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

Derek J. Clark[a] Tapas Kundu[b]

Øystein Myrland[c] Tore Nilssen[d]

August 25, 2026

Abstract

We study a mechanism-design problem in which contestants bid up-front for an advantaged role by forfeiting part of the productive resource in the contest. Since the payment changes the contest technology, bidding prices payoff indifference rather than the allowance that an informed designer would choose to equalise winning probabilities. Within a family of equilibria, weak dominance selects a unique outcome: the winning bid equals the losing bidder's indifference point. Bidding leads to an allocation that no contestant would swap, but does not equalise win probabilities. Under strictly log-concave effort costs, bidding overshoots the equalising allowance for near-equal contestants, making the auction winner an underdog. Armageddon chess provides a direct implementation: players bid the time they will accept as Black with draw odds. Using 477 elite online games, we test the model and structurally compare bidding with random colour, a fixed allowance, and the informed-designer optimum. Ratings predict game outcomes but not bids; bidding improves balance for near-equal pairs but can underperform random assignment once latent heterogeneity is incorporated.

Keywords: mechanism design; contest design; auctions with externalities; endogenous handicaps; role assignment; productive constraints; tie-breaking; structural estimation; Armageddon chess.

JEL codes: C72, D44, D47, D74, L83.

*We thank seminar participants at the Indian Institute of Management Calcutta for comments. We are grateful to the Chess.com company for providing data and to Åsne Nilsen for research assistance. The authors declare no conflicts of interest.

[a]School of Business and Economics, UiT – The Arctic University of Norway. Email: derek.clark@uit.no.

[b]Oslo Business School, Oslo Metropolitan University. Email: tapas.kundu@oslomet.no.

[c]School of Business and Economics, UiT – The Arctic University of Norway. Email: oystein.myrland@uit.no.

[d]Department of Economics, University of Oslo. Email: tore.nilssen@econ.uio.no.

1 Introduction

At the 2010 US Chess Championship, Gata Kamsky and Yury Shulman walked up to the board not to play but to bid (McClain, 2010). They were tied after the regular-play tournament. In front of the cameras, each wrote on a slip the number of minutes he was willing to spend playing Black with draw odds, against the other player playing White with the full clock. Kamsky bid lower; he played Black with reduced time, drew, and took the title. The mechanism is now a standard one at the elite level: a one-shot, sealed-bid, first-price auction in which bidders pay in clock time. The payment is a forfeited share of the resource allowance, not a monetary transfer; it changes the technology of the contest that follows.

The economic interest of this mechanism extends well beyond chess: how should a designer compensate for an asymmetric advantage when she lacks the information required to set the appropriate allowance? Competitive institutions often need to allocate an asymmetric role or resolve an otherwise indecisive contest—most starkly in tie-breaking rules, where the designer wants a decisive outcome without introducing a systematic advantage for either contestant, but does not have information about the contestants' relative strengths.[1] A natural solution is to let the contestants price the advantage themselves. This places the mechanism within a broader class of auctions with allocative externalities: environments in which a bidder's payoff depends on the auction's outcome and not merely on his own allocation, because the auction is embedded in a larger transaction (Jehiel and Moldovanu, 2006). Our mechanism sharpens this general structure: the size of the winning bid, not merely the identity of the winner, resets the terms of the contest that the two rivals subsequently play against each other. Sports institutions supply the cleanest instances of this structure: in Go, auction komi has players bid the score compensation they concede in order to obtain the first-move advantage (Deng and Qi, 2009); in American football, Che and Hendershott (2008) propose auctioning first possession in overtime using field position as the bid, and Granot and Gerchak (2014) study auctions with positive externalities and applications to overtime in football, soccer, and chess. Armageddon bidding belongs to this same family: contestants bid directly in the allowance, appearing to reveal the information the designer lacks.[2] Non-monetary, technology-setting bids also appear outside con-

[1] Equalising winning chances can also be motivated on grounds other than fairness. A contest-design literature shows that levelling the playing field can strengthen competitive incentives and raise aggregate effort (Lazear and Rosen, 1981; Nalebuff and Stiglitz, 1983; Chowdhury et al., 2023); this effort-based rationale applies to contests generally, whereas the one we stress is specific to a tiebreaker, invoked only once regular play has failed to separate the contestants.

[2] When contestants' relative strengths are known, the equalising allowance can be set directly

tests. In highway construction, A+B (cost-plus-time) contracts have contractors bid the completion time added to their dollar price, so the winning bid becomes a binding production deadline (Lewis and Bajari, 2011). There, however, the losing bidders exit once the contract is awarded, so the bid resets only the winner's own production problem, not the terms of a contest against a rival. We ask whether this decentralised mechanism actually implements the allowance an informed designer would choose to equalise win probabilities.

Armageddon chess provides a direct implementation of this self-pricing mechanism. Black is given the rule advantage of winning the match on a draw, while White retains the first-move advantage. In traditional fixed-odds Armageddon, the designer compensates for Black's draw odds by imposing a fixed time allowance. The bidding variant instead delegates that choice to the contestants: each player bids the amount of clock time he is willing to accept in order to play Black, and the lower bidder receives Black at his bid. The crucial feature, as in the broader class of auctions with allocative externalities described above, is that the bid is not a monetary transfer: it directly sets the resource allowance in the subsequent contest, so the payment remains inside the production technology of the game, and a bidder's payoff from losing depends on the opponent's winning bid.

The mechanism raises three questions, and they organise the paper. First, does bidding implement the informed designer's allowance, in the sense of bringing the two roles' winning probabilities together? If not, what allocation property does it implement instead? Second, how does decentralised bidding compare with alternatives available to the designer, including random assignment and a fixed allowance? Third, how does the mechanism distribute the value of the advantage between the contestants: who obtains the advantaged role, what allowance does he accept, and what rent remains?

To address these questions we model both stages: the contest and the bidding stage that precedes it. In the continuation contest, two players of given strengths choose costly inputs. One role carries a baseline advantage, while the other receives a rule-based advantage and can be handicapped through its production technology. In the bidding stage, contestants bid the allowance they are willing to accept in order to obtain the rule-advantaged role. In Armageddon chess, these two advantages are White's first move and Black's draw odds, while the allowance is Black's reduced clock.

Our theoretical analysis yields four findings.

(1) We characterise a family of equilibria of the bidding game, all assigning the

(Franke et al., 2018; Fu and Wu, 2020).

advantaged role to the same contestant but differing in the allowance. Eliminating weakly dominated strategies selects a unique outcome from this family: the winner receives the advantaged role at the losing bidder's indifference point, an allocation with a no-swap property—neither contestant would prefer the other's role at the realised allowance—and, when their cutoffs differ, a strictly positive surplus to the winner relative to the opposite role assignment.

(2) This no-swap allocation generally does not implement the informed designer's preferred allowance: the designer would choose the allowance to equalise winning probabilities, whereas bidding stops where the losing bidder is indifferent between the two roles, and for symmetric contestants the two criteria coincide only when the equilibrium cost burdens of the roles coincide. Under strictly log-concave effort costs, bidding overshoots the equalising allowance for near-equal contestants: the handicap raises the marginal cost of generating input, so the advantaged-role player supplies less input and accepts a lower winning probability for the resulting cost saving, making the auction winner a win-probability underdog in the continuation contest.

(3) This distinction carries over to the comparison of mechanisms. Random assignment generally leaves a role imbalance, a fixed allowance equalises winning probabilities only for special combinations of contestant strengths and allowance levels, and an informed designer eliminates the imbalance whenever the equalising allowance is feasible. Where bidding ranks relative to the coarser mechanisms is therefore a quantitative question.

(4) Bidding determines who bears the cost of obtaining the advantage: the winner accepts the allowance that makes the losing bidder indifferent between the two roles, which—as the overshoot result shows—leaves him the underdog in the continuation contest.

We take the mechanism-design model to 477 elite Armageddon games from Chess.com tournaments played between August 2022 and October 2024, observing both players' bids—unlike most auction data sets, where only the winning bid is recorded—the resulting colour assignment, International Chess Federation (FIDE) ratings, and the game outcome. Player identities are pseudonymous but stable across games, which supports resampling at the player-pair level in the structural inference. The reduced-form evidence reveals a striking contrast: the FIDE rating gap predicts game outcomes, but the bids themselves are essentially flat in the rating gap. We then estimate the model jointly from bidding and outcome behaviour, allowing game-relevant relative ability to deviate from that implied by FIDE ratings through a latent component observed by the players but not by the econometrician. This estimated structure allows us to quantify how bidding performs relative to alternative mechanisms.

Conditional on FIDE ratings alone, which at the estimated coefficient imply a near-zero ability gap for every pair, bidding leaves the two colours' winning probabilities 3.49 percentage points apart on average, compared with 6.04 under random colour and zero under an informed designer. When we integrate over the estimated heterogeneity in latent ability unexplained by those ratings, the comparison changes: bidding produces a slightly larger gap than random colour. This second exercise should be interpreted cautiously because the latent wedge also absorbs rating error and model misspecification. The sample is an ability-compressed elite panel—the likelihood identifies scale-free ability ratios precisely even though levels are weakly identified, and reported rankings are stable across the near-optimal parameter set—and the null that ratings do not predict bids is itself precisely estimated, with a profile interval that excludes the outcome-implied slope by a wide margin. The combination of the reduced-form and structural evidence turns the theoretical implementation problem into an information problem: ratings predict who wins, but they do not appear to be what contestants price when bidding for the advantage, so decentralisation can improve on an uninformed rule without revealing the relative strength an informed designer would use to balance the contest.

The paper connects several literatures. The contest-design literature studies tournament incentives, prize structures, and optimal favouritism when the designer observes relevant asymmetries (Lazear and Rosen, 1981; Nalebuff and Stiglitz, 1983; Moldovanu and Sela, 2001; Konrad, 2009; Franke et al., 2018; Fu and Wu, 2020), with a related strand on asymmetric contests with head starts and other *ex-ante* advantages (Siegel, 2014; Clark and Nilssen, 2018). Our setting starts from the same design problem but removes the designer's information about the contestants' relative strengths: rather than choosing the compensating allowance directly, the designer lets the contestants price the advantaged role themselves, so optimal favouritism provides the informed benchmark against which the bidding mechanism is evaluated, rather than the outcome the mechanism is assumed to implement. On the auction side, a broader literature studies bids made in variables that remain inside the post-auction relationship: least-present-value-of-revenue auctions make the duration of a highway concession endogenous (Engel et al., 2001), and security-bid auctions show more generally that auction design changes when payment is tied to future output rather than made as a fixed transfer (DeMarzo et al., 2005). This contrasts with the canonical money-bid contest auction of Fullerton and McAfee (1999), where the payment leaves the contest as revenue to the principal. Che and Hendershott (2008) and Granot and Gerchak (2014), introduced above, come closest to our setting; the Armageddon application of Granot and Gerchak (2014) takes the post-auction outcome distribution as primitive, whereas we model the continuation contest and the productive effect of the bid explicitly. More generally, our bidding stage belongs to the class of auctions with allocative externalities (Jehiel and Moldovanu, 2006). Because the contestants themselves observe their relative strength, the externality here is one of complete information between the two bidders; the informational limit that motivates the mechanism is the designer's, not the bidders'. The distinctive feature of our setting is the conjunction of these ideas: contestants price the advantage themselves, and the price is a forfeited share of the resource allowance that changes how the ensuing contest is played. Pre-contest investment and endogenous-strength models also allow actions to affect the subsequent contest, but there the actions are chosen independently by the contestants rather than serving as the price paid for an advantage (Münster, 2007; Schaller and Skaperdas, 2020; Beviá and Corchón, 2013; Amegashie, 2012; Kimbrough et al., 2020; Clark et al., 2026).

A separate theoretical literature studies contests in which an unresolved outcome, i.e., a draw, is itself possible. Vesperoni and Yildizparlak (2019) axiomatise contest success functions that admit draws, while Gelder et al. (2022) characterise all-pay auctions in which a tie occurs unless one bid exceeds the other by a sufficient margin. Draws also affect incentives: they can raise the stronger contestant's effort and reduce the weaker contestant's (Deng et al., 2018), while a wider tie interval discourages effort among near-equal contestants (Zhuang et al., 2024). A designer may instead bias the treatment of the draw (Li et al., 2023) or require the contest to be replayed (Franke and Metzger, 2023). That literature asks whether and how a designer should accommodate unresolved outcomes within a contest. We take the underlying tie rule as given and study the subsequent mechanism-design problem created when one role is awarded the unresolved states: how should that advantage be allocated and handicapped when the designer does not know the compensation required to balance the contest?

Finally, the paper relates to the literature on tie-breaking and asymmetric advantage in sports. Existing work studies, among other things, bracket design, home advantage, and alternative tie-breaking procedures (Deck and Kimbrough, 2015; Szech, 2015; Huang, 2016; Chen et al., 2022; Anbarcı et al., 2021; Goel and Goyal, 2023; Pollard et al., 2017). Field evidence shows that such asymmetries can be quantitatively important: Apesteguia and Palacios-Huerta (2010), for example, document a first-mover advantage in penalty shootouts in association football, prompting experimentation with alternative kicking orders. In chess, Vargas-Calderón (2020) uses engine simulation to show that the fair fixed Armageddon time handicap depends on player strength. We instead study a mechanism in which the allowance is determined endogenously by the contestants and take it to observed bids and game outcomes. More broadly, the empirical analysis follows the use of sports institutions as laboratories for economic theory (Palacios-Huerta, 2025). Structural exercises in contest design remain relatively rare; Krishna et al. (2026), for example, structurally evaluate a design problem in which the policy instrument is the pooling of test-rank intervals rather than an auction.

The paper introduces the baseline contest in Section 2, develops and solves the handicapped bidding contest in Section 3, and turns to mechanism design in Section 4, comparing the bidding equilibrium to other mechanisms. Section 5 describes the institution and our new bid-level data set of elite Armageddon chess games, Section 6 translates the theory into testable predictions, and Section 7 takes them to the data, before we close with a discussion and concluding remarks (Section 8). Proofs of the theoretical results are in the Appendix, and the Online Appendix contains supplementary theory and derivations, additional empirical and robustness results, and a parametric example.

2 A contest with an unresolved outcome

We call an outcome *unresolved* when the underlying contest does not produce a decisive winner. In chess, this is a draw; more generally, it is a state in which the institution must apply an additional rule to allocate the prize or determine who advances. We model such a contest using a tournament-style framework in the sense of Lazear and Rosen (1981) and Nalebuff and Stiglitz (1983).[3]

2.1 The baseline contest

Two risk neutral players, $i \in \{1, 2\}$, jointly denoted a *pair*, compete for a prize v. Players have a role in the contest which we label $R \in \{W, B\}$ to keep the later chess application transparent; until that application these are simply role labels. Players have ability A_R and supply a contest input e_R that partially determines the outcome; producing it draws on a resource — in chess, clock time — that we normalise to

[3] Kalwij and De Jaegher (2023) have previously used this type of model to study conventional chess.

a full unit. With unconstrained access to this resource, the cost of producing e_R is given by a twice continuously differentiable function $c(e) \in C^2([0,\infty))$ satisfying $c(0) = c'(0) = 0$ and $c'(e) > 0$ for $e > 0$. We impose a further regularity condition on the cost function: c is *strongly* convex, i.e. there exists $m > 0$ such that $c''(e) > m$ for all $e \geq 0$. Note that strong convexity implies $\lim_{e\to\infty} c'(e) = \infty$, so $(c')^{-1}$ is well defined on $[0,\infty)$.

There is a baseline advantage $\Omega \geq 0$ associated with role W. The score of each role is given by

$$S_W = A_W + \Omega + e_W + \varepsilon_W, \qquad S_B = A_B + e_B + \varepsilon_B,$$

where ε_R is random noise that is i.i.d. Define the difference between the scores of roles W and B as

$$X := A_W - A_B + \Omega + e_W - e_B + \varepsilon_W - \varepsilon_B := \mu + \varepsilon,$$

where the deterministic advantage in favour of role W is given by

$$\mu = (A_W - A_B) + \Omega + (e_W - e_B) := \mu_0 + (e_W - e_B), \tag{1}$$

so that μ_0 captures non-input factors that affect the score difference. The noise component $\varepsilon = \varepsilon_W - \varepsilon_B$ is distributed according to $G(\varepsilon)$. We assume G is continuous and strictly increasing, and that the density $g = G'$ exists and is continuously differentiable. Since ε_W and ε_B are i.i.d., the difference $\varepsilon = \varepsilon_W - \varepsilon_B$ is symmetrically distributed around zero.[4] We additionally assume that g is strictly single-peaked at zero: g is strictly decreasing in $|z|$ wherever it is positive.

