THE NETWORK HHD: QUANTIFYING CYCLIC COMPETITION IN TRAIT-PERFORMANCE MODELS OF TOURNAMENTS *

ALEXANDER STRANG †, KAREN C. ABBOTT, ‡, AND PETER J. THOMAS §

Abstract. Competitive tournaments appear in sports, politics, population ecology, and animal behavior. All of these fields have developed methods for rating competitors and ranking them accordingly. A tournament is intransitive if it is not consistent with any ranking. Intransitive tournaments contain rock-paper-scissor type cycles. The discrete Helmholtz-Hodge decomposition (HHD) is well adapted to describing intransitive tournaments. It separates a tournament into perfectly transitive and perfectly cyclic components, where the perfectly transitive component is associated with a set of ratings. The size of the cyclic component can be used as a measure of intransitivity. Here we show that the HHD arises naturally from two classes of tournaments with simple statistical interpretations. We then discuss six different sets of assumptions that define equivalent decompositions. This analysis motivates the choice to use the HHD among other existing methods. Success in competition is often mediated by the traits of the competitors. A trait-performance model assumes that the probability that one competitor beats another is a function of their traits. We show that, if the traits of each competitor are drawn independently and identically from a trait distribution then the expected degree of intransitivity in the network can be computed explicitly. We show that increasing the number of pairs of competitors who could compete promotes cyclic competition, and that correlation in the performance of A against B with the performance of A against C promotes transitive competition. The expected size of cyclic competition can thus be understood by analyzing this correlation.

Key words. Cyclic competition, intransitivity measures, least squares rating, Helmholtz-Hodge decomposition, trait-performance models

AMS subject classifications. 05C50, 05C20, 05C21

1. Introduction: Tournaments, Ranking, and Intransitivity.

A tournament consists of a group of competitors who compete in pairwise events (head-to-head matches). Tournaments are important across disciplines, from ecology and animal behavior [43, 63], to psychology and sports [6, 35]. Rating and ranking, that is, assigning a measure of quality to the competitors and listing them in order from best to worst, is important in each of these areas. In sports, ranking and rating teams and players is a topic of broad popular interest. In biology, ratings are widely used to evaluate the quality of competitors in social hierarchies. High standing in a competitive hierarchy may be closely related to fitness, as it is often associated with priority access to resources [17, 38, 39, 69], territory maintenance [64], and higher reproductive output [54, 75]. Ranking is especially important in politics, as many electoral systems determine a winner by aggregating votes into a partial ranking of the candidates. Ratings and rankings are often sought since they simplify the description of a tournament by assigning each competitor a single number that purports to measure how good they are.

Not every tournament allows for a consistent ranking of competitors. As a motivating example, consider the 2019–2020 National Basketball Association (NBA) season, which was cut short by the COVID-19 pandemic. Imagine two fans arguing whether the Cleveland Cavaliers (CLE) or Sacramento Kings (SAC) were the better team. The two teams did not play in 2019–2020 due to the abbreviated season, so

*Submitted to the editors August 28, 2022.

Funding: This work was funded by NSF grants DEB-1654989 and DMS-2052109.

†Department of Statistics, University of Chicago, IL (alexstrang@uchicago.edu).

‡Department of Biology, Case Western Reserve University, Cleveland, OH (kcabbott@case.edu).

§Department of Mathematics, Applied Mathematics, and Statistics, Case Western Reserve University, Cleveland, OH (pjthomas@case.edu).

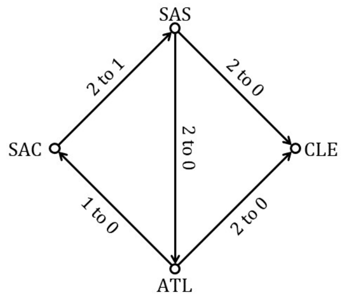

FIG. 1. *A network representing the observed outcomes of games between the Cleveland Cavaliers (CLE), Sacramento Kings (SAC), Atlanta Hawks (ATL), and San Antonio Spurs (SAS) in the 2019–2020 regular season. Arrows point from the team which lost the majority of the games to the team which won the majority. Labels next to the arrows provide the game outcomes.*

they cannot be compared directly. The Cleveland fan points out that CLE beat the San Antonio Spurs (SAS) 2 out of 2 games, and SAS beat SAC 2 out of 3 games, so surely CLE was better than SAC. The SAC fan counters that transitive predictions of this kind are not always valid. For example, the Atlanta Hawks (ATL) beat SAS 2 out of 2 games, and SAS beat SAC 2 out of 3 games, yet SAC still beat ATL in the game they played. Figure 1 illustrates these outcomes as a graph. Notably, the graph contains a mixture of triangles which do and do not allow consistent rankings. A believer in ranking could point to the triangle involving CLE, SAS and ATL as evidence that NBA teams can be consistently ranked, while a skeptic might point to the triangle involving SAS, ATL, and SAC.

The observation that not all tournaments admit consistent rankings motivates classification into transitive and intransitive tournaments. A tournament is *transitive* if knowing that A usually beats B, and B usually beats C, is enough to conclude that A usually beats C. Transitive tournaments are consistent with a global ranking of all the competitors. An *intransitive* tournament is a tournament that is not consistent with any global ranking. Intransitive tournaments must contain at least one cycle where the transitive assumption fails. Figure 2 illustrates examples of transitive and intransitive tournaments.

Intransitive tournaments appear in practically every discipline where tournaments are studied [10, 23, 52, 57, 59], and are the norm rather than the exception when using real data [32, 35, 36, 43, 63, 66, 68]. Intransitivity may arise due to uncertainty in observed data [35, 68], randomness in event outcomes, or may be intrinsic, as in the game of rock-paper-scissors.

Intransitivity is important for two reasons. First, intransitivity presents a challenge when ranking since no ranking is consistent with the tournament. For example, Condorcet's paradox is a voting paradox in cyclic community preferences prevent any fair ranking of candidates, and thus, any choice of winner [23].[1] Second, when intransitivity is intrinsic, then the tournament contains cyclic structure, as in rock-paper-scissors. Cyclic structures can radically alter optimal strategies [10] and long term dynamics [52, 59, 58, 60, 61]. For example, in ecology it is widely hypothesized that intransitive competition between species promotes biodiversity since no species domi-

[1] Suppose there are three candidates in an election and three voters. Suppose that the first voter prefers A to B to C, the second B to C to A, and the third C to A to B. Then A would beat B, B would beat C, and C would beat A in pairwise head-to-head elections.

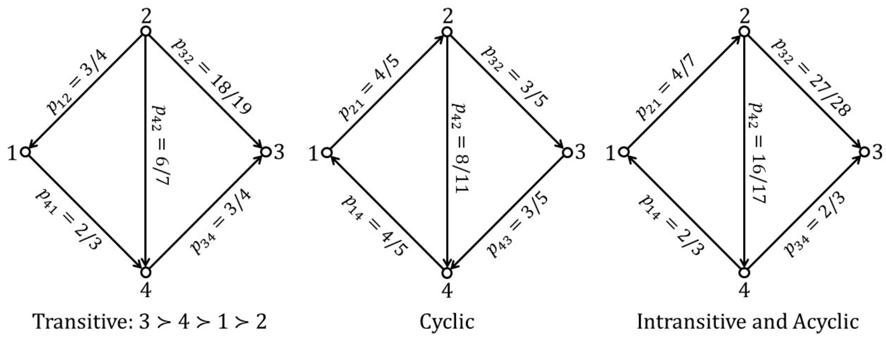

FIG. 2. *Three example networks representing different classes of tournaments. The first is transitive since the win probabilities are consistent with the ranking* $3 \succ 4 \succ 1 \succ 2$. *The second is both intransitive, and, more strongly, is cyclic (see Section 3.1.2 for definitions). The third is neither transitive nor cyclic, and represents a generic tournament, with the same pattern of expected winners and losers as in Figure 1.*

nates. This hypothesis is based on extensive theoretical work [43, 52, 59, 58, 60, 61, 80] and limited case-studies [30, 36, 46, 47, 66].

The importance of intransitivity in real natural communities is controversial [25, 70, 77], in part because there are few robust metrics for measuring intransitivity from incomplete and noisy data. Uncertainty in data can easily be conflated with observed intransitivity, and common sampling methods for filling in missing data can overestimate intransitivity [63]. Thus there is a need for ranking and rating methods that are robust to intransitivity and measures of intransitivity that can handle noisy and incomplete data.

Jiang and Lim introduced the discrete Helmholtz-Hodge Decomposition (HHD) as a general method for ranking objects from incomplete and imbalanced data [32, 50]. The decomposition is a network theoretic tool that we adapt to the study of competitive tournaments. The HHD accomplishes three fundamental tasks. First, it assigns a rating to each competitor. Competitors can be ranked accordingly. Second, it produces a measure of intransitivity that quantifies how far an observed network is from the nearest perfectly transitive network. Third, it represents the observed network as the direct sum of a perfectly transitive and a perfectly cyclic network. This decomposition provides an elegant characterization of intransitivities present in data, and can reveal underlying cyclic tendencies (c.f. [10]).

When compared to existing methods, the discrete HHD has a number of advantages. It is more general than some classical methods since it applies to arbitrary network topologies and can accommodate imbalanced data [32]. It is also more informative because it provides a clear description of both underlying transitive and cyclic structures. Most ranking methods and intransitivity measures focus on the transitive component while the HHD puts the transitive and cyclic components on equal footing. Finally, it remains efficiently computable even for large, incomplete networks [32]. In contrast, Slater's index [68] requires solving an NP hard optimization problem [11, 18], and Kendall's index [35] requires a complete network.

This paper aims to answer two fundamental questions:

1. Why use the HHD to study competition when other methods exist?
2. Having chosen to use the HHD, what do we expect when competitive performance is determined by individuals' traits?

Answering the first question is important since there are many possible methods to choose from, so the choice of method should be made in a principled way. Answering the second question is important since it builds a conceptual bridge from the competitors and competitive event to the overall tournament structure. As in Landau [44], we seek to understand how the underlying distribution of traits among competitors, and the relationship between traits and success, influence the overall tournament.

The latter question is important across disciplines. In some biological settings, success in competition is determined by individual traits, driving selection [76]. For example, competition for social dominance among male elephant seals depends on their body mass [26] and competition among male dwarf Cape chameleons depends on coloration, head size, and body length [76]. Success in these competition events is correlated with reproductive success, suggesting that heritable traits which improve a male's chances of success are strongly selected for [26]. In sports, the relationship between the traits of a player or team and their success is an area of active interest - for athletes, owners, fans, and researchers alike. The rise of sabermetrics, the statistical study of baseball, is a popular example [49, 78].

This paper answers questions 1 and 2 as follows:

1. The HHD arises naturally from the study of ranking and intransitivity. To illustrate this point, we provide a different derivation of the HHD than [32] or [50]. Instead of imposing the decomposition ad hoc, we propose two special classes of tournaments with clear statistical motivation. We then show that any tournament can be uniquely decomposed into a combination of tournaments from these classes. This decomposition is the HHD (see Theorem 3.5). Next we illustrate that the HHD can be reached by six different approaches (Corollary 8.1), and is thus robust to varying motivations.

2. We show that, under simple assumptions on the distribution of traits, the expected sizes of the components of the decomposition can be computed explicitly from the number of competitors, number of pairs who could compete, and the correlation in the performance of A against B with A against C. This correlation is shown to equal the uncertainty in the expected performance of a competitor, linking a decomposition of uncertainty in performance to tournament structure (see Theorem 4.1 and Corollary 9.1).

The answers to the second question prove, under minimal assumptions, a series of intuitive statements about transitive/cyclic competition that appear, as heuristics, across the literature. These include:

(a) The more predictable the performance of A against a randomly drawn competitor (i.e., the less the performance of A depends on their opponent) the more transitive the tournament.

(b) The more correlated the performance of A against B with the performance of A against C, the more transitive the tournament.

(c) The more pairs of competitors who could compete, the more cyclic the tournament is, on average.

(d) Filling in missing data by random sampling overestimates intransitivity.

Statements a, b, and c also hold in reverse. Decreasing a quantity that promotes transivity promotes cyclic competition

The paper is structured as follows. In Section 2 we provide some necessary background. Next, in Section 3, we derive the HHD in the context of tournaments and develop the associated ratings and intransitivity measure. Our derivation complements the cohomological approach used by [50], as it is specially adapted to tournaments, and only requires linear algebra and classical graph theory. In Section 4 we show how assumptions about the statistics underlying competition promote or suppress intransitivity. We focus on trait-performance models in which performance is assumed to be a function of traits sampled from a trait distribution. While win probabilities are not always determined by traits, exploring trait determined performance affords a more realistic and richer perspective than standard null models (c.f. [16]), and demonstrates generic relationships. In particular, we present a theorem (4.1) which allows the expected size of the intransitivity measure to be computed directly from the number of competitors, edges in the network, and correlation in the performance of A against B with A against C. This result is extended by a corollary (9.1) which shows that the correlation in performance is related to a decomposition in the uncertainty of the performance of A against B. These results lead to a deeper conceptual understanding of how cyclic structure can arise from uncertainty in performance, and can be suppressed by correlation in performance. We conclude by generalizing these observations to scenarios where the trait-performance assumptions do not hold.

2. Mathematical Framework. Consider an ensemble of V competitors. Assume that each competition event involves exactly two competitors, and never results in a tie. This standard assumption [35, 43] can be weakened to allow for ties. We will refer to competition of this kind as a tournament.[2]

The probability of any sequence of event outcomes in a tournament is determined by the probabilities that competitors beat each other. If the event outcomes are independent, then for each possible pairing of competitors there is an unambiguous probability one beats the other. Let p_{AB} denote the probability competitor A beats B. The shorthand $A \succ B$ denotes the case when A is expected to beat B ($p_{AB} > 1/2$). In principle, the win probabilities could change in time, and could depend on the history of the process (c.f. [24]). We focus on tournaments with unchanging win probabilities to avoid modeling additional temporal dynamics. Then a fixed set of win probabilities p determine the probability of any sequence of events. Thus the tournament dynamics are realizations of a random process, with probabilities controlled by p and the event order. The event order, i.e. the schedule, could be fixed or random. As in other studies of transitivity, we focus on the structure of the win probabilities p, not the schedule or tournament dynamics, since the win probabilities p determine the distribution of possible tournament outcomes, and whether competition is transitive or intransitive.

The win probabilities may be conveniently represented using a competitive network, $\mathcal{G}_{\rightleftarrows} = (\mathcal{V}, \mathcal{E}, p)$. Assign each competitor a node from the vertex set $\mathcal{V}$. Then $V = |\mathcal{V}|$. Introduce a pair of directed edges between each pair of competitors who could compete. The edge from B to A is assigned the weight p_{AB}. We assume that the tournament is finite, *connected* and *reversible*. That is, there are finitely many competitors, for any pair of competitors A B there is a path from A to B and from B to A through $\mathcal{G}_{\rightleftarrows}$ with probability greater than zero, and that $p_{AB} \neq 0$ or 1 for any pair A, B who could compete.

Sometimes it is preferable to simplify the competition network by rounding all weights less than $1/2$ to 0, and all weights greater than $1/2$ to 1. This can be conveniently represented as an unweighted graph $\mathcal{G}_{\rightarrow}$ which contains all directed edges from $\mathcal{G}_{\rightleftarrows}$ with weights greater than a half, and an undirected edge between all pairs with $p_{AB} = 1/2$. The edges in this graph point from expected losers to expected winners. Most intransitivity measures focus on this graph (see [35], [44], [68]).

[2] This is distinct from a *complete* tournament in which it must be possible for all pairs to compete.

A *ranking* is an ordered list of competitors from best to worst, specified by a rank function R which returns the rank of each competitor. Note that this is distinct from a *rating*, r, which is a function that returns a real number for each competitor [45]. Rankings are often generated by first rating each competitor, then listing them in decreasing order. For example, given competitors A, B, C with ratings $r_A = 10$, $r_B = 20$, $r_C = 0$ the corresponding ranking would be $R_A = 2$, $R_B = 1$, $R_C = 3$ and the competitors would be listed $B \succ A \succ C$. Ratings provide an intuitive description of competition in which some innate competitive ability determines performance.

Ranking methods are diverse, and well studied. Famous examples include the page-rank method used by Google to sort search results [9], the Massey and Colley methods used by the NCAA to rank basketball and football teams [45], and the Elo rating/ranking widely used by chess federations [24, 71]. The rating system produced by the HHD is a kind of log-least squares rating as is frequently used in paired comparison [6, 41, 42]. Examples of least squares rating systems are included in [14, 34, 45, 51, 72, 73].

A competitive network $\mathcal{G}_{\rightleftarrows}$ is consistent with a ranking R if $A \succ B$ whenever $R(A) < R(B)$. If a competitive network is consistent with a ranking then this ranking is unique and the network is *transitive*. Transitive networks satisfy the intuitive property that if we consider some sequence of competitors with $A \succ B \succ C \succ D$ then $A \succ D$. That is, $\mathcal{G}_{\rightarrow}$ contains no cycles, and all the edges in $\mathcal{G}_{\rightarrow}$ point from competitors with worse ranks to competitors with better ranks.

