GOALS, CONSTRAINTS, AND TRANSPARENTLY FAIR ASSIGNMENTS: A FIELD STUDY OF RANDOMIZATION DESIGN IN THE UEFA CHAMPIONS LEAGUE

MARTA BOCZOŃ AND ALISTAIR J. WILSON

Abstract. We analyze the design of a randomization procedure in a field setting with high stakes and substantial public interest: matching sports teams in the UEFA Champions League. While striving for fairness in the chosen lottery—giving teams similar distributions over potential partners—the designers seek to balance two conflicting forces: (i) imposing a series of combinatorially complex constraints on the feasible matches; (ii) designing an easy-to-understand and credible randomization. We document the tournament's solution, which focuses on sequence of simple uniform draws over each element in the final match, assisted by a computer to form the support for each draw. We first show that the constraints' effects within this procedure are substantial, with shifts in expected prizes of up to a million euro. However, examining all possible counterfactual lotteries over the feasible assignments, we show that the generated inequalities are for the most part unavoidable, that the tournament design is close to a constrained-best. In two extensions we outline how substantially fairer randomizations are possible when the constraints are weakened, and how the developed procedure works in more-general settings.

1. Introduction

The fairness of an assignment across participants is a focal feature in many designed solutions. While managers have to make difficult choices balancing out the many factors involved, the perception among customers and employees that they have been fairly treated is one imperative. Recognizing this, the operations literature has begun to explicitly incorporate equity/efficiency trade-offs into assignment optimizations, with examples such as kidney wait lists and air-traffic control landing slots. While fairness may be achievable ex post in these instances, in many other settings indivisibility of the assigned objects necessarily leads to substantial inequality in realized outcomes. Consequently, equitable

Date: August 1, 2021.

Authors: Boczoń: Department of Economics, University of Pittsburgh, 230 South Bouquet Street, Wesley W. Posvar Hall, Pittsburgh PA, 15213. E-mail: mjb249@pitt.edu. Wilson: Department of Economics, University of Pittsburgh, 230 South Bouquet Street, Wesley W. Posvar Hall, Pittsburgh PA, 15213. E-mail: alistair@pitt.edu. For helpful comments and feedback, we thank Eric Budish, Yeon-Koo Che, Dries Gossenns, Ignacio Palacios-Huerta, Fuhito Kojima, Laszló Csátó, Jean-François Richard, M. Utku Ünver, Leeat Yariv, and the participants of the Conference on Economic Design of SED and the International Conference on Economics of Football.

treatment must instead be defined in terms of ex-ante fairness, emphasizing the similar chances of a good or bad outcome across participants. Our paper analyzes a field design for a constrained randomization procedure, where fairness can only be achieved in expectation, and where focus needs to be placed on more-behavioral requirements to randomization design: transparency and credibility.

While a designer might endeavor to make a randomized assignment fair in expectation, a related but separate issue is to ensure that the randomization is *understood* as fair by participants. That is, a worker might accept their bad fortune in drawing the short-straw for an onerous and uncompensated task as long as they can observe the draw, to understand that their peers were at equal jeopardy. In contrast, when only provided with their assigned outcome by a black-box (such as a computer randomization) they suspect they were unfairly selected, that a manager cherry-picked the outcome. While the design of a physical easy-to-follow draw is trivial in many settings, for assignments with many tasks, many workers, and/or constraints, designing such a randomization becomes considerably more complicated. In this paper, we document a field-proven solution designed for a high-stakes sports tournament under huge public scrutiny: a public drawing of a constrained one-to-one matching in the Union of European Football Association's (UEFA) Champions League (UCL).

The UCL is one of the most successful pan-European ventures, and certainly the one with the most enthusiasm from the general public. Selection into the competition is limited to the highest-performing football clubs from across the continent (and beyond). A series of initial qualifying rounds whittle the number of participating teams down to 32 group-stage participants. From there, half of the clubs advance to a knockout stage that begins with the Round of 16 (R16), followed by four quarter-finals (QF), two semi-finals (SF), and a final (F) that determines a European champion. The focus of our paper is on the tournament's design for matching the sixteen teams at the beginning of the knockout phase into eight team pairs. While a fully symmetric draw would be easy to design if all matches were feasible—drawing team-pairs in turn from an urn without replacement—the problem is complicated by three constraints imposed by the tournament's managers: (i) Each pairing must be between a group winner and a group runner-up (*the bipartite constraint*, a coarse form of seeding). (ii) Teams that played one another in the prior group stage cannot be matched (*the group constraint*, increasing the novelty of the matched teams relative to prior games within the tournament). (iii) Teams from the same national association cannot be matched (*the association constraint*, increasing novelty of the matched teams relative to national competitions).

The intensity of interest means that the tournament is under a magnifying glass: from teams, sponsors, fans, and the media. As an organization UEFA must appease these stakeholders, despite their often diametrically opposed interests. The tournament organizers therefore have a clear interest in creating transparent, easy-to-justify procedures. In terms of the imposed matching constraints, these can be motivated as being either meritocratic (the bipartite constraint) or as serving the common interest by maximizing the tournament's entertainment value (all three constraints). While the constraints substantially reduce the number of possible outcomes, there are still thousands of possibilities, where any realized draw must necessarily leave some teams and their fans ecstatic, and others bereft. UEFA's design objective is to ensure that, modulo the constraints, teams are treated fairly ex ante by the randomization. But more than that, UEFA's randomization also needs to be perceived and understood as fair, a task that becomes substantially more complicated under the imposed matching constraints.

In response to these issues the R16 procedure developed by UEFA follows a hybrid approach. The matching is formed through a physical draw of each team from an urn; however, as the draw proceeds, the urn composition is dynamically adapted by a computer algorithm to ensure that all constraints are satisfied. While the draw's computer-assist is undoubtedly a black-box, the calculations are entirely deterministic, and so can be verified during and/or after the draw by more-sophisticated viewers. In contrast, the stochastic component is transparent. All random elements are both easy to comprehend, as a series of discrete uniform draws, and credible, as each selection is an observed physical randomization.

Our paper analyzes the properties of this designed randomization, using the tools of market design: theory, estimation, and simulation (Roth, 2002). We first theoretically characterize the simple to follow (but combinatorically complex) randomization. Next, we focus on measuring the distortions generated by the constraints. After documenting the quantitatively large effects, over both prize money and match likelihoods, we focus on the normative: Does an alternative randomization exist that is *fairer* to the participants? To answer this question, we employ an objective that measures the absolute difference in the match likelihoods for comparable team pairs—where a pair of teams can be compared on a particular match partner if neither are directly excluded. While easy-to-interpret and broadly applicable, one potential downside of our objective is that it is defined over the space of *expected assignments* and thus, might not be implementable as a lottery over discrete assignments. However, utilizing the main theoretical results in Budish et al. (2013), we show that for the UCL R16 assignment this shift in domain is without loss of generality. As such, the search for optimal *expected* assignments satisfying the constraints (30–40 degrees of freedom) is just as informative on the normative implications as the search for an optimal lottery over constrained assignments (2,000–10,000 degrees of freedom).

Our main results examine the constrained assignment draw across the past sixteen seasons of the UCL as well as across an array of complementary simulations. Overall, we show that while marginally better randomizations *are* possible, the tournament's transparency-first procedure is quantitatively similar to the fairest-possible lottery over constrained assignments. Having demonstrated that there exists only minimal scope to improve on the UEFA design when we consider randomized perfect matchings under direct exclusions, we analyze two extensions. First, we examine the extent to which substantially better outcomes are possible when slacking the tournament's constraints. The exercise not only helps demonstrate how a designer can quantify the fairness effects from enforcing the constraints—here representing a trade-off between efficiency and equity concerns—but also illuminates the greater scope for optimal randomizations when softening the applications' hard exclusions. In a second extension, we illustrate how the methodology applies to a more-general setting, using an example of a constrained many-to-many committee randomization problem.

