WHEN DOES NASH AVERAGING IGNORE A DUPLICATE? THE EXACT SCOPE OF REDUNDANCY INVARIANCE

Anonymous authors
Paper under double-blind review

ABSTRACT

Nash averaging reads pairwise results as a zero-sum game and rates each agent by its expected result against an equilibrium. When there is more than one it takes the maximum-entropy equilibrium, the most spread-out one. The method was proposed with the property that an exact copy changes no rating. We show that the property can fail, and that the published proof covers the set of equilibria, not the maximum-entropy equilibrium. A copy tilts the entropy in proportion to the copied agent's probability. We prove that the maximum-entropy equilibrium stays in place exactly when no step inside the equilibrium set gains more from this tilt than it loses in entropy. When it moves, only agents that no equilibrium plays can be re-rated, while the rest tie at the top. A copy never moves a unique equilibrium, including those of the Atari score table and chatbot-arena leaderboard we check. Coarse scores leave more than one equilibrium in forty of the forty-four football seasons studied, often through exact draws. In seventeen of them, adding a copy of one team reorders two others. More than one equilibrium lets a rating jump at an exact draw. Any rule that never jumps is wrong somewhere nearby by at least half that rating's range; a vanishing relaxation of the equilibrium conditions costs at most the whole range and loses invariance to copies. A published Nash average should say which rule picked its equilibrium and whether others exist.

1 INTRODUCTION

Suppose a chess club rates its members by the average score each member obtains against the others. If one member registers a second time under a new name, the two entrants draw each other and play everyone else as before. Yet other members' averages can shift, because that member now counts twice among their opponents. A rating that moves when an entrant is copied can be gamed, and ratings from two pools cannot be compared directly.

Nash averaging (Balduzzi et al., 2018) was proposed to remove this problem from the evaluation of learning agents. It treats the matrix of pairwise results as a two-player zero-sum game, takes an equilibrium of that game, and rates each agent by its expected result against that equilibrium. When the game has more than one equilibrium, it takes the one of maximum entropy. The property claimed for it, invariance to redundancy, is that an exact copy of an agent changes no rating. Later work on rating methods restates it (Lanctot et al., 2023; Marris et al., 2025; Liu et al., 2025).

The property already fails with five agents. Let agents 1, 2 and 3 play rock-paper-scissors, let each of them beat agent 4, and let agent 5 draw against everyone (Figure 1, left). Every win scores $+1$ and every loss -1. The equilibria are exactly the mixtures of the even distribution on agents 1 to 3 with the point mass on agent 5. The maximum-entropy equilibrium puts probability $1/4$ on each of agents 1, 2, 3 and 5. Agent 4 loses to three of them and is rated $-3/4$. Now enter agent 5 a second time. The maximum-entropy equilibrium of the new game gives $1/5$ to each of the five it plays, so agent 5 and its copy hold $2/5$ together and agent 4 is rated $-3/5$. Agent 4's own results are untouched, and the copy draws with it just as agent 5 does.

The published proof covers the set of equilibria: the equilibria of the new game are the old ones with agent 5's probability split between the two entrants in every proportion. It does not show that the maximum-entropy equilibrium stays where it was. Splitting a probability p_5 evenly between two

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

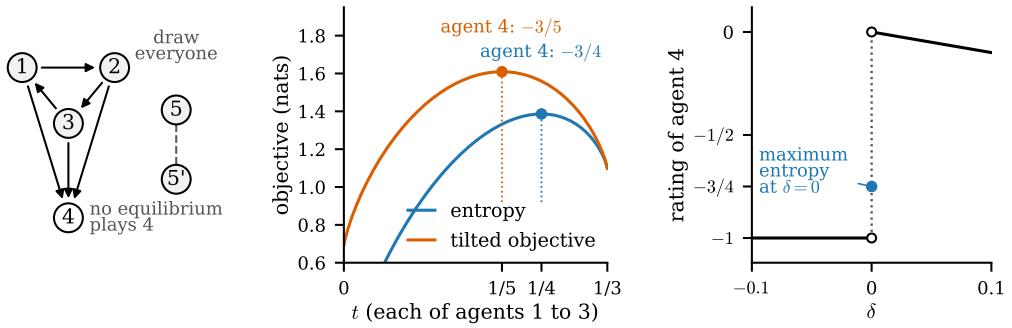


Figure 1: A copy of agent 5 moves the maximum-entropy equilibrium and the rating r_4 of agent 4 (middle); an exact draw makes r_4 jump (right). Left: the five-agent game with the copy $5'$ added, where an arrow points from winner to loser, a missing arrow is a draw, and the dashed edge marks the copy. Middle: the entropy H along the equilibrium set $\{(t, t, t, 0, 1-3t)\}$, and the tilted objective $H + p_5 \log 2$, $p_5 = 1 - 3t$, after a copy of agent 5. Each maximizer and the rating it gives agent 4 are marked. Right: r_4 when agent 5 scores δ against the others and nothing is copied.

identical entrants adds $p_5 \log 2$ to the entropy. Maximizing entropy in the new game then rewards probability on agent 5 more than before, and the maximizer moves toward agent 5 (Figure 1, middle). When the original game has a unique equilibrium there is nothing to move, and the property holds. The condition under which a copy leaves the maximum-entropy equilibrium in place has not been characterized, either by the original proof or by the later work that restates the property. We therefore ask: exactly when does a copy leave it in place, what changes when it does not, and do the populations that people rate have more than one equilibrium?

Our contributions are the following. (1) We prove that a copy of agent k adds $p_k \log 2$ to the entropy over the original equilibria, and that the maximum-entropy equilibrium stays in place exactly when no step inside the equilibrium set gains more from this tilt than it loses in entropy (Theorems 1 and 2). A unique equilibrium always passes this criterion, and one entropy maximization and one linear program decide it in general. Call an agent inessential when no equilibrium plays it. If every inessential agent is rated strictly below zero at the maximum-entropy equilibrium, the criterion holds exactly when the copied agent's probability is the same in every equilibrium (Corollary 1).

(2) We prove what changes when the equilibrium moves. The copied agent and its copy together never lose probability, and gain it whenever the equilibrium moves (Theorem 3). Essential agents, those some equilibrium plays, keep the top rating, zero, under every selection rule. Only inessential agents' ratings can change (Theorem 4), and so can their order. As the copy becomes exact, each quantity that is the same in every equilibrium tends to its exact-copy value: the copied agent's total probability, or an agent's rating (Proposition 1).

(3) We show that more than one equilibrium also lets ratings jump: at an exact draw every rule that picks an equilibrium can be discontinuous in the data (Proposition 2, Figure 1, right). A continuous rule must then be wrong somewhere nearby by at least half a rating's range at the draw. Relaxing the equilibrium conditions comes within a factor of two of that bound as the relaxation shrinks (Theorem 5), and gives up invariance to copies.

(4) We measure where more than one equilibrium occurs and what copies change there. The 2000 Gaussian games we sample, a public Atari score table and a chatbot-arena leaderboard have a unique equilibrium. In 40 of 44 seasons of five European football leagues the equilibrium is not unique, and in 17 of them adding a copy of one team reverses the order of two others. Every rating rule we test, Nash averaging included, either re-rates under an exact copy or jumps at an exact draw. Section 3 proves the results for every antisymmetric outcome matrix and works most of them out on the five-agent example, Section 4 reports the measurements, and Section 5 discusses related work.

2 Preliminaries

Outcome matrices and equilibria. An outcome matrix is an antisymmetric $A \in \mathbb{R}^{n \times n}$, $A^\top = -A$, whose entry A_{ij} is the log-odds that agent i beats agent j. A draw is a zero entry, and since only antisymmetry is used below, score differences and signs are also admissible entries. In the game of

A two players each pick a distribution over the agents and the row player receives $p^\top Aq$. This game is symmetric and zero-sum, and its value is zero because $p^\top Ap = 0$ for every p. A distribution is optimal for both players exactly when no agent has a positive expected result against it (Gale et al., 1950), and the equilibrium set is

$$\mathcal{P}(A) = \{p \in \Delta_n : \; Ap \preceq 0\}, \tag{1}$$

where Δ_n is the probability simplex and $\preceq$ is coordinatewise; the equilibrium set is a nonempty convex polytope. The five-agent game of Figure 1 has the outcome matrix

$$A_5 = \begin{pmatrix} 0 & 1 & -1 & 1 & 0 \\ -1 & 0 & 1 & 1 & 0 \\ 1 & -1 & 0 & 1 & 0 \\ -1 & -1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}, \tag{2}$$

and its equilibrium set is the segment $\{(t, t, t, 0, 1-3t) : t \in [0, 1/3]\}$, because any probability on agent 4, or unequal weights on agents 1 to 3, would give some agent a positive expected result.

Nash averaging. The Shannon entropy $H(p) = -\sum_i p_i \log p_i$ (natural logarithm, $0 \log 0 = 0$) is strictly concave, so it has a unique maximizer on $\mathcal{P}(A)$, the maximum-entropy equilibrium. Nash averaging (Balduzzi et al., 2018) rates each agent by its expected result against it:

$$p^\star(A) = \arg\max_{p \in \mathcal{P}(A)} H(p), \qquad r(A) = A\, p^\star(A). \tag{3}$$

Every rating is at most zero, because $p^\star$ is an equilibrium; zero is the top rating. The rating of every agent that $p^\star$ plays is exactly zero, because ${p^\star}^\top A p^\star = \sum_i p^\star_i (Ap^\star)_i = 0$ is a sum of nonpositive terms. A Nash-average leaderboard therefore distinguishes only the agents that $p^\star$ does not play; the others share the top rating. On the segment of A_5 agent 4 is rated $-3t$ and every other agent 0; entropy is largest at $t = 1/4$, so the ratings are $(0, 0, 0, -3/4, 0)$.

Other rules that pick an equilibrium. We call any rule that picks one point of $\mathcal{P}(A)$ a selection rule. The original paper defines the rating by the maximum-entropy equilibrium but computes it with a linear-programming solver (Balduzzi et al., 2018), while the implementation in OpenSpiel (Lanctot et al., 2019) maximizes entropy. Both kinds of rule are therefore in use. Asking a solver for any point of $\mathcal{P}(A)$ means giving it the constraints $\{Ap \preceq 0,\; p \in \Delta_n\}$ and a zero objective. The objective then does not distinguish one equilibrium from another, and a solver of the simplex family returns a vertex of $\mathcal{P}(A)$, which we call the LP vertex (zero objective). Which vertex it returns depends on the solver and on the order of the agents. The standard minimax program, $\max\{v : \; p \in \Delta_n,\; Ap \preceq -v\mathbf{1}\}$, has value zero and the same optimal set, and a solver run on it may return another vertex, the LP vertex (minimax), which we also test. In A_5 the two vertices are the ends of the segment: $t = 0$, where agent 4 is rated 0, and $t = 1/3$, where it is rated -1.

Essential agents and width. An agent is essential if some equilibrium gives it positive probability, and inessential otherwise (Gale et al., 1950). Some equilibrium q strictly beats every inessential i: $(Aq)_i < 0$. We call the largest range of one agent's probability over $\mathcal{P}(A)$ its width; the width is positive exactly when the equilibrium is not unique. With e essential agents and $A_\mathcal{E}$ their submatrix, the equilibrium set has dimension $e - \operatorname{rank}(A_\mathcal{E}) - 1$, so the equilibrium is not unique exactly when $\operatorname{rank}(A_\mathcal{E}) \le e - 2$ (Bohnenblust et al., 1950; Gale & Sherman, 1950; Gale et al., 1950). This rank condition needs an exact coincidence in the data: tournament matrices, whose off-diagonal entries are all ± 1, have a unique equilibrium (Fisher & Ryan, 1992; Laffond et al., 1993; 1997; Le Breton, 2005), and so does almost every matrix drawn from an absolutely continuous distribution (Brandl, 2017, Proposition 4). In A_5 agent 4 is the only inessential agent, $\operatorname{rank}(A_\mathcal{E}) = e - 2$, and the width is 1, the range of p_5.

Copies and the invariance claim. Duplicating agent k appends a row and a column equal to those of agent k, with a draw between agent k and the new agent, its copy; $\operatorname{dup}_k(A)$ denotes the duplicated game. The supplement of Balduzzi et al. (2018) states the invariance claim formally (its Definition 3 and Lemma 1). Whenever the copied column differs from every other column, the maximum-entropy equilibrium of $\operatorname{dup}_k(A)$ should give agent k and its copy half of $p^\star_k(A)$ each and leave every other probability unchanged. We say that the copy leaves the equilibrium in place when this holds. It then leaves every rating in place, but a rating can also stay unchanged while the equilibrium moves (Theorem 4(b)). Their Lemma 1 proves the statement about the set of equilibria; maximum entropy appears once beside it, in a remark about a copied row that already has copies. A_5 with $k = 5$ meets the hypothesis, yet both the equilibrium and the rating of agent 4 move (Section 1).

3 WHAT A COPY DOES TO THE MAXIMUM-ENTROPY EQUILIBRIUM

3.1 A COPY TILTS THE ENTROPY TOWARD THE COPIED AGENT

Throughout, k is the duplicated agent, A_5 is the five-agent game of Figure 1, and proofs are in Appendix A. To compare distributions before and after a copy, write $\iota_\alpha(p) \in \Delta_{n+1}$ for the distribution that splits p_k into αp_k on agent k and $(1-\alpha)p_k$ on its copy, and $\pi(p') \in \Delta_n$ for the distribution that adds the probability of the copy back to agent k.

Lemma 1 (Equilibria of the duplicated game; first claim due to Balduzzi et al. 2018)**.** *For every antisymmetric A and every k, $\mathcal{P}(\mathrm{dup}_k A) = \{\iota_\alpha(p) : p \in \mathcal{P}(A),\ \alpha \in [0,1]\}$. Duplication also gives $(\mathrm{dup}_k A)\, p' = \big(A\,\pi(p'),\, (A\,\pi(p'))_k\big)$ for every $p' \in \Delta_{n+1}$, so every original agent's rating in the duplicated game is its rating at $\pi(p')$.*

The lemma holds because agent k and its copy score the same against everyone and draw each other; no agent's expected result depends on how their probability is divided. How the probability is divided does change the entropy, and Lemma 2 computes by how much.

Lemma 2 (Entropy of a split)**.** *$H(\iota_\alpha(p)) = H(p) + p_k\, H_2(\alpha)$, where H_2 is the binary entropy. For fixed p the maximum over α is attained at $\alpha = 1/2$, with value $H(p) + p_k \log 2$.*

Theorem 1 (Tilt formula)**.** *$p^\star(\mathrm{dup}_k A) = \iota_{1/2}\big(\arg\max_{p \in \mathcal{P}(A)} [H(p) + p_k \log 2]\big)$, and the maximizer over $\mathcal{P}(A)$ is unique by strict concavity. If agent k is present in m copies that draw one another, the original counted among them, then $\log 2$ becomes $\log m$ and the even split an m-way split. Replacing each agent k of a set K by m_k copies tilts the entropy by $\sum_{k \in K} p_k \log m_k$.*

We call the added term $p_k \log 2$ the tilt. In A_5 with $k = 5$ the tilted objective along the equilibrium set is $H(t, t, t, 0, 1-3t) + (1-3t)\log 2$. Its maximum is at $t = 1/5$ rather than $t = 1/4$ (Figure 1, middle). At $t = 1/5$ the duplicated game gives probability $1/5$ to each of agents 1, 2, 3 and 5 and to the copy of agent 5.

3.2 WHEN THE MAXIMUM-ENTROPY EQUILIBRIUM STAYS IN PLACE

A copy leaves $p^\star$ in place exactly when $p^\star$ also maximizes $H + p_k \log 2$ over $\mathcal{P}(A)$ (Theorem 1). Because the tilted objective is concave, it suffices to look at small steps from $p^\star = p^\star(A)$ that stay in $\mathcal{P}(A)$. A step v is feasible at $p^\star$ if $p^\star + sv \in \mathcal{P}(A)$ for some $s > 0$. The one-sided derivative of entropy along such a step, $D_v H(p^\star)$, is finite because a feasible step vanishes on the inessential agents and $p^\star$ is positive on the essential ones (Theorem 4(a)).

Theorem 2 (Criterion for the equilibrium to stay in place)**.** *Let $p^\star = p^\star(A)$. Then $p^\star(\mathrm{dup}_k A) = \iota_{1/2}(p^\star)$ if and only if $D_v H(p^\star) + v_k \log 2 \le 0$ for every step v that is feasible at $p^\star$. Two conditions are sufficient: p_k is the same in every equilibrium, or $p^\star$ already gives agent k its largest probability over $\mathcal{P}(A)$. Suppose $p^\star$ lies in the relative interior of $\mathcal{P}(A)$, meaning that every step toward another equilibrium can also be taken a little way backwards. Then the copy leaves the equilibrium in place exactly when p_k is the same in every equilibrium.*

In words, the equilibrium stays in place exactly when no step inside the equilibrium set gains more from the tilt than it loses in entropy. The theorem's last claim holds because entropy has zero slope at $p^\star$ along every step that can also be reversed inside $\mathcal{P}(A)$. If such a step changes p_k, the tilt adds a positive slope in the direction that raises p_k, and the equilibrium moves. The three-agent continuity example of Balduzzi et al. (2018) meets the second sufficient condition at the parameter value where two of its agents draw exactly. There the maximum-entropy equilibrium already gives one agent its largest probability over all equilibria, so copying that agent changes nothing.

Checking the criterion takes one entropy maximization for $p^\star$ and one linear program over $\mathcal{P}(A)$, since the essential agents are exactly the agents $p^\star$ plays. Put $w_i = -\log p_i^\star$ on the essential agents and $w_i = 0$ elsewhere, and add $\log 2$ to w_k; because $\partial H / \partial p_i = -\log p_i - 1$ and a feasible step has $\sum_i v_i = 0$, this makes $w^\top v$ the tilted derivative along v, and the criterion holds exactly when $\max_{q \in \mathcal{P}(A)} w^\top q \le w^\top p^\star$ (Appendix A.3).

Corollary 1 (Criterion when inessential agents lose strictly)**.** *Suppose every inessential agent has a strictly negative rating at $p^\star(A)$, that is, $(Ap^\star(A))_i < 0$ for every inessential i. Then $p^\star(\mathrm{dup}_k A) =$*

$\iota_{1/2}(p^\star(A))$ *if and only if* p_k *is the same in every equilibrium. A unique equilibrium is the special case in which every probability is the same in every equilibrium.*

The corollary restates the criterion in terms of the ratings. The assumption says that $(p^\star, p^\star)$ is a quasi-strict equilibrium (Harsanyi, 1973; Jansen, 1981), and it holds exactly when $p^\star$ lies in the relative interior of $\mathcal{P}(A)$ (Appendix A.8). In A_5 agent 4 is rated $-3/4 < 0$, and the corollary applies. A copy of agent 4, whose probability is zero in every equilibrium, changes nothing, while a copy of any of agents 1, 2, 3 or 5 moves the equilibrium. Without the assumption a copy can leave the equilibrium in place although p_k could rise. We find this in 25 of the 153 football duplications of Section 4.3. The constraint of an inessential agent rated zero can make every step that raises p_k impossible to reverse, and entropy can then fall along such a step by at least as much as the tilt adds.

3.3 What changes when the equilibrium moves

Theorem 3 (A copy gains probability)**.** *Let* $p^\dagger = \arg\max_{\mathcal{P}(A)}[H + \lambda p_k]$ *with* $\lambda > 0$*, so that* $p^\star = \arg\max_{\mathcal{P}(A)} H$ *is the case* $\lambda = 0$*. Then* $(p^\dagger)_k \geq (p^\star)_k$*, with strict inequality whenever* $p^\dagger \neq p^\star$*, and* $\lambda \mapsto (\arg\max_{\mathcal{P}(A)}[H + \lambda p_k])_k$ *is nondecreasing. Consequently duplication never lowers the total probability of agent* k *and its copy, and raises it whenever the maximum-entropy equilibrium moves.*

A new maximizer that differs from $p^\star$ gives up entropy, since $p^\star$ already had the most; it can score higher on the tilted objective only by putting more probability on agent k. In A_5 the total probability of agent 5 rises from $1/4$ to $2/5$ at $\lambda = \log 2$.

Theorem 4 (Ratings)**.** *(a) For every* $p, q \in \mathcal{P}(A)$ *and every essential agent* i*,* $(Ap)_i = (Aq)_i = 0$*. Hence the ratings of essential agents are zero under every selection rule, before and after duplication. The maximum-entropy equilibrium also gives positive probability to every essential agent. (b) Let* $p^\dagger = \pi(p^\star(\mathrm{dup}_k\, A))$*. By Theorems 1 and 3 it is the tilted maximizer at* $\lambda = \log 2$*. The rating of an original agent* i *after duplication is* $(Ap^\dagger)_i$*, which is unchanged if and only if row* i *of* A *is orthogonal to the shift* $p^\dagger - p^\star$*, that is* $A_i\,(p^\dagger - p^\star) = 0$*. (c) In* A_5*,* $p^\star(A_5)$ *is uniform on the four essential agents with* $r_4 = -3/4$*, and duplicating agent* 5 *moves* r_4 *to* $-3/5$*.*

Under any selection rule, a copy can move only the ratings of inessential agents, in proportion to how their results weight the shift of probability. The first claim of part (a) is the classical property of essential agents, which earn the value of the game against every equilibrium (Bohnenblust et al., 1950; Gale & Sherman, 1950); here the value is zero. Theorem 4(a) adds that a copy cannot move those ratings either, under any selection rule. In A_5 the shift is $(-1/20, -1/20, -1/20, 0, 3/20)$, which leaves the ratings of agents 1, 2, 3 and 5 at zero and raises agent 4's rating by $3/20$.

A copy can also reverse the order of two inessential agents. Add to A_5 a sixth agent that draws agent 4 and scores -0.7 against every other agent. The sixth agent is inessential with rating -0.7 before and after the copy of agent 5, while agent 4 moves from $-3/4$ to $-3/5$ and passes it.

Lemma 1 and Theorems 1 to 3 also hold when a task is duplicated in the original paper's evaluation of agents on tasks (Appendix A.6). There an equilibrium is a pair of distributions, one over the agents and one over the tasks, and the tilt acts on the task distribution alone.

3.4 Near copies and exact draws

In real data an entrant that imitates another agent differs from it slightly, and we therefore ask whether a near copy behaves like an exact one.

Proposition 1 (Near copies)**.** *Let* A_η *be any antisymmetric* $(n+1) \times (n+1)$ *matrix whose leading* $n \times n$ *block is* A *and whose appended row agrees with row* k *of* A *to within* η *in each of its first* n *entries. Then the following holds as* $\eta \to 0$ *for this fixed* A*, uniformly over such matrices and over all selection rules. If* p_k *is the same in every equilibrium of* A*, the total probability of agent* k *and its near copy converges to that value. If the rating of agent* i *is the same in every equilibrium of* A*, which holds for every essential* i *by Theorem 4, the rating of agent* i *converges to that value. In particular, if the equilibrium of* A *is unique, every total probability and every rating converges to its value under* A*.*

Any limit of equilibria of the near-copy games is an equilibrium of the duplicated game. A quantity that is the same across those equilibria is therefore forced to its value, while a quantity that varies can end at different points of its range. In A_5 both vary, and the limits depend on the near copy. A near copy of agent 5 that scores $+\eta$ against everyone else drives the total probability of agent 5 to 1 and the rating of agent 4 to within η of 0 under every selection rule. A near copy that scores $-\eta$ receives no probability in any equilibrium, so the maximum-entropy probabilities and ratings of A_5 are unchanged, while the exact copy gives $2/5$ and $-3/5$.

The same variation across equilibria causes a discontinuity at an exact draw even when nothing is copied, as Balduzzi et al. (2018) observe in an example. Their continuity result is that an unbeatable distribution over agents stays nearly unbeatable nearby, not that ratings move continuously.

Proposition 2 (Discontinuity at an exact draw)**.** *Let A_δ be the matrix A_5 with agent 5 scoring δ against each of the other four agents. For $\delta > 0$ the only equilibrium of A_δ puts all probability on agent 5, and agent 4 is rated $-\delta$. For $\delta < 0$ the only equilibrium is the uniform distribution on agents 1 to 3, and agent 4 is rated -1. At $\delta = 0$ the maximum-entropy rating of agent 4 is $-3/4$.*

The one-sided limits are the ends of the equilibrium set of A_5, where agent 4 is rated 0 and -1. The maximum-entropy value $-3/4$ at the draw is neither limit. No point chosen at $\delta = 0$ can match both (Figure 1, right). The equilibrium is unique on both sides, so no choice of rule is involved: the rating a rule must report jumps with the data. A rule that never jumps must be wrong somewhere instead.

Theorem 5 (The price of continuity)**.** *Let $\rho_i(A)$ be the range of $(Ap)_i$ over $\mathcal{P}(A)$, and let g_b rate the agents at the entropy maximizer over the relaxed set $\mathcal{P}_b(A) = \{p \in \Delta_n : Ap \preceq b\mathbf{1}\}$ for a slack $b > 0$, so that $b = 0$ would give Nash averaging itself. (a) For every rule continuous at A, its largest difference from agent i's equilibrium rating over matrices with a unique equilibrium within δ of A, as $\delta \downarrow 0$, is at least $\rho_i(A)/2$. (b) g_b is continuous everywhere, and that largest difference for g_b is at most the range $\rho_i^b(A)$ of $(Ap)_i$ over $\mathcal{P}_b(A)$, which falls to $\rho_i(A)$ as b does.*

Where $\rho_i(A) > 0$ the two bounds meet within a factor of two as the slack falls to zero. The slack is not a repair: g_b moves ratings under an exact copy even where the equilibrium is unique, and no single slack serves every matrix (Appendix E).

4 Where more than one equilibrium occurs

4.1 Data and measurements

Section 3 proves its results for every antisymmetric outcome matrix. We measure how often more than one equilibrium occurs and how much a copy we add, or the choice of selection rule, then changes the ratings. More than one equilibrium needs an exact coincidence in the data (Section 2).

Proposition 3 (Two sources of more than one equilibrium)**.** *(a) If two essential agents have the same results against everyone else and draw each other, the equilibrium is not unique. If the two agents are inessential instead, they add no width, although the equilibrium can fail to be unique for other reasons. (b) If some agent draws against everyone and $n \geq 2$, the equilibrium is not unique.*

Neither source is necessary, as a 5×5 matrix with no draw and no repeated row shows (Appendix C).

We plant these two coincidences in two of three generated families of 2000 matrices each, with the number of agents uniform on 4 to 8. We build the games with planted ties from a difference matrix $A_{ij} = s_i - s_j$ with standard normal skills and one random pair of them set equal, or from one cyclic triple or two that draw each other, beside agents that draw everyone. For the games with dominated agents we plant ties of the same kinds and add agents that a cyclic triple beats, or replace the weakest agent by one that draws everyone. Gaussian games are $Z - Z^\top$ with independent standard normal entries of Z.