If $\mathcal{G}_{\rightarrow}$ contains a cycle, then there exists a sequence of competitors such that $A \succ B \succ C \succ .... \succ A$, and the tournament is *intransitive*. If a network is intransitive then it is not consistent with any ranking [57]. Speaking broadly, measures of intransitivity either count the number of intransitive triangles present in $\mathcal{G}_{\rightarrow}$ [35], or measure how far $\mathcal{G}_{\rightarrow}$ is from a nearby transitive network [68]. The Kendall measure [35] counts the number of intransitive triangles in $\mathcal{G}_{\rightarrow}$. This can be done efficiently, however prioritizes triangles over larger loops and does not weight edges equally [2, 68]. The Slater measure of intransitivity is the minimum number of edge directions that need to be reversed in order to transform $\mathcal{G}_{\rightarrow}$ into a transitive network [68]. While conceptually preferable [32], finding the closest transitive network is an NP hard problem [3], [19], [27], [32]. Despite some fast heuristics [18], complexity concerns limit the application of the Slater measure to small networks. The intransitivity measure associated with the HHD is conceptually analogous to the Slater measure, but can be computed efficiently even for very large networks. Note that transitivity and intransitivity are defined relative to the *sign* of $(p_{AB} - 1/2)$, rather than the exact value p_{AB}. In contrast, the intransitivity measure associated with the HHD is continuous in the win probabilities, so uses all the information available in $\mathcal{G}_{\rightleftarrows}$.

3. The Network HHD.

The Network Helmholtz-Hodge Decomposition (HHD) can be derived by defining two special classes of tournaments. These parallel the two classes of games defined in [10].

3.1. Arbitrage Free and Favorite Free Tournaments.

3.1.1. Arbitrage Free Tournaments (Perfectly Transitive).

A currency market is said to be *arbitrage free* if it is impossible to make money by exchanging currencies cyclically [32]. By analogy, we define an *arbitrage free tournament* to be a tournament for which it is impossible to expect to make money by betting on cyclic sequences of events. Specifically, a tournament is arbitrage free if, for any cyclic sequence of competitors $\mathcal{C} = \{i_1, i_2, \ldots, .i_{|\mathcal{C}|}, i_{|\mathcal{C}|+1} = i_1\}$, a sequence of wins where i_j

loses to i_{j+1} is equally likely as a sequence of wins where i_j beats i_{j+1} for all j. Here $|\mathcal{C}|$ denotes the number of competitors in the cycle.

Cycle Condition: A tournament is arbitrage free if and only if, for every cycle $\mathcal{C} = \{i_1, i_2, \dots, i_{|\mathcal{C}|}, i_{|\mathcal{C}|+1} = i_1\}$, the win probabilities satisfy:

$$p_{i_1 i_2} p_{i_2 i_3} ... p_{i_{|\mathcal{C}|} i_1} = p_{i_1 i_{|\mathcal{C}|}} ... p_{i_3 i_2} p_{i_2 i_1}. \tag{3.1}$$

The cycle condition can be simplified by dividing the right hand side across to the left hand side and taking a logarithm. Then:

$$\sum_{j=1}^{|\mathcal{C}|} f_{i_j i_{j+1}} = 0 \tag{3.2}$$

where the f_{ij} is the log-odds that competitor i beats competitor j:

$$f_{ij} = \text{logit}(p_{ij}) \equiv \log\left(\frac{p_{ij}}{1 - p_{ij}}\right). \tag{3.3}$$

The cycle condition is satisfied if and only if the sum of f around any cycle is zero. The log-odds, f, are an example of an *edge flow*: an alternating function, $f_{ij} = -f_{ji}$, on the edges [32]. Note that $\text{logit}(x) = \log(x/(1-x))$ is the inverse of $\text{logistic}(y) = 1/(1 + \exp(-y))$, so no information is lost in moving to f from p.

The sum of f around a cycle is an example of a path sum. A *path sum* against an edge flow is the discrete analog of a path integral against a vector field. Given a sequence of competitors $\mathcal{P} = \left\{i_1, i_2, \dots, i_{|\mathcal{C}|}\right\}$ the path sum against f over the path $\mathcal{P}$ is $\sum_{j=1}^{|\mathcal{C}|-1} f_{i_{j+1} i_j}$. The cycle condition requires that path sums over cycles equal zero.

If path integrals around closed loops equal zero, then the value of path integrals depend only on the endpoints of the path, are otherwise path independent, and the vector field is the gradient of potential. These properties also hold for networks.

Lemma 3.1 (Arbitrage Free). *A tournament is arbitrage free if and only if there exists a unique set of ratings r, with average rating equal to zero, such that the win probabilities satisfy $p_{ij} = \text{logistic}(r_i - r_j)$. Moreover if a tournament is arbitrage free then it is transitive.*

If there exist a set of ratings such that $p_{ij} = \text{logistic}(r_i - r_j)$ then $f_{ij} = r_i - r_j$ so path sums over f are telescoping, and thus cancel around loops. Then the cycle condition holds automatically. The rest of Lemma 6.1 can be proved using a simple spanning tree construction illustrated in of Figure 3 (panel a). We sketch the proof here; the supplement provides further details.

If a network is arbitrage free then the cycle condition requires that the path sum of f around any loop is zero. It follows that path sums over f are path independent. Our goal is to find a rating r such that the difference in r on each edge produces the edge flow f. We recover r by picking a spanning tree[3], and assigning it an arbitrary root, A. Uncentered ratings u are computed by setting u_i equal to the path sum from A to node i along the paths in the tree. Then the ratings r are set equal to u_i minus the average value of u. Path independence guarantees that the choice of tree does not influence u, and centering the ratings eliminates any dependence on the choice of

[3]A spanning tree is a subgraph of the network that contains no loops, includes all competitors, and is connected.

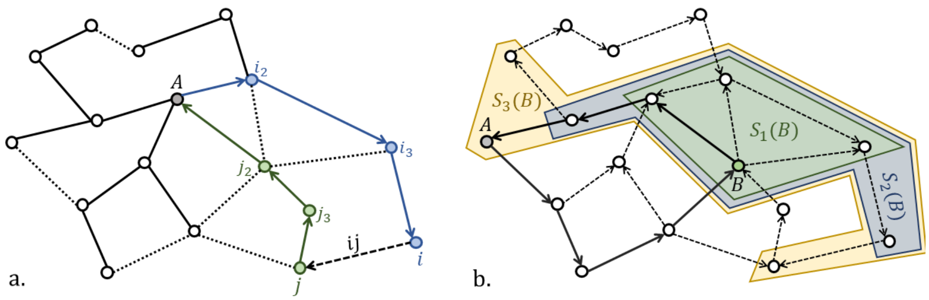

Fig. 3. *Panel a. The spanning tree construction for recovering the ratings for an arbitrage-free tournament. The tree is shown with solid lines, and the chords with dotted lines. The root of the tree, A, is marked in grey. Two vertices, i and j connected by a chord ij, are shown in blue and green respectively. The sequence of nodes leading from A to i and j are labelled. If the ratings, r, are constructed by evaluating path sums over the tree, then the path sum from j to A is $r_j - r_A$, and the path sum from A to i is $r_A - r_i$. Then, by the cycle condition, the sum around the loop marked with arrows is zero, hence $f_{ij} = r_i - r_j$. Panel b. A favorite free tournament must be a cyclic tournament. The arrows represent the direction of competition. If the network is favorite free, then whenever there is an edge pointing into a set there must be an edge pointing out of it. A path from A to B is shown in black. Then the sets $S_1(B), S_2(B), S_3(B)$ are shown as shaded polygons. These contain all competitors distance 1, 2, and 3 (respectively) from B. These sets continue to expand until they include A, hence there is a path from B to A.*

A. Then, by construction, $r_i - r_j = f_{ij}$ on all edges in the tree. The cycle condition guarantees that $r_i - r_j = f_{ij}$ on all edges not in the tree. Since f are the log-odds, $p_{ij} = \text{logistic}(r_i - r_j)$. Transitivity follows automatically since p must be consistent with the ranking induced by r.

Lemma 6.1 shows that arbitrage free tournaments are the only tournaments which match the logistic rating model $p_{ij} = \text{logistic}(r_i - r_j)$ used for Elo rating [1, 29, 45].[4]

Arbitrage free tournaments are also the only tournaments that match the Bradley-Terry model:[5] $p_{ij} = q_i/(q_i + q_j)$ where $q_i \geq 0$ are the Bradley-Terry ratings [8, 7]. If a network is arbitrage free, then setting $q_i = \exp(r_i)$ recovers the Bradley-Terry model. If the tournament satisfies the Bradley-Terry model, then setting $r_i = \log(q_i)$ produces a rating which satisfies $p_{ij} = \text{logistic}(r_i - r_j)$, so the network must be arbitrage free.

Since arbitrage free networks are a special class of transitive networks, we will refer to them as "perfectly" transitive. Note that a perfectly transitive network must satisfy the cycle condition, which is a requirement on the values of p rather than the sign of $(p - 1/2)$. Hence, while all perfectly transitive networks are transitive, not all transitive networks are perfectly transitive. For example, if $p_{AB} = 0.99$, $p_{BC} = 0.99$, and $p_{AC} = 0.51$ then the tournament is transitive, even though p_{AC} is much smaller than might be expected. This example is not perfectly transitive since it does not satisfy the cycle condition. The leftmost network in Figure 2 is perfectly transitive.

[4] The Elo rating system was originally proposed to rate chess players, but is also used to rank Sumo wrestlers [71], English league football teams [29] and international football teams. In the latter example the Elo method was the most predictive out of all methods tested [48]. The Women's World Cup uses a variant on the Elo method [48].

[5] The Bradley-Terry model is widely used in pairwise comparison and to rank competitors in tournaments. Examples include professional tennis [40, 53], Cape dwarf chameleons [76] and northern elephant seals [26]. Bradley-Terry models accounting for surface type, and discounting old games, have been shown to be effective in predicting the outcome of ATP tennis tournaments [53].

3.1.2. Favorite Free Tournaments (Perfectly Cyclic). In contrast, we define a *favorite free tournament* to be a tournament for which it is impossible to make money on average by betting on a favorite competitor over their neighbors. Specifically, A is equally likely to beat all of their neighbors, as to lose to them. Let $\mathcal{N}(i)$ denote the neighborhood of i, the set of all competitors who could compete with i. Then the win probabilities must satisfy a neighborhood condition.

Neighborhood Condition: A tournament is favorite free if and only if, for every competitor i with neighborhood $\mathcal{N}(i)$, the win probabilities satisfy:

$$\prod_{j \in \mathcal{N}(i)} p_{ij} = \prod_{j \in \mathcal{N}(i)} p_{ji}. \tag{3.4}$$

Like the cycle condition, the neighborhood condition can be written directly as a condition on the log-odds edge flow f defined in equation (3.3). A tournament satisfies the neighborhood condition if and only if the sum of f_{ij} over the neighborhood of i is zero for all competitors i:

$$\sum_{j \in \mathcal{N}(i)} f_{ij} = 0. \tag{3.5}$$

If the neighborhood condition is satisfied then it can be extended to all sets of competitors. Let S be a set of competitors and let $\mathcal{N}(S)$ be the set of all competitors not in S who neighbor S. Then the neighborhood condition implies:

$$\sum_{j \in \mathcal{N}(S), i \in S} f_{ij} = 0. \tag{3.6}$$

This identity follows from the discrete analog to the divergence theorem: the sum of f over the neighborhood of S equals the sum of f over the neighborhood of every competitor in S.[6] Then $\sum_{j \in \mathcal{N}(S), i \in S} f_{ij} = \sum_{i \in S} \sum_{j \in \mathcal{N}(i)} f_{ij} = \sum_{i \in S} 0 = 0$.

The cycle condition defined a special subset of transitive tournaments. The neighborhood condition also defines a special class that is a subset of a larger class - the class of cylic tournaments. A *cyclic tournament* is a tournament such that, if there is a path from A to B in $\mathcal{G}_{\rightarrow}$, then there is a path back from B to A in $\mathcal{G}_{\rightarrow}$.

LEMMA 3.2 (favorite free). *A favorite free tournament is cyclic, and is never transitive unless $p_{ij} = 1/2$ for all connected i, j.*

Like Lemma 6.1, Lemma 7.1 can be proved with a simple construction. The proof is sketched here and illustrated in Figure 3 (panel b). See supplement for details.

If there is a path from A to B in $\mathcal{G}_{\rightarrow}$ then we need to construct a path back to A from B. To this end, we define a nested sequence of sets where $S_d(B)$ is all vertices within distance from d of B in $\mathcal{G}_{\rightarrow}$. The neighborhood condition extends to sets of vertices, so if there is an edge into a set S in $\mathcal{G}_{\rightarrow}$ then there must also be an edge leaving S. It follows that, if A is not in $S_d(B)$, then $S_{d+1}(B) \neq S_d(B)$, so we can keep expanding the sequence of nested sets. If the network is finite then the sets cannot expand forever without eventually including A. To finish, a favorite free tournament cannot be transitive unless it is neutral, $p_{ij} = 1/2$ for all i, j, since only neutral tournaments are simultaneously transitive and cyclic.[7]

[6] If i and j are both in S then the sum over the neighborhood of i contributes f_{ij}, and the sum over the neighborhood of j contributes $f_{ji} = -f_{ij}$. Therefore the edge flow on any edge connecting a pair of nodes in S cancels in the sum.

[7] Note that a neutral tournament is considered transitive since it can be consistently ranked - all competitors should be ranked the same.

So, just as the cycle condition (no tendency to cycle) implied transitivity, the neighborhood condition, (no favorites) implies that the network is cyclic. Whether a tournament is cyclic or not depends on the sign of $(p_{ij} - 1/2)$, while the neighborhood condition is a condition on the values of p_{ij}. This motivates the definition: a tournament is *perfectly cyclic* if and only if it is favorite free. As before, all perfectly cyclic tournaments are cyclic, but not all cyclic tournaments are perfectly cyclic. The middle network in Figure 2 is perfectly cyclic.

Note that, unlike perfectly transitive tournaments where f is determined by a set of ratings r, we are not currently equipped to relate the edge flow of a favorite free tournament to a lower dimensional representation. In Subsection 3.2.2 we will show that a favorite free tournament has edge flows f which can always be represented as a sum of cyclic intensities (or vorticities) on a set of loops. This result will parallel the conclusions of Lemma 6.1.

3.2. The Discrete HHD. Given these two classes of tournaments it is natural to ask: can a generic tournament be decomposed into a perfectly transitive (arbitrage free) part and a perfectly cyclic (favorite free) part? We answer in the affirmative. This is the Helmholtz-Hodge decomposition.

3.2.1. Operators. In order to define the decomposition succinctly, it is helpful to have a pair of operators analogous to the gradient and curl operators in the continuum. We simplify the topological presentation in [32] by expressing the decomposition entirely through linear algebra. For a cohomological discussion see [50].

First, define the edge space $\mathbb{R}^E$, where E is the number of pairs i, j who could compete. Index each pair so that edge k points from competitor $j(k)$ to competitor $i(k)$. Note that this requires assigning each edge an arbitrary start and endpoint. Positive f indicates that the competitor at the end is expected to beat the competitor at the start, and negative f indicates the reverse. This is simply a sign convention.

Let the *discrete gradient* operator G be the $E \times V$ matrix which maps from $\mathbb{R}^V$ to $\mathbb{R}^E$ by setting:

$$[Gu]_k = u_{i(k)} - u_{j(k)}. \tag{3.7}$$

Then $g_{kh} = 1$ if $h = i(k)$, equals -1 if $h = j(k)$, and is zero otherwise. The matrix G is sometimes called the edge incidence matrix since it records the start and end point of each edge.

Notice that if r is a rating function on the nodes, then attempting to find r such that $r_i - r_j = f_{ij}$ is equivalent to looking for r such that $Gr = f$. Since any arbitrage free tournament admits a unique rating r satisfying $Gr = f$, the space of perfectly transitive competitive networks is equivalent to the space of competitive networks with edge flow f in the range of the gradient.[8]

The gradient transpose, G^T is the discrete divergence operator. The divergence maps from the space of edges to the space of nodes (competitors) such that:

$$[G^T f]_i = \sum_{j \in \mathcal{N}(i)} f_{ij}. \tag{3.8}$$

The neighborhood condition (3.5) is equivalent to requiring that $G^T f = 0$. That is, the space of favorite free tournaments is equivalent to the space of tournaments with edge flow f in the null space of the divergence.

