

**Manuscript version: Submitted Version**

The version presented here is the submitted version that may later be published elsewhere.

**Persistent WRAP URL:**

<http://wrap.warwick.ac.uk/167286>

**How to cite:**

Please refer to the repository item page, detailed above, for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

**Copyright and reuse:**

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

**Publisher's statement:**

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: [wrap@warwick.ac.uk](mailto:wrap@warwick.ac.uk)

# Football Group Draw Probabilities and Corrections

by

Gareth O. Roberts<sup>1</sup> and Jeffrey S. Rosenthal<sup>2</sup>

(Version of: May 11, 2022)

## 1. Introduction.

Major football (soccer) tournaments such as the FIFA World Cup, European Championships, and UEFA Champions League hold public draws to decide who plays whom. It is customary to employ one or more celebrities to draw balls from pots to sequentially construct the draw, to add excitement and increase interest in the competition. However, such mechanisms can affect the draw probabilities in unexpected ways.

This article will focus on the group draw for the FIFA World Cup, although similar ideas could be applied to other competitions. This draw has various restrictions (based on geographical constraints) about which assignments are permissible, leading to a complicated space of potential draws. In addition, the draw should take place sequentially, to allow for public interest and transparency. The statistical challenge, then, is to simulate from the uniform distribution on a non-symmetric high-dimensional space in a way which is also sequential and entertaining. We shall present several potential solutions that we have developed to address this challenge; they are available for interactive use [18], and have been reported on in the media [20, 13].

---

<sup>1</sup>Department of Statistics, University of Warwick, CV4 7AL, Coventry, U.K. Email: [Gareth.O.Roberts@warwick.ac.uk](mailto:Gareth.O.Roberts@warwick.ac.uk).

<sup>2</sup>Department of Statistics, University of Toronto, Toronto, Ontario, Canada M5S 3G3. Email: [jeff@math.toronto.edu](mailto:jeff@math.toronto.edu). Web: <http://probability.ca/jeff/>

## 1.1. The 2022 FIFA World Cup group draw.

The 2022 FIFA World Cup will take place in November/December 2022, in Qatar. It will involve 32 national teams – 31 who qualified through competition, plus the host team Qatar who qualify automatically. These 32 teams needed to be partitioned into 8 groups of 4 teams each, who will all play each other in the first round of the Cup.

The group assignments were subject to various restrictions, as follows. The 32 teams were first partitioned into 4 seeded pots, with Pot 1 consisting of the hosts Qatar plus the seven most highly ranked teams (according to the official FIFA national team rankings), Pot 2 consisting of the next 8 highest ranked teams, and so on. In addition, each team is affiliated to one of the 6 continental federation regions: UEFA (Europe; 13 teams; henceforth **Eu**), CONMEBOL (South America; 4 or 5 teams; henceforth **SA**), CONCACAF (North and Central America; 3 or 4 teams; henceforth **NA**), AFC (Asia; 5 or 6 teams; henceforth **As**), CAF (Africa; 5 teams; henceforth **Af**), and OFC (Oceania; 0 or 1 team; henceforth **Oc**). The assigned pots were as follows:

**Pot 1:** Qatar[**As**], Belgium[**Eu**], Brazil[**SA**], France[**Eu**], Argentina[**SA**], England[**Eu**], Portugal[**Eu**], Spain[**Eu**].

**Pot 2:** Denmark[**Eu**], Netherlands[**Eu**], Germany[**Eu**], Switzerland[**Eu**], Croatia[**Eu**], Mexico[**NA**], USA[**NA**], Uruguay[**SA**].

**Pot 3:** Iran[**As**], Serbia[**Eu**], Japan[**As**], Senegal[**Af**], Tunisia[**Af**], Poland[**Eu**], KoreaRep[**As**], Morocco[**Af**].

**Pot 4:** Wales/Scot/Ukr[**Eu**], Peru/UAE/Au[**SA,As**], CostaRica/NZ[**NA,Oc**], SaudiArabia[**As**], Cameroon[**Af**], Ecuador[**SA**], Canada[**NA**], Ghana[**Af**].

(The reason for the uncertainty in three of the team names in Pot 4 is that, due to delays caused by Covid-19 and the war in Ukraine, not all teams had been finalised by the time of the draw, so placeholders were used. Two of the placeholders corresponded to two different potential regions, so they had to satisfy the geographical constraints for *both* of the

corresponding regions.)

In terms of these specifications, the restrictions on group formations were that each group needed to include one team from each of the 4 pots, and furthermore include either 1 or 2 teams from **Eu** plus either 0 or 1 teams from each of the other regions.

The FIFA group draw, on 1 April 2022, then proceeded as follows [9]. First, the host team Qatar was automatically placed in Group A. Then, the remaining teams from Pot 1 were selected one at a time, uniformly at random, and each placed into the next group from B through H. Then, the teams from Pot 2 were selected one at a time, uniformly at random, and assigned to the next available group, i.e. the first group which would not cause a conflict with any of the geographical restrictions (either immediately in the group where they were placed, or subsequently by making it impossible to validly fill in the remainder of the draw). This procedure was then repeated with Pot 3, and then with Pot 4 to conclude the draw. Each random selection was performed by a celebrity footballer, who picked a ball uniformly at random from a round bowl, and opened it to reveal the chosen team. (Each team was also randomly assigned a “position” within their group, to determine the order in which the matches would be played, but we do not consider that issue here.)

