

What Is the Correct Way to Seed a Knockout Tournament?

Allen J. Schwenk

1. WHY ARE TOURNAMENTS SEEDED? In an *elimination tournament*, also called a *knockout* tournament, teams (or individual competitors) play head-to-head matches with the loser eliminated from further competition and the winner progressing to the next round of competition. This form of competition is widespread, and includes such popular tournaments as NCAA basketball; professional playoffs in football, baseball, basketball, and hockey; individual competitions such as Olympic boxing and wrestling; and Wimbledon and other tennis tournaments. It is also used in top-level bridge tournaments. In the purest form of knockout tournament, the number of teams entered is precisely a power of two. The NCAA basketball tournament (recently having 64 teams) is an excellent example. But the general method can be used with other numbers by having first round byes as needed to reach a power of two in the second round. For example, in the National Football League, 12 teams qualify for the playoffs. In the first weekend, four division winners are given byes, and the other eight compete so that the field is reduced to eight, the four winners and the four byes. Similarly, in a tennis tournament it would be exceedingly lucky to have the number of entrants just happen to be a power of two. First round byes solve the problem, but they have the disadvantage of leaving some players idle during the first round.

In competitive bridge, such as national team titles (Reisinger, Vanderbilt, and Spingold Championships) and the international team championship (the Bermuda Bowl), it is possible to arrange some three-way matches in the first round. This approach has the advantage of not leaving anyone idle during the first round, but of course it just doesn't work for boxing or wrestling. Depending upon the number of original teams, either one or two teams from each three-way match survive. For this approach, suppose the number of teams is $n = 2^k + r$ with $0 \leq r < 2^k$. When $r < 2^{k-1}$, the schedule uses r three-way matches (with one survivor from each) and $2^{k-1} - r$ two-way (head-to-head) matches to reach 2^{k-1} survivors entering the second round. When $2^{k-1} < r < 2^k$, we need $2^k - r$ three-way matches (with two survivors each) and $2r - 2^k$ head-to-head matches, again reaching 2^{k-1} survivors entering the second round. When r is precisely 2^{k-1} , we can have 2^{k-1} three-way matches perfectly covering the entire field. Here we can choose to allow either one or two survivors from each match, so that the second round field size is either 2^{k-1} or 2^k . When time permits, it seems that the larger field of size 2^k is preferred.

2. AXIOMS FOR FAIR SEEDING. For the rest of this article, we assume that our field size is a perfect power of two, $n = 2^k$. We imagine the NCAA tournament as our model competition, but we use a field of size 16 (or occasionally 8 or even 4) to illustrate the issues we wish to examine. A field of size 64 is just too large to build convenient

examples. For the sake of our analysis, we assume the teams have measurable strengths that permit them to be sorted into a linear order that everyone (or perhaps nearly everyone) finds appropriate. In reality, we understand that rankings determined by polls, for example, are always controversial. The strongest teams are given the highest seeds (smallest numbers, such as first seed, second seed, ...); the weakest team gets the lowest seed (largest number). We assume that whenever T_i faces T_j there is a certain unchanging probability that T_i will win, which is given by $p_{i,j}$. Of course we know that sport events are inherently unpredictable. Some lower seeds may actually be favored in a game against a higher seed. And many repeat matchups of a particular pair lead to dramatically different outcomes. Were this not so, there would be little reason to watch. Nevertheless, for purposes of analysis we assume that the seeding has been assigned in such a way that the seeding is accepted in the sense that for all $i < j < k$, (nearly) everyone agrees that T_i is better than or equal to T_j , which in turn is better than or equal to T_k . Moreover, if these teams face one another, the probabilities are monotonic in that $p_{i,k} \geq p_{i,j} \geq 0.5$ and $p_{i,k} \geq p_{j,k} \geq 0.5$. This fact can be nicely reported in a matrix of probabilities $P = [p_{i,j}]$. To avoid vacancies on the diagonal, we arbitrarily set each $p_{i,i} = 0.5$. Now it is clear that $P + P^T = J$, the matrix of all 1's. The monotonicity property requires P to have nonincreasing rows and nondecreasing columns. Let us call such matrices doubly monotonic.

