**Abstract**

Given a set of competitors, a tournament rule maps the results of games played between pairs of competitors to a single champion of the tournament. Sports competitions have utilized a wide variety of tournament rules, sometimes producing outcomes and incentives which seem unfair or somehow incorrect. We will mathematically formalize this setup and discuss three properties which designers might want a tournament to have (with each property motivated by a real-life tournament which lacked the property and went awry). We will discuss some of the existing literature on designing tournaments with these properties and (time-permitting) discuss some open problems.

Section 2 is a summary of [5]; technical results are from there unless otherwise cited. Some of the sports facts in Section 1 are from [5] as well, but most are just from my memory and Wikipedia. Most sports games are played between just two teams at a time. However, we almost always have more than two teams. A tournament takes in more than two teams, has them play games against each other and eventually crown a champion (the "best" team). We will discuss some basic requirements for what a "good" tournament should be, based on historical incidents where these conditions were not met. We then discuss the mathematical framework usually used to discuss tournaments, and some problems in this area (which as far as I know are open).

# 1 SPORTS!

Most sports games are played between just two teams at a time. However, we almost always have more than two teams. A tournament takes in more than two teams, has them play games against each other and eventually crown a champion (the "best" team).

People have come up with a lot of different ways to run tournaments, but some of them end up offering strange incentives to competitors, or end in results that "feel" wrong. Let's start with some historical examples of things going wrong so we can better discuss our basic requirements for tournaments.

Since it's September, let's start with the sport that is happening right now – college football. There are 128 major college football teams, each of which play only 12 to 14 games in a season. For (fascinating) historical reasons I could talk about at length (but won't) teams self-assembled into pseudo-regional groups which agree to play each other, and pick their "conference" champion. But we want a national champion. Between 1998 and 2013, the solution to this was the Bowl Championship Series. The BCS used a complicated formula to determine who the top 2 teams were and held a championship game to determine the national champion. Great, right? Well, this system wasn't quite perfect. In 2009, three teams (Boise

State, Alabama, and Texas) each won all of their regular season games. Since 3 is more than 2, it wasn't possible for all of these teams to play in the national championship game. Boise State was left out in the cold (Alabama would defeat Texas and was recognized as the champion). The year before was even worse. In 2008, Utah ended the season undefeated, but Florida (who had already lost a game) was selected to play in the championship game, won and was declared the national championship with one loss. This eventually led to Senator Orrin Hatch[1] calling for a Department Of Justice investigation and holding official Senate hearings regarding the way a champion is crowned (seriously) [3]. [2]

For our second incident, we head to the Olympics. The 2012 badminton women's doubles tournament featured 16 teams. The tournament ran as follows: teams were divided into four groups of four (called A,B,C,D). Each group would play a round robin. The top two teams advanced to a single elimination tournament bracketed as follows: A1 vs. C2 B1 vs. D2 on one half and A2 vs. C1 and B2 vs. D1 on the other. Group D finished with a surprise – one of the favorites to win the whole tournament lost one of their matches, causing them to come in second place in the group. This left group A (featuring the project gold and bronze medalists) with interesting incentives. Entering the last match the standings were:

Group A

| | |
|---|---|
| Team 1 | 2-0 |
| Team 2 | 2-0 |
| Team 3 | 1-2 |
| Team 4 | 0-3 |

Both teams were guaranteed to advance, it was only a matter of seeding. Both teams wished to end up on the far side of the bracket from D2. (This increases the chance of D2 losing in an upset before they meet, plus it guarantees them the silver if they can beat everyone else, otherwise your guarantee is just a chance to play for bronze). Thus they actually each wanted to *lose* this match. Both teams tried this (serves were repeatedly sent into the net, no rally went more than 4 shots, it was bad). Hilariously(?) once these matches had completed group C leaders also wanted to manipulate where they were seeded, and those teams too tried to lose. All 4 were disqualified. [3]

We have one more historical incident to discuss (this one quite a bit older). We go back to the 1982 World Cup. The World Cup uses a similar group play to what we just saw in the

[1](not) coincidentally, a senator from Utah

[2]This problem was "fixed" with the introduction of the College Football Playoff, which has four teams. But this doesn't actually completely solve the problem. A dozen (or more) teams could theoretically end the season undefeated (10 conference champs along with some independents), and only 4 playoff spots are available.