The real data are an engine league of Stockfish settings with a second, independent run of the strongest one, the Atari score table published with deviation ratings (Marris et al., 2025), and several arena tables of the LMSYS chatbot arena. The chatbot-arena leaderboard is the 100k human-preference release reduced to 28 models with at least 30 anonymous battles per pair, a tie vote counting one half, and four further arena tables are built the same way from the search, web-development, repository-chat and 140k releases. The expert arena has too few battles for that, so we use it at 5 per

pair. Finally we use the 44 complete seasons from 2015-16 to 2023-24 of the top football divisions of England, Germany, Spain, Italy and France. The abandoned French 2019-20 season is left out.

For every matrix we compute the width of the equilibrium set by linear programming, and on the generated games, the football seasons and the arena pools that are not decided by a strict winner we confirm the verdict by exact rational vertex enumeration. We then compute the maximum-entropy equilibrium before and after a duplication, and the LP vertex with the copy listed last. We duplicate in turn each generated agent, each essential football team, each essential arena model and each Atari agent and task; by Theorem 2 a copy of an inessential agent changes nothing. In the generated games we call one duplication in a game that has more than one equilibrium a test; there are 8708 such tests with planted ties and 10 460 with dominated agents.

Appendices B to D give the generators, the solvers, the hand-built instances, the perturbation studies, and the tolerances and thresholds at which we count a positive width, a move of the equilibrium, a change of rating or a reversal. The counts of generated games with more than one equilibrium are stable from 10^{-5} to 10^{-3}. Appendices F to I document the engine league, the Atari table, the arena tables and the football tables. Appendix J gives the other rating rules, Appendix D.4 the matrices where an equilibrium rating is continuous, and Appendices K to M further related work, a study of the width at the resolution of the data, and the records. Appendices H.7 and J.8 perturb the copy on the engine league and the leaderboard.

4.2 Generated games behave as the theory predicts

In the generated games more than one equilibrium appears only where ties are planted, and none of the 2000 Gaussian games has it, as the classical result of Section 2 predicts. Theorem 2 predicts which tests move: its linear program matches the observed movement on 19 164 of the 19 168 tests, the four exceptions having gaps below the solver tolerance. As Theorems 3 and 4 require, the copied agent's total probability never falls beyond solver noise and essential agents stay at rating zero.

Duplicating an agent moves the maximum-entropy equilibrium by more than 10^{-4} in 94.7% of the tests with planted ties and in 69.9% of those with dominated agents. On the difference matrices among the games with planted ties every equilibrium gives every agent the same rating (Balduzzi et al., 2018), so a copy moves the equilibrium but leaves every rating where it was. No rating moves anywhere in the games with planted ties. In the games with dominated agents a copy changes some inessential rating by more than 10^{-4} in 63.9% of the 10 460 tests. That is 0.56 of all 11 965 duplications (Table 2). Over all 10 460 tests the 90th percentile of that change is 0.19. Those changes never reverse the order of two inessential agents (0 of the 7305 such tests whose copied agent's probability varies across equilibria; the other four of the 7309 the solver counts as moved are the tolerance cases above), although Section 3 shows that a reversal is possible.

In these games the choice of selection rule reorders agents more often than a copy does: the LP vertex and the maximum-entropy equilibrium differ on every game with more than one equilibrium. On the 714 games with dominated agents that have more than one equilibrium and at least two inessential agents, they order those agents differently in 38% of the games under the LP vertex (zero objective) and in 65% under the minimax one. The discontinuity at an exact draw also affects estimated ratings, as Proposition 2 predicts: sampling A_5 breaks its draw, and most estimates of agent 4 then fall near one end of the equilibrium set rather than near $-3/4$ (Appendix D).

4.3 Real data: a unique equilibrium on the main leaderboards, several in most football seasons

Table 1 lists the real data, with the Gaussian games for comparison. The Atari score table, the chatbot-arena leaderboard and the four further arena tables each have a unique equilibrium, and on the arena tables it puts all probability on one model that beats every other.

The engine league holds the one copy in our data that played its own games. The strongest setting and its rerun are deterministic, and their head-to-head is therefore an exact draw by construction. In the estimated matrix they also beat everyone else, which makes the equilibrium set the segment between them. The maximum-entropy equilibrium splits the pair evenly and the LP vertex puts all probability on one of them. They rate the other settings up to $1/30$ apart, which is half the largest

Table 1: Copies and the choice of rule change ratings only where the equilibrium is not unique. Largest values on each data set's own scale; the zeros in the unique rows follow from uniqueness and Theorem 2, and a dash marks the row run under maximum entropy alone. *half the gap between the two recorded rows, 10^{-8} with them equal.

data	size	equilibrium	rating gap, $p^\star$ vs LP	change of other ratings by a copy	copies reversing two agents
Gaussian games	2000, 4–8 agents	unique in all	0	0	none
engine league	16 settings, 300 games per pair	segment on the copy pair	0.033*	0.011	none
Atari score table	20 agents, 53 tasks	unique	0	0	none
chatbot-arena leaderboard and four further arena tables	28 models; 4–44 models	unique	0	0	none
expert arena	6 models, ≥ 5 battles per pair	segment	0.071	0.071	1 of 2
subsampled leaderboard	5–50 battles per pair, 50 replicates each	not unique in 14%–34%	–	0.06–0.12	25%–35%
football seasons	44 seasons, 5 leagues	not unique in 40	1.0	0.275	41 of 153, in 17 seasons

Table 2: Every rule that picks an equilibrium jumps at an exact draw; Bradley–Terry and α-Rank do not jump but change ratings under a copy; Nash averaging does both of these. Fractions of 11 965 generated duplications (10 460 in games with more than one equilibrium, as in Section 4.2) and counts of 153 football duplications (765 relabeled). Last column: the rule's own scale.

rule	of 11 965: changed	of 11 965: reversed	football: changed	football: reversed	agent 4 of A_δ at $\delta = -10^{-3}$ / $+10^{-3}$
Nash averaging (maximum entropy)	0.56	0.00	74	41	-1 / -0.001
LP vertex (zero objective), copy listed last	0.00	0.00	19	19	-1 / -0.001
LP vertex (zero objective), relabeled	0.02	0.00	205	202	-1 / -0.001
deviation ratings	0.00	0.00	0	0	-1 / -0.001
exact copies removed, then maximum entropy	0.00	0.00	0	0	-1 / -0.001
Bradley–Terry (log-strength)	0.96	0.06	153	115	-0.5792 / -0.5796
α-Rank (mass of agent 4)	0.95	0.41	150	124	0.0004 / 0.0011

difference that sampling leaves between the two recorded rows and vanishes when the rows are set equal.

Tables with few battles per pair can have more than one equilibrium. In the expert arena, where two models split their 11 battles evenly because a tie vote counts one half, the equilibrium set is a segment, and one of the two possible copies changes other ratings by up to $1/14$ and reverses two models. Subsampling the leaderboard to 5 to 50 battles per pair leaves more than one equilibrium in 14% to 34% of the replicates. On those replicates a copy reverses two models in 25% to 35% of the duplications, but its change exceeds the standard error of the affected model in at most 2%. On a second pool drawn from the 140k release the share exceeding the standard error reaches 7.3%.

Football seasons are the clearest real case. Each pair of teams meets twice, a win, draw and loss score 1, $1/2$ and 0, and the entry $A_{ij} = 2W_{ij} - 1$ uses the mean score W_{ij} of team i over its two games against j. A series split evenly therefore gives an exact zero entry. In 40 of the 44 seasons the equilibrium is not unique, and in 21 of them every pair of essential teams split its series, which makes their submatrix zero and the rank condition of Section 2 hold. In the other 19 the submatrix loses rank despite nonzero entries, a coincidence common in $\{0, \pm 1/2, \pm 1\}$ data. Duplicating every essential team of each of the 40 seasons gives 153 duplications. In 26 of the 40 seasons the maximum-entropy equilibrium and the LP vertex rate some other team more than 0.01 apart. A copy changes another team's rating by up to 0.275, and by more than 0.01 in 74 of the 153 duplications, which fall in 28 seasons. It reverses the order of two other teams in 41 of the 153 duplications, which fall in 17 seasons, and pooled over all 153 the reversed pairs are 0.7% of the pairs strictly ordered before the copy. These are complete records, and the draws are in the published table, so every rating computed from that table inherits them. Two other ways of writing the entries also give the same picture (Appendix I).

4.4 No other rating rule we test both ignores copies and stays continuous

We compare Nash averaging with the LP vertex (zero objective, copy listed last), deviation ratings (Marris et al., 2025), Bradley–Terry (Bradley & Terry, 1952), α-Rank (Omidshafiei et al., 2019) and removing exact copies before solving (Marris et al., 2022). We apply them to all 11 965 duplications of the games with dominated agents (unique games included), the 153 football duplications and the draw of Proposition 2. Deviation ratings reduce to a selection rule on a symmetric zero-sum game (Lemma 5 in Appendix A.9), so Proposition 2 applies to them.

Table 2 shows the trade-off: every selection rule jumps across the draw, as Proposition 2 requires, while Bradley–Terry and α-Rank change by less than 10^{-3}. Those two rules, however, change some rating on 96% and 95% of the generated duplications and reverse pairs throughout the order, beyond the inessential agents to which Theorem 4 confines Nash averaging, since both rules count a copied agent twice. Deviation ratings and removing exact copies change no rating by more than the table's thresholds, 10^{-4} on the generated duplications and 0.01 on football. With the copy listed last the LP vertex ignores exact copies on the generated duplications, but changes a rating by up to 1.0 in 19 of 153 football duplications and 205 of 765 relabeled ones. On the engine league a copy perturbed by 10^{-3} moves deviation ratings by 0.068, nearly two standard errors, while removing near copies within a threshold moves nothing but jumps at the threshold itself (Appendix J).

5 Related work

Nash averaging and its successors. The original paper (Balduzzi et al., 2018) motivates the maximum-entropy equilibrium by invariance to redundant agents and tasks, and proves Lemma 1's first claim without showing in general that this equilibrium stays in place. The N-player generalization (Marris et al., 2022) led to deviation ratings (Marris et al., 2025) and to a three-player open-ended evaluation game (Liu et al., 2025). Those works cite the invariance without a further proof, as do a social-choice treatment of agent evaluation (Lanctot et al., 2023) and a study of equilibrium selection by learning dynamics (Leal, 2026). Marris et al. (2025) argue that deviation ratings agree with Nash averaging up to an offset; the agreement covers only the agents that the equilibrium plays, where Theorem 4 also finds no change. Marris et al. (2022) also propose removing exact copies before solving and a weighted entropy; Appendix E derives what a threshold version of the first costs.

Clones in social choice. The social-choice treatment of agent evaluation (Lanctot et al., 2023) credits Nash averaging with independence of clones (Tideman, 1987), citing the original paper. For the maximum-entropy equilibrium that credit holds for the ratings of essential agents (Theorem 4) and fails for the order of inessential agents (Section 4.3). Maximal lotteries (Fishburn, 1984) are instead the equilibrium set, whose invariance to clones needs no condition (Brandl et al., 2016; Khalaf et al., 2026). Composition consistency for tournaments (Laffond et al., 1993; 1996) does not cover copies, since a copy draws with its original. A metric-space axiomatization of clone-robust weights credits Nash averaging with invariance to exact copies only (Berriaud & Wattenhofer, 2026a), and its sequel meets the same tie-breaking difficulty (Berriaud & Wattenhofer, 2026b). These works treat set-valued solutions or weights, while we describe what the chosen equilibrium does to the ratings.

6 Conclusion

A copy leaves the maximum-entropy equilibrium in place exactly when no step inside the equilibrium set gains more from the tilt than it loses in entropy (Theorem 2). Every unique equilibrium passes, including those of the Atari score table and the chatbot-arena leaderboard. Otherwise a copy can move and reorder only inessential ratings (Theorems 3 and 4). Coarse scores leave more than one equilibrium in 40 of 44 complete football seasons and in some replicates of the subsampled leaderboard. Every rating rule we test jumps at an exact draw or moves ratings under a copy, and Nash averaging does both (Table 2). Its ratings jump only where they differ between equilibria. There, a continuous rule is wrong somewhere nearby by at least half the rating's range (Theorem 5). A published Nash average should name its selection rule and report whether the equilibrium is unique.

REPRODUCIBILITY STATEMENT

Every proof is in Appendix A. Appendix B specifies the synthetic families, seeds, solvers, thresholds and the exact-arithmetic certificate. Appendix F specifies the engine league (engine versions, openings, bootstrap) and Appendix H the leaderboard construction. The supplementary material holds the code for every experiment, the league score matrix and game records, the filtered leaderboard matrix, the four further arena matrices and the expert-arena matrix, the 44 football season matrices, the transcribed Atari tables and the result files behind every number.

DISCLOSURE OF GENERATIVE AI USE

The authors set the research question, the instruments (synthetic families, an engine league and public tables) and the claims to be tested. The protocol required a success and a stop condition recorded before each experiment, an independent second derivation of every proof, and a check of every literature statement against the cited page. Within those rules, generative AI tools developed the framework of split equilibria and the tilt, formulated and proved the mathematical claims, designed and implemented the experiments and solvers, interpreted the results, and drafted and revised the text and figures. No dataset was generated by a model. The authors verified the proofs through the second derivation and traced every number to its result file. They take full responsibility for every claim, number and sentence.

REFERENCES

David Balduzzi, Karl Tuyls, Julien Perolat, and Thore Graepel. Re-evaluating evaluation. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

Ratip Emin Berker, Sílvia Casacuberta, Isaac Robinson, Christopher Ong, Vincent Conitzer, and Edith Elkind. From independence of clones to composition consistency: a hierarchy of barriers to strategic nomination. In *Proceedings of the 26th ACM Conference on Economics and Computation (EC '25)*, pp. 1109, 2025. doi: 10.1145/3736252.3742673. Full version: arXiv:2502.16973.

Damien Berriaud and Roger Wattenhofer. Clone-robust weights in metric spaces: Handling redundancy bias in benchmark aggregation. In *Proceedings of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, 2026a. doi: 10.65109/MJWJ6521. arXiv:2502.03576.

Damien Berriaud and Roger Wattenhofer. Breaking the illusion of artificial consensus: Clone-robust weighting for arbitrary metric spaces. *arXiv preprint arXiv:2602.24024*, 2026b.

Henri F. Bohnenblust, Samuel Karlin, and Lloyd S. Shapley. Solutions of discrete, two-person games. In Harold W. Kuhn and Albert W. Tucker (eds.), *Contributions to the Theory of Games, Volume I*, number 24 in Annals of Mathematics Studies, pp. 51–72. Princeton University Press, 1950.

Ralph Allan Bradley and Milton E. Terry. Rank analysis of incomplete block designs: I. The method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952.

Florian Brandl. The distribution of optimal strategies in symmetric zero-sum games. *Games and Economic Behavior*, 104:674–680, 2017.

Florian Brandl and Felix Brandt. Consistent probabilistic social choice revisited. *arXiv preprint arXiv:2606.10998*, 2026.

Florian Brandl and Dominik Peters. Approval voting under dichotomous preferences: A catalogue of characterizations. *Journal of Economic Theory*, 205:105532, 2022. doi: 10.1016/j.jet.2022.105532.

Florian Brandl, Felix Brandt, and Hans Georg Seedig. Consistent probabilistic social choice. *Econometrica*, 84(5):1839–1880, 2016. doi: 10.3982/ECTA13337.

Wojciech M. Czarnecki, Gauthier Gidel, Brendan Tracey, Karl Tuyls, Shayegan Omidshafiei, David Balduzzi, and Max Jaderberg. Real world games look like spinning tops. In *Advances in Neural Information Processing Systems*, volume 33, pp. 17443–17454, 2020. arXiv:2004.09468.

Peter C. Fishburn. Probabilistic social choice based on simple voting comparisons. *The Review of Economic Studies*, 51(4):683–692, 1984.

David C. Fisher and Jennifer Ryan. Optimal strategies for a generalized "scissors, paper, and stone" game. *The American Mathematical Monthly*, 99(10):935–942, 1992.

David Gale and Seymour Sherman. Solutions of finite two-person games. In *Contributions to the Theory of Games, Volume I*, number 24 in Annals of Mathematics Studies, pp. 37–49. Princeton University Press, 1950.

David Gale, Harold W. Kuhn, and Albert W. Tucker. On symmetric games. In *Contributions to the Theory of Games, Volume I*, number 24 in Annals of Mathematics Studies, pp. 81–87. Princeton University Press, 1950.

A. J. Goldman and A. W. Tucker. Theory of linear programming. In H. W. Kuhn and A. W. Tucker (eds.), *Linear Inequalities and Related Systems*, number 38 in Annals of Mathematics Studies, pp. 53–98. Princeton University Press, 1956.

John C. Harsanyi. Oddness of the number of equilibrium points: a new proof. *International Journal of Game Theory*, 2(1):235–250, 1973. doi: 10.1007/BF01737572.

M. J. M. Jansen. Regularity and stability of equilibrium points of bimatrix games. *Mathematics of Operations Research*, 6(4):530–550, 1981. doi: 10.1287/moor.6.4.530.

Irving Kaplansky. A contribution to von Neumann's theory of games. *Annals of Mathematics*, 46 (3):474–479, 1945. doi: 10.2307/1969164.

Irving Kaplansky. A contribution to von Neumann's theory of games. II. *Linear Algebra and its Applications*, 226–228:371–373, 1995. doi: 10.1016/0024-3795(95)00167-P.

Hadi Khalaf, Serena L. Wang, Daniel Halpern, Itai Shapira, Flavio du Pin Calmon, and Ariel D. Procaccia. Robust AI evaluation through maximal lotteries. *arXiv preprint arXiv:2602.21297*, 2026.

Gilbert Laffond, Jean-François Laslier, and Michel Le Breton. The bipartisan set of a tournament game. *Games and Economic Behavior*, 5(1):182–201, 1993. doi: 10.1006/game.1993.1010.

Gilbert Laffond, Jean Lainé, and Jean-François Laslier. Composition-consistent tournament solutions and social choice functions. *Social Choice and Welfare*, 13(1):75–93, 1996.

Gilbert Laffond, Jean-François Laslier, and Michel Le Breton. A theorem on symmetric two-player zero-sum games. *Journal of Economic Theory*, 72(2):426–431, 1997.

Marc Lanctot, Edward Lockhart, Jean-Baptiste Lespiau, Vinicius Zambaldi, Satyaki Upadhyay, Julien Pérolat, Sriram Srinivasan, Finbarr Timbers, Karl Tuyls, Shayegan Omidshafiei, Daniel Hennes, Dustin Morrill, Paul Muller, Timo Ewalds, Ryan Faulkner, János Kramár, Bart De Vylder, Brennan Saeta, James Bradbury, David Ding, Sebastian Borgeaud, Matthew Lai, Julian Schrittwieser, Thomas Anthony, Edward Hughes, Ivo Danihelka, and Jonah Ryan-Davis. OpenSpiel: A framework for reinforcement learning in games. *arXiv preprint arXiv:1908.09453*, 2019.

Marc Lanctot, Kate Larson, Yoram Bachrach, Luke Marris, Zun Li, Avishkar Bhoopchand, Thomas Anthony, Brian Tanner, and Anna Koop. Evaluating agents using social choice theory. *arXiv preprint arXiv:2312.03121*, 2023.

Michel Le Breton. On the uniqueness of equilibrium in symmetric two-player zero-sum games with integer payoffs. *Économie publique*, 17:187–195, 2005. doi: 10.4000/economiepublique.2655.

Luis Leal. Which Nash equilibrium? Solver-dependent selection on zero-sum Nash polytopes. *arXiv preprint arXiv:2606.28308*, 2026.

Siqi Liu, Ian Gemp, Luke Marris, Georgios Piliouras, Nicolas Heess, and Marc Lanctot. Re-evaluating open-ended evaluation of large language models. In *International Conference on Learning Representations*, 2025.

Luke Marris, Marc Lanctot, Ian Gemp, Shayegan Omidshafiei, Stephen McAleer, Jerome Connor, Karl Tuyls, and Thore Graepel. Game theoretic rating in N-player general-sum games with equilibria. *arXiv preprint arXiv:2210.02205*, 2022.

Luke Marris, Siqi Liu, Ian Gemp, Georgios Piliouras, and Marc Lanctot. Deviation ratings: A general, clone-invariant rating method. *arXiv preprint arXiv:2502.11645*, 2025.

Quinn McNemar. Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika*, 12(2):153–157, 1947. doi: 10.1007/BF02295996.

Henk Norde. Bimatrix games have quasi-strict equilibria. *Mathematical Programming*, 85(1):35–49, 1999. doi: 10.1007/s101070050045.

Shayegan Omidshafiei, Christos Papadimitriou, Georgios Piliouras, Karl Tuyls, Mark Rowland, Jean-Baptiste Lespiau, Wojciech M. Czarnecki, Marc Lanctot, Julien Perolat, and Remi Munos. α-Rank: Multi-agent evaluation by evolution. *Scientific Reports*, 9:9937, 2019.

R. Tyrrell Rockafellar. *Convex Analysis*. Number 28 in Princeton Mathematical Series. Princeton University Press, 1970.

T. Nicolaus Tideman. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4(3):185–206, 1987. doi: 10.1007/BF00433944.

A PROOFS

What is where. Appendix A proves Proposition 3 and every result of Section 3 except Theorem 5, which is proved in Appendices D.4 and E. It begins with the tilt formula (Appendix A.1), the criterion (Appendix A.2), the linear program that checks it (Appendix A.3) and the direction of the move (Appendix A.4). Then come the ratings, the sources of width, the exact draw and near copies (Appendix A.5), the agent-versus-task version (Appendix A.6), the classical coordinates (Appendix A.7), the classical background (Appendix A.8) and deviation ratings (Appendix A.9). Appendices B to J hold the experimental material behind Section 4, in the order of the body. Appendix B gives the generators, the solvers and the provenance of every count (Table 5); Appendix C the hand-built instances and the error-aware verdict rule (Tables 6 and 7); Appendix D the perturbation and sampling studies behind the exact draw (Figure 2). Appendices F, G, H and I hold the real data: the engine league (Table 8), the published Atari equilibrium, the leaderboard, the further arena tables and the copy effects on the subsampled pools (Tables 9, 10 and 11), and the football seasons (Tables 12 and 13). Appendix J compares the competing rating rules (Tables 14 and 15). Appendix K gives further related work on clone axioms, Appendix L reports the calibration study behind the recommendation of Section 6 (Table 16), and Appendix M the seeds, the software and the compute.

Throughout, inequalities between vectors are coordinatewise, $\mathbf{e}_k$ is the k-th unit vector, A_i denotes row i of A, and coordinates of a vector x are written x_i. We write $\mathcal{E}(A)$ for the essential support of A, the agents that receive positive probability in some equilibrium (Section 2), and $A_{\mathcal{E}}$ for its submatrix. The maps ι_α and π are the split and merge maps of Section 3.1, and a feasible step at a point, in the sense of Section 3, is also called a feasible direction here. Each result of Section 3 and Proposition 3 of Section 4.1 is restated before its proof. The supplement of Balduzzi et al. (2018) prints the equilibrium condition as $Ap \succeq 0$. Under its own payoff convention the correct condition is $Ap \preceq 0$: in the two-agent game with $A_{12} = 1$, $A\mathbf{e}_1 = (0, -1)^\top \preceq 0$ puts the winner in equilibrium and $A\mathbf{e}_2 = (1, 0)^\top$ does not. Every matrix in this paper uses $Ap \preceq 0$.