If we wish, we could design a schedule in which teams are placed at random, but this is seldom done. Why not? Imagine a situation in which the two strongest teams are paired by chance in the first round of competition. One of them must then be eliminated while perhaps the two weakest teams may be meeting in another match and one is assured survival. This strikes people interested in good fair competition as grossly unfair. When the sport is one with great spectator interest and lots of television money, a first round match of the two top contenders squashes further interest in the tournament, and would be financially ruinous. Organizers nearly always assign “seeded” positions in the schedule to prevent such disasters. Seeding refers to the careful placement of the top competitors to avoid premature meetings, and thereby improve both the “fairness” of the schedule and the general interest for the public.

What axioms ought to govern a good seeding method, and what is usually done in practice? In some cases, seeding position is earned by past performance, as in most professional sports playoff systems where the strongest season record is rewarded with the best seed position. In other cases, such as college basketball, professional tennis, and bridge tournaments, a committee of experts assigns seeding positions based on ratings of the various teams. Again let us emphasize that we assume that the rankings have general acceptance among experts and among the public. We want to study how these rankings ought to lead to a fair schedule. We have already discussed the first principle of tournament seeding, and we now state it as:

Axiom DC: Delayed Confrontation. Two teams rated among the top 2^j shall never meet until the field has been reduced to 2^j or fewer teams.

Presumably there will be no dissension about Delayed Confrontation, since it is the *raison d'etre* behind seeding. The better teams should not be required to face one another until it can no longer be avoided. When 2^{j+1} or more teams remain, it is not necessary for any of the top 2^j teams to meet, since enough lower-rated teams remain to fill out the 2^j opposing slots. But when precisely 2^j teams remain, certainly they have to face each other. Of course, we should still keep the top 2^{j-1} separated. Note how seeding serves to favor

and protect the top teams. Teams in the lower half of the field receive no consideration whatsoever. If anything, they are viewed as cannon fodder.

Occasionally, in sports where the seeding position is earned by performance, it has developed that a potential “top-seeded” team viewing its likely seeding position finds that the first opponent assigned by the seeding method is not the one it would most prefer, even though the method was designed to favor the top team. This results when the first opponent, a team with relatively weak record, is deemed stronger than its record would indicate. The top seed might then prefer to drop into second seed position drawing a nominally stronger opponent whom the top team actually judges to be weaker. This leads to some embarrassing situations where a team may perform at less than maximum effort in order to drop to the second seed position and gain (in its own perception) a preferred opponent. In the worst case, opposing teams in the final game of the season both conclude that they would be better off losing, leading to comical competition. Leagues and commissioners sometimes threaten punishment against such perceived efforts to manipulate the system. But why should a team have to face a conflict between ethical behavior and pursuing its best chance of winning the tournament? If such manipulation is even possible, then the system is improperly designed. It should never benefit a competitor to lower his seeding position intentionally. We state this as the second axiom of seeding:

Axiom SR: Sincerity Rewarded. A higher-seeded team should never be penalized by being given a schedule more difficult than that of any lower seed.

Again we have already given the argument supporting this axiom. If our seeding mechanism is well designed, it does not benefit a team to seek to lower its own seeding intentionally. If every team is rewarded for seeking to do its best, we can expect each team to do its best. This surely promotes public confidence in the integrity of the competition. Another way to describe this axiom is to observe that no team should ever envy the schedule being assigned to a lower seed.

Is there any other axiom we need to observe? Well, we’d like the competition to be *fair*, whatever that means. Generally anything we don’t like, anything that militates against our favorite team, we call unfair. In this article we shall take fair to mean that each team’s probability of winning should somehow reflect its inherent strength and not be a result of receiving a favorable schedule. In order to separate the effects of inherent strength and favorable schedule, we imagine replaying the tournament repeatedly using every possible schedule. The average over all schedules of the championship probabilities for a given team is a measure of that team’s expectation based on its ability. Whenever a particular schedule gives that team a probability higher than its average, it is receiving favoritism. The amount of favoritism associated with a given schedule is the maximum over all teams of these probability deviations. Also, if our seeding method has any lower seed with a probability of winning that is higher than that of a higher seed, it must be an effect of the schedule rather than ability. If so, this violates Axiom SR. In general, we can not expect a chosen schedule to perfectly match the average winning chances for each team. But we can dictate that the maximum amount of favoritism that accrues to any single team be minimized.

Axiom FM: Favoritism Minimized. The schedule should minimize favoritism to any particular seed. (Or, if you prefer, Fairness Maximized.)