Along the lines of Gelder et al. (2022), we posit that, in order to win, the higher score must surpass a threshold $\Delta > 0$: if $X > \Delta$, then role W wins, if $X < -\Delta$, then role B wins; intermediate states are unresolved. In conventional chess these states correspond to draws. The probabilities of roles W and B winning are $P_W = \Pr(X > \Delta) = 1 - G(\Delta - \mu)$ and $P_B = \Pr(X < -\Delta) = 1 - G(\Delta + \mu)$, while the probability of an unresolved contest is $P_D = \Pr(|X| \leq \Delta) = G(\Delta - \mu) + G(\Delta + \mu) - 1$. In the baseline contest, a decisive win yields the prize v; if the contest is unresolved, then the prize is awarded to either contestant with equal probability. Hence, the

[4] Hence $G(-z) = 1 - G(z)$, $g(-z) = g(z)$, and $g'(-z) = -g'(z)$.

expected payoff of a player in role $R \in \{W, B\}$ is[5]

$$U_R = \left(P_R + \frac{1}{2}P_D\right)v - c(e_R). \tag{2}$$

The marginal gain from supplying the input moves with the noise density, which can be steep, so U_R need not be concave in e_R. We therefore bound how steep the density can be, relative to the curvature of the cost.

Condition 1. $\exists\, \lambda \in (0,1)$ *such that* $v\dot{g} \leq \lambda m$*, where* $\dot{g} := \sup_{\varepsilon\in\mathbb{R}}\left|g'(\varepsilon)\right|$.

For smooth, unimodal densities that decay to zero in the tails, $\dot{g}$ is finite, so Condition 1 amounts to a joint restriction on the prize v, the maximum slope of the noise density, and the uniform curvature m of the cost.[6] It is also sufficient for the second-order conditions for a maximum to be fulfilled. Unless stated otherwise, the analysis throughout the paper assumes that Condition 1 holds.

Lemma 1 (Symmetric input in the baseline contest)**.** *The baseline contest has a unique pure strategy equilibrium in which the input profile is symmetric:* $e_W = e_B = e^\star$*, where*

$$e^\star = (c')^{-1}(M(\mu_0)v), \quad M(\mu_0) := \frac{g(\Delta - \mu_0) + g(\Delta + \mu_0)}{2}. \tag{3}$$

Consequently, the equilibrium probabilities depend only on the observable primitives μ_0 *and* Δ*:* $P_W = 1 - G(\Delta - \mu_0)$*,* $P_B = 1 - G(\Delta + \mu_0)$*, and* $P_D = G(\Delta - \mu_0) + G(\Delta + \mu_0) - 1$*.*

Equilibrium probabilities are invariant to costs: any change in c affects the level of the input supplied but not the probability of each outcome. An increase in the non-input advantage μ_0 (W's ability edge plus the baseline advantage) raises W's probability of winning and lowers B's, while an unresolved outcome is most likely at an even matchup ($\mu_0 = 0$) and becomes more likely as the threshold Δ rises, which also lowers input supply in equilibrium (collected formally in Online Appendix O.1). Section 3 makes role B's access to the productive resource the object of the bidding stage.

[5] In conventional chess, the unresolved outcome is a draw and the same expected payoff arises from the usual scoring rule of one point for a win and one-half point for a draw.

[6] The condition is easier to satisfy the larger m and the smaller $v\dot{g}$. For the Gaussian noise of the structural model in Section 7, $\dot{g} \approx 0.242$ at $N(0,1)$.

3 Bidding for the advantaged role

We now assume that unresolved states from the baseline contest are awarded to role B (known in Armageddon chess as draw-odds). This makes B attractive, so the mechanism restricts role B's access to the resource: role B receives an allowance $\tau \in (0, 1]$ of the resource, capped at the full amount role W holds. Players simultaneously bid the allowance they are willing to accept to obtain role B, and the lower bid wins. A smaller allowance raises the cost of producing the contest input. We use the multiplicatively separable specification $C(\tau, e) = H(\tau)c(e)$, which in Armageddon chess captures the compression of play into a shorter clock; $H(\tau)$ is the resulting handicap. $H(\cdot)$ is continuously differentiable with $H(\tau) \geq 1, H'(\tau) < 0$. We normalise $H(1) = 1$ as the state with no handicap, so that a smaller allowance τ means a larger handicap $H(\tau)$. For allowance τ, let $\Gamma(\tau)$ denote the resulting continuation contest with roles W and B already assigned.

3.1 The continuation contest

In $\Gamma(\tau)$, role B receives the unresolved states of the baseline contest, so the win probabilities are: $\rho_B = G(\Delta - \mu)$, and $\rho_W = 1 - G(\Delta - \mu)$. Expected payoffs are then

$$U_W = v\rho_W - c(e_W), \qquad U_B = v\rho_B - H(\tau)c(e_B). \tag{4}$$

The first-order conditions for optimal supply of the input yield the working equation

$$g(\Delta - \mu)v = c'(e_W) = H(\tau)c'(e_B). \tag{5}$$

Define $\hat{e} := e_W - e_B$; the system (5) can be reduced to a scalar fixed-point equation that identifies the pure strategy equilibrium candidate:

$$\hat{e} = \eta(\hat{e}), \quad \textit{where} \tag{6}$$

$$\eta(e) := (c')^{-1}(vg(\Delta - \mu_0 - e)) - (c')^{-1}\left(\frac{vg(\Delta - \mu_0 - e)}{H(\tau)}\right). \tag{7}$$

Condition 1 ensures that η in (6) is a strict contraction with modulus $v\dot{g}/m$. The fixed point is therefore unique, and the same condition delivers the second-order inequalities for both players. Economically, the contraction modulus is the ratio of the maximum responsiveness of marginal benefit to cost curvature. If marginal benefits are sufficiently responsive relative to cost curvature, the contraction argument, and hence the uniqueness result, can fail.

Proposition 1 (Equilibrium of the handicapped contest)**.** *Consider the game* $\Gamma(\tau)$*.*

(i) Existence. *The game has a unique pure strategy equilibrium in which the input profile* (e_W^*, e_B^*) *satisfies* (5)*.*

(ii) Contest-input gap. *The equilibrium input gap* $\hat{e}^* = e_W^* - e_B^*$ *is positive for* $\tau < 1$ *and strictly decreasing in* τ *on* $(0, 1]$*, and satisfies* $\hat{e}^*(1) = 0$*.*

(iii) Cost burden. *If* c *is log-concave (log-convex) on* $(0, \infty)$*, then* $H(\tau)\, c(e_B^*) \leq (\geq)\ c(e_W^*)$*, with strict inequality if* c *is strictly log-concave (log-convex) and* $\tau < 1$*.*

The input gap is positive for $\tau < 1$ irrespective of the players' relative abilities: the player in role W supplies a larger input. As role B's resource allowance is increased, the input gap narrows and disappears when $\tau = 1$.

Who bears the higher cost of supplying the input depends on the shape of the cost function. Log-concavity makes the proportional marginal cost $c'(e)/c(e)$ decreasing in e, and then role W supplies more input and also bears the larger cost. If instead $c'(e)/c(e)$ is non-decreasing, role B, though supplying a lower level of the input, bears the higher effective cost $H(\tau)c(e_B)$: the handicap cost is a genuine welfare penalty rather than a normalisation.

Let $(e_W(\theta), e_B(\theta))$ be the unique equilibrium input profile satisfying (5) for any $\theta \in \{v, \tau, \Delta, \mu_0\}$, and let $\hat{e}(\theta) = e_W(\theta) - e_B(\theta)$ denote the equilibrium input gap. The equilibrium win probabilities in the handicapped contest are

$$\rho_B(\theta) = G(\Delta - \mu_0 - \hat{e}(\theta)), \qquad \rho_W(\theta) = 1 - G(\Delta - \mu_0 - \hat{e}(\theta)).$$

The following corollary records the comparative statics of the win probability; the input-gap responses that drive them are derived in the proof of the corollary in Appendix A.

Corollary 1 (Win-probability comparative statics)**.** *In the equilibrium of Proposition 1, role* W*'s win probability satisfies*

$$\frac{d\rho_W}{d\mu_0} > 0, \qquad \frac{d\rho_W}{d\Delta} < 0, \qquad \frac{d\rho_W}{d\tau} < 0,$$

while the effect of the prize v *is in general ambiguous. Since* $\rho_W + \rho_B = 1$*, the effects on* ρ_B *are of the opposite sign.*

A primitive moves role W's win probability through two channels: a direct shift of the effective threshold $\Delta - \mu_0 - \hat{e}$, and an indirect effect operating through the equilibrium input gap $\hat{e}$. For μ_0 and Δ, the direct channel dominates, so ρ_W rises with the baseline role advantage and ability gap μ_0 and falls with the threshold Δ. The prize v and the allowance τ act only through the input gap. That channel is signed for τ: by Proposition 1(ii), the input gap falls in τ, so an increase in role B's allowance lowers W's chances; it is ambiguous for v, whose effect on the input gap can take either sign.

3.2 The bidding stage

We now solve for the players' optimal first-period bids, which settle the role assignment for the continuation contest. Since roles are not yet assigned in period 1, player i's expected payoff combines values from two *different* second-period games according to who has role W. For $i, j \in \{1, 2\}$, let $\Gamma^i(\tau)$ and $\Gamma^j(\tau)$ denote these games, in which player i and player j, respectively, has role W while the opponent has role B with allowance τ. In $\Gamma^i(\tau)$, the equilibrium inputs are $e^i_W(\tau)$, supplied by player i in role W, and $e^j_B(\tau)$, supplied by player j in role B; then role W's deterministic advantage is

$$\mu^i(\tau) := A_i - A_j + \Omega + e^i_W(\tau) - e^j_B(\tau), \tag{8}$$

equivalent to W's advantage (1) applicable to $\Gamma^i(\tau)$. The game $\Gamma^j(\tau)$ is obtained by interchanging roles of i and j, with role W's advantage $\mu^j(\tau) = A_j - A_i + \Omega + e^j_W(\tau) - e^i_B(\tau)$.

Given a bid profile (τ_i, τ_j), the lower bid wins role B at that allowance, so player i's stage-1 expected payoff is

$$EU_i = \begin{cases} U_W\Big(\mu^i(\tau_j), \tau_j\Big), & \tau_i > \tau_j \ \ (i \text{ loses, has role } W \text{ in } \Gamma^i(\tau_j)), \\ U_B\Big(\mu^j(\tau_i), \tau_i\Big), & \tau_i < \tau_j \ \ (i \text{ wins, has role } B \text{ in } \Gamma^j(\tau_i)), \\ \frac{1}{2}\Big(U_W(\mu^i(\tau), \tau) + U_B(\mu^j(\tau), \tau)\Big), & \tau_i = \tau_j = \tau \ \ (\text{random assignment}), \end{cases}$$

with U_W, U_B given by (4).

For each player, we identify the allowance τ that leaves him indifferent between the two roles, against the opponent in the opposite role. At a common τ, define the payoff-difference function

$$\phi_i(\tau) := U_B\Big(\mu^j(\tau), \tau\Big) \ - \ U_W\Big(\mu^i(\tau), \tau\Big). \tag{9}$$

So $\phi_i(\tau) > 0$ means player i prefers to have role B at τ, and $\phi_i(\tau) = 0$ defines his point of indifference. Since the two terms in $\phi_i(\tau)$ live in two different games, we examine how $\phi_i(\tau)$ behaves with respect to τ by computing $dU_B/d\tau$ and $dU_W/d\tau$ along the equilibrium path, each inside its own game.

Lemma 2 (Marginal effect of τ on payoffs)**.** *Along the equilibrium path of the game $\Gamma^j(\tau)$,*

$$\frac{d}{d\tau}\,U_B\Big(\mu^j(\tau),\tau\Big) = -H(\tau)\,c'\Big(e_B^i(\tau)\Big)\,\frac{de_W^j(\tau)}{d\tau} - H'(\tau)\,c\Big(e_B^i(\tau)\Big), \tag{10}$$

and along the equilibrium path of the game $\Gamma^i(\tau)$,

$$\frac{d}{d\tau}\,U_W\Big(\mu^i(\tau),\tau\Big) = -H(\tau)\,c'\Big(e_B^j(\tau)\Big)\,\frac{de_B^j(\tau)}{d\tau}. \tag{11}$$

Notably, the marginal effect of τ on each player's payoff is mediated by the opponent's input response to τ. We now isolate the two channels through which these responses operate and impose a restriction that binds only when the channels conflict, yielding monotonicity of ϕ_i. The derivatives in Lemma 2 yield

$$\phi_i'(\tau) = \underbrace{H(\tau)c'(e_B^j(\tau))\,\frac{de_B^j(\tau)}{d\tau}}_{-\,dU_W/d\tau\ \text{ in }\Gamma^i(\tau)}\ \underbrace{-\,H(\tau)c'(e_B^i(\tau))\,\frac{de_W^j(\tau)}{d\tau} - H'(\tau)c(e_B^i(\tau))}_{dU_B/d\tau\ \text{ in }\Gamma^j(\tau)}. \tag{12}$$

The input responses within each game are, in general, ambiguous in sign. A larger allowance for B moves the equilibrium supply of the input through two channels. The first is a marginal cost relief: a higher τ lowers B's marginal cost of input supply $H(\tau)c'(e)$ and therefore raises his supply of the input. The second is strategic and operates through the input gap, which from Proposition 1 is decreasing in τ within each game. Writing the input gaps as $\hat e^i(\tau) = e_W^i(\tau) - e_B^j(\tau)$ and $\hat e^j(\tau) = e_W^j(\tau) - e_B^i(\tau)$ in games $\Gamma^i(\tau)$ and $\Gamma^j(\tau)$, respectively, and differentiating the first-order conditions (5) within each game give (suppressing argument for brevity)

$$\frac{de_B^j}{d\tau} = \frac{1}{H(\tau)c''(e_B^j)}\Big[\underbrace{v\,g'(\Delta-\mu^i)\Big(-\tfrac{d\hat e^i}{d\tau}\Big)}_{\text{strategic}}\ \underbrace{-\,H'(\tau)c'(e_B^j)}_{\text{marginal cost relief}>0}\Big], \quad \text{(Game } \Gamma^i(\tau)) \tag{13}$$

$$\frac{de_W^j}{d\tau} = \frac{1}{c''(e_W^j)}\ \underbrace{v\,g'(\Delta-\mu^j)\Big(-\tfrac{d\hat e^j}{d\tau}\Big)}_{\text{strategic}}. \quad \text{(Game } \Gamma^j(\tau)) \tag{14}$$

The marginal cost relief term in (13) is always positive ($H' < 0$) and raises role B's supply of the input. The strategic term, present in both, carries the sign of g'

at that game's decisive margin: as the gap narrows, μ^i (resp., μ^j) falls and $\Delta - \mu^i$ (resp., $\Delta - \mu^j$) rises, shifting the marginal return to supplying the input, which is the critical source of indeterminacy. In equation (13), the two channels therefore pull against each other precisely when role B is the favourite at that game's margin ($\Delta - \mu^i > 0$, so that $g' < 0$), and reinforce each other when he is the underdog.

Although changing τ has no direct effect on W's input response, it nonetheless mediates role B's equilibrium payoff, where it is traded off against the total cost relief afforded to B. In the following analysis, we assume that this cost relief, marginal in the input response and total in the payoff, is sufficiently strong that it dominates the strategic effect even when the two operate in opposite directions.

Assumption 1. *For any pair $i, j \in \{1, 2\}$ with $i \neq j$, in the game $\Gamma^i(\tau)$ in which player i has role W and player j has role B, an increase in the allowance τ:*

(i) raises the equilibrium input supplied by player j: $\dfrac{de_B^j(\tau)}{d\tau} > 0$ for all $\tau \in (0, 1]$; and

(ii) provides a total cost relief to player j that at least offsets the strategic effect of player i's input response: $-H'(\tau)\, c\Big(e_B^j(\tau)\Big) \;\geq\; H(\tau)\, c'\Big(e_B^j(\tau)\Big)\, \dfrac{de_W^i(\tau)}{d\tau}$ for all $\tau \in (0, 1]$.

