Article

Sports Tournaments and Social Choice Theory

Rory Smead

College of Social Sciences and Humanities, Northeastern University, Boston, MA 02115, USA;
r.smead@northeastern.edu

check for
updates

Abstract: Sports tournaments provide a procedure for producing a champion and ranking the contestants based on game results. As such, tournaments mirror aggregation methods in social choice theory, where diverse individual preferences are put together to form an overall social preference. This connection allows us a novel way of conceptualizing sports tournaments, their results, and significance. I argue that there are genuine intransitive dominance relationships in sports, that social choice theory provides a framework for understanding rankings in such situations and that these considerations provide a new reason to endorse championship pluralism.

Keywords: sports tournaments; team rankings; intransitive dominance; win-loops; social choice theory; Condorcet's paradox; championship pluralism

1. Introduction

The New York Giants won the Superbowl in 2008, but most fans do not consider them the best NFL team that year. In fact, it is relatively clear that the team the Giants defeated, the New England Patriots, were better even though the Giants outplayed them on the day of the big game. So it goes. Most fans would agree that the best team does not always win the game. Poor refereeing, bad luck, an uncharacteristic mistake, or just a great day from a scrappy underdog can cost the better team victory. Upsets happen [1], and since tournaments are just collections of games, it follows that the best team is not always declared a champion.

Of course, it is more complicated than that. While the Giants were not the best team that year, they arguably had the best overall performance in the NFL playoffs. Afterall, they won the Superbowl! Also, if we knew what team was best, why do we need the playoffs? Why do we need the tournament at all? Performance in a tournament is often taken to be a sign of quality, as evidence about the true nature of the team or athlete. If the label "winner" was not supposed to indicate quality, we could simply flip coins or roll dice to decide the results. When we, as spectators or participants, buy into a sports tournament we do not think about it as a fancy slot-machine. We think about it as a decision procedure. It is a method for choosing a champion. There are reasons to think that the victorious team and the better team do not always coincide [2], but it is generally supposed that there is a correlation between the two. And success in games (and tournaments) is at least taken as an indication of team quality, even if not always a perfect one.

But what does a tournament really tell us about the contestants? If the aim is to identify the "best" team, then what kind of tournament structure should we have? And, how does a tournament's structure, apart from the quality of the contestants, impact the results? These questions have close parallels in social choice theory, which explores how individual preferences or judgements should be aggregated to produce a single preference or judgement for a social group [3]. The analogous question in sports is how specific performances and game results should be compiled to produce a champion. Here, I will explore some central results in social choice theory and their implications for how we ought to interpret and understand sports tournaments.

After drawing parallels between tournaments and social choice (Sections 2 and 3), I argue that there are genuine intransitive dominance relationships in sports, as there are in social choice theory (Section 4). I then argue that key impossibility results from social choice theory do not apply to sports tournaments as readily as they do democratic choices (Section 5). I conclude with a discussion on how this impacts our understanding of tournaments and suggest that we should embrace championship pluralism.

2. Condorcet's Paradox: World Cup Edition

The 1994 FIFA World Cup Group E consisted of four teams: Mexico, Republic of Ireland, Italy, and Norway. They played a round-robin, each playing against every other team. As has become standard in association football, a win earns a team 3 points, a draw 1 point. The teams are then ranked by total number of points, the top 2 teams advancing to the next round, with the third-place team possibly advancing depending on how it compared with other third-place teams.

After all the games were played, every team had one win, one draw, and one loss, for a total of 4 points each. Table 1 shows the complete standings.

Table 1. FIFA 1994 World Cup Group E Final Standings [4].

TEAM	Games	Win	Draw	Loss	GS	GA	GD	PTS
Mexico	3	1	1	1	3	3	0	4
Rep of Ireland	3	1	1	1	2	2	0	4
Italy	3	1	1	1	2	2	0	4
Norway	3	1	1	1	1	1	0	4

Who should win this group? According to FIFA rules in 1994, Mexico won the group on the strength of goals scored (GS) and Republic of Ireland defeated Italy in their match and so was ranked above them on the head-to-head tiebreaker. So, one might argue that Mexico should have won, as those were the rules established. The tournament rules are simply an agreed upon convention for picking a champion, and the team that ought to win is simply the team that actually wins (assuming they followed the rules).