A.1 EQUILIBRIA OF THE DUPLICATED GAME AND THE TILT FORMULA

Restatement (Lemma 1, equilibria of the duplicated game)**.** *For every antisymmetric A and every k, $\mathcal{P}(\mathrm{dup}_k A) = \{\iota_\alpha(p) : p \in \mathcal{P}(A),\ \alpha \in [0, 1]\}$. Adding the copy's probability back sends*

$\mathcal{P}(\mathrm{dup}_k\, A)$ *onto* $\mathcal{P}(A)$*, the equilibria that add up to* p *form the segment* $\{\iota_\alpha(p) : \alpha \in [0,1]\}$*, which degenerates to a point when* $p_k = 0$*, and* $(\mathrm{dup}_k\, A)\, p' = \big(A\,\pi(p'),\, (A\,\pi(p'))_k\big)$ *for every* $p' \in \Delta_{n+1}$*.*

Proof. Agent k and its copy score the same against everyone and draw each other, so no agent's expected result depends on how their probability is divided between them. Write $A' = \mathrm{dup}_k(A)$, let $p' \in \Delta_{n+1}$, and set $p = \pi(p')$, so that $p_j = p'_j$ for $j \neq k$ and $p_k = p'_k + p'_{n+1}$. For an original row $i \leq n$ the appended column of A' equals column k of A in rows $i \leq n$, so

$$\begin{aligned}(A'p')_i &= \sum_{j\leq n,\, j\neq k} A_{ij}p'_j + A_{ik}p'_k + A_{ik}p'_{n+1} \\ &= \sum_{j\leq n,\, j\neq k} A_{ij}p_j + A_{ik}(p'_k + p'_{n+1}) \\ &= (Ap)_i. \end{aligned} \tag{4}$$

The appended row has the entries of row k against the original agents, and its entries against agent k and against itself are $A'_{n+1,k} = A'_{n+1,n+1} = A_{kk} = 0$, so

$$\begin{aligned}(A'p')_{n+1} &= \sum_{j\leq n,\, j\neq k} A_{kj}p'_j + 0\cdot p'_k + 0\cdot p'_{n+1} \\ &= \sum_{j\leq n,\, j\neq k} A_{kj}p_j + A_{kk}p_k = (Ap)_k, \end{aligned} \tag{5}$$

where the second equality uses $A_{kk} = 0$. Together (4) and (5) give $A'p' = (Ap, (Ap)_k)$, the rating identity, and $(A'p')_k = (Ap)_k$ as well. Hence $A'p' \preceq 0$ if and only if $Ap \preceq 0$, that is, $p' \in \mathcal{P}(A')$ if and only if $\pi(p') \in \mathcal{P}(A)$. Conversely, for $p \in \mathcal{P}(A)$ and $\alpha \in [0,1]$ we have $\pi(\iota_\alpha(p)) = p$, which places $\iota_\alpha(p)$ in $\mathcal{P}(A')$. Every p' with $\pi(p') = p$ is $\iota_\alpha(p)$ with $\alpha = p'_k/p_k \in [0,1]$ when $p_k > 0$ (since $0 \leq p'_k \leq p'_k + p'_{n+1} = p_k$). When $p_k = 0$ the same identity forces $p'_k = p'_{n+1} = 0$ and only one point of $\mathcal{P}(A')$ adds up to p. □

Restatement (Lemma 2, entropy of a split)**.** $H(\iota_\alpha(p)) = H(p) + p_k\, H_2(\alpha)$*, where* H_2 *is the binary entropy. For fixed* p *the maximum over* α *is attained at* $\alpha = 1/2$ *with value* $H(p) + p_k \log 2$*. Splitting* p_k *across* m *entries in the proportions* $\alpha \in \Delta_m$ *gives* $H(p) + p_k H(\alpha)$ *by the same computation, with* H_2 *the case* $m = 2$*.*

Proof. Only the two split coordinates differ between $\iota_\alpha(p)$ and p, so

$$\begin{aligned}H(\iota_\alpha(p)) - H(p) &= -\alpha p_k \log(\alpha p_k) - (1-\alpha)p_k \log\big((1-\alpha)p_k\big) + p_k \log p_k \\ &= -\alpha p_k \log\alpha - \alpha p_k \log p_k \\ &\qquad - (1-\alpha)p_k \log(1-\alpha) - (1-\alpha)p_k \log p_k + p_k \log p_k \\ &= p_k\big(-\alpha\log\alpha - (1-\alpha)\log(1-\alpha)\big) = p_k\, H_2(\alpha), \end{aligned} \tag{6}$$

where the second equality expands the logarithms of products and the third cancels the three $p_k \log p_k$ terms. H_2 is strictly concave on $[0,1]$ with maximum $\log 2$ at $\alpha = 1/2$, which gives the maximum $H(p) + p_k \log 2$, uniquely when $p_k > 0$. □

Restatement (Theorem 1, tilt formula)**.** $p^\star(\mathrm{dup}_k\, A) = \iota_{1/2}\big(\arg\max_{p\in\mathcal{P}(A)} [H(p) + p_k \log 2]\big)$*, and the maximizer over* $\mathcal{P}(A)$ *is unique;* m *copies in total that draw one another, the original counted among them, replace* $\log 2$ *by* $\log m$ *and the even split by an* m*-way split, and replacing each agent* k *of a set* K *by* m_k *copies tilts by* $\sum_{k\in K} p_k \log m_k$*.*

Proof. By Lemma 1, maximizing H over $\mathcal{P}(\mathrm{dup}_k\, A)$ is maximizing over pairs (p, α) with $p \in \mathcal{P}(A)$ and $\alpha \in [0,1]$, and by Lemma 2

$$\begin{aligned}\max_{p'\in\mathcal{P}(\mathrm{dup}_k\, A)} H(p') &= \max_{p\in\mathcal{P}(A)} \max_{\alpha\in[0,1]} \big[H(p) + p_k H_2(\alpha)\big] \\ &= \max_{p\in\mathcal{P}(A)} \big[H(p) + p_k \log 2\big], \end{aligned} \tag{7}$$

the inner maximum being attained at $\alpha = 1/2$ (uniquely when $p_k > 0$; when $p_k = 0$ every α gives the same point). The objective $H(p) + p_k \log 2$ is strictly concave on the nonempty compact convex set $\mathcal{P}(A)$ (a strictly concave function plus a linear one), so its maximizer $p^\dagger$ is unique, and the maximizer of H over $\mathcal{P}(\mathrm{dup}_k A)$ is $\iota_{1/2}(p^\dagger)$. For m copies in total, induct: writing $B = \mathrm{dup}_{n+1}(A')$ for the twice-duplicated matrix, the second duplication appends a row equal to row $n+1$ of A', so $B_{n+2,k} = B_{n+1,k} = B_{n+2,n+1} = 0$ and the copies draw one another, which lets (4) and (5) apply to A' in place of A. The equilibria are therefore the splittings of p_k into m nonnegative parts. Writing such a splitting as $p_k\alpha$ with $\alpha \in \Delta_m$, entropy on these is $H(p) + p_k H(\alpha) \le H(p) + p_k \log m$, with equality at the uniform split. For a set K of agents the splits of distinct agents occupy disjoint coordinates, entropy is additive over them, and adding each agent's copies back leaves the other constraints unchanged. The tilts add. □

A.2 The criterion for the equilibrium to stay in place

Lemma 3 (Support and directional derivative of the maximum-entropy equilibrium)**.** *Let $p^\star = p^\star(A)$. (a) $\mathrm{supp}(p^\star) = \mathcal{E}(A)$. (b) H is differentiable at $p^\star$ along every direction v feasible at $p^\star$ in $\mathcal{P}(A)$, with $D_v H(p^\star) = -\sum_{i\in\mathcal{E}(A)}(1 + \log p_i^\star)\, v_i$, which is linear in v. Maximality gives $D_v H(p^\star) \le 0$ for every feasible v, with equality whenever $-v$ is feasible as well.*

Proof. (a) Suppose $i \in \mathcal{E}(A)$ but $p_i^\star = 0$. Pick $q \in \mathcal{P}(A)$ with $q_i > 0$ and set $v = q - p^\star$. Then $v_i = q_i > 0$, $v_j = q_j \ge 0$ at every j with $p_j^\star = 0$, and $p^\star + tv \in \mathcal{P}(A)$ for $t \in [0, 1]$ by convexity. Let $C = \sup\{|1 + \log x| : x \in [\min_{j\in\mathrm{supp}(p^\star)} p_j^\star/2,\, 1]\}$, a bound on the absolute derivative of $x \mapsto -x \log x$ on that interval. Then

$$\begin{aligned} H(p^\star + tv) - H(p^\star) &= \sum_{j:p_j^\star=0} \big(-tv_j \log(tv_j)\big) \\ &\quad + \sum_{j\in\mathrm{supp}(p^\star)} \big(-(p_j^\star + tv_j)\log(p_j^\star + tv_j) + p_j^\star \log p_j^\star\big) \\ &\ge -tv_i \log(tv_i) - Ct \sum_{j\in\mathrm{supp}(p^\star)} |v_j| \qquad \text{for } 0 < t \le 1/2. \end{aligned} \tag{8}$$

The terms with $p_j^\star = 0$ and $j \ne i$ are nonnegative because $v_j \ge 0$ and $tv_j \le 1$. The terms in $\mathrm{supp}(p^\star)$ are bounded by the mean value theorem, since $p^\star + tv = (1-t)p^\star + tq$ gives $p_j^\star + tv_j \ge (1-t)p_j^\star \ge p_j^\star/2$ for $t \le 1/2$, which keeps the argument inside the interval on which C bounds the derivative. The right side is positive for small t because $-t\log t$ dominates every linear term, contradicting the maximality of $p^\star$.

(b) By (a), and because every $q \in \mathcal{P}(A)$ vanishes outside $\mathcal{E}(A)$, every feasible direction v at $p^\star$ is supported on $\mathcal{E}(A)$. Thus H is differentiable at $p^\star$ along v with

$$D_v H(p^\star) = -\sum_{i\in\mathcal{E}(A)} (1 + \log p_i^\star)\, v_i, \tag{9}$$

which is linear in v. Maximality of $p^\star$ gives $D_v H(p^\star) \le 0$ for every feasible v. When $-v$ is feasible as well, the two inequalities give equality. Let x lie in the relative interior. Every direction v in the affine hull then has $x \pm tv$ in the set for all small $t > 0$, so v and $-v$ are both feasible. □

Restatement (Theorem 2, criterion for the equilibrium to stay in place)**.** *Let $p^\star = p^\star(A)$. Then $p^\star(\mathrm{dup}_k A) = \iota_{1/2}(p^\star)$ holds if and only if $D_v H(p^\star) + v_k \log 2 \le 0$ for every v feasible at $p^\star$ in $\mathcal{P}(A)$. In particular the copy leaves the equilibrium in place if p_k is constant on $\mathcal{P}(A)$, and also if $p^\star$ maximizes p_k over $\mathcal{P}(A)$. If $p^\star \in \mathrm{relint}\,\mathcal{P}(A)$, it leaves the equilibrium in place only if p_k is constant on $\mathcal{P}(A)$.*

Proof. Entropy has zero slope at $p^\star$ along every step that can be reversed inside $\mathcal{P}(A)$, so the tilt $p_k \log 2$ decides the question. Write $G(p) = H(p) + p_k \log 2$. By Theorem 1, $p^\star(\mathrm{dup}_k A) = \iota_{1/2}(p^\star)$ if and only if $p^\star$ maximizes G over $\mathcal{P}(A)$. Since G is concave and $\mathcal{P}(A)$ convex, $G(q) -$

$G(p^\star) \le D_{q-p^\star} G(p^\star)$ for every $q \in \mathcal{P}(A)$, so $p^\star$ maximizes G if and only if $D_v G(p^\star) \le 0$ for every feasible v at $p^\star$, where

$$D_v G(p^\star) = D_v H(p^\star) + v_k \log 2. \tag{10}$$

If p_k is constant on $\mathcal{P}(A)$, then $v_k = 0$ for every feasible v and (10) reduces to $D_v H(p^\star) \le 0$, which holds by Lemma 3(b). If $p^\star$ maximizes p_k over $\mathcal{P}(A)$, then $v_k \le 0$ for every feasible v, so

$$D_v G(p^\star) = D_v H(p^\star) + v_k \log 2 \le D_v H(p^\star) \le 0. \tag{11}$$

If $p^\star \in \operatorname{relint} \mathcal{P}(A)$ and p_k is not constant on $\mathcal{P}(A)$, pick $q \in \mathcal{P}(A)$ with $q_k \ne p^\star_k$ and let $v = \pm(q - p^\star)$ with the sign chosen so that $v_k > 0$. Both v and $-v$ are feasible at a relative-interior point, so $D_v H(p^\star) = 0$ and

$$D_v G(p^\star) = 0 + v_k \log 2 > 0, \tag{12}$$

hence $p^\star$ does not maximize G. □

Restatement (Corollary 1, criterion when inessential agents lose strictly)**.** *If $(Ap^\star(A))_i < 0$ for every inessential i, then $p^\star(\mathrm{dup}_k A) = \iota_{1/2}(p^\star(A))$ if and only if p_k is constant on $\mathcal{P}(A)$.*

Proof. Inside the hyperplane $\mathbf{1}^\top p = 1$ the faces of $\mathcal{P}(A)$ are cut out by the constraints $p_i \ge 0$ and $(Ap)_i \le 0$. Relative to the affine hull cut out by the constraints that hold with equality throughout, a point lies in the relative interior exactly when every constraint tight at it is tight on all of $\mathcal{P}(A)$ (Rockafellar, 1970, §6); a redundant inequality tight at a relative-interior point is tight on the whole set, so redundancy does not affect this test. A constraint $p_i \ge 0$ is tight at $p^\star$ exactly when $i \notin \mathcal{E}(A)$, by Lemma 3(a), and those constraints are tight on all of $\mathcal{P}(A)$ by the definition of $\mathcal{E}(A)$. A constraint $(Ap)_i \le 0$ with $i \in \mathcal{E}(A)$ is tight on all of $\mathcal{P}(A)$ by Theorem 4(a), proved below without reference to this corollary. The only constraints that could be tight at $p^\star$ without being tight on all of $\mathcal{P}(A)$ are $(Ap)_i \le 0$ with $i \notin \mathcal{E}(A)$, which the assumption excludes. Hence $p^\star \in \operatorname{relint} \mathcal{P}(A)$, which for matrix games is the regular case of Jansen (1981) (Appendix A.8), and Theorem 2 gives the equivalence. When the equilibrium is unique, every coordinate is constant on $\mathcal{P}(A)$. □

A.3 Checking the criterion by one linear program

By (9) the directional derivative is linear in v, and the feasible directions are the nonnegative multiples of $q - p^\star$ with $q \in \mathcal{P}(A)$; hence the criterion of Theorem 2 is

$$\max_{q \in \mathcal{P}(A)} \Big[- \sum_{i \in \mathcal{E}(A)} (1 + \log p^\star_i)(q_i - p^\star_i) + (q_k - p^\star_k) \log 2 \Big] \le 0$$

$$\iff \qquad \max_{q \in \mathcal{P}(A)} w^\top q \le w^\top p^\star, \tag{13}$$

with $w_i = -\log p^\star_i$ for $i \in \mathcal{E}(A)$, $w_i = 0$ otherwise, and $\log 2$ added to w_k, the constant terms cancelling because q and $p^\star$ both vanish outside $\mathcal{E}(A)$, so $\sum_{i \in \mathcal{E}(A)} q_i = \sum_{i \in \mathcal{E}(A)} p^\star_i = 1$. No entropy problem on the duplicated game has to be solved.

A.4 A copy gains probability

Restatement (Theorem 3, a copy gains probability)**.** *Let $p^\dagger = \arg\max_{\mathcal{P}(A)} [H + \lambda p_k]$ with $\lambda > 0$. Then $(p^\dagger)_k \ge p^\star_k$, with strict inequality whenever $p^\dagger \ne p^\star$, and $\lambda \mapsto (\arg\max_{\mathcal{P}(A)} [H + \lambda p_k])_k$ is nondecreasing. Consequently duplication never lowers the total probability of agent k and its copy, and raises it whenever the maximum-entropy equilibrium moves.*

Proof. Both maximizers exist and are unique because H and $H + \lambda p_k$ are strictly concave and continuous on the nonempty compact convex set $\mathcal{P}(A)$. Comparing each maximizer with the other point,

$$H(p^\star) \ge H(p^\dagger) \qquad (p^\star \text{ maximizes } H), \tag{14}$$

$$H(p^\dagger) + \lambda (p^\dagger)_k \ge H(p^\star) + \lambda p^\star_k \qquad (p^\dagger \text{ maximizes } H + \lambda p_k), \tag{15}$$

and therefore

$$\begin{aligned} \lambda\big((p^\dagger)_k - p^\star_k\big) &\geq H(p^\star) - H(p^\dagger) && \text{(rearranging (15))} \\ &\geq 0 && \text{(by (14)).} \end{aligned} \tag{16}$$

If $p^\dagger \neq p^\star$ then (14) is strict by uniqueness of the maximum-entropy equilibrium, so the chain is strict. For two tilts $\lambda_1 < \lambda_2$ with maximizers q_1, q_2, the same two inequalities give

$$\begin{aligned} (\lambda_2 - \lambda_1)\big((q_2)_k - (q_1)_k\big) = \big[H(q_2) + \lambda_2(q_2)_k - H(q_1) - \lambda_2(q_1)_k\big] \\ + \big[H(q_1) + \lambda_1(q_1)_k - H(q_2) - \lambda_1(q_2)_k\big] \geq 0, \end{aligned} \tag{17}$$

each bracket being nonnegative by the optimality of q_2 and of q_1 for its own tilt. For duplication the consequence is Theorem 1 with $\lambda = \log 2$: the total probability of agent k and its copy after duplication is $(p^\dagger)_k$ because $\pi(\iota_{1/2}(p^\dagger)) = p^\dagger$, and the maximum-entropy equilibrium moves exactly when $p^\dagger \neq p^\star$ because $\iota_{1/2}$ is injective. □

A.5 Ratings, sources of width, exact draws and near copies

Restatement (Theorem 4, ratings)**.** *(a) For every $p, q \in \mathcal{P}(A)$ and $i \in \mathcal{E}(A)$, $(Ap)_i = (Aq)_i = 0$, ratings on the essential support are invariant under duplication for every selection rule, and $\operatorname{supp}(p^\star(A)) = \mathcal{E}(A)$. (b) With $p^\dagger = \pi(p^\star(\operatorname{dup}_k A))$, the rating of an original agent i after duplication is $(Ap^\dagger)_i$, unchanged if and only if $A_i(p^\dagger - p^\star) = 0$. (c) In A_5, $p^\star(A_5)$ is uniform on the four essential agents with $r_4 = -3/4$, and duplicating agent 5 moves r_4 to $-3/5$.*

Proof. (a) Let $p, q \in \mathcal{P}(A)$, so $Ap \preceq 0$ and $Aq \preceq 0$ by (1).

$$\begin{aligned} q^\top(Ap) &\leq 0 && (Ap \preceq 0,\, q \succeq 0), \\ q^\top(Ap) &= -p^\top(Aq) \;\geq\; 0 && \text{(antisymmetry, then } Aq \preceq 0,\, p \succeq 0). \end{aligned} \tag{18}$$

Hence $\sum_j q_j(Ap)_j = 0$ is a sum of nonpositive terms, so $(Ap)_j = 0$ for every $j \in \operatorname{supp}(q)$. Taking the union over all $q \in \mathcal{P}(A)$ gives $(Ap)_i = 0$ for all $i \in \mathcal{E}(A)$ and every $p \in \mathcal{P}(A)$. By Lemma 1 the ratings of the original agents in the duplicated game are $A\,\pi(p')$ with $\pi(p') \in \mathcal{P}(A)$. The essential support is preserved: an original agent i with $p_i > 0$ for some $p \in \mathcal{P}(A)$ has $(\iota_{1/2}p)_i > 0$ (or $(\iota_{1/2}p)_k > 0$ when $i = k$); conversely $\pi(p')_i \geq p'_i$; and the copy is essential exactly when k is. Hence the ratings of essential agents are zero before and after. That $\operatorname{supp}(p^\star(A)) = \mathcal{E}(A)$ is Lemma 3(a).

(b) is (4) taken for $i \leq n$ at $p' = p^\star(\operatorname{dup}_k A)$.

(c) Write a generic equilibrium of A_5 as $x = (x_1, \ldots, x_5)$. The three cycle constraints and their sum are

$$\begin{aligned} (A_5x)_1 = x_2 - x_3 + x_4 \leq 0, \\ (A_5x)_2 = -x_1 + x_3 + x_4 \leq 0, \\ (A_5x)_3 = x_1 - x_2 + x_4 \leq 0, \\ (A_5x)_1 + (A_5x)_2 + (A_5x)_3 = 3x_4 \leq 0, \end{aligned} \tag{19}$$

so $x_4 = 0$, and then the three constraints force $x_1 = x_2 = x_3 =: t$. The remaining constraints $(A_5x)_4 = -(x_1 + x_2 + x_3) \leq 0$ and $(A_5x)_5 = 0$ hold. Thus $\mathcal{P}(A_5) = \{(t, t, t, 0, 1 - 3t) : t \in [0, 1/3]\}$. On the segment,

$$\begin{aligned} H(t, t, t, 0, 1 - 3t) &= -3t \log t - (1 - 3t)\log(1 - 3t), \\ \frac{d}{dt} H(t, t, t, 0, 1 - 3t) &= -3\log t + 3\log(1 - 3t) = 3\log\frac{1 - 3t}{t}, \end{aligned} \tag{20}$$

which vanishes at $t = 1/4$ (the uniform distribution on four entries), and $(A_5p^\star)_4 = -3t = -3/4$. After duplicating agent 5 the equilibrium set is $\{(t, t, t, 0, \alpha(1-3t), (1-\alpha)(1-3t))\}$ by Lemma 1, and by Theorem 1 the tilted objective on the segment is

$$\begin{aligned} F(t) &= H(t, t, t, 0, 1 - 3t) + (1 - 3t)\log 2, \\ F'(t) &= 3\log\frac{1 - 3t}{t} - 3\log 2 = 3\log\frac{1 - 3t}{2t}, \end{aligned} \tag{21}$$

which vanishes at $t = 1/5$ (the uniform distribution on the five nonzero entries), so $r_4 = -3t = -3/5$. □

Restatement (Proposition 3, sources of width)**.** *(a) If two essential agents fare identically against everyone else and draw each other, $\mathcal{P}(A)$ contains more than one equilibrium. For a pair of inessential agents with identical columns and a drawn head-to-head, the draw contributes no width. (b) If some agent draws against everyone, $\mathcal{P}(A)$ contains more than one equilibrium whenever $n \geq 2$.*

Proof. (a) Two agents that draw each other and fare identically against everyone else are an agent and its copy up to relabeling, so Lemma 1 applies. Then $\mathcal{P}(A)$ is the union of the splits of the equilibria of the merged game. The splits of an equilibrium that gives the pair positive probability form a segment of positive length, and such an equilibrium exists exactly when the pair lies in the essential support. An essential pair therefore puts a segment inside $\mathcal{P}(A)$. The splits of an inessential pair are single points and add no width, whatever the equilibrium set of the merged game is.

(b) If $A_{jl} = 0$ for all l then $A\mathbf{e}_j = 0 \preceq 0$ and $\mathbf{e}_j \in \mathcal{P}(A)$. Without agent j the reduced matrix is antisymmetric, so the game it defines has value zero and an equilibrium (Gale et al., 1950); call it p. Embedded with $p_j = 0$, it satisfies $(Ap)_l \leq 0$ for $l \neq j$ and $(Ap)_j = 0$, so it is a second point of $\mathcal{P}(A)$ when $n \geq 2$. □

Restatement (Proposition 2, discontinuity at an exact draw)**.** *Let A_δ be A_5 with agent 5 scoring δ against each other agent. For $\delta > 0$, $\mathcal{P}(A_\delta) = \{\mathbf{e}_5\}$ and $r_4 = -\delta$; for $\delta < 0$, $\mathcal{P}(A_\delta) = \{(1/3, 1/3, 1/3, 0, 0)\}$, the uniform distribution on agents 1 to 3, and $r_4 = -1$; at $\delta = 0$ the maximum-entropy rating of agent 4 is $-3/4$. Since the equilibrium is unique on both sides, no rating rule that selects an exact equilibrium is continuous at $\delta = 0$.*

Proof. Each A_δ is antisymmetric, with $(A_\delta)_{5j} = \delta$ and $(A_\delta)_{j5} = -\delta$ for $j \leq 4$, so $\mathcal{P}(A_\delta)$ is nonempty (Gale et al., 1950) and each case below only has to rule out every other point. For $\delta > 0$,

$$(A_\delta p)_5 = \delta(p_1 + p_2 + p_3 + p_4) = \delta(1 - p_5) \leq 0 \quad \Longleftrightarrow \quad p_5 = 1, \tag{22}$$

so $\mathbf{e}_5$ is the only point that can lie in $\mathcal{P}(A_\delta)$, and being nonempty the set equals $\{\mathbf{e}_5\}$, with $(A_\delta \mathbf{e}_5)_4 = (A_\delta)_{45} = -\delta$. For $\delta < 0$, the three cycle constraints and their sum are

$$\begin{gathered} x_2 - x_3 + x_4 - \delta x_5 \leq 0, \qquad -x_1 + x_3 + x_4 - \delta x_5 \leq 0, \qquad x_1 - x_2 + x_4 - \delta x_5 \leq 0, \\ 3x_4 + 3|\delta| x_5 \leq 0, \end{gathered} \tag{23}$$

which force $x_4 = x_5 = 0$. The cycle constraints then force $x_1 = x_2 = x_3 = 1/3$, so that point is the only candidate and, the set being nonempty, it is the equilibrium, with

$$(A_\delta p)_4 = -(x_1 + x_2 + x_3) + 0 \cdot x_4 + (A_\delta)_{45} x_5 = -1. \tag{24}$$

At $\delta = 0$ the equilibrium set is the segment of Theorem 4(c) and the maximum-entropy rating of agent 4 is $-3/4$. For $\delta \neq 0$ the equilibrium set is a single point. Every selection rule returns it, and every rating rule built on an exact equilibrium reports $r_4 = -\delta$ for $\delta > 0$ and -1 for $\delta < 0$. The two one-sided limits differ, which is the last claim. □

Restatement (Proposition 1, near copies)**.** *Let A_η be antisymmetric with leading block A and appended row within η of row k of A in its first n entries. As $\eta \to 0$ for this fixed A, uniformly over such matrices and all selection rules: if p_k is constant on $\mathcal{P}(A)$, the total probability of agent k and its copy converges to that constant; if p_i is constant on $\mathcal{P}(A)$ for some $i \neq k$, that probability converges to its constant; and if the rating of agent i is constant on $\mathcal{P}(A)$, that rating converges to the constant. A unique equilibrium of A therefore forces every probability and every rating to its value under A.*

Proof. Every admissible A_η is antisymmetric. Thus the symmetric zero-sum game it defines has value zero and an equilibrium, and $\mathcal{P}(A_\eta)$ is nonempty. The appended column is the negative of the appended row, and the head-to-head entry is within η of $A_{kk} = 0$, so the two matrices differ by at most η in every entry, $\|A_\eta - \operatorname{dup}_k A\|_{\max} \leq \eta$. Suppose the total-probability claim fails for some k with $p_k \equiv c$ on $\mathcal{P}(A)$: there are $\eta' > 0$, $\eta_j \to 0$, admissible A_{η_j}, and points $p^{(j)} \in \mathcal{P}(A_{\eta_j})$ (chosen by any selection rules) with $|p_k^{(j)} + p_{n+1}^{(j)} - c| \geq \eta'$ for all j. By compactness of the simplex a subsequence converges to some $\bar{p} \in \Delta_{n+1}$, and

$$\begin{aligned} (\operatorname{dup}_k A)\bar{p} &= \lim_j A_{\eta_j} p^{(j)} && \text{(joint continuity of } (M, p) \mapsto Mp, \|A_{\eta_j} - \operatorname{dup}_k A\|_{\max} \leq \eta_j \to 0) \\ &\preceq 0 && \text{(each } A_{\eta_j} p^{(j)} \preceq 0 \text{; the closed orthant is closed),} \end{aligned} \tag{25}$$

so $\bar{p} \in \mathcal{P}(\mathrm{dup}_k\, A)$ and by Lemma 1 $\pi(\bar{p}) \in \mathcal{P}(A)$, whence

$$\bar{p}_k + \bar{p}_{n+1} = \pi(\bar{p})_k = c, \tag{26}$$

contradicting $|\bar{p}_k + \bar{p}_{n+1} - c| \geq \eta'$. The rating claim is identical: along the same construction,

$$\lim_j (A_{\eta_j} p^{(j)})_i = ((\mathrm{dup}_k\, A)\bar{p})_i = (A\,\pi(\bar{p}))_i \qquad \text{(rating identity (4))}, \tag{27}$$

which equals the constant rating of agent i on $\mathcal{P}(A)$. For an original agent $i \neq k$ the same limit gives $\bar{p}_i = \pi(\bar{p})_i$, which is the constant, and a unique equilibrium makes every coordinate constant. Here $p^{(j)}$ itself need not converge: the split of the total probability between k and its copy is free in the limit. □

A.6 The agent-versus-task version

Section 3 states that the results transfer to the agent-versus-task game of the original paper. The original paper evaluates agents on tasks through the antisymmetric block matrix $\begin{pmatrix} 0 & S \\ -S^\top & 0 \end{pmatrix}$; each meta-player chooses a pair of distributions $(p_a, p_t) \in \Delta_{n_a} \times \Delta_{n_t}$ over agents and tasks. A pair is optimal if and only if $\min_j (S^\top p_a)_j \geq \max_i (S p_t)_i$, the saddle-point condition for this block game, and the optimal pairs therefore form the product of the two one-sided optimal sets. Call an agent or a task essential if some optimal pair gives it positive mass. At equilibrium both sides equal a common value. Agents in the support of p_a have skill equal to that value and every other agent at most that value, and dually for tasks. That paper selects the maximum-entropy pair without writing the objective. The feasible set is the set of optimal pairs, the product of the two equilibrium sets inside $\Delta_{n_a} \times \Delta_{n_t}$, not the simplex over the $n_a + n_t$ rows of the block matrix. Its two natural objectives differ by an increasing affine map and select the same point: the sum of the two entropies, and the entropy of the pair as one distribution with each block carrying mass one half.