In terms of the paper's organization, Section 2 provides a brief review of related literature. Section 3 describes the application and outlines the UCL R16 draw procedure. Section 4 discusses the constraint effects on teams' match likelihoods and expected prize money from the tournament. Section 5 shows near-optimality of the UEFA procedure. Sections 6 and 7 outline the extensions, and finally, Section 8 concludes.[1]

2. Literature review

Our paper contributes to two main strands of literature: the issue of fairness for constrained assignment problems (an emerging issue in operations) and randomization design over assignments (a primarily theoretical literature in market design). While our paper's application focuses on a tournament design feature,[2] the main thrust of the analysis is to (i) examine the design of a lottery over the constrained set of assignments and (ii) motivate a more-behavioral design consideration.

Similar to a growing body of applied work, our paper exploits the structure of a sports tournament as a precise field setting to outline/identify an economic idea and method of

[1]The paper's Online Appendices present proofs together with additional (theoretical and empirical) results. Full data, programs, and the Online Supplementary Material are available at https://www.martaboczon.com.

[2]For the substantial literature examining the incentive effects of tournaments see Prendergast (1999).

analysis. Where the applied literature typically centers on positive aspects of individual behavior,[3] our focus is instead on a market-design concern embedded in the tournament design. In this regard, our work is related to a handful of applied papers examining designed markets. Key examples here are: Fréchette et al. (2007), demonstrating the problem of inefficient unraveling in a decentralized market for US college-football bowls; Anbarci et al. (2015), designing a fairer mechanism for penalty shootouts in football tournaments; Baccara et al. (2012), investigating spillovers and inefficiency in a faculty office-assignment procedure; and Budish and Cantillon (2012), studying the superiority of a manipulable mechanism to a strategy-proof one for allocating courses in a business school. In these papers and ours, an applied market-design question is addressed through a mix of theory and structural analysis.

The problem of finding optimal solutions to combinatorial problems has an extensive history in the operations literature (see Von Neumann, 1953; Kuhn, 1955; Orden, 1956; Koopmans and Beckmann, 1957, and references therein), where a number of modern texts offer more comprehensive treatments (see Burkard et al., 2012). More recently, the operations literature has begun to examine the trade-offs between efficiency and fairness in allocation problems, discussing procedures and methods to incorporate equity concerns into the optimization (see Bertsimas et al., 2011, 2012).[4] However, fairness there is typically achievable ex post, through bundles of goods in a combinatorial assignment or through continuous variables such as wait time. In contrast, our paper focuses on finding fair solutions in an ex-ante sense, through a lottery over a set of assignments satisfying a series of constraints. In particular, notions of efficiency for the designer are here integrated into imposed constraints on the set of allowable outcomes, where the randomization is used to generate fairness across this set (in expectation).

Our paper's focus on the ex-ante properties of lottery over assignments is closely related to the literature in mechanism design that goes back to Hylland and Zeckhauser (1979). Problems of fair treatment and efficiency in the realized assignments are complicated by strategic requirements that agents reveal their preferences to the mechanism, often through a pseudo-market approaches (also see Budish, 2011). However, a literature stemming from Bogomolnaia and Moulin (2001) examines mechanisms with simpler implementations: (i) the *random priority* mechanism, analogous to the uniform draw we discuss

[3]Data from football to cricket to golf have been used to illustrate notions from both standard theory (Walker and Wooders, 2001; Chiappori et al., 2002; Palacios-Huerta, 2003) and behavioral biases (Bhaskar, 2008; Apesteguia and Palacios-Huerta, 2010; Pope and Schweitzer, 2011; Foellmi et al., 2016).

[4]Examples include applications in computer networking (Shreedhar and Varghese, 1996; Radunovic and Le Boudec, 2004), air-traffic control procedures (Vossen et al., 2003; Bertsimas and Gupta, 2016), and kidney wait lists (Bertsimas et al., 2013).

in the paper but with agent choice over the partner after selection; and (ii) the *probabilistic serial* mechanism, where agents build up an *expected assignment* by simultaneously 'eating' probability shares across the different outcomes.

While our setting removes any strategic considerations, the main normative insights over randomizations are possible through a relatively new result in the market-design literature. Budish et al. (2013) show that the probabilistic serial mechanism approach extends to a much wider array of problems, so long as a separability condition holds for the constraints.[5] While this result primarily serves as a constructive tool in the market-design literature—allowing a transition from an expected assignment assembled by mechanisms like the probabilistic serial to the lotteries over assignments required for implementation—we employ it as a tool to simplify an optimization problem, to exhaustively search across alternative randomizations. To our knowledge, we are the first to demonstrate the power of this market-design tool in a normative assessment of a field application.[6]

Finally, our paper outlines an implementation issue for randomization design: that the principal may not be fully trusted. Elements of this idea are related to the concerns outlined in Akbarpour and Li (forthcoming), examining the credibility problem for a principal implementing an auction rule. While our setting does not have strategic issues, the concern is similarly over the principal, here over a cherry-picking over possible realizations. Our field application addresses this credibility issue through a physical draw procedure (a common feature to many randomizations like state-lotteries and high-stakes gambling games). While having a physical draw facilitates credibility, a physical, easy-to-follow randomization also helps to ensure that fairness is understood by participants.[7] Hence, the transparent randomization procedure we analyze both mitigates credibility issues and aids understanding of equal treatment. While strategic mechanisms based on random priority can be readily adapted to such requirements, the extent to which other mechanisms like the probabilistic serial have easy-to-follow implementations remains an interesting open design question.

[5]The Budish et al. (2013) result extends the Birkhoff–Von-Neumann theorem (that an expected-assignment matrix can be implemented as a lottery over feasible assignments) to settings with constraints, many-to-many assignments, etc. Also see Akbarpour and Nikzad (2020), who provide a weakened version of a constraint condition that guarantees *approximate* implementability.

[6]While our paper primarily serves as a clear field setting to use the Budish et al. (2013) result as an optimization tool, our results also contribute to the literature on optimal tournament design. See Dagaev and Sonin (2018), Guyon (2018, 2015), Ribeiro (2013), Scarf and Yusof (2011), Scarf et al. (2009), and Vong (2017).

[7]See also Bó and Chen (2019), who document the importance of simplicity and transparency in a historical random assignment for civil servants in Imperial China.

3. Application Background

The UCL is the most-prestigious club competition in football. Its importance within Europe is similar to that of the Superbowl in the United States, though with stronger global viewership figures.[8] Introduced in 1955 as a European Champion Club's Cup (and consisting only of the national champion from each association) the tournament has evolved over the years to admit multiple entrants from each national association (at most five). Since the last major change to the tournament's design took place in the 2004 season,[9] in our empirical analysis we focus on the 16 seasons during 2004–19.

Since the 2004 season, the UCL consists of a number of pre-tournament qualifying rounds followed by a group and then a knockout stage.[10] In the group stage, 32 teams are divided into eight groups of four, where each team plays the other three group members twice (once at home, once away).[11] At the end of the group stage, the two lowest-performing teams in each group are eliminated, while the group winner and runner-up advance to the knockout stage. The knockout stage (except for the final game) follows a two-legged format, in which each team plays one leg at home, one away. Teams that score more goals over the two legs advance to the next round, where the remaining teams are eliminated.[12]

The focus of our paper is on the assignment problem of matching the sixteen teams at the beginning of the knockout phase into eight mutually disjoint pairs.[13] If the problem consisted simply of matching two equal-sized sets of teams under the bipartite constraint, the assignment could be conducted with two urns (one for group winners, one for runners-up) by sequentially drawing team pairs without replacement. However, the presence of the group and association constraints prohibits such a simple procedure for

[8] The UCL final game is globally the most-watched annual sporting event. For example, the 2015 final had an estimated 400 million viewers across 200 countries, with a live audience of 180 million. For comparison, the 2015 Super Bowl had 114 million viewers.

[9] Since each UCL season spans across two calendar years, for clarity and concision, we refer to a particular season by the year of its final game; so 2019 would indicate the 2018–19 season. For more details regarding the format changes, see Table C.1 in Online Appendix C.