Without the geographical constraints, just drawing the teams in random order from the four pots in sequence, there would be  $7! \times (8!)^3 \doteq 3.3 \times 10^{17}$  possible draws that could be produced. We shall see below that about 1 in 560 of these unconstrained draws satisfy the geographical constraints. Hence, the number of valid draws is approximately  $5.9 \times 10^{14}$ . A *uniform* draw is one for which all  $5.9 \times 10^{14}$  valid draws have an equal chance of materialising. This article will explore the non-uniformity in the FIFA 2022 draw procedure, and also propose various methods of refining the draw to achieve complete uniformity.

## 1.2. Previous literature.

It has long been known that sequential draws such as those adopted by FIFA and UEFA led to non-uniform probabilities, see e.g. [14, 4, 16]. Various papers have looked at different

mechanisms for carrying out different sorts of sequential draws; see e.g. [4]. Much of the literature has focused on how to obtain a *balanced* draw rather than a uniform one, i.e. trying to make each group roughly equal in strength [10, 11]. (Indeed, balance is what inspired FIFA to create the seeded pots based on world rankings rather than continental affiliation, for both the 2018 and 2022 World Cup group draws. But that is separate from the question of uniformity.)

Various papers have proposed solutions to the non-uniformity. Some produce sequential procedures which are *closer* to uniform than existing methods [2], but are still not completely uniform. Others propose completely new draw mechanisms, e.g. [12]. The paper [16] even briefly postulates a possible Markov chain procedure (specifically an exclusive Gibbs sampler) for the UEFA Champions League draw. However, none of these methods respect the desired sequential nature of the draw while also achieving perfect uniformity, which is our goal here.

## 2. Comparing Uniform and FIFA Probabilities.

Before proposing alternative solutions, we investigate the extent to which the FIFA Sequential Algorithm (described in Section 1.1 above) is non-uniform.

A preliminary look illustrates the nature of the problem. For example, consider the question of whether the USA is assigned to Group A with Qatar. (This is an important question, since Qatar is weaker than the other teams in Pot 1, so Group A is the most desirable placement.) Under the FIFA method, any of the 8 teams in Pot 2 is equally likely to be selected first, and none of them have any regional conflict with Qatar, so the USA has probability exactly  $1/8$  or 12.5% of being placed in Group A. But under the uniform distribution, the USA should be *less* likely to land in Group A, because that leaves fewer ways for the numerous **Eu** teams to be placed in the other groups. (Indeed, we shall see below that the uniform distribution gives probability 9.0% to the USA being in Group A, which is significantly smaller.)

To compute the differences between the two distributions, we shall use a Monte Carlo approach. Thus, we need to repeatedly simulate from both the uniform distribution  $\mathbf{U}$ , and the distribution  $\mathbf{P}$  of draws created by the FIFA sequential method. We now consider each of these problems in turn, in Sections 2.1 and 2.2 below.

## 2.1. Uniform Simulation by Rejection Sampling.

To obtain draws with equal probabilities, the simplest way is to use a Rejection Sampler. This algorithm proceeds as follows:

1. First, assign a draw completely randomly, without regard to the geographical constraints. That is, select uniformly at random from the  $7! \times (8!)^3 \doteq 3.3 \times 10^{17}$  possible unconstrained draws discussed in Section 1.1 (which is straightforward).
2. Then, compute the number of teams from each region in each group, and check if the geographic constraints are satisfied (i.e., 1 or 2  $\mathbf{Eu}$  teams plus 0 or 1 from each other region in each group).
3. If the constraints are satisfied, then output the chosen draw.
4. If the constraints are *not* satisfied, then reject the chosen draw, and start again at step 1 above with a fresh unconstrained draw.

It is well-known (e.g. [5, Section II.3]) that this Rejection Sampler produces a draw which is uniform over all possible valid draws, i.e. which is distributed as the uniform distribution  $\mathbf{U}$ . Then, by sampling many times and averaging the results, we can get a good estimate (by the Law of Large Numbers) of the expected value according to  $\mathbf{U}$  of any functional, or the probability of any particular event [19].

In addition, our simulations [19] show that about 1 in 560 of the proposed draws were accepted, i.e. satisfied the geographical constraints. This means that the total number of valid draws is approximately  $7! \times (8!)^3 / 560 \doteq 5.9 \times 10^{14}$ .

## 2.2. Simulating the FIFA Sequential Draw.

Although the FIFA Sequential Draw method is easy to describe (see Section 1.1), and can usually be implemented without difficulty, it is surprisingly challenging to simulate with a computer program. This is because potential group assignments have to be skipped whenever they would conflict with the geographical constraints.

Now, immediate conflicts are simple to detect. Indeed, if adding a team to a group would create a third **Eu** team, or a second team from any other region, then obviously that group must be skipped. However, subsequent conflicts are much more complicated. It may be that assigning a certain team to a certain group would then make it impossible to fill in the rest of the draw in a valid way, either due to being forced to later add a third **Eu** team or second team from another region to a group, or failing to add at least one **Eu** team to a group. These subsequent conflicts are usually relatively simple to resolve (and indeed, they did not occur at all in the actual FIFA 2022 draw), but this is not guaranteed.

