Pricing advantage in contests: Mechanism design and evidence from Armageddon chess

Derek J. Clark Tapas Kundu Øystein Myrland Tore Nilssen

Online Appendix

August 25, 2026

This online appendix accompanies *Pricing advantage in contests: Mechanism design and evidence from Armageddon chess*. Numbering of sections, results, equations, and footnotes is prefixed O; references to unprefixed objects point to the main text. Bibliographic citations in this appendix refer to the reference list of the main paper. Supplementary theory is in Section O.1, additional empirical results are in Section O.2, and a parametric example is given in Section O.3.

O.1 Theory

This section collects supporting theory, in the order the main text invokes it: the comparative statics of the baseline contest (Section 2.1), the algebra omitted from the appendix proofs of Proposition 1 and Corollary 1, the primitive content of Assumption 3 (Section 3), the outcome-uniqueness result for the bidding game, the proofs of Propositions 7 and 8, and the derivations behind Section 6.1.

O.1.1 Comparative statics in the baseline contest

By Lemma 1, the equilibrium probabilities depend only on the non-input advantage μ_0 and the threshold Δ. Differentiating with respect to μ_0 gives

$$\begin{aligned}\frac{dP_W}{d\mu_0} =& g(\Delta-\mu_0) > 0, \quad \frac{dP_B}{d\mu_0} = -g(\Delta+\mu_0) < 0,\\ \frac{dP_D}{d\mu_0} =& -g(\Delta-\mu_0) + g(\Delta+\mu_0).\end{aligned}$$

A larger μ_0 raises role W's winning probability and lowers that of role B. Since g is even and single-peaked at zero, $dP_D/d\mu_0$ has the sign of $-\mu_0$, so the prevalence of an unresolved outcome is maximised at $\mu_0 = 0$ and falls with $|\mu_0|$. Moreover, $\partial P_D/\partial\Delta = g(\Delta-\mu_0) + g(\Delta+\mu_0) > 0$, so a higher threshold makes an unresolved outcome more likely.

O.1.2 Omitted algebra for the appendix proofs of Proposition 1 and Corollary 1

Differentiating the map in (7) gives

$$\eta'(x) = -\,v\,g'(\Delta - \mu_0 - x)\left(\frac{1}{c''(e_W(x))} - \frac{1}{H(\tau)c''(e_B(x))}\right),$$

and since $c'' \geq m$ and $H(\tau) \geq 1$, the absolute value of the term in parentheses is at most $1/m$. Differentiating $\hat{e}^*(\tau) = \eta(\hat{e}^*(\tau), \tau)$ yields

$$\frac{d\hat{e}^*}{d\tau} = \frac{\eta_\tau(\hat{e}^*(\tau), \tau)}{1 - \eta_x(\hat{e}^*(\tau), \tau)},$$

where the denominator is positive because $|\eta_x| < 1$. Directly from (7),

$$\eta_\tau(x, \tau) = \frac{vg(\Delta - \mu_0 - x)H'(\tau)}{H(\tau)^2 c''(e_B(x))} < 0.$$

For Corollary 1, with $\rho_W(\theta) = 1 - G(\Delta - \mu_0 - \hat{e}(\theta))$ for $\theta \in \{v, \tau, \Delta, \mu_0\}$,

$$\frac{d\rho_W}{d\theta} = g(\Delta - \mu_0 - \hat{e})\left(\frac{d\hat{e}(\theta)}{d\theta} - \mathbf{1}_{\{\theta=\Delta\}} + \mathbf{1}_{\{\theta=\mu_0\}}\right).$$

Since the fixed-point map depends on (Δ, μ_0, e) only through $\Delta - \mu_0 - e$,

$$\frac{d\hat{e}}{d\mu_0} = -\frac{d\hat{e}}{d\Delta} = \frac{\eta_e}{1 - \eta_e},$$

and therefore

$$\frac{d\rho_W}{d\mu_0} = g(\Delta - \mu_0 - \hat{e})\left(1 + \frac{d\hat{e}}{d\mu_0}\right) > 0, \qquad \frac{d\rho_W}{d\Delta} = g(\Delta - \mu_0 - \hat{e})\left(\frac{d\hat{e}}{d\Delta} - 1\right) < 0.$$

Finally,

$$\eta_v = g(\Delta - \mu_0 - \hat{e})\left(\frac{1}{c''(e_W^*)} - \frac{1}{H(\tau)c''(e_B^*)}\right),$$

whose sign the maintained assumptions do not pin down.

O.1.3 Interior cutoffs in the bidding game

Lemma O.1 (Assumption 3 as a restriction on primitives)**.** *Maintain Condition 1 and full-support noise, i.e., $g > 0$ on $\mathbb{R}$. Then, at every ability gap and for each $i \in \{1, 2\}$,*

$$\operatorname{sign} \phi_i(1) = \operatorname{sign}(\Delta - \Omega).$$

In particular, Assumption 3 is equivalent to the single primitive inequality $\Delta > \Omega$: it holds at some ability gap if and only if it holds at every ability gap, and it cannot hold for one player while failing for the other. If $\Delta = \Omega$, then $\phi_i(1) = 0$ for both players at every gap; if $\Delta < \Omega$, then $\phi_i(1) < 0$ for both.

Proof. Set $x := \Delta - \Omega$ and, without loss, $d := A_1 - A_2 \geq 0$. At $\tau = 1$, $H(1) = 1$ and $\hat{e}^*(1) = 0$, so in each $\Gamma^k(1)$, both players exert the common input e^{Γ^k} solving $c'(e^{\Gamma^k}) = v\, g(\Delta - \mu^k(1))$, with margins $x - d$ and $x + d$. From (4),

$$\phi_i(1) = v\, S(d) + T_i(d), \qquad S(d) := G(x - d) + G(x + d) - 1,$$

where $T_i(d) := c(e^{\Gamma^i}) - c(e^{\Gamma^j})$ and $T_1 = -T_2$. It suffices to prove the claim for $x \geq 0$, since replacing x by $-x$ flips the signs of both S and T_i.

For fixed d, $S_x = g(x - d) + g(x + d) > 0$ and $S = 0$ at $x = 0$, so $S > 0$ for $x > 0$. Let $F := c \circ (c')^{-1}$. Strong convexity gives $F'(y) = y/c''((c')^{-1}(y)) \leq y/m$. With $y_1 := v\, g(x - d) \geq y_2 := v\, g(x + d) \geq 0$,

$$|T_i(d)| \leq \frac{y_1^2 - y_2^2}{2m} = \frac{v^2}{m} \int_{|x-d|}^{x+d} g(z)\big(-g'(z)\big)\, dz \leq \frac{v^2 \dot{g}}{m} \big[G(x + d) - G(|x - d|)\big],$$

and $G(x + d) - G(|x - d|) \leq S(d)$, with equality when $d > x$; when $d \leq x$, this follows from $S(d) = G(x+d) + G(x-d) - 1$ and symmetry around zero. Condition 1 therefore yields $|T_i(d)| \leq \lambda v S(d)$. Hence, for $x > 0$, $\phi_i(1) \geq (1 - \lambda) v S(d) > 0$; for $x = 0$, $S = T_i = 0$; and the case of $x < 0$ follows by the sign-flip symmetry. Thus $\operatorname{sign} \phi_i(1) = \operatorname{sign}(\Delta - \Omega)$. $\square$

O.1.4 Outcome uniqueness in the bidding game

The following discussion strengthens the selection in Proposition 4: within the class of regular, admissible strategies, and under one additional shape condition on the payoff primitives, the outcome of the bidding game is unique.