[10] A major redesign of the UCL is planned for the 2025 season, with the group-stage being replaced by a league of 36 teams, where current documents suggest no plans for making changes to any of the knockout stages. Consequently, the R16 matching will no longer be affected by the group constraint, but both the bipartite and association constraints will continue to hold.

[11] Seeding in the group stage is determined by the teams' current league ranking and the value of their UEFA club coefficients calculated based on clubs' historical performance.

[12] Ties are broken with the number of goals scored away from home, where a further draw on goals away from home results in extra time, and subsequently penalties as the final tiebreaker.

[13] The QF and SF draws are free from any constraints, and are conducted by drawing balls from an urn without replacement.

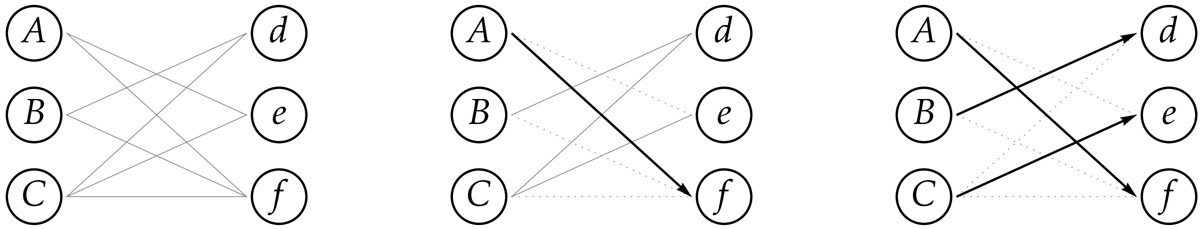

(A) Initial matches (B) Af selected (C) Final matching

FIGURE 1. Perfect one-to-one constrained draw from two urns

two reasons. First, after drawing a team, the urn containing eligible partners cannot contain any directly excluded teams. Second, a match to a non-excluded partner cannot force an excluded match at a later point in the draw. While the first concern is easy to address, the second one requires more-complicated combinatoric inference.

For illustration, consider the example in Figure 1. Suppose that we want to randomly construct a perfect one-to-one matching between teams A, B, and C on one side, to teams d, e, and f on the other. Moreover, assume that match-ups Ad and Be are directly excluded, so there are seven feasible match-ups, as illustrated in Figure 1(A). In a first random draw we select A on the left-hand side. Since d is directly excluded from a match with A, we then randomly choose between e and f on the right-hand side. Suppose f is drawn and the match Af is formed, as shown in bold in Figure 1(B), making the three matches shown with the dotted lines infeasible. At the next stage, suppose we randomly select C on the right-hand side. Team C has no *directly* excluded partners—both Cd and Ce were initially feasible. However, a perfect matching requires that Cd is inhibited from forming, as doing so would leave B with no feasible partner (only e would remain on the right, and Be is directly excluded). Therefore, in the second round, C must be *indirectly* excluded from matching to d. In fact, as soon as Af is selected the only feasible final matching is $\{Af, Bd, Ce\}$, as we illustrate in Figure 1(C).

Although the above example is easy to follow, with eight teams on each side and many more constraints, the combinatorics become involved. While matchings could be formed via fully computerized draws, UEFA has instead opted to use a physical random draw aided by a deterministic computer algorithm to make sure the constraints are satisfied. Specifically, the UEFA draw procedure randomizes the tournament's R16 matching as follows: (i) balls representing the unmatched runners-up (eight to begin with) are placed in the first urn, and one runner-up is drawn uniformly without replacement; (ii) the computer determines the maximal feasible set of group winners that can match with the drawn runner-up, given the constraints and any previous draws; (iii) balls representing the feasible group winners are placed in the second urn, with one drawn uniformly; (iv) a pairing of the two drawn teams (one winner and one runner-up) is added to the aggregate R16 matching. The procedure repeats until all eight matches are formed.

This procedure has three useful design features. First, all randomizations are conducted using a physical draw, and thus are credible.[14] Second, the draw emphasizes the identity of the match, rather than the aggregate matching. This choice not only simplifies the scale of the draw (no more than eight possible realizations), but also highlights that the chances of each element being drawn are perfectly equal. Hence, given the urn composition, it is much easier for the viewer to appreciate their team's fair treatment (though here at the conditional step, rather than overall). Finally, even though the urn compositions are determined in an opaque manner (using a computer to check the maximal set of valid partners), all calculations are entirely deterministic. As such, the calculations can always be checked and explained by more-sophisticated viewers.

These design features transform what could be otherwise a highly esoteric randomization to an easy-to-follow procedure for public consumption. Indeed, the R16 draw ceremony is streamed live by UEFA over the Internet, broadcast by many national media companies, and live blogged by almost every sports page. A rerun of the 2020 UCL R16 draw ceremony currently shows over 1.3 million views on UEFA's YouTube channel, where to the best of our knowledge this (along with the prior group-stage) is likely the most-ardently followed constrained randomization in existence.

4. Constraint Effects in the UEFA Draw

This section highlights the effects of the matching constraints on teams' tournament outcomes. Section 4.1 theoretically characterizes a generalized version of the UEFA draw. Section 4.2 discusses the nature of the distortions generated by the tournament's matching constraints. Finally, Section 4.3 defines two measures of the effects from the constraints and quantifies them across the 16 UCL seasons under consideration.

4.1. **Theory for the Draw.** Let $\mathcal{W} = \{w_1, ..., w_K\}$ and $\mathcal{R} = \{r_1, ..., r_K\}$ denote the sets of group winners and runners-up, respectively, and $\mathcal{V}$ the set of all *possible* perfect (exhaustive one-to-one) matchings between $\mathcal{W}$ and $\mathcal{R}$. We examine a random assignment $\psi : 2^{\mathcal{V}} \to \Delta\mathcal{V}$

[14]Unlike many state lotteries that use mechanical randomization devices to draw outcomes, the UEFA draw is conducted by human third parties (typically famous footballers). While in some sense this might increase the draw's credibility, pointing to football fans' distrust in the process, the human element has led to allegations of UEFA cherry-picking outcomes for favored teams with hot/cold balls (here made by a former FIFA president Sepp Blatter in an interview with Argentine newspaper *La Nacion* on June 13th, 2016).

that takes as input $\Gamma \subseteq \mathcal{V}$ (a set of *feasible* matchings) and provides as output a probability distribution over the elements in Γ.

Algorithm (Γ-Constrained $\mathcal{R}$-first element-uniform draw)**.** *Given an input set of admissible matchings $\Gamma \subseteq \mathcal{V}$, the algorithm selects a matching $\psi(\Gamma)$ in K steps, where at each step a team pair in $\mathcal{R} \times \mathcal{W}$ is formed via two sequential uniform draws.*

***Initialization**: Set $\mathcal{R}_0 = \mathcal{R}$ and $\Gamma_0 = \Gamma$.*
***Step-k** (for $k = 1$ to K):*

(i) Choose R_k through a uniform draw over $\mathcal{R}_{k-1}$;

(ii) Choose W_k through a uniform draw over $\mathcal{W}_k := \{w \in \mathcal{W} \,|\, \exists V \in \Gamma_{k-1} \text{ s.t. } R_k w \in V\}$ (the feasible partners for R_k at step k);

(iii) Define a set of the currently unmatched runners-up $\mathcal{R}_k := \mathcal{R}_{k-1} \setminus \{R_k\}$ and a set of valid assignments given the current draw, $\Gamma_k := \{V \in \Gamma_{k-1} \,|\, R_k W_k \in V\}$.