The two parts of Assumption 1 restrict two different equilibrium objects: a marginal cost set against role B's input response in part (i), and a total cost set against role B's payoff response in part (ii). The two clauses bind on complementary regions of the margin $\Delta - \mu^i$. Part (i) requires the marginal cost relief to overcome the strategic term where role B is the favourite ($\Delta - \mu^i > 0$); part (ii) requires the total cost relief to overcome the strategic loss through W's input response where role B is the underdog ($\Delta - \mu^i < 0$); elsewhere each part holds automatically. Stated for every ordered pair, the assumption restricts $\Gamma^j(\tau)$ as well, at the swapped pair (j, i): part (i) then gives $dU_W/d\tau < 0$ in $\Gamma^i(\tau)$, and part (ii) at (j, i) makes $dU_B/d\tau$ non-negative in $\Gamma^j(\tau)$. Both brackets in equation (12) therefore push the same way, so ϕ_i is increasing in τ.

Proposition 2 (Preference monotonicity)**.** *Suppose Assumption 1 holds. Then for each $i \in \{1, 2\}$, the function $\phi_i(\tau)$ is continuously differentiable and strictly increasing on* $(0, 1]$.

By Proposition 2, ϕ_i crosses zero at most once. An interior cutoff $\hat{\tau}_i \in (0, 1)$ therefore exists for each player if and only if ϕ_i is negative near $\tau = 0$ and positive near $\tau = 1$:

$$\lim_{\tau \to 0^+} \phi_i(\tau) < 0 \qquad \text{and} \qquad \phi_i(1) > 0. \tag{15}$$

The behaviour of ϕ_i at the lower end of the bid space is captured by the following assumption.

Assumption 2 (Lower boundary)**.** *There exists* $\underline{\tau} \in (0,1)$ *such that* $\phi_i(\tau) < 0$ *for all* $\tau \leq \underline{\tau}$ *and each* $i \in \{1,2\}$*.*

Economically, the assumption requires the resource restriction to become sufficiently severe at low allowances that the advantage of role B no longer compensates for the resulting increase in the cost of supplying contest input. $\underline{\tau}$ is the most restrictive admissible allowance at which both contestants prefer role W, i.e., where $\phi_i(\underline{\tau}) < 0$. We take the *feasible* bid space to be the closed interval $\mathcal{T} := [\underline{\tau}, 1]$. In Armageddon chess, this boundary has a literal interpretation: sufficiently little clock time causes forfeiture.

At $\tau = 1$, the handicap disappears ($H(1) = 1$) and the resulting equilibrium input gap closes ($\hat{e}^*(1) = 0$), so the remaining asymmetry between the two roles, over and above their ability gap, is role B's rule-based advantage Δ set against role W's baseline advantage Ω. Assumption 3 requires role B to be preferred when unconstrained. Under Condition 1, this is equivalent to $\Delta > \Omega$, so the rule-based advantage of receiving unresolved states dominates role W's baseline advantage.

Assumption 3 (Interior cutoff)**.** $\phi_i(1) > 0$ *for each* $i \in \{1,2\}$*.*

Assumption 3 is a restriction on primitives rather than on behaviour: the equivalence with $\Delta > \Omega$ holds at every ability gap and for both players, so the assumption cannot fail for one player alone (see Lemma O.1, online appendix).

The two boundary conditions in (15), together with Proposition 2, imply that each ϕ_i has a unique zero in the feasible bid space $\mathcal{T} = [\underline{\tau}, 1]$, which we refer to as player i's bid cutoff:

$$\hat{\tau}_i \;:=\; \phi_i^{-1}(0) \;\in\; \mathcal{T}. \tag{16}$$

Without loss of generality, we assume that $\hat{\tau}_1 \leq \hat{\tau}_2$. The equilibria of the bidding game can now be characterised in terms of the pair $(\hat{\tau}_1, \hat{\tau}_2)$.

Proposition 3 (Equilibrium characterisation of the bidding game)**.** *Suppose that Assumptions 1–3 hold and* $\hat{\tau}_1 \leq \hat{\tau}_2$*, and let* $\mathcal{T} = [\underline{\tau}, 1]$ *be each player's period-1 bid space.*

(a) For each player i*, every bid* $\tau < \hat{\tau}_i$ *is weakly dominated by* $\hat{\tau}_i$*.*

(b) If $\hat{\tau}_1 = \hat{\tau}_2 = \hat{\tau}$*, then the pure action profile* $(\hat{\tau}_1, \hat{\tau}_2) = (\hat{\tau}, \hat{\tau})$ *is part of a Nash equilibrium, with winning bid* $\tau^* = \hat{\tau}$*.*

(c) If $\hat{\tau}_1 < \hat{\tau}_2$, then for any $\tau \in (\hat{\tau}_1, \hat{\tau}_2]$, there exists $\delta > 0$ such that the profile in which player 1 bids τ with certainty and player 2 mixes uniformly on the interval $[\tau, \tau + \delta]$ is part of a Nash equilibrium. In every such equilibrium, player 1 wins the bidding and is assigned role B, the winning bid is $\tau^* = \tau \in (\hat{\tau}_1, \hat{\tau}_2]$, and expected payoffs are $U_1 = U_B\big(\mu^2(\tau), \tau\big)$, $U_2 = U_W\big(\mu^2(\tau), \tau\big)$.

Part (c) of Proposition 3 exhibits a continuum of equilibria, all assigning role B to the same player (player 1, the player with the smaller indifference point) but differing in the level of the winning bid. The multiplicity raises an immediate selection question: which of these equilibria should we expect to be played? Proposition 4 shows that only one member of this family—the one with the largest sustainable winning bid—survives elimination of weakly dominated strategies. This is because, in every equilibrium with $\tau^* < \hat{\tau}_2$, player 2 assigns positive probability to bids below $\hat{\tau}_2$, and each such bid is weakly dominated by the bid $\hat{\tau}_2$.

Proposition 4 (Weak-dominance refinement)**.** *Suppose that Assumptions 1–3 hold and that $\hat{\tau}_1 \le \hat{\tau}_2$. Among the family of equilibria identified in Proposition 3, only the equilibrium with $\tau^* = \hat{\tau}_2$ survives deletion of weakly dominated strategies.*

Whether the game admits Nash equilibria outside this family is a question we cannot settle in full generality. Online Appendix O.1 shows that, when the cutoffs differ, within the class of regular strategies that place no mass on weakly dominated bids, every equilibrium *outcome* coincides with the selected one: player 1 is assigned role B at $\hat{\tau}_2$ (Theorem O.1); at the knife-edge $\hat{\tau}_1 = \hat{\tau}_2$, the winning bid is still the common cutoff, and only the role assignment is left to the tie-break. The selection is thus not an artefact of the constructed family.

The undominated equilibrium has an auction-theoretic interpretation. Player 2, who has the larger indifference point, is bid up to his reservation, so the bidding game leaves him no rent over his reservation value of role W: his payoff remains strictly positive, but no larger than what the opposite role at the same allowance would give him. Player 1 (with the smaller indifference point) wins at his opponent's reservation and is assigned role B at $\hat{\tau}_2$, capturing the strictly positive surplus $\phi_1(\hat{\tau}_2) > 0$ over the swapped assignment. The auction "price" is the accepted allowance $\hat{\tau}_2$, i.e. the rival's reservation allowance.

We close the bidding analysis by returning to Assumption 1. Beyond making ϕ_i single-crossing, the assumption makes the role payoffs U_B and U_W globally monotone in τ, and this payoff monotonicity is what yields a unique bid cutoff for each player and the characterisation of Propositions 3–4. Were it to fail, U_B or U_W could be non-monotone even when ϕ_i has a single zero, and a player might strictly prefer a deviation to a distant allowance at which his expected second-period payoff happens to be higher. A mixed-strategy equilibrium of the bidding game continues to exist by standard arguments,[7] but the player-specific cutoff parameterization, the dominance refinement, and with them the model's testable structure would break down.

4 Mechanism design

The bidding mechanism is one of several ways to resolve role assignment and set the allowance in our general model. To see what bidding adds, we compare it with an informed designer and two coarser benchmarks:[8]

- **Random assignment (RC).** Role assignment is randomised and no handicap is associated with role B.
- **Fixed allowance (FA).** Role assignment is randomised and role B gets a fixed allowance $\tau_{\mathsf{FA}} \in \mathcal{T}$.
- **Designer-optimal (DO).** A designer who observes the players' abilities selects an allowance for role B that equalises the win probability. The rule takes two forms: one that treats the role assignment as given (DO_τ) and one that also selects it ($\mathsf{DO}_{c\tau}$), both developed together in Section 4.2.
- **Bidding mechanism (BM).** Our mechanism. Players bid the allowance to receive role B, lower bid wins.

Recall that, in the realised game, role W's deterministic advantage and role B's win probability are

$$\mu^W(\tau) = (A_W - A_B) + \Omega + \hat{e}(\tau), \qquad \rho_B = G(\Delta - \mu^W(\tau)), \tag{17}$$

with the equilibrium input gap $\hat{e}(\tau)$ strictly decreasing in τ and $\hat{e}(1) = 0$. A larger allowance for role B lowers μ^W and raises ρ_B. The *signed* win-probability margin

[7] The bidding-game payoff is discontinuous along the diagonal $\tau_i = \tau_j$, where the role assignment switches, although the underlying continuation values U_B and U_W are continuous in τ. Existence of a mixed-strategy equilibrium follows from the theorem of Dasgupta and Maskin (1986) for games with payoff discontinuities. Unfortunately, the argument yields no structural information about the equilibrium.

[8] To be transparent with the Armageddon application, we use FA to denote the fixed allowance, known in the Armageddon application as Fixed Odds, and RC for random assignment (Random Colour). The subscript c in $\mathsf{DO}_{c\tau}$ has the same origin: in the application the designer's role instrument is the colour assignment.

in a realised game is

$$\rho_B - \rho_W = 2\rho_B - 1 = 2G(\Delta - \mu) - 1, \tag{18}$$

which is fair at $\rho_B = \frac{1}{2}$ (equivalently, $\mu = \Delta$: the deterministic advantage of role W exactly offsets the rule-based advantage assigned to role B), positive when role B is favoured ($\rho_B > \frac{1}{2}$, $\mu < \Delta$) and negative when role W is favoured ($\rho_B < \frac{1}{2}$, $\mu > \Delta$).

4.1 The designer's objective

To decide an unresolved state of the contest, a desirable tiebreak rule should not tilt the balance in any realised game in which the roles are already assigned: neither role B nor role W should hold a systematic edge. A mechanism induces a distribution over realised games, each a role assignment paired with an allowance. This distribution is degenerate under the designer-optimal and bidding mechanisms, both of which select a single realisation (for bidding, away from the knife-edge of equal cutoffs, where the winning bid is still the common cutoff and only the role assignment is randomised). Under the random-assignment and fixed-allowance mechanisms it is instead two-point, a fair coin placing equal weight on the two role assignments.

Definition 1 (Role gap). *For a mechanism* $M \in \{\mathsf{RC}, \mathsf{FA}, \mathsf{DO}, \mathsf{BM}\}$*, the role gap is*

$$\Delta\rho^M \ := \ \mathbb{E}_M \left|\rho_B - \rho_W\right|,$$

the expected absolute difference in winning probabilities of roles B and W in the realised game, where the expectation is taken over the mechanism's randomisation.

A different notion asks whether a mechanism favours a *player* rather than a role. Writing π_i for player i's ex-ante probability of winning the contest, the player gap is $\Delta\pi^M := |\pi_1 - \pi_2|$, where the mechanism's randomisation is averaged out *before* the absolute value is considered. The two notions are not interchangeable, and in our setting only the first can serve as the designer's objective. The player gap is blind to the design instrument itself: with roles assigned by a coin flip and an even matchup, $\Delta\pi = 0$ whatever allowance role B is given. So every fixed-allowance rule would be considered as ex-ante fair even though each realised game can be biased toward one role. The role gap, by contrast, precisely records this bias. For the designer and for bidding, the role assignment is deterministic, so $\Delta\rho^M$ coincides with the realised game's own gap and $\Delta\pi^M$ carries the same information. For random role assignment and fixed allowance the two notions diverge, as Sections 4.4–4.5 describe below.

4.2 The informed designer

An informed designer observes contestants' abilities and chooses the allowance to minimise the role gap. How close it gets depends on whether she can decide role assignment, giving two versions:

(a) **Role given – the primary designer DO_τ.** The designer conditions τ on the *realised* role assignment and solves $\tau_{\mathsf{DO}} \in \arg\min_{\tau \in [\underline{\tau},1]} \left|\Delta - \mu^W(\tau)\right|$. With one game to fix, one instrument equalises it whenever feasible. This is the version the structural counterfactual computes.

(b) **Role chosen – the two-instrument designer $\mathsf{DO}_{c\tau}$.** The designer also picks which player has role W, $i \in \{1,2\}$, solving $\min_{i\in\{1,2\},\tau} |\Delta - \mu^i(\tau)|$; the role instrument shifts the equalising target by $2(A_1 - A_2)$.

We characterise both designers in the following proposition.

Proposition 5 (Equalising allowance)**.** *Each designer's programme has a solution, unique for DO_τ; for $\mathsf{DO}_{c\tau}$, ties across the two role assignments at boundary solutions are broken in favour of the realised role. The solution is set by the equalising target against the feasibility window $\left(0, \hat{e}(\underline{\tau})\right)$ of the game being evaluated: the realised game for DO_τ, and $\Gamma^k(\cdot)$ with input gap $\hat{e}^k$ (as in (26)) for each candidate role assignment k under $\mathsf{DO}_{c\tau}$.*

(i) ***One instrument DO_τ.*** *The realised role gives one target; DO_τ equalises interiorly,*

$$\mu^W(\tau_{\mathsf{DO}}) = \Delta \iff \hat{e}(\tau_{\mathsf{DO}}) = \Delta - \Omega - (A_W - A_B), \tag{19}$$

giving $\rho_B = \frac{1}{2}$, when this target lies in the window, $0 < \Delta - \Omega - (A_W - A_B) < \hat{e}(\underline{\tau})$, and otherwise rests on the nearer boundary of $\mathcal{T}$.

(ii) ***Two instruments $\mathsf{DO}_{c\tau}$.*** *The role choice gives two targets $\Delta - \Omega \mp (A_1 - A_2)$, symmetric about $\Delta - \Omega$ and $2|A_1 - A_2|$ apart; $\mathsf{DO}_{c\tau}$ equalises if and only if, for some role assignment k, the corresponding target lies in that game's window $\left(0, \hat{e}^k(\underline{\tau})\right)$.*

When the target lies outside the window, the boundary is explicit: DO_τ sets $\tau_{\mathsf{DO}} = 1$ if $\Delta - \Omega - (A_W - A_B) < 0$ (role W is still ahead and B has no handicap) and $\tau_{\mathsf{DO}} = \underline{\tau}$ if it exceeds $\hat{e}(\underline{\tau})$ (role B is still ahead at the most restrictive feasible allowance); for $\mathsf{DO}_{c\tau}$, the same rule applies to each target against its own window

$(0, \hat{e}^k(\underline{\tau}))$, and the designer keeps the role with the smaller residual. Two instruments therefore weakly dominate one: $\mathsf{DO}_{c\tau}$ equalises a larger set of pairs than DO_{τ}, and coincides with it at symmetric pairs. When the ability gap is wide enough that each target falls outside its own window, even $\mathsf{DO}_{c\tau}$ rests on a boundary with a positive gap. The feasible set is $[\underline{\tau}, 1]$. As $\underline{\tau}$ falls, each window widens, so the share of pairs that can be equalised is governed by the floor rather than by the model itself.