[3]In 2016 this was "fixed" by placing the second place teams randomly against a first place team in the draw. This doesn't actually fix the incentives though. Given this scenario, one still increases the chance of deferring facing the most formidable opponent by losing

Olympics for the first round. Before the last game these were the standings:

| Team | Rec. | GD |
|---|---|---|
| Austria | 2-0 | +3 |
| Algeria | 2-1 | +0 |
| West Germany | 1-1 | +2 |
| Chile | 0-3 | -5 |

Again, two teams would advance from this group. In the event of a tie in the standings, goal differential would be used as the tiebreaker. Note that both West Germany and Austria advance if Germany wins by only a few goals. But if Germany wins by 4 Austria is eliminated, and Germany is eliminated in a loss or tie. Funnily enough, a 1-goal victory for Germany is exactly what happened. Indeed, they scored a goal in the first 10 minutes, and then neither side did anything for the next 80 minutes.[4]

When each of these incidents occured, the tournament designers admitted they had a problem, and tried to fix it, but (if you've read the footnotes), their fixes have never altered the setup enough to completely fix the underlying problem. We want a tournament to pick the team that is the best at the sport we're playing, not the ones who are the best at winning tournaments – i.e. teams should always be incentivized to do their best to win, and we shouldn't have situations like these. Let's try to discuss formally how we might handle this.

## 2 Mathematical Formalism

A "tournament" is a round-robin tournament. I.e. it is a result (X beat Y or Y beat X) for all pairs of teams in the tournament (or more generally a probability that X would beat Y, which is resolved to a result when the tournament rule examines it). A tournament-rule is a way of assigning a champion from a given tournament.

With our current definition, we can do really silly things (player 1 is always the champion, champion is selected uniformly at random ignoring results, etc.) thus we will have a few requirements on our rule to enforce our "just win the games" idea.

1. A rule should be fair. (i.e. the BCS was bad)
   A rule is called **Condorcet-consistent** if a team that defeats every other team in the tournament is always the champion.

[4]This problem was also "fixed" by having the last games for a group start simultaneously. Bad situations can still arise. See 2010 UEFA incident.

2. A rule should encourage teams to win(i.e. the badminton setup was bad)
   A rule is **monotone** if $\Pr[i$ becomes champion $|i$ beats $j \wedge$ partial tournament $T] \geq \Pr[i$ becomes champion $|T]$

3. A rule should be non-manipulable. (i.e. the World Cup setup was bad)
   Suppose before a match is played, $\Pr[i$ becomes champion or $j$ becomes champion$] = p$. A rule is **2SNM**[5] if $i$ and $j$ cannot conspire to change the result of the match such that after the match $\Pr[i$ becomes champion or $j$ becomes champion$] > p$

There are two ways we can view non-manipulability. One intuition is that it is trying to eliminate bribery concerns. Let $r_i(T), r_j(T)$ be the probability that $i$ and $j$ win without colluding respectively, and let $r_i(T'), r_j(T')$ be the probabilities after colluding. 2-SNM says that $r_i(T) + r_j(T) = r_i(T') + r_j(T')$ Thus if this is false we have (WOLOG) $r_i(T) + r_j(T) > r_i(T') + r_j(T')$ or $r_i(T) - r_i(T') > r_j(T) - r_j(T')$ that is, the probability of being the champion gained by $i$ for winning is greater than the probability lost by $j$ for losing. Intuitively: $i$ is rewarded more for winning that $j$ is punished, and thus $i$ is more incentivized to win. If we were playing in a tournament with a cash prize, for example, $i$ could pay $j$ a certain amount of money, and have both improve their winnings in expectation.