For an extreme example, suppose a draw in progress has teams from the various regions as follows:

<b>Group:</b>	A	B	C	D	E	F	G	H	To Go
<b>Pot 1</b>	<b>Af</b>	<b>NA</b>	<b>NA</b>	<b>NA</b>	<b>NA</b>	<b>NA</b>	<b>NA</b>	<b>NA</b>	
<b>Pot 2</b>	<b>As?</b>								<b>SA,SA,SA,SA,SA,SA,SA</b>
<b>Pot 3</b>									<b>Eu,Eu,Eu,Eu,Eu,Eu,NA,SA</b>
<b>Pot 4</b>									<b>Oc,Oc,Oc,Oc,Oc,Oc,Oc,Oc</b>

Then placing the **As** team from Pot 2 into Group A already creates a subsequent conflict, since the seven **SA** teams from Pot 2 must then be placed in Groups B through H, and then in Pot 3 the **NA** and **SA** teams cannot both be placed in valid groups without violating their respective geographic constraints. (Instead, the **As** team in Pot 2 must be placed in some other group besides Group A, thus leaving e.g. Group A for the **NA** team from Pot 3, and the other group for the **SA** team from Pot 3.) However, this subsequent conflict might not become apparent until the very end of the Pot 3 placements.

This example illustrates that a general program to simulate the FIFA method must be

robust enough to detect subsequent conflicts many steps later. Our approach [19] was to program this recursively. We created a subroutine “`placerest()`” which, given a partial draw, randomly selects a team to place next, and then attempts to place that team in the next available slot. It then recursively calls itself with that one additional placement, to see if the rest of the draw can then be filled in successfully. If it can, then the draw is complete. If it cannot, then it instead attempts to place that team in the subsequent available slot. Continuing in this way, it eventually successfully places every team. The recursive nature of the program ensures that any subsequent conflicts will be dealt with, no matter how much further along they occur; for details see the computer program at [19].

Thus, this algorithm produces a draw with probabilities as in the current FIFA sequential method, i.e. which is distributed as the FIFA distribution  $\mathbf{P}$ . Then, by sampling many times and averaging the results, we can get a good estimate (by the Law of Large Numbers) of the expected value according to  $\mathbf{P}$  [19] of any functional, or the probability of any event.

### 2.3. Specific Probability Comparisons.

Armed with our computer program [19], we can now compare the probabilities of various events according to the uniform distribution  $\mathbf{U}$  and the FIFA Sequential Method  $\mathbf{P}$ .

For example, the probability that England is in the same group as Germany should be 10.6% under  $\mathbf{U}$ , but increases to 11.8% under  $\mathbf{P}$ . The probability that Germany is in the same group as Qatar should be 13.6% under  $\mathbf{U}$ , but decreases to 12.5% under  $\mathbf{P}$ . The probability that Canada is with Qatar should be 15.4% under  $\mathbf{U}$ , but increases to 16.5% under  $\mathbf{P}$ . And, the probability that USA is with Qatar should be just 9.0% under  $\mathbf{U}$ , but increases to 12.5% under  $\mathbf{P}$  (a relative increase of 39%). Probabilities of other events can similarly be computed [19].

These probability differences are not huge, but they are large enough to illustrate that the FIFA method is significantly different from a uniform draw, and could lead to significantly different group assignments. We now consider how to fix this problem, to achieve uniform



group draws while still preserving the interest and excitement of the FIFA method.

### 3. A Motivating Example.

To see more clearly why FIFA's sequential procedure (described in Section 1.1 above) fails to achieve uniformity, and how we might correct that, consider a simplified set up with only 6 teams to be allocated to 3 groups (of 2 teams each). Suppose there are two seeded pots (a subset of the actual World Cup 2022 pots):

**Pot 1:** Qatar (Q) [**Af**], France (F) [**Eu**], Brazil (B) [**SA**]

**Pot 2:** Mexico (M) [**NA**], Switzerland (S) [**Eu**], Uruguay (U) [**SA**]

Qatar, as hosts, are automatically assigned to Group A. Without loss of generality, assume that France is placed in Group B, and Brazil in Group C. Then the two **SA** teams, Brazil in Pot 1 and Uruguay in Pot 2, must be kept apart, so Uruguay cannot be placed in Group C. There are thus four possible draws for assigning the three different groups:

$D_1$ : QM, FU, BS.

$D_2$ : QS, FU, BM.

$D_3$ : QU, FM, BS.

$D_4$ : QU, FS, BM.

Let  $\mathbf{P}$  be the probability measure resulting from FIFA's sequential procedure described above (adapted to this simplified setting), and let  $\mathbf{U}$  be the probability measure which assigns equal probability to each valid draw. Thus  $\mathbf{U}(D_i) = 1/4$  for  $i = 1, 2, 3, 4$ . However, a simple calculation gives that  $\mathbf{P}(D_1) = \mathbf{P}(D_2) = 1/3$  while  $\mathbf{P}(D_3) = \mathbf{P}(D_4) = 1/6$ . In particular, if  $QU$  is the event that Qatar is paired with Uruguay, then  $\mathbf{P}(QU) = 1/3$  whereas  $\mathbf{U}(QU) = 1/2$ . This clearly illustrates the potential non-uniformity arising from FIFA's sequential procedure.