But there are other ways to understand the question of who ought to win this group. Perhaps we want our tournament to identify the objectively "best" team, or at least the one that had the "best" set of performances according to some true standard of the sport. In which case, there may be questions about whether some tournament rules or structures are better at identifying quality. Such considerations assume that there is something about the teams beyond the mere results that tournament rankings should track. Setting aside concerns about cheating or foul-play, debates over whether the "right team won" would seem to imply such a view.

Assume, for now, that there is some objective "ranking" among teams that tournaments are supposed to help us discern. Mexico won the group by total goals scored, which may arguably track offensive performance. But why favor offense rather than defense? If defensive performances were treated as more "important" than offensive performances in the tiebreaking rules, Norway would have switched places with Mexico and won the group. The argument may be that we want to incentivize scoring because it makes the games more exciting, but is excitement an indicator of team quality?

Taking a closer look at the specific results of the group games makes this even more interesting. Here are all 6 matches:

1. Republic of Ireland defeats Italy 1-0
2. Norway defeats Mexico 1-0
3. Italy defeats Norway 1-0
4. Mexico defeats Republic of Ireland 2-1
5. Italy and Mexico draw 1-1
6. Republic of Ireland and Norway draw 0-0

It is natural to think that a victorious team ought to be ranked higher than the defeated team. But if so, the first four matches generate an intransitive cycle: Republic of Ireland > Italy > Norway > Mexico > Republic of Ireland. I will call such a set of results a "win-loop". Such loops mirror the Condorcet Paradox in social choice theory [5].

Condorcet's paradox occurs when there are three or more candidates (e.g., A, B and C) in an election and where a majority of voters prefer A to B, B to C and C to A. This can occur if one-third of the population prefers A to B to C, one-third prefers B to C to A and the remaining third prefers C to A to B. Then, a total of two-thirds prefer A to B, two-thirds also prefer B to C, and two-thirds prefer C to A. There is no majority winner and an intransitive cycle occurs if we vote on pairwise comparisons.

To avoid the paradox in the context of the World Cup group one of the first four results must be effectively ignored or weighted less than others. In this case the Norway defeat of Mexico is not counted as significantly as others. Norway defeated Mexico, tied them on points, and yet it was Mexico that won the group and Norway that was eliminated. It would be easy to make the case that the most surprising result, however was Ireland defeating Italy. Italy was the top seeded team in that group and Ireland was the only team of the 4 that did not finish at the top of their groups during the qualifying phase. If we ignore game #1 instead and respect the other results, the ranking becomes Italy > Norway > Mexico > Ireland. However, then Ireland is in the situation where they defeated the winner, tied them on points, and yet are eliminated. What the paradox shows us is that some team was necessarily going to be aggrieved.

3. Condorcet's Jury Theorem and Series of Games

One might think that the problem is just an artifact of not enough games. Condorcet paradoxes are typically more likely to occur when there is a relatively small number of voters [6]. Perhaps the same is true with a small number of games played between sports teams. The thought is that if each team had played each other 10 times, we would better avoid this kind of problem. This is the rationale behind having teams play a series of games against one another. Any individual game may be subject to a chance outcome, but the victor in a best-of-five or a best-of-seven series (e.g., the playoffs in American Baseball, Hockey, or Basketball) is less likely to be determined by accidental outcomes.

This relates to another result from social choice theory: Condorcet's Jury Theorem (Condorcet 1785, Boland 1989). The theorem states that if you have several independent people making a judgement of some fact, provided each is (even slightly) more likely to make the correct determination than the incorrect one, the more people you have the more likely the aggregate is to be correct. The parallel in sports is that the more games are played, the more accurate the overall results will be.

However, this requires that the results be independent of one another and that the different games played between the teams be roughly equivalent, which is not generally true in sports. Some "bias" in game outcomes can be counter-acted and balanced between different games, such as having both teams play at their home field, as is done in two-legged contests in European soccer. In the best-of-seven series in many American sports, one team (usually the one with the better record) gets 4 games at home, the other 3. Regardless, the core idea is that the more results we have between teams, the more confident we can be in declaring one team better.