Alternative intuition: if going into a match, $i$ and $j$ have probability $p$ of being champion, and leaving they have probability $p + q$, that probability $q$ came from somewhere. Where did that $q$ come from? It must have been from the teams not in that match. Intuitively it's "unfair" for large amounts of probability to suddenly transfer from $k$ to $i$ and $j$ when $k$ is watching TV, not playing a game. Such things happen in real-life tournaments (e.g. an upset somewhere else in a bracket can reduce ones chances of winning, a game which affects tiebreakers can also affect other teams chances of winning) but minimizing this differential seems like a good idea.

It turns out it is impossible to have a rule that follows both Condorcet consistency and 2-SNM, so we have to weaken one [1, 2]. In some (non-sport related) instances (if you don't actually know results explicitly) weakening Condorcet-consistency may work. In sports, lacking Condorcet-consistency is a really bad idea,[6] so we have to weaken non-manipulability.

Thus we weaken the non-manipulability requirement to the following: A rule is 2SNM-$\alpha$ if for all $i$ and $j$ $|r_i(T) + r_j(T) - (r_i(T') + r_j(T'))| \leq \alpha$. That is the probability of a player in this game winning increases by at most $\alpha$ by manipulating the result of the match between $i$ and $j$.

Why is this relevant? I'm risking something by trying to get you to collude. You might tell

[5]short for "2- strongly non-manipulable" where the 2 refers to this being a group of 2 players colluding and not more

[6]Orrin Hatch is still a senator, and you don't want to cross him

the tournament organizers (for example). Even if you don't, I'll know in my heart I didn't win fair and square. If the probability of me winning can go from, say, 1/3 to 1 by getting you to throw a match, (and yours just goes from 1/3 to 0 say) I might be willing to look past these moral objections. On the other hand, if my chances of winning go from 1/3 to $1/3 + \epsilon$ for a really small $\epsilon$ I would probably prefer to try and win fair and square.

Thus we should try to find the smallest $\alpha$ for which we can say we have a Condorcet-consistent (i.e. fair) rule which is also 2SNM-$\alpha$.

It turns out that 1/3 is the best possible, and that 1/3 is attainable.

To see that it is the best possible, consider a tournament on $n$ people, where $a$ beats $b$, $b$ beats $c$ and $c$ beats $a$ and $a, b, c$ beat all of the $n-3$ remaining opponents. Note there is no Condorcet winner, but all of $a, b, c$ are one match away from being the Condorcet winner. There is some size two subset of $\{a, b, c\}$ (WOLOG $a, b$) which has combined probability at most 2/3 of winning under the rule. By throwing its match, $a$ can cause $b$ to be a Condorcet winner, so in the altered tournament, $b$ wins with probability 1, and the increase is $1 - 2/3 = 1/3$. [7]

On the other hand, a randomized single-elimination tournament produces $\alpha = 1/3$ (this isn't obvious, it requires a few pages of argument). There's a catch here though, the decrease comes because of randomization – essentially 2/3 of the time, that tournament will make it so we don't play each other, or by the time we play each other it's too late for us to affect the tournament outcome. But once the seeding is fixed, we haven't really solved anything. Indeed, if we have that transitive group of three construction, whichever two meet first will have the incentive to collude.

# 3 Possible Extensions

This is a really nice result, but it leaves something to be desired. I see openings on two levels:

[7]One might wonder whether this transitive set is realistic. In 2013 football the top three teams in the Big XII were Baylor, Oklahoma, and Oklahoma State who beat each other in a cycle (though Oklahoma and Oklahoma State each had an additional conference loss). 2012 Big East football also had a cycle at the top (Louisville, Rutgurs, Syracuse) with Cincinatti joining the tie by giving Syracuse their second loss, and the others losing unrelated games. In the same year, the Mountain West truly had this setup. Boise State, Fresno State and San Diego State beat each other in a cycle, winning all their remaining games. Indeed, "circles of suck" (directed cycles with each team beating the next along the cycle) encompassing almost all of 12+ team conferences are common by the end of the season. Some version of a cycle also apparently occurs in tennis, among Novak Djokavic, Rafael Nadal, and Roger Federer. `http://grantland.com/features/the-tennis-triangle/`