### 3.1. In search of debiasing: random order sequential procedures.

It is natural to ask how to fix the biases of the above example while retaining the sequential nature of the draw. Perhaps the most obvious try is to randomise the order in which opponents for the Pot 1 teams are found. In the usual FIFA method, the first drawn team in Pot 2 is first attempted to be placed in Group A with Qatar. To emphasise that we start with Qatar, write  $\mathbf{P}_Q$  for the probabilities induced by this method (so  $\mathbf{P}_Q = \mathbf{P}$  from above). If instead the first drawn team in Pot 2 is attempted to be placed in Group B with France, then this leads to different probabilities  $\mathbf{P}_F$ , and similarly  $\mathbf{P}_B$  if the first team is attempted to be placed in Group C with Brazil. Alternatively, if the first drawn team in Pot 2 is attempted to be placed in Group A or B or C with probability  $1/3$  each, then such a procedure could be called a *random order sequential draw*, with resulting probabilities given by  $\mathbf{P}_{rand} = \frac{1}{3}(\mathbf{P}_Q + \mathbf{P}_F + \mathbf{P}_B)$ . (In principle, we could also specify that the *second* team drawn from Pot 2 should have their attempted group randomised among the two remaining teams, but this example is sufficiently simple that we can ignore that.)

In this example, it can be checked that  $\mathbf{P}_Q$  and  $\mathbf{P}_F$  are both non-uniform, but  $\mathbf{P}_B$  happens to be uniform, i.e.  $\mathbf{P}_B = \mathbf{U}$ . Moreover, a simple calculation yields that

$$\mathbf{P}_{rand} := \frac{1}{3}(\mathbf{P}_Q + \mathbf{P}_F + \mathbf{P}_B) = \mathbf{U},$$

i.e. the random order sequential draw is uniform in this case. This leads to the question of whether random order sequential draws are always uniform.

To answer this question, consider a modification of our simple example. Suppose that we also forbid pairings between two **Eu** teams, thus disallowing  $D_4$  (which pairs France with Switzerland). In that case,

$$\mathbf{U}(D_1) = \mathbf{U}(D_2) = \mathbf{U}(D_3) = 1/3.$$

For this modified example,  $\mathbf{P} = \mathbf{P}_Q$  is easily seen to be non-uniform, but  $\mathbf{P}_F$  and  $\mathbf{P}_B$  both turn out to be uniform. Moreover, for the random order sequential draw  $\mathbf{P}_{rand} :=$

$\frac{1}{3}(\mathbf{P}_Q + \mathbf{P}_F + \mathbf{P}_B)$ , we calculate that

$$\mathbf{P}_{rand}(D_1) = 5/18; \quad \mathbf{P}_{rand}(D_2) = 13/36; \quad \mathbf{P}_{rand}(D_3) = 13/36 .$$

So, in this case, FIFA’s fixed order draw  $P_Q$  is uniform, while the randomised order draw  $P_{rand}$  is non-uniform. This shows that randomised order draws do not solve the non-uniformity problem. In fact, they can sometimes make the non-uniformity worse, or even introduce non-uniformity in case where the original FIFA algorithm happened to be uniform.

### 3.2. In search of debiasing: multiple ball procedures.

We now return to the original example, where the only constraint is on the two **SA** teams. Then we see from the above that two of the possible draws ( $D_3$  and  $D_4$ ) put Uruguay in Group A with Qatar, while only one ( $D_1$ ) puts Mexico in Group A, and one ( $D_2$ ) puts Switzerland in Group A. This suggests that, to achieve the uniform distribution  $\mathbf{U}$ , when selecting the team from Pot 2 to put in Group A, we could use a bowl with *two* Uruguay balls, and only one from each of Mexico and Switzerland. The next team in the draw is then selected uniformly at random from those four balls, just as before. This simple “multiple balls solution” thus achieves the correct conditional probability (in terms of  $\mathbf{U}$ ) for the team from Pot 2 to be placed in Group A.

More formally, this solution could be described as follows. When selecting the team from Pot 2 to put in Group A, we first count, for each team in Pot 2, the number of valid draws which have that team in that position (while keeping all of the previously-selected teams in their previously-selected positions, too). In the above example, we have  $n_U = 2$  for Uruguay, and  $n_M = 1$  for Mexico, and  $n_S = 1$  for Switzerland. We then place  $n_U$  balls for Uruguay,  $n_M$  balls for Mexico, and  $n_S$  balls for Switzerland, all into a bowl, and then sample one of the balls uniformly at random. In this way, we select Uruguay with probability  $n_U/(n_U + n_M + n_S)$ , and so on. This ensures the correct conditional probability for the next position. Hence, if this procedure is repeated for each new position, then it ensures the

correct full uniform probability  $\mathbf{U}$  for the entire draw.

Can this procedure be extended to larger draws? In principle, yes. However, the number of balls would soon get completely out of hand. For instance in the World Cup 2022 draw, when choosing the Pot 2 team to play Qatar, we would need to put about  $(8!)^3/560 \doteq 10^{11}$  balls into the bowl, clearly impossible in practice. Nevertheless, we shall see in Section 5 below that we can exhibit a practical and completely uniform multiple balls solution to this problem, using far smaller total numbers of balls.

## 4. A Metropolis Algorithm Solution.

The Rejection Sampler algorithm of Section 2.1 provides a perfectly uniform group draw distributed according to  $\mathbf{U}$ , so in some sense it completely solves the problem. However, use of this algorithm would require the public to “trust” the computer to sample correctly, and would not provide any drama or entertainment value during the draw. So, we next consider ways to achieve a more interesting and entertaining and transparent draw while still preserving the uniform probabilities  $\mathbf{U}$ .