4.3 Bidding versus the designer

Bidding, like $\mathsf{DO}_{c\tau}$, controls both role and allowance, yet sets the allowance by a different criterion. Whenever equalisation is feasible, the informed designer chooses the allowance so that the deterministic advantage of role W exactly offsets the rule-based advantage of role B. We call this the *equalising allowance*, at which $\rho_B = \frac{1}{2}$. Bidding instead sets the allowance at the losing bidder's payoff-indifference point: the undominated bid leaves the loser indifferent between the two role assignments. To isolate the source of the difference, consider first a symmetric pair. In this case the one- and two-instrument designer problems coincide, the players' payoff-difference functions are identical, and so are their bid cutoffs. We can write the common function as:

$$\phi(\tau) = \Big[2G(\Delta - \Omega - \hat{e}(\tau)) - 1\Big]v - \Big[H(\tau)\, c\Big(e_B(\tau)\Big) - c\Big(e_W(\tau)\Big)\Big], \tag{20}$$

where the first term is the signed win-probability margin in prize units and the second is the cost-burden difference between the two roles. Equation (20) exposes the implementation problem. The informed designer sets the first term to zero: she chooses the allowance that eliminates the win-probability advantage of either role. Bidding instead sets the entire payoff difference to zero, so at the equilibrium bid the win-probability margin must offset the difference in equilibrium cost burdens. Thus, even when the contestants possess the information about relative strength that the designer lacks, willingness to pay reveals the allowance that equalises the marginal bidder's payoffs across roles, not the allowance that equalises winning probabilities.

By Proposition 1(iii), the sign of the cost-burden difference is governed by the log-curvature of the cost function. Under strictly log-concave costs and an interior bid, including the quadratic specification estimated in Section 7.2, the cost burden of role B is lower. The equilibrium bid therefore pushes role B past the equalising allowance and leaves B a win-probability underdog.

Proposition 6 (Bidding misses the equalising allowance)**.** *Suppose $A_1 = A_2$, so that $\tau^* = \hat{\tau}_1 = \hat{\tau}_2$. Bidding sets the equalising allowance, $\tau^* = \tau_{\mathsf{DO}}$, if and only if $H(\tau^*)\, c(e_B^*) = c(e_W^*)$. If c is strictly log-concave and $\tau^* < 1$, then role B is a strict win-probability underdog at the bid and $\tau_{\mathsf{DO}} > \tau^*$.*

The driving force is the cost side of the indifference condition. The handicap raises the marginal cost of producing the input in role B, so the contestant assigned that role supplies less input than the contestant in role W. Under strictly log-concave costs, the reduction in input more than compensates for the handicap multiplier, leaving the equilibrium cost burden $H(\tau)c(e_B)$ below $c(e_W)$. A contestant can therefore accept a win-probability disadvantage in role B and still obtain the same payoff as in role W. The losing bidder consequently bids past the equalising allowance, whereas the designer, concerned with the win-probability margin alone, stops earlier. Under strictly log-convex costs, these comparisons reverse, and bidding stops short of the equalising allowance. Although Proposition 6 is stated at $A_1 = A_2$, the cutoffs and equilibrium objects vary continuously in abilities (Lemma O.6, Online Appendix), so the strict overshoot and the underdog property extend to near-symmetric pairs; Section 6.2 invokes them in that form.

Away from symmetry, a second force enters. The winning bid is now the loser's cutoff $\hat{\tau}_2$ (players are indexed by cutoffs, $\hat{\tau}_1 \leq \hat{\tau}_2$, as in Section 3; the ability ordering $A_1 \geq A_2$ of Section 4.4 is not maintained here), and writing out the loser's indifference $\phi_2(\hat{\tau}_2) = 0$ expresses the realised game's margin at the bid, exactly the quantity the designer zeroes out, as the sum of two wedges:

$$\begin{aligned}\Big[2G(\Delta - \mu^2(\hat{\tau}_2)) - 1\Big]v \;&= \Big[H(\hat{\tau}_2)\, c\Big(e_B^2(\hat{\tau}_2)\Big) - c\Big(e_W^2(\hat{\tau}_2)\Big)\Big] \\ &\quad + \Big[G(\Delta - \mu^2(\hat{\tau}_2)) - G(\Delta - \mu^1(\hat{\tau}_2))\Big]v. \end{aligned} \tag{21}$$

The first term is the cost-burden difference of the symmetric case, now taken across the loser's two role-assignment games, and near symmetry its sign still follows the log-curvature of the cost function. The second term is an ability wedge, reflecting the payoff difference generated by the contestants' unequal strengths across the two role assignments. Its sign depends on which contestant wins the bidding and therefore on the ability regime characterised above. When the winner is the weaker player, the two wedges share a sign and role B remains a win-probability underdog at the bid; when the winner is the stronger player, they oppose, and for wide ability gaps the theory no longer signs the direction of the miss. Away from symmetry, therefore, the theory no longer generally signs the direction of the implementation error. Coincidence between the bid and the equalising allowance nevertheless remains a knife-edge condition: it requires the cost-burden and ability wedges in (21) to offset exactly.

What, then, does bidding achieve? It aligns the role assignment with the players' own preferences. At the winning bid, neither player prefers the opposite assignment at the realised allowance. The proposition below shows that this property in fact characterises the winning bids of the equilibrium family.[9]

Proposition 7 (No-swap property of bidding). *In every equilibrium of the family characterised in Proposition 3, and in particular in the weak-dominance-refined equilibrium, neither player prefers the opposite role assignment at the winning bid. Conversely, every allowance with this property under the equilibrium role assignment, except possibly the bidding winner's own cutoff, is the winning bid of some equilibrium in the family.*

The bidding mechanism and the informed designer therefore implement two different notions of balance. Bidding delivers preference-based balance: at the realised allowance, neither contestant would exchange roles. The informed designer delivers outcome-based balance: neither role has a win-probability advantage. For symmetric contestants, these notions coincide only when the allowance creates no cost-burden wedge between the roles.

The characterised equilibria thus deliver exactly the outcomes that no player would swap, and the undominated equilibrium sits at the boundary of that set. Contrast this with random role assignment, where the value of role B goes unpriced in every realisation and whoever is assigned role W would swap. The same logic explains why the designer-optimal allowance is generally not an equilibrium outcome. At and near symmetric pairs, (20) gives $\phi(\tau_{\mathsf{DO}}) > 0$ under strictly log-concave cost: in the equalised game the player assigned role W strictly prefers role B and would underbid.

The two allowances, the designer-optimal and the bidding equilibrium, also differ in the inputs they induce in the continuation contest.

Proposition 8 (Contest input under bidding and design). *Suppose $A_1 = A_2$ and $c(\cdot)$ is strictly log-concave. Then $e_W(\tau_{\mathsf{DO}}) > e_W(\tau^*)$ and $e_B(\tau_{\mathsf{DO}}) > e_B(\tau^*)$: both players supply strictly more of the contest input under the designer-optimal allowance than at the equilibrium bid.*

The equalising allowance is interior without further assumption: under Assumption 3, role B is favoured in the absence of a handicap and, by Proposition 6, is an

[9]The proofs of Propositions 7 and 8 are in Online Appendix O.1.5.

underdog at the equilibrium bid. Hence τ_{DO} lies strictly between the bid and the no-handicap benchmark. Input responds to an increase in the allowance through two channels. First, moving the decisive margin towards the centre of the noise distribution raises the marginal return to input for both contestants. Second, a larger allowance directly lowers role B's marginal cost. Between the equilibrium bid and the equalising allowance these channels operate in the same direction, yielding the strict comparison in Proposition 8. Beyond the equalising allowance they oppose each other, so input at the bid cannot in general be ranked relative to the no-handicap benchmark.

It remains to place bidding against the two coarser mechanisms. Sections 4.4 and 4.5 show that random assignment leaves the role advantage unpriced and that a fixed allowance closes the role gap only for particular combinations of contestant strengths and allowance levels; where bidding lies relative to these benchmarks is not generally signed by the theory and is settled quantitatively in Section 7.

4.4 Random assignment

The least-controlled benchmark is random role assignment with no handicap. There is no input gap ($\hat{e}(1) = 0$), so the advantage is purely structural: $\mu^i(1) = (A_i - A_j) + \Omega$. Without loss of generality, assume that player 1 has higher ability, $A_1 \geq A_2$. Each player is assigned role W with probability one half, so player 1 wins the contest with probability $\pi_1 = \frac{1}{2}\Big[1 - G(\Delta - \mu^1(1))\Big] + \frac{1}{2}\,G(\Delta - \mu^2(1))$ and player 2 with $\pi_2 = 1 - \pi_1$. The player gap is then

$$\Delta\pi^{\mathsf{RC}} \;=\; \pi_1 - \pi_2 \;=\; G\Big(\Delta - \mu^2(1)\Big) - G\Big(\Delta - \mu^1(1)\Big) \;\geq\; 0, \tag{22}$$

with equality if and only if $A_1 = A_2$. So the stronger player wins more often.

The role gap behaves differently. If player 1 is assigned role B, then the realised gap is $\rho_B - \rho_W = 2G(\Delta - \mu^2(1)) - 1$, whereas it becomes $2G(\Delta - \mu^1(1)) - 1$ when the roles are reversed. Averaging the absolute margin over the two equally likely draws gives

$$\Delta\rho^{\mathsf{RC}} \;=\; \Big|G\Big(\Delta - \mu^1(1)\Big) - \tfrac{1}{2}\Big| + \Big|G\Big(\Delta - \mu^2(1)\Big) - \tfrac{1}{2}\Big|. \tag{23}$$

The two terms inside the absolute values in (23), $G(\Delta - \mu^k(1)) - \frac{1}{2}$, measure how far role B's win probability sits from one half under each role assignment. They combine in two natural ways. Their difference is the player gap (22). Their sum turns out to be the mean value of role B, measured per unit of the prize. It follows that

$$\bar{\phi}(1) \;:=\; \tfrac{1}{2}\Big(\phi_1(1) + \phi_2(1)\Big) \;=\; \Big[G\Big(\Delta - \mu^1(1)\Big) + G\Big(\Delta - \mu^2(1)\Big) - 1\Big]\, v, \tag{24}$$

which is strictly positive under Assumption 3, largest at an even matchup, and falls as the ability gap widens so that lopsided pairs value the rule-based advantage less; the expression in (24) is derived in the proof of Proposition 9. The following proposition expresses the role gap in terms of these two combinations.

Proposition 9 (Random role assignment)**.** *Under random role assignment with no handicap,*

$$\Delta\rho^{\mathsf{RC}} \;=\; \max\Big\{\, \Delta\pi^{\mathsf{RC}},\; \big|\bar{\phi}(1)\big|/v \,\Big\}. \tag{25}$$

Furthermore, for every pair $\Delta\rho^{\mathsf{RC}} \geq G\Big(2(\Delta - \Omega)\Big) - \frac{1}{2}$*. Under Assumption 3, the lower bound is strictly positive and is attained at* $|A_1 - A_2| = \Delta - \Omega$*.*

Random role assignment does not close the role gap. By (25), the gap is bounded below by the player gap: the stronger player wins more often whichever role he draws, a component that vanishes only at $A_1 = A_2$. Yet even at an equal matchup, every realised game hands role B the same advantage, worth $\bar{\phi}(1) = [2G(\Delta - \Omega) - 1]v$, so the game stays lopsided whenever the rule-based and baseline advantages fail to cancel.

4.5 Fixed allowance

A fixed-allowance mechanism randomises role assignment but gives role B an allowance τ_{FA}, common to both realised games. For the role gap to vanish, both games must be fair at the same allowance: $\mu^1(\tau_{\mathsf{FA}}) = \mu^2(\tau_{\mathsf{FA}}) = \Delta$. Fairness requires two things at once, that the two advantages coincide and that the common value equals Δ. The first requirement yields

$$2(A_1 - A_2) \;=\; \hat{e}^{\,2}(\tau) - \hat{e}^{\,1}(\tau), \tag{26}$$

where $\hat{e}^k(\tau)$ is the equilibrium input gap in game $\Gamma^k(\tau)$; the right-hand side is a bounded difference of input gaps, set against a left-hand side that grows with the ability gap. However, at a symmetric pair the two input gaps coincide, $\hat{e}^1(\tau) = \hat{e}^2(\tau) =: \hat{e}(\tau)$, so (26) holds for every allowance, and fairness turns entirely on the level: $\hat{e}(\tau_{\mathsf{FA}}) = \Delta - \Omega$.

Proposition 10 (Fixed allowance)**.** *Under a fixed allowance at* τ_{FA}*,* $\Delta\rho^{\mathsf{FA}} = 0$ *if and only if* $\mu^1(\tau_{\mathsf{FA}}) = \mu^2(\tau_{\mathsf{FA}}) = \Delta$*. If* $A_1 = A_2$*, then the condition reduces to*

$\hat{e}(\tau_{\mathsf{FA}}) = \Delta - \Omega$*; if* $A_1 \neq A_2$*, then it requires* (26) *in addition, which no* τ *satisfies once the ability gap is sufficiently large.*

The key result is that a single unconditional allowance is fair only at a knife-edge. Even for equal players, the fixed-allowance mechanism equalises the game only when τ_{FA} coincidentally equates role W's deterministic advantage exactly to B's rule-based advantage. For unequal players, the randomisation makes matters worse: a single allowance cannot undo an ability gap that enters the two role assignments with opposite sign. The best single allowance therefore trades off across the pair distribution and leaves a residual gap almost everywhere. Even so, because $\tau_{\mathsf{FA}} = 1$ reproduces random role assignment, the best fixed allowance weakly dominates it: $\Delta\rho^{\mathsf{FA}} \leq \Delta\rho^{\mathsf{RC}}$.

5 Armageddon chess: institution and data

5.1 Institution

Armageddon chess is used to produce a decisive result when a match or a tournament is tied. One player has White and moves first, while the player with Black wins the match both by winning the game and by drawing it. The draw odds therefore create a rule-based advantage for Black, which is offset by giving Black less time than White. In traditional *fixed-odds Armageddon*, the organiser sets this time allowance in advance and allocates the colours independently of player preferences. The format was first used as a tiebreaker at the FIDE World Chess Championship knockout in Groningen in 1997 and subsequently appeared in a number of elite competitions.[10]

The bidding variant replaces the organiser's fixed allowance with an auction. For a game with base time T, each player i simultaneously submits the amount of time $t_i \in (0, T]$ that he is willing to accept in order to play Black with draw odds. The lower bid wins: that player receives Black and plays with the time he bid, while the losing bidder receives White and the full base time T.[11] We normalise bids by the base time, $\tau_i = \frac{t_i}{T}$, so the realised allowance is $\tau = \min\{\tau_1, \tau_2\}$. Thus $\tau = 1$

[10] The Groningen format gave White six minutes and Black five (Chessgames.com, 1997); the FIDE Knockout retained this convention through 2004 (Weeks, 2004). Other elite events have used different ratios, including five minutes against four at the World Championship (Doggers, 2018; Vargas-Calderón, 2020) and ten against seven at Norway Chess (Norway Chess, 2018). We use the Groningen ratio, $\tau = 5/6$, as the fixed-odds benchmark in Section 7.3.

[11] The bidding format was used at the 2010 US Chess Championship (McClain, 2010; McShane, 2023) and is the standard Armageddon procedure on the Champions Chess Tour, which supplies most of the games in our data (Chess.com, 2026).

gives Black the full clock, while a lower τ represents a smaller allowance, hence a more severe handicap.

This institution maps directly into the model of Sections 2–3. role W is White, with the baseline first-move advantage Ω; role B is Black, with the rule-based advantage of receiving all unresolved states, which in chess are draws. The threshold Δ governs the scope for those unresolved states, while τ is the fraction of the base clock available to Black. The clock concession does not itself represent the input e_B. Rather, shortening Black's clock makes the input more costly to generate, represented in the model by the multiplier $H(\tau)$. The maintained condition $\Delta > \Omega$ has the natural Armageddon interpretation that, in the absence of a clock handicap, Black's draw odds dominate White's first-move advantage; otherwise there would be no interior reason to bid away time for the Black role.

5.2 Data

Chess.com supplied the data: 477 Armageddon games played between August 2022 and October 2024, the large majority of them on the Champions Chess Tour.[12] The distinguishing feature of the data set is that we observe both bids, not just the winning one, letting us test the bid-spread bound derived in Section 6 (Hypothesis H1.D). The raw feed identified players by name. We used the names to attach FIDE ratings from the closest monthly list, and pseudonymised in accordance with European privacy regulation before any econometric work began.

Ties are common in these events. On the Tour, each event runs qualification stages and then three divisions, with the largest prizes in Division I; the remaining games come from other elite Chess.com events. Within a division, play is double elimination, so a player leaves only after losing twice. Matches are multi-game, so they can finish level, and a level match goes straight to a single Armageddon game at a base time of five, ten, or fifteen minutes, with no increment. The bidding then runs exactly as in Section 3.