***Finalization**: After K steps the algorithm assembles a vector of K runner-up–winner pairs, $\mathbf{v} = (R_1 W_1, \dots, R_K W_K)$, where the realization of $\psi(\Gamma)$ is given by $\{R_1 W_1, R_2 W_2, \dots, R_K W_K\} \in \Gamma$.*

In order to characterize the probability of a specific matching $V \in \Gamma$ we define: (i) $\mathcal{P}(V)$, the set of possible sequence permutations for matching V; and (ii) $\mathcal{W}_k(\mathbf{v})$, the set of admissible match partners for runner-up R_k selected at Step-$k(i)$ in the permutation $\mathbf{v}$.[15]

Proposition 1. *Under the Γ-constrained $\mathcal{R}$-first element-uniform draw the probability of any perfect matching $V \in \Gamma$ is given by*

$$\Pr\{\psi(\Gamma) = V\} = \frac{1}{K!} \sum_{\mathbf{v} \in \mathcal{P}(V)} \prod_{k=1}^{K} \frac{1}{|\mathcal{W}_k(\mathbf{v})|}.$$

Proof. See Section A.1 in Online Appendix A. □

Proposition 1 indicates that more than $K! \times |\Gamma|$ calculations are required to understand the chosen lottery over $\Delta\Gamma$. Therefore, even though the cardinality of Γ can be substantially lower than $K!$ due to constraints, the exact computation of $\Pr\{V\}$ involves between $K!$ and $(K!)^2$ steps.[16]

[15]That is, for the permutation $\mathbf{v} = (R_1 W_1, \dots, R_K W_K)$ the set of partners at step k is $\mathcal{W}_k(\mathbf{v})$: $\left\{ w \in \mathcal{W} \middle| \exists V \in \Gamma \text{ s.t. } R_k w \in V \text{ and } \wedge_{j=1}^{k-1} \left(R_j W_j \in V \right) \right\}$.

[16]This can be computationally taxing even for our application with $K = 8$. The main takeaway from the result is that Monte Carlo simulations are best suited for our applied section.

Given the characterization, a remaining question is the extent to which the above calculation can be simplified. Defining two randomization procedures as *distinct* if they induce different probabilities over the matchings in $\mathcal{V}$ we find that:

Proposition 2. *The* Γ*-constrained* $\mathcal{R}$*-first element-uniform draw is distinct from:*

(i) A uniform draw over Γ*;*

(ii) A sequential uniform draw of Γ*-feasible team pairs;*[17]

(iii) The same draw where we switch the labeling of $\mathcal{R}$ *and* $\mathcal{W}$ *(the* Γ*-constrained element-uniform draw where we draw from* $\mathcal{W}$ *first).*

Proof. See Section A.2 in Online Appendix A for counter-examples. □

The first two parts of Proposition 2 indicate that the environment is not equivalent to two computationally simpler uniform draws, while the third part shows that the procedure is asymmetric, in the sense that it does matter which side you draw from first.[18]

The above characterizes a generalized version of the procedure developed by UEFA to assemble the R16 matching at $K = 8$, where the procedure imposes the bipartite constraint by construction. We now provide more detail on how UEFA selects the feasible set Γ from a set of match exclusions $H = H_A \cup H_G \subset \mathcal{R} \times \mathcal{W}$, where the match exclusions are the union of the association-level exclusions H_A and group-level exclusions H_G.[19]

The feasible matching set for the UEFA implementation of the Γ-constrained $\mathcal{R}$-first element-uniform draw is given by

$$\Gamma_H := \left\{ V \in \mathcal{V} \middle| V \cap H = \emptyset \right\}.$$

In the absence of the association constraint, the tournament has 14,833 possible R16 matchings, where each same-nation exclusion significantly reduces the number of valid assignments in Γ_H.[20] Across the 16 seasons under consideration, the number of valid assignments ranged from 2,988 in the 2009 season to 6,304 in 2011 to 9,200 in 2006.

[17]That is, we consider a sequential uniform draw of feasible match pairs, where if the draw has already selected pairs in the set V_t the draw is a uniform over: $\mathcal{M}(V_t) := \{rw \in \mathcal{R} \times \mathcal{W} | \exists V \in \Gamma \text{ with } rw \in V \wedge V_t \subset V\}$.

[18]While distinct, in our particular setting the three draw procedures lead to only marginally different outcomes. Consequently, expected assignments under the UEFA procedure can be approximated fairly well by a uniform draw over Γ, which generates assignment probabilities in fractions of second.

[19]The precise set H therefore varies across seasons depending on the group-level assignment and the composition of teams in the R16.

[20]A political constraint also excludes Russian teams from being drawn against Ukrainian teams. In what follows, we re-interpret this restriction as a part of the association constraint.

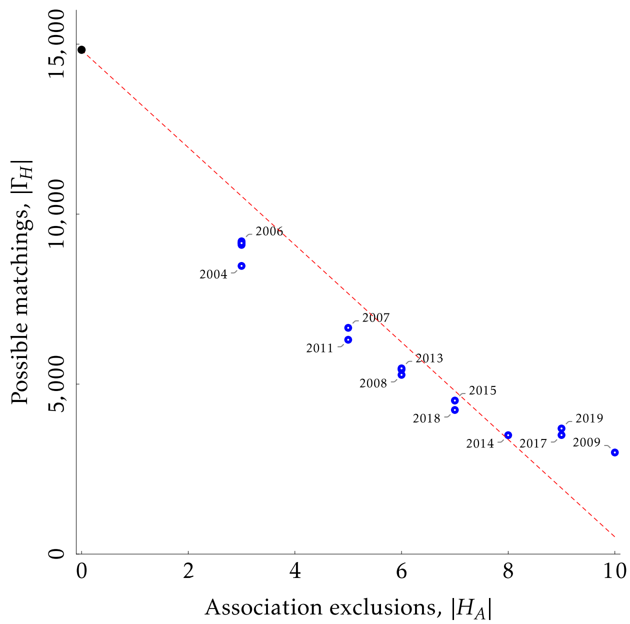

FIGURE 2. Possible assignments against same-nation exclusions
Figure details: Red dashed line indicates a fitted linear relationship; intercept constrained to 14,833.

We graph the relationship between the number of possible matchings and the number of same-nation exclusions within the association constraint H_A in Figure 2. While the number of feasible assignments is not purely a function of the number of exclusions (it depends on their arrangement too) the relationship in question can be approximated by a linear function that decreases by 1,400 matchings for each same-nation exclusion.[21]

4.2. **Example Expected Assignment in the R16.** In Table 1 we provide an example of the expected assignment matrix under the UEFA procedure for the R16 in the 2018 season. Each row represents a group winner, and each column a runner-up, so the row-i–column-j cell indicates the probability that the (ij)-pair is selected within the R16 matching.[22]

The constraints in the 2018 draw are as follows: First, along the diagonal, the probability for each match is zero, reflecting the eight exclusions implied by the group constraint

[21] See Table C.4 in Online Appendix C for the constraints in the remaining seasons between 2004 and 2019.
[22] We calculate all probabilities with a Monte Carlo simulation of size $N = 10^6$. At this size, 95 percent confidence intervals for each probability are within ±0.001 of the given coefficient (see Proposition 4 in Online Appendix A).

Table 1. Expected assignment matrix for the 2018 R16 draw

	Basel	*Bayern Munchen*	*Chelsea*	*Juventus*	*Sevilla*	*Shakhtar Donetsk*	*Porto*	*Real Madrid*
Manchester United	0 (H_G)	0.148	0 (H_A)	0.183	0.183	0.155	0.148	0.182
Paris Saint-Germain	0.109	0 (H_G)	0.294	0.128	0.128	0.108	0.105	0.128
Roma	0.159	0.151	0 (H_G)	0 (H_A)	0.189	0.160	0.152	0.189
Barcelona	0.149	0.144	0.413	0 (H_G)	0 (H_A)	0.150	0.144	0 (H_A)
Liverpool	0.159	0.151	0 (H_A)	0.189	0 (H_G)	0.160	0.152	0.189
Manchester City	0.156	0.148	0 (H_A)	0.183	0.184	0 (H_G)	0.148	0.183
Besiktas	0.109	0.105	0.293	0.128	0.128	0.108	0 (H_G)	0.129
Tottenham Hotspur	0.160	0.152	0 (H_A)	0.189	0.189	0.159	0.151	0 (H_G)

H_G.[23] Second, seven same-nation matches are excluded reflecting the 2018-specific association constraint H_A. Finally, all rows and columns sum to exactly one, as each represents the marginal match distribution for the respective team through the bipartite constraint.[24]

Despite having uniform selections at each point in the draw, the match likelihoods are far from equal, due to asymmetries generated by the association constraint. For illustration, consider Paris Saint-Germain in the 2018 season, the second row of Table 1. As the only French team in the 2018 knockout stage they have no same-nation exclusions, and thus seven feasible match partners. However, the likelihoods of the seven match-ups vary substantially, where the probability that Paris Saint-Germain plays Chelsea is almost three times larger than the probability that they play either Basel, Shaktar Donetsk, or Porto (columns 1, 6, and 7).