To illustrate this, suppose team A has a 60% chance to beat team B on any given day. For simplicity, let us ignore any homefield advantage or other exogenous factors. Here are the chances that A will win a "best-of" series: best-of-1 (60%), best-of-3 (64.8%), best-of-5 (68.3%), and best-of-7 (71%). This trend continues. The more games played, the more likely it is that the better team will be the winner of the series.

If this is correct, the tournaments that have more games will be more likely to truly select the best team as a champion. Single-elimination tournaments are prone to flukes and upsets, which is why the NCAA collegiate basketball tournament is so hard to predict every year. On the other hand, in best-of-7 tournaments, such as the NBA playoffs, it seems that the better teams are almost always the

ones that come out on top. Likewise, double round-robin tournaments should produce more reliable results than single-round robin tournaments.

4. Win Loops and Intransitive Dominance

The jury theorem reasoning works well when there are two teams involved, but it can break down when there are three or more teams. If team A defeats team B and team B defeats team C, this does not mean that A will defeat B, or even that it is probable that A will defeat B. Transitivity may fail accidentally or systematically.

Let A > B represent that team A is in fact more likely than not to beat team B (although team B may happen to win on any given day). Accidental failures of transitivity occur when we have A > B, B > C and A > C, but it just so happens that A defeats B (as expected), B defeats C (as expected), and C defeats A (in an upset). With respect to accidental failures of transitivity, the task of a ranking system is (arguably) to identify which results were accidents and remove them from consideration. By playing multiple games those "accidental" results become less likely and we can be more confident that the final rankings represent the true rankings. However, when win-loops are non-accidental, this no longer works.

It may be that the "true" relationship between the three teams is intransitive. Just as there is no best-move in rock-paper-scissors, teams or competitors can find themselves in situations where transitivity fails systematically. Seth Bordner argues that such genuine win-loops probably occur and that if they do occur "it poses immense problems for ranking conventions" [7] (p. 223). To explore this possibility here, consider the 2017 National League East Division in Major League Baseball. There are 5 teams and each played every other 19 times during the season. Table 2 shows the win-loss splits among those teams.

Table 2. National League East head-to-head records for 2017 [8].

TEAM	vs. ATL	vs. MIA	vs. NYM	vs. PHI	vs. WSH
ATL	N/A	11-8	7-12	6-13	9-10
MIA	8-11	N/A	12-7	8-11	6-13
NYM	12-7	7-12	N/A	12-7	6-13
PHI	13-6	8-11	7-12	N/A	8-11
WSH	10-9	13-6	13-6	11-8	N/A

The official final standings, based on total win-loss records were WSH (97-65), MIA (77-85), ATL (72-90), NYM (70-92) and PHI (66-96) [8]. However, looking at the team vs. team splits a more interesting relationship emerges. WSH is still clearly the top team, having a winning record against all other teams. However, the remaining four teams are more difficult to place in an ordering due to an intransitive relationship: MIA > NYM > PHI > ATL > MIA.

Is this win-loop representative of some genuine intransitive relationship between teams, or is it merely the result of some fluke in results? First, consider the latter possibility. Which ordering is the "fluke"? The only single result that could be removed to break the loop is the MIA > NYM. So, maybe that result is an accident and should be disregarded. But it is not so simple. That result is also consistent with the overall rankings and only has about an 0.18 probability of occurring by chance alone (assuming each team is expected to win half of their games). Moreover, these kinds of win-loops occur regularly in Major League Baseball. In the 2018 season for example, every team in the National League was involved in at least one such loop.

Could such an intransitivity represent the true relationship between the teams? The following simple model shows that such intransitive relationships are genuine possibilities and should not always be understood as mere statistical flukes. Suppose we have three tennis players, A, B, and C who each have distinct skills in serving, returning, and volleying. Suppose further that we can precisely quantify these skills on a 1 (poor) to 4 (excellent) scale and that each player has the skill levels in Table 3.

Table 3. Skills of three tennis players in a simple model.

Player	Serving	Returning	Volleying
Player A	4	2	1
Player B	1	4	2
Player C	3	1	3

Consider an idealized representation of a tennis match, where player i is playing against player j and the net expected points earned are calculated as follows.