Proposition 4 (Agent-versus-task transfer)**.** *Duplicating task j appends a copy of column j of S. Merging the two task coordinates maps the optimal pairs of the duplicated game onto those of the original, and on that set the objective becomes the sum of the two entropies plus $(p_t)_j \log 2$, the tilt acting on the task distribution alone. Hence Lemma 1, the tilt formula, the criterion and Theorem 3 hold with p_t in the role of p. Skills and difficulties on the supports are fixed at the common value, and the others move with the selected pair.*

Proof. Write S' for S with column j copied and p'_t for a task distribution of the duplicated game. Then $S' p'_t = S\,\pi(p'_t)$, as the copied column adds the mass of the copy to column j. Also $S'^\top p_a$ repeats $(S^\top p_a)_j$ in the copied coordinate, which leaves its minimum unchanged. Hence (p_a, p'_t) is optimal for the duplicated game exactly when $(p_a, \pi(p'_t))$ is optimal for the original, which is Lemma 1 for the task coordinate. Along the splits of a task distribution the agent entropy is unchanged and the task entropy is $H(\pi(p'_t)) + (\pi(p'_t))_j\, H_2(\alpha)$ by Lemma 2. The even split maximizes it with value $H(\pi(p'_t)) + (\pi(p'_t))_j \log 2$, which is the tilt formula. Lemma 3 holds here with the same proof: part (a) uses only that the feasible set is convex and that $-x \log x$ has infinite slope at 0, both of which hold coordinatewise in each block, and part (b) follows from (a). The proofs of Theorems 2 and 3 then use only the strict concavity of the objective, the compactness and convexity of the feasible set, and that feasible directions vanish outside the essential support. The product set and the sum of entropies have these three properties, so the criterion and Theorem 3 follow with p_t in the role of p. Skills and difficulties on the supports equal the common value at every optimal pair: the value is $c = \sum_i (p_a)_i (S p_t)_i$ with $(S p_t)_i \leq c$ for every i, so $(S p_t)_i = c$ on $\mathrm{supp}(p_a)$, and since the optimal pairs form a product this holds against every optimal p_t; the argument for tasks is dual. They do not move. All others are evaluated at the selected pair and may move with it. □

A two-agent, three-task example shows the skill of an inessential agent moving. Let agent 1 score $1/2$, 1 and $1/2$ on the three tasks and agent 2 score 0, 0 and 1. The value is $1/2$: agent 1 alone guarantees it. Thus $p_a = \mathbf{e}_1$, and the optimal task distributions are those that hold agent 1 at $1/2$ and agent 2 at most $1/2$, the segment $(t, 0, 1-t)$ with $t \in [1/2, 1]$ (the skill $1 - t$ of agent 2 must not exceed the value). Here the maximum-entropy pair takes $t = 1/2$, where the skill of agent 2 is $1/2$. Duplicating task 1 adds $t \log 2$ to the task entropy, so its maximum on the segment moves to $t = 2/3$. The skill of agent 2 falls to $1/3$ while agent 1 stays at the value.

A.7 Duplication in the classical coordinates

Lemma 4 (Duplication in the classical coordinates)**.** *Write d for the affine dimension of $\mathcal{P}(A)$. Duplication preserves $e - d$ and $\operatorname{rank}(A_{\mathcal{E}})$: duplicating an essential agent increases e and d by one each, and duplicating an inessential agent changes none of the three quantities.*

Proof. By Lemma 1 the equilibria of $\operatorname{dup}_k A$ are exactly the splits $\iota_\alpha(p)$ of the equilibria of A. Suppose agent k is essential, so some $p \in \mathcal{P}(A)$ has $p_k > 0$. The map π is linear and surjective with one-dimensional kernel spanned by $\mathbf{e}_k - \mathbf{e}_{n+1}$. Hence the set of splits, the preimage of $\mathcal{P}(A)$ inside the simplex, has affine dimension at most $d + 1$. It has exactly $d + 1$: the preimage maps onto $\mathcal{P}(A)$ under π, so its direction space has an image of dimension d, and the kernel direction is realized inside the simplex along the splits of such a p, which form a segment of positive length, adding one more. The copy receives positive mass on the interior of that segment and is essential. Agent k stays essential (any $\iota_\alpha(p)$ with $\alpha \in (0, 1)$ gives it positive mass). Every other original agent keeps its status because $\pi(p')_i \geq p'_i$ and $(\iota_\alpha p)_i = p_i$ for $i \neq k$. Hence e increases by exactly one. The essential submatrix of the duplicated game is $A_{\mathcal{E}}$ with the row and column of agent k repeated, since the three new entries (copy versus k, k versus copy, copy versus itself) all equal $A_{kk} = 0$. A repeated row and column add no rank. Suppose instead agent k is inessential, so $p_k = 0$ on all of $\mathcal{P}(A)$. The splits of each equilibrium reduce to a single point, so the equilibrium set is carried isomorphically onto the set of splits. The copy never receives mass and is inessential, and none of e, d, $\operatorname{rank}(A_{\mathcal{E}})$ changes. In both cases $e - d$ and $\operatorname{rank}(A_{\mathcal{E}})$ are preserved, so the classical identity for the duplicated game is the identity for A. In particular the parity of $e - d$ is preserved even though duplication changes the parity of n. □

A.8 Classical background

An equilibrium is quasi-strict when every pure best reply is played with positive probability. Harsanyi (1973) calls this notion quasi-strong and Jansen (1981) regular, a word Harsanyi reserves for a nonvanishing Jacobian. At $p^\star$ the pure best replies are the agents with $(Ap^\star)_i = 0$ and the support is $\mathcal{E}(A)$. Thus the assumption of Corollary 1 is exactly quasi-strictness of $(p^\star, p^\star)$. For matrix games Jansen (1981) shows that the regular equilibria are exactly the relative interior of the product of the two equilibrium sets, in our symmetric setting $\operatorname{relint}\mathcal{P}(A) \times \operatorname{relint}\mathcal{P}(A)$. The proof of Corollary 1 uses the same fact in one coordinate. His degree-of-freedom count for a regular maximal Nash subset (a maximal convex set of equilibria), for positive payoffs, is $e - d = \operatorname{rank}(A_{\mathcal{E}})$ per factor. For a game of value zero the normalization $\mathbf{1}^\top q = 1$ is no longer implied by the equilibrium equations $A_{\mathcal{E}} q = \nu \mathbf{1}$, because the implication divides by the value ν. That equation becomes independent, the extra 1 in the identity $e - d = \operatorname{rank}(A_{\mathcal{E}}) + 1$. Since $e - d$ is then determined by the rank, more than one equilibrium, that is $d \geq 1$, occurs exactly when $\operatorname{rank}(A_{\mathcal{E}}) \leq e - 2$, which is the criterion of Section 2. Bohnenblust et al. (1950) print both forms side by side, indexed by whether the value vanishes. Adding a constant to every entry of an antisymmetric matrix moves a game from one form to the other without changing e or d. Bohnenblust et al. (1950) first proved this identity. Gale & Sherman (1950) gave an algebraic proof and the terms essential and superfluous, and Gale et al. (1950) the symmetric form. Since an antisymmetric matrix has even rank, $e - d$ is odd. For $d = 0$ this parity is the corollary of Gale & Sherman (1950) that a unique solution has an odd number of essential agents, after the even-order impossibility of Kaplansky (1945) (general criterion in Kaplansky, 1995). Quasi-strict equilibria exist in every bimatrix game (Norde, 1999), and for matrix games this is the statement that $\operatorname{relint}\mathcal{P}(A)$ is nonempty (Jansen, 1981). Whether $p^\star$ gives the copied agent its largest probability is a property of Nash averaging rather than of the set: $\operatorname{relint}\mathcal{P}(A)$ is never empty, but $p^\star$ need not lie in it. Genericity results for general bimatrix games (Harsanyi, 1973) do not transfer: symmetric games are a measure-zero class (Norde, 1999), and the almost-sure statement we use is the one proved inside the antisymmetric class (Brandl, 2017).

A.9 Deviation ratings

Deviation ratings (Marris et al., 2025) come from a lexicographic construction, which in the symmetric two-player zero-sum case reduces to the lexicographic rule of Lemma 5. The rule proceeds in rounds: it lowers the largest unfixed rating as far as the equilibrium constraints allow, fixes every rating that takes that value in every optimal solution, and repeats on the rest. Over $\mathcal{P}(A)$ the rule minimizes the largest unfixed coordinate θ of Aq with the fixed coordinates held at their values. It

then fixes at θ exactly those coordinates whose minimum over the optimal face equals θ, and repeats. We solve that fixing test as one linear program per coordinate at tolerance 10^{-4} (Table 5). Each round fixes at least one coordinate, so the rule terminates and returns a single rating vector. If every unfixed coordinate a' admitted some $q_{a'}$ in the optimal face with $(Aq_{a'})_{a'} < \theta$, their average would lie in that face, which is convex, and would hold every unfixed coordinate below θ, contradicting the value θ of the round. The rating of agent i is $(Aq)_i$ at the selected q, zero on the essential support.

Lemma 5 (Deviation ratings on a symmetric zero-sum game)**.** *Let A be antisymmetric and consider the two-player game with payoffs $G_1(a,b) = A_{ab}$ and $G_2(a,b) = -A_{ab}$. Run Algorithm 1 of Marris et al. (2025), the deviation-rating procedure, on joint distributions σ over pairs (a,b). In each round, minimize over σ the largest deviation gain of either player over the actions not yet fixed, while the fixed gains keep their values. Then fix every unfixed gain that equals the minimum on the whole optimal face. Both players' rating vectors equal the rating vector of the lexicographic rule on A.*

That procedure freezes the constraints that are active at the optimum, and Marris et al. read active as carrying a nonzero value in the optimal dual solution their solver returns. We use the primal form instead: a gain is fixed when it equals the value of the round at every point of the optimal face. By complementary slackness a gain with a nonzero dual value meets the primal condition, and a gain that meets the primal condition has a nonzero value in some optimal dual solution (Goldman & Tucker, 1956), though not necessarily in the one a solver returns. A gain that the returned solution misses keeps the same value on the same optimal face in the next round and is fixed there, so the two readings give the same ratings.

Proof. Write u and q for the marginals of σ on the first and second coordinate and $V_\sigma = \sum_{a,b} \sigma(a,b) A_{ab}$. Switching player 1 to a' and player 2 to b' gains

$$\delta_1^\sigma(a') = \sum_{a,b} \sigma(a,b)(A_{a'b} - A_{ab}) = (Aq)_{a'} - V_\sigma,$$
$$\delta_2^\sigma(b') = \sum_{a,b} \sigma(a,b)(-A_{ab'} + A_{ab}) = (Au)_{b'} + V_\sigma \qquad \text{(antisymmetry)}. \tag{28}$$

The gains depend on σ only through (u, q, V_σ), and every pair (u,q) of distributions is the marginal pair of the product $\sigma = u \otimes q$, for which $V_\sigma = u^\top A q$.

We show by induction over the rounds that at the start of round ℓ the fixed pairs of both players are the same set F_ℓ of agents with the same values $\rho(a')$ for $a' \in F_\ell$. We also show that F_ℓ and ρ agree with the lexicographic rule after $\ell - 1$ rounds.

Base case. This holds trivially for $\ell = 1$, where F_1 is empty.

Inductive step. In round ℓ the joint program is

$$\min_\sigma \max \left\{ \max_{a' \notin F_\ell} \delta_1^\sigma(a'),\ \max_{b' \notin F_\ell} \delta_2^\sigma(b') \right\} \quad \text{subject to} \quad \delta_1^\sigma(a') = \delta_2^\sigma(a') = \rho(a') \ \text{ for } a' \in F_\ell. \tag{29}$$

Let θ be the value of the round of the lexicographic rule, attained at some $\hat{q} \in \mathcal{P}(A)$. Let $Q^\star \subseteq \mathcal{P}(A)$ be its optimal face, the feasible q with $\max_{a' \notin F_\ell} (Aq)_{a'} = \theta$. The value θ is at most 0: the first round has value 0 because $Aq \preceq 0$ on $\mathcal{P}(A)$ and $\max_{a'} (Aq)_{a'} \geq q^\top A q = 0$, and the optimizer of each round stays feasible in the next with every unfixed coordinate at most the previous value. The product $\hat{q} \otimes \hat{q}$ has $V_\sigma = \hat{q}^\top A \hat{q} = 0$ and gains $(A\hat{q}, A\hat{q})$. It is feasible for the joint program with objective θ, hence the joint value is at most θ. Conversely, let σ be feasible with objective at most θ. Then every gain of σ is at most 0 (the fixed gains equal $\rho \leq 0$), and hence $\max_{a'} (Aq)_{a'} \leq V_\sigma$ and $\max_{b'} (Au)_{b'} \leq -V_\sigma$. Since $\max_{a'} (Aq)_{a'} \geq q^\top A q = 0$ and likewise for u, this forces $V_\sigma = 0$ and $u, q \in \mathcal{P}(A)$. The gains of σ are then (Aq, Au), with q and u feasible for the one-sided program and each with objective at most θ. That program has value θ, so neither objective can be below it and both equal θ. Hence the joint value is θ, and the optimal face of the joint program is the set of feasible σ with $V_\sigma = 0$ and $u, q \in Q^\star$. Its marginal pairs are exactly $Q^\star \times Q^\star$: every product $u \otimes q$ with $u, q \in Q^\star$ has $V_\sigma = u^\top A q = 0$, because $Aq \preceq 0$ and $u \succeq 0$ give $u^\top A q \leq 0$ while antisymmetry and $Au \preceq 0$ give $u^\top A q = -q^\top A u \geq 0$, so the product lies in that face. A gain $(1, a')$ is therefore tight on the whole optimal face if and only if $(Aq)_{a'} = \theta$ for every $q \in Q^\star$, and the same holds for $(2, a')$ with u in place of q. This is the fixing condition of the lexicographic rule.

Both players therefore fix the same agents at the value θ, which completes the induction. When every agent is fixed, both rating vectors equal ρ, the output of the lexicographic rule. □

B SYNTHETIC FAMILIES, SOLVERS AND PROVENANCE

Table 4 gives the per-family duplication statistics, Table 3 gives the threshold sweep and Table 5 the solver, slack and threshold behind each body number. We recorded the protocol and the success and stop condition of each run before generating the first instance. Conditions and their outcomes are in a dated text file in the supplementary material.

B.1 FAMILIES

Every instance is an antisymmetric matrix of size $n \in \{4, \dots, 8\}$ with n drawn uniformly. There are 2000 instances per family from the random seed 1000 (NumPy default generator). The family with planted ties draws one of three shapes with equal probability. One shape is a skill-difference matrix $A_{ij} = s_i - s_j$ with standard Gaussian skills and one random pair of skills forced equal. Another is a cyclic triple with log-odds 1 among three random agents. Every other agent draws against everyone. The last is a cyclic triple with log-odds 1. When $n \geq 6$ it adds a second disjoint cyclic triple with log-odds $1/2$ on three further agents, with no entries between the two triples. All remaining agents draw against everyone. The family with dominated agents draws one of three other shapes with equal probability, each of which plants a tie and adds dominated agents. One shape is a cyclic triple that beats one or two added agents, their number drawn uniformly. It beats each added agent at a log-odds drawn independently from $\{1/2, 1, 2\}$ and used against all three members of the triple; the rest draw against everyone. Another is a skill matrix with the two strongest skills forced equal and the weakest agent replaced by an agent that draws everyone. The last is one cyclic triple with log-odds 1, joined when $n \geq 7$ by a second disjoint triple with log-odds $1/2$ and no entries between them; one further agent is then beaten at log-odds 1 by every member of those triples, and the rest draw against everyone. In the Gaussian games the matrix is $Z - Z^\top$ with Z i.i.d. standard normal.

Each instance draws n first, then the shape, then the shape's own quantities in the order written above. Seeds for every study are listed in Appendix M.

A table of marks has the shape of the skill-difference matrix above. Mark n agents right or wrong on each of N common items, and let A_{ij} be the share of items on which i is right and j wrong, less the share on which j is right and i wrong. An item the two answer alike contributes nothing. An item they answer differently contributes $\pm 1/N$, with a sign set by which of them was right and by nothing else. Summing over the items gives $A_{ij} = (r_i - r_j)/N$, where r_i counts the items agent i answered right. The ordinal form of this is classical: on dichotomous ballots the majority relation is the order of the approval scores (Brandl & Peters, 2022), and for a single pair the identity is the one behind McNemar's test (McNemar, 1947). A benchmark marked item by item therefore gives Nash averaging a difference matrix. On such a matrix it reproduces the ranking by score, it is invariant to copies, and the choice of equilibrium makes no difference. All three of these depend on the margin $(r_i - r_j)/N$ alone. A rule fitted instead to the per-item wins and losses separately, which is how Bradley–Terry is usually fitted to such data, uses more than the margin: on 349 sampled tables of 4 agents and 60 items it disagreed with the ranking by score on 11, while the same rule fitted to the margin agreed on all 349.

B.2 EQUILIBRIUM SET AND WIDTH

We compute the width of Section 2 over the slackened set $\{p \in \Delta_n : Ap \preceq \tau \mathbf{1}\}$ with $\tau = 10^{-9}$: the largest range of one agent's probability over that set. It is computed by two linear programs per coordinate (HiGHS). An instance counts as having more than one equilibrium when the width exceeds 10^{-4}. We sweep the detection threshold from 10^{-9} to 10^{-3} (Table 3). Counts are stable from 10^{-5} on in the games with planted ties and from 10^{-8} on in the games with dominated agents. At 10^{-9} the counts are inside the linear-programming tolerance. In the Gaussian games the detections vanish from 10^{-5} on.

Table 3: The counts are stable from 10^{-5} on in the two families with planted ties and the Gaussian detections vanish there. Counts of instances with more than one equilibrium out of 2000 per family as the width threshold is swept; the 10^{-9} column is inside the linear-programming tolerance.

threshold	10^{-9}	10^{-8}	10^{-7}	10^{-6}	10^{-5}	10^{-4}	10^{-3}
planted ties	1876	1492	1446	1440	1439	1439	1439
dominated agents	2000	1643	1643	1643	1643	1643	1643
Gaussian	1962	920	129	15	0	0	0

B.3 Dimension

For the instances with more than one equilibrium the affine dimension $d = e - \operatorname{rank}(A_{\mathcal{E}}) - 1$ is computed with the essential support from the linear programs and a rank tolerance of 10^{-6}. It ranges over 1 to 5 in the games with planted ties and 1 to 4 in the games with dominated agents. On every instance $e - d$ is odd.

B.4 Maximum-entropy equilibrium

Entropy is maximized over $\{p \succeq 0,\ \mathbf{1}^\top p = 1,\ Ap \preceq \tau\}$ by sequential quadratic programming (SLSQP in SciPy) with tolerance 10^{-12}. Our solver starts from three points, the uniform distribution and two draws from the symmetric Dirichlet with concentration 1 seeded 1 and 2, and keeps the best of the three runs. On 150 fresh instances of each family and on the same instances after duplicating one random agent, the resulting points agree to 1.1×10^{-4} in the sup norm with the independent implementation shipped in OpenSpiel. That implementation (nash_averaging) maximizes the same entropy under the same constraints through CVXPY. Ratings agree to 8.2×10^{-5}.

B.5 Duplication test

For every instance and every agent k, the duplicated matrix appends row and column k with a zero copy-versus-original entry. We compare the maximum-entropy equilibrium of the duplicated matrix with the even split of the original equilibrium (deviation in the sup norm). We compare the original agents' rating vectors before and after and record the total probability of the duplicated agent and its copy. A test is counted only when the original instance has width above the threshold 10^{-4}.

B.6 Convex-solver recomputation

All duplication statistics of the first two families are recomputed with the convex solver (CVXPY with Clarabel), solving every instance and every agent and then counting, as above, only the tests whose original instance has width above the threshold. At the 10^{-4} threshold the convex solver moves the point on 8244 of the 8708 tests of the first family and on 7309 of the 10 460 tests of the second. The criterion of Theorem 2 predicts movement exactly on the tests whose copied agent has a probability that varies across equilibria, and those number 8244 and 7305; the four extra tests in the second family are the solver artefacts described next. An inessential rating moves on 0 tests in the games with planted ties and on 6680 in the games with dominated agents. At 10^{-6} the counts rise to 8683 and 10 361, because that threshold lies within the precision of the convex solver. Four tests counted at 10^{-4} have a duplicated coordinate that is constant (range below 10^{-8}). Their deviations are 1.0 to 1.7×10^{-4} at the default solver tolerance, at most 1.5×10^{-5} at tolerance 10^{-10}, and at most 4×10^{-7} under the sequential-quadratic solver. They are solver artefacts at the threshold, and Section 4.2 quotes the counts of tests whose copied agent has a probability that varies across equilibria.

Table 4: A copy moves the maximum-entropy equilibrium on most tests of the first two families. One test per agent of each game with more than one equilibrium; fractions from the convex solver, with the four artefacts included, and medians and maxima from the sequential-quadratic run; dev. and rating: sup-norm distance from the even split and sup-norm change of the original ratings.

family	more than one equilibrium	tests (on them)	frac. dev. $> 10^{-4}$	dev. med. / max	frac. rating $> 10^{-4}$	rating med. / max
planted ties	1439 (72.0%)	8708	0.947	0.027 / 0.167	0.000	$0.000 / 3 \times 10^{-9}$
dominated agents	1643 (82.2%)	10460	0.699	0.019 / 0.083	0.639	0.036 / 0.88
Gaussian	0 (0.0%)	0	–	–	–	–

No constant coordinate moves the point beyond solver precision. Under the sequential-quadratic solver the largest deviation over the 464 constant-coordinate tests of the games with planted ties and the 3155 of the games with dominated agents is 2.5×10^{-6}. Of these, 5 and 23 exceed 10^{-6} and none exceeds 3×10^{-6}. The sequential-quadratic solver at the 10^{-6} threshold gives the same counts of instances with more than one equilibrium (1439, 1643, 0). Deviation fractions are 0.947 and 0.701 (medians 0.027 and 0.019, maxima 0.167 and 0.083). Rating-moved fractions are 0.000 and 0.628 (medians 0.000 and 0.036, maxima 0.000 and 0.880). It finds no fall in the total probability of a copied agent, an LP vertex support different from the maximum-entropy support on every instance with more than one equilibrium, and median LP vertex ℓ_1 distances 1.000 and 1.143. These agree with the convex-solver figures of Table 4 to within two percentage points. The median and maximum columns of that table and the 3×10^{-9} rating invariance of Table 4 are the sequential-quadratic values. For Section 4.2 the reordering fractions are 271 and 461 of the 714 instances. At the 90th percentile the largest inessential-rating change under duplication is 0.19 in the games with dominated agents.

B.7 Exact certificate for the counts

Every double-precision entry is a dyadic rational, so the equilibrium set of each sampled matrix can be enumerated exactly. Vertex enumeration in rational arithmetic (cddlib) on all 6000 instances gives 1439, 1643 and 0 instances with more than one equilibrium in the games with planted ties, the games with dominated agents and the Gaussian games. These equal the linear-programming counts at the 10^{-4} threshold on every instance. Because the smallest positive exact width is $1/3$, no instance has a width between 0 and 10^{-4}. In exact arithmetic the dimension distribution equals the one reported above, so the identity $e - d = \operatorname{rank}(A_{\mathcal{E}}) + 1$ holds on every one of the 3082 instances with more than one equilibrium. Exact essential supports agree with the linear-programming supports on 5937 instances. The 63 exceptions (7 with planted ties, 56 Gaussian, 0 with dominated agents) contain agents of exact mass zero that the 10^{-9} constraint slack lifts to a maximum mass between 10^{-7} and 4.7×10^{-6}, above the support tolerance 10^{-7}. Each of the 63 has one such agent, except two that have two. None of the reported counts changes.

B.8 Accuracy of the entropy solutions

On all 22 250 entropy problems of the two families the convex-solver solutions satisfy the equilibrium constraints to 6×10^{-8} and the normalization to 8×10^{-9}. A first-order solver (SCS) run on the first 3487 and 3493 problems of the two families agrees in objective value to 1.2×10^{-4} and in the maximizer to 10^{-4}, the precision of that solver.

B.9 Negative control

On instances with a unique equilibrium the even split must be reproduced to 10^{-6}. Testing every instance with a unique equilibrium finds four crossings of that bound in the games with planted ties and fourteen in the Gaussian games. The sequential-quadratic solver fails on 24 of the roughly 12 040 Gaussian duplications (2 original games, 22 duplicated games), and we skip those 24. Exact enumeration gives every Gaussian instance a single point. Every crossing has width below 6×10^{-6} and deviation below 6×10^{-6}, inside solver precision. In the games with dominated agents there are none. The tolerance τ in the width scan and the entropy maximization enlarges the feasible set

by τ in every constraint, and the enlargement grows with the dimension. That mechanism produces these solver-precision crossings.

Table 5: Every count of Section 4.2, Table 14 and the ε-sets comes from a named solver at a stated slack and threshold; the other appendices state theirs at the point of use.

result	solver	slack	threshold
counts of instances with more than one equilibrium and widths (all sections)	HiGHS linear programs; exact rational vertex enumeration (cddlib) as certificate	10^{-9}; 0	width $> 10^{-4}$; width > 0
Table 4, Section 4.2 fractions	CVXPY/Clarabel (default tolerance)	10^{-9}	deviation $> 10^{-4}$
counts of tests whose copied agent has a probability that varies or is constant (Section 4.2)	LP range of p_k; SLSQP for deviations	10^{-9}	range $> 10^{-4}$; deviation $> 10^{-6}$
sequential-quadratic figures (this appendix)	SLSQP	10^{-9}	10^{-6}
LP form of the criterion (Appendix A.3)	HiGHS	10^{-9}	gap $> 10^{-3}$
perturbed draws between identical agents (Section 4.2)	Clarabel, tolerance 10^{-10}	0	none (distances reported)
competing rules (Table 14)	Clarabel; lexicographic LP; MM; OpenSpiel α-Rank	10^{-9}	10^{-4}; reversals at 10^{-6}
league and leaderboard ε-sets	HiGHS	ε	width ≥ 0.2

B.10 The linear-programming check on the first two families

On the 19 168 duplication tests of the first two families the gap $\max_q w^\top q - w^\top p^\star$ is at least 0.057 on every test whose duplicated point moves by more than 10^{-4}, except the four solver artefacts of the convex-solver recomputation above. It is at most 3.6×10^{-4} on every test with a constant coordinate. At a gap threshold of 10^{-3} the linear-programming verdict agrees with the observed movement on 19 164 of 19 168 tests. The four exceptions are those artefacts. In Table 4 the moved/unmoved classification rests on the observed movement of the entropy solutions, and the agreement above is the cross-check.