4.3. **Quantifying the Association Constraint's Effect.** Below, we quantify the distortions generated by the association constraint that drives the unequal match chances illustrated in Table 1. We first quantify the monetary effect of imposing the association constraint in terms of expected prize money—though the monetary distortions are not easily interpretable as fairness distortions, since the size of the effect vary with the teams' underlying ability. We then measure the distortive effect of the constraints by focusing on the difference in match chances for teams with a common partner, providing a better proxy for the inequality caused by the constraints as it is agnostic as to the team's identity.

We find that:

[23] The bipartite and group constraints impose symmetric restrictions, leading to an equal probability of matching with every non-excluded partner. Consequently, without the association constraint, the expected assignment would have a one-in-seven chance (0.143) for each off-diagonal entry.

[24] The expected assignment matrices for the R16 draw in the remaining seasons can be found in the Online Supplementary Material.

Result 1. *The association constraint imposed on the R16 matching generates substantial effects by: (i) altering expected tournament prizes by up to a million euro; and (ii) creating large inequalities in the match chances for otherwise comparable teams.*

Evidence for Result 1: In order to measure the association constraint's effect on expected tournament prizes we first estimate a commonly used structural model for football-game outcomes (the bivariate Poisson, see Maher, 1982; Dixon and Coles, 1997). The model produces season-specific estimates of attacking and defensive performance of each R16 team in each season in 2004–19. Armed with these estimates and knowledge of the team prizes awarded for reaching each stage of the competition,[25] we can simulate the later stages of the tournament and calculate the *expected* prize for each team i in each season t, under any specified randomization over the R16. As the parameter estimation is standard we relegate the details to Online Appendix B.

To estimate the total effect of the association constraint we calculate differences in the expected prize money under the current UEFA draw procedure (input set $\Gamma_{H_A \cup H_G}$) and a counterfactual procedure that drops the association constraint (input set Γ_{H_G}).[26]

Teams with a positive association-constraint effect are those benefiting from the association constraint, whereas those with a negative value are being disadvantaged. Across all 16 UCL seasons, the association-constraint effect has a standard deviation of 0.3 million euro (it is mean-zero within each season by construction) and a range of 2 million euro: a cost of 0.8 million euro to Arsenal in the 2014 season (eliminated that year in the R16) and a subsidy of 1.2 million euro to Real Madrid in the 2017 season (where they won the tournament that year). The expected effect from enforcing the association constraint is substantial. In Figure 3(A) we illustrate the range in the association-constraint (the difference between the maximal and minimal effect across the sixteen teams) for each season on the vertical axis, against the number of same-nation exclusions on the horizontal.

We next quantify the constraint effects on the chances of each match pair, on the inequality over team's treatment in an ex-ante sense. In particular, our fairness objective

[25] We use tournament prize amounts from the 2019 season. Prizes range from approximately 19 million euro for a team exiting at the R16, to just over 48 million euro for the team winning the tournament. Actual earnings are substantially larger as they also include media payments, so our figures underplay the size of the effects.

[26] In detail, we first draw $J = 1{,}000$ R16 matchings, $\left\{V_j\right\}_{j=1}^{J}$, under each of the two draw procedures. Then, for each realized R16 matching V_j we simulate the remaining tournament outcomes $S = 1{,}000$ times (the R16 home/away games, QF and SF home/away games, and the final game on neutral soil) using the estimated bivariate Poisson model.

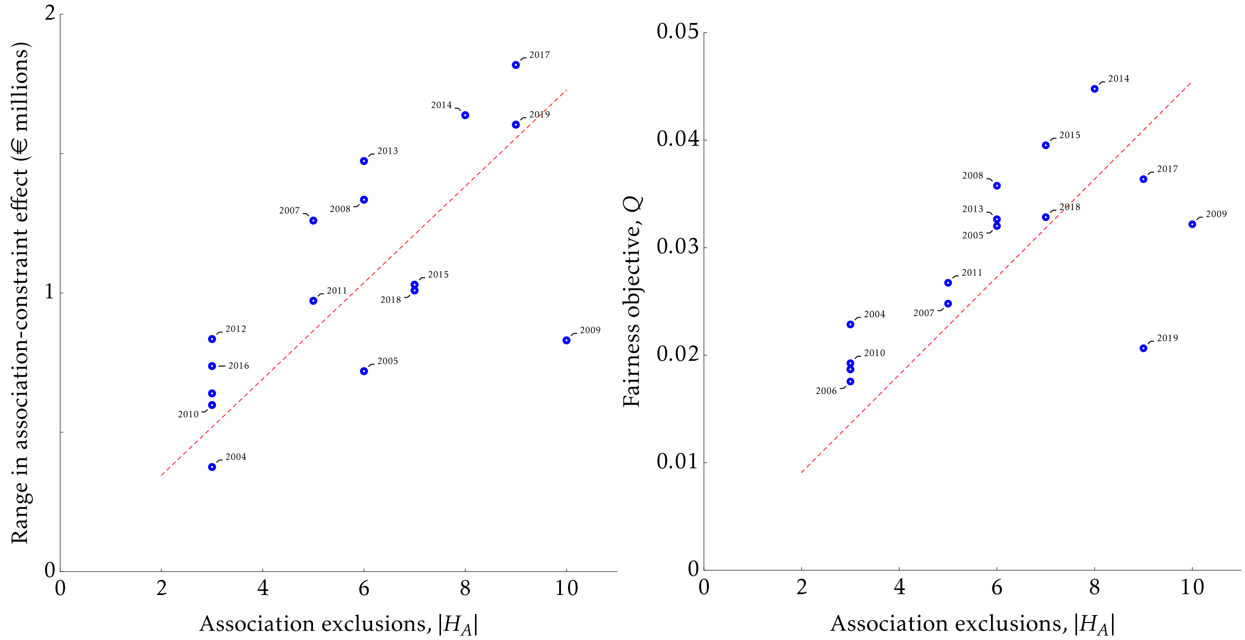

(A) Total effect (prize money)

(B) Indirect effect (match likelihoods)

FIGURE 3. Effects from imposing the association constraint

Figure details: Red dashed line indicates fitted linear relationship.

measures the average absolute difference in match likelihoods across all pairwise comparisons that are not directly excluded by the constraints. That is, teams i and j can be compared on their chance of matching with another team k if neither ik nor jk are directly excluded. For any expected assignment matrix $\mathbf{A}$ like Table 1, where a_{ik} indicates the probability that ik is selected, we can measure the distortion between these two teams using $\left|a_{ik} - a_{ij}\right|$. Taking averages across all possible comparisons our fairness distortion measure is defined by:

$$Q(\mathbf{A};H) = \frac{1}{|\Upsilon_H|} \sum_{(ik,jk)\in\Upsilon_H} \left|a_{ik} - a_{jk}\right|$$

where the set of pairwise comparisons that are not directly excluded is given by:

$$\Upsilon_H := \{(ik,jk)|i,j \in \mathcal{W}, k \in \mathcal{R}, ik, jk \notin H\} \bigcup \{(ki,kj)|k \in \mathcal{W}, i,j \in \mathcal{R}, ki,kj \notin H\}.$$

The distortion measure Q therefore has a minimum at zero when all comparable matches have the same exact chances (for example, if there were no constraints at all each match has a $1/8$ chance) and achieves a maximum with a degenerate assignment (for example, with no constraints, and every pair comparable and the average difference is $\overline{Q}_\emptyset = 1/4$).