As in the proof of Proposition 3, write

$$\mathcal{B}_i(\tau) = U_B(\mu^j(\tau), \tau), \qquad \mathcal{W}_i(\tau) = U_W(\mu^i(\tau), \tau), \qquad \phi_i(\tau) = \mathcal{B}_i(\tau) - \mathcal{W}_i(\tau),$$

and label the players so that $\hat{\tau}_1 \leq \hat{\tau}_2$. Assumption 1 and Lemma 2 imply that $\mathcal{B}_i$ is weakly increasing, $\mathcal{W}_i$ is strictly decreasing, and ϕ_i is strictly increasing with unique zero $\hat{\tau}_i$. If player i bids $b \in \mathcal{T}$ against an opponent bid t, the lower bid wins role B:

$$u_i(b,t) = \begin{cases} \mathcal{B}_i(b), & b < t, \\ \mathcal{W}_i(t), & b > t, \\ \frac{1}{2}\{\mathcal{B}_i(b) + \mathcal{W}_i(b)\}, & b = t. \end{cases}$$

Let F_j denote player j's bid distribution. Away from atoms of F_j, player i's expected payoff from bid b is

$$V_i(b) = \mathcal{B}_i(b)\{1 - F_j(b)\} + \int_{\underline{\tau}}^{b} \mathcal{W}_i(t)\, dF_j(t).$$

If both players mix with densities on an open interval, indifference implies $V_i'(b) = \mathcal{B}_i'(b)\{1 - F_j(b)\} - \phi_i(b) f_j(b) = 0$ on that interval for $i = 1, 2$; equivalently, the opponent's hazard rate must satisfy

$$\frac{f_j(b)}{1 - F_j(b)} = \frac{\mathcal{B}_i'(b)}{\phi_i(b)}. \tag{O.1}$$

Each $\mathcal{B}_i'$ and $\mathcal{W}_i'$ is continuous on $\mathcal{T}$. The map η in (7) is continuously differentiable in both arguments, since $c \in C^2$ with $c'' > m$, and g and H are continuously differentiable. Condition 1 keeps $1 - \eta_x \geq 1 - \lambda > 0$, so the implicit function theorem makes the equilibrium input gap, and with it both equilibrium inputs, continuously differentiable (Online Appendix O.1.2). Then (10)–(11) write $\mathcal{B}_i'$ and $\mathcal{W}_i'$ as sums of products of continuous functions.

Lemma O.2 (Bids below a player's cutoff are weakly dominated)**.** *For player i, any bid $b < \hat{\tau}_i$ is weakly dominated by $\hat{\tau}_i$.*

Proof. Fix the opponent's bid t. If $t < b$, then both bids lose and give $\mathcal{W}_i(t)$. If $b < t < \hat{\tau}_i$, then bid b wins and gives $\mathcal{B}_i(b) \leq \mathcal{B}_i(t) < \mathcal{W}_i(t)$, where the strict inequality follows from $\phi_i(t) < 0$; bid $\hat{\tau}_i$ loses and gives $\mathcal{W}_i(t)$. If $t > \hat{\tau}_i$, then both bids win, and monotonicity of $\mathcal{B}_i$ gives $\mathcal{B}_i(\hat{\tau}_i) \geq \mathcal{B}_i(b)$. At $t = \hat{\tau}_i$, the tie payoff at $\hat{\tau}_i$ equals $\mathcal{B}_i(\hat{\tau}_i) = \mathcal{W}_i(\hat{\tau}_i) \geq \mathcal{B}_i(b)$. At $t = b$, bidding b gives the tie lottery $\frac{1}{2}\big(\mathcal{B}_i(b) + \mathcal{W}_i(b)\big) \leq \mathcal{W}_i(b)$, since $\phi_i(b) < 0$, while $\hat{\tau}_i$ loses and gives $\mathcal{W}_i(b)$. Thus, $\hat{\tau}_i$ is never worse. $\square$

Proposition O.1 (Regular equilibria cannot have two-sided mixing)**.** *Suppose strategies have no singular continuous component and every non-degenerate mixed part has interval support with a density, with finitely many such parts. Suppose in addition that each $\mathcal{B}_i$ is strictly increasing on $(\hat{\tau}_2, 1]$. If $\hat{\tau}_1 < \hat{\tau}_2$, then no Nash equi-*

librium has both players mixing on a common interval on which both have positive residual probability.[O1]

Proof. On a common mixing interval on which both residual probabilities are positive, each player's indifference makes the hazard equations (O.1) necessary for $i = 1, 2$.

On an interval below $\hat{\tau}_1$, both ϕ_is are negative, so (O.1) requires non-negative hazard rates to equal non-positive ratios, which is impossible when $\mathcal{B}_i' > 0$. On an interval inside $(\hat{\tau}_1, \hat{\tau}_2)$, the same contradiction applies to the equation for $i = 2$, because $\phi_2 < 0$ there. Where $\mathcal{B}_i' = 0$, (O.1) instead forces the opponent density to be zero, so that portion cannot be part of a common mixing interval with density. Thus no common regular mixing interval can lie below the relevant cutoff.

Now let $I = (a, z)$, with $\hat{\tau}_2 < a < z \leq 1$, be a common mixing interval on which both residual probabilities are positive. The finiteness supposition confines each absolutely continuous component to a finite union of intervals; take I to be the one with the largest upper endpoint. On I both ϕ_is are positive, so (O.1) defines non-negative hazard rates, and since $\mathcal{B}_i'$ and ϕ_i are continuous with $\phi_i > 0$ on $[a, z] \subset (\hat{\tau}_2, 1]$, the ratios $\mathcal{B}_i'/\phi_i$ are bounded on I. Solving (O.1) gives

$$1 - F_j(b) = \Big(1 - F_j(a)\Big) \exp\left[- \int_a^b \frac{\mathcal{B}_i'(s)}{\phi_i(s)}\, ds \right].$$

The integral stays finite as $b \uparrow z$, so $q_j := 1 - F_j(z^-) > 0$ for both players: each places probability $q_j > 0$ on $[z, 1]$. Write $\bar{V}_i$ for player i's equilibrium payoff; indifference on I and continuity from the left give $\lim_{b\uparrow z} V_i(b) = \bar{V}_i$. We show that no placement of the two residual masses on $[z, 1]$ is consistent with equilibrium, allowing explicitly for a support that resumes in a disjoint interval above z.

No common atom at z. If both players had atoms at z, then for each i, $V_i(z) = \lim_{b\uparrow z} V_i(b) - \frac{1}{2} p_j\, \phi_i(z) < \bar{V}_i$, where $p_j > 0$ is the opponent's atom: bidding z converts a win against that atom into a tie at the same allowance. An atom is a support point and must be optimal, a contradiction.

One support ends at z. Suppose F_j has no mass above z; the displayed solution then forces the terminal atom $p_j = q_j > 0$. Player i's residual $q_i > 0$ lies at z or above. An atom of i at z is excluded by the previous paragraph, so i has support points in $(z, 1]$. Every bid $b > z$ loses against the whole of F_j, so $V_i(b) = \int_{[\underline{\tau}, z)} \mathcal{W}_i\, dF_j + q_j\, \mathcal{W}_i(z)$, a constant, while $\lim_{b\uparrow z} V_i(b) - V_i(b) = q_j\, \phi_i(z) > 0$. Every support point of i above z therefore falls short of $\bar{V}_i$ by the same positive amount,

[O1] The supposition that each $\mathcal{B}_i$ is strictly increasing on $(\hat{\tau}_2, 1]$ is a sharpened form of Condition (A1) of Theorem O.1 below.

wherever it lies; in particular, a support resuming in a disjoint interval above z is suboptimal throughout, a contradiction.

Neither support ends at z. In this case, both players place mass strictly above z. Write $r_j := 1 - F_j(z) > 0$ and $b_i^* := \inf\{\operatorname{supp} F_i \cap (z, 1]\}$. For $b \in (z, b_j^*)$, the distribution F_j is constant, so $V_i(b) = \mathcal{B}_i(b)\, r_j + \int_{[\underline{\tau}, z]} \mathcal{W}_i \, dF_j$, which is strictly increasing in b because $\mathcal{B}_i$ is strictly increasing on $(\hat{\tau}_2, 1]$. Hence, i has no support point in (z, b_j^*), so $b_i^* \geq b_j^*$; by symmetry, $b_1^* = b_2^* =: b^*$. Moreover, $b^* > z$: were $b^* = z$, regularity would give each player an absolutely continuous piece on an interval with left endpoint z, extending the common mixing interval past z and contradicting the maximality of I. At b^*, the analysis at z repeats, with $\phi_i(b^*) > 0$. A common atom at b^* contradicts optimality, as in the first case. If neither player has an atom at b^*, regularity starts an absolutely continuous piece for each player at b^*, creating a common mixing interval with positive residuals above z, which is against the choice of I. If exactly one player, say j, has an atom at b^*, then player i enters $(b^*, \cdot)$ with a density. Player j cannot also carry a density there, and F_j is constant on some $(b^*, b^* + \delta)$: if player j retains mass above b^*, then $V_i(b) = \mathcal{B}_i(b)\big(1 - F_j(b^*)\big) + \text{const.}$ is strictly increasing on the gap, contradicting player i's indifference on a set of positive F_i-measure; if player j does not, then j's support ends at b^* with an atom and the second case applies verbatim at b^*. Hence, a common regular mixing interval above $\hat{\tau}_2$ cannot be completed into an equilibrium. These cases exhaust common regular mixing intervals. □

Theorem O.1 (Regular admissible outcome uniqueness)**.** *Suppose Assumptions 1–3 hold and $\hat{\tau}_1 < \hat{\tau}_2$. Assume in addition that:*

(A1) *for each player i, the mapping $\mathcal{B}_i(\tau) := U_B\big(\mu^j(\tau), \tau\big)$ is strictly increasing on $(\hat{\tau}_i, 1]$;*[O2]

(A2) *strategies are* regular*: each bid distribution has finitely many atoms and an absolutely continuous component supported on a finite union of intervals, with no singular continuous part;*[O3]

(A3) *equilibrium supports are* admissible*: they contain no weakly dominated bids.*

Then, in every Nash equilibrium satisfying these regularity and admissibility

[O2] Assumption 1 makes $\mathcal{B}_i$ weakly increasing on $\mathcal{T}$; the strict increase required here is an additional shape condition, with strict inequality in Assumption 1(ii), taken at the swapped pair, being a sufficient primitive condition by (10).