For each Armageddon game, we observe both bids, both ratings, the colour assignment, the outcome, the match's position in the bracket, and whether the players had already met in an earlier Armageddon game. Ability is measured throughout by the FIDE *standard* Elo rating, for three reasons. First, it exists for every player in the sample, so the analysis is not selected on rating coverage. Second, it rests on the largest number of rated games, so it is the most precisely estimated of the FIDE ratings; on a data set where ability differences are small, that precision matters.

[12]See www.chess.com/about.

Third, where both ratings exist, standard and rapid correlate at 0.85 (White) and 0.89 (Black), and the reduced-form conclusions do not change when the ability gap is measured in rapid ratings.[13] We also record the *stake* of each game, an index on a 1–100 scale that rises with the consequence of the round and with the division; its construction and moments are set out in Online Appendix O.2.1.

Table 1 reports the moments. The ability gap is centred almost exactly on zero (mean −1 Elo, median −7), and the widest gap in the sample is 251 Elo.[14] That compression is a limitation: on a data set this strong, the comparative statics have very little room to move. The bid gaps, by contrast, are large. The winning Black bid averages 73.5% of the base time and the losing White bid 83.2%.[15] Bid levels are flat in the ability gap: sorting the games into equal-count bins on the signed gap, both bid distributions barely move (Online Appendix O.2.1).[16]

Table 1: Summary statistics, 477 elite Armageddon games (August 2022 to October 2024). Bids are expressed as a share of the base time. Sample composition, the rating tertiles, and the corresponding rapid-rating moments are in Online Appendix O.2.1.

Variable	Mean	SD	Median	Min	Max
FIDE standard rating					
White Elo	2653	79.9	2655	2085	2853
Black Elo	2655	88.0	2662	2298	2853
Elo gap (W−B)	−1	90.7	−7	−213	251
Bid, as a share of base time					
White (losing bid)	0.832	0.094	0.828	0.564	1.000
Black (winning bid)	0.735	0.083	0.743	0.400	0.993

6 Implications for the Armageddon application

This section translates the general theory into testable predictions for Armageddon chess: continuation-game and bidding-stage predictions first (Section 6.1), then predictions for the mechanism comparison (Section 6.2).

[13] Rapid ratings cover both players in 472 of the 477 games, with similar moments; Online Appendix O.2.1 reports the comparison.

[14] On the Elo rating system see Elo Rating System.

[15] A few losing bids equal the base time exactly; we treat them as right-censored at the base time when estimating the bid distribution. Some specifications in Section 7 drop games with a missing outcome or a singleton event-fixed-effect cell, so the active sample is smaller than 477; each table reports its own.

[16] Table 1 also gives us a picture of the players' strengths in our sample. In October 2024, FIDE's Top 100 had a cutoff rating of 2637, below the average in our sample.

6.1 Continuation-game and bidding-stage predictions

The equilibrium analysis yields a short list of predictions relating player abilities and game outcomes. We document them as sign-consistency checks against the model. The supporting lemmas and all derivations are collected in Online Appendix O.1.6.

6.1.1 Continuation-game predictions

The equilibrium of the Armageddon game and its comparative statics determine how the second-period outcome varies with the ability gap. Because μ_0 moves one-for-one with $A_W - A_B$, Corollary 1 signs the effect of the ability gap at a fixed τ.

Hypothesis H1.A (Win probability and the ability gap)**.** *When Black plays with a fixed time τ, the probability that White wins the Armageddon game is strictly increasing in $A_W - A_B$, and the probability that Black wins is strictly decreasing in $A_W - A_B$.*

The same comparative statics also determines the sign of the effect of Black's available time at a fixed ability gap: White's win probability falls in τ, since more time for Black both narrows the equilibrium input gap and lowers his marginal cost of input. We flag this as an observation rather than a formal hypothesis, because τ is not exogenous in the data: it is chosen in the bidding stage, so regressing the outcome on the realised τ need not recover the structural sign absent an identification argument that conditions on bidding.

6.1.2 Bidding-stage predictions

The bidding-stage predictions turn on how a player's cutoff $\hat{\tau}_i$ responds to ability, and two structural facts organise them. First, abilities enter the model only through the difference $A_i - A_j$ (Lemma O.3), so ϕ_i is a function of the ability gap alone, and $\partial\phi_i/\partial A_j = -\partial\phi_i/\partial A_i$; comparative statics "in the gap" are therefore well defined. Second, differentiating ϕ_i along the equilibrium paths of the two role-assignment games Γ^i, Γ^j (Lemma O.4) gives

$$\frac{\partial\phi_i}{\partial A_i} = \underbrace{v\Big[g(\Delta-\mu^j) - g(\Delta-\mu^i)\Big]}_{\text{direct}} \underbrace{- v\, g(\Delta-\mu^j)\frac{de^j_W}{dA_i} + v\, g(\Delta-\mu^i)\frac{de^j_B}{dA_i}}_{\text{strategic}}, \tag{27}$$

where de^j_W/dA_i, de^j_B/dA_i are the rival's supply of input responses (Lemma O.5). The direct channel is the role differential in the marginal value of ability: it is positive when own ability is more pivotal in the Black role than in the White role. The strategic channel is positive when the game favourite is whoever holds Black in both games and ambiguous when one player is the favourite regardless of colour. The two channels oppose each other for the stronger player, and the Armageddon specialization of the model does not sign their sum. At an even matchup, the response has the sign of $\Delta - \Omega - \hat{e}(\tau)$: it is positive when draw odds dominate the first-mover edge at that clock—as they do at $\tau = 1$ under Assumption 3—but it can reverse for large ability gaps or short clocks. We therefore leave the direction open.

Definition 2 (Ability regimes)**.** *The players are in regime* (P+) *if* $\partial\phi_i/\partial A_i > 0$ *and in regime* (P−) *if* $\partial\phi_i/\partial A_i < 0$*. By equation* (27)*, the regime at a symmetric match-up is (P+) if and only if* $\Delta - \Omega - \hat{e}(\tau) > 0$*; away from symmetry, it depends on the players, the clock, and the ability gap.*

We do not impose a regime and read its sign off the data. Applying the implicit-function theorem to $\phi_i(\hat{\tau}_i) = 0$ with $\partial\phi_i/\partial\tau > 0$ (Proposition 2; Lemma O.6, Online Appendix O.1) gives

$$\frac{d\hat{\tau}_i}{dA_i} = -\frac{\partial\phi_i/\partial A_i}{\partial\phi_i/\partial\tau}, \qquad \mathrm{sign}\Big(\frac{d\hat{\tau}_i}{dA_i}\Big) = -\,\mathrm{sign}\Big(\frac{\partial\phi_i}{\partial A_i}\Big) \tag{28}$$

Therefore, under (P+) a stronger player accepts less time to obtain Black, under (P−) he concedes less and keeps White. With the equilibrium characterisation (Propositions 3–4), which assigns Black to the player with the smaller cutoff and selects the winning bid $\tau^* = \hat{\tau}_2$, (28) yields the following.

Hypothesis H1.B (Who plays Black)**.** *Whenever* $A_1 \neq A_2$ *and the regime sign is stable over the relevant ability-gap range, the identity of the Armageddon Black player is a deterministic function of the ability ranking: the higher-ability player under regime (P+), the lower-ability player under (P−). Locally, the sign of the empirical relationship between ability rank and the colour assignment* identifies *the governing regime.*

Hypothesis H1.C (The winning bid and the ability gap)**.** *Let* $A_B - A_W$ *be the realised ability gap. The winning bid* τ^* *depends on abilities only through* $A_B - A_W$ *and, within a regime, moves monotonically with it: increasing under (P+) and decreasing under (P−), in the same direction as H1.B. Because the governing regime can itself change as the gap widens, the relation need not be monotone globally; the sign is a local comparative static that the data identify.*

The dispersion of the two bids is governed by the interval over which the losing bidder randomises in the mixed-strategy equilibria of Proposition 3(c). The width the equilibrium can sustain is capped, at the lower endpoint of the mixing interval, by the winner's surplus: at the undominated bid $\tau^* = \hat{\tau}_2$,

$$\bar{\delta}(\hat{\tau}_2) \;=\; \frac{\phi_1(\hat{\tau}_2)}{\frac{d}{d\tau}U_B\Big(\mu^2(\hat{\tau}_2), \hat{\tau}_2\Big)}, \tag{29}$$

where the denominator is non-negative by Assumption 1 and the cap is taken as infinite where it vanishes. Equation (29) is a local, lower-endpoint cap: it bounds the sustainable spread globally only under the additional condition that $\frac{d}{d\tau}U_B(\mu^2(\cdot), \cdot)$ does not rise on the mixing interval (Lemma O.7, Online Appendix O.1).

Hypothesis H1.D (Bid spread)**.** *When $\frac{d}{d\tau}U_B(\mu^2(\hat{\tau}_2), \hat{\tau}_2) > 0$, the within-match bid spread $|\tau_1 - \tau_2|$ is capped near the winning bid by $\bar{\delta}(\hat{\tau}_2)$ in (29); this local cap is zero when $A_1 = A_2$; near equality it rises with $|A_1 - A_2|$ at a rate proportional to the ability response $|\partial\phi_i/\partial A_i|$, and is zero to first order where that response vanishes. Near-equal players must bid almost identically; the realised spread within the cap is fixed by equilibrium selection, and the Armageddon model does not bound the spread globally without a further condition on the shape of $dU_B/d\tau$.*

6.2 Empirical implications for Armageddon chess

In the Armageddon application, the generic role gap of Definition 1 is the colour gap, and we use the latter term when reporting the chess counterfactuals. Section 4 establishes which mechanisms can close this gap and shows that bidding misses the designer benchmark near symmetry under strictly log-concave costs, while leaving the ranking of bidding against fixed odds and random colour open away from symmetry. The magnitude of these gaps and the ranking of the mechanisms are therefore quantitative objects for Section 7.3.

The comparison between bidding and the designer yields two testable predictions. The first is evaluated through the structural counterfactuals of Section 7.3; the second both against realised outcomes in Section 7.1 and within the estimated model in Section 7.3. Throughout this subsection, we maintain strictly log-concave input costs, as in the quadratic specification of the structural model.

Hypothesis H2.A (Bidding does not close the win-probability gap)**.** *For near-symmetric pairs, the win-probability gap of the realised Armageddon game strictly exceeds that under the equalising designer: $\Delta\rho^{\mathsf{BM}} > \Delta\rho^{\mathsf{DO}_\tau}$.*

Equality would require the winning bid to coincide with the equalising allowance, which in Armageddon is implemented through the clock ratio. Proposition 6 rules this out except at a knife-edge. That the designer does at least weakly better follows by construction, since she minimises the same win-probability gap over a set of feasible allowances that includes the equilibrium bid. The empirical content is therefore the size of the shortfall. Away from symmetry, the cost-burden and ability wedges in (21) can oppose one another, and the theory no longer signs the direction of the implementation error; the structural counterfactual quantifies it across ability gaps.

Hypothesis H2.B (The bidding winner is a win-probability underdog)**.** *For near-symmetric pairs, the player who wins the bidding and plays Black wins the match strictly less than half the time.*

This is the reduced-form counterpart of the implementation shortfall. At the equilibrium bid, White's deterministic advantage exceeds the draw threshold, so $\rho_B < \frac{1}{2}$ by Proposition 6. The prediction can therefore be tested directly using the match-win rate of the bidding winner among similar-ability pairs. It also provides a particularly transparent departure from the intuition that the contestant who values the advantaged role most should subsequently be more likely to win.

7 Empirical analysis

The reduced form of Section 7.1 tests the model's directional predictions. The structural analysis of Sections 7.2 and 7.3 estimates the primitives by maximum likelihood and quantifies the mechanism comparison.

7.1 Reduced-form evidence

The reduced form takes the model's directional content and its structural restrictions to the data set directly. Standard errors are clustered at the tournament event throughout: they allow arbitrary correlation of outcomes within an event, so with 21 events in the data set the inference rests on 21 effectively independent blocks. The choice is conservative: with only 21 clusters, the standard errors are wide, and significance is hard to reach. The data set is compressed in ability, with a quarter of the games separated by fewer than 30 Elo points, so the tests concern signs and restrictions rather than magnitudes, which are left to the structural estimation of Section 7.2. Table 2 collects the five reduced-form tests; the full regression tables are in Online Appendix O.2.

Table 2: Summary of reduced-form tests. Event-clustered standard errors for regressions; exact binomial confidence intervals for shares. $N = 477$.

Hyp.	Kind	Prediction	Estimate	Verdict
H1.A	Sign test	Black wins less as White's Elo edge grows ($\beta < 0$)	-0.277 (0.136), $p = 0.042$; wild-cluster $p \approx 0.10$	Directional support
H1.B	Regime identification	Sign of the rank–colour relation identifies (P+)/(P−)	share $= 0.526$ $[0.480, 0.572]$	Regime not identified; weak (P+) lean
H1.C	Conditional sign / sufficiency	Level term zero given stakes; pooled gap slope null under regime mixing	gap 0.0048 (0.0031), $p = 0.140$; level $p = 0.644$	Consistent
H1.D	Bound	Spread bound vanishes as $\lvert\text{gap}\rvert \to 0$	Q1 contrast -0.0158 (0.0089), $p = 0.090$; fitted spread at zero gap 0.076 (0.007)	Mixed: smaller near symmetry, but does not vanish; structural bound check required
H2.B	Underpowered test	Bid winner wins $< \frac{1}{2}$ near symmetry	0.471 $[0.380, 0.564]$, $N = 121$	Direction consistent; underpowered (MDE $= 0.127$)

Note. Gaps per 100 Elo points; spreads as shares of base time; the H2.B window is $\lvert\text{gap}\rvert$ at or below the 25th percentile. H1.A's wild-cluster correction is reported in the text.

H1.A says that a stronger White should win the Armageddon game more often. This is the one well-powered test. In a logit of the probability that Black wins on the signed FIDE gap, with stakes, base-time and winning-bid quintile controls, the gap coefficient is -0.277 per 100 Elo points (s.e. 0.136, $p = 0.042$). The estimate stays negative and of stable magnitude in every specification, from no controls to event fixed effects, and under the finer session-level clustering (Table 4, Online Appendix O.2; a session is an event–round–date block). With 21 event clusters, conventional cluster-robust inference can over-reject; corrections designed for few clusters place the two-sided p-value near 0.10 (wild cluster bootstrap-t on the linear-probability specification, $p = 0.107$; effective-score wild bootstrap for the logit, $p = 0.098$), so we read the test as directional support rather than decisive evidence (Table 5, Online Appendix O.2). We should mention one caveat: because the time control is itself chosen in the bidding stage, the coefficient confirms only the direction of the ability effect but not its size.

H1.B says that the identity of the Black player reveals the regime of Definition 2: if the value of playing Black rises with a player's own ability (P+), the stronger player wins the bid; if it falls (P−), the weaker player does. The higher-rated player takes Black in 52.6% of matches (95% CI $[0.480, 0.572]$; Table 6, Online Appendix O.2). The share sits above one half in both halves of the rating-gap split and at the longest base time, though it falls below one half at the two shorter base times (Table 6, Online Appendix O.2); no subsample share is statistically distinguishable from one half. The data lean towards (P+), the regime in which the stronger player values Black more, and they cannot settle the question. Section 7.2

returns to it with the estimated model.

H1.C says that, within a regime, the winning bid should move with the ability gap, and that rating levels should not matter once the gap is fixed (gap sufficiency). In a regression of the winning bid on the signed gap, the average rating, the stake and base-time controls, the pooled gap slope is null (+0.0048 per 100 Elo points, $p = 0.140$; Table 8, Online Appendix O.2). This is consistent with what the theory predicts for a sample that pools the two regimes: in (P+) the winning bid rises with the gap, in (P−) it falls, and the pooled slope averages the two. Gap sufficiency is the sharper test: the average-rating level enters with a coefficient indistinguishable from zero, with the stake control ($p = 0.644$) and without it ($p = 0.26$). The stake, the model's prize shifter and itself correlated with pair strength, enters negatively and significantly (-0.042, $p = 0.022$): higher-stake games attract more aggressive bids.

H1.D says that near-equal players should post nearly identical bids. The prediction is a bound, not a slope: the equilibrium caps the spread between the two bids by the bid winner's surplus, the cap vanishes as abilities equalise, and the theory does not restrict the spread's level away from equality. The data agree in direction: the spread is smallest among the most evenly matched fifth of pairs (contrast -0.016, $p = 0.090$) and rises with the gap over the lower half of the gap range ($p = 0.002$; Table 9, Online Appendix O.2). The level does not vanish, however: the fitted spread at a zero gap is 0.076 of the base time, about 46 seconds on a 600-second clock. This may reflect latent ability differences, equilibrium selection/bid dispersion, or forces outside the model.