In Figure 3(B), we graph the fairness measure Q for the UEFA randomization's expected assignment matrix in each season in 2014–19 against the number of same-nation exclusions. The illustrated relationship indicates that the association constraint substantially distorts the fairness in the ensuing draw. In particular, we find that ten association exclusions cause an expected difference in the match chances for two comparable teams of approximately 5 percentage points. This represents a large relative swing of approximately one-third, when we consider that the common match chance were we to drop the association constraints is 14 percent (a uniform draw over the seven non-group partners).

While the results point to quantitatively large spillovers from the association constraint, we next show that there is only limited scope to ameliorate these effects through better randomization.

5. Near-Optimality of the UEFA Procedure

A natural question raised by the fairness distortions in Result 1 is whether there exists a *better* randomization, one that can reduce the inequality in match chances. However, optimizing over assignment lotteries is computationally complex, as $\Delta\Gamma_H$ is high dimensional. To address this we turn to the core result in Budish et al. (2013) to make the search for alternative randomizations computationally feasible. In particular, the main theorem in Budish et al. states that so long as the constraint structure can be separated into a bihierarchy, then any expected assignment matrix satisfying the constraints is implementable as a lottery over constrained assignments.[27]

Proposition 3 (Implementability). *For any bi-stochastic expected assignment matrix* $\mathbf{A}$ *satisfying a series of match exclusions* H *there exists an equivalent lottery over the perfect matchings in* $\Delta\Gamma_H$.

Proof. Per Theorem 1 in Budish et al., providing any bihierarchy construction over the constraints is sufficient here. Our matching constraints can be decomposed into a bihierarchy over: (i) $\mathcal{H}_1$, the bi-stochastic constraint that each group runner-up is matched to exactly one winner; and (ii) $\mathcal{H}_2$, the bi-stochastic constraint that each group winner is matched to exactly one runner-up, as well each of the singleton exclusions in H. (For a formal construction see Section A.3 in Online Appendix A.) □

[27] A constraint structure $\mathcal{H}$ is termed a *hierarchy* if all of the component constraints are either nested (for example, the sum to one constraint across a team i, and the singleton exclusion ij) or disjoint (for example, the sum to one constraint for team k and the exclusion ij for $k \neq, i, j$). A constraint structure $\mathcal{H}$ is termed a bihierarchy if it can be expressed as the union of two disjoint hierarchies. See Definition 3 in Budish et al. for a precise statement.

This result implies that as long as one can define the optimization objective over expected assignments, any optimization problem over Γ_H (a space with $O(K!)$ degrees of freedom) can be relaxed without loss of generality to an optimization problem over expected assignment matrices satisfying the constraints ($O(K^2)$ degrees of freedom). For our specific UEFA application with $K = 8$, the result reduces the degrees of freedom by two orders of magnitude. That is, across the 16 UCL seasons between 2004–19 the degrees of freedom in the optimization problem are reduced from 2,000–10,000, when searching over $\Delta\Gamma_H$, to 30–40, when optimizing over expected assignment matrices.

5.1. **Examining the 2004–19 UCL Seasons.** In order to investigate whether the UEFA procedure is close to a constrained-best we use Proposition 3 to conduct a computationally tractable optimization, with the fairness distortion measure Q as an objective. Specifically, we define an optimal expected assignment as one that solves the following:

$$\mathbf{A}^{\star} := \arg\min_{\mathbf{A}} Q(\mathbf{A}; H),$$

subject to the matching constraints: (i) $\forall ij \in H^t : a_{ij} = 0$; (ii) $\forall ij : 0 \leq a_{ij} \leq 1$; (iii) $\forall i : \sum_k a_{ik} = \sum_k a_{ki} = 1$.

By comparing the optimal expected assignment $\mathbf{A}_t^{\star}$ in each season t with the expected assignment under the current UEFA draw $\hat{\mathbf{A}}_t$ we arrive at the following result:

Result 2. *While the UEFA randomization is not optimal with respect to the fairness measure Q, it comes quantitatively very close to a constrained-best.*

Evidence for Result 2: In Figure 4 we graph the fairness-distortion measure Q for the fairness-optimized expected assignment $\mathbf{A}_t^{\star}$ on the vertical axis, against the value under the current draw procedure $\hat{\mathbf{A}}_t$ for each season t between 2004–19. While some improvements are possible across the realized constraints in each season, the gain from a fairness-optimal randomization is marginal. On average, we find that optimization can reduce the fairness distortions by approximately a tenth.[28] One potential objection here is that we should be using an alternative fairness measure. However, similar results hold when we repeat the analysis over the following: minimizing the square differences in match probabilities; minimizing the differences between the maximal and minimal positive-probability matches for each team; minimizing the average Kullback-Leibler divergence for each team. Given the similar results across these different measures, we have instead focused on our distortion measure Q, which has the benefit of a simple interpretation.

[28]The size of the reduction is given by the estimated slope coefficient from a regression of $Q(\mathbf{A}^{\star})$ on $Q(\hat{\mathbf{A}})$ for the UCL seasons between 2004 and 2019.

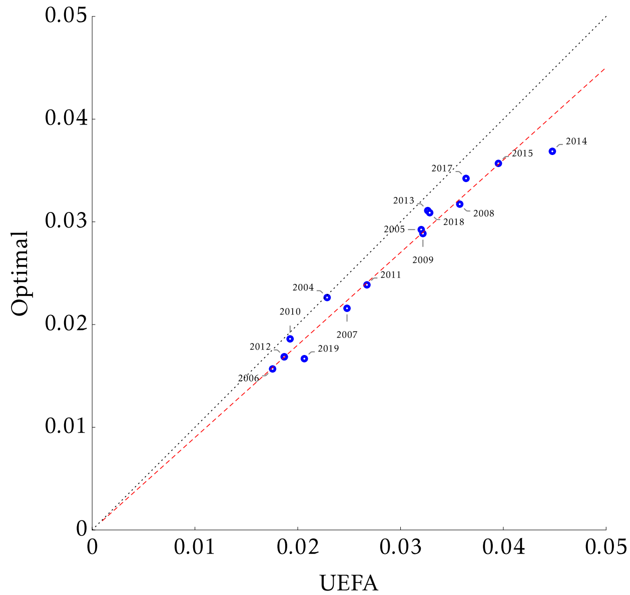

FIGURE 4. Comparison of fairness distortion measure Q: optimal vs. current UEFA procedure

Figure details: Red dashed line indicates fitted linear relationship.

Against the small potential benefits from a fairness-optimized randomization there are large prospective costs in giving up the simple implementation. Procedures yielding the optimal expected assignment $\mathbf{A}^{\star}$ as a lottery over Γ_H are potentially complex in comparison to the current draw, and are assembled as draws over complete matchings, rather than over the component parts.[29] Put against the implementation cost of reduced transparency, a reduction in the average match-chance difference from 5 percent points to 4.5 percentage points seems marginal.[30,31]

5.2. **Examining Simulated Draws.** Above we show that the UEFA procedure is close to a constrained-best for all UCL seasons across 2004–19. We now augment that result by demonstrating that the same property holds under a number of simulated alternatives:

Result 3. *The Γ-constrained $\mathcal{R}$-first element-uniform randomization for a perfect one-to-one matching continues to be close to a constrained-best under a direct exclusion set H as we shift:*

[29] See Online Appendix B to Budish et al. (2013) for a construction.

[30] Note that while there may exist a simple modification of the current assignment rule that would result in fairer match-ups, none of the distinct procedures detailed in Proposition 2 achieve such an end (see Figures C.1-C.3 in Online Appendix C).

[31] The inability to improve upon the expected assignments generated under the UEFA procedure is not driven by a limited scope in moving the expected assignments under the constraints. Taking the 2018 season as an example, we can obtain any value for our spillover measure from a minimum of 0.03 up to a maximum of 0.32.