B.11 Inessential ratings and the LP vertex

For every instance we compute the essential support by maximizing each coordinate over the equilibrium set. Two further linear programs compute the range of every inessential rating over the equilibrium set. The LP vertex (zero objective) is what a simplex solve returns and is computed by a linear program with zero objective, while the LP vertex (minimax) comes from the standard minimax linear program. We record support and ℓ_1 distance to the maximum-entropy equilibrium.

C Hand-built instances, and an error-aware verdict rule

Four hand-built instances illustrate Theorem 2 and Theorem 4. Table 6 gives the cases near a degeneracy on which the linear-programming check and the convex solver are tested, and Table 7 an error-aware version of the verdict rule, which reports an error term with each answer. Instances 1 and 2 meet the second sufficient condition of Theorem 2. They have more than one equilibrium and the duplicated coordinate is not constant over their equilibrium sets, yet the copy leaves the equilibrium in place, exactly because the maximum-entropy equilibrium already maximizes that coordinate. In each, an inessential agent has rating exactly zero at the maximum-entropy equilibrium. That is the situation Corollary 1 excludes. Instance 3 shows width with no zero off-diagonal entry and no repeated row. Instance 4 shows a reordering of two inessential agents.

C.1 Instance 1 (integer 4×4)

Take the integer matrix with rows $(0,-1,-1,-1)$, $(1,0,1,0)$, $(1,-1,0,1)$, $(1,0,-1,0)$. Its equilibrium set has $p_4 \in [0,1/2]$ and the maximum-entropy equilibrium is $(0,1/2,0,1/2)$. Duplicating agent 4 gives $(0,1/2,0,1/4,1/4)$ to 2×10^{-9} under the sequential-quadratic solver (to 3.0×10^{-6} under the convex solver of Table 6).

C.2 Instance 2 (the continuity example)

The continuity example of the original paper is $A = C + \gamma T$ in its notation, with the cyclic C having first row $(0,1,-1)$ and the transitive T having first row $(0,1,2)$ and second row $(-1,0,1)$. Its published closed form for the maximum-entropy equilibrium is $((1+\gamma)/3, (1-2\gamma)/3, (1+\gamma)/3)$ for $\gamma \le 1/2$ and $(1,0,0)$ beyond. Our solver reproduces it on $[0,1]$. At $\gamma = 1/2$ the equilibrium set is the segment $p_1 \in [1/2,1]$, $p_2 = 0$, $p_3 = 1 - p_1$, on which the maximum-entropy equilibrium $(1/2,0,1/2)$ maximizes p_3. Duplicating agent 3 gives $(1/2,0,1/4,1/4)$ to 6×10^{-8}. The invariance claim therefore holds in that example, but by the second sufficient condition of Theorem 2 and not by uniqueness.

C.3 Instance 3 (no zero off-diagonal entry, no repeated row)

The matrix with rows $(0,-3,2,-1,2)$, $(3,0,-3,-3,-2)$, $(-2,3,0,3,-2)$, $(1,3,-3,0,2)$, $(-2,2,2,-2,0)$ has every off-diagonal entry nonzero and no two rows alike. Its essential support is $\{1,2,3,4\}$ with $\operatorname{rank}(A_{\mathcal{E}}) = 2 = e - 2$. The ranges of the five agents' probabilities are 0.09, 0.18, 0.15, 0.24, 0 and the maximum-entropy equilibrium is $(0.429, 0.142, 0.285, 0.145, 0)$. Its two vertices $(9/22, 2/11, 7/22, 1/11, 0)$ and $(1/2, 0, 1/6, 1/3, 0)$ have supports of size four and three. Duplicating agent 1, 2, 3 or 4 moves the maximum-entropy equilibrium away from the even split by 0.020 to 0.043 in the sup norm, and the rating of the inessential agent 5 by 0.11 to 0.31. Duplicating agent 5 changes nothing.

C.4 Instance 4 (a reordering of two inessential agents)

Append to A_5 a sixth agent with row $(-0.7,-0.7,-0.7,0,-0.7,0)$. It loses to agents 1, 2, 3 and 5 at log-odds 0.7 and draws agent 4. It receives no mass in any equilibrium, so $\mathcal{P}(A)$ is the segment $\{(t,t,t,0,1-3t,0)\}$. Its rating is -0.7 at every equilibrium, while $r_4 = -3t$ ranges over $[-1,0]$. At the maximum-entropy equilibrium $t = 1/4$ and $r_4 = -3/4 < -0.7 = r_6$. Duplicating agent 5 moves t to $1/5$ and r_4 to $-3/5 > -0.7$. The two inessential agents change order. The solver check gives $r_4 = -0.7500 \to -0.6000$ and $r_6 = -0.7000 \to -0.7000$. Reordering depends on the rule. With the sixth agent losing at log-odds $x \in \{0.6, 0.65, 0.7, 0.8, 0.9\}$ and agent 5 duplicated, the five rules behave as follows. The maximum-entropy equilibrium moves r_4 from -0.75 to -0.6 at every x. Thus the pair changes order for $x = 0.65$ and 0.7, ties at $x = 0.6$, and keeps agent 4 ahead for $x \ge 0.8$. Bradley–Terry also changes their order at $x \le 0.7$ (agent 4 rises from -0.48 to -0.39 while agent 6 falls). The LP vertex (zero objective) and deviation ratings select $t = 1/3$, where $r_4 = -1$ before and after duplication, so agent 6 stays ahead at every x. α-Rank gives agent 6 no mass and agent 4 mass 0.001 to 0.003 at every x, and agent 4 stays ahead.

C.5 Numerical behaviour near a degeneracy

The generated families have no instance of the second sufficient condition of Theorem 2 and a smallest positive width of $1/3$, so this check interpolates around Instances 1 and 2 to test the linear-programming check and the convex solver at small distance to a degeneracy. Table 6 covers two families. One is $A = C + \gamma T$ at $\gamma = 1/2$ and at the twelve signed distances $\gamma = 1/2 \pm 10^{-7}$ to $\pm 10^{-2}$ (duplicating agent 3). The other is Instance 1 with its zero entry A_{24} replaced by the same twelve signed distances (duplicating agent 4). For each case the table reports the width of the equilibrium set at slack 10^{-9}, the gap of the linear program of Appendix A.3, and the observed sup-norm deviation of the duplicated game's maximum-entropy equilibrium from the even split. Table 6 uses a requested tolerance of 10^{-12}; the looser settings below are a separate check on how far the status flag of the solver can be trusted. At the degeneracy itself the gap and the deviation are both of order 10^{-6}, the invariant case in which the maximum-entropy equilibrium already gives the copied

agent its largest probability. At distance 10^{-4} and beyond on either side the set is a single point, the gap is at most 2×10^{-8} in absolute value, and the deviation at most 6×10^{-9}. The width the solver reports there is set by the slack rather than by the geometry: it is the slack divided by the distance for Instance 1 (10^{-5} and 10^{-7}) and half that for $C + \gamma T$ (5×10^{-6} and 5×10^{-8}), because a change of γ moves the entry that vanishes at the degeneracy by twice that change. Close to the degeneracy on the side where the maximum-entropy equilibrium jumps ($\gamma > 1/2$, or $A_{24} > 0$) the solver returns the status "inaccurate". Its point violates an equilibrium constraint by 10^{-7} and gives the vanishing coordinate mass 0.04 to 0.05. Observed deviation at distance 10^{-6} is 0.031 and 0.014, and the linear-programming gap is negative there (-0.18 and -0.16). That is impossible for a feasible input, since $q = p^\star$ is itself feasible. A negative gap therefore does not decide invariance; it flags the input as violating the premises of the check. The rule accordingly uses only the original entropy solution and the linear program:

1. a solver status other than optimal, or a gap below -10^{-6}, returns undetermined;

2. a gap above 10^{-3} returns moved;

3. any other gap returns no movement detected at the gap threshold 10^{-3}.

For the 26 cases of both families at signed distances 0 and $\pm 10^{-7}$ to $\pm 10^{-2}$ the exact answer is invariant. The rule at the default solver tolerance returns no movement detected on 21 and undetermined on 5, never moved; the five are the jumping side of both families at 10^{-7} and 10^{-6}, and of Instance 1 at 10^{-5}. In the table the duplicated solve agrees with every verdict. At a requested tolerance of 10^{-10} the solver reports an inaccurate status on 19 of the 26 instances, and at 10^{-12} on every instance. The status flag then stops being informative and the negative gap is the remaining signal. It flags the same five cases. Within about 10^{-5} of a degeneracy on the jumping side the numerical point is inaccurate and the rule returns undetermined.

C.6 A FAMILY WITH A NON-INVARIANT ANSWER NEAR A DEGENERACY

The cases above are all invariant. A further family tests the rule where the exact answer is a movement of controllable size. Append to A_5 a sixth agent that loses to agents 1, 2 and 3 at log-odds 1, draws agent 4, and beats agent 5 at log-odds κ. Its equilibrium set is the segment $t \in [\kappa/(3(1+\kappa)), 1/3]$ of the five-agent game A_5. For $\kappa < 3$ the maximum-entropy equilibrium $t = 1/4$ is interior, and duplicating agent 5 moves it toward $t = 1/5$ until the lower bound stops it (a movement that shrinks as $\kappa \to 3$). For $\kappa \geq 3$ the lower bound is at or above $1/4$, the maximum-entropy equilibrium lies on it, and duplication is invariant by the second sufficient condition of Theorem 2. We apply the same rule. Here we compute the duplicated solve only for reference. On $\kappa \in \{2, 2.9, 2.99, 2.999, 2.9999, 3, 3.0001, 3.001, 3.01, 3.1, 4\}$ the rule is correct on eight cases and never returns undetermined or a false movement. It returns no movement detected on the three non-invariant cases $\kappa = 2.99$, 2.999 and 2.9999. There the duplicated solve displaces the point by 3.0×10^{-4}, 2.1×10^{-5} and 9.9×10^{-6} (exact values 3.1×10^{-4}, 3.1×10^{-5} and 3.1×10^{-6}), with gaps 4.2×10^{-4}, 2.9×10^{-5} and 1.4×10^{-5}. On this family the gap and the displacement are proportional (ratio 1.3 to 1.4 on every moved case, 0.058 against 0.042 at $\kappa = 2$ and 0.0044 against 0.0032 at $\kappa = 2.9$). By the strong concavity of the entropy on the support, a gap bounds the displacement at the order of the gap's square root. The invariant cases of the second sufficient condition produce gaps of order 2×10^{-5} from solver precision alone.

Table 6: On the jumping side the solver returns an inaccurate point and a large negative gap at 10^{-7} and 10^{-6} in both families and at 10^{-5} in Instance 1; at distance 10^{-4} and beyond, gap and deviation stay below 2×10^{-8} in absolute value. Width at slack 10^{-9}; gap of the linear-programming check; deviation from the even split.

			jumping side		other side	
family	distance	width	LP gap	deviation	LP gap	deviation
$C + \gamma T$	0	0.5	1.9×10^{-6}	2.8×10^{-6}		
$C + \gamma T$	10^{-7}	5×10^{-3}	-0.187	0.0335	3.8×10^{-6}	5.4×10^{-6}
$C + \gamma T$	10^{-6}	5×10^{-4}	-0.178	0.0312	1.7×10^{-7}	2.4×10^{-7}
$C + \gamma T$	10^{-5}	5×10^{-5}	3.2×10^{-11}	2.3×10^{-10}	2.6×10^{-7}	3.7×10^{-7}
$C + \gamma T$	10^{-4}	5×10^{-6}	-3.6×10^{-11}	3.5×10^{-10}	3.3×10^{-9}	5.9×10^{-9}
$C + \gamma T$	10^{-3}	5×10^{-7}	3.9×10^{-10}	1.3×10^{-9}	4.3×10^{-9}	3.0×10^{-9}
$C + \gamma T$	10^{-2}	5.1×10^{-8}	2.7×10^{-10}	4.1×10^{-11}	2.6×10^{-9}	2.7×10^{-10}
Instance 1	0	0.5	2.0×10^{-6}	2.9×10^{-6}		
Instance 1	10^{-7}	10^{-2}	-0.164	0.0188	6.2×10^{-7}	8.7×10^{-7}
Instance 1	10^{-6}	10^{-3}	-0.157	0.0145	4.7×10^{-7}	6.5×10^{-7}
Instance 1	10^{-5}	10^{-4}	-1.6×10^{-3}	1.8×10^{-4}	2.1×10^{-8}	1.2×10^{-8}
Instance 1	10^{-4}	10^{-5}	-2.5×10^{-10}	1.6×10^{-9}	1.3×10^{-8}	3.6×10^{-9}
Instance 1	10^{-3}	10^{-6}	-4.4×10^{-11}	7.0×10^{-12}	8.3×10^{-9}	8.5×10^{-10}
Instance 1	10^{-2}	10^{-7}	5.0×10^{-11}	8.3×10^{-11}	5.3×10^{-9}	6.6×10^{-11}

C.7 An error-aware version of the rule

The rule above classifies by a fixed gap threshold; this part gives a version whose verdicts come with an error term. Entropy is 1-strongly concave in the Euclidean norm on the simplex, since its Hessian is $-\operatorname{diag}(1/p_i)$ and $1/p_i \geq 1$. Let G be either H or the tilted $H + p_k \log 2$, with maximizer p_G over $\mathcal{P}(A)$, and let q be feasible and positive on the essential support. Concavity and optimality give $\|q - p_G\|_2^2 \leq 2\,\mathrm{gap}(q)$, where $\mathrm{gap}(q) = \max_{z \in \mathcal{P}(A)} \nabla G(q) \cdot (z - q)$ is the linear program of Appendix A.3 evaluated at q. Write $\hat{p}$ for the numerical maximum-entropy equilibrium, $\hat{m}$ for the merged numerical point of the duplicated game, $\hat{\Delta} = \|\hat{m} - \hat{p}\|_\infty$, and e_0 and e_1 for the square roots of twice the two gaps. For a feasible point and an exact gap, the true displacement then lies in $[\hat{\Delta} - e_0 - e_1,\ \hat{\Delta} + e_0 + e_1]$. This interval controls the merged displacement $\|\pi(p^\star(\mathrm{dup}_k\, A)) - p^\star(A)\|_\infty$. The deviation of Section 4.2 and Table 6 is the sup-norm distance of the duplicated point from the even split $\iota_{1/2}(p^\star(A))$. It is zero when the equilibrium stays in place, but when the point moves it records the change of the copied coordinate at half its size because the two copies share that change. A gap below 10^{-8} is floored at that value, which is the accuracy of the linear program and the constraint slack. The smallest movement the interval can resolve is 2.8×10^{-4}. The verdict is bounded moved when the lower end is positive, bounded below 10^{-3} when the upper end is at most 10^{-3}, and undetermined otherwise. A solver status other than optimal, a constraint violation above 10^{-7}, or a numerical point that vanishes on an essential coordinate also returns undetermined. Table 7 reports it on the κ family (the eleven values above and $\kappa = 2.5$), on $C + \gamma T$ and on A_δ at the twelve signed distances $\pm 10^{-7}$ to $\pm 10^{-2}$, where the exact displacements are known, on 200 fresh games with dominated agents that have more than one equilibrium, and on the 153 football duplications (reference displacements from a solve at tolerance 10^{-12}). The 200 instances have one random essential agent duplicated per instance. None of them has a constant coordinate. Thus the below-10^{-3} verdict is exercised on the exact families and the football table. No determinate verdict is false: 370 of the 390 cases are decided and all 370 agree with the exact or reference displacement. Twenty are undetermined. Two are the movements of the κ family below 10^{-3} (6×10^{-4} and 6×10^{-5}). There the computed gap near the degeneracy is 10^{-5} and the error terms are 7×10^{-3}. Four are at distance 10^{-7} to 10^{-6} from a degeneracy, where the solve is inaccurate or misses the support. Fourteen are on the generated and football tables, where the error terms (0.0042 to 0.03 at the default solver tolerance) exceed the reference displacements (3×10^{-7} to 0.02). The fixed-gap rule reports no movement detected on the three small κ movements and is never wrong on the rest. The bound-based rule replaces that undetected movement by an undetermined verdict at the cost of two linear programs and the duplicated solve. Its resolution is set by the solver tolerance rather than by a threshold.

Table 7: The bound-based rule decides 370 of 390 cases with no false verdict. Verdicts: bounded moved, bounded below 10^{-3}, undetermined; false: determinate verdicts that disagree with the exact or reference displacement; error term $e_0 + e_1$, median per family.

family	cases	moved	below 10^{-3}	undetermined	false	error term
κ family (exact)	12	3	7	2	0	2.8×10^{-4}
$C + \gamma T$ (exact)	13	0	10	3	0	2.8×10^{-4}
A_δ (exact)	12	0	11	1	0	2.8×10^{-4}
dominated agents (reference)	200	195	0	5	0	0.014
football (reference)	153	74	70	9	0	0.003

D Discontinuity at exact draws

For Section 4.2 the finding of discontinuity at an exact draw rests on the three perturbation studies reported here. Each perturbs exact draws by a small δ and records how far the maximum-entropy equilibrium jumps. Results are quoted in the text and in Figure 2.

D.1 Identical agents

For this study we select 200 games with planted ties that contain an essential pair with identical rows and a drawn head-to-head. They are the first 200 among 20 000 fresh draws from the generator of games with planted ties that also contain an inessential agent. We replace the draw by $A_{ij} = \delta$ for $\delta \in \{10^{-4}, 10^{-3}, 10^{-2}, 0.03, 0.1\}$. We then recompute the maximum-entropy equilibrium, the LP vertex and the ratings. The convex solver (Clarabel at tolerance 10^{-10}) converges on all 200 instances at every δ, and the numbers in Section 4.2 are from that run. On 198 of the 200 instances the even split collapses onto one member of the pair (median ℓ_1 distance to the $\delta = 0$ point 1.0, tenth percentile 0.99999); on the other two the perturbation pushes both members out of the support, each below 10^{-10}, and those two are excluded from the counts below. Whichever member wins the perturbed head-to-head is the one that remains, since the other becomes weakly dominated. Neither LP vertex consistently picks the surviving member: the LP vertex (zero objective) puts its mass on that member on 79 of the 198 instances where the collapse is complete, and the LP vertex (minimax) on none. Inessential ratings change by a median of 7×10^{-11}, as the identical columns require.

D.2 Agents that draw everyone, in dimension one and above

The jump is not confined to one-dimensional equilibrium sets. In the games with dominated agents, 1560 instances with more than one equilibrium contain an agent that draws everyone. On each, the zero row of that agent is replaced by δ times a random sign vector for $\delta \in \{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}\}$, and the maximum-entropy equilibrium is recomputed with the convex solver at tolerance 10^{-10} and slack 10^{-9}. When the set has dimension at least two (1307 instances) the perturbation moves the maximum-entropy equilibrium by a median ℓ_1 distance of 0.67 (tenth percentile 0.29) and the largest inessential rating by a median 0.42. When the set is a segment (253 instances) the two medians are 0.50 and 0.50. The ℓ_1 medians do not depend on the size of δ. In dimension two and above the rating median falls from 0.42 to 0.37 at $\delta = 0.1$.

D.3 The five-agent sign sweep

The sampling experiment of Figure 2 draws g Bernoulli outcomes for each pair with the logistic win probability $\sigma(A_{ij})$, $\sigma(x) = 1/(1 + e^{-x})$, of the five-agent game A_5. It clips the empirical win rate to $[1/(2g), 1 - 1/(2g)]$, takes its log-odds, antisymmetrizes as $\frac{1}{2}(\hat{A} - \hat{A}^\top)$, and solves the estimated game with the convex solver. There are 400 replicates per g. In Figure 2 the histograms use 40 equal bins on $[-1.3, 0.3]$ for every sample size. We make two runs of this design from the same seed, and they differ only in where each run starts in the random stream: the run behind Figure 2 begins at $g = 300$, while the run reported in Section 4.2 begins at $g = 100$. In that run the estimated rating of agent 4 is within 0.1 of a vertex value (-1 or 0) in 61% of replicates at $g = 300$ and in 94% at $g = 3000$; in the plotted run the same shares are 62% and 93%. Neither run puts any estimate within 0.05 of $-3/4$. Counting every replicate at or below -0.9 or at or above -0.1 gives 67% at $g = 300$ (45% and 22%), since 6% of replicates fall below -1.1.

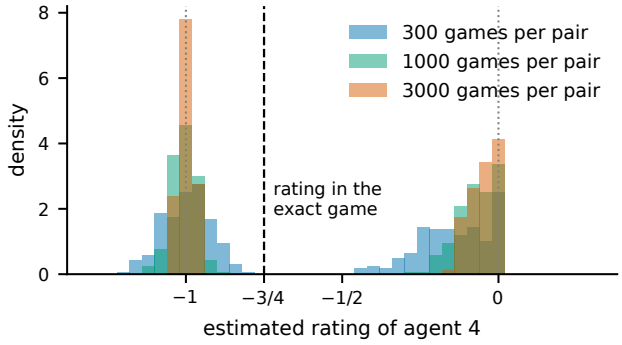


Figure 2: Estimated ratings of agent 4 of the five-agent game fall near the ratings at the two ends of the equilibrium set, -1 and 0, and not near $-3/4$. Each pairwise result is sampled from g logistic games; 400 replicates per g.

D.4 Where an equilibrium rating is continuous

Proposition 2 exhibits one matrix at which every rule that selects an exact equilibrium jumps. This part says where that happens in general, and what a rule must give up to avoid it. Write $\rho_i(A) = \max_{p\in\mathcal{P}(A)}(Ap)_i - \min_{p\in\mathcal{P}(A)}(Ap)_i$ for the range of the rating of agent i over the equilibrium set. Call a rating rule g continuous at A when $g(A_j) \to g(A)$ for every sequence $A_j \to A$ of antisymmetric matrices. A rule selects when it reports $(Ap)_i$ at some $p \in \mathcal{P}(A)$. Nash averaging, the LP vertex and deviation ratings all select; Bradley–Terry and α-Rank do not.

Tilting a matrix by a small multiple of B_c rewards the agents with large c_i, so an arbitrarily small perturbation can single out any chosen vertex of the equilibrium set.

Lemma 6 (Steering the equilibrium to a chosen vertex)**.** *Let A be antisymmetric, $c \in \mathbb{R}^n$, and $B_c = c\mathbf{1}^\top - \mathbf{1}c^\top$. Let $\varepsilon_j \downarrow 0$, let R_j be antisymmetric with entries at most 1 in absolute value, put $A_j = A + \varepsilon_j B_c + \varepsilon_j^2 R_j$, and take any $p_j \in \mathcal{P}(A_j)$. Then $c^\top p_j \geq \max_{q\in\mathcal{P}(A)} c^\top q - \varepsilon_j$, and every limit point of p_j maximizes $c^\top q$ over $\mathcal{P}(A)$.*

Proof. Fix $q \in \mathcal{P}(A)$. Since $A_j p_j \preceq 0$ and $q \succeq 0$, we have $q^\top A_j p_j \leq 0$, that is, $q^\top A p_j + \varepsilon_j q^\top B_c p_j + \varepsilon_j^2 q^\top R_j p_j \leq 0$. By antisymmetry the first term equals $-p_j^\top A q$, which is at least 0 because $Aq \preceq 0$ and $p_j \succeq 0$. In the third term $|q^\top R_j p_j| \leq 1$, so that term is at most ε_j^2 in absolute value. Hence $q^\top B_c p_j \leq \varepsilon_j$. For p, q in the simplex, $q^\top B_c p = (c^\top q)(\mathbf{1}^\top p) - (\mathbf{1}^\top q)(c^\top p) = c^\top q - c^\top p$. Maximizing over q gives the inequality. If $p_j \to \bar{p}$ along a subsequence, then $A\bar{p} \preceq 0$, so $\bar{p} \in \mathcal{P}(A)$, and the inequality passes to the limit. □

Brandl et al. (2016) prove the corresponding statement for maximal lotteries, in a form that already makes a chosen vertex a limit of unique equilibria of nearby games, and Brandl & Brandt (2026) give a new proof of it; the proof above covers an arbitrary real antisymmetric matrix. Proposition 2 is the case $c = \pm\mathbf{e}_5$: the matrix A_δ is $A_5 + \delta B_{\mathbf{e}_5}$ up to the sign convention.

Restatement (Continuity of an equilibrium rating)**.** *Let g be a rule that selects, and fix an agent i. Then g_i is continuous at A if and only if $\rho_i(A) = 0$. The forward implication uses only that g returns the equilibrium of every matrix whose equilibrium is unique. If h is any rating rule continuous at A, then for every neighbourhood V of A,*

$$\sup_{A'\in V,\ \mathcal{P}(A')\ \text{a point}} \left|h_i(A') - (A'p(A'))_i\right| \ \geq\ \rho_i(A)/2,$$

The constant $1/2$ cannot be replaced by anything larger.

Proof. Suppose $\rho_i(A) > 0$. The linear function $p \mapsto (Ap)_i$ attains its maximum over the polytope $\mathcal{P}(A)$ at a vertex v^+ and its minimum at a vertex v^-, and $(Av^+)_i - (Av^-)_i = \rho_i(A)$. A vertex of a polytope is an exposed point, so there is c^+ with $\arg\max_{q\in\mathcal{P}(A)} (c^+)^\top q = \{v^+\}$, and likewise c^-.

Matrices with a unique equilibrium are dense (Brandl, 2017), so for $\varepsilon_j = 1/j$ we may choose A_j^+ with a unique equilibrium and $\|A_j^+ - (A + \varepsilon_j B_{c^+})\|_{\max} \le \varepsilon_j^2$; then A_j^+ has the form required by Lemma 6. That lemma gives $p(A_j^+) \to v^+$, so $g_i(A_j^+) \to (Av^+)_i$, and the same construction with c^- gives $g_i(A_j^-) \to (Av^-)_i$. Both sequences converge to A and the two limits differ by $\rho_i(A) > 0$, so no value of $g_i(A)$ makes g_i continuous. The argument uses only the matrices $A_j^\pm$, each of which has a unique equilibrium, so the weaker hypothesis suffices.