Finally, H2.B says that among near-equal pairs the bid winner should win the match less than half the time. Among the 121 pairs in the lowest quartile of the absolute gap, the bid winner's match-win rate is 0.471 (95% CI $[0.380, 0.564]$), and it sits below one half in every symmetry window we form, from the median split down to the tightest window (Table 11, Online Appendix O.2); no interval excludes one half. The power of the test is the constraint. Both magnitudes are deviations of the bid winner's win rate from one half: the model implies a deviation of one to two percentage points, while the smallest deviation detectable with 121 games is roughly 13 points. A deviation of the model-implied size would take several thousand games to detect.

7.2 Structural estimation

In the theoretical model, the bidding outcome is deterministic given abilities, yet the observed bids and colour assignments vary across games. The estimation must account for that dispersion. Each game γ contributes three observations: the colour assignment, the winning bid τ^*_γ, and the outcome y_γ (one if Black wins the match). With X_γ collecting the covariates (the signed FIDE rating gap, the stake, and the base time), the likelihood contribution of game γ can be expressed as

$$L_\gamma(\theta) \;=\; \underbrace{\Pr\Big(\text{colour}_\gamma,\, \tau^*_\gamma \mid X_\gamma;\theta\Big)}_{\text{auction factor}} \times \underbrace{\Pr\Big(y_\gamma \mid \tau^*_\gamma, \text{colour}_\gamma, X_\gamma;\theta\Big)}_{\text{outcome factor}}.$$

The colour assignment and the winning bid are jointly determined by the latent ability gap through the two cutoffs, so the auction data enter as one joint probability.

The auction factor describes the colour assignment and winning bid that the model determines conditional on abilities and the maintained equilibrium selection. To generate the dispersion observed in the data, we introduce one source of latent heterogeneity. We write $A_W - A_B = \beta_E\,\text{EloGap}_\gamma + \nu_\gamma$, with $\nu_\gamma \sim N(0,\sigma^2_\nu)$, where ν_γ is the latent deviation of the game-relevant ability gap from that implied by the FIDE rating gap. Within the structural model, players observe the realised ability gap when they bid, whereas the econometrician observes ratings but not ν_γ. Integrating over ν_γ therefore gives the colour assignment and winning bid their joint distribution.

Bids are submitted in whole seconds and heap on round numbers. A recorded bid therefore corresponds to an interval of model bids, and the bid-interval component of the likelihood is the probability assigned to that interval. The auction factor rests on one further maintained assumption: equilibrium selection. We condition on the weak-dominance-refined outcome in which the winning bid is the losing bidder's indifference point. The selection itself is not a source of randomness; it is imposed explicitly rather than treating the equilibrium multiplicity as an identified set. The losing bid, whose equilibrium distribution the theory restricts without determining uniquely, remains outside the likelihood and enters only through the bid-spread check of Hypothesis H1.D, in the incomplete-model tradition of Haile and Tamer (2003).

The outcome factor requires no assumption beyond the model: the theory itself makes the chess result stochastic. At the realised abilities and the observed clock, the Black-win probability is $\rho_B = G(\Delta - \mu)$, with μ measuring White's equilibrium advantage. The recorded result is one Bernoulli draw at that probability, with no auxiliary error and no free scale.

For estimation, we adopt standard forms: quadratic input cost $c(e) = \frac{\kappa}{2}e^2$, Gaussian noise $\varepsilon \sim N(0,\sigma^2)$, handicap-cost multiplier $H(\tau) = \exp\big(\zeta(\tau-1)\big)$ with $\zeta < 0$, and prize $v = \exp(\beta_v\, s/100)$ at stake s. The ability gap follows the specification above and carries no average-rating term. Abilities enter only through their difference, a property the additive contest implies. Section 7.1 found no residual level effect once stakes were controlled, so the restriction is imposed in estimation. The likelihood sees the prize and the cost curvature only through v/κ, so we normalise $\kappa = 1$. Seven parameters remain, $\theta = (\beta_E, \beta_v, \Omega, \Delta, \sigma, \zeta, \sigma_\nu)$, estimated by joint maximum likelihood; the latent ν_γ is integrated numerically on a fixed grid.[17]

The two factors of L_γ divide the identification of θ between them. The outcome factor identifies the ratios Δ/σ, Ω/σ, and β_E/σ. Within the auction factor, the bid level and its stake gradient identify Δ, Ω, ζ, and β_v jointly. The residual bid width pins σ_ν, and the colour component carries the regime information. One restriction spans the factors and makes the specification refutable: the same ν_γ must move both the bids and the outcomes, in a proportion the theory fixes.

Table 3: Structural estimates: joint MLE at the best-known point, 477 games, $\log L = -2856.12$. Panel B reports the scale-free ratios; interval methods are tagged by row (see note).

Panel A: parameter estimates			
Parameter		Estimate	95% CI
β_E	rating-gap loading	-0.00026	$[-0.00077,\ +0.00031]^{\mathrm{p}}$
β_v	stake elasticity of the prize	-0.395	$[-0.750,\ +0.072]^{\mathrm{b}}$
Ω	White first-mover advantage	0.312	$[0.066,\ 0.394]^{\mathrm{b}}$
Δ	draw threshold	0.390	$[0.161,\ 0.470]^{\mathrm{b}}$
σ	contest noise scale	1.032	$[0.525,\ 1.151]^{\mathrm{b}}$
ζ	time-cost slope	-2.301	$[-4.257,\ -0.800]^{\mathrm{b}}$
σ_ν	ability-wedge scale	0.920	$[0.496,\ 1.213]^{\mathrm{b}}$
κ	cost curvature	1 (normalisation)	
Panel B: identified ratios			
Δ/σ	draw threshold, noise units	0.378	$[0.320,\ 0.412]^{\mathrm{p}}$
Ω/σ	first-mover advantage, noise units	0.302	$[0.082,\ 0.624]^{\mathrm{b}}$
σ_ν/σ	wedge scale, noise units	0.892	$[0.803,\ 1.102]^{\mathrm{b}}$

Note: $^{\mathrm{p}}$ marks a 95% profile-likelihood interval; $^{\mathrm{b}}$ a percentile interval from a 300-draw player-pair block bootstrap, with ratios computed draw by draw. Because the resampled optima drift along the common-scale valley (Section 7.4), the bootstrap intervals for the six parameter levels in Panel A measure sampling dispersion and carry no formal coverage claim.

Table 3 reports the joint estimates. The ability-gap coefficient is effectively zero: $\hat{\beta}_E = -0.00026$, against $+0.00107$ when the outcome factor is maximised alone. The rating gap moves the outcomes and does not move the bids. Decompos-

[17] The indicator factors make the likelihood surface finely ridged, so the maximiser is L-BFGS-B wrapped in a deterministic pattern search from multiple starts. Replication code and the full estimation record are in the replication package.

ing the maximised log-likelihood (-2856.12) locates the source: outcome -331.29, colour -370.02, bid -2154.80. The positive slope comes from the outcome factor; the colour and bid-interval components, which carry most of the likelihood, hold the joint estimate at zero. Fixing β_E at the outcome-implied $+0.00107$ and re-optimising the remaining parameters lowers the maximised log-likelihood by 2.60 points (likelihood ratio 5.20 on one degree of freedom, $p = 0.023$; Table 12, row S4, Online Appendix O.2). The outcome factor improves by 3.6 points. The bid and colour factors worsen by 4.5 and 1.7 and more than offset the gain. The result reproduces, inside the structure, the reduced-form nulls of H1.B and H1.C. It also traces the reduced form's silence on the regime question of H1.B to its source. With $\hat{\beta}_E$ this close to zero, the model places every pair near symmetry, and the assignment of Black flips between optima the likelihood cannot tell apart. No regime claim survives.

One property of the likelihood disciplines how these numbers are read: in this sample, the auction factor identifies the parameter levels only weakly. The surface is maximised not at a point but along a common-scale valley: moving Δ, Ω, σ_ν and σ together barely changes the log-likelihood. The ratios, by contrast, are precisely determined. Panel B of Table 3 reports these ratios. Inference accordingly uses ratios and profile-likelihood intervals rather than levels. The profile for β_E is nonetheless sharp: its 95% interval, $[-0.00077, +0.00031]$, contains zero and excludes the outcome-implied $+0.00107$ by a wide margin. A 300-draw player-pair block bootstrap corroborates it: the percentile interval is $[-0.00090, +0.00098]$.

The counterpart of the flat bid map is the estimated dispersion in the latent component, $\hat{\sigma}_\nu = 0.92$, or 0.89 in noise units: the same order as the contest noise itself. Empirically, this dispersion is composite. It may reflect genuine unobserved differences in game-relevant strength, error in the rating proxy, misspecification of the bid equation, or some combination of the three; the likelihood cannot apportion them. What the estimates establish more cleanly is the negative result: the FIDE rating gap predicts outcomes but does not account for the observed variation in bidding.

At $\hat{\theta}$, the worst-case contraction ratio of Condition 1 across the 477 games is 0.22, against the 0.95 cap imposed in estimation. The equilibrium characterisation behind the likelihood therefore applies with slack throughout the sample. Assumption 1 holds at $\hat{\theta}$ as well: across the feasible clock range and the full estimated ability range, the input response in its part (i) is strictly positive, and the cost-relief inequality in its part (ii) holds with a relative margin of at least 55 percent. No admissible specification makes the bids respond to the rating gap (see Section 7.4).

7.3 Mechanism comparison

We now take the estimated primitives to the four mechanisms: random colour, fixed odds at $\tau = 5/6$, bidding at the model's equilibrium bid, and the designer DO_τ on the institutional window $[\underline{\tau}, 1]$, with $\underline{\tau} = 0.40$, the lowest winning bid in the data set. Each mechanism is scored by its colour gap (Section 6.2). Random colour and fixed odds are accordingly averaged over their own fair coin rather than evaluated at the realised colour.

We evaluate the mechanisms conditional on each game's observed FIDE rating gap under two treatments of the latent component ν_γ. In the first, we set $\nu_\gamma = 0$, so that the ability gap is the value implied directly by the FIDE gap, $\hat{\beta}_E\,\mathrm{EloGap}_\gamma$. We call this the *FIDE-implied* reading. In the second, we retain the same observed FIDE gap but integrate over $\nu_\gamma \sim N(0, \hat{\sigma}_\nu^2)$. We call this the *FIDE-conditioned integrated* reading. The latter therefore incorporates the estimated heterogeneity in game-relevant ability left unexplained by FIDE ratings. Because $\hat{\beta}_E$ is close to zero, the FIDE-implied ability gaps are themselves close to zero, while under the integrated reading most of the variation in the model's ability gap comes from ν_γ.

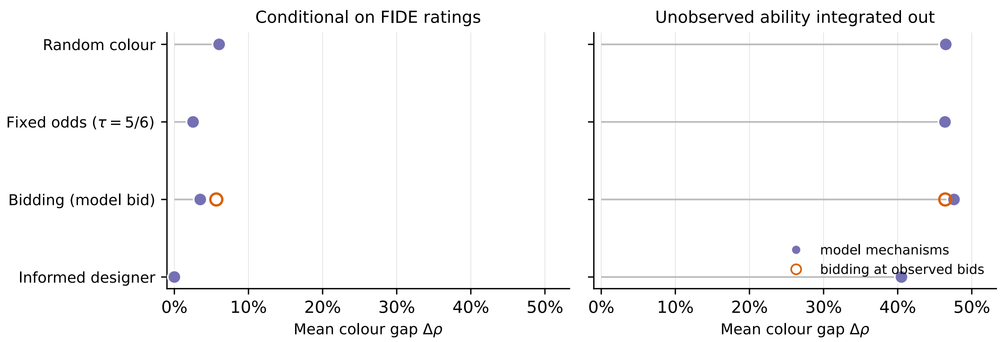

Figure 1: Mean colour gap $\Delta\rho^M$ under the four mechanisms at the structural estimates: conditional on the observed FIDE ratings. Left: FIDE-implied ability $\nu = 0$. Right: integrating over the estimated latent heterogeneity ν. Open markers: the realised games at the observed bids. $N = 477$.

Under the FIDE-implied reading (Figure 1, left panel), the designer equalises every game exactly. The equalising allowance (clock ratio) is interior in all 477 games, and $\Delta\rho^{\mathsf{DO}_\tau} = 0$. Bidding leaves 3.49 percentage points against 6.04 under random colour. H2.A says that the realised game stays more lopsided than the informed designer's. The estimates confirm this in every game. The full ranking $\mathsf{DO}_\tau < \mathsf{BM} < \mathsf{RC}$ holds in 465 of the 477. The exceptions all lie on the bidding–random comparison. At the model's bid the winner is a strict underdog, with a Black win probability of 0.483. This is the structural counterpart of H2.B. It sits close to the observed rate among the most evenly matched quarter of pairs, where Black won 0.471 of matches (Section 7.1). Every pair is near-symmetric under this reading, so both hypotheses cover the whole data set. The magnitudes should be read with caution. The model's scale is set by the draw-odds wedge in noise units, $(\Delta - \Omega)/\sigma = 0.076$. This is a level-type object of the kind the likelihood identifies only weakly. The rankings, by contrast, are stable across the near-optimal parameter points. Fixed odds leave 2.52 points, below the bidding number at the reported point. The ordering of this pair, however, reverses across those near-optimal points. The theory does not rank this pair. Neither, with any stability, do the estimates.

Under the FIDE-conditioned integrated reading (Figure 1, right panel), the same observed FIDE gaps are retained, but the counterfactual averages over the estimated residual heterogeneity ν around the FIDE-implied ability gap. This produces much larger gaps because most latent draws place the contestants far from symmetry. Under the structural interpretation of ν as game-relevant ability known to the players, the random-colour gap rises to 46.49 points. The designer still observes the abilities. The equalising clock is nonetheless interior on only 10.9% of the mass of ν, and those draws are equalised exactly. The rest sits at a corner, 42.5% at $\underline{\tau}$ and 46.6% at $\tau = 1$, where even the most extreme feasible clock leaves a gap. Her residual of 40.51 points therefore accumulates entirely at the boundary. Widening the window to $\underline{\tau} = 0.25$ lowers her residual only to 39.96 and changes nothing else in the comparison (Table 13, Online Appendix O.2). Bidding does not improve on random colour. Its gap of 47.60 points exceeds the random-colour gap in every one of the 477 games, by 1.11 points on average. The excess holds across the near-optimal parameter points. At the reported parameter point, away from symmetry the weaker player wins the bid, so bidding assigns that player Black deterministically at a conceded clock. Random colour assigns him Black, at the full clock, half the time. Bidding improves on random colour only in a narrow band of near-equal latent ability gaps. At $\hat{\sigma}_\nu = 0.92$ almost all of the estimated mass lies outside that band. Section 6.2 left the sign of BM against RC open. The estimates resolve it against bidding. The full ranking under this reading, $\mathsf{DO}_\tau < \mathsf{FA} < \mathsf{RC} < \mathsf{BM}$, holds pairwise in all 477 games. Because ν may also absorb rating error and model misspecification, the integrated magnitudes should be interpreted as upper bounds on the contribution of latent ability.

Three numbers summarise the bidding column. The gap at the model bid (3.49 conditional on FIDE; 47.60 integrated over ν) is what the model predicts. The gap at the observed bids (5.66; 46.42) is a policy statement about the games as played, not a test of the model. The fit between the two bid series is a mean absolute deviation of 0.107 of base time. Relative to the observed mean bid (0.735), the $\nu = 0$ model bid overshoots (0.837) and the ν-integrated mean bid undershoots (0.710). What bidding is worth thus depends on the information over which the counterfactual is evaluated. Conditional on FIDE ratings, it narrows the gap for the near-symmetric pairs implied by the joint estimate. Integrating over the estimated distribution of ν, it widens the average gap.