Of course this axiom needs to be observed subject to the implications of the other axioms. A totally random schedule has no Favoritism and does not punish Sincerity, but it

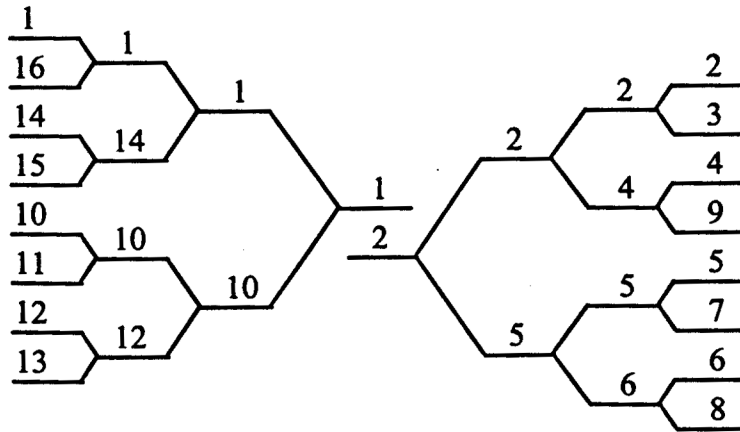


Figure 1: A very biased schedule.

can grossly violate Delayed Confrontation. Meanwhile, DC automatically favors the top half of the field in the first round, guaranteeing them bottom half opponents. We need to implement DC with Minimal Favoritism. Three examples illustrate the importance of Minimizing Favoritism.

Look at the schedule imagined in Figure 1. Here we have filled in the higher seed as probable winner in each succeeding round to indicate the likely opponents faced by a team on its way to the finals. Could we possibly design a schedule that is more favorable to team T_1 and more biased against T_2 ? Most likely, T_1 will face T_{16} , T_{14} , and T_{10} on its way to the championship game. Meanwhile if T_2 is going to survive, it will probably have to defeat T_3 , T_4 , and T_5 ! Or if T_3 wins the opening game, it will have to defeat T_2 , T_4 , and T_5 . Surely team T_1 , is being favored over T_2 , and also T_3 .

Now a cynic might assert that the tournament organizers really don't care about "fairness". What they really desire is to maximize the probability that the top two seeds survive to meet in the climactic championship game. If that were truly so, the seeded schedule would look something like Figure 2. Here T_1 and T_2 are practically given a free ride into the finals while teams T_3 through T_6 have it tough. Surely such a schedule would stir up outrage and not be tolerated. While the sponsors cross their fingers and hope for top seed survival, they need to conduct a tournament perceived to be fair. Thus, we contend that Axiom FM is essential.

A mythical third example shows how sensitive the outcome can be to selection of schedule; four teams suffice to illustrate. Suppose three teams, the Broncos, the Chippewas, and the Eagles, are nearly equally strong, but the fourth team, the Wolverines, is really outclassed. Probabilities for single games are given in the doubly monotonic matrix

$$P = \begin{bmatrix} .50 & .51 & .52 & .99 \\ .49 & .50 & .51 & .98 \\ .48 & .49 & .50 & .97 \\ .01 & .02 & .03 & .50 \end{bmatrix}$$

The Broncos are favored over the Chippewas by 51% to 49%, and the Chippewas in turn are

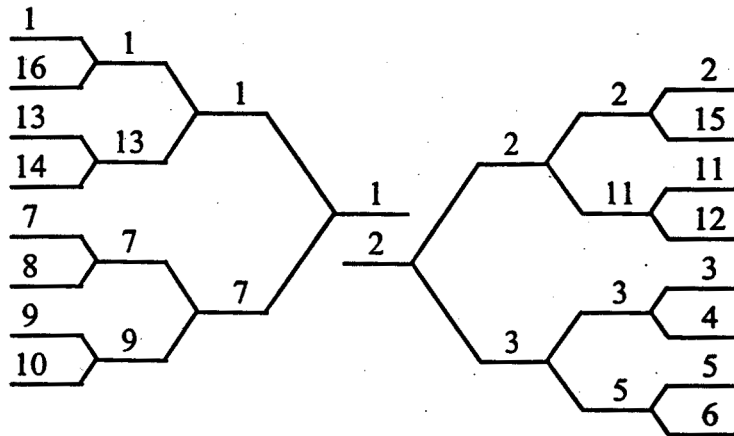


Figure 2: A schedule designed to favor a top two final game.