One solution is as follows. Begin with a Rejection Sampler uniform sample as above. Then, repeatedly choose two teams at random from the same pot, and “swap” their group assignments provided that swap does not violate any of the geographical constraints. (If the swap would violate any constraint, then the group assignments are left unchanged.) The validity of this solution follows from:

**Proposition 1** *If we begin with a valid draw chosen from  $\mathbf{U}$  (e.g. using a Rejection Sampler as in Section 2.1), and then repeatedly perform a fixed number of swap moves as described above, then the distribution of the draw remains equal to the uniform distribution  $\mathbf{U}$ .*

**Proof.** First note that the proposed swaps are symmetric, since making the same swap twice is equivalent to not changing at all. Furthermore, they are accepted if the swapped draw is still valid, otherwise they are rejected. It follows that the swap moves correspond

to a Metropolis algorithm [17] with stationary distribution  $\mathbf{U}$ . Hence,  $\mathbf{U}$  is a stationary distribution for the Markov chain corresponding to this method, i.e. this method induces a Markov chain Monte Carlo (MCMC) algorithm [3] which preserves the distribution  $\mathbf{U}$  as it runs. Hence, the overall distribution of the group assignment remains uniform no matter how many swap moves are attempted. ■

The swap moves used by this method can easily be performed manually, by choosing pairs of teams by drawing balls from bowls, and checking directly if any geographic constraints would be violated by swapping them. Unlike the FIFA Sequential Draw of Section 2.2, there are no subsequent conflicts, just immediate conflicts, so they can be checked very easily.

As more swaps are performed, the overall group assignment continues to change randomly, in unexpected and entertaining ways which could make for an exciting spectacle. If desired, a large number of initial swaps could be performed quickly by a computer. Then, a certain fixed number of final swaps could be performed manually, by physically selecting balls from urns to determine which two teams are selected next for possible swap. The final assignment would then be whatever configuration remains after the final manual swap has been performed.

An interactive simulation of this method is available at [18].

## 5. A Multiple-Balls Solution.

We now investigate how to generalise the multiple-ball method introduced in Section 3.1 to a full World Cup draw.

This method fills in the groups one team at a time. At each step, the computer generates a collection of balls corresponding to all the teams who could potentially occupy the next spot. One of those balls is then chosen uniformly at random, and that team is placed in the next spot. Once all spots are filled, it provides a complete group draw.

Recall that, for this method to generate the uniform distribution  $\mathbf{U}$ , the next team needs

to be selected from its correct conditional probability according to  $\mathbf{U}$ . That is, given a partial draw, if  $n_j$  is the number of ways of completing the draw which put team  $j$  in the next position, then we should select team  $j$  with probability proportional to  $n_j$ .

This approach immediately presents several challenges. How could we compute the  $n_j$  values, or at least the corresponding probabilities  $p_j = n_j / \sum_i n_i$ ? And even if we knew the  $p_j$ , how could we sample with probabilities proportional to them, in a practical way without needing a massive collection of  $\sum_i n_i$  different balls? We first consider the second problem, of sampling with probabilities  $p_j$  (Section 5.1). We then explain how it suffices to use *estimates* of the  $p_j$ , so that they do not need to be computed analytically (Section 5.2). This leads to a practical algorithm for conducting the draws (Section 5.3).

## 5.1. Discrete Random Rational Simulation.

Consider the problem of simulating a discrete random event from  $J$  possible values with given non-negative probabilities  $p_1, \dots, p_J$  summing to 1. Call this distribution  $\mathbf{D}$ . We wish to sample from  $\mathbf{D}$  using a *rational simulation algorithm*, i.e. select some small non-negative integer number  $m_j$  of balls representing each value  $j$ , and then draw each ball with probability  $1/M$  where  $M = \sum_{j=1}^J m_j$ .

Now, if the  $p_j$  were small integer multiples of each other, then this would be easy. For example, if  $p_1 = 1/2$ ,  $p_2 = 1/3$ , and  $p_3 = 1/6$ , then we could choose  $m_1 = 3$ ,  $m_2 = 2$ , and  $m_3 = 1$ , and then drawing uniformly at random from the  $3 + 2 + 1 = 6$  balls would accomplish our task. However, we do not wish to assume that the  $p_j$  are rationally related, and we want the total number of balls  $M$  to remain moderate even if the ratios of the  $p_j$  have no simple fractional form. We therefore propose the following algorithm. (Our solution uses the stratified sampling strategy common in Sequential Monte Carlo [15, 8]. Other resampling strategies could also be used; see [6] for a discussion of options.)

1. Let  $M = \lceil \max\{(1/p_j) : p_j > 0\} \rceil$  be the ceiling of the reciprocals of the non-zero  $p_j$ .

2. Set  $r_j = M p_j$ , for  $1 \leq j \leq J$ . (Note that our choice of  $M$  ensures that  $r_j \geq 1$  for each  $j$  with  $p_j > 0$ , which guarantees at least one ball of type  $j$  below.)
3. For each  $1 \leq j \leq J$ , place  $a_j := \lfloor r_j \rfloor$  balls of type  $j$  into the bowl.
4. Set  $u_j = r_j - a_j$ , and  $v_j = \sum_{\ell=1}^j u_\ell$  for  $1 \leq j \leq J$ , with  $v_0 = 0$ . Also set  $K = v_J$ . (Thus  $K = \sum_j u_j = \sum_j r_j - \sum_j a_j = M - \sum_j a_j$  which is a non-negative integer.)
5. Simulate independent uniform random variables  $U_1, \dots, U_K$  with  $U_i \sim \text{Uniform}[i-1, i)$ .
6. For each  $1 \leq j \leq J$ , let  $b_j = \#\{i : U_i \in [v_{j-1}, v_j)\}$  be the number of random variables  $U_i$  which lie in the interval  $[v_{j-1}, v_j)$ , and add  $b_j$  additional balls of type  $j$  to the bowl. (Note that we must have  $0 \leq b_j \leq 2$ , and  $\sum_j b_j = K = M - \sum_j a_j$ . Furthermore, the total number of balls of type  $j$  is  $m_j = a_j + b_j$ , so the total number of balls in the bowl is  $\sum_j m_j = \sum_j a_j + \sum_j b_j = \sum_j a_j + (M - \sum_j a_j) = M$ .)
7. Select a ball uniformly at random from the  $M$  balls in the bowl.