7.4 Fit and robustness

The most demanding test of fit is whether the model can reconcile its predicted Black-win rate with the observed one. At the observed clocks the FIDE-conditional model implies a Black win rate of about 47% (the 5.66-point gap at the observed bids, Section 7.3). Black won 50.3% of the 477 games. Integrating over ν closes the distance. The colour and bid choices are informative about ν. Conditioning on them pulls the prediction back to the observed rate. The outcome factor at the joint optimum, -331.29, sits within 0.7 points of the fair-coin benchmark $-477 \ln 2 = -330.63$. A misspecified clock effect in the outcome equation could be absorbed by the same latent component and cannot be separated from it on these data. On the bid side the misfit is the mean absolute deviation of 0.107 reported in Section 7.3.

The data determine the scale-free ratios (Table 3, Panel B) and the gap coefficient β_E. The nuisance levels are only weakly identified along the common-scale valley. Profile-likelihood intervals are therefore the primary inference. At each value of the profiled parameter they re-maximise the likelihood over the remaining parameters, so the levels adjust freely rather than distort the interval. A 300-draw player-pair block bootstrap corroborates the profile intervals (Figure 2, Online Appendix O.2). The levels are reported with bootstrap dispersion bands that carry no formal coverage claim. The reported point is the best point found by the profile search (Online Appendix O.2).

Across every admissible specification the bids stay flat in the rating gap. Uniform ± 2.5-second rounding intervals leave $\hat{\beta}_E$ at zero (-0.00015). Allowing the variance of ν to depend on the rating level or the absolute gap, the only theory-consistent channels for a level effect, yields no evidence of a variance channel (zero coefficient, LR 0.15; Table 12, Online Appendix O.2). An average-rating level term yields a likelihood ratio of 1.20 ($p = 0.27$). Gap sufficiency survives inside the structure. The outcome-anchored diagnostic is the 5% rejection reported in Section 7.2. The unrestricted spline in the gap yields a likelihood ratio of 211.6, but the fitted shape is even and U-shaped. An ability effect must be antisymmetric under relabelling of the players, so the improvement is a specification diagnostic, not an ability effect.

The odd projection, the only admissible ability shape, is flat: a likelihood ratio of 3.3 on four degrees of freedom ($p \approx 0.5$).

8 Discussion and concluding remarks

This paper studies a mechanism-design problem in which contestants bid for an advantaged role by forfeiting part of the productive resource in the contest that follows. The attraction of such a mechanism is informational. When the designer does not know the contestants' relative strengths, allowing them to price the advantage themselves appears to decentralise the choice of allowance to those who possess the relevant information. Our analysis shows why this logic is incomplete when the bid changes the technology of the continuation contest. The informed designer and the bidders solve different problems: the designer chooses an allowance to balance winning probabilities, whereas bidding stops at the marginal bidder's payoff-indifference point.

This distinction gives the bidding mechanism a precise allocation property. The equilibrium assignment is a no-swap allocation: at the realised allowance, neither contestant would prefer the other's role. Bidding therefore produces preference-based balance rather than outcome-based balance. The two coincide only when the equilibrium cost burdens associated with the two roles coincide. With strictly log-concave costs, they do not. For near-symmetric contestants, the productive handicap lowers the equilibrium cost burden of the advantaged role sufficiently that its holder is willing to accept a lower winning probability. Bidding consequently overshoots the equalising allowance and makes the auction winner a win-probability underdog. Away from symmetry, differences in ability introduce an additional wedge and the direction of the implementation error is no longer generally signed, but coincidence with the designer's allowance remains a knife-edge condition.

Armageddon chess provides an unusually direct setting in which to study this mechanism. Black receives the rule advantage of winning the match on a draw, while bidding determines both who receives Black and how much clock time that player sacrifices. The clock concession is not a transfer that leaves the contest once the auction is over. It changes the conditions under which the subsequent game is played by altering the cost of generating input. This is the source of the divergence between willingness to pay for the advantage and the allowance that balances the contest. It is also what distinguishes the mechanism from a standard auction in which payment enters utility additively and does not itself modify the continuation technology.

The empirical results reinforce this distinction. FIDE ratings predict game outcomes but carry essentially no relationship with the bids themselves, and the structural estimates reproduce the same tension: the outcome component of the likelihood ties game-relevant strength to ratings, while the auction component holds that relationship close to zero, so reconciling the two requires substantial dispersion in latent ability unexplained by ratings. How that dispersion is interpreted matters directly for the mechanism comparison reported in Section 7.3: read literally as private information available to the players, it reverses bidding's ranking relative to random assignment, though the same dispersion may also absorb rating error and model misspecification, so the reversal is best read as an upper bound on what latent ability alone can explain.

Several limitations follow directly from this interpretation. The sample contains elite players whose ratings are tightly compressed, leaving limited variation with which to identify the ability comparative statics of the bidding model. The common scale of several structural parameters is also weakly identified, which is why our inference emphasises scale-free ratios, profile likelihoods, and rankings that remain stable across near-optimal parameter values. The losing-bid distribution is only partially disciplined by the theory, and the structural likelihood therefore uses the selected winning-bid outcome rather than attempting to fit the complete pair of bids. Finally, the latent dispersion cannot be decomposed on these data into genuine unobserved strength, rating error, and model misspecification. A broader sample spanning a substantially wider range of contestant abilities would provide a much sharper test of both the bidding comparative statics and the mechanism comparison.

A further limitation concerns external validity: every estimate in this paper comes from one institution, elite online Armageddon chess. The mechanism's two defining features—the bid is paid in a resource allowance, and it changes the technology of the contest that follows—are not both present in the closest real-world precedents. Auction komi in Go is a score transfer: it shifts the winning threshold but leaves the game's technology untouched, so it carries no cost-burden channel of the kind that drives the overshoot result here. A+B construction contracts do price a productive quantity, but the winning bid buys the contract itself rather than an in-contest advantage over a rival who continues to compete under the same rules. Corroborating evidence from another productive-currency, in-contest setting—a comparably structured tie-break or seeding auction in another sport, for instance—would let us test whether the overshoot and reversal results documented here are a general feature of this class of mechanisms or an artefact of Armageddon chess's particular cost and prize structure.

The broader lesson is that decentralising an allowance to informed contestants does not by itself implement the choice of an informed designer. When the price of an advantage is a forfeited share of the resource allowance, the bid changes both the contestants' winning probabilities and the costs they incur in the subsequent contest. Willingness to pay therefore reveals payoff indifference, not the allowance that equalises winning chances. Bidding can still perform a useful allocation function: it assigns the advantage at terms that neither contestant wants to swap. But preference-based balance and outcome-based balance are different objectives. When the payment stays inside the game and changes the technology of competition itself, that distinction is fundamental.

Appendix

The appendix contains the complete proofs of the results stated in the body (Appendix A), except for Propositions 7 and 8, whose proofs are in the Online Appendix. The Online Appendix also contains supplementary theory, including the derivations behind Section 6.1, a parametric example, and additional empirical results.

A Proofs

A.1 Proof of Lemma 1

Using evenness of g, each player's marginal return to supplying input is $M(\mu)v$, where $\mu = \mu_0 + e_W - e_B$. Hence any interior equilibrium satisfies $c'(e_W) = M(\mu)v = c'(e_B)$. Since c' is strictly increasing, $e_W = e_B$, and therefore $\mu = \mu_0$, giving $e_W = e_B = e^\star = (c')^{-1}\big(M(\mu_0)v\big)$. If $M(\mu_0) = 0$, the same expression gives the corner $e^\star = 0$. If $M(\mu_0) > 0$, then the candidate is interior because $c'(0) = 0$ and the marginal return at zero is positive.

It remains only to check sufficiency and uniqueness. Since $M'(\mu) = \frac{1}{2}[g'(\Delta+\mu) - g'(\Delta-\mu)]$, we have $|M'(\mu)| \leq \dot{g}$. The own-input second derivatives, $vM'(\mu) - c''(e_W)$ and $-vM'(\mu) - c''(e_B)$, are both bounded above by $v\dot{g} - m < 0$ under Condition 1. Thus, each payoff is strictly concave in own input, so the first-order conditions are sufficient. They also force $e_W = e_B$, hence the candidate above is the unique pure-strategy equilibrium. Substituting $e_W = e_B$ into the outcome probabilities gives the expressions in the lemma. □

A.2 Proof of Proposition 1

Part (i). The first-order conditions are equivalent to the fixed point (6). Differentiating the map in (7) and using $c'' \geq m$ and $H(\tau) \geq 1$ bounds its slope (Online Appendix O.1.2): Condition 1 gives $|\eta'(x)| \leq v\dot{g}/m \leq \lambda < 1$, so η is a contraction and the contraction-mapping theorem makes the solution of (6) unique. This pins down a unique candidate (e_W, e_B) through (5). The same condition gives $-vg'(\Delta - \mu) - c''(e_W) < 0$ and $vg'(\Delta - \mu) - H(\tau)c''(e_B) < 0$, so each first-order condition is sufficient. In the maintained non-degenerate case $g(\Delta - \mu) > 0$ at the solution, the marginal return at zero is positive and the candidate is interior.

Part (ii). Since c' is strictly increasing and $H(\tau) \geq 1$, (5) gives $e_W^* \geq e_B^*$, with strict inequality when $\tau < 1$. The contraction argument above also gives differentiability of the fixed point in τ. Differentiating $\hat{e}^*(\tau) = \eta(\hat{e}^*(\tau), \tau)$ shows that $d\hat{e}^*/d\tau$ carries the sign of η_τ, since $|\eta_x| < 1$; and $\eta_\tau < 0$ follows directly from (7) because $H'(\tau) < 0$. The three differentiation displays are in Online Appendix O.1.2. Thus, $d\hat{e}^*/d\tau < 0$. At $\tau = 1$, $H(1) = 1$ makes the right-hand side of (7) identically zero, so the unique fixed point is $\hat{e}^*(1) = 0$; continuity gives $\lim_{\tau \uparrow 1} \hat{e}^*(\tau) = 0$.

Part (iii). Since $c'(e_W) = H(\tau)c'(e_B)$, we have

$$H(\tau)\, c(e_B) \leq c(e_W) \quad \iff \quad \frac{c(e_B)}{c'(e_B)} \leq \frac{c(e_W)}{c'(e_W)}.$$

Because $e_W \geq e_B$, this inequality holds when $e \mapsto c(e)/c'(e)$ is non-decreasing, equivalently when $\log c$ is concave; it reverses under log-convexity. Strict log-curvature and $\tau < 1$ give strict inequalities. □

A.3 Proof of Corollary 1

The contraction argument in Proposition 1 makes the equilibrium input gap differentiable in each primitive $\theta \in \{v, \tau, \Delta, \mu_0\}$, with $d\hat{e}/d\theta = \eta_\theta/(1-\eta_e)$ evaluated at the fixed point and $1-\eta_e > 0$. The fixed-point map depends on (Δ, μ_0, e) only through $\Delta - \mu_0 - e$, so $-\eta_\Delta = \eta_{\mu_0} = \eta_e$, and $d\hat{e}/d\mu_0 = -\,d\hat{e}/d\Delta = \eta_e/(1-\eta_e)$, which exceeds $-1/2$ because $|\eta_e| < 1$. Applying the chain rule to $\rho_W(\theta) = 1 - G(\Delta - \mu_0 - \hat{e}(\theta))$ and collecting the direct threshold terms then sign the two threshold effects, $d\rho_W/d\mu_0 > 0$ and $d\rho_W/d\Delta < 0$; the four chain-rule displays are in Online Appendix O.1.2. For τ and v, there is no direct threshold term. Proposition 1(ii) gives $d\hat{e}/d\tau < 0$, hence $d\rho_W/d\tau < 0$. The sign of η_v is not pinned down by the maintained assumptions, so the prize effect is ambiguous. Since $\rho_B = 1 - \rho_W$, all signs reverse for role B. □

A.4 Proof of Lemma 2

In each game, a player's first-order condition cancels the total effect of his own input: this is the envelope theorem applied to (4). For role B in game $\Gamma^j(\tau)$, it leaves W's input response and the direct handicap term, $\frac{d}{d\tau}U_B = -v\, g(\Delta - \mu^j)\,\frac{de^j_W}{d\tau} - H'(\tau)c(e^i_B)$. For W in $\Gamma^i(\tau)$, τ enters only through equilibrium inputs, and the same argument leaves only B's input response, $\frac{d}{d\tau}U_W = -v\, g(\Delta - \mu^i)\,\frac{de^j_B}{d\tau}$. Substituting the first-order conditions $v\, g(\Delta - \mu^j) = H(\tau)c'(e^i_B)$ and $v\, g(\Delta - \mu^i) = H(\tau)c'(e^j_B)$ gives (10)–(11). □

A.5 Proof of Proposition 2

The implicit function theorem applied to (5) gives a C^1 equilibrium input profile under Condition 1; hence, ϕ_i is C^1 with derivative (12) (Lemma 2). Assumption 1, applied to the ordered pairs (i, j) and (j, i), makes both terms in (12) non-negative and at least one strictly positive, so $\phi_i'(\tau) > 0$ on $(0, 1]$. □

A.6 Proof of Proposition 3

Write

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

so that

$$\phi_i(\tau) = \mathcal{B}_i(\tau) - \mathcal{W}_i(\tau).$$

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 $\hat{\tau}_i$.

For part (a), fix a bid $b < \hat{\tau}_i$ for player i and compare it with $\hat{\tau}_i$ against an arbitrary opponent bid t. If $t < b$, then both bids lose and give $\mathcal{W}_i(t)$. If $t = b$, then bidding b gives the tie lottery, worth no more than $\mathcal{W}_i(b)$ because $\phi_i(b) < 0$, while $\hat{\tau}_i$ loses and gives $\mathcal{W}_i(b)$. If $b < t < \hat{\tau}_i$, then bid b wins and gives $\mathcal{B}_i(b) \le \mathcal{B}_i(t) < \mathcal{W}_i(t)$, the strict inequality because $\phi_i(t) < 0$, while $\hat{\tau}_i$ loses and gives $\mathcal{W}_i(t)$. If $t = \hat{\tau}_i$, then bidding $\hat{\tau}_i$ ties and gives $\mathcal{B}_i(\hat{\tau}_i) = \mathcal{W}_i(\hat{\tau}_i) \ge \mathcal{B}_i(b)$. If $t > \hat{\tau}_i$, then both bids win role B, and monotonicity of $\mathcal{B}_i$ gives $\mathcal{B}_i(\hat{\tau}_i) \ge \mathcal{B}_i(b)$. Thus, $\hat{\tau}_i$ is never worse than b, proving part (a).

If $\hat{\tau}_1 = \hat{\tau}_2 = \hat{\tau}$, then the profile $(\hat{\tau}, \hat{\tau})$ is a Nash equilibrium. At the tie, each player receives $\mathcal{B}_i(\hat{\tau}) = \mathcal{W}_i(\hat{\tau})$. A lower bid gives the role B player at some $y < \hat{\tau}$ and payoff $\mathcal{B}_i(y) \le \mathcal{B}_i(\hat{\tau})$; a higher bid gives the player role W at $\hat{\tau}$ and payoff $\mathcal{W}_i(\hat{\tau})$. Thus no unilateral deviation is strictly profitable, proving part (b).

For part (c), suppose $\hat{\tau}_1 < \hat{\tau}_2$, and fix $\tau \in (\hat{\tau}_1, \hat{\tau}_2]$. We verify that, for small enough $\delta > 0$, player 1 bidding τ and player 2 mixing uniformly on $[\tau, \tau + \delta]$ is a Nash equilibrium. Against player 1's pure bid τ, player 2 weakly prefers role W at every $x \leq \tau$, since $\mathcal{B}_2(x) \leq \mathcal{B}_2(\tau) \leq \mathcal{W}_2(\tau)$. Any bid $x > \tau$ gives exactly $\mathcal{W}_2(\tau)$, while the tie bid $x = \tau$ gives no more because $\phi_2(\tau) \leq 0$. Hence every bid in $(\tau, \tau + \delta]$ is a best response; the continuously uniform mix assigns zero probability to the endpoint τ.