[8]Assuming that the competitive network is connected, the gradient has a one-dimensional null-

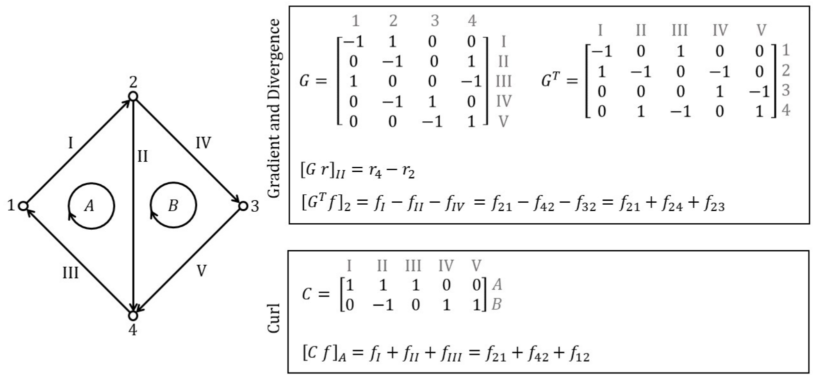

FIG. 4. *The gradient, divergence, and curl for the example networks in Figure 2. A spanning tree for networks of this form could consist of edges I, II, and IV. Then the edges III and V are the chords, and the associated loops are the triangles labelled A and B.*

In order to build a parallel description for perfectly cyclic tournaments, we need a space of loops. First define the sum of two cycles $\mathcal{C}_1$, $\mathcal{C}_2$ to be all edges included in either $\mathcal{C}_1$ or $\mathcal{C}_2$ but not both. Equipped with this addition operation, the space of cycles is a vector space, which can be represented with a cycle basis. A *cycle basis* is a collection of linearly independent cycles $\mathcal{C}_1, \mathcal{C}_2, \ldots, \mathcal{C}_L$ such that any other cycle $\mathcal{C}$ can be expressed as a linear combination of cycles in the collection [21].

Any connected graph admits a cycle basis. A simple construction follows. First, pick a spanning tree of the network. Then the spanning tree includes $V-1$ edges, and $E-(V-1)$ edges are left out. The latter are the *chords*. By construction, the tree does not contain any loops. If one chord is added to the tree then the network contains exactly one cycle. Note that no two chords can produce the same cycle, and that the set of cycles produced by adding the chords is necessarily linearly independent since no chord appears in more than one of these cycles. Let L be the number of chords. If we enumerate the chords from $1, 2, \ldots, L = E - V + 1$ then the set of cycles $\mathcal{C}_1, \ldots, \mathcal{C}_L$ associated with each chord is a cycle basis. The Figure 4 caption provides an example.

A basis generated by a spanning tree is a *fundamental cycle basis* [5, 21]. Cycle bases are rarely unique, since there are often many possible spanning trees, and not all bases are fundamental. An alternate basis for the network shown in Figure 4 could be the outer square consisting of edges I, IV, V and III, and either of the triangles.

Next, define the cycle space $\mathbb{R}^L$ to be the space of real vectors with one entry for each cycle in a chosen cycle basis. The dimension of the cycle space $L = E - V + 1$ is the *cyclomatic number* of the network [5, 21]. We define the *discrete curl* operator to be the matrix which maps from $\mathbb{R}^E$ to $\mathbb{R}^L$ (edges to cycles) by performing the path sum around each loop. If $\{k_1, k_2, \ldots, k_{|\mathcal{C}_l|}\} = \mathcal{C}_l$ then:

$$[Cf]_l = \sum_{k \in \mathcal{C}_l} f_k. \tag{3.9}$$

Note that in order to perform this sum, each loop must be assigned an arbitrary direction of traversal. This is another sign convention.

---

space spanned by the vector $[1; 1; \ldots 1]$. It follows that $G(r + c) = Gr$ if c is some constant. This motivates the constraint $\sum_i r_i = 0$ used throughout.

We limit our attention to curl operators such that there exists an invertible matrix T for which $TC = \tilde{C}$, where $\tilde{C}$ is the curl defined with respect to a fundamental basis.

This curl is analogous to the curl in continuous space, which is a path integral over infinitesimally small loops. Note that the discrete curl defined in this way is more general than the discrete curl defined in [32, 50], where the curl is restricted to act on triangles. Restricting the curl can lead to unintuitive conclusions. For example, if $p_{AB} = p_{BC} = p_{CD} = p_{DA} = 0.99$ then there is clearly a cyclic tendency in competition, but if the curl is restricted to only act on triangles, then the curl would be zero. Here we extend the curl to act on loops of arbitrary length since, like [68], we do not see a fundamental distinction between cyclic structure on triangles and cyclic structure on larger loops. If desired, we could partition the curl operator into blocks, each according to loops of a fixed length, and treat each block as the curl operator restricted to loops of a given size. In this way our approach is distinct from the approaches developed from cohomology, and is closer to the methods developed by Kirchoff to study electric circuits [5].

Figure 4 provides examples of these operators.

LEMMA 3.3 (Orthogonality). *The curl C and gradient G satisfy $CG = 0$.*

Proof. Consider the product CGu for some arbitrary vector $u \in \mathbb{R}^V$. The product Gu produces a perfectly transitive edge flow, so the product CGu evaluates the path sum of that edge flow around a set of loops. All perfectly transitive edge flows are arbitrage free, so the path sum of Gu over any loop is zero. It follows that $CGu = 0$ for all $u \in \mathbb{R}^V$ so:

$$CG = 0. \tag{3.10}$$ ∎

LEMMA 3.4. *Let f be an edge flow, C be a curl operator, and G be the gradient. If $Cf = 0$, then there exists a set of ratings r such that $Gr = f$.*

Proof. This Lemma is a direct consequence of Lemma 6.1. If C is a curl operator, then there exists an invertible transform T such that $C = T\tilde{C}$ where $\tilde{C}$ is the curl operator with respect to some fundamental cycle basis. Then $Cf = T\tilde{C}f = 0$ if and only if $\tilde{C}f = 0$. Since $\tilde{C}$ is defined with respect to a fundamental cycle basis, $\tilde{C}$ is defined with respect to a spanning tree $\mathcal{T}$ which generates the cycle basis. Requiring that $\tilde{C}f = 0$ is equivalent to requiring that the sum of f around every loop formed by adding one chord into the tree is zero. This condition is sufficient to reconstruct r such that $Gr = f$ using the spanning tree construction given in the proof of Lemma 6.1, where the chosen tree is $\mathcal{T}$. ∎

Lemma 3.3 and Lemma 3.4 establish that, if the edge flow is the gradient of some set of ratings then its curl is zero, and if the curl of the edge flow is zero then it can be expressed as the gradient of some set of ratings. Therefore the range of the gradient is the nullspace of the curl. The equivalence of these two spaces and the orthogonality of the operators allows us to decompose f into unique perfectly transitive and perfectly cyclic components. This decomposition is the HHD.

3.2.2. The Discrete Helmholtz-Hodge Decomposition.

THEOREM 3.5 (The HHD). *Any $f \in \mathbb{R}^E$ can be decomposed such that:*

$$f = f_t + f_c \tag{3.11}$$

where f_t is arbitrage free (perfectly transitive) and f_c is favorite free (perfectly cyclic) and both are unique. In addition, there exists a unique rating r satisfying $\sum_i r_i = 0$

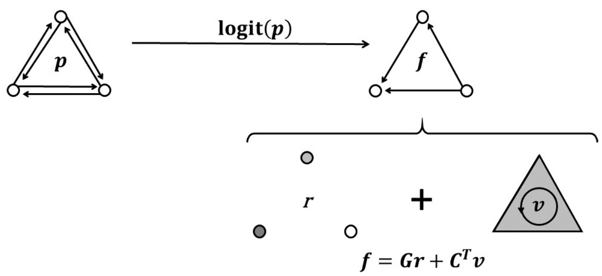

FIG. 5. *A schematic representation of the decomposition for a complete tournament on three competitors. The edge flow f is set equal to logit(p), and then broken into a set of ratings r and vorticities v, such that $f = Gr + C^T v$.*

such that $f_t = Gr$ and, for any choice of C, a unique vorticity $v \in \mathbb{R}^L$ exists such that $f_c = C^T v$. Thus the original edge flow f can be uniquely decomposed:

$$f = Gr + C^T v. \tag{3.12}$$

Proof. By the fundamental theorem of linear algebra $\mathbb{R}^E = \text{null}(C) \oplus \text{range}(C^\intercal)$ [74]. Lemma 3.3 and Lemma 3.4 guarantee that $\text{range}(G) = \text{null}(C)$, so:

$$\mathbb{R}^E = \text{range}(G) \oplus \text{range}(C^T). \tag{3.13}$$

Thus any edge flow can be uniquely decomposed into the sum of a perfectly transitive and perfectly cyclic edge flow, and those edge flows are the projections of f onto the perfectly transitive and cyclic subspaces.

Equation (3.13) establishes that there exists an r such that $Gr = f_t$, and a v such that $C^T v = f_c$. We have already proved r was unique. Equation (3.13) guarantees $E = \text{rank}(G) + \text{rank}(C^T)$. In general, G has nullity equal to the number of connected components in the network. We assumed the network is connected, so G has a one-dimensional nullspace and rank $V - 1$. Therefore, $\text{rank}(C^T) = E - (V - 1) = L$. By construction, C^T has L columns, so is full rank. It follows that the linear system $C^T v = f$ has a unique solution if $f \in \text{range}(C^T)$. ∎

Therefore, any arbitrary tournament can be decomposed into a perfectly transitive and a perfectly cyclic tournament, where the perfectly transitive tournament is specified by a set of ratings, and the perfectly cyclic tournament is specified by a set of vorticities. The ratings associated with the HHD are the Hodge ratings proposed by [32]. Figure 5 provides a schematic representing the decomposition.

The three example networks displayed in Figure 2 are actually an example of an HHD. Reading left to right, the first network is perfectly transitive, the second is perfectly cyclic, and they add to produce the generic network shown on the right. The edge flows, ratings r, and vorticities v are shown in Figure 6.

The gradient G has exactly 2 nonzero entries per edge, so it becomes sparser as the number of competitors increases. Consequently, the decomposition can be performed efficiently, even for large, fully connected networks. Methods are discussed in [10, 32].

The intransitivity measure associated with the HHD is the size of the cyclic component $||f_c||_2$. Because the HHD is a decomposition onto orthogonal subspaces, this measure is equal to the distance from f to the closest perfectly transitive tournament. Therefore the Helmholtz-Hodge intransitivity measure is conceptually analogous to the Slater intransitivity measure [68], and its variants [57], [70], [77]. Similarly, the

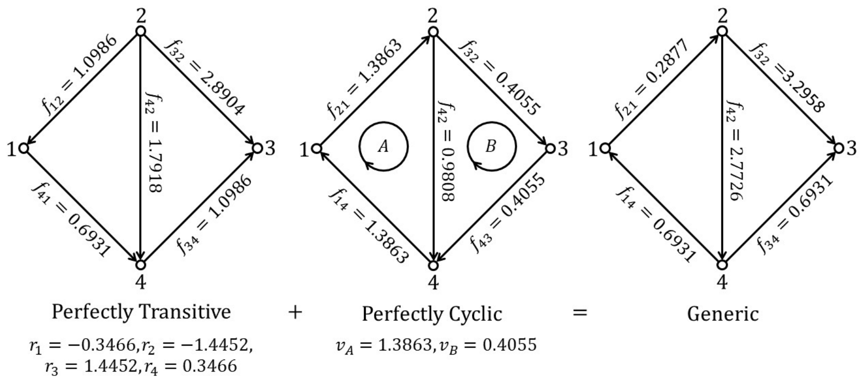

Fig. 6. *An example HHD using the three networks from Figure 2. From left to right: the leftmost network is perfectly transitive, the middle network is perfectly cyclic, and the network on the right is the sum of the perfectly transitive and cyclic networks. The ratings associated with the perfectly transitive graph are provided beneath it. Notice that the difference in the ratings recover the edge flow on each edge. For example,* $r_3 - r_4 = 1.4452 - 0.3466 = 1.0986 = f_{34}$. *Also notice that the curl of the edge flow around any loop is zero. For example,* $f_{41} + f_{12} = 0.6931 + 1.0986 = 1.7918 = f_{42}$ *so* $f_{41} + f_{12} + f_{24} = f_{41} + f_{12} - f_{42} = 0$. *The vorticities associated with the perfectly cyclic network are provided beneath it. Notice that the perfectly cyclic edge flow satisfies the neighborhood condition. For example, the total flow into node* 2 *is* $1.3863 - 0.9808 - 0.4055 = 0$. *Finally, notice that the values of the edge flow in the rightmost network are the sum of the edge flows in the perfectly transitive and cyclic networks. For example, looking at the edge connecting nodes* 1 *and* 2, $-1.0986 + 1.3863 = 0.2877$.

transitivity measure associated with the HHD is the size of the transitive component $||f_t||_2$, and is the distance from f to the closest perfectly cyclic tournament.

Note that these measures are continuous in p. In contrast, classical methods such as the Kendall [35] or Slater [68] measures only depend on $\mathcal{G}_{\rightarrow}$ so are discrete in p. This distinction is important, since it means that the Helmholtz-Hodge measure distinguishes between the cases $p_{AB} = p_{BC} = p_{CA} = 0.99$ and $p_{AB} = p_{BC} = p_{CA} = 0.51$ (intransitivity 7.96 and 0.07 respectively). Using the discrete measures, these two tournaments are equally intransitive. Thus the Helmholtz-Hodge measure reflects the absolute strength of cyclic competition by distinguishing strong and weak cycles. The discrete measures reflect the relative strength of cyclic competition since they only depend on the sign of f, which depends on both f_c and f_t. If the transitive part is large then it may mask weaker cyclic competition when using a discrete measure. For example, if $p_{AB} = 0.99, p_{BC} = 0.99$ and $p_{CA} = 0.49$ then the probability that C beats A is much larger than might be expected. However, in this example competition is transitive so all discrete measures of intransitivity would return their minimal value, 0. In contrast, the Helmholtz-Hodge measure returns intransitivity 5.29. These examples are illustrated in Figure 7. Normalizing the Helmholtz-Hodge measures by $||f||_2$ produces the equivalent relative measures: $||f_c||_2/||f||_2$ and $||f_t||_2/||f||_2$.

3.2.3. Equivalent Formulations. Here we present six different approaches that arrive at the same decomposition. These provide different, useful, perspectives on the HHD, and illustrate that it is robust to varying motivations. The ensuing Corollary follows directly from standard properties of projection onto orthogonal subspaces, so we omit the proof.

Corollary 3.6 (Equivalent Formulations). *The following six decompositions are equivalent:*

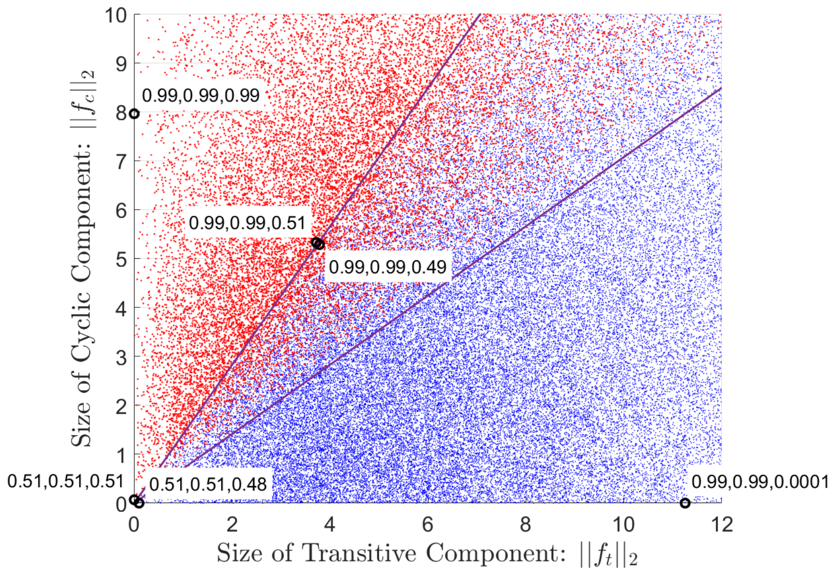

FIG. 7. *Transitivity and intransitivity of* 10^4 *triangular networks with randomly drawn win probabilities. The horizontal axis is the size of the transitive component and the vertical axis is the size of the cyclic component. Each scatter point is a sampled network. Smaller blue scatter points are transitive, larger red points are intransitive. The upper and lower purple lines (slope* $\sqrt{2}$ *and* $\sqrt{0.5}$*) divide regions where competition on triangles is always cyclic, either transitive or cyclic, and always transitive. The large black circles represent example networks. The text next to each example gives the probability A beats B, B beats C, and C beats A. If all of these numbers are greater than* 0.5 *then the network is intransitive. Note that the classification into transitive and intransitive draws a sharp distinction between networks whose win probabilities are nearly identical, while networks with similar win probabilities remain close to each other when using the Hodge measures. Also note that the boundary between transitive and intransitive networks is an angular sector, hence this classification is based on the relative sizes of the transitive and cyclic components, not their absolute sizes. In contrast the Hodge measures reflect the absolute size of each component. Thus the example with win probabilities* $0.99, 0.99, 0.49$ *can be transitive and the example* $0.51, 0.51, 0.51$ *can be intransitive, even though the former has a larger cyclic component than the latter.*

1. $f = f_t + f_c$ *where* f_t *is arbitrage free and* f_c *is favorite free;*
2. $f = f_t + f_c$ *where* $f_t = Gr$ *for ratings* r *and* $f_c = C^T v$ *for vorticity* v*;*
3. *the ratings* r *solve the constrained least squares problem:*

$$\text{Minimize } {||Gu - f||_2}^2 \quad \text{given } u \in \mathbb{R}^V \text{ and } \sum_{i=1}^{V} u_i = 0 \tag{3.14}$$

   and $f_t = Gr, f_c = f - f_t$*;*
4. *the vorticities* v *solve the least squares problem:*

$$\text{Minimize: } {||C^T w - f||_2}^2 \quad \text{given } w \in \mathbb{R}^L \tag{3.15}$$

   and $f_c = C^T v, f_t = f - f_c$*;*
5. $f = f_t + f_c$ *where* $f_t = Gr$ *for the unique ratings* r *such that the circulant* $f - f_t$ *is favorite free;*
6. $f = f_t + f_c$ *where* $f_c = C^T v$ *for the unique vorticities* v *such that* $f - f_c$ *is arbitrage free.*

Each of these approaches provides a different perspective on the HHD. We might seek to decompose f into components that do not circulate and do not converge, into components defined by a set of ratings and vorticities, according to the best perfectly transitive or perfectly cyclic approximation, so that the residue left over when approximating f does not circulate, or so that the residue left over when approximating f does not converge anywhere. In each case the resulting decomposition is the same. The fact that the HHD is equivalent to all of these approaches motivates its use.