- The fraction of serves that result in a volley (*v*) depends on the difference in serving and returning skills such that smaller differences mean volleys are more likely:

$$v = 0.8 - |S_i - R_j| * 0.1 - |S_j - R_i| * 0.1. \tag{1}$$

- Player i's net expected points earned on volleys is proportional to the difference in volleying skill and the percentage of volleys that occur:

$$v * (V_i - V_j)/4. \tag{2}$$

- Player i's net expected points earned on serves is:

$$(1 - v) * (S_i - R_j)/4. \tag{3}$$

- Player i's net expected points earned on returns is:

$$(1 - v) * (R_i - S_j)/4. \tag{4}$$

S_i is i's serving skill, R_i is i's returning skill, and V_i is i's volleying skill. Note that the "4" is simply a normalizing term representing the maximum possible difference in skill levels as it relates to expected points earned. Likewise, the other specific numbers chosen are not essential for the point, nor are they based on any precise data from tennis, they are merely chosen for ease of illustration and to capture the qualitative descriptions.

In this model, we can calculate the expected points for any given serve of a tennis match between any two players. Table 4 shows the expected outcome of each matchup.

Table 4. Expected points and results from each matchup in a simple tennis model. Net points are the average expected points earned for the first player listed in the matchup.

Matchup	Net Points for P1 on Serves	Net Points for P1 on Returns	Net Points for P1 from volleys	Total	Expected Result
A vs. B	0	0.075	−0.175	−0.1	B defeats A
A vs. C	0.45	−0.15	−0.2	0.1	A defeats C
B vs. C	0	0.075	−0.175	−0.1	C defeats B

Note that these expected points reflect the true skill levels of the players in the different aspects of the game and how they compare. There simply is not a transitive ordering of the players with respect to their quality. It will not help to play more games. Because the intransitivity is real and not an artifact of luck, the more games we play, the more probable it is that we see an intransitive cycle.

Given the frequency of win-loops occurring in sports and the fact that relatively simple models can be constructed that generate plausible intransitive relationships, it is reasonable to believe that there are genuine intransitive loops in sports. Note that intransitive dominance occurs in other domains as well. In ecology, for example, there is an ecological version of rock-paper-scissors among males

of the side-blotched lizard *Uta stansburiana* [9]. These lizards have three distinct mating strategies: a harem-building strategy, a mate-guarding strategy, and a cuckolding strategy. Harem-building does well amongst mate-guarders, cuckolding does well against harem-building, and mate-guarding does well against cuckolding. Which strategy is best (from a biological perspective) depends on the rest of the population. There simply is no true biological ranking of the behaviors. Another example is seen in games of chance with Nontransitive dice [10]. These are sets of dice which form intransitive win-probability cycles when played against one another. Every die can be beaten by some other die in these games.

If dominance is not a transitive relationship in sports, then perhaps there simply is not a best team or athlete [7]. This may be an acceptable conclusion for those that think we ought to participate in or watch sports simply for the love of the game. However, such a conclusion will be unsatisfying for the rest of us who like to think that sports tournaments and the crowning of a champion carry a special significance. Social choice theory shows us a way forward despite these difficulties.

5. Arrow's Theorem

Social choice theory has long had to deal with the intransitive relations that occur in Condorcet's paradox. It is easy to imagine cases where candidate A easily beats candidate B in an election, B easily beats candidate C, and C easily beats candidate A. Nevertheless, we cannot simply say "there is no true best candidate" and call it a day. Even if we think that's true, we still need to pick a winner. Many probably feel the same way about sports—the whole point is to have a winner, so we need to have a way of determining one.

This turns out to be much more difficult than simply settling on a conventional method for aggregating preferences. Arrow's Impossibility Theorem is a result that shows what kind of constraints social choice functions face [11,12]. It states that no social choice function can satisfy five desirable conditions: (i) have an unrestricted domain of individual preferences; (ii) have a positive association of social and individual values; (iii) be independent of irrelevant of alternatives; (iv) be non-imposed; and (v) be non-dictatorial.