**Proposition 2** *Given a collection  $p_1, \dots, p_J$  of non-negative probabilities summing to 1, the above procedure selects a ball of type  $j$  with probability  $p_j$ .*

**Proof.** The interval  $[v_{j-1}, v_j)$  has length  $v_j - v_{j-1} = u_j < 1$ . If it lies entirely inside an interval  $[i-1, i)$ , then  $\mathbf{P}(U_i \in [v_{j-1}, v_j)) = u_j$ , so  $\mathbf{E}(b_j) := \mathbf{E}[\#\{i : U_i \in [v_{j-1}, v_j)\}] = u_j$ . Or, if there is an integer  $i$  with  $i-2 < v_{j-1} \leq i-1 < v_j \leq i$ , then  $\mathbf{P}(U_{i-1} \in [v_{j-1}, v_j)) = (i-1) - v_{j-1}$  and  $\mathbf{P}(U_i \in [v_{j-1}, v_j)) = v_j - (i-1)$ , so we still have  $\mathbf{E}(b_j) := \mathbf{E}[\#\{i : U_i \in [v_{j-1}, v_j)\}] = (i-1) - v_{j-1} + v_j - (i-1) = u_j$ . Hence, in any case,  $\mathbf{E}(b_j) = u_j$ , whence  $\mathbf{E}(m_j) = a_j + u_j = a_j + (r_j - a_j) = r_j = M p_j$ . That is, the expected number of balls of type  $j$  is proportional to  $p_j$ . Hence, the probability that a ball drawn uniformly at random will be of type  $j$  is also proportional to  $p_j$ . Then, since  $\sum_j p_j = 1$ , this probability is actually equal to  $p_j$ , as claimed. ■

## 5.2. Estimating the Probabilities.

To use the previous algorithm for group draws would require that we know the conditional probabilities  $p_j = n_j / \sum_i n_i$  for the next team to be chosen in a partially-completed group draw, where  $n_j$  is the number of ways of completing the draw with team  $j$  in the next position.

Now, it might perhaps be possible to compute the  $n_j$  directly. We first observe that the number of possible completions depends only on the geographic region match-ups of the various teams, not on the actual team names. So, in the 2022 FIFA World Cup as described in Section 1.1, Pot 1 can be distributed arbitrarily without affecting the subsequent  $n_j$ . Then, all we need to know about the Pot 2 teams is how many **Eu** teams were put in the same group as a **Eu** team from Pot 1, whether the two **NA** teams were put with **Eu** or **As** or **SA** teams, and so on. This leads to cascading combinatorial counts for the numbers of ways of putting different regions into different positions in the draw. Then, multiplying by corresponding factorials gives the numbers of ways of placing actual teams into the draw.

Such combinatorial problems can eventually be solved. However, they quickly become rather complicated and messy. Furthermore, they have to be re-computed for each possible partial draw, and the calculations are entirely different for different regional distributions of the teams in each pot. So, this does not appear to be a feasible way of proceeding.

Fortunately, there is a practical alternative. Proposition 2 remains true if the above algorithm instead used values  $\hat{p}_j$  which were unbiased estimates of the  $p_j$ . That is:

**Proposition 3** *Given a collection  $\hat{p}_1, \dots, \hat{p}_J$  of non-negative unbiased estimators summing to 1, with  $\mathbf{E}(\hat{p}_j) = p_j$ , if we replace  $p_j$  by  $\hat{p}_j$  throughout in the above procedure, then it will still select a ball of type  $j$  with probability  $p_j$ .*

**Proof.** Conditional on the values of the estimators  $\hat{p}_j$ , the proof of Proposition 2 shows that the probability of selecting a ball of type  $j$  is equal to  $\hat{p}_j$ . Hence, taking expectations over the  $\hat{p}_j$ , it follows that the probability that the algorithm using the estimators  $\hat{p}_j$  will



select a ball of type  $j$  is given by  $\mathbf{E}(\widehat{p}_j) = p_j$ , as claimed. ■

**Remark.** The approach of using an estimator  $\widehat{p}_j$  for simulation, without knowing the true value  $p_j$  is the defining feature of a collection of simulation techniques known as retrospective sampling, see [1]. However, our use of these methods here is rather different from those in the literature which have focused largely on Bayesian inference and exact diffusion simulation problems.

The advantage of Proposition 3 is that there is a natural way to estimate the  $p_j$  in an unbiased way: classical Monte Carlo samples. That is, given a partial draw, we can generate a large number  $N$  of valid draw completions using a Rejection Sampler similar to Section 2.1 above. (In fact it is even easier, since part of the draw is already chosen and does not need to be sampled.) Then, we can estimate  $p_j$  by the fraction of those valid draw completions which have team  $j$  in the next position. This gives a good unbiased estimator  $\widehat{p}_j$  of the true conditional probability  $p_j$  that team  $j$  would be placed in the next position according to the uniform distribution  $\mathbf{U}$ . That provides the final piece of the puzzle for us to produce an effective Multiple-Balls uniform draw sample, which we now present.