It is worth highlighting the third and fourth approach, which show that f_t is the nearest perfectly transitive edge flow to f, and f_c is the nearest perfectly cyclic edge flow to f. Decomposition 3 shows that the ratings produced by the HHD are a type of least squares rating. Least squares ratings methods are widely used [6, 14, 34, 41, 42, 45, 51, 72, 73]. Although the literature has focused almost exclusively on Decomposition 3, Decompositions 3 and 4 are dual to one another. This parity in approach sets the HHD apart from existing methods.

4. Null Models and the Trait-Performance Theorem.

How intransitive is a typical tournament?

Answering this question requires defining a statistical model for sampling tournaments - in particular, for sampling edge flows. How do assumptions about the distribution of possible edge flows affect the expected strength of cyclic competition? What statistical features tend to promote or suppress cyclic competition?

We initially explore these questions for a generic null model where the edge flow, F, is sampled randomly from an unspecified distribution. This analysis identifies which features of the edge flow and the network topology influence the degree of cyclic competition. These conclusions set the stage for the following insight.

If the edge flow is sampled using a trait-performance model, then the covariance of the edge flow takes on a canonical form which depends only on *two* statistical quantities: the variance in the flow on each edge, and the correlation in the flow on pairs of edges that share an endpoint. This simplified structure leads to an elegant closed form expression for the expected sizes of the cyclic and transitive components that separates the influence of the network topology from the trait-performance statistics.

We generalize this result in two ways. First, the relations between correlation and network structure derived under the trait-performance assumptions hold for any complete network - whether or not the trait-performance assumptions are valid. Second, we show that the canonical form for the covariance can be used to design null models for tournaments with tunable transitive structure. These models can be easily adjusted to promote or suppress cycles, and could be used to define more nuanced transitivity tests than the standard randomization tests [2, 15, 35].

4.1. Generic Null Models.

We start by considering generic null models where the edge flow $F \in \mathbb{R}^E$ is drawn randomly from some distribution. For now we introduce no assumptions on the distribution other than that it has finite first and second moments. Denote the expected edge flow $\bar{f} = \mathbb{E}[F]$ and the covariance $\mathrm{Cov}(F) = \mathbb{E}[(F-\bar{f})(F-\bar{f})^T]$.

Let P_c be the orthogonal projector onto the space of perfectly cyclic (favorite free) tournaments. Then the expected squared strength of cyclic competition is:

$$\begin{aligned} \mathbb{E}[||F_c||^2] &= \mathbb{E}[F^T P_c^T P_c F] = \mathbb{E}[F^T P_c F] = \sum_{kl} (P_c)_{kl}\, \mathbb{E}[F_k F_l] \\ &= \sum_{kl} (P_c)_{kl} \,(\bar{f}_k \bar{f}_l + \mathrm{Cov}(F)_{kl}) = ||\bar{f}_c||^2 + \mathrm{trace}(P_c \mathrm{Cov}(F)) \end{aligned} \tag{4.1}$$

where $||\bar{f}_c||^2 = \bar{f}^T P_c \bar{f}$ is the cyclic component of the expected edge flow.

Therefore, no matter the underlying distribution of edge flows, the expected strength of cyclic competition is determined exclusively by three quantities: the *expected edge flow*, the *covariance in the edge flow*, and the *topology of the network* (which determines P_c).

The matrix inner product, $\text{trace}(P_c \text{Cov}(F))$, can be simplified if the flows on each edge are independent. Then $\text{Cov}(F)$ is diagonal with entries $\sigma_k^2 = \mathbb{E}[(F_k - \bar{f}_k)^2]$. It follows that $\text{trace}(P_c V) = \sum_{k=1}^{E} (P_c)_{kk} \sigma_k^2$.

The nonzero eigenvalues of a projector all equal one, so its trace equals the dimension of the space it projects onto. The projector P_c projects onto the space of perfectly cyclic tournaments, which has dimension $L = E - (V-1)$. Therefore $\sum_k (P_c)_{kk} = L$. Rewrite the expected strength of cyclic competition:

$$\mathbb{E}[||F_c||^2] = ||\bar{f}_c||^2 + L \sum_{k=1}^{E} \left(\frac{(P_c)_{kk}}{L} \right) \sigma_k^2. \tag{4.2}$$

Since the diagonal entries of an orthogonal projector are always nonnegative, the right hand term can be interpreted as a weighted average of the variance on each edge. Therefore, when the edges are independent, the expected strength of cyclic competition is given by the strength of the cyclic component of the expected edge flow, plus the dimension of the loop space times a weighted average of the variance on each edge. Similarly, the expected strength of transitive competition is:

$$\mathbb{E}[||F_t||^2] = ||\bar{f}_t||^2 + (V-1) \sum_{k=1}^{E} \left(\frac{(P_t)_{kk}}{V-1} \right) \sigma_k^2 \tag{4.3}$$

and the expected total strength of competition is:

$$\mathbb{E}[||F||^2] = ||\bar{f}||^2 + E\bar{\sigma}^2 \tag{4.4}$$

where $\bar{\sigma}^2$ is the average of the variance in the flow on each edge. Equation (4.4) is valid even if the edges are not independent, as the projector onto the full space is simply the identity.

Equations (4.2) - (4.4) show that the contribution to the expected strength of competition from the variances is not distributed equally between the transitive and cyclic spaces. Instead, the amount that is cyclic is proportional to the number of cycles, while the amount that is transitive is proportional to the number of competitors. As a result, adding edges to a network will typically increase the expected degree to which competition is cyclic. It follows that sparse networks with randomly drawn edge flows will be relatively more transitive than would be expected given $\bar{f}$, while dense networks will typically be more cyclic. It also follows that, for a posterior distribution of possible edge flows given observed data, uncertainty will likely lead to an overestimate of the degree to which competition is cyclic when the network is dense. If a tournament is complete, then $E = V(V-1)/2$ so $(V-1)/E = 2/V$ and $L/E = 1 - 2/V$. It follows that, for a complete tournament with more than four competitors, any uncertainty in the edge flow will typically bias competition to appear more cyclic than transitive.[9]

[9]This result does not contradict Shizuka's result that the proportion of transitive triangles in a network with uniformly randomly sampled dominance relations is independent of the network topology [63], since our measure accounts for the global structure of the edge flow, thus incorporates cyclic structure over longer cycles.

Numerical studies have suggested that filling in missing edges with randomly drawn F typically overestimates the degree to which competition is cyclic, thereby weakening transitivity tests [63]. Our result provides a clear explanation for this observation. When the edge flow F is drawn randomly to fill in missing data, it is usually drawn independently and identically distributed, cf. [15]. If edges are added until the network is complete, then, for any tournament with more than four competitors, the resulting "imputed" tournament will likely be significantly more cyclic than the original tournament. Therefore, unless the edge flows are well modeled by assuming that the F_k are independent and identically distributed, *and* that all pairs of competitors could compete, this procedure is not valid for estimating the strength of cyclic competition in a partially observed tournament. This observation underscores the need for intransitivity measures that can be applied to incomplete tournaments.

Unfortunately the projectors P_t and P_c may be expensive to compute, and cannot always be constructed directly without performing a matrix decomposition. This makes it challenging to identify exactly how the topology of the network and covariance structure promote or suppress cyclic competition. Nevertheless, as we show in the next section, using a more principled model for sampling F, ensures that the covariance matrix $\text{Cov}(F)$ takes on a canonical form. This form clarifies the interaction between the topology of the network and the distribution of edge flows.

4.2. Trait-Performance. The outcomes of real-world competition events are typically influenced by a constellation of underlying competitor traits. Examples of trait-influenced competition abound, ranging from sports[10] to simulated competitive events to biology.[11] In some cases, trade-offs inherent in certain traits have been observed to lead to cyclic competition between organisms [36, 66].[12] In such examples, trade-offs lead to advantages against certain opponents, and weaknesses that are exploited by others. In evolutionary biology, trade-offs of this kind challenge the notion that members of intransitive communities can be consistently ranked according to fitness. Intransitivity can lead to deeply counterintuitive evolutionary dynamics [20, 33], and may promote biodiversity since no single species has an absolute advantage over all competitors [59, 58, 60, 61, 70]. These considerations motivate a study of how the

[10]Some predictive tennis models estimate the probability that one competitor will beat another based on a parameterized model for the probability that each player will win a point, where the underlying parameters depend on traits of the players [40]. Similarly, considerable effort has been devoted to predictive models for baseball based on team and player statistics [78].

[11]Ecological studies of competition for dominance in social hierarchies have analyzed how traits confer success, because selection acts on heritable traits contributing to reproductive success. Examples include competition among male northern elephant seals [26] and male Cape dwarf chameleons [76]. Relevant traits for elephant seals include body mass, length, age, and time of arrival on the beach [26]. Relevant traits for chameleons include body mass, length from snout to base of tail, length of the tail, jaw length, head width, casque size, and size of a pink colored flank patch used in signaling [76].

[12]Two particularly famous examples are side-blotched lizards and colicin producing *E. coli* [36, 66]. In the former example, large orange-throated males maintain large territories, medium blue-throated males defend small territories, while small yellow-throated 'sneaker' males resemble females and do not maintain territories. Orange-throated males typically defeat the smaller blue-throated males, who defeat the even smaller yellow throated males, who defeat the orange throated males by sneaking into their territories [66]. In the latter example, three strains of *E. coli* were grown in direct competition in a laboratory setting. The first strain produced a colicin toxin, the second was susceptible to the toxin, and the third was resistant to the toxin but not toxin-producing. In the absence of the resistant strain, the toxic strain could outcompete the susceptible strain. In the absence of the toxic strain, the susceptible strain could outcompete the resistant strain, which reproduced more slowly because resistance is costly. But, in the absence of the susceptible strain, the resistant strain could outcompete the toxic strain by reproducing more quickly [36].

distribution of traits, and the way traits confer success, either promote or suppress cyclic competition.

To study this scenario, suppose that win probabilities p can be modeled as a function of some underlying traits x of each competitor. Let $X(i) = [X_1(i), \ldots, X_T(i)]$ denote the T randomly sampled traits of the i^{th} competitor. Then let $f(x, y)$ be a performance function, such that $f(x, y)$ is the log-odds that a competitor with traits x would beat a competitor with traits y.

To construct a trait-performance model assume that:

1. The trait vectors of the competitors are drawn independently and identically from a trait distribution π_x.
2. There exists a performance function $f(x, y)$ that maps from $\mathbb{R}^T \times \mathbb{R}^T$ to $\mathbb{R}$. We require that the performance function is alternating, $f(x, y) = -f(y, x)$, for any trait vectors x and y in the support of π_x. This ensures that f can be used to generate an edge flow. It also ensures that the performance function is fair, $\mathbb{E}[f(X, Y)] = 0$, since when X and Y are drawn i.i.d then $\mathbb{E}[f(X, Y)] = \mathbb{E}[f(Y, X)] = -\mathbb{E}[f(X, Y)]$ which implies $\mathbb{E}[f(X, Y)] = 0$.
3. There exists a connected competitive network $\mathcal{G}_{\rightleftarrows}$ with edges representing possible competition events, and the network is either fixed a priori or sampled independently from the traits.

The first assumption holds if all competitors are drawn from the same trait pool. Different pools can be incorporated into the model by adding a trait which indexes which pool each competitor is sampled from, provided that trait can be sampled independently of the graph. For example, Major League Baseball team budgets vary widely. In 2018 the Yankees' total value was over 4.6 billion dollars, which was more than the total value of the bottom six teams combined [56]. This difference in resources gives high value teams the opportunity to pay higher salaries[13] and attract stars. Thus wealth could be incorporated as a trait.

The second assumption is valid whenever the probability that one competitor beats another can be conditioned on the traits of the competitors, independent of their location on the network, and of the outcomes of past events. Note that in some biological contexts, such as social hierarchies, event outcomes are not necessarily independent, and may be influenced by past events. For example, winner, loser, and bystander effects, in which winners are more likely to win again, losers are more likely to lose again, and bystander behavior is influenced by observed events between other competitors, play an important role in the self-organization of certain social hierarchies [12, 13, 28, 55, 65]. The assumption that competition outcomes are mediated by traits is also not supported in convention based societies where rank is determined by a social convention, such as matrilineal rank inheritance (c.f. [69, 75]). Nevertheless, other hierarchies can be explained by traits (c.f. [31, 62]), and even in situations when competition outcomes are influenced by past events, competitor attributes typically influence competition outcomes as well [4, 13].

The third assumption treats the network topology (who competes with whom) as independent from the traits of the competitors. This may not be realistic if competitors avoid competing when they are likely to lose [67]. This also limits our ability to model systems where traits or rank are heritable (c.f. [69, 75]), or distributed differently across different clusters of competitors (different divisions, or local populations).

While these assumptions do not hold in all situations, they provide a tractable paradigm that lays the foundation for a more general understanding.

[13]For example, in 2019 the Yankees' combined payroll was three times larger than the Marlins'.

Under assumptions 1-3, we define a trait-performance model as follows. First, sample $X(i) \sim \pi_x$ for all competitors i. Then, set $F_k = f(X(i(k)), X(j(k)))$, where $i(k), j(k)$ are the endpoints of edge k.

THEOREM 4.1 (Trait-Performance). *Let $\mathcal{G}_{\rightleftarrows}$ be a competitive network with V competitors, E edges and L loops, satisfying assumption* 3. *If the traits of each competitor are drawn independently from π_x, and the edge flow is defined by $F_k = f(X(i(k)), X(j(k)))$ where $f(x, y)$ is an alternating performance function, then the covariance* $\mathrm{Cov}(F)$ *of the edge flow has the form:*

$$\mathrm{Cov}(F) = \sigma^2 \left[I + \rho\left(GG^T - 2I\right)\right] \tag{4.5}$$

where σ^2 is the variance in F_k for arbitrary k, and ρ is the correlation coefficient between $f(X, Y)$ and $f(X, W)$ for X, Y, W drawn i.i.d from π_x.

Moreover:

$$\mathbb{E}\left[\frac{1}{E}||F||^2\right] = \sigma^2 \xrightarrow{decompose} \begin{cases} \mathbb{E}\left[\frac{1}{E}||F_t||^2\right] = \sigma^2\left[\frac{(V-1)}{E} + 2\rho\frac{L}{E}\right] \\ \mathbb{E}\left[\frac{1}{E}||F_c||^2\right] = \sigma^2\left(1 - 2\rho\right)\frac{L}{E} \end{cases} \tag{4.6}$$

The correlation ρ ranges from 0 to $1/2$, and if $\rho = 1/2$ then competition is perfectly transitive.

Proof. First consider the covariance matrix $\mathrm{Cov}(F)$.

Since the trait vectors are drawn i.i.d from the trait distribution, the diagonal entries of the covariance are given by:

$$\mathrm{Cov}(F)_{kk} = \mathbb{E}\left[f(X(i(k)), X(j(k)))^2\right] = \mathbb{E}\left[\left(f(X, Y)\right)^2\right] \equiv \sigma^2 \tag{4.7}$$

where X, Y are drawn i.i.d from the trait distribution, and σ^2 is the variance in $f(X, Y)$. Thus, the diagonal entries of the covariance are identical.

The off-diagonal entries are $\mathbb{E}\left[f(X(i(k)), X(j(k))) \cdot f(X(i(l)), X(j(l)))\right]$.

Suppose the edges k and l do not share an endpoint. Then $i(k) \neq i(l)$ or $j(l)$ and $j(k) \neq i(l)$ or $j(l)$. Then $f(X(i(k)), X(j(k)))$ is a function of two random vectors, and $f(X(i(l)), X(j(l)))$ is a function of two other random vectors, where the pair of random vectors are independent. It follows that $f(X(i(k)), X(j(k)))$ is independent of $f(X(i(k)), X(j(k)))$. Then, since competition is fair for all alternating performance functions, $\mathrm{Cov}(F)_{kl} = \mathbb{E}\left[f(X(i(k)), X(j(k))) \cdot f(X(i(l)), X(j(l)))\right] = \mathbb{E}\left[f(X(i(k)), X(j(k)))\right]\mathbb{E}\left[f(X(i(l)), X(j(l)))\right] = 0$. It follows that the support of the covariance matches the adjacency structure of the edges of the competition network.