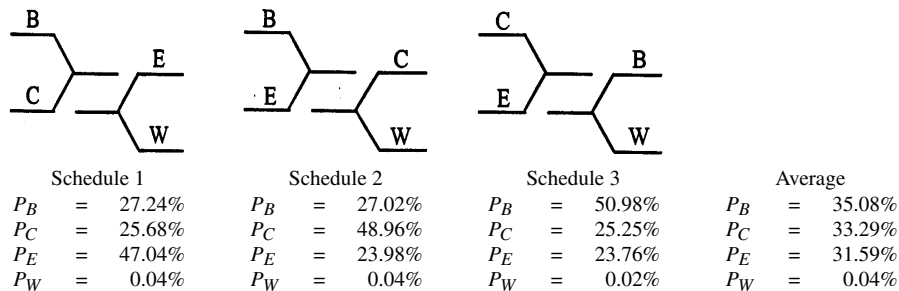


Figure 3: Three possible schedules and the resulting outcomes.

favored over the Eagles by 51%-49%. For Broncos-Eagles it is 52%-48%, and when the top three face the hapless Wolverines it's 99%, 98%, 97%, respectively. For four teams, there are only three ways to schedule the tournament. All three are shown in Figure 3 along with the resulting probabilities for winning the tournament. The fourth column, averaging the first three, is taken as an indication of each team's winning expectation based on its strength. The Bronco's probability of winning the tournament under Schedule 1 is computed as follows: They have a 51% of beating the Chippewas in the opening game. Then 97% of the time they will face the Eagles in the final, and 3% of the time the Wolverines will pull off the upset over the Eagles to reach the final. Thus

$$P_B = .51(.97 \cdot .52 + .03 \cdot .99) = 0.2724.$$

The other eleven entries in Figure 3 are computed similarly. Notice that the three nearly equal teams have vastly different chances of winning the tournament. In Schedule 1, the Eagles benefit from the first round assignment facing the Wolverines. They practically get a free pass into the final game, where they then have nearly a 50% chance. Overall, they win 47.04% of the time. On the other hand, the slightly better Broncos and Chippewas beat up on each other in the opener. Only one can survive to face the expected Eagles, and so they

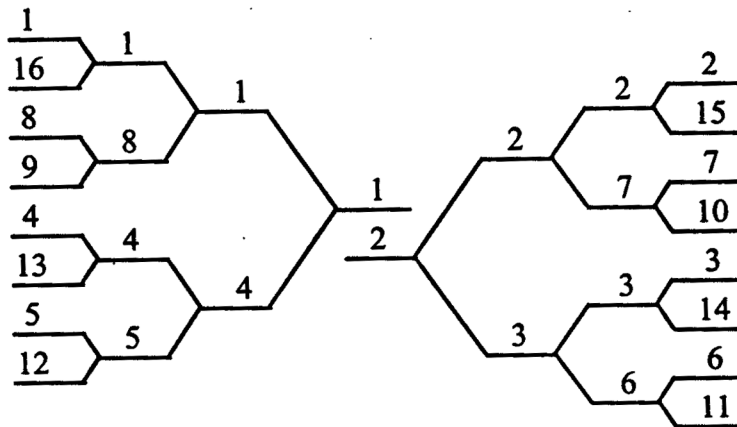


Figure 4: The Standard Method for seeding a tournament with 16 teams.

get probabilities of 27.24% and 25.68%. Clearly the Eagles have received a windfall in the form of the weak opening opponent. Certainly the other teams would envy their position in the schedule. Schedule 1 thus violates Axiom SR as well as Axiom DC. Similarly, Schedule 2 violates SR for the top rated Broncos. Only Schedule 3 fails to violate Axiom SR. But consider the favoritism associated with each schedule. It is 15.45% for Schedule 1, 15.57% for Schedule 2, and 15.90% for Schedule 3. We are disturbed by the excessive benefit that the Broncos receive under Schedule 3. The method proposed in Section 4 reduces this favoritism, but first we examine the shortcomings of the method in standard use.