1. In the tournament that produces the lower bound, all matches are won with probability 1 or 0. A more realistic scenario has that everyone wins or loses a match with probability in $[\epsilon, 1-\epsilon]$ for some $\epsilon > 0$. The lower bound doesn't hold here – could we do something better in this regime? (This is one of the open questions mentioned in [5]) Moreover, if we're not entirely convinced this is the correct notion of non-manipulable (i.e. we really want one that applies right before the game starts, say, or at least after all randomness has been applied), it's possible some other definition handles that bad example better.

2. As a sports fan this is disappointing – single-elimination tournaments are about the simplest thing you could think of if you wanted a tournament. That hasn't stopped people from designing other tournament rules! There are reasons to not use single elimination:

   - It's devastating to half of the teams that show up to only play one game and then go home. This is probably why the Olympics are so heavy on two stage (group play then single elimination) rules. [8]
   - Relatedly, in real life all teams will lose with some probability – it's bad when one of the best/most high profile teams is knocked out immediately on a fluke. A design that allows for a team to lose a game is often more fair, especially in sports with high randomness (e.g. baseball).

Moreover, the existing mathematical framework leaves something to be desired – tournaments where teams play more than once (e.g. any two phase Olympic tournament, a double elimination tournament, a double round robin, etc.) are not directly encodable in the round-robin tournament framework. Thus we'd like to propose the following question:

Suppose we require a tournament have one of the following extra conditions:

1. Every team is guaranteed to play in at least 2 meaningful games (where meaningful means "they have probability $> 0$ of becoming champion entering the game")
2. A single loss cannot eliminate a team

[Note that the second condition is strictly stronger than the first – traditional pool-then-single-elimination tournaments meet condition one but not two.]

[8] In summer 2016 group play was used for: Badminton [despite the 2012 kerfuffle], Basketball, Beach and regular volleyball, Field hockey, Handball, Rugby, Soccer, and Water Polo. (and it's not like the Olympic organizers don't know about single elimination tournaments, they're used in tennis, table tennis, fencing, etc.)

First, what does it mean to be Condorcet consistent in this new setting?

Can we still get a Condorcet-consistent and monotone tournament? Can we get 2SNM-1/3? Can we get close?

A paper by Pauly [4] showed that if you have a two-phase tournament, where phase 1 was a round robin, you could not find a second phase such that the entire tournament was symmetric (i.e. being in group A is not better than being in group B), condorcet-consistent and monotone (i.e. you are never incentivized to lose). Could we replace the round robin with something meeting our condition 1 say, and finish with a single elimination and still meet all of our desired properties?

What about college baseball? Is there another option, other than the format they use now?

# References

[1] Alon Altman and Robert Kleinberg. Nonmanipulable randomized tournament selections. In *AAAI*, 2010.

[2] Alon Altman, Ariel D Procaccia, and Moshe Tennenholtz. Nonmanipulable selections from a tournament. In *Dagstuhl Seminar Proceedings*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2010.

[3] Orrin Hatch. Hatch requests doj investigation into bcs. `http://www.hatch.senate.gov/public/index.cfm/2009/10/hatch-requests-doj-investigation-into-bcs`, 2009. Accessed: 2016-09-01.

[4] Marc Pauly. Can strategizing in round-robin subtournaments be avoided? *Social Choice and Welfare*, 43(1):29–46, 2014.

[5] Jon Schneider, Ariel Schvartzman, and S. Matthew Weinberg. Condorcet-consistent and approximately strategyproof tournament rules. *CoRR*, abs/1605.09733, 2016.