Suppose instead $\rho_i(A) = 0$, so $(Ap)_i$ equals a constant κ on $\mathcal{P}(A)$ and $g_i(A) = \kappa$. Let $A_j \to A$ and write p_j for the equilibrium g selects at A_j. Any subsequence has a further subsequence along which $p_j \to \bar{p}$, and then $\bar{p} \in \mathcal{P}(A)$ as in Lemma 6, so $g_i(A_j) = (A_j p_j)_i \to (A\bar{p})_i = \kappa$. Every subsequence therefore has a further subsequence converging to κ. Hence $g_i(A_j) \to \kappa$. This direction is the upper hemicontinuity of the equilibrium correspondence (Bohnenblust et al., 1950) and holds for every selection rule.

For the bound, let h be continuous at A and take the two sequences above. Along them $h_i(A_j^\pm) \to h_i(A)$ while the equilibrium rating tends to $(Av^\pm)_i$, so the two limiting deviations are $|h_i(A) - (Av^+)_i|$ and $|h_i(A) - (Av^-)_i|$. Their sum is at least $\rho_i(A)$, so one of them is at least $\rho_i(A)/2$. Both sequences eventually lie in V and consist of matrices with a unique equilibrium. The constant cannot be raised. Let h report the fixed value $\frac{1}{2}\big((Av^+)_i + (Av^-)_i\big)$ at every matrix; this rule is continuous because it is constant. Upper hemicontinuity places $(A'p(A'))_i$ within ε of $[(Av^-)_i, (Av^+)_i]$ for every A' in a small enough neighbourhood, so the deviation of h there is at most $\rho_i(A)/2 + \varepsilon$. Continuity at A is what forces a deviation: a rule that reports the midpoint of the range at each A' separately agrees with the equilibrium rating wherever the equilibrium is unique, and is not continuous at A. □

By Theorem 4(a) every essential agent has $\rho_i(A) = 0$, so the discontinuities are confined to the inessential agents. These are also the agents whose ratings a copy can move. For agent 4 in A_5 the range is $\rho_4 = 1$ and the bound is $1/2$. The smallest matrix with a discontinuity has three agents. Let agents 1 and 2 draw and beat agent 3 at log-odds 1 and 2, so that $\mathcal{P} = \{(t, 1-t, 0)\}$ and agent 3 is rated $-2 + t$. Tilting the draw towards agent 1 makes $(1, 0, 0)$ the only equilibrium, where agent 3 is rated -1; tilting it towards agent 2 gives $(0, 1, 0)$ and rating -2. With two agents the rating is continuous everywhere.

E THE PRICE OF CONTINUITY

Appendix D.4 shows that a rule which returns the equilibrium rating whenever the equilibrium is unique cannot be continuous wherever an agent's rating varies over the equilibrium set, and that any continuous rule deviates from it by at least half that range. This appendix pairs that bound with a rule that pays close to it. Write $\mathcal{P}_b(A) = \{p \in \Delta_n : Ap \preceq b\mathbf{1}\}$ for the equilibrium conditions relaxed by a slack $b > 0$, $p_b(A)$ for the maximizer of H over $\mathcal{P}_b(A)$, unique because H is strictly concave and the set is compact and convex, and $g_b(A) = A\,p_b(A)$. Write $\rho_i^b(A)$ for the range of $(Ap)_i$ over $\mathcal{P}_b(A)$, so that $\rho_i^0 = \rho_i$.

Restatement (Theorem 5, the price of continuity)**.** *Fix $b > 0$. Write $\pi_i(h; A)$ for the largest difference between a rule h's rating of agent i and the equilibrium rating, over matrices with a unique equilibrium within δ of A, in the limit $\delta \downarrow 0$. (a) $\pi_i(h; A) \ge \rho_i(A)/2$ for every rule h continuous at A. (b) g_b is continuous everywhere and $\pi_i(g_b; A) \le \rho_i^b(A)$, which falls to $\rho_i(A)$ as b does.*

A third property makes the slack readable as a dial. Once b reaches $\max_i(A\mathbf{1})_i/n$, the uniform distribution is feasible and $g_b(A) = A\mathbf{1}/n$, the average result of each agent against the whole population, so the slack runs from Nash averaging at $b = 0$ to plain averaging.

Proof. Part (a) is the quantitative half of Appendix D.4, proved there for every rule continuous at A.

For the continuity in (b), let $A_j \to A$ and put $p_j = p_b(A_j)$. Along any subsequence with $p_j \to \bar{p}$, taking limits in $A_j p_j \preceq b\mathbf{1}$ gives $\bar{p} \in \mathcal{P}_b(A)$. Fix $p \in \mathcal{P}_b(A)$ and $s \in \mathcal{P}(A)$, which is nonempty, and put $\lambda_j = \min(1, \|A_j - A\|_{\max}/b)$ and $p^{(j)} = (1-\lambda_j)p + \lambda_j s$. Then $A_j p^{(j)} \preceq Ap^{(j)} + \|A_j - A\|_{\max}\mathbf{1} \preceq (1-\lambda_j)b\mathbf{1} + \|A_j - A\|_{\max}\mathbf{1} \preceq b\mathbf{1}$ as soon as $\|A_j - A\|_{\max} \le b$, so $p^{(j)} \in \mathcal{P}_b(A_j)$, and

$\lambda_j \to 0$ gives $p^{(j)} \to p$. Optimality of p_j gives $H(p_j) \geq H(p^{(j)})$, so $H(\bar{p}) \geq H(p)$ in the limit. As p was arbitrary, $\bar{p} = p_b(A)$, and since every subsequence has a further subsequence converging to that point, $p_b(A_j) \to p_b(A)$ and $g_b(A_j) \to g_b(A)$. The point s is a Slater point, since $As \preceq 0 \prec b\mathbf{1}$, and this is the whole difference between $b > 0$ and $b = 0$: at $b = 0$ no antisymmetric matrix has a strictly feasible point, since $p^\top Ap = 0$, so this argument breaks and whether the exact selection jumps is decided by $\rho_i(A)$ instead.

For the bound in (b), $\mathcal{P}(A) \subseteq \mathcal{P}_b(A)$ and $p_b(A) \in \mathcal{P}_b(A)$, so $(Ap)_i$ for $p \in \mathcal{P}(A)$ and $g_b(A)_i$ both lie in an interval of length $\rho_i^b(A)$. That bounds the difference at A; to bound it near A, note that $\|A' - A\|_{\max} \leq \delta$ gives $\mathcal{P}_b(A') \subseteq \mathcal{P}_{b+\delta}(A)$ and $|(A'p)_i - (Ap)_i| \leq \delta$, so $\rho_i^b(A') \leq \rho_i^{b+\delta}(A) + 2\delta$. Applying the bound at A' and letting $\delta \downarrow 0$ gives $\pi_i(g_b; A) \leq \rho_i^b(A)$, since $\rho_i^{b+\delta}(A) \downarrow \rho_i^b(A)$ by the next paragraph. For the limit, $\mathcal{P}_b(A)$ is nondecreasing in b and $\bigcap_{b>0} \mathcal{P}_b(A) = \mathcal{P}(A)$ because $Ap \preceq b\mathbf{1}$ for every $b > 0$ forces $Ap \preceq 0$. If the maximum of $(Ap)_i$ over $\mathcal{P}_b(A)$ stayed a fixed amount above its maximum over $\mathcal{P}(A)$ for every b, the maximizers would have a limit point in the intersection with the same property, which is impossible; the minimum is the same argument, so $\rho_i^b(A) \downarrow \rho_i(A)$.

For the remark, the uniform distribution $\mathbf{1}/n$ is the unique maximizer of H over Δ_n, and it lies in $\mathcal{P}_b(A)$ exactly when $(A\mathbf{1})_i/n \leq b$ for every i. When it does, it is also the maximizer over the smaller set, so $g_b(A) = A\mathbf{1}/n$, the average result of each agent against the whole population. The threshold is nonnegative because $\mathbf{1}^\top A\mathbf{1} = 0$. □

The bound of part (b) is stated with ρ_i^b rather than with a multiple of b because no such multiple exists.

Proposition 5 (The price does not fall at a uniform rate)**.** *Let A_3^δ be the three-agent matrix with $A_{12} = \delta$, $A_{13} = 1$ and $A_{23} = 2$, for $\delta > 0$. Its equilibrium is unique and $\rho_3 = 0$, while*

$$\rho_3^b(A_3^\delta) = \min\left(1, \frac{b}{\delta}\right) + \min\left(b, \frac{b+\delta}{2+\delta}, 1\right).$$

Proof. The constraints on $\mathcal{P}_b$ are $\delta p_2 + p_3 \leq b$, $-\delta p_1 + 2p_3 \leq b$ and $-p_1 - 2p_2 \leq b$, the last always true, and agent 3 is rated $-p_1 - 2p_2 = -1 - p_2 + p_3$. Minimizing that rating takes $p_3 = 0$, and the first constraint then allows $p_2 \leq \min(1, b/\delta)$. Maximizing it takes $p_2 = 0$ and $p_1 = 1 - p_3$, where the first constraint gives $p_3 \leq b$, the second gives $p_3 \leq (b+\delta)/(2+\delta)$ and the simplex gives $p_3 \leq 1$. At $b = 0$ the first constraint forces $p_2 = p_3 = 0$, so the equilibrium is $\mathbf{e}_1$ and $\rho_3 = 0$. □

For $b \leq \delta$ the first term is b/δ while $\rho_3 = 0$, so no bound of the form $\rho_i^b \leq \rho_i + Cb$ holds with a C that does not depend on the matrix. Theorem 5(a) forces this. Take A with $\rho_i(A) > 0$. Parts (a) and (b) give matrices A' with a unique equilibrium arbitrarily near A at which $\rho_i^b(A')$ is at least $\rho_i(A)/2$ up to any ε, while $\rho_i(A') = 0$; a bound $\rho_i^b \leq \rho_i + Cb$ with a fixed C would then give $Cb \geq \rho_i(A)/2 - \varepsilon$ for every $b > 0$.

Removing near copies within a threshold, the other repair Appendix J tests, trades on a different axis. When no two of the original agents are already within the threshold, the agent it removes is the copy, which carries the largest index, so ratings of the original agents are exactly unchanged by any copy inside the threshold. It is not continuous. Append to A_5 an agent that scores $\eta > 0$ against each of the five. The six-agent game has a unique equilibrium, which rates agent 4 at $-\eta$, while for $b < 3/4$, which keeps the five original agents outside the threshold, the rule reports $-3/4$ for $\eta \leq b$ and $-\eta$ above it, a jump of $3/4 - b$. It therefore buys invariance to near copies and not continuity, and Theorem 5(a) says nothing at the threshold, where it is not continuous.

F ENGINE LEAGUE

The engine-league finding of Section 4.3 is that an inserted exact copy makes the estimated equilibrium set a segment on which the two selection rules disagree. The parts below specify the instrument, the calibration of the top band, the engine league and the three bootstraps behind the quoted gaps and standard errors. Table 8 lists the sixteen settings.

F.1 INSTRUMENT

Twenty-one settings of the open-source engine Stockfish (versions 18 and 17.1) are arranged as a ladder of seven strength bands. Each band has three weakening sources: Stockfish 18 at a small node budget, Stockfish 18 under its UCI rating limit (which plays stochastically), and Stockfish 17.1 at a node budget. Every setting runs single-threaded with a 16 MB hash, the network shipped with its binary and skill level 20. Rating-limited settings set the UCI strength limit at the band's rating (Table 8). All games use one shared pool of 214 random opening prefixes (eight random legal plies, kept when the engine's evaluation at 2^{12} nodes is within 150 centipawns of equality). Prefix o serves game o of every pair with both color assignments, so the 300 games of a pair use 150 prefixes. A game whose move sequence repeats an earlier one of the same pair and colors is replayed from a reserve prefix. A game ends by the rules of chess (checkmate, stalemate, insufficient material, threefold repetition or the fifty-move rule, claimed automatically) or is scored a draw at 400 plies. A node-limited setting is deterministic: from a given opening and color it plays the same moves every time. Two settings with identical engine, version and node budget therefore have identical population outcome probabilities against every opponent. This determinism makes the inserted copy exact at the population level. Because each prefix is played with both color assignments and the two settings are identical, the two games of a prefix repeat the same moves with the colors exchanged, so a decisive prefix gives one win to each side and the 300 games of the copy pair score exactly 150 to 150. Only the rating-limited settings inject randomness. Outcomes are scored 1, $1/2$, 0 and antisymmetrized as $W - W^{\top}$.

Table 8: The engine league has sixteen settings, one an exact copy of the strongest. Columns: engine version, node budget per move and the UCI rating limit for the rating-limited settings. Bands number the ladder from the weakest (1) to the strongest (7); the top band is calibrated separately and sits above band 7.

setting	engine	nodes per move	rating limit
band 4, node-limited	Stockfish 18	512	–
band 4, rating-limited	Stockfish 18	4 096	2600
band 4, older version	Stockfish 17.1	1 024	–
band 5, node-limited	Stockfish 18	4 096	–
band 5, rating-limited	Stockfish 18	32 768	2800
band 5, older version	Stockfish 17.1	8 192	–
band 6, node-limited	Stockfish 18	16 384	–
band 6, rating-limited	Stockfish 18	262 144	3000
band 6, older version	Stockfish 17.1	32 768	–
band 7, node-limited	Stockfish 18	65 536	–
band 7, rating-limited	Stockfish 18	262 144	3190
band 7, older version	Stockfish 17.1	131 072	–
top band, strongest	Stockfish 18	262 144	–
top band, older version	Stockfish 17.1	262 144	–
top band, $2^{17.5}$	Stockfish 18	185 364	–
exact copy of the strongest	Stockfish 18	262 144	–

F.2 CALIBRATION OF A MATCHED TOP BAND

Real populations are nearly transitive at the top and cyclic in a middle band (Czarnecki et al., 2020), and therefore the league calibrates a matched top band rather than taking the ladder as found. Two calibration rounds (100 and 60 games per pair) select the strongest setting (the newer engine at 2^{18} nodes), the older engine at 2^{18} and the newer at $2^{17.5}$. Their pairwise scores are 0.525, 0.542 and 0.542. The engine league consists of this band, an exact copy of the strongest setting and the four upper bands of the ladder: 16 settings, 120 pairs, 300 games per pair.

F.3 FORMAL LEAGUE

The copy is exact in the engine configuration and hence in the latent outcome matrix. In the estimated matrix the row of the copy equals that of the original against the eleven deterministic opponents (the original among them) and differs against the four stochastic ones by independent sampling. The league has 300 games per pair and 36 000 games in total, played on two machines

with the same opening pool and color balance as the calibration rounds. Every pair has exactly 300 completed games. Interrupted games are requeued and only completed games are recorded.

F.4 Solved quantities and resolution

The solved quantities in Section 4.3 come from the antisymmetrized score matrix of this league. The ε-equilibrium set is $\{p \in \Delta_n : Ap \preceq \varepsilon \mathbf{1}\}$ with ε measured in the entry bound $b = 1/\sqrt{300} = 0.058$. b is the no-draw upper bound on the standard error of an entry of $A = W - W^\top = 2W - 1$ from 300 games (a match score W_{ij} has half that bound). The set is computed at $b/2$, b and $2b$ (0.029, 0.058, 0.115). The observed draw rates (22% over the league, 68% inside the top band) put the true standard error of a top-band entry near $0.56\,b$; hence b is about 1.8 true standard errors of an entry. The ε-equilibrium set at $\varepsilon = b$ has top-band widths 1.0, 0.54 and 0.60, against 0.20 when the score matrix is shuffled. At $b/2$ the width criterion also holds (widths 1.0, 0.27, 0.30). At $2b$ all four top settings reach width 1.0.

F.5 Three bootstraps

Each of three bootstraps over match outcomes (200 replicates) resamples the recorded outcomes within every pair independently and answers one question. The width bootstrap resamples the 300 games of every pair, the copy head-to-head included, and recomputes the ε-widths at $\varepsilon = b$. It asks what the data resolve. The rating bootstrap resamples every pair but sets the copy's row and column equal to the original's in every replicate. It asks what the selection rules do on an exact copy. The unconditional bootstrap resamples every pair, the head-to-head included, and recomputes the exact equilibrium set and both selection rules. It asks how often the exact draw is stable under resampling of single games.

F.6 Width bootstrap

The width bootstrap gives 5th to 95th percentile intervals for the top-band widths at $\varepsilon = b$. They are $[1.0, 1.0]$ for the copy pair, $[0.38, 0.69]$ (median 0.50) for the older engine at 2^{18} nodes and $[0.41, 0.73]$ (median 0.54) for the setting at $2^{17.5}$. On every replicate the width criterion holds.

F.7 Rating bootstrap

The rating bootstrap gives the 10^{-8} rating gap quoted in Section 4.3. Between the LP vertex and the maximum-entropy equilibrium, the largest rating difference of a setting outside the support is then 1.1×10^{-8} (median over 200 replicates, maximum 1.4×10^{-7}). With the two rows as played the difference is 0.033, which equals half of the largest entry difference between the two rows (0.067, attained against a rating-limited opponent).

F.8 Unconditional bootstrap

Resampling the head-to-head breaks the exact draw by sampling noise on almost every replicate. On 96.5% of replicates the exact equilibrium set is then a single point and the maximum-entropy equilibrium collapses onto one copy (median ℓ_1 distance 1.0 from the point estimate with the exact draw, 90th percentile 1.0). Then the LP vertex and the maximum-entropy equilibrium have the same support. On the remaining 3.5% the resampled head-to-head is again an exact draw, the draw holds, and the maximum-entropy support keeps both copies. The 300 recorded games are 204 draws and 48 wins each way, and a resample averages exactly one half with probability 0.041. Disagreement between the two selection rules on the matrix with the exact draw is a property of that draw, which resampling single games does not preserve. The ε-widths of the width bootstrap are stable under resampling.

G Published Atari equilibrium and the Atari score table

Two Atari items are recorded here: the equilibrium printed in the supplement of the original paper, and the later Atari score table, with its own roster, on which the check of Section 4.3 is run. In the

section reporting this evaluation, Balduzzi et al. (2018) say that they find a Nash equilibrium with a linear-programming solver, and that finding the maximum-entropy Nash should be possible with an algorithm they cite. The supplement of the original paper prints the equilibrium found by that solver for 20 agents and 54 environments as two page images. Its support is five agents and five environments, and the remaining 15 agents and 49 environments are outside its support.

The Atari score table used for the check in Section 4.3 is the normalized amalgamation in the appendix of the deviation-ratings preprint (20 agents, 53 tasks, each normalized to $[0, 1]$). We extract the table from the text layer of the appendix of the preprint (Table 3 of arXiv:2502.11645v1) and check every entry against the LaTeX source of that preprint: all 53 task rows and all 1060 entries agree. The maximum of every task row is exactly 1.000, as the normalization requires. The equilibrium is solved by the same convex program as Appendix A.6 (CVXPY/Clarabel, slack 10^{-9}), with the value from a HiGHS linear program. On that table the common value $c^\star$ of $\min_j (S^\top p_a)_j$ and $\max_i (S p_t)_i$ at equilibrium is 0.415. Under both selection rules the support is four agents and four tasks, with identical probabilities to 2×10^{-7}, and every coordinate of the equilibrium set has range below 5×10^{-8}. The entries of the table are three-decimal rationals, so this verdict also has an exact certificate. On the numerical support the two linear systems that hold every supported skill and difficulty at the common value have unique rational solutions, which agree on $c^\star = 144807281707/348596153500$; every agent outside the support falls at least 0.112 below that value and every task outside it stands at least 0.020 above, so no equilibrium can give them mass. The equilibrium is therefore unique in exact arithmetic. Every one of the 73 single duplications reproduces the even split to 10^{-5}.

H Public arena and leaderboard tables

Section 4.3 reports one equilibrium for the human-preference leaderboard at the exact level. The parts below specify the leaderboard construction, its controls, the near-copy check and the subsampled pools. Tables 9, 10 and 11 list the leaderboard variants, the further arena tables and the copy effects on those pools.

H.1 Data

The dataset is the 100 000-battle human-preference release of the LMSYS arena (Hugging Face dataset lmarena-ai/arena-human-preference-100k; 106 134 battles, columns model_a, model_b, winner, anonymity flag). Each battle contributes one score to the unordered pair: a win 1, a tie $1/2$, a loss 0 for the first-named model. Our primary analysis keeps anonymous battles only.

H.2 Model set

The model set is a subset in which every pair has at least $N_{\min}$ battles. It is found by greedy peeling: remove the model with the fewest qualifying pairs (ties broken by name) until the remaining set is complete. Below, the $N_{\min}$ variants are the sensitivity check. This leaves 28 models at $N_{\min} = 30$, 31 at $N_{\min} = 15$ and 22 at $N_{\min} = 60$. The all-battles variant coincides with the primary analysis because the release contains no non-anonymous battles between these models.

H.3 Matrix and resolution

The matrix is $A = W - W^\top$ with W_{ij} the mean score of i against j. The entry bound b is the median over pairs of $1/\sqrt{n_{ij}}$, where n_{ij} is the number of battles recorded for the pair; this is the no-draw upper bound on the standard error of an entry of $A = 2W - 1$. It is 0.089 for the primary set, whose median pair has about 125 battles. The ε-equilibrium set is $\{p \in \Delta_n : Ap \preceq \varepsilon \mathbf{1}\}$ at $b/2$, b and $2b$. Two further sets replace the median by the bound $s_{ij} = 1/\sqrt{n_{ij}}$ of each pair: the row-bound set $\{p \in \Delta_n : (Ap)_i \le \max_j s_{ij}\}$ and the weighted set $\{p \in \Delta_n : (Ap)_i \le (\Sigma p)_i\}$ with $\Sigma = (s_{ij})$ and $s_{ii} = 0$. The weighted set contains the exact equilibrium set of every matrix within the entry bounds.

H.4 SELECTION RULES

The LP vertex on that set is the point returned by the zero-objective linear program of Appendix B.11. The ε-maximum-entropy equilibrium maximizes entropy over that set.

H.5 SHUFFLE AND BOOTSTRAP CONTROLS

The shuffle control permutes the upper-triangle entries of the score-difference matrix and mirrors them twenty times (structure destroyed, entry multiset kept). It then recomputes the ε-widths and the rating gap, the largest rating difference between the LP vertex and the ε-maximum-entropy equilibrium. At $\varepsilon = b$ the shuffled matrices have all 28 coordinates above 0.2 with median ε-width 0.43, against the observed 24 coordinates and 0.39. Their rating gap is $2.3\,b$ (median; $3.0\,b$ at the 90th percentile) against the observed $1.15\,b$. Neither the widths nor the gap is therefore evidence of ties. A bootstrap over battles (200 replicates) puts the number of coordinates above 0.2 between 24 and 26 and the largest width between 0.87 and 1.0 (5th to 95th percentiles).

H.6 PER-PAIR BOUNDS

The row bounds $\max_j s_{ij}$ run from $0.94\,b$ to $2.1\,b$ over the 28 models. On the row-bound set all 28 coordinates have ε-width above 0.2, the median ε-width is 0.49 and the rating gap is $1.3\,b$. The same twenty shuffled matrices (entries permuted together with their battle counts) have all 28 coordinates above 0.2 in every case, median widths from 0.55 to 0.64 and rating gaps from $1.7\,b$ to $3.6\,b$ (median $2.7\,b$). The weighted set is the single point on the strict winner w: every margin A_{wj} of the winner exceeds its own pair's bound s_{wj} by at least 0.032. The row-w inequality then forces $p_j = 0$ for every $j \neq w$. On the shuffled matrices the weighted set has all 28 coordinates above 0.2 (median ε-widths 0.39 to 0.47, rating gaps $1.1\,b$ to $3.5\,b$). Under the common bound and under both per-pair bounds the observed set is no wider than that of a shuffled matrix.

Table 9 lists the variants: the numerical width of each is at the solver slack 10^{-9}, its width counts and rating differences are at $\varepsilon = b$ unless stated, and the rating differences are the rating gap in units of b.

Table 9: Every leaderboard variant is a single point at the exact level. The all-battles variant coincides with the primary row; the comparison with a shuffled matrix is in Appendix H.5.

variant	models	battles	numerical width	models with width ≥ 0.2	rating difference (b)
primary ($N_{\min} = 30$)	28	60 253	10^{-8}	24	1.15
primary at $\varepsilon = b/2$	28	60 253	10^{-8}	13	0.66
$N_{\min} = 15$	31	54 968	2×10^{-8}	31	0.72
$N_{\min} = 60$	22	47 116	10^{-8}	20	0.9

H.7 NEAR COPIES ON THE TWO EMPIRICAL MATRICES

The near-copy question of Section 3 is answered on the score matrix of the league (16 settings) and the leaderboard matrix (28 models). For every agent k, every $\eta \in \{0, 10^{-3}, 3 \cdot 10^{-3}, 10^{-2}, 3 \cdot 10^{-2}, 0.1\}$ and five noise seeds, a copy is appended whose row is the original's plus η times a random sign vector. There are two modes: the copy draws its original, or the copy-versus-original entry is perturbed by η as well. The maximum-entropy equilibrium q of the near-copy game is recomputed and two quantities are recorded: the change of the total probability of agent k and the largest change of an original agent's rating $(A_\eta q)_i$ $(i \leq n)$ in the near-copy game. The rating in the near-copy game includes a direct term, the result against the perturbed copy. Folding the copy into its original first, $A\,\pi(q)$, drops that direct term and understates the change. The rating in the near-copy game is the one reported.

Our comparison scale is the resampling standard error of the ratings. For the league it comes from 200 outcome bootstraps (largest per-setting standard error 0.036, median 0.021). For the leaderboard it comes from a parametric bootstrap of 200 replicates that redraws every pair's score from a binomial at the median pair count (largest 0.10, median 0.080).

On the leaderboard the equilibrium set is a single point with one essential model. The maximum-entropy equilibrium is then the point mass on that model, the strict winner w. Every other model's rating is its score difference against the winner alone, $r_i = A_{iw} = 2W_{iw} - 1$. For η below the strict winner's smallest margin (0.096 here, hence for the tested values up to 0.03), a near copy of that model with a drawn head-to-head changes the other ratings by exactly $\eta/2$ (the copy takes half the mass and each opponent meets its perturbed row half the time). With a perturbed head-to-head the change is η (all mass moves to the winner of the pair). Copies of the other models change nothing while the strict winner stays a strict winner. Up to $\eta = 0.03$ no change exceeds the median standard error. At $\eta = 0.1$ the perturbation exceeds that margin, so the copy can lose its strict win against the closest model and the even split no longer holds. The largest change is then 0.12 to 0.15 (about one entry standard error on the score scale) and 2.1% of cases exceed the largest standard error.