3. THE STANDARD METHOD AND ITS FLAWS. What is actually done in most seeded tournaments? Let us recall that we assume that the seeding is accepted in the sense that for all $i < j < k$, (nearly) everyone agrees that T_i is better than or equal to T_j , which in turn is better than or equal to T_k . Moreover, we assume that the probability matrix is doubly monotonic. Think of filling in the chart with the anticipated winners, starting in the middle at the championship match. Here we can place T_1 and T_2 arbitrarily. Next we predict that the semifinal opponents will be T_3 and T_4 . If we were to pair T_3 against T_1 and T_4 against T_2 , it would probably violate Axiom SR since T_1 might very much prefer to face T_4 in the semifinals. So we must have T_1 vs. T_4 and T_2 vs. T_3 . Now in the quarterfinal round we need to place teams 5 through 8. Again, seeking to satisfy Axiom SR, we place T_1 vs. T_8 , T_2 vs. T_7 , T_3 vs. T_6 , and T_4 vs. T_5 . The easy way to remember the pairings is to note that $i + j = 9$ in all matches. Similarly, in the round of 16 we have $i + j = 17$. The result is the Standard Method for seeding a tournament shown in Figure 4. If there are more rounds, we continue in this fashion. In the round of size 2^r opponents satisfy $i + j = 2^r + 1$. Filling in the later rounds as we have described does not imply that these pairings will necessarily occur. Upsets are possible. It only indicates the predicted pairings.

This method clearly satisfies Axiom DC. At least superficially, it appears to satisfy Axiom SR since any team scheming to lower its own seed buys only an opponent with a higher seed in the first round of competition. How could that possibly benefit the insincere team? Consider an imaginary tournament with only eight teams having the probability

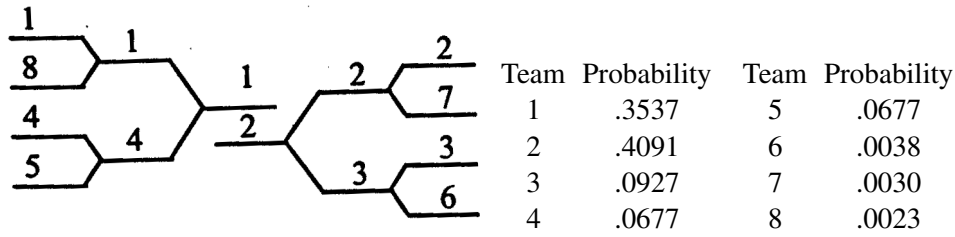


Figure 5: An eight team tournament with probabilities for each team winning.

matrix

$$P = \begin{bmatrix} .5 & .5 & .67 & .67 & .67 & .95 & .95 & .95 \\ .5 & .5 & .67 & .67 & .67 & .95 & .95 & .95 \\ \hline .33 & .33 & .5 & .5 & .5 & .67 & .67 & .67 \\ .33 & .33 & .5 & .5 & .5 & .67 & .67 & .67 \\ .33 & .33 & .5 & .5 & .5 & .67 & .67 & .67 \\ \hline .05 & .05 & .33 & .33 & .33 & .5 & .5 & .5 \\ .05 & .05 & .33 & .33 & .33 & .5 & .5 & .5 \\ .05 & .05 & .33 & .33 & .33 & .5 & .5 & .5 \end{bmatrix}$$

Teams T_1 and T_2 are truly equal. The pundits are unable to break the tie rationally, and the seeding committee does so reluctantly and at random. Likewise the second tier of T_3 , T_4 , and T_5 are also equal teams. Finally teams T_6 , T_7 , and T_8 are equals and also-rans. No one expects them to win. Now let us suppose that whenever a team faces an opponent in its own tier it has a 50-50 chance of winning. When teams from neighboring tiers meet, the higher seed has a $\frac{2}{3}$ chance of surviving. And when a top team faces an also-ran, it wins 95% of the time. The Standard schedule is shown in Figure 5. If you are the coach of one of the top two contenders, where would you like to be positioned, as first or as second seed? Both have patsies as first round opponents, and expect to survive with probability 95%. In the second round they are likely to have mid-tier opponents. Well, not quite. Team 1 is guaranteed a mid-tier opponent, either T_4 or T_5 . Team 2 faces a mid-tier opponent $\frac{2}{3}$ of the time but $\frac{1}{3}$ of the time T_6 upsets T_3 and permits T_2 to face the weaker T_6 . Isn't it preferable to be positioned as T_2 rather than T_1 ? It is a simple matter to write survival equations for each round, where each team receives the sum of probabilities for each possible opponent appearing multiplied by its probability of beating that particular opponent. For example, T_2 survives the first round with probability .95 and then $\frac{2}{3}$ of the time faces T_3 and $\frac{1}{3}$ of the time faces T_6 . Consequently T_2 's likelihood of reaching the final match is

$$.95 \left[\frac{2}{3} \cdot \frac{2}{3} + \frac{1}{3} \cdot .95 \right] \cong .7231.$$