[O3] Regularity is invoked through one consequence: a support point with no atom whose every right-neighbourhood carries positive probability starts an absolutely continuous piece, because finitely many atoms and finitely many component intervals cannot otherwise supply mass arbitrarily close above it.

requirements, player 1 wins role B with probability one and the winning bid is $\tau^* = \hat{\tau}_2$. Player 2's exact mixing distribution above $\hat{\tau}_2$ need not be unique.

Proof. First, player 2 places no probability on bids below $\hat{\tau}_2$. By Lemma O.2, every such bid is weakly dominated by $\hat{\tau}_2$, and (A3) excludes weakly dominated bids from equilibrium support.

Secondly, player 1 also places no probability on bids below $\hat{\tau}_2$. Fix any $b < \hat{\tau}_2$, and choose $b' \in (\hat{\tau}_1, \hat{\tau}_2)$ with $b' > b$, which is possible because $\hat{\tau}_1 < \hat{\tau}_2$. Since player 2's equilibrium support lies weakly above $\hat{\tau}_2$, player 1 wins role B against exactly the same opponent bids with b' as with b, and the only payoff difference is the allowance. If $b > \hat{\tau}_1$, then the strict increase of $\mathcal{B}_1$ on $(\hat{\tau}_1, 1]$ in (A1) gives $\mathcal{B}_1(b') > \mathcal{B}_1(b)$ directly; if $b \le \hat{\tau}_1$, then any $b'' \in (\hat{\tau}_1, b')$ gives $\mathcal{B}_1(b) \le \mathcal{B}_1(b'') < \mathcal{B}_1(b')$ by weak monotonicity and (A1). In either case, b is not a best response for player 1.

Thus, both supports lie in $[\hat{\tau}_2, 1]$. Let a_i be the lower endpoint of player i's support. If $a_1 < a_2$, then player 1 wins role B for sure at bids near a_1 and can raise his bid slightly without changing that outcome; since $a_1 \ge \hat{\tau}_2 > \hat{\tau}_1$; (A1) makes the deviation profitable. The case $a_2 < a_1$ is identical. Hence $a_1 = a_2 =: a$.

We next show that $a = \hat{\tau}_2$. Suppose instead that $a > \hat{\tau}_2$. A common atom at a is impossible: since $\phi_i(a) > 0$ for both players, either player would strictly gain by undercutting the opponent's atom and winning role B at an allowance arbitrarily close to a. If exactly one player has an atom at a, the other player has support arbitrarily close and above a and can undercut that atom, again profitably. If neither player has an atom at a, then regularity implies that both absolutely continuous supports contain a right-neighbourhood of a, creating a common mixing interval on which both residual probabilities are positive; Proposition O.1 rules this out. Therefore $a > \hat{\tau}_2$ is impossible, and the common lower endpoint is $a = \hat{\tau}_2$.

Suppose F_2 has an atom $p > 0$ at $\hat{\tau}_2$. If player 1 has an atom there, then his payoff is $\mathcal{B}_1(\hat{\tau}_2) - \frac{p}{2}\,\phi_1(\hat{\tau}_2)$; if he is atomless, then his support points $b \downarrow \hat{\tau}_2$ earn $\mathcal{B}_1(\hat{\tau}_2) - p\,\phi_1(\hat{\tau}_2)$ in the limit. In either case, the deviation to $\hat{\tau}_2 - u \in (\hat{\tau}_1, \hat{\tau}_2)$, feasible because $\hat{\tau}_1 < \hat{\tau}_2$, wins with probability one and pays $\mathcal{B}_1(\hat{\tau}_2 - u) \to \mathcal{B}_1(\hat{\tau}_2)$, a strict improvement for small u. Hence, F_2 is atomless at $\hat{\tau}_2$.

It follows that the common lower endpoint $\hat{\tau}_2$ must carry an atom for player 1. Suppose instead that player 1 had no atom at $\hat{\tau}_2$. By the coincidence of the lower endpoints and regularity, both players would then mix with densities on a common interval $(\hat{\tau}_2, \hat{\tau}_2+\delta)$, and player 2's indifference there makes (O.1) with $i = 2$ necessary for F_1: $1 - F_1(b) = \lim_{a' \downarrow \hat{\tau}_2}\big(1 - F_1(a')\big)\exp[-\int_{a'}^{b} \mathcal{B}_2'(s)/\phi_2(s)\,ds]$. If $\int_{\hat{\tau}_2} \mathcal{B}_2'/\phi_2\,ds$ diverges, then, since F_1 has no atom at $\hat{\tau}_2$ and no mass below it, $1 - F_1(b) = 0$ for every $b > \hat{\tau}_2$, so F_1 jumps by one at $\hat{\tau}_2$, contradicting the absence of an atom. If the integral converges, then the common interval extends to a maximal $(\hat{\tau}_2, z)$ with both residual probabilities positive at z^-, and the terminal analysis in the proof of Proposition O.1 applies at $z > \hat{\tau}_2$ unchanged. Hence, player 1 carries an atom at $\hat{\tau}_2$.

Player 1 cannot also have support strictly above $\hat{\tau}_2$. A continuous support piece of player 1 above $\hat{\tau}_2$ that overlaps with player 2's continuous support is ruled out by Proposition O.1. If player 2 has no mass in a neighbourhood of a support point $b > \hat{\tau}_2$ of player 1 but retains mass above it, then a slightly higher bid leaves the winning event unchanged and, by (A1), strictly raises $\mathcal{B}_1$; if player 2 has no mass above b at all, then every player-1 bid at or above b loses surely and earns at most $\mathcal{W}_1(\hat{\tau}_2) < \mathcal{B}_1(\hat{\tau}_2) \leq \bar{V}_1$, where $\bar{V}_1$ is player 1's equilibrium payoff. An atom of player 1 at a point b that player 2's support meets from above is also impossible: player 2's support points just above b earn strictly less than a bid just below b, by approximately $\phi_2(b)$ times the atom, contradicting their optimality. Hence, player 1 bids $\hat{\tau}_2$ with probability one.

Player 2's admissible best responses may place probability on bids weakly above $\hat{\tau}_2$, subject to the condition that player 1 has no profitable upward deviation. These choices affect the losing bid but not the realised outcome. Therefore, every regular admissible equilibrium outcome assigns role B to player 1 at the allowance equal to the winning bid $\tau^* = \hat{\tau}_2$. The exact distribution of player 2's losing bid above $\hat{\tau}_2$ need not be unique. □

Theorem O.1 is a regularity-and-admissibility result, not full Nash uniqueness: irregular supports and singular strategies lie outside its scope. Admissibility does substantive work, because weakly dominated bids can be best replies against opponent strategies that place no probability where the domination is strict; within the constructed family, that is exactly the step Proposition 4 supplies. Equilibria that differ only in player 2's losing-bid distribution do exist: for any two sufficiently small mixing widths below the bound of Lemma O.7, the construction in Proposition 3(c) at $\tau^* = \hat{\tau}_2$ yields regular equilibria with distinct player-2 distributions and the same outcome. Their supports contain no weakly dominated bids. Under (A1), any alternative to a bid at or above a player's own cutoff does strictly worse against some opponent bid.