The estimated matrix of the league contains the exact draw between the copy pair. There an exact copy ($\eta = 0$) of a member of the drawn pair moves its total probability by 0.167 (three-way split) and the other settings' ratings by at most $1/90$. That bound is exact: the equilibrium is supported on the pair, so every other rating is a combination of the pair's two rows, the even split becomes a two-thirds share for the duplicated member, and each rating moves by one sixth of the difference between the two rows. None of the 90 strictly ordered pairs of the other settings reverses, their smallest separation being $1/120$. The pair's estimated rows differ by at most 0.067, six times the $1/90$ bound. Up to $\eta = 0.03$ the largest rating change grows from 0.033 to 0.049 (drawn head-to-head) or 0.064 (perturbed head-to-head). It exceeds the largest standard error on at most 2.6% of cases. At $\eta = 0.1$ the largest change is 0.083 or 0.21 and 7.5% to 12.5% of cases exceed the largest standard error.

The tilt and the direct contribution separate. The tilt contribution is the change of $A\,\pi(q)$. It is zero on the leaderboard at every η up to 0.03 and at most $1/30$ on the league, where the total probability of agent k moves to one member of the pair that draws in the estimated matrix. The pair's columns differ by up to 0.067. The direct contribution is the remainder, of the order of η.

H.8 Further arena tables

Under the protocol above, this part checks the four further arena tables of Section 4.3: the search arena (24k battles), the web-development arena (10k), the repository-chat arena (4k) and the overlapping 140k release of the same LMSYS arena. We check the expert arena (5k) separately. The releases are the Hugging Face datasets lmarena-ai/search-arena-24k, webdev-arena-preference-10k, repochat-arena-preference-4k, arena-expert-5k and arena-human-preference-140k. We use anonymous battles only where a release has that flag (the web-development and repository-chat tables). The other releases have none. Our model set is the subset found by greedy peeling in which every pair has at least 30 battles (15 as a sensitivity check), as for the leaderboard. A win, a tie and a loss score 1, $1/2$ and 0. Votes of the 140k release that call both answers bad count as ties. Table 10 lists the outcomes; a dash there marks a pool with a unique equilibrium, where no duplication was run. At 15 and 30 battles per pair every table with a complete pool is a single point with one essential model, the strict winner of its pool, and the two selection rules coincide. The two thresholds need not peel to nested pools: on the repository-chat arena both give four models, but not the same four, so the looser threshold there leaves fewer battles inside the pool. In the expert arena there are too few battles for a complete pool at those thresholds. At 5 battles per pair its pool of six models has more than one equilibrium in exact arithmetic (entries are rational, score over battle count). Two models split their 11 battles exactly, and no other model beats any mixture of the two on a segment. The equilibrium set is that segment on them, from $(5/11, 6/11)$ to $(5/14, 9/14)$, with width $15/154$. Its maximum-entropy equilibrium is the first endpoint, held there by a third model rated exactly zero. The maximum-entropy equilibrium already gives the first essential model its largest probability over the set, so duplicating that model changes nothing. Duplicating the second moves the equilibrium to the other endpoint and changes the other models' ratings by up to $1/14$. Two of them change order. At 3 battles per pair the pool grows to twelve models and is a single point again. Here the mechanism is the same as in the football tables, coarse scores from few outcomes. With the pool held fixed the equilibrium is not unique in 14% to 16% of replicates at 50 battles per pair ("Subsampling a fixed pool" below) and is unique at the recorded counts.

Table 10: Four further arena tables have a unique equilibrium at 15 and 30 battles per pair, and the expert arena is a segment at 5. n: models; battles: within the pool; e: essential models; width: largest range of one model's probability; gap: largest rating difference between the two rules; copy: largest rating change of another model under one copy.

table	min. battles per pair	n	battles	e	width	gap	copy
search arena	30	11	10 653	1	0	0	–
search arena	15	12	11 604	1	0	0	–
web-development arena	30	10	9 452	1	0	0	–
web-development arena	15	11	10 241	1	0	0	–
repository-chat arena	30	4	1 575	1	0	0	–
repository-chat arena	15	4	1 317	1	0	0	–
general arena, 140k release	30	37	105 474	1	0	0	–
general arena, 140k release	15	44	122 533	1	0	0	–
expert arena	5	6	101	2	15/154	0.071	0.071
expert arena	3	12	343	3	0	0	–

H.9 Changes against each agent's own uncertainty

The near-copy comparison of Appendix H.7 uses the largest standard error across agents; this part repeats it against each agent's own standard error, with the leaderboard's actual per-pair battle counts (30 to 1699) in the bootstrap instead of a common median count. Per-agent standard errors of the ratings at the maximum-entropy equilibrium range from 0.003 to 0.106 on the leaderboard and from 0.003 to 0.038 on the engine league. On the leaderboard no agent's change exceeds its own standard error for η up to 0.03, and at $\eta = 0.1$ some agent does so in 1.4% of the 140 cases. Those cases also exceed the largest per-agent error of this part. On the engine league the per-agent convention flags more cases. Some agent's change exceeds its own standard error in 7.5% of cases at $\eta = 0.01$ (2.5% against this part's largest per-agent error, 0.038), 8.8% at 0.03 (5%) and 10% at 0.1 (10%). The largest ratio of change to own error rises from 1.5 to 3.8 to 23. Those flagged are the best-estimated settings, whose ratings move by a few hundredths on a standard error of a few thousandths. On the league, a change hidden by the largest standard error can still exceed the standard error of the setting that moves.

H.10 Subsampling a fixed pool

This part supports the subsampled-pool statement of Section 4.3. The battle thresholds above change the pool together with the coverage. This part instead fixes the pool and subsamples the battles. It fixes the 28-model leaderboard pool and the 37-model pool of the 140k release at 30 battles per pair. For each pair it subsamples the recorded battles without replacement to $g \in \{5, 10, 20, 50\}$ battles. Pairs with fewer battles keep all of them. It then rebuilds the score matrix and analyses it as above, with 50 replicates per g. On the leaderboard pool the equilibrium is not unique in 34%, 26%, 20% and 14% of the replicates at $g = 5$, 10, 20 and 50, with median essential support 7, 6, 5 and 2. On the 140k pool the fractions are 14%, 26%, 34% and 16% with median support 9, 6, 3 and 3. At the full battle counts both pools are single points with one essential model. Before the run we expected the replicates with more than one equilibrium to disappear by $g = 50$. They do not, and the fraction at $g = 50$ is 14% and 16%. With the pool held fixed, sampling alone produces more than one equilibrium in 14% to 34% of replicates up to fifty battles per pair. As the battles accumulate the essential support shrinks toward the strict winner.

H.11 Copies on the subsampled tables

This part supports the same statement, copy effects inside sampling noise. More than one equilibrium alone does not say whether a copy changes anything. Hence the subsampling replicates above that have more than one equilibrium are duplicated next. On every replicate with more than one equilibrium each essential model is duplicated and the duplicated table is rated by the maximum-entropy equilibrium. A parametric bootstrap of the subsampled table (each pair's score redrawn from its battle count, 50 resamples) gives each model's standard error on that replicate. Table 11 reports both pools. On the leaderboard pool a copy changes a rating on 23% to 36% of the duplications and reverses a strictly ordered pair about as often, at every sample size. On the 140k pool

both fractions are larger at $g = 5$, 20 and 50 and smaller at $g = 10$, and at $g = 50$ they still stand at 50% and 55%. A change exceeds the affected model's own standard error in at most 2% of the duplications at any g on the leaderboard pool and in at most 7.3% (at $g = 20$) on the 140k pool.

Table 11: Copies move ratings and reorder models on both subsampled pools, and on the 140k pool the fractions at $g = 50$ are still 50% and 55%. Every essential model is duplicated on every replicate with more than one equilibrium. moved: some other rating changes by more than 0.01; reversed: some strictly ordered pair reverses; last two columns on the score scale.

pool	g	duplications	moved	reversed	max change	median margin
leaderboard	5	89	36%	35%	0.11	0.015
leaderboard	10	58	24%	28%	0.12	0.010
leaderboard	20	26	23%	27%	0.12	0.025
leaderboard	50	16	25%	25%	0.06	0.020
140k release	5	36	42%	56%	0.09	0.013
140k release	10	51	22%	24%	0.10	0.014
140k release	20	41	46%	46%	0.13	0.025
140k release	50	20	50%	55%	0.10	0.013

Before the run we expected copies to change ratings on most replicates with more than one equilibrium, by amounts of the order of the score resolution and mostly below the standard error at $g = 5$ and 10, and few duplications to change anything at $g = 50$. In the outcome the changes are below the standard error in all but a few percent of the duplications at every g, and a quarter (leaderboard) to a half (140k) of the duplications still change a rating at $g = 50$. On subsampled chatbot-arena tables copies do move ratings and reorder models, but the moves stay inside sampling noise in all but a few percent of duplications. The football tables differ because their records are complete.

H.12 Calibration of the bootstrap standard error

This part supports the noise scale used throughout Section 4.3. The comparisons above use the bootstrap standard error of the observed table as the noise scale, and the estimated rating is not a regular estimator near an exact draw (Appendix D), so this part checks that scale on populations with known ties. Three populations are built from the leaderboard's full score matrix: the matrix itself (a single point), the matrix with an exact copy of the top model inserted (two identical models) and the matrix with the top model's row replaced by the average of the top two (two nearly identical models). From each, 200 tables are sampled at $g \in \{5, 10, 20, 50\}$ battles per pair. On each table we compute the ratings at the maximum-entropy equilibrium, their bootstrap standard errors (50 resamples) and the paired change of every original model's rating under a copy of the model with the largest weight. The bootstrap standard error tracks the true sampling standard deviation across the 200 tables: the median ratio over models is between 0.94 and 1.04 on every population and g, and every model's ratio lies between 0.66 and 1.25, including the neighbours of the identical and nearly identical models. The paired change exceeds the affected model's true sampling standard deviation on at most 2% of tables on every population and g (largest change 0.18, against sampling standard deviations of 0.08 to 0.19). We had predicted that the bootstrap would understate the spread for the neighbours of the identical and nearly identical models by a factor above two. It does not. Thus the noise scale of the subsampling comparisons is calibrated on these populations, and their conclusion stands against the true sampling spread as well. The three populations above tie identical rows. The same check is repeated on the population A_δ of Proposition 2 at $\delta \in \{0, 0.01, 0.1\}$ and $g \in \{300, 1000, 3000\}$ games per pair with 200 sampled tables per cell. Its tied agents have different rows and its estimated rating of agent 4 is bimodal (Figure 2). At $\delta \leq 0.01$ the true sampling standard deviation of agent 4's rating is 0.43 to 0.48 and the median bootstrap standard error is 0.36 to 0.41, a ratio of 0.82 to 0.89. In 12% to 17% of tables the bootstrap error is below half the true spread. At $\delta = 0.1$ the ratio is 0.84 at $g = 300$ and 1.1 to 1.3 beyond. Near a tie between agents with different rows the bootstrap understates the spread by about a fifth at the median and by more than half on one table in seven, less than the factor of two predicted before the run. An exact copy of the agent that draws everyone changes no rating on any sampled table at any δ or g: sampling breaks the tie, the estimated set is a single point, and the copy is invariant there while the estimated rating itself is off by up to 1.

H.13 Other complete pools

This part tests how far the verdicts of Section 4.3 depend on which complete pool the peeling lands on: every alternative pool of the leaderboard release is a single point, while 13 of 222 alternative pools of the 140k release are segments on one split pair. The pools above come from one greedy peeling with ties broken by name. The peeling is rerun with random tie-breaking in ten independent streams of twenty peels each per table and threshold, on the leaderboard and 140k releases at 30 and 15 battles per pair. Ten peels of every stream start from a random subset of four fifths of the models. The reruns give 106 and 103 (leaderboard) and 121 and 101 (140k) complete pools of 19 to 44 models. Every one of the 209 pools of the leaderboard release is a single point. Of these, 197 have one essential model, the pool's strict winner. The other twelve, all without the top models, are single points as well: eleven with a unique mixed equilibrium on three models and one with a single essential model that is not a strict winner. On the 140k release 13 of the 222 pools are segments, between zero and two per stream at each threshold, and the rest are single points (177 strict winners, 31 pools with a unique mixed equilibrium on three models and one with a single essential model that is not a strict winner). Every segment lies on the same pair, chatgpt-4o-latest-20250326 and deepseek-r1-0528. They split their 143 battles exactly. These segments come in two widths. In eight pools the width is exactly $728/2081$ (0.350) and the maximum-entropy equilibrium gives the pair weights $(0.650, 0.350)$, the second at the top of its range. The maximum-entropy equilibrium already gives deepseek-r1-0528 its largest probability, so duplicating it moves nothing. Duplicating chatgpt-4o-latest-20250326 moves the pair to $(2/3, 1/3)$, changes the ratings of other models by at most 0.005 and reverses one to three pairs. In five pools the width is exactly $225/368$ (0.611) and the maximum-entropy equilibrium splits the pair evenly. Duplicating either model moves other ratings by up to 0.028 (chatgpt-4o-latest-20250326) or 0.042 (deepseek-r1-0528) and reverses 8 to 25 pairs. The bootstrap standard errors of the ratings in these pools are 0.06 to 0.08 at the median and 0.09 to 0.14 at most, and the largest copy effect is about half a standard error. Every pool whose verdict is not a strict winner (57 pools) is enumerated in exact rational arithmetic from the battle counts, restricted to the numerical candidate support with every inequality kept. The 13 segments have essential support two and dimension one. The three-model pools have dimension zero and $\operatorname{rank}(A_{\mathcal{E}}) = 2$, the odd support of a single point required by the dimension identity of Section 2. Every one of the 1584 models excluded from these enumerations has an enumerated vertex q with $(Aq)_i < 0$ in rational arithmetic, which certifies that model i is inessential. Hence the numerical candidate supports are the exact supports. A strict winner is itself an exact certificate of a single point, so every pool of this part is decided exactly. Before the run we expected every complete pool to be a single point with a strict winner. The answer depends on the peeling path. About six percent of the 140k pools are segments, all on one pair that split its series, the football mechanism at over a hundred battles per pair.

I Football league tables

Section 4.3 reports that unmodified league seasons have more than one equilibrium in most years. In most of those seasons the choice of selection rule, a copy, or both change reported ratings. Tables 12 and 13 give the duplications under the criterion and the season list.

I.1 Round-robin league tables

These tables are the pairwise score matrices of the top divisions of England, Germany, Spain, Italy and France over the nine years 2015-16 to 2023-24, 44 league seasons in all. They are built from the openfootball match lists (GitHub repositories openfootball/england, deutschland, espana, italy and france, master branches as of 2026-09-06). The French 2019-20 season was abandoned after 279 of 380 matches and is excluded. Every pair of teams plays twice, once at each ground. A win, a draw and a loss score 1, $1/2$ and 0. Both games are averaged, and the matrix is $A = W - W^{\top}$ as for the leaderboard. Every entry is therefore a multiple of $1/2$ and the table is an exact rational matrix. Awarded matches are counted at their recorded score. For each season the width and the essential support are computed by linear programming as in Section 4.1. The vertices of the equilibrium set are enumerated in exact rational arithmetic with cddlib, with the variables restricted to the essential support and every team's inequality retained. Restricting the variables certifies the enumerated set only if the excluded teams receive no mass in any equilibrium. If some equilibrium q has $(Aq)_i < 0$,

then team i gets no mass in any equilibrium p, because $p^\top Aq = 0$ and every term of that sum is nonpositive. An enumerated vertex q with $(Aq)_i < 0$ in rational arithmetic is then an exact witness that team i is inessential. Every one of the 689 excluded teams across the 44 seasons has such a witness. The numerical support is the exact support, and the restricted enumeration is the full equilibrium set in every season. The enumeration gives the exact width and dimension. The two computations agree on every season, and no season has a width between 0 and 10^{-4}. Table 13 lists the seasons, with a dash where the equilibrium is unique and no duplication was run.

Forty of the 44 seasons have more than one equilibrium at the exact level, with widths from $2/41$ to 1 and equilibrium sets of dimension one (24 seasons), two (10), three (5) or four (1) supported on two to eight top teams. The source of the width is the scoring rule. A split series, one win each or two draws, is an exact zero entry. With entries in $\{0, \pm 1/2, \pm 1\}$ the rank condition $\operatorname{rank}(A_{\mathcal{E}}) \leq e - 2$ of Section 2 is met often. In 21 seasons every pair of essential teams split its series, so $A_{\mathcal{E}} = 0$. In 19 the essential block meets the rank condition with some nonzero entries. A zero block removes the constraints among the essential teams, but the inequalities of the other teams still cut the simplex on them. Those cuts are why zero-block seasons can have widths below 1 (England 2019-20: two essential teams, width $1/3$). The full segment occurs only when no other team's inequality binds. England 2017-18 is that case. Manchester City and Manchester United split their series and no other team beats any mixture of the two, and the equilibrium set is the whole segment between them. The maximum-entropy equilibrium splits it evenly and the LP vertex puts all mass on one of them. Between the two selection rules the ratings of the other eighteen teams then differ by up to 0.5 on the score scale. The largest rating gap of Section 4.3 is in France 2017-18 and the largest copy effect in Spain 2021-22. These are complete season records, so the widths have no sampling uncertainty and the dependence on the selection rule is a property of the recorded ratings.

I.2 The football duplications under the criterion

The 153 duplications also test Theorem 2 case by case. For each duplicated team we compute the range of its coordinate over the equilibrium set, the gap of the linear program of Appendix A.3, how far the maximum-entropy equilibrium moves once the copy's probability is added back, and the change of the other teams' ratings. Table 12 sorts the 153 duplications into the cases of Theorem 2. The two invariant cases of the theorem account for 42 of them, and the criterion decides the 107 whose coordinate varies and lies below its largest value: whenever the gap exceeds 10^{-3} the point moves, and whenever it is below 4×10^{-6} nothing moves. The remaining four are classified mathematically and sit below the numerical resolution. The four duplications in the third row are movements below the numerical resolution; they lie in quasi-strict seasons, where by Corollary 1 a probability that varies across equilibria cannot be at its maximum at the maximum-entropy equilibrium, and their coordinate is within 7×10^{-5} of that maximum. The 25 invariant duplications of the last row all lie in seasons that are not quasi-strict. These 25 of the 153 are the case described after Corollary 1: the copy leaves the equilibrium in place although the duplicated team's probability could rise.

Table 12: The criterion decides the 107 football duplications whose coordinate varies and lies below its largest value; the two invariant cases of Theorem 2 account for 42 more, and four sit below the numerical resolution. #: duplications; displacement is of the merged maximum-entropy equilibrium. The last two columns give a bound over the case where nothing moves and a count where something does.

case	#	LP gap	displacement	rating change
coordinate constant	6	$\leq 5 \times 10^{-7}$	$\leq 3 \times 10^{-5}$	$\leq 4 \times 10^{-5}$
coordinate at its largest value	36	$\leq 2 \times 10^{-5}$	$\leq 3 \times 10^{-5}$	$\leq 6 \times 10^{-5}$
within resolution of that value	4	$\leq 5 \times 10^{-5}$	$\leq 7 \times 10^{-5}$	$\leq 1.3 \times 10^{-4}$
varies, below it, gap $> 10^{-3}$	82	$> 10^{-3}$	median 0.072, up to 0.167	all 82
varies, below it, gap ≈ 0	25	$< 4 \times 10^{-6}$	$\leq 4 \times 10^{-5}$	none

In all 25 the criterion holds to solver precision: the optimum of the linear program is zero to 4×10^{-6}, and to that precision no feasible direction has a positive tilted derivative. A zero optimum does not say whether the fall in entropy exactly cancels the tilt along some direction that raises the coordinate, or outweighs it along all of them, since the zero direction is always feasible. Maximizing the tilted

derivative over the feasible directions that raise the coordinate by a fixed amount separates the two: in 6 of the 25 the best such direction has tilted derivative zero (to 3×10^{-5}), and in 19 it is strictly negative, between -0.05 and -5.0 per unit of mass moved onto the coordinate. England 2019-20 is one of the six and shows the mechanism analytically. Liverpool and Manchester City split their series and form the essential support, so the equilibrium set is a segment on the two, parametrized by Liverpool's weight t. Manchester United took a draw and a loss against Liverpool and two wins against Manchester City, so its rating is $-t/2 + (1 - t)$, and the constraint that this rating is at most zero cuts the segment to $t \geq 2/3$. The entropy is maximized at the cut, $t = 2/3$, where its slope along the segment is $\log((1 - t)/t) = -\log 2$. Duplicating Liverpool adds exactly $\log 2$; the tilted objective is stationary at the cut and the point does not move although t could rise to 1. Quasi-strictness holds in 18 of the 40 seasons. In those seasons every probability that varies across equilibria and lies more than 10^{-4} below its maximum moves the point, as Corollary 1 requires.

I.3 The payoff representation

This part supports the football statement of Section 4.3: more than one equilibrium and reversals occur under each of three payoff representations. The league tables use the score difference $A = 2W - 1$, with W the mean score over the two games of a pair. The 44 seasons are recomputed under two other representations: the clipped log-odds $\log(\tilde{W}/(1 - \tilde{W}))$ with $\tilde{W}$ the mean score clipped to $[1/8, 7/8]$, and the sign of $W - 1/2$. Both maps are odd. The matrices stay antisymmetric, and an evenly split series stays a zero entry. Under the score, log-odds and sign representations 40, 38 and 43 seasons have more than one equilibrium, and 21, 21 and 19 seasons have an all-zero essential block. The sign map keeps the zero pattern. The drop to 19 comes from a changed essential support. Duplicating each essential team (153, 139 and 170 duplications) reverses at least one pair of other teams in 41, 54 and 31 of them, spread over 17, 26 and 14 seasons. The largest change of another team's rating is 0.275, 0.507 and 0.333. We had written down before this run that the zero-block seasons should stay unchanged and that the reversal counts change but stay positive. Of these, the second holds, and the first holds for the log-odds map only. That most seasons have more than one equilibrium and that copies reorder teams does not depend on the representation.

J Competing rating rules

Section 4.4 compares Nash averaging with five other rules. This appendix defines each rule as implemented, reports the counts behind Tables 14 and 15, and compares the rules under near copies.

J.1 Definitions

The lexicographic construction is defined in Appendix A.9, where Lemma 5 identifies it with the deviation ratings of Marris et al. (2025) on a symmetric zero-sum game. The rule returns -1 for agent 4 of the five-agent game A_5 under every one of the 120 relabelings, before and after duplicating agent 5. On A_δ it returns $-\delta$ for $\delta > 0$ and -1 for $\delta < 0$. Every comparison here is on ratings. The two rules that do not pick an equilibrium are the standard paired-comparison model (Bradley & Terry, 1952) and the evolutionary ranking of multi-agent evaluation (Omidshafiei et al., 2019). Bradley–Terry strengths β are the maximum-likelihood fit (minorization iterations) of the model in which i beats j with probability $\sigma(\beta_i - \beta_j)$, applied to the win-probability matrix $\sigma(A)$ with equal game counts per pair on the generated games, and to the mean scores on football. They are reported as log-strengths centred on the original agents. α-Rank is the single-population implementation in OpenSpiel on the same win-probability matrix with $\alpha = 100$ and population size 50. It ranks agents by the stationary distribution of a Markov chain on agents: the population switches from its current agent to a challenger with a probability that rises with the win probability of the challenger against it, at selection intensity α. The mass of an agent is the long-run fraction of time the chain spends at it. The pair-reversal, top-set and reversed-fraction columns are invariant to any common shift or rescaling of the originals' values. The copied agent's own value follows the convention fixed next. Renormalizing the masses of α-Rank over the original agents can register movement even when the total mass of the copy and its original is unchanged (masses $(1/2, 1/2)$ split into $(1/4, 1/2, 1/4)$ renormalize to $(1/3, 2/3)$). We therefore report α-Rank under the merged convention, in which the mass of the copy is added to that of its original without renormalization. The same 11 399 of 11 965 duplications move some original mass by more than 10^{-4} under either convention, with median

Table 13: Forty of 44 football seasons have more than one equilibrium in exact arithmetic. width and dim.: the largest range of one team's probability over the equilibrium set and that set's dimension; e: essential teams; zero block: every essential pair split its series; gap and copy: the largest rating difference between the two rules and the largest change under one duplication.

season	width	dim.	e	zero block	rating gap	copy
England 2015-16	$1/4$	1	4	no	0.375	0.125
England 2016-17	$1/4$	1	4	no	0	0
England 2017-18	1	1	2	yes	0.5	0.167
England 2018-19	$5/11$	1	4	no	0	0
England 2019-20	$1/3$	1	2	yes	0.667	0
England 2020-21	$2/21$	2	5	no	0	0
England 2021-22	1	1	4	no	0.25	0.15
England 2022-23	0	0	5	no	0	–
England 2023-24	$3/14$	1	4	no	0	0.039
Germany 2015-16	$1/3$	1	4	no	0.174	0.076
Germany 2016-17	$2/7$	1	4	no	0	0.063
Germany 2017-18	$4/7$	1	4	no	0	0
Germany 2018-19	1	1	2	yes	0.5	0.167
Germany 2019-20	$7/9$	3	4	yes	0.986	0.202
Germany 2020-21	$2/3$	1	4	no	0	0.116
Germany 2021-22	$2/11$	1	4	no	0	0
Germany 2022-23	$6/13$	1	6	no	0.113	0.105
Germany 2023-24	$2/3$	1	2	yes	0.25	0.25
Spain 2015-16	$1/3$	2	3	yes	0.667	0.222
Spain 2016-17	0	0	3	no	0	–
Spain 2017-18	$1/3$	1	2	yes	0	0
Spain 2018-19	$1/2$	2	3	yes	0.25	0.024
Spain 2019-20	$1/9$	1	4	no	0	0
Spain 2020-21	$1/3$	2	5	no	0.131	0.076
Spain 2021-22	$18/19$	3	4	yes	0.5	0.275
Spain 2022-23	$16/203$	1	6	no	0	0
Spain 2023-24	0	0	3	no	0	–
Italy 2015-16	$1/2$	2	3	yes	0.667	0.019
Italy 2016-17	$2/3$	2	3	yes	0.667	0.167
Italy 2017-18	$2/3$	1	2	yes	0.75	0.25
Italy 2018-19	0	0	7	no	0	–
Italy 2019-20	$1/2$	4	5	yes	0.733	0.142
Italy 2020-21	1	2	3	yes	0.5	0.167
Italy 2021-22	$1/2$	2	3	yes	0.75	0.167
Italy 2022-23	1	3	4	yes	0.5	0.25
Italy 2023-24	$2/37$	1	6	no	0	0
France 2015-16	$2/3$	2	3	yes	0.846	0.205
France 2016-17	$2/3$	1	2	yes	0.75	0.25
France 2017-18	$8/9$	3	4	yes	1	0.188
France 2018-19	$6/7$	3	4	yes	0.875	0.2
France 2020-21	$1/8$	1	4	no	0.187	0.062
France 2021-22	$8/77$	1	6	no	0	0
France 2022-23	$1/3$	2	3	yes	0	0.222
France 2023-24	$2/41$	1	8	no	0.023	0.006

change 0.079 merged against 0.033 renormalized. Ranking columns differ and are given below. Removing exact copies before solving (Marris et al., 2022, Appendix C.2) is invariant to every exact duplication by construction, since it returns the original game (or the same reduced game before and after the duplication when the original already contains exact copies). It shares the maximum-entropy jump at the exact draw, since A_δ contains no repeated row. Table 14 lists it.