By comparison, T_1 survives with probability

$$.95 \left[\frac{1}{2} \cdot \frac{2}{3} + \frac{1}{2} \cdot \frac{2}{3} \right] \cong .6333.$$

Since T_2 has nearly a 9% advantage of reaching the final game, this translates into a 5.54% higher probability of winning the tournament. Why isn't it just half of 9%? On those

Team	Standard Method	Randomized Method
1	35.37%	36.93%
2	40.91%	36.93%
3	9.27%	8.99%
4	6.77%	8.99%
5	6.77%	6.97%
6	0.38%	0.40%
7	0.30%	0.40%
8	0.23%	0.40%

Table 1: Limited randomization reduces favoritism

occasions when T_1 fails to make the final, T_2 usually inherits a $\frac{2}{3}$ chance of becoming champion. Thus, the advantage swells beyond half of 9%. For the record, the probabilities for all eight teams are shown in Figure 5. We have demonstrated that the standard method can violate Axiom SR by a good 5.5%.

The problem with the standard schedule is that T_1 is assured a moderately tough opponent in the second round while T_2 benefits from a significantly possible upset by T_6 over T_3 . Might not randomization be used to counterbalance the schedule’s advantage favoring T_2 ? Let us randomize as follows. First T_1 and T_2 are entered in opposite brackets. Next T_3 and T_4 are randomized in the two semifinal positions. Finally, the remaining four teams are placed at random. We average the winning probabilities over all $2!4! = 48$ randomized schedules. In this particular instance, symmetries between T_3 and T_4 and among T_6 , T_7 , and T_8 allow us to escape with only four schedules to compute, specifically the result of moving T_5 to each possible slot. Table 1 shows the improved fairness in winning probabilities for the various teams. Notice how the insincerity provoking favoritism of T_2 has been eliminated while the overall balance within each tier has improved, perfectly in the top and bottom tier, and imperfectly in the middle tier. Here DC places T_3 and T_4 in the preferred positions, but the lot of T_5 has improved regardless, narrowing the gap with its peers.

4. COHORT RANDOMIZED SEEDING. Let us now formalize the method we used in Section 3, and which we contend meets the three axioms of fair seeding in a tournament of size $n = 2^k$. We call this *cohort randomized seeding*. We assign seeds from T_1 down to T_n . The first two teams comprise the first *cohort*, $C_1 = \{T_1, T_2\}$, the second cohort is $C_2 = \{T_3, T_4\}$, the third cohort is $C_3 = \{T_5, T_6, T_7, T_8\}$, and from here on the cohorts continue to double in size. In general, we have k cohorts with $C_i = \{T_{2^{i-1}+1}, \dots, T_{2^i}\}$ for $2 \leq i \leq k$. In Figure 6 the seeding positions of teams in cohorts C_2 , C_3 , and C_4 are shown as Roman numerals II, III, and IV. When the schedule is created, Teams T_1 and T_2 are placed in opposite brackets. Next cohort C_2 is placed at random in the two II slots, cohort C_3 is randomized among the four III slots, and so on.

Recall our mythical four team tournament of Figure 3. Cohort Randomization directs us to randomize the placement of the second cohort, the Eagles and the Wolverines. Thus, we randomize between Schedules 2 and 3. The effect is to rein in the windfall to the Broncos under the standard schedule and give more nearly equal probabilities of $P_B = 39.00\%$,

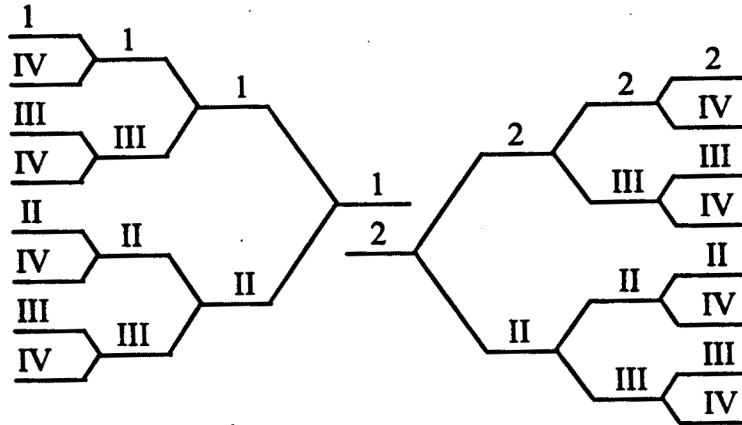


Figure 6: Cohort Randomized Seeding for 16 teams.