O.1.5 Proofs of Propositions 7 and 8

Proof of Proposition 7. Take $\hat{\tau}_1 \leq \hat{\tau}_2$. In every equilibrium of the family characterised in Proposition 3, the winning bid satisfies $\tau^* \in [\hat{\tau}_1, \hat{\tau}_2]$ (parts (b)–(c), with role B assigned to player 1 when the cutoffs differ). Under that role assignment, player 1 weakly prefers his own assignment to the swap if and only if $\phi_1(\tau^*) \geq 0$, i.e., $\tau^* \geq \hat{\tau}_1$, and player 2 does so if and only if $\phi_2(\tau^*) \leq 0$, i.e. $\tau^* \leq \hat{\tau}_2$, each ϕ_i being strictly increasing with root $\hat{\tau}_i$ (Proposition 2). Every winning bid in the family satisfies both, which proves the first claim; under the reverse role assignment, the two requirements are $\phi_2 \geq 0 \geq \phi_1$, jointly possible only if $\hat{\tau}_1 = \hat{\tau}_2$. Conversely, Proposition 3(c) sustains every $\tau \in (\hat{\tau}_1, \hat{\tau}_2]$ as the winning bid of some equilibrium in the family; Proposition 3(b) covers coincident cutoffs. □

Proof of Proposition 8. Under Assumption 3, role B is preferred when no handicap is applied ($\tau = 1$) and, by Proposition 6, is an underdog at the bid $\tau^* = \hat{\tau}$, so the equalising allowance is interior: $\hat{\tau} < \tau_{\mathsf{DO}} < 1$. Write $a(\tau) := vg(\Delta - \bar{\mu}(\tau))$ for the common marginal benefit along the symmetric-pair equilibrium. Thus, $c'(e_W) = a$, and $c'(e_B) = a/H(\tau)$. Differentiating, $a'(\tau) = vg'(\Delta - \bar{\mu}(\tau)) \cdot \big(-\hat{e}\,'(\tau)\big)$, with $\hat{e}\,' < 0$ (Proposition 1(ii)), so a' has the sign of $g'(\Delta - \bar{\mu})$, which is non-negative when $\bar{\mu} \geq \Delta$ and, since $g(\Delta - \bar{\mu}) > 0$ at the equilibrium margin, strictly positive when $\bar{\mu} > \Delta$ by strict single-peakedness of g. Then, $e_W' = a'/c''(e_W) \geq 0$, where the inequality is strict when $\bar{\mu} > \Delta$; and $e_B' = \big[a'H - aH'\big]/\big[H^2 c''(e_B)\big] > 0$ on the same region, since $-aH' > 0$. By Proposition 6, $\bar{\mu}(\hat{\tau}) > \Delta = \bar{\mu}(\tau_{\mathsf{DO}})$ under strictly log-concave cost with the equalising allowance interior, and $\bar{\mu}$ is strictly decreasing. So $\bar{\mu} \geq \Delta$ on $[\hat{\tau}, \tau_{\mathsf{DO}}]$, with strict inequality below τ_{DO}. Integrating over that interval gives $e_W(\tau_{\mathsf{DO}}) > e_W(\hat{\tau})$, and $e_B(\tau_{\mathsf{DO}}) > e_B(\hat{\tau})$.

□

O.1.6 Derivations for Section 6.1

Throughout this subsection, we use the Armageddon labels of Section 6.1: White is role W at the full clock, Black is role B at allowance τ. $\Gamma^k(\tau)$ is then the second-period game in which player k holds White at full time and the opponent holds Black at time τ; its equilibrium inputs are (e_W^k, e_B^{-k}), the input gap is $\hat{e}^k = e_W^k - e_B^{-k} > 0$, and White's deterministic advantage is $\mu^k = \mu_0^k + \hat{e}^k$ with $\mu_0^k = A_k - A_{-k} + \Omega$.

Lemma O.3 (Gap-sufficiency)**.** *Fix Ω, the primitives (c, g, H, v, Δ), and the allowance τ. The equilibrium of each game $\Gamma^k(\tau)$ depends on (A_1, A_2) only through $\mu_0^k = A_k - A_{-k} + \Omega$. Consequently, $\phi_i(\tau; A_i, A_j) = \Phi(\tau;\, A_i - A_j)$ for a single function Φ, writing Φ_d for the derivative in the gap argument $d := A_i - A_j$, and*

$$\frac{\partial \phi_i}{\partial A_i} = \Phi_d, \qquad \frac{\partial \phi_i}{\partial A_j} = -\,\Phi_d = -\,\frac{\partial \phi_i}{\partial A_i}.$$

Proof. The first-order conditions and input-gap fixed point depend on abilities only through $\Delta - \mu_0^k - e$, hence through $A_k - A_{-k}$. Thus (e_W^k, e_B^{-k}), $\hat{e}^k$, μ^k, the win probabilities, and the payoffs entering ϕ_i are functions of the ability gap alone. Differentiating $\Phi(\tau; A_i - A_j)$ gives the two partials. $\square$

Lemma O.4 (Ability decomposition of ϕ_i)**.** *Fix the opponent j, the allowance τ, the rival ability A_j, and the primitives. Along the equilibria of $\Gamma^i(\tau)$ and $\Gamma^j(\tau)$,*

$$\frac{d\phi_i(\tau)}{dA_i} = v\Big[g(\Delta - \mu^j) - g(\Delta - \mu^i)\Big] \; - \; v\, g(\Delta - \mu^j)\,\frac{de_W^j}{dA_i} \; + \; v\, g(\Delta - \mu^i)\,\frac{de_B^j}{dA_i},$$

where $g(\Delta - \mu^i), g(\Delta - \mu^j)$ are the noise densities at the decisive arguments of Γ^i, Γ^j, de_W^j/dA_i is rival White's input response in Γ^j, and de_B^j/dA_i rival Black's input response in Γ^i.

Proof. Write, as in the proof of Proposition 3, $\mathcal{W}_i(\tau) = U_W(\mu^i(\tau), \tau)$ and $\mathcal{B}_i(\tau) = U_B(\mu^j(\tau), \tau)$, so that $\phi_i = \mathcal{B}_i - \mathcal{W}_i$. For player i as White in Γ^i, the envelope cancellation gives

$$\frac{d\mathcal{W}_i}{dA_i} = vg(\Delta - \mu^i)\Big(1 - \tfrac{de_B^j}{dA_i}\Big).$$

For player i as Black in Γ^j, A_i enters μ^j with the opposite sign, and the Black first-order condition gives

$$\frac{d\mathcal{B}_i}{dA_i} = vg(\Delta - \mu^j)\Big(1 - \tfrac{de_W^j}{dA_i}\Big).$$

Indeed, in Γ^j, an increase in A_i lowers White's deterministic advantage μ^j, thereby raising player i's Black win probability. Subtracting $d\mathcal{W}_i/dA_i$ from $d\mathcal{B}_i/dA_i$ gives $d\phi_i/dA_i$ as in the lemma. Own input drops out in each game; the rival's input responses remain because they are chosen in the opposite role-assignment games. $\square$

Lemma O.5 (Signs of the rival contest-input responses)**.** *Maintain Condition 1. Differentiating the equilibrium first-order conditions within each game in A_i,*

$$\frac{de_W^j}{dA_i} = \frac{v\, g'(\Delta - \mu^j)\, H(\tau)\, c''(e_B^i)}{D_j}, \qquad \frac{de_B^j}{dA_i} = \frac{-\,v\, g'(\Delta - \mu^i)\, c''(e_W^i)}{D_i},$$

with $D_j, D_i > 0$ under Condition 1. Hence, $\operatorname{sign}(de_W^j/dA_i) = \operatorname{sign} g'(\Delta - \mu^j)$ *and* $\operatorname{sign}(de_B^j/dA_i) = -\operatorname{sign} g'(\Delta - \mu^i)$*. Since g is single-peaked at 0, $g'(\Delta - \mu^k)$ is non-positive when $\mu^k < \Delta$ and non-negative when $\mu^k > \Delta$; the inequalities are strict on regions where the density is strictly decreasing away from zero.*

Proof. In game Γ^j, player j is White and i Black, with $c'(e_W^j) = H(\tau)c'(e_B^i) = vg(\Delta - \mu^j)$, and $\partial\mu^j/\partial A_i = -1$. Write $p = de_W^j/dA_i$, $q = de_B^i/dA_i$. Differentiating both equalities,

$$c''(e_W^j)\,p = vg'(\Delta - \mu^j)(1 - p + q), \qquad H(\tau)c''(e_B^i)\,q = vg'(\Delta - \mu^j)(1 - p + q).$$