J.2 Measure

For each rule the change is the largest absolute change of any original agent's rating after duplicating one agent, over all agents of all 2000 instances of the games with dominated agents. "Moved" means a change above 10^{-4}. Above 10^{-6} the maximum-entropy count rises from 6680 to 9475 of 11 965, because the precision of the convex solver is about 10^{-6}. Above 10^{-2} it is 6511. At 10^{-4} the count is 6680 of 11 965, that is 0.56; the same 6680 is 0.639 of the 10 460 tests of Table 4, which count one test per agent of each game with more than one equilibrium. The continuity check uses the matrices A_δ of Proposition 2 at $\delta \in \{\pm 10^{-3}, \pm 10^{-2}, \pm 0.1\}$ and reports the rating of agent 4 under each rule. At $\delta = -10^{-3}$ and $+10^{-3}$ the maximum-entropy equilibrium, the LP vertex and deviation ratings give -1 and -0.001; Bradley–Terry gives -0.5792 and -0.5796 (a change of 4×10^{-4}). The mass of α-Rank on the agent is 0.0004 and 0.0011 (a change of 7×10^{-4}). At $\pm 10^{-2}$ and ± 0.1 the equilibrium rules give -1 and $-\delta$, Bradley–Terry moves by at most 0.05 and the mass of α-Rank stays below 0.002.

Table 14: On these generated games only the two rules that do not pick an equilibrium reorder agents under duplication. Seven rows over 11 965 duplications; columns: the fraction moving some rating by more than 10^{-4} on the rule's own scale (unit column), reversing some pair, changing the top set, and the mean reversed-pair fraction.

rule	moved	unit	pair reversed	top set changed	reversed-pair fraction
Nash averaging (maximum entropy)	0.56	rating (log-odds)	0.00	0.00	0
LP vertex (zero objective), copy listed last	0.00	rating	0.00	0.00	0
LP vertex (zero objective), relabeled	0.02	rating	0.00	0.02	0
deviation ratings	0.00	rating	0.00	0.00	0
exact copies removed, then maximum entropy	0.00	rating	0.00	0.00	0
Bradley–Terry	0.96	log-strength	0.06	0.40	0.0061
α-Rank	0.95	stationary mass	0.41	0.67	0.069

J.3 α-Rank parameters

The α-Rank counts change little over the tested parameter grid. With $\alpha \in \{10, 100\}$, the infinite-α limit, and population sizes 50 and 200, the first 500 instances give identical counts under the renormalized convention: 2841 of 2993 duplications move some original agent's mass above 10^{-4} and 531 reverse a pair. With $\alpha = 1000$, 160 solves fail; a failure on an original matrix skips every duplication of that instance, so 1975 of the 2993 tests remain, of which 1860 move some original agent's mass. The change of the inessential agent's mass at $\delta = \pm 10^{-3}$ stays below 0.005 in every configuration.

J.4 Relabeling test for the LP vertex

The LP vertex (zero objective) of Appendix B.11 is the first feasible vertex the solver returns, so its invariance in Table 14 may depend on where the copy is placed. We repeat the duplication tests after applying a random permutation to the agents of each instance (the same permutation to the original and the duplicated game) and inserting the copy at a random position, for five permutation seeds. The LP vertex (zero objective) then moves some original agent's rating by more than 10^{-4} on 242 to 262 of the 11 965 duplications (2.0% to 2.2%), by up to 4.4 log-odds. The LP vertex (minimax) of Appendix B.11 does so on 425 to 496 (3.5% to 4.2%). Neither rule reverses a pair that was strictly ordered before the duplication. Every move of the LP vertex (zero objective) changes the top set, because the vertex reached after duplication gives some inessential agent rating exactly zero where the previous vertex did not; for the first seed that is 243 top-set changes, one per moved test of that rule. The LP vertex row of Table 14 therefore describes one solver path. The maximum-entropy row

is label independent by Theorem 1, and deviation ratings are symmetric in the agents by construction and return the same ratings under all 120 relabelings of the five-agent game A_5.

J.5 Scale-free measures

For each duplication the scale-free measure counts the ordered pairs of original agents whose rating order is reversed (pairs tied within 10^{-6} before the duplication are skipped, and a pair tied after it does not count as reversed). The measure divides the reversed pairs by the pairs that were strictly ordered before the duplication. The last column averages this fraction over duplications. The measure also records whether the set of top-rated original agents (within 10^{-6} of the maximum) changes. With the copy appended last, the maximum-entropy equilibrium, the LP vertex and deviation ratings reverse no pair and change no top set on any of the 11 965 duplications. Bradley–Terry reverses some pair on 699 duplications (mean reversed fraction 0.6%) and changes the top set on 4788. α-Rank, with the copy's mass merged into its original's, reverses some pair on 4918 (mean 6.9%) and changes the top set on 7984. With the originals' masses renormalized instead the counts are 2358 (3.3%) and 4788. The two conventions agree on every pair that does not involve the copied agent (4035 reversed pairs). They differ on the copied agent's own order, whose total mass typically grows and whose renormalized mass is halved. The top set is computed over the original agents only, and the copy itself cannot enter it. Under Bradley–Terry all 4788 top-set changes are duplications of a member of a tied top set (out of 6116 such duplications; no instance of the family has a unique leader), and none of the 5849 other duplications changes the top set. Duplicating one member of a tie at the top breaks the tie among the original leaders: the duplicated agent and its copy drop below one of the other members. Such a change breaks an exact tie, which the pair-reversal column skips by construction. That column is the comparable measure of instability. Under the renormalized convention α-Rank gives the same top-set count, 4788. Under the merged convention of Table 14 it gives 7984, since merging can lift the copied agent's family above a tied leader. In Table 14 the top set of the merged α-Rank row is taken within 10^{-4} of the maximum mass and every other row within 10^{-6}.

J.6 Rating rules on the football duplications

The comparison of Section 4.4 gets its examples from the football duplications: which rules move, reorder, or change the top set when one essential team is duplicated. For each of the 40 seasons with more than one equilibrium and each essential team we append an exact copy with a drawn head-to-head (153 duplications). The duplicated table is rated by seven rules: the maximum-entropy equilibrium, the LP vertex (zero objective) with the copy appended last, the same LP vertex after a random relabeling of the duplicated table (five relabelings per duplication, 765 tests), deviation ratings, removing exact copies before solving followed by the maximum-entropy equilibrium, Bradley–Terry and α-Rank with the copy's mass merged into its original's. Bradley–Terry and α-Rank receive the observed score matrix $W = (A + 1)/2$, which is exact here. The counts of α-Rank are unchanged under $\sigma(A)$, the convention of the synthetic comparison, and Bradley–Terry then reverses a pair in 148 rather than 115 duplications with a largest change of 0.10 rather than 0.35. Therefore the Bradley–Terry counts depend on the payoff representation used, whereas a common rescaling of the output ratings leaves every ranking column unchanged. In Table 15 the reversed-pair column is a fraction of the strictly ordered pairs, and a pair counts as reversed only when it is separated by more than 10^{-3} before and after. A tolerance of 10^{-5} would count 141 pairs that are exact ties before duplication and are separated by solver noise of up to 8.1×10^{-5}. The smallest exact margin between two strictly ordered teams is 0.006, so 10^{-3} separates the two. Table 15 reports for each rule the duplications on which some other team's rating changes by more than 0.01, those on which some strictly ordered pair of other teams reverses, the reversed pairs as a fraction of all strictly ordered pairs, and those on which the top set changes (teams within 10^{-4} of the maximal rating). The maximum-entropy equilibrium moves a rating by more than 0.01 in 74 duplications and reverses a strictly ordered pair in 41. All 82 movements of the equilibrium counted under the criterion in Appendix I change some rating, eight by at most 0.01. The 0.7% of Table 15 is 170 of the 24 063 strictly ordered pairs of other teams summed over the 153 duplications. Among the 170 reversed pairs the margin before duplication has quartiles 0.036, 0.083 and 0.125 (range 0.006 to 0.375) and the margin after has quartiles 0.028, 0.083 and 0.10. Over all strictly ordered pairs the median margin is 0.333. Before duplication 78 reversed pairs (in 22 duplications and 12 seasons) have a margin above 0.1, and 14 have one above 0.2. The maximum-entropy equilibrium changes the top set in 31

duplications: an inessential team can be rated exactly zero and duplication moves such teams onto or off zero (Germany 2023-24: Heidenheim from -0.25 to 0 when Stuttgart is duplicated). The set of essential teams itself is unchanged (Lemma 4). The LP vertex (zero objective) with the copy appended last moves in 19 duplications, by up to 1.0. On these tables even the fixed-order solver path is not invariant. Under relabeling it moves in 205 of 765. Deviation ratings change no rating by more than 7×10^{-6} and reverse nothing above the tolerance, and removing exact copies before solving changes nothing by construction. Bradley–Terry moves a rating in every duplication (up to 0.35) and reverses a pair in 115. α-Rank moves in 150 and reverses a pair in 124 (152 at separation 10^{-5}: the stationary masses of weaker teams are of order 10^{-4}, so the coarser separation hides reversals among them). On the unmodified football seasons the conclusions of Table 14 therefore change in one respect: the maximum-entropy equilibrium reorders other teams here, in 17 of 40 seasons, and so do both LP vertex paths.

Table 15: On the football duplications the maximum-entropy equilibrium moves and reorders other teams while deviation ratings and removing exact copies do not. Counts over 153 duplications, 765 for the relabeled row; moved: some other team's rating changes by more than 0.01; reversal: some strictly ordered pair reverses; top set: the top-rated set changes.

rule	moved > 0.01	max change	reversal	reversed pairs	top set
Nash averaging (maximum entropy)	74	0.275	41	0.7%	31
LP vertex (zero objective), copy listed last	19	1.0	19	1.0%	19
LP vertex (zero objective), relabeled (of 765)	205 (27%)	1.0	202	2.4%	190
deviation ratings	0	7×10^{-6}	0	0	0
exact copies removed, then maximum entropy	0	0	0	0	0
Bradley–Terry	153	0.35	115	1.0%	12
α-Rank	150	0.16	124	1.6%	10

J.7 Maximum entropy over a relaxed set as a comparator

This part puts the trade-off of Section 4.4 inside a single rule: maximum entropy over a relaxed equilibrium set. The ratings of Section 4.2 come from the exact equilibrium set of the estimated game. A natural alternative maximizes entropy over the relaxed set $\{p : \hat{A}p \preceq \varepsilon \mathbf{1}\}$ with a slack chosen from the data. Here the slack is the median over pairs of the delta-method standard error of the estimated log-odds, $1/\sqrt{g\,\hat{W}(1-\hat{W})}$ with $\hat{W}$ the estimated win probability of the pair (0.13 at $g = 300$, 0.041 at $g = 3000$). On a separate 200-replicate run of the sampling design of Figure 2 (agent 4's rating), at both tested sample sizes the relaxed selection rule lands nearer the exact rating at the exact draw than the exact selection rule does. At $\delta = 0$ its median error is 0.086 at $g = 300$ and 0.030 at $g = 3000$, against 0.37 and 0.31 for the exact selection rule. Away from the exact draw it is biased: at $\delta = 0.1$ its median error is 0.56 and 0.32, against 0.10 and 0.028 for the exact selection rule. It also loses exact-copy invariance. Duplicating agent 5 in the estimated game moves its rating of agent 4 by a median 0.12 to 0.15 at $\delta = 0$ and 0.11 (at $g = 300$) at $\delta = 0.1$. The exact selection rule, whose estimated set is a single point, moves by 0. The relaxed selection rule is therefore continuous where the exact one is not and non-invariant where the exact one is invariant, the trade-off of Section 4.4 in one rule. A matched comparison appends a near copy of agent 5 to the same estimated games, with its row perturbed by 0, half the slack, or twice the slack. It compares three rules: the exact selection rule, removal of near copies within the slack followed by the exact selection rule, and the relaxed selection rule. The removal rule as implemented here and in the comparison below scans the pairs in agent order, removes the later member of the first pair whose head-to-head and rows against the others are all within the slack, and gives it its partner's rating. Since the copy is appended last, the removal rule always removes the copy. A population-level version (connected components of the near graph, one representative per component) and its behaviour under permutations, several copies and chains are in the part "A population-level removal rule" below. For an exact copy the first two coincide and move nothing while the relaxed selection rule moves by 0.12 to 0.15 at $\delta = 0$. For a copy within the slack the exact selection rule can tie or untie the estimated set, and its rating of agent 4 jumps on part of the replicates (90th percentile of the change 0.74 to 0.82 at $\delta = 0$). Removal of the near copy holds the change to at most 0.03 at the 90th percentile. Here the relaxed selection rule moves by 0.12 to 0.15 as before. For a copy beyond the slack the removal rule no longer merges it and coincides with the exact selection rule again (90th

percentile 0.77 to 0.93 at $\delta = 0$). At $\delta = 0.1$ the exact and removal rules agree to within 0.02 in every cell, with errors 0.10 at $g = 300$ and 0.026 at $g = 3000$, while the relaxed selection rule's error is 0.45 and 0.30. Removal within a data-chosen slack therefore extends exact-copy invariance to copies inside the slack at the price of a threshold. Its own discontinuity now lies there.

J.8 Four rules under near copies on the two empirical matrices

The closing comparison of Section 4.4 asks which rule keeps its ratings on real tables when a near copy is added. It runs the near-copy protocol of Appendix H.7 on the engine league (slack $b = 0.058$, largest bootstrap standard error 0.037 in this run) and the leaderboard ($b = 0.089$, largest 0.097 in this run). Every agent is copied, at $\eta \in \{0, 10^{-3}, 10^{-2}, 0.03, 0.1\}$, with drawn and perturbed head-to-heads and five fresh noise seeds. Four rules are compared: the exact maximum-entropy equilibrium, deviation ratings, near-copy removal within b followed by the maximum-entropy equilibrium, and the maximum-entropy equilibrium over the relaxed set $\{p : Ap \preceq b\mathbf{1}\}$, with the one slack b of the table serving both to merge and to relax. Our measure is the largest change of an original agent's rating in the game with the copy appended. On the engine league an exact copy ($\eta = 0$) moves the exact selection rule by 0.011 (the tilt on the estimated draw of Section 4.3), the relaxed selection rule by 0.012, and the other two rules by nothing. Removal moves nothing on any case for η up to 0.03, inside the slack. The exact selection rule moves at most 0.049, the relaxed selection rule at most 0.028, and deviation ratings 0.066 to 0.097 with 10% to 22% of cases above the largest standard error. A perturbation of 10^{-3} already moves deviation ratings by 0.068, nearly two standard errors. The copy is no longer merged at $\eta = 0.1$, beyond the slack. The exact rule, deviation ratings and the removal rule then coincide at 0.14 to 0.15 (3% to 22% of cases above the largest standard error), while the relaxed selection rule moves 0.049 to 0.075. On the leaderboard, which has a unique equilibrium, the exact rule and deviation ratings agree within 0.004, and removal moves nothing up to $\eta = 0.03$. All three move 0.14 to 0.18 at $\eta = 0.1$ (2% to 5% above the largest standard error). The relaxed selection rule moves 0.015 under an exact copy and at most 0.078 at $\eta = 0.1$. Removal within the slack, with the copy appended last, is thus the only rule invariant to copies inside the slack on both tables; deviation ratings are the least stable on the league; and the relaxed selection rule trades invariance at $\eta = 0$ for the smallest changes at large η.

J.9 A population-level removal rule

Whether near-copy removal works without knowing which row was added is the last question. Two agents are near when their head-to-head is within the slack b and their rows against the other agents differ by less than b in every entry. The rule takes the connected components of this relation and replaces each component by one representative whose entry against another component is the mean of the entries of its members against that component (a double average, so the reduced game stays antisymmetric). It then computes the maximum-entropy equilibrium of the reduced game and gives every member the rating of its representative and an equal share of its mass. Equality of rows is tested to 10^{-9} in the implementation. On exact copies of a single agent it is exact. We run five tests on the engine league ($b = 0.058$) and the leaderboard ($b = 0.089$), with the same bootstrap standard errors as above. Under 20 random orderings of the agents the ratings of the original tables change by at most 5×10^{-6} and 2×10^{-9}: the rule is permutation invariant. Copying two agents at once (every pair, 120 and 378 tests) changes no rating. On these tables every component is a set of exact copies or a single agent. Averaging is not invariant to an exact copy inside a component that already contains distinct rows, because the copy enters the mean twice. With three agents 1, 2, 3, entries $A_{12} = 0$, $A_{13} = 0.8$, $A_{23} = 0.7$ and threshold 0.2, the component $\{1, 2\}$ beats agent 3 by 0.75. A copy of agent 1 moves this to 0.767 and the rating of agent 3 from -0.75 to -0.767. The repair is to collapse exact copies inside each component before averaging (rows equal to 10^{-9} against the others, drawn head-to-head), which restores -0.75 in the example. After a near copy of an agent at $0.6b$ has formed a two-member component on both tables, an exact copy of that agent moves the averaging rule by $\eta/6$ at the median (0.0058 and 0.0089; at most 0.0085 and 0.0090) and the repaired rule by 0 in all 80 and 140 tests. With two exact copies the averaging rule moves by $\eta/4$ at the median (at most 0.0136). Both rules coincide on every test below, since there a component either has no exact copies or consists of them alone. Two independent near copies of one agent at $\eta = 0.6b$, which can lie up to $1.2b$ apart, change ratings by at most 0.057 and 0.036. A chain (a copy at $0.6b$ and a copy of the copy a further $0.6b$ away, so that the ends are $1.2b$ apart) joins all three into one component in every case and changes ratings by at most 0.068 and 0.054. Both changes

are the direct effect of averaging perturbed rows and are of the order of η (0.035 and 0.053). The multiplicity-weighted entropy named by Marris et al. (2022) for noisy payoffs (their Appendix D) is stated without a rule for which rows count as copies, so an implementation would have to supply that rule. Our component rule above is one such rule, and its behaviour on heterogeneous components is the relevant finding for that choice. The pairwise rule used above, applied to the same chains under 20 orderings, changes ratings by up to 0.035 (league) and 0.22 (leaderboard) depending on the order. The repaired component rule removes the order dependence and is invariant to exact copies by construction, at the price of merging chains of near agents whose ends are not near. Changes it makes under near copies are those of averaging perturbed rows.

K Further related work on clone axioms

Marris et al. (2022) propose removing exact copies before solving (their Appendix C.2). For noisy payoffs they propose an entropy that divides each term by a similarity count (their Appendix D). At the exact level the removal returns the original game, so Theorem 2 says when it changes the maximum-entropy equilibrium. The weighted entropy is not the tilt of Theorem 1, and its effect on the ratings is not derived there. Composition consistency was defined for tournament solutions by Laffond et al. (1993, Definition 4), and the bipartisan set satisfies it (their Proposition 6). Laffond et al. (1996) study it as an axiom. These tournament axioms do not cover duplication. Composition consistency replaces a vertex by a sub-tournament whose members still beat one another, whereas a copy creates a drawn pair and leaves the class of tournaments, the matrices without draws. For social choice functions, composition consistency is stronger than independence of clones (Brandl et al., 2016; Proposition 8 with Theorem 1 of the arXiv version of Berker et al., 2025). A mixture of rules that each satisfy the axiom can still break it (Berker et al., 2025, footnote 5). The metric-space axiomatization of clone-robust weights (Berriaud & Wattenhofer, 2026a) also shows that, together with symmetry, summing weights over near-clone families is incompatible with fairness among the clones. Its sequel meets the difficulty of this paper in a neighbouring setting. Maximizing a graph entropy, concave but not strictly concave, returns a convex set of weightings. Breaking the tie by Shannon entropy then violates locality, because that entropy depends on the number of copies in each class (Berriaud & Wattenhofer, 2026b).

L Finite-sample calibration of the resolution diagnostic

Section 6 recommends reporting whether the equilibrium is unique, that is, whether the exact width of the equilibrium set is zero, and this appendix treats the width at the resolution of the data as descriptive. It asks whether that resolution width, or the range of the ratings over the same set, predicts the finite-sample error of the ratings at the maximum-entropy equilibrium. We wrote down four criteria before the run. Each is listed with its outcome, and Table 16 lists the medians.

L.1 Design

Three kinds of population are sampled: the matrix A_δ of Proposition 2 at $\delta \in \{0, 0.01, 0.03, 0.1, 0.3\}$, the estimated score matrix of the engine league of Appendix F taken as a population, and one Gaussian game $Z - Z^\top$ with $n = 8$ (standard normal Z). For each population and each number of games per pair $g \in \{30, 100, 300, 1000, 3000\}$, 200 replicates draw g Bernoulli outcomes per pair with the population's expected score. For the league this replaces its win, draw and loss frequencies by a two-point outcome with the same mean. Each replicate then estimates the matrix (clipped log-odds for the first and third kinds as in Figure 2, $W - W^\top$ for the league) and solves the estimated game. Here the slack is one standard error of an entry of the estimated matrix on its own scale. For the score-difference matrix it is $1/\sqrt{g}$. For the log-odds matrices it is the median over pairs of the delta-method standard error $1/\sqrt{g\,p(1-p)}$ at the population probability p. For A_δ this is 0.41 at $g = 30$ and 0.041 at $g = 3000$. Three numbers are recorded per replicate: the rating error $\max_i |\hat{r}_i - r_i|$ against the population's exact rating; the width, the largest range of one agent's probability over $\{p : \hat{A}p \preceq \varepsilon \mathbf{1}\}$; and the rating range, the largest over i of the range of $(\hat{A}p)_i$ over the same set. Table 16 reports the medians of the three.

Table 16: In the engine league the rating range falls with the error across sample sizes; the width does not, and neither diagnostic meets what we wrote down before the run (text below). Median rating error / width / rating range over 200 replicates per population and games per pair g, at a slack of one entry standard error.

population	$g = 30$	100	300	1000	3000
A_δ, δ=0	0.40/1.00/1.58	0.38/1.00/1.07	0.36/1.00/1.01	0.53/1.00/1.00	0.62/1.00/1.00
A_δ, δ=0.01	0.73/1.00/1.59	0.67/1.00/1.12	0.28/1.00/1.00	0.13/1.00/0.99	0.086/1.00/0.98
A_δ, δ=0.03	0.66/1.00/1.54	0.38/1.00/1.07	0.24/1.00/0.97	0.11/1.00/0.98	0.053/1.00/0.96
A_δ, δ=0.1	0.59/1.00/1.54	0.29/1.00/1.02	0.15/1.00/0.90	0.088/0.81/0.74	0.049/0.46/0.42
A_δ, δ=0.3	0.39/1.00/1.53	0.30/0.99/1.00	0.16/0.56/0.75	0.089/0.30/0.53	0.048/0.16/0.30
league (exact draw)	0.21/1.00/0.62	0.14/1.00/0.34	0.087/1.00/0.19	0.048/1.00/0.11	0.036/1.00/0.089
Gaussian $n = 8$	0.74/0.99/2.98	0.38/0.83/1.91	0.20/0.55/1.12	0.11/0.35/0.62	0.062/0.20/0.35

L.2 Criteria and outcomes

1. For the Gaussian population the median error at $g = 3000$ was to fall below 0.05 and the median width below 0.2. Neither does (0.062 and 0.20, after 0.11 and 0.35 at $g = 1000$; the error falls as $1/\sqrt{g}$).
2. At $\delta = 0$ the median error was to stay at or above 0.2 at every g and the width at or above 0.2. Both hold. The error does not fall with g (0.36 at $g = 300$, 0.62 at $g = 3000$ in this run) because the estimates concentrate on the two vertex values of Figure 2, whose errors are 0.25 and 0.75. For that two-point mixture the median depends on which mode holds more than half of a run's replicates.
3. At intermediate δ the error was to fall once two slacks are smaller than δ, and the rating range was then to fall below 0.2. Its error part holds at every δ (monotone in g). The rating-range part does not (0.42 at $\delta = 0.1$ and $g = 3000$, where two slacks are 0.08).
4. In the transition cells (those with median width between 0.05 and 0.5) the rank correlation across replicates between a diagnostic and the error was to exceed 0.3 for that diagnostic to count as a predictor of error. The five such cells give median correlations of 0.22 for the width and 0.22 for the rating range, and neither meets the criterion.

L.3 Reading

Across replicates at a fixed sample size neither quantity predicts the rating error. Across sample sizes the rating range does track it where the width does not: in the league the exact draw keeps the width at 1.0 at every g, while the rating range falls from 0.62 to 0.089 alongside the error (0.21 to 0.036). The league population here is the estimated matrix as played, in which the rows of the copy pair differ against the four stochastic opponents. Over the population's equilibrium set (the segment between the pair) twelve ratings are constant to 3×10^{-9}. Four vary, by 0.0033, 0.0033, 0.0067 and 0.067 (the rating-limited opponent). The ratings are therefore nearly but not exactly independent of how the probability is split between the pair, and the residual error at $g = 3000$ (0.036) lies inside that range. At $\delta = 0$ the width, the rating range and the error all stay large, because moving along the equilibrium set moves probability between agents with different columns. A positive width at the data's resolution says that the point a selection rule picks is one choice among several. Over the same set the rating range says whether that choice can reach the ratings. For an exact copy at an exact equilibrium, whether the choice reaches the ratings is decided by Theorem 4(b), not by the size of either quantity.

M Seeds, software and compute

M.1 Seeds and released files

The seeds are 1000 for the three families, 7000 for the 150 fresh instances compared with OpenSpiel, 33 for the 200 instances of the bound-based check (Appendix C), 3000 for the identical-agent draws and 7 for the sign sweep (Appendix D), 21 for the opening prefixes of the engine league (Appendix F), and 0 and 17 for the Gaussian game and the sampling of the calibration study (Appendix L). Every other random draw behind these appendices, the bootstraps, the shuffle controls,

the subsampling, the peeling streams, the relabelings and the near-copy noise, is seeded inside the released script that produces it. Every per-case file behind the appendices (the scripts, the per-test gaps, the full tables, the score matrix and game records of the league, the leaderboard matrix and model list, the per-case and per-replicate results and the transcribed Atari values) is in the supplementary material.

M.2 Software

Python 3.10.20 with NumPy 2.2.6, SciPy 1.15.3 (HiGHS, SLSQP), CVXPY 1.7.5 with Clarabel 0.11.1, pycddlib 3.0.2 and OpenSpiel 1.6.11; every script runs on a CPU.

M.3 Compute

Every analysis except the league games takes about one CPU-day in total, most of it in the 22 250 entropy solves of the two families and the football enumerations; the league games took about six hours of wall time on two 48-core machines with the two Stockfish versions.