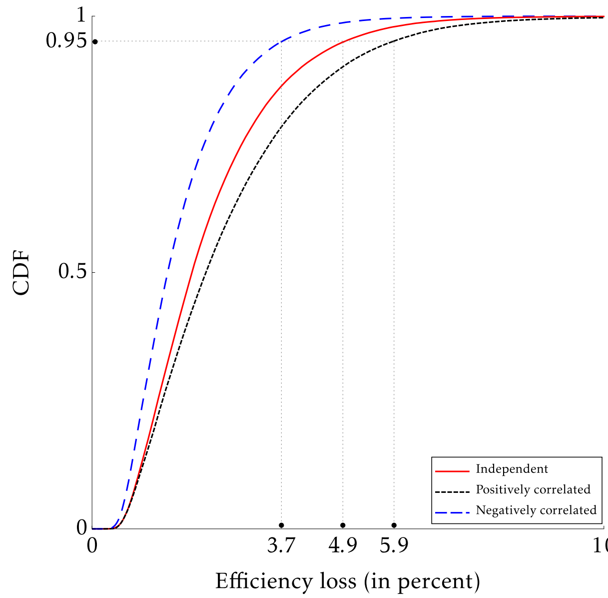

Figure 5. Efficiency loss CDFs across simulations

(i) the number of constraints $|H|$, (ii) the location of the constraints within the matrix, and (iii) the underlying dimension of the assignment problem K.

Evidence for Result 3: We use an array of Monte-Carlo simulations, covering hundreds of thousands of constraint structures. For ease of interpretation, for each simulated constraint structure we measure the efficiency loss associated with choosing the UEFA element-uniform draw over a fairness-optimized randomization via:

$$\Phi = \frac{Q(\mathbf{A};H) - Q(\mathbf{A}^{\star};H)}{\bar{Q}_{\emptyset} - Q(\mathbf{A}^{\star};H)},$$

where $\bar{Q}_{\emptyset} = 1/4$ denotes the maximum value of our fairness measure for an arbitrary assignment under no constraints at $K = 8$.[32]

In our simulations we consider constraint structures where we fix the bipartite and group constraints and randomly vary the number and arrangement of the association exclusions. In particular, we vary: (i) the number of exclusions, from $|H_A| = 5$ to $|H_A| = 20$; and (ii) the distribution of the exclusions across the assignment matrix.[33]

[32]Since both $Q(\mathbf{A}^{\star})$ and $Q(\hat{\mathbf{A}})$ can be zero, expressing Φ in terms of $\bar{Q}_{\emptyset} - Q\left(\mathbf{A}^{\star};H\right)$ prevents Φ from being undefined, where overall level differences are comparable to those in Result 2.

[33]For each chosen value of $|H_A|$ we generate 30,000 constraint structures: one-third using conditional independence across each sequentially drawn exclusion; another third under a positive correlation in placement, making sequentially drawn exclusions in the same row or column more likely; and the final third under negative correlation in placement, with subsequently drawn exclusions in the same row or column less likely. For full details of the simulations see the Online Supplementary Material.

Figure 5 illustrates the empirical CDF for the efficiency loss ϕ at $K = 8$ pooled across values of $|H_A|$, as we do not observe any relative differences in the efficiency-loss measure across $|H_A|$ (per Figure 4 the effect is instead proportional). We do differentiate the results by the process we use to locate the constraints within the matrix, where the figure indicates a clear stochastic dominance relationship. The simulations therefore indicate that the estimated efficiency loss of element-uniform draw is largest when exclusions are likely to fall in the same row or column—the case for the UCL R16 assignment problem—rather than spread across different rows and columns. However, while the inefficiency of the element-uniform randomization does increase when the constraint locations are interrelated, the quantitative level of the effect is still small. Superimposed on Figure 5 we indicate that even when constraint locations are positively correlated the 95 percent of the simulated constraint structures have relative efficiency losses of less than 5.9 percent (compare to the ten percent loss in Figure 4).

In addition to checking for near-optimality of the UEFA procedure at $K = 8$, we also conduct simulations for $K = 6$ and $K = 7$. Again, using randomly generated constraint structures where $|H_A|$ is varying (where exclusion locations are sequentially independent draws) we find that the UEFA assignment rule continues to be close to a constrained-best.[34] Using a linear regression model to examine how the efficiency loss ϕ responds to changes in K we find that the efficiency loss decreases by approximately 3 percentage points for every unit increase in the problem size. The simulation results therefore point to the efficiency loss for the element-uniform draw declining as the dimension of the problem increases.[35]

6. Weakening the constraints

Thus far we have shown that while the fairness distortions generated by the association constraint are substantial, the scope for reducing them through better randomization design is limited. For one-to-one perfect matchings under exclusion constraints the UEFA procedure that draws each element of the matching uniformly is close to a constrained-best. The conclusion then is that the inequality here is an unavoidable consequence of the imposed constraints.

[34]See Online Supplementary Material for most details and comparisons at differing K under the independently drawn same-nation exclusions.
[35]See Table C.7 in Online Appendix C for the estimation results.

A natural question for the designer—particularly if we consider the constraints a design channel to impose efficiency—is to understand and quantify the fairness costs from enforcing the constraints. In our first extension, we therefore quantify what improvements are possible from weakening the association constraint. Given that the constraints are imposed on discrete assignments, we relax the association constraint marginally by allowing for *at most one* same-nation pair in the R16 matching. The relaxed constraint therefore protects the tournament from *excessive* same-nation match-ups.[36]

The current UEFA draw procedure is defined over a valid set of matchings, Γ. One way of weakening the constraint is to do so by holding constant the current draw procedure (the constrained element-uniform draw) but expanding the set of matchings to:

$$\tilde{\Gamma}_{H_A,H_G} := \Big\{ V \in \mathcal{V} \Big| V \cap H_G = \emptyset \text{ and } |V \cap H_A| \leq 1, \Big\}.$$

The modified draw therefore retains its transparent design features, but admits more feasible matchings.

The other option to assess the effect of weakening the constraint is to search for an optimal randomization in the expanded space $\Delta\tilde{\Gamma}_{H_A,H_G}$. Finding an expected-assignment matrix satisfying the relaxed constraints is again computationally simple.[37] However, a major downside of this approach is that the bihierarchy requirement on the constraint structure is no longer satisfied. Indeed, from Theorem 2 in Budish et al. (2013) we know that a feasible expected assignment exists for this problem that is *not* implementable as a lottery over constrained assignments. As such, when optimizing the expected assignment over Q, the results only provide a bound on what is possible to achieve with a lottery over $\Delta\tilde{\Gamma}_{H_A,H_G}$. Fortunately, as we discuss below, the process is still informative, as for some of the most-interesting cases we can constructively show that the bound is attained with a constructive implementation. Nevertheless, while we are able to provide some insight in this particular setting, the exercise illustrates the problems that may arise when universal implementability is not guaranteed.

Analyzing the weakened constraints we find that:

Result 4. *Weakening the association constraint generates a quantitatively large reduction in fairness distortions within the UEFA draw procedure. However, the element-uniform procedure is no longer close to a constrained-best in relative terms. In particular, in cases with*

[36]One indirect benefit of weakening the constraint at the R16 is that it also reduces the likelihood of same-nation games at later stages of the tournament. See Figure C.4 in Online Appendix C.

[37]The optimal expected assignment under the relaxed set of constraints satisfies the following conditions: (i) $\forall ij \in H_G : a_{ij} = 0$; (ii) $0 \leq \sum_{H_A} a_{ij} \leq 1$; (iii) $\forall ij : 0 \leq a_{ij} \leq 1$; (iv) $\forall i : \sum_k a_{ik} = \sum_k a_{ki} = 1$.