Solving gives the first formula with

$$D_j = H(\tau)c''(e_B^i)c''(e_W^j) + vg'(\Delta - \mu^j)[H(\tau)c''(e_B^i) - c''(e_W^j)].$$

The same calculation in game Γ^i, where $\partial\mu^i/\partial A_i = 1$, gives the second formula with

$$D_i = H(\tau)c''(e_B^j)c''(e_W^i) + vg'(\Delta - \mu^i)[H(\tau)c''(e_B^j) - c''(e_W^i)].$$

Condition 1 implies $|vg'(\Delta - \mu)| \le v\dot{g} \le \lambda m < m \le c''(e)$, so $D_i, D_j > 0$. $\square$

Lemma O.6 (Cutoff response)**.** *Maintain the assumptions of Proposition 2 and Assumption 3. For each i, $\hat{\tau}_i$ is continuously differentiable in A_i with*

$$\frac{d\hat{\tau}_i}{dA_i} = -\frac{\partial\phi_i/\partial A_i}{\partial\phi_i/\partial\tau}, \qquad \operatorname{sign}\Big(\frac{d\hat{\tau}_i}{dA_i}\Big) = -\operatorname{sign}\Big(\frac{\partial\phi_i}{\partial A_i}\Big).$$

The magnitude $|d\hat{\tau}_i/dA_i|$, equal to $|\partial\phi_i/\partial A_i|/(\partial\phi_i/\partial\tau)$, is nonnegative, and strictly positive when $\partial\phi_i/\partial A_i \neq 0$.

Proof. ϕ_i is C^1 with $\partial\phi_i/\partial\tau > 0$ by Proposition 2. The implicit-function theorem applied to $\phi_i(\hat{\tau}_i; A_i) = 0$ gives the displayed expression; the sign statement and the magnitude formula follow since the denominator is strictly positive; the magnitude vanishes exactly where the numerator does. $\square$

Lemma O.7 (Sustainable mixing width)**.** *Maintain the assumptions of Proposition 3, with player 1 the bidding winner (smaller cutoff) bidding τ^* with certainty and player 2 mixing uniformly on $[\tau^*, \tau^* + \delta]$. Player 1's expected payoff*

$$EU_1(b) = \int_{\tau^*}^{b} U_W(\mu^1(s), s)\, f_2(s)\, ds + \Big(1 - F_2(b)\Big)\, U_B(\mu^2(b), b)$$

has derivative $dEU_1/db = -f_2(b)\,\phi_1(b) + \Big(1 - F_2(b)\Big)\frac{d}{d\tau}U_B(\mu^2(b), b)$. The lower-endpoint no-deviation condition is

$$\delta < \bar{\delta}(\tau^*) := \frac{\phi_1(\tau^*)}{\frac{d}{d\tau}U_B(\mu^2(\tau^*), \tau^*)}.$$

We interpret this ratio in the extended sense: if the denominator is zero, then $\bar{\delta}(\tau^*) = +\infty$. A sufficient condition for τ^* to be player 1's best response on the whole support is

$$\delta < \frac{\phi_1(\tau^*)}{\sup_{b\in[\tau^*,\tau^*+\delta]} \frac{d}{d\tau} U_B(\mu^2(b), b)}.$$

When the supremum in the denominator is zero, this condition is understood as vacuous. This sufficient condition coincides with the endpoint expression $\bar{\delta}(\tau^*)$ whenever $\frac{d}{d\tau}U_B(\mu^2(b), b) \leq \frac{d}{d\tau}U_B(\mu^2(\tau^*), \tau^*)$ on the mixing interval.

Proof. On the support, we have $\phi_1 > 0$ (since $\tau^* > \hat{\tau}_1$) and $\frac{d}{d\tau}U_B \geq 0$ (by Assumption 1). With $f_2 = 1/\delta$ and $F_2(b) = (b - \tau^*)/\delta$, the derivative at the lower endpoint is $dEU_1/db = -\phi_1(\tau^*)/\delta + \frac{d}{d\tau}U_B(\mu^2(\tau^*), \tau^*) < 0$ if and only if $\delta < \bar{\delta}(\tau^*)$, with the inequality automatic if the derivative term is zero. For any $b \in [\tau^*, \tau^* + \delta]$,

$$\frac{dEU_1}{db} \leq -\frac{\phi_1(\tau^*)}{\delta} + \sup_{x\in[\tau^*,\tau^*+\delta]} \frac{d}{d\tau} U_B(\mu^2(x), x),$$

because ϕ_1 is increasing and $1 - F_2(b) \leq 1$. Hence, the displayed supremum condition makes EU_1 strictly decreasing on the support. If the derivative of U_B is no larger than its lower-endpoint value on the interval, then this condition reduces to $\delta < \bar{\delta}(\tau^*)$. Evaluated at the undominated bid $\tau^* = \hat{\tau}_2$, using $\phi_1(\hat{\tau}_2) = \int_{\hat{\tau}_1}^{\hat{\tau}_2} \phi_1'(s)\, ds$ ($\phi_1(\hat{\tau}_1) = 0$, $\phi_1' > 0$), the bound is zero at $A_1 = A_2$ and, near equality, opens at rate $2|d\hat{\tau}/dA|$ with leading-order integrand ϕ_1', so $\phi_1(\hat{\tau}_2) \approx 2\,|\partial\phi/\partial A|\,|A_1 - A_2|$—governed by the magnitude, not the sign, of the ability response, and zero at first order when that response vanishes. $\square$

O.2 Additional empirical results

Here we present the reduced-form tables behind Section 7.1, the structural inference and robustness material behind Sections 7.3–7.4, and the descriptive material behind Section 5.

Reduced-form tables. The tables behind Section 7.1 follow. Table 4 reports the full specification ladder behind the headline reduced-form estimate. Table 6 gives the share of matches in which the higher-rated player takes Black, overall and by subsample; every interval straddles one half. Table 7 adds the reversal check, logits of that propensity on the absolute gap, the clock, and their interaction. The winning-bid regressions behind the gap-sufficiency and regime-mixing readings are in Table 8. Tables 9 and 10 report the bid-spread contrasts and the spread profile by absolute-gap quintile. Table 11 closes with the bid winner's match-win rate by symmetry window, together with each window's minimum detectable deviation from one half.

Table 4: H1.A. Logit of Pr(Black wins the Armageddon) on the signed FIDE standard Elo gap (White − Black, per 100 points). Specifications: (1) baseline; (2) adds base-time fixed effects; (3) adds stakes and winning-bid quintile fixed effects; (4) adds event fixed effects; (5) repeats (3) with session-level clustering. Event-clustered standard errors elsewhere. $N = 477$; column (4) uses $N = 471$ after dropping singleton event cells.

	(1)	(2)	(3)	(4)	(5)
Elo gap (White − Black) / 100	-0.289**	-0.283**	-0.277**	-0.309**	-0.277**
	(0.133)	(0.138)	(0.136)	(0.137)	(0.113)
Stakes / 100			0.625	0.977**	0.625*
			(0.426)	(0.449)	(0.377)
Base time = 600s		0.374	0.353	0.340	0.353
		(0.378)	(0.316)	(0.388)	(0.436)
Base time = 900s		0.405	0.419		0.419
		(0.386)	(0.369)		(0.450)
Intercept	0.008	-0.361	-0.787*		-0.787
	(0.094)	(0.343)	(0.442)		(0.503)
Tau-min quintile 2			0.120	0.154	0.120
			(0.223)	(0.219)	(0.333)
Tau-min quintile 3			0.266	0.327	0.266
			(0.217)	(0.235)	(0.312)
Tau-min quintile 4			0.376	0.370	0.376
			(0.358)	(0.393)	(0.307)
Tau-min quintile 5			0.229	0.416	0.229
			(0.351)	(0.383)	(0.347)
Num.Obs.	477	477	477	471	477
R2	0.012	0.013	0.020	0.047	0.020
R2 Adj.	0.009	0.004	-0.004	-0.018	-0.004

* p < 0.1, ** p < 0.05, *** p < 0.01

Table 5: Few-cluster inference for H1.A. Both procedures are null-imposed with Rademacher weights and $B = 9{,}999$ draws; the wild cluster bootstrap-t operates on the linear-probability version of specification (3), the effective-score wild bootstrap on the logit itself.