### 5.3. Generating Draws with Multiple Balls.

Combining the previous ideas together gives the following algorithm for choosing the team in the next position of a uniform group draw with distribution  $\mathbf{U}$ , by drawing uniformly at random from a modest number of balls, as follows:

1. Select a positive integer valued algorithm parameter  $N$ .
2. Simulate  $N$  different uniformly-distributed completions of the current partial draw, using a Rejection Sampler.

3. For each team  $j$ , let  $n_j$  be the number of such completions which have team  $j$  in the next position, and set  $\hat{p}_j = n_j/N$ .
4. Let  $M = \lceil \max\{(1/\hat{p}_j) : \hat{p}_j > 0\} \rceil$  be the ceiling of the reciprocals of the non-zero  $\hat{p}_j$ .
5. Set  $r_j = M\hat{p}_j$ , for  $1 \leq j \leq J$ . (So,  $r_j \geq 1$  whenever  $n_j > 0$ .)
6. For each  $1 \leq j \leq J$ , place  $a_j := \lfloor r_j \rfloor$  balls of type  $j$  into the bowl.
7. Set  $u_j = r_j - a_j$ , and  $v_j = \sum_{\ell=1}^j u_\ell$  for  $1 \leq j \leq J$ , with  $v_0 = 0$ . Also set  $K = v_J$ .
8. Simulate independent uniform random variables  $U_1, \dots, U_K$  with  $U_i \sim \text{Uniform}[i-1, i)$ .
9. For each  $1 \leq j \leq J$ , let  $b_j = \#\{i : U_i \in [v_{j-1}, v_j)\}$  be the number of random variables  $U_i$  which lie in the interval  $[v_{j-1}, v_j)$ , and add  $b_j$  additional balls of type  $j$  to the bowl (so the total number of balls of type  $j$  is  $m_j = a_j + b_j$ ).
10. Select a ball uniformly at random from the  $M$  balls in the bowl.

Then, it follows from Proposition 3 that:

**Proposition 4** *If the above procedure is used sequentially for each position of a group draw, then the final group draw will have distribution  $\mathbf{U}$ , i.e. will be equally likely to be any of the potential valid draws.*

An interactive simulation of this group draw generation method is available at [18].

**Remark.** The above procedure can be improved in a few minor ways. For example, the true conditional probabilities  $p_j$  must be the same for all teams in the same geographical region, so it is possible after step 3 to replace each  $\hat{p}_j$  by the average of the  $\hat{p}_i$  over all teams in the same region as team  $j$ . Also, if it happens after step 9 that the final numbers of balls  $m_j$  have a non-trivial common factor, i.e.  $\gcd\{m_j\} > 1$ , then each  $m_j$  can be divided by this common factor to produce a simpler draw which still maintains the same probabilities.

## 6. A Multiple-Rejections Solution.

Finally, we present a somewhat different solution, which still selects the teams sequentially one position at a time, and still uses uniform draws from bowls with modest-sized numbers of balls representing the different teams, and still produces a completely uniform valid draw having distribution  $\mathbf{U}$ , but with details which are somewhat different as we now describe.

Given a partial draw, suppose we know the number  $n_j$  of valid ways of completing the draw with team  $j$  in the next position. In terms of this, let  $\mathcal{T}$  be the set of all teams with  $n_j > 0$ , and set  $n_{max} = \max_{j \in \mathcal{T}} n_j$ . Then, for each team  $j \in \mathcal{T}$ , we sample a geometric random variable  $G_j \sim \text{Geometric}(n_j/n_{max})$ . (The  $G_j$  could be produced automatically by a computer in advance.) Thus,  $G_j \geq 1$ . Also, if  $n_j = n_{max}$ , then  $G_j = 1$ . Furthermore, if the  $n_j$  are all roughly equal, then most (if not all) of the  $G_j$  will be equal to 1.

Given these  $G_j$  values, the selection process proceeds as follows:

1. Choose one of the teams  $j \in \mathcal{T}$ , uniformly at random (e.g. from balls in a bowl).
2. If team  $j$  has now been chosen a total of  $G_j$  times, then select team  $j$  for the next position in the draw.
3. Otherwise, return to step 1 and again choose a team uniformly at random.

In this way, the above procedure produces a “race” in which the different teams are each trying to be the first to be chosen  $G_j$  times. Eventually one team will be chosen  $G_j$  times, and will then be selected for the next position in the draw.

The usefulness of this procedure is given by:

**Proposition 5** *If the above procedure is used sequentially for each position of a group draw, then the final group draw will have distribution  $\mathbf{U}$ , i.e. will be equally likely to be any of the potential valid draws.*

**Proof.** It suffices to show that each new position of the draw is filled with the correct conditional probability according to  $\mathbf{U}$ . This follows because the above procedure is actually a way of carrying out a corresponding Rejection Sampler. Indeed, consider a Rejection Sampler with proposal distribution which is uniform on  $\mathcal{T}$ , with acceptance probability  $n_j/n_{max}$  for each team  $j$ . Then each choice of team in step 1 corresponds to one proposal from the Rejection Sampler. And,  $G_j$  corresponds to the number of times that team  $j$  must be proposed before it is finally accepted. Hence, team  $j$  being the first to be chosen  $G_j$  times is precisely equivalent to team  $j$  being the team selected by a Rejection Sampler with target probabilities proportional to  $n_j/n_{max}$ , i.e. proportional to  $n_j$ . The result follows. ■

Although this multiple-rejections solution produces a valid uniform draw, it requires knowledge of the actual  $n_j$  values, which might be difficult in practice (as discussed in Section 5.2 above). Hence, for actual group draws, we believe that the methods of Sections 4 and 5 above are preferable.