$P_C = 37.10\%$, $P_E = 23.87\%$, $P_W = 0.03\%$. Cohort Randomization has improved the fairness considerably by reducing favoritism of the Broncos from 15.90% to 3.92%. Of course, total randomization would balance the probabilities perfectly, but at the cost of abandoning Delayed Confrontation.

These two cases show that Cohort Randomization can improve fairness, but can some other method ever give better results? We now show that the answer is “No!” by verifying that Cohort Randomization satisfies all three seeding axioms. Thus, no other method can be “more fair”. First consider Axiom DC. Up to and including the round of size 2^{j+1} , the top 2^j seeds in cohorts C_1 through C_j have been positioned to avoid one another, so we have achieved Axiom DC. Next, we need to confirm Axiom SR. If any higher seed ever has a tougher schedule than another lower seed, there must be some consecutive pair T_i and T_{i+1} in which T_{i+1} is favored over T_i . Now if T_i and T_{i+1} are in the same cohort, no favoritism is possible because they are placed at random in precisely the same set of positions. So what happens if $i = 2^j$ and our teams are in consecutive cohorts C_j and C_{j+1} ? Prior to the round of size 2^{j+1} these teams are treated identically, drawing randomly chosen opponents from lower cohorts. Similarly, after the round of size 2^{j+1} they are again treated identically, drawing random opponents from higher cohorts. The only time they are treated differently is in the round of 2^{j+1} , when T_i gets an opponent from C_{j+1} or lower while T_{i+1} is likely to get one from C_j or higher. Clearly the schedule rewards T_i , so Axiom SR is satisfied.

The most difficult axiom to verify is Axiom FM. Delayed confrontation requires that teams be assigned a position from among the slots reserved for its own cohort. However, within those cohort positions there may be alternatives to the uniform randomization we have chosen. Indeed, the standard method uses a specified assignment of $i + j = 2^r + 1$, and we have seen examples to show that this can violate Axioms SR and FM. Conceivably a different assignment, or using some nonuniform probability distribution within the cohort, could lead to improved fairness. We need to eliminate these possibilities. We start with teams T_1 and T_2 placed by fiat. Next we need to position T_3 and T_4 . Imagine that T_3 is the equal of teams T_1 and T_2 , and that all lower teams have no chance of upsetting any of these three. Suppose T_3 is given the semifinal position against T_1 with probability p and it faces

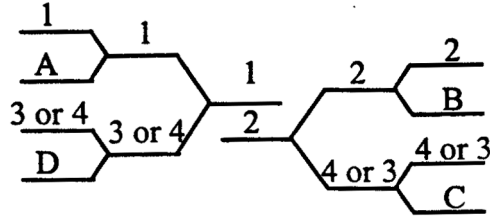


Figure 7: Schedule for analysis of placement of team T_5 .

T_2 with probability $1 - p$. Now if $p > \frac{1}{2}$ we violate SR since T_1 would prefer T_2 's seeding. But if $p < \frac{1}{2}$ we compute T_1 's winning chances as $1/2 - p/4$, while T_2 gets $1/4 + p/4$ and T_3 gets $\frac{1}{4}$. To minimize favoritism we need to take $p = \frac{1}{2}$.

Now consider the placement of T_5 into the positions labeled A, B, C, and D in Figure 7. Assume the field has five equally strong teams and three also-rans. Suppose the largest probability for any of the four open slots for T_5 is $p \geq \frac{1}{4}$; we need to show $p = \frac{1}{4}$. If the A slot has $p > \frac{1}{4}$, we violate SR for T_1 since T_1 would prefer whichever slot opposes T_5 with $p < \frac{1}{4}$. Next, if the B slot has $p > \frac{1}{4}$, team T_2 would prefer the seed of either T_3 or T_4 , violating SR. Now if C has $p > \frac{1}{4}$, then imagine a field in which T_5 is a notch below the top four teams, but could upset any of them. Now team T_1 would like to exchange with T_2 so that T_1 might benefit from an upset advancing T_5 giving T_1 the softer second round opponent, namely T_5 . Finally, if D has $p > \frac{1}{4}$, consider T_1 's view when T_5 is the equal of T_3 and T_4 , but the three other teams are a step back, but capable of upsets. Again swapping with T_2 gains the hope of a potential upset yielding a softer second round opponent, this time one of the three bottom teams. After discarding all other variations, we are forced to place T_5 in each slot with equal probability $p = \frac{1}{4}$.

