceptionally high, at 0.9999913, indicating near-perfect agreement between the two methods.