It remains to check player 1. A bid $y < \tau$ wins for sure but yields $\mathcal{B}_1(y) \leq \mathcal{B}_1(\tau)$. A bid $y > \tau + \delta$ loses for sure and gives at most $\mathcal{W}_1(\tau) < \mathcal{B}_1(\tau)$, because $\tau > \hat{\tau}_1$. For $y \in [\tau, \tau + \delta]$, player 1's expected payoff is

$$EU_1(y) = \frac{1}{\delta} \int_{\tau}^{y} \mathcal{W}_1(s)\, ds + \left(1 - \frac{y - \tau}{\delta}\right) \mathcal{B}_1(y).$$

Differentiating,

$$EU_1'(y) = -\frac{1}{\delta} \phi_1(y) + \left(1 - \frac{y - \tau}{\delta}\right) \mathcal{B}_1'(y).$$

Since $\phi_1(\tau) > 0$ and both ϕ_1 and $\mathcal{B}_1'$ are continuous, $EU_1'(y) < 0$ throughout the interval once δ is small enough; Lemma O.7 (Online Appendix O.1) records the explicit sustainable-width bound. Thus, $y = \tau$ is player 1's best response, which proves part (c), and the realised winning bid in the constructed equilibria lies in $[\hat{\tau}_1, \hat{\tau}_2]$. □

A.7 Proof of Proposition 4

Consider any bid $s < \hat{\tau}_2$ by player 2 and compare it with $\hat{\tau}_2$ against an arbitrary opponent bid t. The five-case comparison in the proof of Proposition 3(a), read at $i = 2$ and $b = s$, applies unchanged: $\hat{\tau}_2$ is never worse than s, and it is strictly better against $t \in [s, \hat{\tau}_2)$, where $\phi_2(t) < 0$.

Thus $\hat{\tau}_2$ weakly dominates every $s < \hat{\tau}_2$, with a strict improvement against opponent bids in $[s, \hat{\tau}_2)$ whenever the cutoff is interior. Any equilibrium in Proposition 3 with $\tau^* < \hat{\tau}_2$ requires player 2 to put positive probability on such weakly dominated bids. The weak-dominance refinement therefore eliminates all those equilibria and leaves only $\tau^* = \hat{\tau}_2$ within the family. □

A.8 Proof of Proposition 5

By Proposition 1(ii), $\hat{e}$ is continuous and strictly decreasing on $\mathcal{T}$ with $\hat{e}(1) = 0$; since $\mu^W(\tau)$ differs from $\hat{e}(\tau)$ by the τ-independent constant $(A_W - A_B) + \Omega$, $\mu^W(\cdot)$ is continuous and strictly decreasing with range $[(A_W - A_B) + \Omega,\ (A_W - A_B) + \Omega + \hat{e}(\underline{\tau})]$.

The objective $|\mu^W(\tau) - \Delta|$ is the absolute value of the strictly monotone $h(\tau) = \mu^W(\tau) - \Delta$, hence strictly quasiconvex, with a unique minimiser. If $h(\underline{\tau}) > 0 > h(1)$, equivalently $0 < \Delta - \Omega - (A_W - A_B) < \hat{e}(\underline{\tau})$, then the intermediate value theorem and strict monotonicity give a unique interior zero τ_{DO}, where $|h| = 0$, and substituting $\mu^W = \Delta$ into (17), we have that (19) holds with $\rho_B = \frac{1}{2}$. Otherwise, h keeps one sign on $\mathcal{T}$, and $|h|$ is minimised at the nearer endpoint. This proves part (i). For part (ii), the assignment of role W enters only through $A_W - A_B \in \{A_1 - A_2,\, A_2 - A_1\}$, giving the two targets $\Delta - \Omega \mp (A_1 - A_2)$; part (i) applied game by game shows that an interior equaliser exists for role assignment k if and only if its target lies in $(0, \hat{e}^{\,k}(\underline{\tau}))$, the window of Γ^k, and $\mathsf{DO}_{c\tau}$ takes the role with the smaller residual. Its option set contains the realised role, so $\mathsf{DO}_{c\tau}$ weakly dominates DO_{τ}; if each target lies outside its own window, both roles sit at a boundary and $\mathsf{DO}_{c\tau}$ leaves a positive gap. □

A.9 Proof of Proposition 6

At $A_1 = A_2$, the relabelling $1 \leftrightarrow 2$ maps the game to itself, so $\phi_1 \equiv \phi_2$, the cutoffs coincide at $\hat{\tau}$, and $\mu^1(\tau) = \mu^2(\tau) = \bar{\mu}(\tau) = \Omega + \hat{e}(\tau)$. The indifference $\phi(\hat{\tau}) = U_B(\bar{\mu}(\hat{\tau}), \hat{\tau}) - U_W(\bar{\mu}(\hat{\tau}), \hat{\tau}) = 0$ rearranges to

$$\Big[\, 2G(\Delta - \bar{\mu}(\hat{\tau})) - 1 \,\Big]\, v \;=\; H(\hat{\tau})\, c(e_B^*) - c(e_W^*). \tag{A.1}$$

The left side of equation (A.1) vanishes if and only if $\bar{\mu}(\hat{\tau}) = \Delta$, equivalently $\hat{\tau} = \tau_{\mathsf{DO}}$, by strict monotonicity of $\bar{\mu}$; the right side vanishes if and only if $H(\hat{\tau})c(e_B^*) = c(e_W^*)$. This proves the first claim. For strictly log-concave c and $\hat{\tau} < 1$, Proposition 1(iii) gives $H(\hat{\tau})c(e_B^*) < c(e_W^*)$, so the right side of (A.1) is negative. Hence $G(\Delta - \bar{\mu}(\hat{\tau})) < 1/2$, or $\bar{\mu}(\hat{\tau}) > \Delta$. Since $\bar{\mu}(\tau_{\mathsf{DO}}) = \Delta$ and $\bar{\mu}$ is strictly decreasing, $\tau_{\mathsf{DO}} > \hat{\tau}$. □

A.10 Proof of Proposition 9

Write $x_k := 2G(\Delta - \mu^k(1)) - 1$ for the signed margin under the role assignment with player k in role W. Random role assignment gives $\Delta\rho^{\mathsf{RC}} = \frac{1}{2}(|x_1| + |x_2|)$, and $|x_1| + |x_2| = \max\{|x_1 + x_2|, |x_1 - x_2|\}$. Now $\pi_1 - \pi_2 = G(\Delta - \mu^2(1)) - G(\Delta - \mu^1(1)) = \frac{1}{2}(x_2 - x_1)$, so $\frac{1}{2}|x_1 - x_2| = \Delta\pi^{\mathsf{RC}}$. For the sum, $H(1) = 1$ implies $e_W = e_B =: e^{\Gamma^k}$ in each $\Gamma^k(1)$, so $\phi_i(1) = \Big[G(\Delta - \mu^j(1))\, v - c(e^{\Gamma^j})\Big] - \Big[\Big(1 - G(\Delta - \mu^i(1))\Big)v - c(e^{\Gamma^i})\Big]$; averaging over i cancels costs and gives $\bar{\phi}(1) = v(x_1 + x_2)/2$, hence (25).

For the lower bound, set $x := \Delta - \Omega$ and $S(d) := G(x - d) + G(x + d) - 1$

for $d = A_1 - A_2 \geq 0$. Assumption 3 implies $S(d) > 0$ and therefore $x > 0$. Equation (25) becomes $\Delta\rho^{\mathsf{RC}} = \max\{G(x+d) - G(x-d), S(d)\}$. The first branch is increasing in d, the second is non-increasing by single-peakedness, and they cross when $S(d) - [G(x+d) - G(x-d)] = 2G(x-d) - 1 = 0$, i.e., at $d = x$. The minimum of the maximum is therefore $G(2x) - \frac{1}{2} > 0$. □

A.11 Proof of Proposition 10

Under a fixed allowance, the role is randomised and role B receives the common allowance τ_{FA}. Hence $\Delta\rho^{\mathsf{FA}} = \frac{1}{2}\Big(|x_1(\tau_{\mathsf{FA}})| + |x_2(\tau_{\mathsf{FA}})|\Big)$ with $x_k(\tau) = 2G(\Delta - \mu^k(\tau)) - 1$, which vanishes if and only if $x_1 = x_2 = 0$, i.e. if and only if $\mu^1(\tau_{\mathsf{FA}}) = \mu^2(\tau_{\mathsf{FA}}) = \Delta$. Writing the baselines out, $\mu^1(\tau) - \mu^2(\tau) = 2(A_1 - A_2) - (\hat{e}^{\,2}(\tau) - \hat{e}^{\,1}(\tau))$, so equality of the two advantages is (26). At a symmetric pair, (26) holds trivially, and the remaining condition is $\hat{e}(\tau_{\mathsf{FA}}) = \Delta - \Omega$. For $A_1 \neq A_2$, each input gap lies in $[0, \hat{e}^{\,k}(\underline{\tau})]$ (Proposition 1(ii)), so the right-hand side of (26) is bounded while the left-hand side grows linearly in the ability gap; once $2|A_1 - A_2|$ exceeds $\max_k \hat{e}^{\,k}(\underline{\tau})$, no τ satisfies it. □

References

Amegashie, J. A. (2012). Productive versus destructive efforts in contests. *European Journal of Political Economy*, 28(4):461–468.

Anbarcı, N., Sun, C.-J., and Ünver, M. U. (2021). Designing practical and fair sequential team contests: The case of penalty shootouts. *Games and Economic Behavior*, 130:25–43.

Apesteguia, J. and Palacios-Huerta, I. (2010). Psychological pressure in competitive environments: Evidence from a randomized natural experiment. *American Economic Review*, 100(5):2548–2564.

Beviá, C. and Corchón, L. C. (2013). Endogenous strength in conflicts. *International Journal of Industrial Organization*, 31(3):297–306.

Che, Y.-K. and Hendershott, T. (2008). How to divide the possession of a football? *Economics Letters*, 99(3):561–565.

Chen, B., Jiang, X., and Wang, Z. (2022). Single- and double-elimination tournaments under psychological momentum. *The B.E. Journal of Theoretical Economics*, 22(2):509–525.

Chess.com (2026). Chess.com official 2026 event rulebook. `https://www.chess.com/article/view/chesscom-event-rulebook`. Accessed 2026-08-19.

Chessgames.com (1997). Groningen candidates. https://www.chessgames.com/perl/chess.pl?tid=38596. Accessed 2026-08-19.

Chowdhury, S. M., Esteve-González, P., and Mukherjee, A. (2023). Heterogeneity, leveling the playing field, and affirmative action in contests. *Southern Economic Journal*, 89(3):924–974.

Clark, D. J., Kundu, T., and Nilssen, T. (2026). Instrument choice in pre-contest investment. Unpublished manuscript, available at SSRN: https://ssrn.com/abstract=6432640.

Clark, D. J. and Nilssen, T. (2018). Keep on fighting: The dynamics of head starts in all-pay auctions. *Games and Economic Behavior*, 110:258–272.

Dasgupta, P. and Maskin, E. (1986). The existence of equilibrium in discontinuous economic games, I: theory. *The Review of Economic Studies*, 53(1):1–26.

Deck, C. and Kimbrough, E. O. (2015). Single- and double-elimination all-pay tournaments. *Journal of Economic Behavior & Organization*, 116:416–429.

DeMarzo, P. M., Kremer, I., and Skrzypacz, A. (2005). Bidding with securities: Auctions and security design. *American Economic Review*, 95(4):936–959.

Deng, S., Wang, X., and Wu, Z. (2018). Incentives in lottery contests with draws. *Economics Letters*, 163:1–5.

Deng, X. and Qi, Q. (2009). Priority right auction for komi setting. In Leonardi, S., editor, *Internet and Network Economics*, pages 521–528, Berlin, Heidelberg. Springer Berlin Heidelberg.

Doggers, P. (2018). World Chess Championship: Carlsen vs Caruana. Chess.com. https://www.chess.com/article/view/world-chess-championship-2018-carlsen-caruana. Accessed 2026-08-19.

Engel, E., Fischer, R., and Galetovic, A. (2001). Least-present-value-of-revenue auctions and highway franchising. *Journal of Political Economy*, 109(5):993–1020.

Franke, J., Leininger, W., and Wasser, C. (2018). Optimal favoritism in all-pay auctions and lottery contests. *European Economic Review*, 104:22–37.

Franke, J. and Metzger, L. P. (2023). Repeated contests with draws. Ruhr Economic Papers 1016, RWI – Leibniz-Institut für Wirtschaftsforschung.

Fu, Q. and Wu, Z. (2020). On the optimal design of biased contests. *Theoretical Economics*, 15(4):1435–1470.

Fullerton, R. L. and McAfee, R. P. (1999). Auctioning entry into tournaments. *Journal of Political Economy*, 107(3):573–605.

Gelder, A., Kovenock, D., and Roberson, B. (2022). All-pay auctions with ties. *Economic Theory*, 74(4):1183–1231.

Goel, S. and Goyal, A. (2023). Optimal tie-breaking rules. *Journal of Mathematical Economics*, 108:102872.

Granot, D. and Gerchak, Y. (2014). An auction with positive externality and possible application to overtime rules in football, soccer, and chess. *Operations Research Letters*, 42(1):12–15.

Haile, P. A. and Tamer, E. (2003). Inference with an incomplete model of English auctions. *Journal of Political Economy*, 111(1):1–51.

Huang, L. (2016). Prize and incentives in double-elimination tournaments. *Economics Letters*, 147:116–120.

Jehiel, P. and Moldovanu, B. (2006). Allocative and informational externalities in auctions and related mechanisms. In Blundell, R., Newey, W. K., and Persson, T., editors, *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress*, pages 102–135. Cambridge University Press.

Kalwij, A. and De Jaegher, K. (2023). The age-performance relationship for a cognitive-intensive task: Empirical evidence from chess grandmasters. *Sports Economics Review*, 2:100010.

Kimbrough, E. O., Laughren, K., and Sheremeta, R. (2020). War and conflict in economics: Theories, applications, and recent trends. *Journal of Economic Behavior & Organization*, 178:998–1013.

Konrad, K. A. (2009). *Strategy and Dynamics in Contests*. Oxford University Press.

Krishna, K., Lychagin, S., Olszewski, W., Siegel, R., and Tergiman, C. (2026). Pareto improvements in the contest for college admissions. *Review of Economic Studies*, 93(1):629–663.

Lazear, E. P. and Rosen, S. (1981). Rank-order tournaments as optimum labor contracts. *Journal of Political Economy*, 89(5):841–864.

Lewis, G. and Bajari, P. (2011). Procurement contracting with time incentives: Theory and evidence. *Quarterly Journal of Economics*, 126(3):1173–1211.

Li, B., Wu, Z., and Xing, Z. (2023). Optimally biased contests with draws. *Economics Letters*, 226:111076.

McClain, D. L. (2010). New way to crown winners in games that end in ties. *New York Times*, Chess column, May 30.

McShane, L. (2023). Bidding one's time. The Spectator. https://spectator.com/article/bidding-ones-time/. Accessed 2026-08-19.

Moldovanu, B. and Sela, A. (2001). The optimal allocation of prizes in contests. *American Economic Review*, 91(3):542–558.

Münster, J. (2007). Contests with investment. *Managerial and Decision Economics*, 28(8):849–862.

Nalebuff, B. J. and Stiglitz, J. E. (1983). Prizes and incentives: towards a general theory of compensation and competition. *The Bell Journal of Economics*, 14(1):21–43.

Norway Chess (2018). New and exciting playing format for 2019. https://norwaychess.no/en/2018/10/10/new-and-exciting-playing-format-for-2019/. Accessed 2026-08-19.

Palacios-Huerta, I. (2025). The beautiful dataset. *Journal of Economic Literature*, 63(4):1363–1423.

Pollard, R., Prieto, J., and Gómez, M.-Á. (2017). Global differences in home advantage by country, sport and sex. *International Journal of Performance Analysis in Sport*, 17(4):586–599.

Schaller, Z. and Skaperdas, S. (2020). Bargaining and conflict with up-front investments: How power asymmetries matter. *Journal of Economic Behavior & Organization*, 176:212–225.

Siegel, R. (2014). Asymmetric contests with head starts and nonmonotonic costs. *American Economic Journal: Microeconomics*, 6(3):59–105.

Szech, N. (2015). Tie-breaks and bid-caps in all-pay auctions. *Games and Economic Behavior*, 92:138–149.

Vargas-Calderón, V. (2020). Are armageddon chess games implemented fairly? *ICGA Journal*, 41(4):180–190.

Vesperoni, A. and Yildizparlak, A. (2019). Contests with draws: Axiomatization and equilibrium. *Economic Inquiry*, 57(3):1597–1616.

Weeks, M. (2004). FIDE World Championship 2004: Background. https://www.mark-weeks.com/aboutcom/aa04f19.htm. Accessed 2026-08-19.

Zhuang, L., Chen, Z., and Keppo, J. (2024). The impact of tie intervals on rank-based contests. Working paper, National University of Singapore.