	Statistic	p-value
Logit, event-clustered (paper)	−0.277 (s.e. 0.136)	0.042
LPM, event-clustered	$t = -2.10$	0.049
LPM, wild cluster bootstrap-t ($B = 9999$)		0.107
Logit, effective-score wild bootstrap ($B = 9999$)	$t = -1.63$	0.098

Structural inference and robustness. Figure 2 shows the geometry behind the inference treatment of Section 7.4: the block-bootstrap draws populate the common-

Table 6: H1.B. Share of matches in which the higher-rated player wins the bid and plays Black, overall and by subsample. Exact binomial 95% confidence intervals; two-sided p-values against one half. $N = 477$.

Sample	N	Share	95% CI	p (vs $\frac{1}{2}$)
All pairs	477	0.526	[0.480, 0.572]	0.272
$\lvert gap \rvert \leq$ median	240	0.533	[0.468, 0.598]	0.333
$\lvert gap \rvert >$ median	237	0.519	[0.453, 0.584]	0.603
Base time 300s	25	0.480	[0.278, 0.687]	1.000
Base time 600s	228	0.496	[0.429, 0.562]	0.947
Base time 900s	224	0.562	[0.495, 0.628]	0.071
First half (by date)	243	0.519	[0.454, 0.583]	0.608
Second half (by date)	234	0.534	[0.468, 0.599]	0.327

Table 7: H1.B. Logit of Pr(higher-rated player plays Black) on the absolute Elo gap (per 100 points), base time, and their interaction. Event-clustered standard errors. $N = 477$.

	(1) Baseline	(2) + base time	(3) Gap × clock	(4) Event FE
Absolute Elo gap / 100	-0.170	-0.173	-0.543	-0.180
	(0.129)	(0.122)	(0.353)	(0.128)
Base time = 600s		0.041	-0.166	-0.439*
		(0.335)	(0.535)	(0.225)
Base time = 900s		0.315	-0.157	
		(0.313)	(0.419)	
Abs. gap × base 600s			0.235	
			(0.462)	
Abs. gap × base 900s			0.597	
			(0.387)	
Intercept	0.230	0.065	0.368	
	(0.155)	(0.306)	(0.394)	
Num.Obs.	477	477	477	469
R2	0.001	0.005	0.007	0.025
R2 Adj.	-0.002	-0.004	-0.008	-0.021

* p < 0.1, ** p < 0.05, *** p < 0.01

Table 8: H1.C. OLS of the winning bid τ^* (share of base time) on the signed ability gap (Black − White, per 100 Elo), the ability level, stakes, and base-time controls. Event-clustered standard errors; column (4) clusters by session. $N = 477$.

	(1) No stakes	(2) Stakes	(3) Absolute gap	(4) Session cl.
Elo gap (Black − White) / 100	0.006*	0.005		0.005
	(0.003)	(0.003)		(0.004)
Absolute Elo gap / 100			-0.014	
			(0.009)	
(Average Elo − 2650) / 100	-0.013	-0.006	-0.007	-0.006
	(0.011)	(0.013)	(0.014)	(0.007)
Stakes / 100		-0.042**	-0.041**	-0.042***
		(0.017)	(0.017)	(0.015)
Base time = 600s	-0.052***	-0.056***	-0.056***	-0.056***
	(0.018)	(0.016)	(0.016)	(0.021)
Base time = 900s	-0.102***	-0.102***	-0.101***	-0.102***
	(0.019)	(0.016)	(0.016)	(0.022)
Intercept	0.809***	0.825***	0.835***	0.825***
	(0.018)	(0.019)	(0.019)	(0.022)
Num.Obs.	477	477	477	477
R2	0.155	0.170	0.175	0.170
R2 Adj.	0.148	0.161	0.166	0.161

* p < 0.1, ** p < 0.05, *** p < 0.01

Table 9: H1.D. Within-pair bid spread and the absolute ability gap: lowest-|gap|-quintile contrasts, in shares of base time (1) and seconds (2); spread on |gap|/100 (3); local regression on the $|\text{gap}| \leq$ median subsample (4). Event-clustered standard errors. $N = 477$ ($N = 240$ in column (4)).

	(1)	(2)	(3)	(4)
Absolute Elo gap / 100			0.006	0.053***
			(0.007)	(0.015)
Base time = 600s	0.003	28.391***	0.006	
	(0.010)	(3.989)	(0.010)	
Base time = 900s	0.007	60.819***	0.011	
	(0.011)	(5.951)	(0.011)	
Lowest |gap| quintile	-0.016*	-10.932		
	(0.009)	(6.775)		
Intercept	0.095***	30.696***	0.084***	0.076***
	(0.009)	(3.592)	(0.011)	(0.007)
Num.Obs.	477	477	477	240
R2	0.007	0.081	0.003	0.013
R2 Adj.	0.000	0.075	-0.004	0.009

* p < 0.1, ** p < 0.05, *** p < 0.01

Table 10: H1.D. Within-pair bid spread $|\tau_1 - \tau_2|$ by absolute Elo-gap quintile, in seconds and as a share of base time. $N = 477$.

		Seconds		Share of base time	
|Gap| quintile (Elo)	N	Mean	Median	Mean	Median
Q1 (1–25)	96	58.7	45.0	0.084	0.067
Q2 (25–51)	96	73.8	51.5	0.100	0.068
Q3 (51–78)	95	76.9	55.0	0.102	0.073
Q4 (78–118)	95	72.5	58.0	0.098	0.083
Q5 (118–251)	95	71.4	55.0	0.100	0.075

Table 11: H2.B. Match-win rate of the bid winner (Black) by symmetry window. Exact binomial 95% confidence intervals; two-sided p-values against one half; MDE is the minimum detectable deviation from one half at 80% power (two-sided 5%). $N = 477$.

Window	N	P(Black wins)	95% CI	p (vs $\frac{1}{2}$)	MDE
All pairs	477	0.503	[0.457, 0.549]	0.927	0.064
$\|gap\| \leq 63$ (median)	240	0.496	[0.431, 0.561]	0.949	0.090
$\|gap\| \leq 29$ (25th pct)	121	0.471	[0.380, 0.564]	0.586	0.127
$\|gap\| < 25$	93	0.484	[0.379, 0.590]	0.836	0.145

scale valley rather than scattering around the reported point. Table 12 collects the functional-form battery behind the robustness discussion of Section 7.4. Table 13 reports the mechanism comparison of Section 7.3 at the alternative institutional floor $\underline{\tau} = 0.25$.

The likelihood surface shapes the search itself. The indicator factors make the surface a washboard. Plain perturbation multistarts land tens of log-likelihood points above the optimum, trapped in grooves. The profile grid is what explores the surface. Fixing one coordinate and re-optimising the rest is exactly the move that walks the valley. We therefore treat the profile chain as the global search. The reported point of Table 3 is the best it has found. The bootstrap run behind Figure 2 is mechanically clean. Every draw re-optimised away from its warm start. None stuck there. Pointwise profiling of the parameter levels fails on this surface. Five level profiles excluded known near-optimal points. The main text accordingly reports no level profiles.

Three details complete the record of the battery. Each specification in Table 12 is re-fitted from the reported point embedded exactly. Every alternative therefore nests the baseline, which makes the likelihood ratios meaningful. The S8 spline gain is an observed-White bonus in the absolute gap: +80.9 points in the colour factor against a 17.8-point worsening of the outcome factor. Refining the integration grid cuts the S6 curvature likelihood ratio of 3.52 to about one. At the baseline

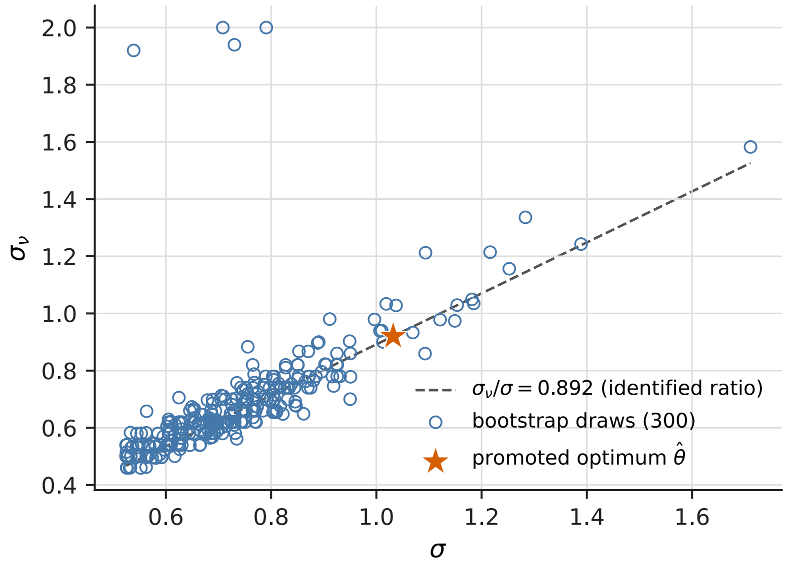

Figure 2: Block-bootstrap draws in the (σ, σ_ν) plane. Each of the 300 player-pair resamples is re-estimated from the reported point (star); the dashed ray marks the identified ratio $\sigma_\nu/\sigma = 0.892$. The draws string along the ray while the common scale slides.

specification the finer 801-node grid moves the total log-likelihood from −2856.12 to −2857.41 and changes no conclusion.