If the edges do share an endpoint, then there are four possibilities. Either $i(k) = i(l)$, $j(k) = j(l)$, $i(k) = j(l)$, or $j(k) = i(l)$. We say that the edges are *consistently oriented* if they share either the same starting point or the same ending point, and are *inconsistently oriented* if the endpoint of one is the start of another. Since all the trait vectors are drawn i.i.d., we suppress the indices and let the three trait vectors Y, W, Z be drawn i.i.d. from π_x. The performance function is alternating, so:

$$\begin{aligned} &\mathbb{E}[f(Y, W)f(Y, Z)] = \mathbb{E}[f(W, Y)f(Z, Y)] \equiv \rho\sigma^2 \\ &\mathbb{E}[f(Y, W)f(Z, Y)] = \mathbb{E}[f(W, Y)f(Y, Z)] = -\mathbb{E}[f(Y, W)f(Y, Z)] = -\rho\sigma^2 \end{aligned} \tag{4.8}$$

where ρ is the correlation coefficient between $f(Y, W)$ and $f(Y, Z)$. Notice that a positive correlation indicates that the probability that A beats B is increased by conditioning on the event that A beats C.

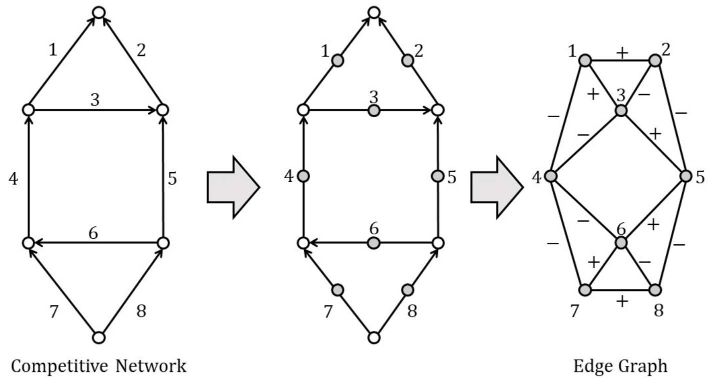

FIG. 8. *The edge graph (right) associated with a competitive network (left). The middle panel shows an intermediate graph where a node has been introduced for each edge. The edges of the competitive network become the nodes of the edge graph. The edges of the edge graph correspond to nodes in the competitive network that are the shared endpoint of a pair of edges. These are labelled with a + or − to indicate whether the edges are consistently or inconsistently oriented with respect to the shared endpoint.*

The *edge graph* is the graph with a node for each edge in the competition network, and with an undirected edge between nodes corresponding to connected edges in the competition network (Figure 8). Let A_E be the weighted adjacency matrix for the edge graph with $a_{Ekl} = +1$ or -1 if edges k and l are consistently or inconsistently oriented with respect to a shared endpoint. Then:

$$\mathrm{Cov}(F) = \sigma^2 \left[I + \rho A_E\right]. \tag{4.9}$$

The weighted adjacency matrix A_E for the edge graph is equal to $GG^T - 2I$ since:

$$[GG^T]_{kl} = (e_{i(k)} - e_{j(k)})^T(e_{i(l)} - e_{j(l)}) = \left\{\begin{matrix} 2 \text{ if } k = l \\ 1 \text{ if } i(k) = i(l) \text{ or } j(k) = j(l) \\ -1 \text{ if } i(k) = j(l) \text{ or } j(k) = i(l) \\ 0 \text{ else} \end{matrix}\right\} \tag{4.10}$$

where $e_i \in \mathbb{R}^V$ is the indicator vector for node i. Thus we establish (4.5).

All of the absolute measures of the strength of competition (squared) are given by the squared length of the orthogonal projection of the edge flow onto some subspace. Let P_S be an arbitrary orthogonal projector onto some subspace S. By construction, the edge flow is zero mean, therefore, by equation (4.1), the expected value of the associated measure is:

$$\mathbb{E}\left[||F_S||^2\right] = \mathrm{trace}(P_S \mathrm{Cov}(F)). \tag{4.11}$$

The intensity of competition, $||F||^2$, corresponds to the projector I, $||F_t||^2$ corresponds to the projector P_t, and $||F_c||^2$ corresponds to the projector P_c. Then, by equation (4.11):

$$\mathbb{E}\left[\frac{1}{E}||F||^2\right] = \frac{1}{E}\mathrm{trace}(\mathrm{Cov}(F)) = \frac{E}{E}\sigma^2 = \sigma^2. \tag{4.12}$$

This formula establishes that the absolute strength of competition only depends on the variance σ^2 in each individual performance function.

To compute $||F_t||^2$, use equation (4.11) with projector P_t:

$$\mathbb{E}\left[\frac{1}{E}||F_t||^2\right] = \frac{1}{E}\text{trace}(P_t\text{Cov}(F)) = \frac{\sigma^2}{E}\text{trace}\left(P_t[I+\rho(GG^T-2I)]\right)$$
$$= \frac{\sigma^2}{E}\text{trace}\left(P_t\right) + \frac{\rho\sigma^2}{E}\text{trace}\left(P_t(GG^T)\right) - \frac{2\rho\sigma^2}{E}\text{trace}\left(P_t\right). \tag{4.13}$$

The trace of an orthogonal projector equals the dimension of the subspace it projects onto, so trace$(P_t) = V-1$. The range of GG^T is in the range of G, which is the subspace P_t projects onto. It follows that $P_tGG^T = GG^T$ so trace$(P_tGG^T) =$ trace$(GG^T) = 2E$ (see equation (4.10)). Therefore:

$$\mathbb{E}\left[\frac{1}{E}||F_t||^2\right] = \sigma^2\left[\frac{V-1}{E} + 2\rho\frac{E-(V-1)}{E}\right] = \sigma^2\left[\frac{V-1}{E} + 2\rho\frac{L}{E}\right]. \tag{4.14}$$

Since $L \geq 0$, $\mathbb{E}[\frac{1}{E}||F_t||^2]$ increases monotonically in ρ: the larger ρ, the more A beating B is correlated with A beating C, implying transitive competition.

Then, by the orthogonality of the decomposition $f = f_c + f_t$:

$$\mathbb{E}\left[\frac{1}{E}||F_c||^2\right] = \mathbb{E}\left[\frac{1}{E}||F||^2\right] - \mathbb{E}\left[\frac{1}{E}||F_t||^2\right] = \sigma^2\left[1-2\rho\right]\frac{L}{E}. \tag{4.15}$$

It follows that the expected absolute strength of cyclic competition is monotonically decreasing in the correlation coefficient ρ. Note that, as when considering the generic null models, dense networks promote cyclic competition.

To conclude, we show that $\rho \in [0, 1/2]$, so the expected measures are maximized and minimized when ρ is 0 or 1/2, respectively.

The correlation ρ is nonnegative since W and Z are i.i.d., thus $f(y, W)$ and $f(y, Z)$ are also i.i.d. for all y. Then:

$$\sigma^2\rho = \mathbb{E}_{Y,W,Z}[f(Y,W)f(Y,Z)] = \int_{\mathbb{R}^T}\mathbb{E}_{W,Z}[f(y,W)f(y,Z)]\pi_x(y)dy$$
$$= \int_{\mathbb{R}^T}\mathbb{E}_W[f(y,W)]\mathbb{E}_Z[f(y,Z)]\pi_x(y)dy = \int_{\mathbb{R}^T}\mathbb{E}_W[f(y,W)]^2\pi_x(y)dy \geq 0. \tag{4.16}$$

Here expectation is taken with respect to the variables in the subscript.

To prove that $\rho \leq 1/2$, note that all covariance matrices are positive semi-definite, so, for any vector u:

$$u^T\text{Cov}(F)u = \sigma^2u^T(I+\rho(GG^T-2I))u = \sigma^2(1-2\rho)||u||^2 + \rho u^TGG^Tu \geq 0. \tag{4.17}$$

If $E > V-1$, then the network has at least one loop, so the range of C^T is non-empty, hence the null-space of G^T is non-empty. Choosing u perfectly cyclic sets $G^Tu = 0$ so $\sigma^2(1-2\rho)||u||^2 \geq 0$ which requires $\rho \leq \frac{1}{2}$. If $E = V-1$ then the network is a tree, so all competition is necessarily perfectly transitive.

It follows that the expected absolute strength of *transitive* competition is minimized when $\rho = 0$, and maximized when $\rho = 1/2$. In contrast, the expected strength of *cyclic* competition is maximized when $\rho = 0$, and minimized when $\rho = 1/2$.

If $\rho = 1/2$ then $\mathbb{E}[||F_c||^2] = 0$. The measure is nonnegative for all edge flows. Therefore, its expected value is only zero if the probability that $||F_c||^2 \neq 0$ is zero.

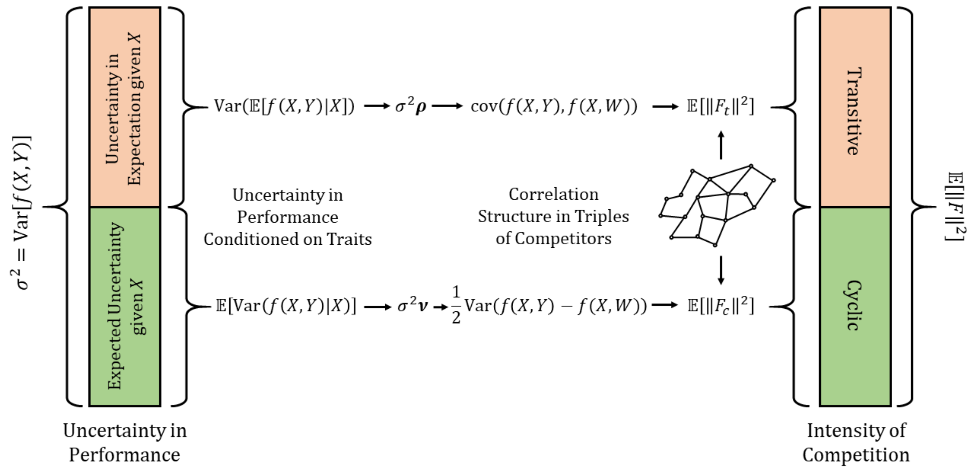

FIG. 9. *A schematic representing the conclusions of Theorem 4.1 and Corollary 9.1. The left hand side decomposes the uncertainty in performance into the uncertainty in the expected performance given X, and the expected uncertainty in the performance, given X. These uncertainties are converted into ρ and ν which describe the correlation structure of triples of competitors. The sizes of ρ and ν, plus the topology of the network, determine the expected sizes of the transitive and cyclic components. Thus we convert a decomposition of the uncertainty in the performance into a decomposition of the intensity of the edge flow representing competition.*

In this case, the tournament is arbitrage free. It follows that, if $\rho = 1/2$, then the tournament must be perfectly transitive.[14] □

Theorem 4.1 establishes that the expected degree to which competition is transitive or cyclic depends principally on the density of the network, and the correlation structure of F. In particular, the degree to which a network is cyclic or transitive depends on the correlation between the performance of A against B with the performance of A against C. The larger this correlation, the more consistently each competitor performs, hence the more consistent the network is with a set of ratings.

The variance σ^2 and the correlation coefficient ρ could be computed given an assumed trait distribution π_x and performance function $f(x, y)$. This could be done analytically if π_x and f lead to simple calculations. Otherwise, σ^2 and ρ can be approximated numerically by sampling or quadrature. The analytic method follows.

Suppose that X, Y are drawn from a sample space Ω which is a subset of $\mathbb{R}^T$. Then, for trait distribution π_x:

$$\rho = \frac{\mathbb{E}_X[\mathbb{E}_Y[f(X,Y)]^2]}{\mathbb{E}_{X,Y}[f(X,Y)^2]} = \frac{\int_\Omega \left(\int_\Omega f(x,y)\pi_x(y)dy\right)^2 \pi_x(x)dx}{\int_\Omega \int_\Omega f(x,y)^2 \pi_x(y)\pi_x(x)dydx}. \tag{4.18}$$

Note that the correlation coefficient is only large if it is possible to find some set of traits which are expected to perform either well or poorly on average, and if these

[14]Note that $\rho = 1/2$ guarantees perfect transitivity but $\rho = 0$ does not guarantee that the tournament is perfectly cyclic. A counterexample suffices to explain why. Suppose each competitor chooses rock, paper, or scissors uniformly and independently. Suppose there are three competitors and the tournament is complete. Then, in order for the tournament to be perfectly cyclic, rock must be chosen by one competitor, scissors by another, and paper by the last. There are 6 ways this can happen but 27 possible tournaments, so there is a 21/27 chance the tournament is perfectly transitive. Note that if the network is dense and $\rho = 0$ the network may be predominantly, if not perfectly, cyclic.

traits occur sufficiently often. That is, there must be some x such that $|\mathbb{E}_Y[f(x,Y)]|$ is large, and $\pi_x(x)$ is not too small. From this expression, it is not surprising that the expected strength of transitive competition is monotonically increasing in ρ. If there is a set of traits x which, on average, either overperform or underperform against randomly drawn opponents, and are frequently sampled, then a random sample of V competitors is expected to include some who perform well, and some poorly, against their neighbors. If, on the other hand, the expected performance conditioned on traits x is close to neutral, then ρ is small and competition is expected to be cyclic. In a rock-paper-scissors style game in which competitors are randomly and uniformly assigned rock, paper, or scissors, conditioning on receiving a particular trait does not change the probability that an individual with that trait will win most contests, hence the tournament is expected to be highly cyclic if L is large relative to V.

Another way to read (9.3) is as follows. Define the expected performance of traits x to be $\mathbb{E}_Y[f(x,Y)]$. Then, since $\mathbb{E}_X[\mathbb{E}_Y[f(X,Y)]] = \mathbb{E}_{X,Y}[f(X,Y)] = 0$, $\mathbb{E}_X[\mathbb{E}_Y[f(X,Y)]^2]$ is the variance in the expected performance given X. Therefore ρ is the ratio of the variance in the expected performance given X to the variance in performance. A large variance in the expected performance means we are likely to sample some competitors who perform well, or poorly, against most opponents. Consequently, the sampled edge flow is expected to be more transitive than cyclic.

Rereading Theorem 4.1 in this way leads to the following insight:

COROLLARY 4.2. *If the traits W, X, Y are sampled independently from π_x and $F = f(X,Y)$ then the correlation coefficient ρ is proportional to the variance in the expected performance:*

$$\rho = \frac{1}{\sigma^2}\mathrm{Cov}(f(X,Y), f(X,W)) = \frac{1}{\sigma^2}\mathrm{Var}\left(\mathbb{E}[F|X]\right). \tag{4.19}$$

Let ν be the expected variance in the performance:

$$\nu = \frac{1}{\sigma^2}\mathbb{E}\left[\mathrm{Var}(F|X)\right]. \tag{4.20}$$

Then $\nu = \frac{1}{\sigma^2}\mathrm{Var}[f(X,Y) - f(X,W)] = 1 - \rho$, so $\mathbb{E}[||F_c||^2]$ is monotonically increasing in ν and $\mathbb{E}[||F_t||^2]$ is monotonically decreasing in ν.

The proof is provided in the supplement and follows from the law of total variance,

$$\sigma^2 = \mathrm{Var}(F) = \mathbb{E}\left[\mathrm{Var}(F|X)\right] + \mathrm{Var}\left[\mathbb{E}(F|X)\right] = \sigma^2(\rho + \nu). \tag{4.21}$$

Theorem 4.1 identifies which statistical feature of the trait distribution and performance function promotes transitive and suppresses cyclic competition. Corollary 9.1 identifies which feature suppresses transitive and promotes cyclic competition. Transitive competition is promoted by uncertainty in expected performance, $\mathrm{Var}[\mathbb{E}(F|X)]$, and suppressed by expected uncertainty, $\mathbb{E}[\mathrm{Var}(F|X)]$. Conversely, cyclic competition is suppressed by uncertainty in expected performance, and promoted by expected uncertainty. If the expected uncertainty in performance is large, then performance is competitor dependent, hence competition is mostly cyclic.

Theorem 4.1 and Corollary 9.1 provide conceptual bridges between uncertainty in the edge flow, correlation structure on adjacent edges, and network structure (see Figure 9). They establish the intuitive statements that conclude the introduction (p. 4). For example, the expected uncertainty in the performance of A against a random competitor is $\sigma^2\nu = \frac{1}{2}\mathbb{E}_X[\mathrm{Var}_Y(f(X,Y)|X)]$. Thus, *"the less predictable the performance of A against a randomly drawn competitor, the more cyclic the tournament"*.

By the equivalence of $\mathbb{E}_X[\text{Var}_Y(f(X,Y)|X)]$ to $\text{Var}(f(X,Y) - f(X,W))$, *"the more the performance of A depends on their opponent, the more cyclic the tournament."*

4.3. Generalization. The trait-performance assumptions are not valid for all tournaments of interest.

Nevertheless, the conclusions of the trait-performance can be generalized to situations where the assumptions do not hold. We propose three generalizations. First we consider a situation where performance is only partially determined by traits. Second, if the network is complete, then the established relationship between expected structure and correlation holds when ρ is replaced with its empirical estimate. The empirical correlation depends only on the observed network, so the relation is an algebraic fact that is true for all complete networks, whatever the underlying distribution. Third, the trait-performance results hinged on a canonical form for the covariance in the edge flow (4.5). If an edge flow distribution has covariance in the canonical form, then the expected structure of the network satisfies (4.6). Thus, the conclusions relating structure to correlation hold for any edge flow distribution with covariance in the canonical form, whether or not that distribution came from a trait-performance model. If we assume a priori that our distribution has a covariance in this form, then ρ is a single parameter that tunes the sampled networks structure.