Arrows theorem has corollaries in judgement aggregation [13,14], as well as in contexts such as scientific theory choice [15]. It also can be interpreted in the context of sport, where a particular set of game results is analogous to a voter's preferences and the social choice function, which aggregates preferences, is the tournament. The five conditions then become: (i) able to accommodate any possible set of game results and produce a complete ranking; (ii) winning games does not hurt a team's ranking in the tournament; (iii) the relative ranking of two teams should not depend on some other team; (iv) the ranking should not be determined by factors outside the results of the games and performances; and (v) we should not privilege one set of games over all others. Perhaps these are all desirable conditions for sports tournaments. However, each of them is violated in some contexts (see the Table 5), and considering the examples reveals that these are not usually taken as serious requirements for sports tournaments.

Table 5. Arrow's conditions and sports tournaments.

Condition	Example of Violation
(i) All possible game results are allowed, and the final ranking is complete.	Many tournaments do not allow for ties in match results, disallowing logically possible team comparisons. Additionally, not all tournaments require that the inputs (the individual games) be transitive or complete, they can be partial and intransitive. Likewise, the final ranking need not be complete (e.g., elimination tournaments often settle 1st, 2nd, 3rd, and 4th place, but do not settle lower rankings).
(ii) Winning never hurts a team's ranking.	Winning can hurt a team's chances of overall tournament success if the win causes a pairing in the next round that is significantly more difficult. The 2018 world cup featured such a game between Belgium and England. The winner (Belgium) faced a significantly more difficult slate of games in the elimination round as a result of defeating England.
(iii) The relative ranking of teams depends only on games between those teams.	Most round-robin style tournaments will violate this criteria as the overall ranking depends on how a team performs against the field rather than any specific team. For example, Arsenal defeated eventual Premier League champions Leicester City in both of their matches in 2015/16 but finished in second place.
(iv) Rankings are not imposed by anything other than game results.	The NCAA college football rankings dictate the entrants to the playoffs in ways (arguably) that do not depend only on the previous games played.
(v) No set of game results trumps all others.	The results of the NFL playoffs serve as the effective "dictator" of a team's final ranking, supplanting any regular season results that may have been different.

Cardinal vs. Ordinal Comparisons

What does Arrow's impossibility theorem mean for the rankings determined by sports tournaments? Arrow's result will apply to any tournament that relies only on wins/losses and has more than two contestants. However, given the examples of violations above, the conditions may be less important in sport than they are in democratic choice. Furthermore, rankings in sport need not be based merely on wins/losses. For instance, margin of victory in a given game may be considered.

The parallel in social choice theory is to consider the differences in preference rather than simply the preference ordering. This requires cardinal preferences rather than ordinal preferences, and a social choice function must assume some meaningful way of comparing such preferences. This allows a way around Arrow's theorem [3,16]. Whether these exist and whether they can be meaningfully compared is a contentious issue in social choice theory. In sports contests, however, margins of victory certainly exist and can be easily compared. Indeed, point differences are often used in tiebreaking procedures for round-robin style tournaments. This, combined with the several ways in which tournaments violate the five conditions, suggests that Arrow's theorem is not as serious a constraint on sports tournaments as it may be on democratic choice mechanisms.

6. Conclusions: Implications for Sport Tournaments

The above considerations have numerous implications for how we should think about sports tournaments, weight alternative tournament structures, and understand tournament results. Here, I will briefly comment on different types of tournaments before concluding with an argument for championship pluralism.

6.1. Elimination Tournaments

Interestingly, certain tournament structures preclude the possibility of intransitive cycles in results. Because losers of single-elimination tournaments are removed from the competition, there is no opportunity for that team to generate a win-loop. This allows us to take the results of such tournaments as definitive. We can guarantee that there are no results from within the tournament that would contradict the champion's claim. This comes with costs, however. Seeding in elimination tournaments becomes paramount. If the two best teams in the tournament are paired in the first round, we will be unable to identify that contest as a kind of "true" championship game. Thus, we should not be confident that the loser of the championship game is actually the second-best team (or even third, fourth, etc.). Additionally, elimination tournaments will not and do not provide a complete ranking of teams (e.g., a ranking between all the teams that were eliminated in the first round). Finally, elimination tournaments are prone to upsets having a significant impact on overall rankings, something that some fans like about such tournaments. While this may be desirable for excitement purposes, it certainly is detrimental to identifying any "true" ranking of the teams in so far as there is one.