Table 12: Functional-form robustness of the ability factor: the specification battery, each alternative re-fitted from the reported point embedded exactly, 477 games.

Specification	Added df	LR	Bid factor	Verdict
S2: wedge scale in the rating level	1	0.15	+0.12	no channel
S3: wedge scale in the absolute gap	1	0.15	+0.12	no channel
S4: β_E fixed at the outcome value	−1	5.20	−4.48	rejected ($p = 0.023$)
S6: signed-quadratic curvature term	1	3.52	−0.59	flat
S7: average-rating level term	1	1.20	+0.77	flat
S8: unrestricted spline in the gap	4	211.58	+42.67	diagnostic
S8-odd: antisymmetric spline in the gap	4	3.3	–	flat

Note: LR is twice the log-likelihood gain over the baseline ($\log L = -2856.12$) from the embedded start; S4 fixes β_E at the outcome-implied +0.00107 rather than adding a parameter, so its LR prices one restriction. The bid-factor column is the chain-decomposition change in the bid-interval factor, in log-likelihood points. The S8 fit is a specification diagnostic, not an ability effect: its shape is an even function of the gap, an observed-White bonus concentrated in the colour factor (Section 7.4). S8-odd re-fits the same spline with antisymmetry imposed, the only admissible shape for an ability effect.

Table 13: Mechanism comparison at the alternative institutional floor: mean colour gap $\Delta\rho^M$, percentage points, at the structural estimates, $\underline{\tau} = 0.40$ against $\underline{\tau} = 0.25$. $N = 477$.

	Conditional on FIDE ($\nu = 0$)		Integrated over ν	
Mechanism	$\underline{\tau} = 0.40$	$\underline{\tau} = 0.25$	$\underline{\tau} = 0.40$	$\underline{\tau} = 0.25$
Random colour (RC)	6.04	6.04	46.49	46.49
Fixed odds	2.52	2.52	46.38	46.38
Bidding (BM, model bid)	3.49	3.49	47.60	47.60
Designer (DO_{τ})	0.00	0.00	40.51	39.96

O.2.1 The data

This section collects the descriptive material relegated from Section 5: the composition of the sample, the comparison between the standard and rapid FIDE ratings, the construction of the stake index, and the bid distributions by ability gap.

Sample composition and ratings. Panel A of Table 14 reports the distribution of base times, the number of unique pseudonymous players, and how games split across tertiles of average FIDE standard rating. Panel B cross-classifies games by tertiles of the White and Black standard ratings at the list date, on equal-count bins in each margin. Panel C reports the rapid-rating moments referred to in Section 5: their closeness to the standard-rating moments is what licenses using the standard rating throughout.

The stake index. We map each game's position in the bracket to a round-progression score on a 1–100 scale, rising with the consequence of the round: a grand-final reset scores 100, a semifinal roughly 70–80, a first knock-out round roughly 20, and a qualifier roughly 5. The score is then multiplied by a division weight (Division I $\times 1.2$, Division II $\times 1.0$, Division III $\times 0.9$) and clipped to $[1, 100]$. In the analysis sample the median is 22.8 and the inter-quartile range is [17.1, 50.0], which reflects the uneven coverage of the high-stakes rounds.

Bid distributions by ability gap. Figure 3 partitions the sample into six equal-count sextiles of the signed FIDE Elo gap (White minus Black) and, within each sextile, overlays histograms of the realised Black (auction-winning) and White (auction-losing) time shares. Within a match the Black share is mechanically the lower of the two, so the vertical separation of the two histograms records the within-match spread in realised offers, not a free-standing test of equilibrium selection. What the figure does show is that the centre and the spread of both cross-sectional marginals

Table 14: Sample composition and rating tertiles. Panel A: base times, unique players, and games by average-Elo tertile. Panel B: games by tertiles of White and Black FIDE standard ratings (equal-count bins). Panel C: FIDE rapid ratings, computed on the games in which rapid ratings are available for both players.

Panel A. Sample composition

Base time (s)		Players / avg. strength	
300 (5 min)	25	Unique players	204
600 (10 min)	228	Games, low avg. Elo tertile	159
900 (15 min)	224	Games, mid avg. Elo tertile	159
Total games	**477**	Games, high avg. Elo tertile	159

Panel B. FIDE standard rating tertiles (White × Black, equal-count bins)

	Black FIDE standard (tertile)			
	Black T1	Black T2	Black T3	Total
White FIDE T1 (low)	74	54	31	159
White FIDE T2	50	60	49	159
White FIDE T3 (high)	35	45	79	159
Total	159	159	159	477

Panel C. FIDE rapid rating (N = 475 White, 474 Black, 472 gap)

Variable	Mean	SD	Median	Min	Max
White Elo	2618	97.9	2627	2024	2839
Black Elo	2617	110.1	2629	2102	2839
Elo gap (W−B)	1	113.2	−3	−344	325

move only modestly across sextiles, with most of the mass in a comparable middle band, even as the signed advantage changes sign across the panels. This is the flatness reported in Section 5.

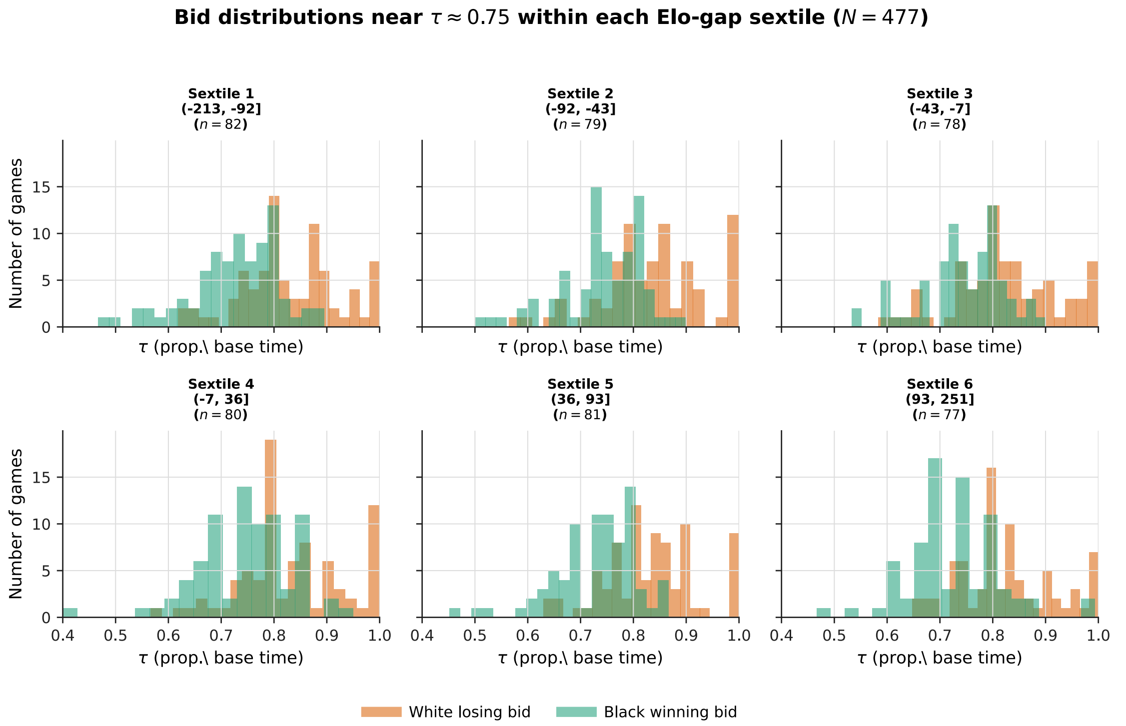

Figure 3: Bid distributions conditional on the signed FIDE Elo gap (White minus Black), full sample ($N = 477$). Six equal-count sextiles on the signed gap. Orange bars: losing (White) bid; green bars: winning (Black) bid. Elite Chess.com Armageddon games, 2022-08 to 2024-10.