To start, what if performance is influenced by some random factors (such as unmeasured traits) in addition to a set of measured traits? Decompose $\text{Cov}(F)$ using the law of total variation. The first term in the decomposition would be the covariance in the the expected log-odds given the traits, which is a function of randomly drawn traits, so would take the canonical form (4.5) where the performance function $f(x,y)$ is replaced with $\mathbb{E}[F|x,y]$. Then, since $\mathbb{E}[||F_t||^2]$ and $\mathbb{E}[||F_c||^2]$ are linear in $\text{Cov}(F)$, the expected sizes of the transitive and cyclic components could each be expressed as a combination of a term contributed by the uncertainty in traits, and a term contributed by the uncertainty in performance given traits. The first term would simplify in the standard way, so the influence of the measured traits on expected network structure would follow as in the trait-performance theorem.

Second, we define the empirical correlation $\rho(\mathcal{G}_{\rightleftarrows})$ and variance $\sigma^2(\mathcal{G}_{\rightleftarrows})$ associated with a particular competitive network $\mathcal{G}_{\rightleftarrows}$. The empirical variance and correlation are estimators for the variance and correlation given the observed network. The empirical correlation is the covariance in the edge flow over all pairs of edges sharing an endpoint, divided by the empirical variance in the edge flow. Note that we only have one observation of f per edge, so we need to make some assumption about the expected value of the edge flow. We compute both the covariance and variance under the assumption that the expected edge flow is zero on each edge k. The assumption is valid provided that we would have no way to predict the sign of f_k (whether $i(k)$ or $j(k)$ usually wins) from the network topology alone. Then, $\rho(\mathcal{G}_{\rightleftarrows})$ is the average value of $a_{Ekl} f_{i(k)j(k)} f_{i(l)j(l)}$ over all pairs of edges k, l that share an endpoint, where $a_{Ekl} = 1$ if the edges are consistently oriented, and $a_{Ekl} = -1$ if the edges are inconsistently oriented. The empirical variance $\sigma^2(\mathcal{G}_{\rightleftarrows})$ is simply $\frac{1}{E}||f||^2$.

LEMMA 4.3. *If the competitive network is complete, has V vertices, E edges, L loops, empirical variance $\sigma^2(\mathcal{G}_{\rightleftarrows})$, and correlation $\rho(\mathcal{G}_{\rightleftarrows})$ then:*

$$\frac{1}{E}||f||^2 = \sigma^2(\mathcal{G}_{\rightleftarrows}) \xrightarrow{decompose} \begin{cases} \frac{1}{E}||f_t||^2 = \sigma^2(\mathcal{G}_{\rightleftarrows}) \left[\frac{(V-1)}{E} + 2\rho(\mathcal{G}_{\rightleftarrows})\frac{L}{E}\right] \\ \frac{1}{E}||f_c||^2 = \sigma^2(\mathcal{G}_{\rightleftarrows})\left(1 - 2\rho(\mathcal{G}_{\rightleftarrows})\right)\frac{L}{E} \end{cases} \tag{4.22}$$

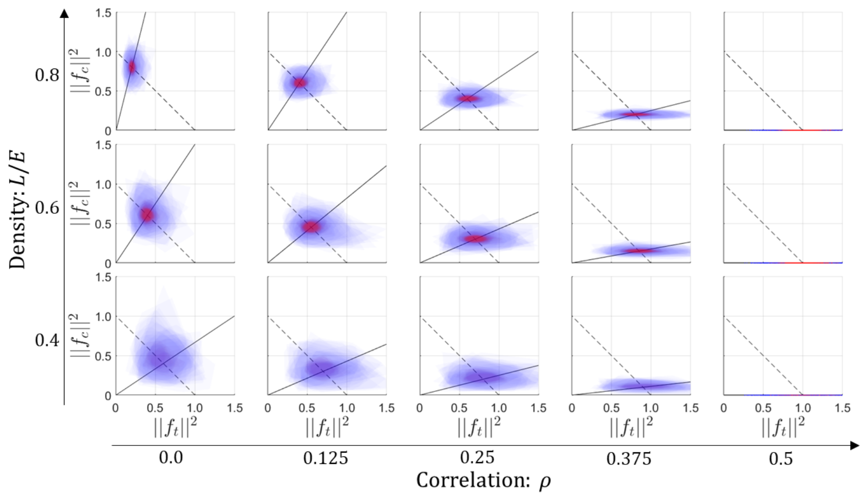

FIG. 10. *Transitivity and intransitivity of sampled networks with varying edge density, number of competitors, and correlation ρ. Each row represents networks with a fixed ratio L/E where L is the number of loops, $E - (V - 1)$, and E is the number of edges. Each column represents a fixed correlation ρ. When $\rho = 0$ the edge flows on all edges are independent. When $\rho = 0.5$ the randomly sampled networks are all perfectly transitive. The blue shaded region is a heat map representing 10^4 sampled networks with 20 competitors. The red shaded region is a heat map representing 10^4 sampled networks with 300 competitors. The topology of each network is sampled randomly from the family of connected Erdos-Renyi networks. The edge flows are sampled from the multivariate Gaussian distribution with mean zero and covariance of form (4.5). The solid black line represents the expected relative sizes of the transitive and intransitive component predicted by equation (4.6). The dashed black line represents the expected total intensity of competition, σ^2. The intersection of these two lines gives the expected absolute sizes of the transitive and intransitive components. Notice that the trait-performance theorem correctly predicts the relative and absolute sizes of the transitive and intransitive components as a function of L/E, σ, and ρ. Moreover, the more competitors in the network, the tighter the agreement to the expected sizes.*

The proof is provided in the supplement.

Third, the conclusions of the trait-performance theorem relating correlation and topology to structure hold as long as the edge flow F has covariance in the canonical form (4.5). The trait-performance assumptions guarantee that the covariance takes this form, but an edge flow F may have a covariance in this form whether or not it is related to an underlying trait-performance model. Thus the conclusions of the theorem generalize to all edge flow distributions with covariance of the form (4.5).

It follows that we can use the trait-performance results to design families of null models with tunable structure. For example, suppose that we are given a specific network topology. Then we could sample F from the multivariate Gaussian distribution with mean zero and covariance chosen to match (4.5). By choosing σ^2 and ρ we uniquely specify the edge flow distribution. Then the expected absolute and relative sizes of the transitive and cyclic components would be directly controlled by the choice of σ^2 and ρ. We could tune the overall intensity of competition by varying σ^2, and the relative degree of intransitivity by varying ρ. Results from null models of this kind are demonstrated in Figure 10. The figure demonstrate that it is possible to define null models with a chosen degree of transitivity by tuning the correlation ρ.

Null models of this kind could be useful since many empirical studies involve complex competition events where reasonable statistical modelling of sampling error is difficult [16, 79]. Absent a statistical error model, the observed edge flow must be treated as truth, so significance must be computed with respect to a null distribution. The standard test approximates significance relative to a uniform distribution of dominance relationships (sign of the edge flow) on a complete network [2, 15, 35]. This significance is only useful so far as the uniform null model is a plausible model for competition, or as it restricts the space of possible competition structures. The fact that most studies identify significant transitivity suggests that the uniform distribution is rarely plausible. Failure to match a uniform distribution also does not limit the competitive structure significantly, since, as demonstrated above, it is easy to construct null models that produce intermediate levels of transitivity.

In fact, complete networks with edge flow drawn uniformly are the *most* cyclic edge flow distribution with covariance of the form (4.5) since they are simultaneously as dense and uncorrelated as possible. Complete networks with uniform i.i.d. edge flow live in the upper left-hand corner of Figure 10. It is not surprising that most empirical networks are more transitive than the most cyclic ensemble. For this reason, significance computed against the uniform complete null model may depend primarily on the number of imputed edges, as observed in [63, 37, 22], rather than true structure.

The family of null models proposed here could generalize the standard randomization test in two useful ways. First, it allows for arbitrary network topology, so does not require imputing missing edges which reduces the strength of the test [63]. Second, the expected degree of transitivity in the null model can be tuned using one parameter, ρ. Once ρ is chosen, we could compute the probability of observing a network that is more or less transitive or intransitive relative to random networks with correlation ρ. Thus significance could be measured against a flexible range of networks with varying degrees of transitivity. Then it would be possible to search over $\rho \in [0, 0.5]$ to find the largest and smallest ρ which produce random networks with significantly different structure than the observed network. The interval between these upper and lower bounds on ρ would define an interval in each transitivity measure that could plausibly correspond to the observed network. Thus, expanding the family of null models would allow more flexible, informative, significance testing, as well as interval estimation of the measures of competitive structure.

5. Discussion.

The discrete HHD provides a natural, unified method for ranking and measuring intransitivity via a decomposition into perfectly transitive and cyclic components. The expected size of these components can be computed from the correlation structure of the edge flow. Using a trait-performance model simplifies this structure. We provide an illustrative example in the supplement. Note that the trait-performance conclusions are valid whenever the assumptions hold, whether or not the relevant traits or performance function are known. Thus the assumptions can be tested by checking whether the observed correlation structure matches (4.5).

Further theoretical work could address random network topologies. If the network is sampled independently of the edge flow then the results of Theorem 4.1 are largely unchanged, so one might consider random networks whose topology is coupled to the competitor traits. For example, neighbors in the network might have correlated traits. Future work could also investigate null models with different covariance structures.

We emphasize that the HHD can be applied to analyze a tournament independent of a null model. Code for implementing our methods are available on github. In particular, our methods can be extended to analyze data from real tournaments. By studying win-loss records it is possible to infer the log odds, and thus estimate the components of the HHD. The estimation problem is saved for future work.

Acknowledgments. We would like to thank Lek-Heng Lim, Robin Snyder, and Michael Hinczewski for helpful conversations. We also thank Gilbert Strang for his help revising. This work was funded by NSF grants DEB-1654989 and DMS-2052109.

REFERENCES

[1] D. Aldous et al., *Elo ratings and the sports model: A neglected topic in applied probability?*, Statistical Science, 32 (2017), pp. 616–629.

[2] M. C. Appleby, *The probability of linearity in hierarchies*, Animal Behaviour, 31 (1983), pp. 600–608.

[3] J. Bartholdi, C. A. Tovey, and M. A. Trick, *Voting schemes for which it can be difficult to tell who won the election*, Social Choice and Welfare, 6 (1989), pp. 157–165.

[4] J. L. Beacham, *The relative importance of body size and aggressive experience as determinants of dominance in pumpkinseed sunfish, Lepomis gibbosus*, Animal Behaviour, 36 (1988), pp. 621–623.

[5] B. Bollobás, *Modern graph theory*, vol. 184, Springer Science & Business Media, 2013.

[6] S. Bozóki, L. Csató, and J. Temesi, *An application of incomplete pairwise comparison matrices for ranking top tennis players*, European Journal of Operational Research, 248 (2016), pp. 211–218.

[7] R. A. Bradley, *Incomplete block rank analysis: On the appropriateness of the model for a method of paired comparisons*, Biometrics, 10 (1954), pp. 375–390.

[8] R. A. Bradley and M. E. Terry, *Rank analysis of incomplete block designs: I. The method of paired comparisons*, Biometrika, 39 (1952), pp. 324–345.

[9] S. Brin and L. Page, *Reprint of: The anatomy of a large-scale hypertextual web search engine*, Computer networks, 56 (2012), pp. 3825–3833.

[10] O. Candogan, I. Menache, A. Ozdaglar, and P. A. Parrilo, *Flows and decompositions of games: Harmonic and potential games*, Mathematics of Operations Research, 36 (2011), pp. 474–503.

[11] P. Charbit, S. Thomassé, and A. Yeo, *The minimum feedback arc set problem is NP-hard for tournaments*, Combinatorics, Probability and Computing, 16 (2007), pp. 1–4.

[12] I. D. Chase, C. Bartolomeo, and L. A. Dugatkin, *Aggressive interactions and inter-contest interval: how long do winners keep winning?*, Animal Behaviour, 48 (1994), pp. 393–400.

[13] I. D. Chase, C. Tovey, D. Spangler-Martin, and M. Manfredonia, *Individual differences versus social dynamics in the formation of animal dominance hierarchies*, Proceedings of the National Academy of Sciences, 99 (2002), pp. 5744–5749.

[14] W. Colley, *Colley's bias free college football ranking method*, PhD thesis, Princeton University, 2002.

[15] H. De Vries, *An improved test of linearity in dominance hierarchies containing unknown or tied relationships*, Animal Behaviour, 50 (1995), pp. 1375–1389.

[16] H. De Vries, *Finding a dominance order most consistent with a linear hierarchy: a new procedure and review*, Animal Behaviour, 55 (1998), pp. 827–843.

[17] C. Drews, *The concept and definition of dominance in animal behaviour*, Behaviour, 125 (1993), pp. 283–313.

[18] P. Eades, X. Lin, and W. F. Smyth, *A fast and effective heuristic for the feedback arc set problem*, Information Processing Letters, 47 (1993), pp. 319–323.

[19] U. Endriss and R. de Haan, *Complexity of the winner determination problem in judgment aggregation: Kemeny, Slater, Tideman, Young.* (downloaded from https://eprints.illc.uva.nl/534/), 2016.

[20] M. Frean and E. R. Abraham, *Rock–scissors–paper and the survival of the weakest*, Proceedings of the Royal Society of London. Series B: Biological Sciences, 268 (2001), pp. 1323–1327.

[21] G. Galbiati, *On optimum cycle bases*, Electronic Notes in Discrete Mathematics, 10 (2001), pp. 113–116.

[22] F. Galimberti, A. Fabiani, and L. Boitani, *Socio-spatial levels in linearity analysis of dominance hierarchies: a case study on elephant seals*, Journal of ethology, 21 (2003), pp. 131–136.

[23] W. V. Gehrlein, *Condorcet's paradox and the Condorcet efficienty of voting rules*, Mathematica Japonica, 45 (1997), pp. 173–199.

[24] M. E. Glickman, *Parameter estimation in large dynamic paired comparison experiments*, Journal of the Royal Statistical Society: Series C (Applied Statistics), 48 (1999), pp. 377–394.

[25] O. Godoy, D. B. Stouffer, N. J. Kraft, and J. M. Levine, *Intransitivity is infrequent and fails to promote annual plant coexistence without pairwise niche differences*, Ecology, 98 (2017), pp. 1193–1200.

[26] M. P. Haley, C. J. Deutsch, and B. J. Le Boeuf, *Size, dominance and copulatory success in male northern elephant seals, Mirounga angustirostris*, Animal Behaviour, 48 (1994), pp. 1249–1260.

[27] E. Hemaspaandra, H. Spakowski, and J. Vogel, *The complexity of Kemeny elections*, Theoretical Computer Science, 349 (2005), pp. 382–391.

[28] Y. Hsu and L. L. Wolf, *The winner and loser effect: integrating multiple experiences*, Animal Behaviour, 57 (1999), pp. 903–910.

[29] L. M. Hvattum and H. Arntzen, *Using Elo ratings for match result prediction in association football*, International Journal of forecasting, 26 (2010), pp. 460–470.

[30] J. Jackson and L. Buss, *Alleopathy and spatial competition among coral reef invertebrates*, Proceedings of the National Academy of Sciences, 72 (1975), pp. 5160–5163.

[31] W. M. Jackson and R. L. Winnegrad, *Linearity in dominance hierarchies: a second look at the individual attributes model.*, Animal Behaviour, (1988).

[32] X. Jiang, L.-H. Lim, Y. Yao, and Y. Ye, *Statistical ranking and combinatorial Hodge theory*, Mathematical Programming, 127 (2011), pp. 203–244.

[33] C. R. Johnson and I. Seinen, *Selection for restraint in competitive ability in spatial competition systems*, Proceedings of the Royal Society of London. Series B: Biological Sciences, 269 (2002), pp. 655–663.

[34] J. P. Keener, *The Perron–Frobenius theorem and the ranking of football teams*, SIAM review, 35 (1993), pp. 80–93.

[35] M. G. Kendall and B. B. Smith, *On the method of paired comparisons*, Biometrika, 31 (1940), pp. 324–345.

[36] B. Kerr, M. A. Riley, M. W. Feldman, and B. J. Bohannan, *Local dispersal promotes biodiversity in a real-life game of rock–paper–scissors*, Nature, 418 (2002), pp. 171–174.

[37] K. Klass and M. Cords, *Effect of unknown relationships on linearity, steepness and rank ordering of dominance hierarchies: simulation studies based on data from wild monkeys*, Behavioural processes, 88 (2011), pp. 168–176.

[38] ———, *Agonism and dominance in female blue monkeys*, American journal of primatology, 77 (2015), pp. 1299–1315.

[39] A. Koenig, *Competition for resources and its behavioral consequences among female primates*, International Journal of Primatology, 23 (2002), pp. 759–783.

[40] S. A. Kovalchik, *Searching for the goat of tennis win prediction*, Journal of Quantitative Analysis in Sports, 12 (2016), pp. 127–138.

[41] M. Kwiesielewicz, *The logarithmic least squares and the generalized pseudoinverse in estimating ratios*, European Journal of Operational Research, 93 (1996), pp. 611–619.