## 7. Discussion.

This paper has considered the challenge of designing football group draw mechanisms which have the uniform distribution over all valid draw assignments, but are also entertaining, practical, and transparent. We have explained (Section 2) how to simulate the FIFA Sequential Draw method, to compute the non-uniformity of its draws by comparison to a uniform Rejection Sampler. We have then proposed two practical methods (Sections 4 and 5) of achieving the uniform distribution while still using balls and bowls in a way which is suitable for a televised draw. These two solutions can be tried interactively at [18].

Our approach also has implications for other football draws, such as the upcoming UEFA (European) tournament in August 2022. UEFA uses a slightly different procedure from FIFA, first determining the set of teams which are eligible for the next slot and then selecting one

of them uniformly at random. This procedure, like the FIFA procedure, again leads to non-uniform draws. In fact, it can be seen as an approximation to a Sequential Monte Carlo algorithm (e.g. [7]) in which all of the non-zero incremental particle weights are assumed to be equal. We plan to study the UEFA draw in greater detail in subsequent work.

**Acknowledgements.** We thank David Firth for helpful discussions of these topics. GOR was supported by EPSRC grants Bayes for Health (R018561) and CoSInES (R034710). JSR was supported by NSERC discovery grant RGPIN-2019-04142.

## References

- [1] A. Beskos, O. Papaspiliopoulos, and G. O. Roberts. Retrospective exact simulation of diffusion sample paths with applications. *Bernoulli*, 12(6):1077–1098, 2006.
- [2] Marta Boczon and Alistair J. Wilson. Goals, constraints, and public assignment: A field study of the UEFA Champions League. Technical report, Technical Report 18/016, University of Pittsburgh, Kenneth P. Dietrich School of Arts and Sciences, 2018.
- [3] S. Brooks, A. Gelman, G.L. Jones, and X.-L. Meng (eds). *Handbook of Markov chain Monte Carlo*. Chapman & Hall, 2011.
- [4] László Csató. On the fairness of the restricted group draw in the 2018 FIFA World Cup. *arXiv preprint arXiv:2103.11353*, 2021.
- [5] L. Devroye. *Non-Uniform Random Variate Generation*. Springer-Verlag, New York, 1986. Available at: <http://www.nrbook.com/devroye/>.
- [6] Randal Douc and Olivier Cappé. Comparison of resampling schemes for particle filtering. In *ISPA 2005 – Proceedings of the 4th International Symposium on Image and Signal Processing and Analysis*, pages 64–69. IEEE, 2005.

- [7] Arnaud Doucet, Nando De Freitas, and Neil James Gordon, editors. *Sequential Monte Carlo methods in practice*. Springer, New York, 2001.
- [8] Paul Fearnhead. Sequential Monte Carlo methods in filter theory. *PhD thesis, Oxford University*, 1998.
- [9] FIFA. FIFA World Cup Qatar 2022 Draw Procedures. <https://digitalhub.fifa.com/m/2ef762dcf5f577c6/original/Portrait-Master-Template.pdf>, 2022.
- [10] Julien Guyon. Rethinking the FIFA World Cup final draw. *Journal of Quantitative Analysis in Sports*, 11(3):169–182, 2015.
- [11] Julien Guyon. Pourquoi la coupe du monde est plus équitable cette année. <https://theconversation.com/pourquoi-la-coupe-du-monde-est-plus-equitable-cetteannee-97948>, 2018.
- [12] Julien Guyon. ‘Choose your opponent’: A new knockout design for hybrid tournaments. *Journal of Sports Analytics*, pages 1–21, 2021.
- [13] C. Jones. Toronto professor claims FIFA draw contains hidden bias that skews the odds. <https://www.cbc.ca/sports/soccer/canada-world-cup-toronto-professors-better-draw-chris-jones-april-1-1.6405610>, CBC News, April 1, 2022.
- [14] M.C. Jones. The World Cup draw’s flaws. *The Mathematical Gazette*, 74(470):335–338, 1990.
- [15] Genshiro Kitagawa. Monte carlo filter and smoother for non-gaussian nonlinear state space models. *Journal of computational and graphical statistics*, 5(1):1–25, 1996.
- [16] Stefan Klößner and Martin Becker. Odd odds: the UEFA Champions League round of 16 draw. *Journal of Quantitative Analysis in Sports*, 9(3):249–270, 2013.

- [17] Nicholas Metropolis, Arianna W Rosenbluth, Marshall N Rosenbluth, Augusta H Teller, and Edward Teller. Equation of State Calculations by Fast Computing Machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953.
- [18] G.O. Roberts and J.S. Rosenthal. Football Group Draw Methods and Probabilities. <http://probability.ca/fdraw/>, 2022.
- [19] J.S. Rosenthal. C program to simulate Uniform and FIFA group draws. <http://probability.ca/fdraw/drawprobs.c>, 2022.
- [20] J.S. Rosenthal. World Cup groups – what are the odds? <https://www.thestar.com/opinion/contributors/2022/03/30/for-kicks-we-came-up-with-a-fairer-way-to-determine-the-world-cup-draw.html>, Toronto Star, page A17, March 31, 2022.