Some of these concerns can be addressed by making the tournament more complex, such as double-elimination tournaments, or elimination tournaments where teams play multiple games between one another. These can, at best, only partially address the costs and can also re-introduce the possibility of intransitive win-loops (e.g., win-loops are possible in double elimination tournaments). Finally, if genuine intransitive relationships do exist among the teams, it can be rationally optimal to make predictions for an elimination tournament that are logically inconsistent, such as predicting team A to lose in the first round but win in the second [17].

6.2. Round-Robin Tournaments

Round-robin style tournaments are susceptible to win-loops and other related concerns. They do, however, provide a method for generating a complete ranking of teams. Additionally, round-robin tournaments produce more game results than elimination tournaments, meaning that the results are more likely to be representative of the true relationships among the teams. However, as we have seen, this can mean they are more likely to generate intransitive win-loops and complicate finding a "true" ranking. For these tournaments to generate a full ranking, there must be an established way of weighting some results as more important than others (e.g., factoring in win-margins, or overall points scored). This is usually done in the form of tie-breaking procedures. It is also important to note that because these tournaments effectively score performance against the field, the winners of round-robin tournaments can be teams that will systematically lose to other teams in the tournament.

Mixed tournament styles, such as the FIFA World Cup finals, involve both round-robin elements (in the group stage) and elimination elements (in the knockout stage). Mixed tournaments can help address some of the seeding concerns in elimination tournaments, but only at the cost of facing the potential complexities generated by round-robin style results, such as win-loops.

There are also group-style tournaments that are distinct from round-robin style, such as a Swiss-style tournament, where competitors are paired with others that have a similar ranking or aggregate score. The idea is to match contestants that are "closely" ranked, and thereby facilitate coming to a complete ranking more quickly with fewer matches being played. There are several different kinds of Swiss-style tournament systems that vary in exactly how contestants are matched. The core idea of these systems is to get a complete ranking without a complete set of games by differentiating competitors of similar strength and extrapolating the entire rankings from those results. Note that such a method assumes a kind of transitivity among competitors, and yet still does not eliminate the possibility of win-loops in results.

6.3. Alternative Ranking Criteria

Social choice theory offers a plethora of different ways to rank options given voters preferences [18,19]. Different voting procedures have different strengths and weaknesses. There is arguably not a single correct or best method. Rather, there are a plurality of methods, some of which may be preferred to others on a variety of grounds depending on what the aim of the vote is. For example, do we want to identify the option that is most preferred by the most people, or find the one that minimizes dissatisfaction?

The parallel with social choice theory and sports tournaments allows us to conceive of novel ways to conduct tournaments and produce final rankings. For example, the Ranked Pairs method [20] considers each pairwise comparison between options in an election and ranks the pairings in order of largest margin of victory by votes among the pair. Then, it produces a ranking among the options by locking in the results with the largest margins first. If an intransitive cycle occurs, it discards the result with the lowest margin. Such a procedure would allow for a novel alternative way of ranking teams in a round-robin tournament.

The Ranked Pair procedure satisfies the Condorcet criterion, which holds that if one option (team) is preferred (wins) to all others in pairwise contests (games) that option ought to be preferred on aggregate. This condition, which seems highly intuitive in the context of sport, is not satisfied by many league structures. Take the Premier League in English association football for example. The Premier League follows a standard process in European soccer of having teams play a double-round robin, awarding three points for a win and one point for a tie. Teams are then ranked by the number of points they accumulate during the season. If there were a team that somehow won every game they played in their home stadium and tied every game they played away from that stadium they would amass 76 points across 38 games, never be defeated, and have beaten every other team on aggregate. However, 76 points would have been enough to win the league only once in its recent history. In many years, this total would have ranked third place or lower. The setup of the Premier League allows for a team to beat all other teams, never lose to any team, and yet not win the league. This violates the equivalent of the Condorcet criterion.

This is but one example. Social choice theory has produced a vast array of possible choice mechanisms and explored their effects and tradeoffs. Corresponding tournament rules or structures could be implemented that parallel virtually any voting or social choice mechanism. The parallel between social choice and sports tournaments allows us to draw on these ideas to re-consider how we conduct tournaments.