And so it goes. One after another we determine that each member of cohort C_3 must be distributed equally to the four slots assigned for C_3 . And then we repeat the argument on C_4 forcing equal probabilities of $\frac{1}{8}$. Thus Cohort Randomized Seeding not only satisfies Axiom FM, but it is the only method that does so while preserving Axioms DC and SR.

5. IMPLEMENTATION. How might cohort randomization be implemented in, say, the NCAA basketball tournament? I see only advantages. The seeding committee need not determine specific seed numbers for every team. All they need is the cohort numbers, C_1 through C_6 . This means that they select two cofavorites in C_1 and two other "semi-favorites" in C_2 . If the organizers wish, they can merge these two classes and have "four regional top seeds" with little impact on the method. Next, four teams are chosen for cohort C_3 , and one is randomly assigned to each regional. This is followed by eight teams in C_4 , two assigned to slots in each region. Next we have 16 teams in C_5 and finally 32 teams in C_6 . I would think the seeding committee would be happy to place these 32 in a single category and not have to answer questions about why a particular team was seeded 13th in its region and not 11th as its fans think it deserved. A lot of minor quibbles of precise seeding are thus avoided. However I admit it now becomes quite significant whether my favorite team is rated 16th overall (a C_4 seed) or 17th (a C_5 seed). The announcement of cohort assignments followed immediately by bouncing balls being selected to create the

precise schedule would make an hour-long television show of great interest, particularly if the official announcements are both preceded by speculation and followed by reactions.

Finally, it is now the practice to make a strong effort to distribute several teams from the same conference about the country in different regionals. This could easily be accommodated by stating that when a second team from a particular conference is about to be placed, two of the regions are forbidden to it so that it can meet its conference compatriot only in the final game. When placing the third team from this conference, we forbid two regions, the two already occupied by the conference; similarly the fourth goes into the only region still lacking representation from this conference. With the fifth team we can return to any region, but forbid slots that would face a conference foe prior to the regional final. The sixth goes to one of the other three regions with the same regional final restriction, and so on. And so we get a schedule that is simultaneously more fair and also more interesting to the fans, a win-win situation.

Various sports commentators have tried to interpret cohort randomization by asking “Whom does it favor? Whose ox gets gored?” There is considerable suspicion that there is a hidden agenda to favor a certain class of teams. Since the high seeds are happy with the current setup, they suspect that this is intended to improve things for the lower seeds, and who needs that? Be aware that many doubly monotonic matrices are possible. In the four team tournament analyzed in Figure 3, team T_2 benefits from cohort randomization when T_1 's excessive favoritism is reined in. But in the eight team example in Figure 5, team T_1 would want cohort randomization to reward sincerity relative to T_2 . So for some matrices cohort randomization benefits T_1 at the expense of T_2 and for others it benefits T_2 at the expense of T_1 . The method doesn't inherently favor either team. Instead, it tends to create better balance in all cases. If we exercise our judgment prior to estimating the probability matrix P , there is no reason for either team to oppose cohort randomization. Similarly, if we carefully examine the lower seeds, I suspect that teams 13 through 16 would usually welcome cohort randomization, while 9 through 12 might well oppose it. Similarly, teams 7 and 8 might favor it while 5 and 6 oppose it. Roughly, the lower seeds in each cohort usually benefit from randomization while the higher ones suffer. The key is to adopt the method for a future season before anyone knows his own position. This way organizers can exercise judgement without having self interest cloud the decision making process.

ALLEN SCHWENK won the E. T. Bell Undergraduate Research Prize at Caltech. He earned his Ph.D. from the University of Michigan. Before accepting a professorship at Western Michigan University in 1985, he taught for nine years at the U.S. Naval Academy and spent single years at Michigan State and at the University of Waterloo in Ontario. He remains addicted to mathematical games and puzzles. He used to give frequent lectures on Rubik's cube, but has replaced it with the Lights Out Cube and nontransitive dice. This analysis of seeding was motivated by his interest in duplicate bridge as well as NCAA basketball.

Western Michigan University, Kalamazoo, MI 49008-5152
schwenk@wmich.edu