O.3 A parametric specification of the handicapped continuation contest

In this appendix we present a parametric specification of the handicapped continuation contest with a log-concave triangular noise density. The triangular density is not globally smooth, so the example should be read as a local closed-form benchmark on the maintained central branch. On that branch, it delivers closed-form equilibrium supply of the input and payoffs and verifies the key monotonicity properties used in the main text.

O.3.1 Primitives

We retain the basic structure from Section 3 with the following definitions.

Assumption O.1. *In the handicapped contest:*

(i) The shock to the score difference, ε, is distributed symmetrically on $[-1, 1]$

with triangular density

$$g(x) = \begin{cases} 1 - |x|, & x \in (-1, 1), \\ 0, & \textit{otherwise}, \end{cases}$$

and cumulative distribution function

$$G(x) = \begin{cases} 0, & x \leq -1, \\ \dfrac{(x+1)^2}{2}, & -1 < x \leq 0, \\ 1 - \dfrac{(1-x)^2}{2}, & 0 < x < 1, \\ 1, & x \geq 1. \end{cases}$$

This density is symmetric, unimodal at $x = 0$, and log-concave.

(ii) *The cost of producing the input is quadratic:*

$$c(e) = \frac{1}{2}e^2, \qquad c'(e) = e, \qquad c''(e) = 1.$$

(iii) *The handicap specification for role B is*

$$H(\tau) = \frac{1}{\tau}, \qquad \tau \in (0, 1],$$

so that a larger allowance (higher τ) reduces B's marginal cost of supplying input.

(iv) *The parameters (Δ, μ_0, v) and the feasible set of allowances $\mathcal{T} \subset (0, 1]$ are such that*

$$\Delta - \mu(\tau) \in (0, 1) \quad \textit{for all } \tau \in \mathcal{T},$$

where $\mu(\tau)$ denotes the equilibrium mean of the score difference X defined below. This ensures that we remain in the region where $g(\Delta - \mu) = 1 - (\Delta - \mu)$ and $G(\Delta - \mu) = 1 - \frac{1}{2}(1 - \Delta + \mu)^2$.

The last restriction is purely technical and guarantees that the triangular density operates on its central branch for all relevant allowances.

O.3.2 Equilibrium contest inputs and mean score difference

As in the main text, the score difference is

$$\mu = \mu_0 + e_W - e_B,$$

and the equilibrium win probabilities are

$$\rho_W = 1 - G(\Delta - \mu), \qquad \rho_B = G(\Delta - \mu),$$

so that $\rho_W + \rho_B = 1$. The expected payoff to each role is

$$U_W = \rho_W v - c(e_W), \quad U_B = \rho_B v - H(\tau)c(e_B).$$

The first-order conditions for supplying optimal input are

$$g(\Delta - \mu)v = c'(e_W) = H(\tau)c'(e_B).$$

Under Assumption O.1 and the restriction $\Delta - \mu(\tau) \in (0, 1)$, we have

$$g(\Delta - \mu) = 1 - (\Delta - \mu) = 1 - \Delta + \mu.$$

With quadratic costs, $c'(e) = e$, so the FOCs reduce to

$$e_W = v\Big(1 - \Delta + \mu\Big), \qquad e_B = \tau\, v\Big(1 - \Delta + \mu\Big). \tag{O.2}$$

Substituting into the definition of μ gives

$$\mu = \mu_0 + e_W - e_B = \mu_0 + v(1 - \Delta + \mu)(1 - \tau).$$

Solving for μ yields the closed-form score difference

$$\mu(\tau) = \frac{\mu_0 + v(1-\tau)(1-\Delta)}{1 - v(1-\tau)}. \tag{O.3}$$

Define

$$D(\tau) := 1 + v(\tau - 1) = 1 - v(1 - \tau).$$

For $\tau \in \mathcal{T}$, we assume $D(\tau) > 0$, which holds for example whenever $v < 1$. Substituting (O.3) back into (O.2) yields equilibrium inputs

$$e_W^*(\tau) = \frac{v\big(1+\mu_0-\Delta\big)}{D(\tau)}, \qquad e_B^*(\tau) = \frac{\tau v\big(1+\mu_0-\Delta\big)}{D(\tau)}. \tag{O.4}$$

In particular,

$$\frac{\partial e_W^*}{\partial \tau} = \frac{v^2(\Delta-\mu_0-1)}{D(\tau)^2}, \qquad \frac{\partial e_B^*}{\partial \tau} = \frac{v(1+\mu_0-\Delta)(1-v)}{D(\tau)^2}.$$

For parameter values satisfying $1+\mu_0-\Delta>0$ and $v<1$, we have $\partial e_B^*/\partial\tau>0$ and $\partial e_W^*/\partial\tau<0$. Thus, in this parametric specification, role B's equilibrium supply of the input is strictly increasing in his allowance τ, while role W's input is strictly decreasing in τ.

O.3.3 Win probabilities and payoffs

Let

$$x(\tau) := \Delta - \mu(\tau).$$

Under Assumption O.1, we have $x(\tau)\in(0,1)$ for all $\tau\in\mathcal{T}$, and thus

$$G(x(\tau)) = 1-\frac{(1-x(\tau))^2}{2}, \qquad g(x(\tau)) = 1-x(\tau).$$

Using (O.3), it is straightforward to verify that

$$x(\tau) = \frac{\Delta-\mu_0+v(\tau-1)}{D(\tau)}, \qquad g(x(\tau)) = \frac{1+\mu_0-\Delta}{D(\tau)}.$$

Hence, the win probabilities can be written as

$$\rho_B(\tau) = G(x(\tau)) = 1-\frac{1}{2}\left(1-\frac{\Delta-\mu_0+v(\tau-1)}{D(\tau)}\right)^2, \qquad \rho_W(\tau) = 1-\rho_B(\tau),$$

and the payoffs are

$$U_W(\tau) = \rho_W(\tau)v - \frac{(e_W^*(\tau))^2}{2}, \qquad U_B(\tau) = \rho_B(\tau)v - \frac{1}{\tau}\frac{(e_B^*(\tau))^2}{2}.$$

Although the expressions for ρ_R are non-linear in τ, the derivatives of the payoffs simplify considerably. Using (O.4) and the fact that $D(\tau)>0$, a short calculation shows

$$\frac{\partial U_B}{\partial \tau} = \frac{v^2\big(\Delta-\mu_0-1\big)^2\big(\tau v+v+1\big)}{2\,D(\tau)^3} > 0, \tag{O.5}$$

and

$$\frac{\partial U_W}{\partial \tau} = \frac{v^2(v-1)\Big(\Delta - \mu_0 - 1\Big)^2}{D(\tau)^3} < 0 \quad \text{for } v < 1. \tag{O.6}$$

Thus, under the mild condition $v < 1$, relaxing role B's constraint by increasing τ strictly increases his expected payoff and strictly decreases role W's expected payoff. This verifies the payoff-monotonicity content of Assumption 1 on the maintained central branch of a fully specified log-concave example.

O.3.4 The advantage of role *B*

In this symmetric parametric specification the gain from choosing role B rather than role W at allowance τ is

$$\phi(\tau) = U_B(\tau) - U_W(\tau).$$

From (O.5) and (O.6) we obtain

$$\frac{\partial \phi}{\partial \tau} = \frac{\partial U_B}{\partial \tau} - \frac{\partial U_W}{\partial \tau} > 0 \quad \text{for all } \tau \in \mathcal{T}.$$

Hence, $\phi(\tau)$ is continuous and strictly increasing in τ, and the (unique) indifference point $\hat{\tau}$ solving $\phi(\hat{\tau}) = 0$ exists whenever ϕ changes sign on $\mathcal{T}$. This verifies the payoff-monotonicity content of Proposition 2 on the maintained branch of the parametric example and yields an explicit benchmark in which the qualitative comparative statics of the main model can be seen transparently.