6.4. Championship Pluralism

Aaron Harper argues for what he calls "Championship Pluralism" [21], saying we should accept that one team may be the regular season champion, and another may be the playoff champion, and neither is strictly better. He argues for pluralism on the grounds that "no single format for measuring and comparing athletic excellence successfully captures all conceptions of excellence" [21] (p. 307). Considerations from social choice theory and the parallels with sports tournaments offer a distinct reason for embracing championship pluralism. Even if there were a tournament that could capture all aspects of athletic excellence, there is good reason to suspect there may be intransitive dominance among teams. In which case, the method of compiling results—the rules and structure of the tournament—may be the deciding factor. There is arguably no best tournament structure, just as there is no best social choice mechanism, but we (often) want to produce a winner nonetheless. Therefore, we should accept a plurality of tournaments and the plurality of champions that comes with it.

References

1. Fry, J.P. Underdogs, upsets, and overachievers. *J. Philos. Sport* **2017**, *44*, 15–28. [CrossRef]
2. Dixon, N. On winning and athletic superiority. *J. Philos. Sport* **1999**, *26*, 10–26. [CrossRef]
3. List, C. Social choice theory. *Stanf. Encycl. Philos.* **2013**. Available online: https://plato.stanford.edu/entries/social-choice/ (accessed on 28 March 2019).
4. FIFA.com. Available online: https://www.fifa.com/worldcup/archive/usa1994/groups/index.html (accessed on 28 April 2019).
5. Condorcet, M. *Essai sur L'application de L'analyse à la Probabilité des Décisions Rendues à la Probabilité des Voix*; De L'imprimerie Royale: Paris, France, 1785.
6. Gehrlein, W.V. Condorcet's paradox and the likelihood of its occurrence: different perspectives on balanced preferences. *Theory Decis.* **2002**, *52*, 171–199. [CrossRef]
7. Bordner, S.S. 'All-things-considered,' 'better-than,' and sports rankings. *J. Philos. Sport* **2016**, *43*, 215–232. [CrossRef]
8. ESPN.com. Available online: http://www.espn.com/mlb/standings/grid/_/year/2017 (accessed on 28 April 2019).
9. Alonzo, S.H.; Sinervo, B. Mate choice games, context-dependent good genes, and genetic cycles in the side-blotched lizard, Uta stansburiana. *Behav. Ecol. Sociobiol.* **2001**, *49*, 176–186. [CrossRef]
10. Tenney, R.L.; Foster, C.C. Non-transitive dominance. *Math. Mag.* **1976**, *49*, 115–120. [CrossRef]
11. Arrow, K.J. A difficulty in the concept of social welfare. *J. Polit. Econ.* **1950**, *58*, 328–346. [CrossRef]
12. Arrow, K.J. *Social Choice and Individual Values*; Yale University Press: New Haven, CT, USA, 1963.
13. List, C.; Polak, B. Introduction to judgment aggregation. *J. Econ. Theory* **2010**, *145*, 441–466. [CrossRef]
14. List, C. The theory of judgment aggregation: An introductory review. *Synthese* **2012**, *187*, 179–207. [CrossRef]
15. Okasha, S. Theory choice and social choice: Kuhn versus Arrow. *Mind* **2011**, *120*, 83–115. [CrossRef]
16. Sen, A. The possibility of social choice. *Am. Econ. Rev.* **1999**, *89*, 349–378. [CrossRef]
17. Smead, R. On the rationality of inconsistent predictions: The March Madness paradox. *J. Philos. Sport* **2016**, *43*, 163–169. [CrossRef]
18. Brams, S.J.; Fishburn, P.C. Voting procedures. In *Handbook of Social Choice and Welfare*; Arrow, K.A., Sen, A., Suzumura, K., Eds.; Elsevier: Amsterdam, The Netherlands, 2002; Volume 1, pp. 173–236.
19. List, C. The logical space of democracy. *Philos. Public Aff.* **2011**, *39*, 262–297. [CrossRef]
20. Tideman, T.N. Independence of clones as a criterion for voting rules. *Soc. Choice Welfare* **1987**, *4*, 185–206. [CrossRef]
21. Harper, A. "You're the best around": An argument for playoffs and tournaments. *J. Philos. Sport* **2016**, *43*, 295–309. [CrossRef]