[42] M. Kwiesielewicz and E. Van Uden, *Ranking decision variants by subjective paired comparisons in cases with incomplete data*, in International Conference on Computational Science and Its Applications, Springer, 2003, pp. 208–215.

[43] R. A. Laird and B. S. Schamp, *Competitive intransitivity promotes species coexistence*, The American Naturalist, 168 (2006), pp. 182–193.

[44] H. G. Landau, *On dominance relations and the structure of animal societies: I. Effect of inherent characteristics*, The bulletin of mathematical biophysics, 13 (1951), pp. 1–19.

[45] A. N. Langville and C. D. Meyer, *Who's# 1?: the science of rating and ranking*, Princeton University Press, 2012.

[46] R. A. Lankau and S. Y. Strauss, *Mutual feedbacks maintain both genetic and species diversity in a plant community*, science, 317 (2007), pp. 1561–1563.

[47] R. A. Lankau, E. Wheeler, A. E. Bennett, and S. Y. Strauss, *Plant–soil feedbacks contribute to an intransitive competitive network that promotes both genetic and species diversity*, Journal of Ecology, 99 (2011), pp. 176–185.

[48] J. Lasek, Z. Szlávik, and S. Bhulai, *The predictive power of ranking systems in association football*, International Journal of Applied Pattern Recognition, 1 (2013), pp. 27–46.

[49] M. Lewis, *Moneyball: The art of winning an unfair game*, WW Norton & Company, 2004.

[50] L.-H. Lim, *Hodge laplacians on graphs*, SIAM Review, 62 (2020), pp. 685–715.

[51] K. Massey, *Statistical models applied to the rating of sports teams*, Bluefield College, (1997).

[52] R. M. May and W. J. Leonard, *Nonlinear aspects of competition between three species*, SIAM Journal on Applied Mathematics, 29 (1975), pp. 243–253.

[53] I. McHale and A. Morton, *A Bradley-Terry type model for forecasting tennis match results*, International Journal of Forecasting, 27 (2011), pp. 619–630.

[54] L. Muniz, S. Perry, J. H. Manson, H. Gilkenson, J. Gros-Louis, and L. Vigilant, *Male*

dominance and reproductive success in wild white-faced capuchins (Cebus capucinus) at Lomas Barbudal, Costa Rica, American Journal of Primatology, 72 (2010), pp. 1118–1130.
[55] R. F. Oliveira, P. K. McGregor, and C. Latruffe, *Know thine enemy: fighting fish gather information from observing conspecific interactions*, Proceedings of the Royal Society of London. Series B: Biological Sciences, 265 (1998), pp. 1045–1049.
[56] M. Ozanian and K. Badenhausen, *Baseball team values 2019*, Forbes, (2019).
[57] P. S. Petraitis, *Competitive networks and measures of intransitivity*, The American Naturalist, 114 (1979), pp. 921–925.
[58] T. Reichenbach and E. Frey, *Instability of spatial patterns and its ambiguous impact on species diversity*, Physical Review Letters, 101 (2008), p. 058102.
[59] T. Reichenbach, M. Mobilia, and E. Frey, *Coexistence versus extinction in the stochastic cyclic Lotka-Volterra model*, Physical Review E, 74 (2006), p. 051907.
[60] ———, *Mobility promotes and jeopardizes biodiversity in rock–paper–scissors games*, Nature, 448 (2007), pp. 1046–1049.
[61] ———, *Noise and correlations in a spatial population model with cyclic competition*, Physical Review Letters, 99 (2007), p. 238105.
[62] R. M. Sapolsky and L. J. Share, *Rank-related differences in cardiovascular function among wild baboons: Role of sensitivity to glucocorticoids*, American Journal of Primatology, 32 (1994), pp. 261–275.
[63] D. Shizuka and D. B. McDonald, *A social network perspective on measurements of dominance hierarchies*, Animal Behaviour, 83 (2012), pp. 925–934.
[64] H. H. Shoemaker, *Social hierarchy in flocks of the canary*, The Auk, (1939), pp. 381–406.
[65] J. B. Silk, *Male bonnet macaques use information about third-party rank relationships to recruit allies*, Animal Behaviour, 58 (1999), pp. 45–51.
[66] B. Sinervo and C. M. Lively, *The rock–paper–scissors game and the evolution of alternative male strategies*, Nature, 380 (1996), pp. 240–243.
[67] Y. Sismanis, *How I won the "Chess Ratings – Elo vs the Rest of the World" competition*, arXiv preprint arXiv:1012.4571, (2010).
[68] P. Slater, *Inconsistencies in a schedule of paired comparisons*, Biometrika, 48 (1961), pp. 303–312.
[69] L. Smale, L. G. Frank, and K. E. Holekamp, *Ontogeny of dominance in free-living spotted hyaenas: juvenile rank relations with adult females and immigrant males*, Animal Behaviour, 46 (1993), pp. 467–477.
[70] S. Soliveres, F. T. Maestre, W. Ulrich, P. Manning, S. Boch, M. A. Bowker, D. Prati, M. Delgado-Baquerizo, J. L. Quero, I. Schöning, et al., *Intransitive competition is widespread in plant communities and maintains their species richness*, Ecology letters, 18 (2015), pp. 790–798.
[71] R. Stefani, *The methodology of officially recognized international sports rating systems*, Journal of Quantitative Analysis in Sports, 7 (2011).
[72] R. T. Stefani, *Football and basketball predictions using least squares*, IEEE Transactions on systems, man, and cybernetics, 7 (1977), pp. 117–21.
[73] ———, *Improved least squares football, basketball, and soccer predictions*, IEEE transactions on systems, man, and cybernetics, 10 (1980), pp. 116–123.
[74] G. Strang, *The fundamental theorem of linear algebra*, The American Mathematical Monthly, 100 (1993), pp. 848–855.
[75] E. D. Strauss and K. E. Holekamp, *Social alliances improve rank and fitness in convention-based societies*, Proceedings of the National Academy of Sciences, 116 (2019), pp. 8919–8924.
[76] D. M. Stuart-Fox, D. Firth, A. Moussalli, and M. J. Whiting, *Multiple signals in chameleon contests: designing and analysing animal contests as a tournament*, Animal Behaviour, 71 (2006), pp. 1263–1271.
[77] W. Ulrich, S. Soliveres, W. Kryszewski, F. T. Maestre, and N. J. Gotelli, *Matrix models for quantifying competitive intransitivity from species abundance data*, Oikos, 123 (2014), pp. 1057–1070.
[78] C. S. Valero, *Predicting Win-Loss outcomes in MLB regular season games – A comparative study using data mining methods*, International Journal of Computer Science in Sport, 15 (2016), pp. 91–112.
[79] H. d. Vries and M. C. Appleby, *Finding an appropriate order for a hierarchy: a comparison of the I&SI and the BBS methods*, Animal Behaviour, 59 (2000), pp. 239–245.
[80] C. Xue and N. Goldenfeld, *Coevolution maintains diversity in the stochastic "Kill the Winner" model*, Physical Review Letters, 119 (2017), p. 268101.

Supplementary Materials

6. Proof of Lemma 3.1.

LEMMA 6.1 (Arbitrage Free). *A tournament is arbitrage free if and only if there exists a unique set of ratings r with average rating equal to zero such that the win probabilities satisfy $p_{ij} = \text{logistic}(r_i - r_j)$* [15]. *Moreover if a tournament is arbitrage free then it is transitive.*

Proof. Suppose that a tournament is arbitrage free. Then it must satisfy the cycle condition. The cycle condition requires that the path sum of f around any cycle is zero. Consider two paths $\mathcal{P}_1$ and $\mathcal{P}_2$ both starting at A and ending at B. The value of the path sum over $\mathcal{P}_1$ minus the path sum over $\mathcal{P}_2$ equals the path sum around a cycle following $\mathcal{P}_1$ from A to B, then following the path $\mathcal{P}_2$ backwards from B to A. The path sum around any cycle is zero, thus the path sum over $\mathcal{P}_1$ and $\mathcal{P}_2$ must be equal. It follows that, for any pair of endpoints A, B, the value of the path sum of f over a path connecting A to B only depends on A and B and is otherwise path independent.

To recover the associated ratings, pick an arbitrary spanning tree of the network and an arbitrary starting competitor A.[16] Then let u_B equal the path sum of f over the path connecting A to B in the tree. Then u are ratings relative to competitor A. Path independence guarantees that the values u depend only on the choice of A, not the choice of spanning tree. To eliminate the dependence on A, center the ratings by subtracting off their average. Let $r_B = u_B - \frac{1}{V}\sum_{i=1}^{V} u_i$. Then r are unique and independent of the choice of tree and A, and, by construction, $\sum_i r_i = 0$. It remains to show that $r_i - r_j = f_{ij}$ for all connected pairs i, j. This equality holds by construction for all i, j that are connected through an edge in the spanning tree. Consider an edge not in the spanning tree (a chord) connecting i and j. Let $i_1 = A, i_2, \ldots, i_l = i$ and $j_1 = A, j_2, \ldots, j_k = j$ be the paths from A to i and j through the spanning tree (see Figure 11). Then, the path sum from j to i in the tree equals $r_i - r_j$:

$$\underbrace{r_i - r_j = u_i - u_j}_{\text{Rating difference}} = \underbrace{\sum_{n=1}^{l-1} f_{i_{n+1} i_n}}_{\text{sum } A \text{ to } i} - \underbrace{\sum_{n=1}^{k-1} f_{j_{n+1} j_n}}_{\text{sum } A \text{ to } j} = \underbrace{\sum_{n=k}^{2} f_{j_{n-1} j_n} + \sum_{n=1}^{l-1} f_{i_{n+1} i_n}}_{\text{sum } j \text{ to } i}$$

The chord connecting j and i also defines a path from j to i. Since path sums are path independent when the network is arbitrage free, the path sum over the chord ij equals the path sum through the tree. The path sum over the chord is f_{ij} so $f_{ij} = r_i - r_j$. Therefore, if a tournament is arbitrage free then there exist a set of ratings r such that $r_i - r_j = f_{ij}$. Then, since $f_{ij} = \text{logit}(p_{ij})$, $p_{ij} = \text{logistic}(r_i - r_j)$.

Suppose that $p_{ij} = \text{logistic}(r_i - r_j)$. Then $f_{ij} = r_i - r_j$ for all connected i, j, and, given a path $i_1, i_2, \ldots, i_n$ the sum $f_{i_2 i_1} + f_{i_3 i_2} + \ldots f_{i_n i_{n-1}} = r_{i_n} - r_{i_1}$ as it telescopes. If the path is a loop then $i_n = i_1$ so the sum equals zero. But then f satisfies the cycle condition, so the tournament is arbitrage free.

Suppose the tournament is arbitrage free. Then $p_{ij} = \text{logistic}(r_i - r_j)$ for a unique set of ratings r. This means that $p_{ij} > 1/2$ if and only if $r_i > r_j$. It follows that $A \succ B$ if and only if $r_A > r_B$, so the win probabilities are consistent with the ranking induced by the ratings r, thus the tournament is transitive. ∎

[15] $\text{logistic}(x) = \text{logit}^{-1}(x) = 1/(1 + \exp(-x))$.

[16] A spanning tree is a subgraph of the network that contains no loops, includes all competitors, and is connected.

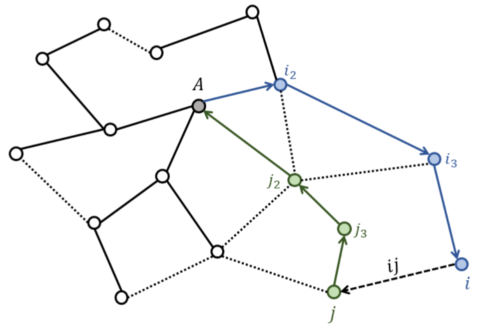

Fig. 11. *The spanning tree construction for recovering the ratings for an arbitrage-free tournament. The tree is shown with solid lines, and the chords with dotted lines. The root of the tree, A is marked in grey. Two vertices, i and j connected by a chord ij, are shown in blue and green respectively. The sequence of nodes leading from A to i and j are labelled. Then, by the cycle condition, the sum around the loop marked with arrows is zero, hence $f_{ij} = r_i - r_j$.*

7. Proof of Lemma 3.2.

Lemma 7.1 (favorite free). *A favorite free tournament is cyclic, and is never transitive unless $p_{ij} = 1/2$ for all connected i, j.*

Proof. Suppose that a tournament is favorite free. Then $\sum_{j \in \mathcal{N}_i} f_{ij} = 0$ for all i. This leaves two distinct possibilities, either $f_{ij} = 0$ for all $j \in \mathcal{N}(i)$, or there is some j such that $f_{ij} \neq 0$. The former case requires $p_{ij} = 1/2$ for all $j \in \mathcal{N}(i)$. We will refer to this case as the *neutral* case. If the neighborhood of i is not neutral then $f_{ij} \neq 0$ for some $j \in \mathcal{N}(i)$. Since the sum over all j is zero this means that there must be at least one other edge ik such that $\text{sign}(f_{ij}) = -\text{sign}(f_{ik})$. Thus, if there is an edge into competitor i in $\mathcal{G}_{\rightarrow}$ there must also be an edge out of i in $\mathcal{G}_{\rightarrow}$.

Since the neighborhood condition can be extended from the neighborhood of competitors to the neighborhood of sets this property also extends to sets. If there is an edge into the set S in $\mathcal{G}_{\rightarrow}$ then there must also be an edge out of the set.

Now suppose that there is a path from A to B in $\mathcal{G}_{\rightarrow}$. It remains to construct a path back to A.

Define the nested sets $S_0(B), S_1(B), \ldots,$, where $S_d(B)$ is the set of all nodes that can be reached from B with a path in $\mathcal{G}_{\rightarrow}$ of length less than or equal to d. Since there is a path from A to B in $\mathcal{G}_{\rightarrow}$ there is an edge in $\mathcal{G}_{\rightarrow}$ arriving at $\{B\} = S_0(B)$. Thus there is a path from A to all competitors in $S_1(B)$. Now there are two possibilities, either A is in $S_1(B)$, or A is not in $S_1(B)$. If A is in $S_1(B)$ then we are done. If not, then there is an edge entering $S_1(B)$ in $\mathcal{G}_{\rightarrow}$ since there is a path from $A \notin S_1(B)$ to $B \in S_1(B)$. Then the neighborhood condition implies that there is an edge out of $S_1(B)$, which means that $S_2(B) \neq S_1(B)$. Now the logic repeats. Either A is in $S_2(B)$, in which case we are done, or it is not. If it is not then there must be an edge entering $S_2(B)$ so there must be an edge leaving $S_2(B)$ so $S_3(B) \neq S_2(B)$. As long as $A \notin S_d(B)$ there is a larger set $S_{d+1}(B) \neq S_d(B)$ which can be reached from B. Since we assumed that there are finitely many competitors this can only continue until A is contained in $S_d(B)$ for some B. See Figure 12 for illustration.

Suppose that the tournament is transitive, favorite free, and not neutral. Since it isn't neutral there must be at least one pair ij such that $p_{ij} > 1/2$. This means that $R_i < R_j$ and there is an edge from j to i in $\mathcal{G}_{\rightarrow}$. But, if the tournament is favorite

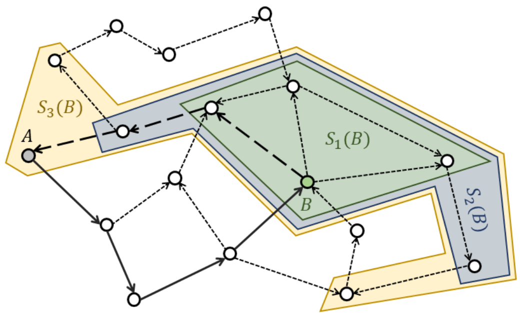

FIG. 12. *A favorite free tournament must be a cyclic tournament. The arrows represent the direction of competition. If the network is favorite free then if there is an edge pointing into a set there must be an edge pointing out of it. A path from A to B is shown in black. Then the sets $S_1(B), S_2(B), S_3(B)$ are shown as shaded polygons. These contain all competitors distance 1, 2, and 3 (respectively) from B. These sets continue to expand until they include A, hence there is a path from B to A.*

free then there must be some other path from i back to j in $\mathcal{G}_{\rightarrow}$. Then $R_j < R_i$ since there is a path in $\mathcal{G}_{\rightarrow}$ from j to i. This is clearly a contradiction. Therefore, a cyclic tournament is not transitive unless it is neutral: $p_{ij} = 1/2$ for all ij.[17] □

8. Interpretation of Corollary 3.6.

COROLLARY 8.1 (Equivalent Formulations). *The following six decompositions are equivalent:*

1. $f = f_t + f_c$ *where* f_t *is arbitrage free and* f_c *is favorite free;*
2. $f = f_t + f_c$ *where* $f_t = Gr$ *for some rating* r *and* $f_c = C^T v$ *for some vorticity* v*;*
3. *the ratings* r *solve the constrained least squares problem:*

$$\text{Minimize } ||Gu - f||_2{}^2 \quad \text{given } u \in \mathbb{R}^V \text{ and } \sum_{i=1}^{V} u_i = 0 \tag{8.1}$$

   and $f_t = Gr, f_c = f - f_t$*;*
4. *the vorticities* v *solve the least squares problem:*

$$\text{Minimize: } ||C^T w - f||_2{}^2 \quad \text{given } w \in \mathbb{R}^L \tag{8.2}$$

   and $f_c = C^T v, f_t = f - f_c$*;*
5. $f = f_t + f_c$ *where* $f_t = Gr$ *for the unique ratings* r *such that the circulant* $f - f_t$ *is favorite free;*
6. $f = f_t + f_c$ *where* $f_c = C^T v$ *for the unique vorticities* v *such that* $f - f_c$ *is arbitrage free.*

The first decomposition separates f into a pair of flows each defined by what it is not: namely, one is not circulatory, and the other has no tendency to diverge or converge. The second decomposition separates f into a pair of flows each defined by what they are: namely, one is perfectly transitive, and the other is perfectly cyclic. The equivalence of these two decompositions was established by Theorem 3.5.

[17] This shows that the two classes of tournaments are distinct, as their only overlap is the neutral case. Note that a neutral tournament is considered transitive since it can be consistently ranked - all competitors should be ranked the same.

The next two decompositions are based on fitting problems. In each case the goal is to represent f as nearly as possible when restricted to the range of an operator. Decomposition 3 searches for a set of ratings r such the error, $Gr - f$, is minimized in the least squares sense. This means that the ratings produced by the HHD are a type of least squares rating, in particular, log least squares rating [6, 41, 42]. Least squares ratings methods are widely used [14, 34, 45, 51, 72, 73]. Decomposition 3 also shows that the HHD is equivalent to finding the nearest perfectly transitive edge flow.

Similarly, Decomposition 4 searches for a set of vorticities v such that the error $C^T v - f$ in approximating f with $C^T v$ is minimized in the least squares sense. This is equivalent to finding the nearest perfectly cyclic edge flow. Although the literature has focused almost exclusively on Decomposition 3, Decompositions 3 and 4 are dual to one another. This parity in approach sets the HHD apart from existing methods.

The final two decompositions are defined by enforcing a constraint on the residue when approximating f with either the gradient of a set of ratings or the curl transpose of a set of vorticities. These approaches can be motivated as follows. Suppose one sought a rating r such that Gr approximated f. The error in this approximation (the circulant) is $Gr - f$. As long as the divergence of the circulant is nonzero the approximation has not captured a tendency of the edge flow to either point inwards towards, or outwards from, a competitor. If the net flow into a competitor is positive, then that competitor tends to outperform their neighbors in a way that the ratings fail to capture. Therefore it would be natural to adjust the ratings until the net flow into or out of any set of competitors is zero. That is, until the divergence of the circulant is zero, or equivalently, the circulant is favorite free.

The final decomposition can be motivated similarly. Define the *divergent*, $C^T v - f$ to be the error upon approximating f with vorticity v. As long as the curl of the divergent is nonzero, the approximation has failed to capture some tendency of f to circulate. This tendency to circulate is exactly what the vorticities are meant to capture, so it is natural to look for a v such that the curl of the divergent is zero on every loop. That is, until the divergent is arbitrage free.

9. Proof of Corollary 4.2.

COROLLARY 9.1. *If the traits W, X, Y are sampled independently from π_x and $F = f(X, Y)$ then the correlation coefficient ρ is proportional to the variance in the expected performance:*

$$\rho = \frac{1}{\sigma^2}\mathrm{Cov}(f(X,Y), f(X,W)) = \frac{1}{\sigma^2}\mathrm{Var}\left(\mathbb{E}[F|X]\right). \tag{9.1}$$

Let ν be the expected variance in the performance:

$$\nu = \frac{1}{\sigma^2}\mathbb{E}\left[\mathrm{Var}(F|X)\right]. \tag{9.2}$$

Then $\nu = \mathrm{Var}[f(X,Y) - f(X,W)] = 1 - \rho$, so $\mathbb{E}[||F_c||^2]$ is monotonically increasing in ν, $\mathbb{E}[||F_t||^2]$ is monotonically decreasing in ν.

Proof. The proof of equation (9.1) follows from the explicit expression for ρ:

$$\rho = \frac{\int_\Omega \left(\int_\Omega f(x,y)\pi_x(y)dy\right)^2 \pi_x(x)dx}{\int_\Omega\int_\Omega f(x,y)^2\pi_x(y)\pi_x(x)dydx} = \frac{\mathbb{E}_X[\mathbb{E}_Y[f(X,Y)]^2]}{\mathbb{E}_{X,Y}[f(X,Y)^2]}. \tag{9.3}$$

Then, since $\mathbb{E}[F] = 0$, $\mathbb{E}_X[\mathbb{E}_Y[f(X,Y)]^2] = \mathrm{Var}(\mathbb{E}_Y[f(X,Y)]) = \mathrm{Var}(\mathbb{E}[F|X])$.

Next, $\nu = 1 - \rho$ follows from the law of total variance:

$$\sigma^2 = \text{Var}(F) = \mathbb{E}\left[\text{Var}(F|X)\right] + \text{Var}\left[\mathbb{E}(F|X)\right] = \sigma^2(\rho + \nu). \tag{9.4}$$

Since $\mathbb{E}[||F_c||^2]$ is decreasing in ρ, it is increasing in ν. Similarly, since $\mathbb{E}[||F_t||^2]$ is increasing in ρ, it is decreasing in ν.

The intermediate expression for ν follows from $\sigma^2\nu = \sigma^2(1-\rho) = \text{Var}[f(X,Y)] - \text{cov}[f(X,Y), f(X,W)]$. Since Y and W are i.i.d., $\text{Var}[f(X,Y)] = \frac{1}{2}(\text{Var}[f(X,Y)] + \text{Var}[f(X,W)])$. Substituting in gives $\sigma^2\nu = \frac{1}{2}\mathbb{E}[(f(X,Y) - f(X,W))^2]$. Since $\mathbb{E}[f(X,Y)]$ equals $\mathbb{E}[f(X,W)]$ this raw second moment is the variance in $f(X,Y) - f(X,W)$. ☐

10. Proof of Lemma 4.3.

LEMMA 10.1. *If the competitive network is complete, has m vertices, E edges, $L = E - (m-1)$ loops, empirical variance $\sigma^2(\mathcal{G}_{\leftrightarrows})$, and correlation $\rho(\mathcal{G}_{\leftrightarrows})$ then:*

$$\frac{1}{E}||f||^2 = \sigma^2(\mathcal{G}_{\leftrightarrows}) \xrightarrow{decompose} \begin{cases} \frac{1}{E}||f_t||^2 = \sigma^2(\mathcal{G}_{\leftrightarrows})\left[\frac{(V-1)}{E} + 2\rho(\mathcal{G}_{\leftrightarrows})\frac{L}{E}\right] \\ \frac{1}{E}||f_c||^2 = \sigma^2(\mathcal{G}_{\leftrightarrows})\left(1 - 2\rho(\mathcal{G}_{\leftrightarrows})\right)\frac{L}{E} \end{cases} \tag{10.1}$$

Proof. The empirical correlation $\rho(\mathcal{G}_{\leftrightarrows})$ is given by averaging $s_{k,l} f_k f_l$ over all pairs of edges k and l that share an endpoint, then normalizing by the average of f_k^2. The prefactor $s_{k,l} = 1$ if edges k and l both start or both end at the same node, and equals -1 otherwise. The prefactor $s_{k,l}$ is the k,l entry of the weighted adjacency matrix for the edge graph, A_E. The weighted adjacency matrix equals $GG^\intercal - 2I$ where G is the gradient operator. Therefore:

$$\begin{aligned} \rho(\mathcal{G}_{\leftrightarrows}) &= \frac{E}{\sum_{k,l}|[GG^\intercal - 2I]_{k,l}|}\frac{f^\intercal(GG^\intercal - 2I)f}{f^\intercal f} \\ &= \frac{E}{\sum_{k,l}|[GG^\intercal - 2I]_{k,l}|}\left(\frac{||G^\intercal f||^2}{||f||^2} - 2\right) \end{aligned} \tag{10.2}$$

The sum in the denominator is twice the total number of pairs of edges sharing an endpoint. The factor of two cancels since each pair of edges is counted twice in the quadratic product in the numerator.

For a complete tournament the projector from f to its transitive component is $\frac{1}{V}GG^\intercal$ [74]. Therefore $||G^\intercal f||^2 = f^\intercal GG^\intercal f = V f^\intercal f_t$. But $f = f_t + f_c$ where f_c is orthogonal to f_t since it is the projection of f onto the cyclic subspace, which is perpendicular to the transitive subspace. Therefore $f^\intercal f_t = f_t^\intercal f_t = ||f_t||^2$ and $f^\intercal GG^\intercal f = V||f_t||^2$.

For a complete tournament with V competitors there are $V-1$ edges leaving each competitor and $V(V-1)/2$ edges total. Therefore, each edge shares an endpoint with $2(V-2)$ other edges, so there are $V(V-1)(V-2)/2$ distinct pairs of edges sharing an endpoint. The cyclomatic number in a complete graph is $V(V-1)/2 - (V-1) = (V-1)(V-2)/2$. Therefore $L = (V-1)(V-2)/2$, and $\sum_{k,l}|[GG^\intercal - 2I]_{k,l}| = V(V-1)(V-2) = 2VL$.

Thus:

$$\rho(\mathcal{G}_{\leftrightarrows}) = \frac{E}{2VL}\left(\frac{V||f_t||^2}{||f||^2} - 2\right) \tag{10.3}$$

Solving for $||f_t||^2$ gives:

$$||f_t||^2 = ||f||^2 \left(\frac{2}{V} + 2\rho(\mathcal{G}_{\leftrightarrows})\frac{L}{E}\right) \tag{10.4}$$

In a complete network $(V-1)/E = 2/V$ since $E = V(V-1)/2$. Then, substituting in $||f||^2 = E\sigma^2(\mathcal{G}_{\leftrightarrows})$ yields the desired result:

$$\frac{1}{E}||f_t||^2 = \sigma^2(\mathcal{G}_{\leftrightarrows})\left(\frac{V-1}{E} + 2\rho\frac{L}{E}\right) \tag{10.5}$$

Since $f_c + f_t = f$ and f_c is orthogonal to f_t, $||f_c||^2 = ||f||^2 - ||f_t||^2$. Therefore:

$$\frac{1}{E}||f_c||^2 = \sigma^2(\mathcal{G}_{\leftrightarrows})(1 - 2\rho(\mathcal{G}_{\leftrightarrows}))\frac{L}{E} \tag{10.6}$$

□

11. A Trait-Performance Example. Suppose that each competitor has a set of T traits. Assume that the traits are chosen so that the performance function $f(x, y)$ is non-decreasing in x_j, and non-increasing in y_j, for all j. This amounts to choosing a sign convention for each trait so that increasing any trait improves performance. Then a competitor with traits x has an advantage (in trait j) over an opponent with traits y if $x_j > y_j$.

In some events, competitors with a large advantage in a given trait can dominate, so that the event is primarily mediated by that trait. That is, competitors press their advantages. For example, a performance function of this type is the extremal performance function $f(x, y) = x_j - y_j$, where j is the dimension in which this difference is largest in magnitude, $j = \text{argmax}_j|x_j - y_j|$. In the extremal performance model, the performance is completely controlled by the largest advantage, so competitive events are as one-sided as possible, given the competitor's traits.

Consider, in contrast, a competitive event in which competitors cannot press their advantages. For example: $f(x, y) = x_j - y_j$ for the dimension $j = \text{argmin}_j|x_j - y_j|$ that minimizes the advantage. This rule could model a contest in which competitors are required to reach a consensus about how to compete in advance or, where the weaker competitor controls which traits primarily mediate the competitive event. Competitors could be motivated or compelled to compete without pressing advantages by an external mediating body. For example, a sports league is motivated to keep teams evenly matched, even if the individual teams are motivated to win.

Suppose that the traits are drawn i.i.d from either an exponential, Gaussian, or uniform distribution. In each case, the variance of the trait distribution has no effect on ρ so, without loss of generality, each distribution is chosen to have variance one.

We estimated the correlation coefficient ρ for all six models (two performance functions, three distributions) with trait dimension varying from 1 to 25. To estimate the correlation coefficient for a given model and trait dimension we sampled 10^6 triples of trait vectors X, Y, W and computed $f(X, Y)f(X, W)$. Averaging over all 10^6 triples gave an empirical estimate for the covariance, which was then normalized by an empirical estimate of the variance σ^2. Figure 13 shows the results.

For all three choices of trait distribution, $\rho(T)$ was larger if the extremal advantage model was used instead of the fair-fight model. This indicates that, the more competitors can press their advantages, the more transitive competition is, on average. This is not surprising, since in the fair-fight model, the traits mediating performance for competitor A against competitor B are likely different from the traits mediating

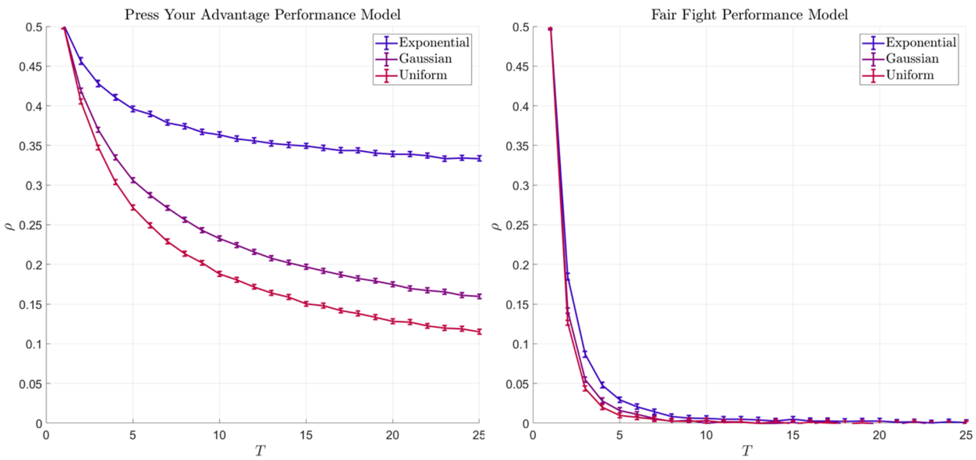

Fig. 13. *The correlation coefficient* ρ *for two different performance functions and three different trait distributions as a function of the number of competitive traits. Error bars represent three standard deviations in the estimated correlation coefficient. The "Press Your Advantage" panel shows* $\rho(T)$ *for the extremal performance model:* $f(x, y) = x_j - y_j$ *for* j *that maximizes the difference. The "Fair Fight" panel shows* $\rho(T)$ *for* $f(x, y) = x_j - y_j$ *for* j *that minimizes the difference. In all cases the correlation coefficient is higher in the 'Press-your-Advantage" model than in the "Fair-Fight" model. In both panels the correlation coefficient is larger for exponential than Gaussian traits, and Gaussian than uniform traits. In all cases* $\rho(T)$ *decreases with increasing trait dimension. The corresponding variances* σ^2 *are computed in the supplement.*

competition between A and C. As a result, the success of competitor A is highly competitor dependent. Thus competition is more cyclic.

Note that this conclusion is much easier to test using the trait-performance theorem than by sampling a series of random edge flows. We only needed to sample trait vectors for triples of competitors to evaluate ρ. This simplification greatly reduces the sampling cost.

In all six models tested, $\rho(T)$ is decreasing in T, so the expected proportion of competition that is cyclic is increasing. This matches the results in [44], where increasing the trait dimension typically decreased the expected degree of transitivity. This is intuitive, since larger T allows more ways for two competitors to compete, so it is harder to assign a single rating to a competitor.[18]

When using the extremal performance model the correlation $\rho(T)$ decays much faster in T for Gaussian and uniform traits than for exponential traits. This is because exponentially sampled traits are more likely to include large outliers. Since the extremal performance model sets f to the largest trait difference, the performance is more likely to depend on the outlier traits of each competitor. If a competitor has one particularly large trait, and T is large, then it is unlikely that any other competitor has a comparably large trait value in the same dimension. As a result, the competitor with the largest trait usually competes along that dimension and their performance against other competitors is fairly consistent. This leads to a relatively high ρ.

On the other hand, if the traits are drawn uniformly from $[0, 1]$ then no competitor can achieve a universal advantage by having one extremely large trait value. Instead, as the dimension of the trait space increases, competitors succeed by having a large trait value where their opponent has a small trait value - that is, by exploiting their

[18] Note that while this is often true it is *not* true for all trait-performance models.

opponents' weaknesses. In this situation, the relevant trait dimension that determines the outcome of competition depends on whom each competitor competes with. Consequently the correlation ρ becomes very small as T becomes large, so competition becomes predominantly cyclic.

In the fair-fight model all three trait distributions produce nearly identical correlations, since outlier traits do not mediate performance. Instead, performance is mediated by average traits, since the smallest advantage $X_j - Y_j$ is likely to come from a trait dimension where both X_j and Y_j are close to their expected values.

This example illustrates the explanatory power of the trait-performance theorem. By separating the influence of network topology from statistical assumptions about competition, the theorem facilitates numerical hypothesis testing and affords deeper insights by focusing the questions we ask about competitive tournaments.

12. Code Repository. A code repository is available at https://github.com/AlexRunsAway/HHD_and_Trait_Performance. The repository contains a read me file which explains the contents.