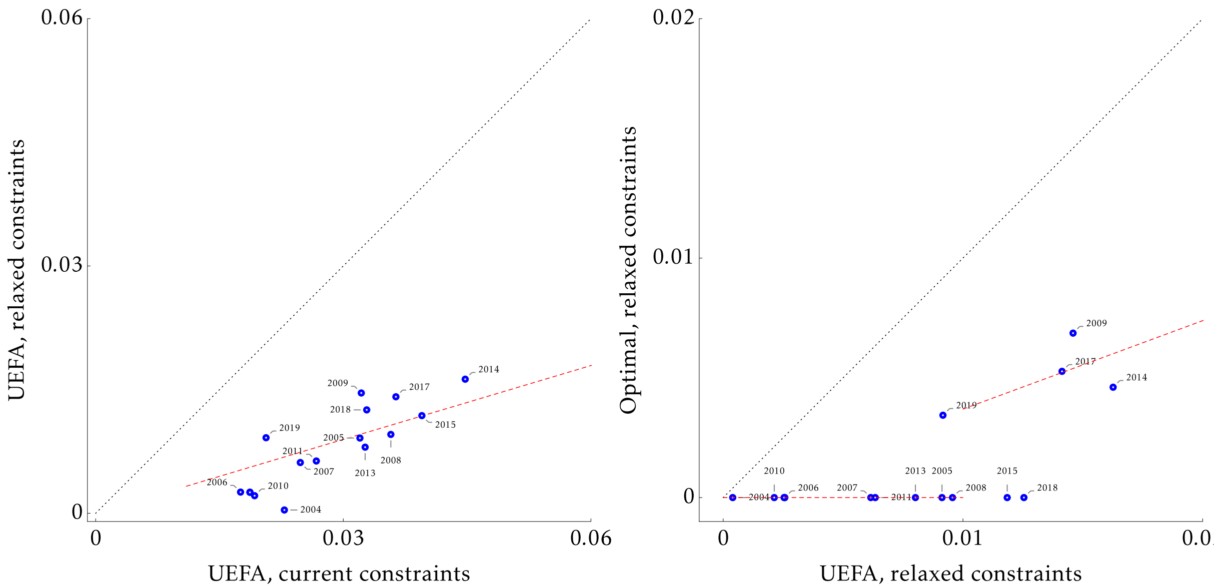

(A) Effect of weakened constraint (UEFA procedure)

(B) Optimal versus UEFA procedure (weakened constraint)

Figure 6. Fairness comparisons with the weakened constraint

Figure details: Red dashed lines indicate fitted linear relationship, where Panel-B relationship allows for boundary solution in years where $|H_A| \leq 7$.

seven or fewer association exclusions, a perfectly fair randomization exists under the weakened constraint.

Evidence for Result 4. We start our analysis by constructing analogs to the results presented in Section 5.1. In Figure 6(A) we illustrate the fairness distortion measure for the current UEFA draw under the expanded feasible set $\tilde{\Gamma}_{H_A,H_G}$ (the weakened constraint) on the vertical axis, where the horizontal axis indicates the same UEFA procedure under the actual constraints Γ.[38] The illustrated results indicate that weakening the association constraint to allow a single same-nation match in the R16 decreases the relative size of the fairness distortions by more than 70 percent. This is obviously a sizable reduction, especially when compared to the 10 percent reduction under an optimal randomization over $\Delta\Gamma_H$ illustrated in Figure 4.

However, optimal randomizations under the weakened constraints can now achieve perfect equity in the match chances for many of the tournament seasons. In Figure 6(B) we illustrate the fairness distortions for a Q-optimal expected assignment satisfying the weakened constraint on the vertical axis, against the UEFA procedure for the weakened

[38] For comparability, in all comparisons we hold constant the pairs compared in the objective Q, excluding all comparisons that include any pair in $H = H_A \cup H_G$.

constraint on the horizontal. Readers attention should first be drawn to the different scales for the axes in Figure 6(A) and (B). Both plots indicate substantially greater fairness when the constraint is weakened. However, the optimal expected assignments here often obtain the first-best outcome. Figure 6(B) suggests that a fairness-optimizing expected assignment can achieve perfect equity under the relaxed constraint in twelve out of the sixteen seasons (where in the remaining four, it reduces the distortions in the UEFA procedure by approximately 50 percent).[39]

A concern with jumping to conclusions on the optimal expected assignment result here is that the weakened constraint structure violates the bihierarchy condition, and universal implementation is not possible.[40] Without the equivalence in the spaces, the optimization results in Figure 6(B) can only be viewed as a lower bound on the value of Q possible in a lottery over $\Delta\tilde{\Gamma}_{H_A,H_G}$. As bounds go, those coinciding with a hard boundary on the range are of questionable content. However, violation of the bihierarchy condition does not imply that a particular expected assignment is not implementable. In fact, each of the boundary solutions are in seasons with seven or fewer association exclusions. In each case it is possible to create a partition of the 56 feasible matched pairs into seven disjoint matchings in $\tilde{\Gamma}_{H_A,H_G}$. A uniform randomization over these seven matchings then leads to a perfectly fair expected assignment.

The analysis of the weakened constraint suggests that our result that the UEFA procedure is close to a constrained-best in fairness terms for this matching environment is related to the fact that the constraints are direct exclusions. However, the analysis also indicates the complications that can arise for the analyst when the constraint structure does not satisfy the bihierarchy restriction.

7. Beyond the UEFA Application

While our tournament application is of standalone interest, the draw procedure and methodology map to other constrained assignment problems. To illustrate this, in a final exercise we briefly outline its use in an example many-to-many constrained assignment—here framed as matching faculty members to committees.

[39] We illustrate the pattern using a piecewise linear fit, where we allow for a different slope coefficient in seasons with more than seven same-nation exclusions.

[40] This is true for all seasons except for 2004, where all same-nation restrictions are imposed on a single team.

Consider the problem of matching eight faculty members—five seniors (S_1 through S_5, where S_1 is the department's chair) and three juniors (J_1 through J_3)—to three committees (A, B, and C). Each committee needs to be composed of three faculty, where we assume the department's objectives for the randomization are to have: (a) a fair expected division of workload across faculty ranks and (b) equitable chances of assignment to each committee within the randomization.

Complicating the assignment, the department wishes to impose a series of constraints: (i) Steering committee A requires exactly two seniors and one junior, and cannot include the department chair S_1. (ii) Hiring committee B requires one senior to chair it, but is otherwise unrestricted. (iii) Tenure committee C requires all three members to be seniors. Beyond the composition constraints, the department imposes a series of constraints aimed at minimizing ex-post differences in workload: (iv) each junior can serve on at most one committee, and each senior must serve on at least one and at most two committees (where we further restrict the department chair S_1 to serve on exactly one). Finally, an idiosyncratic constraint imposed to minimize acrimony requires that: (v) seniors S_2 and S_3 cannot serve together on the same committee.

We illustrate the problem's expected assignment matrix in Figure 7(A), where each entry p_i^j denotes the probability of faculty i being assigned to committee j, and each block of entries represents a constraint.[41]

Even though the considered problem is many-to-many—and with more-complicated constraints than the direct exclusions in the UCL—a randomized assignment can still be constructed with an element-uniform draw. In Figure 7(B) we illustrate the expected assignment matrix resulting from such a constrained draw, where the assignment is assembled sequentially, selecting a committee uniformly from those with faculty slots left to fill, and then choosing uniformly among the feasible faculty for that committee slot (given the prior selections and constraints).[42] As in our field application, this randomization

[41] Vertical blocks indicate the individual workload restrictions, and horizontal blocks the committee's composition restrictions, where the idiosyncratic restriction on S_2 and S_3 is marked in red. Beyond the restriction that each singleton satisfies $0 \geq p_i^j \leq 1$ the constraints are: (i) Committee A: $\sum_{i=2}^{5} p_{S_i}^A = 2$, $\sum_{i=1}^{3} p_{J_i}^A = 1$ and $p_{S_1}^A = 0$; (ii) Committee B: $\sum_{i=1}^{5} p_{S_i}^B + \sum_{i=1}^{3} p_{J_i}^B = 3$, $3 \geq \sum_{i=1}^{4} p_{S_i}^B \geq 1$ and $0 \leq p_{S_2}^B + p_{S_3}^B \leq 1$; (iii) Committee C: $\sum_{i=1}^{5} p_{S_i}^C = 3$, $0 \leq p_{S_2}^C + p_{S_3}^C \leq 1$ and $\forall i\ p_{J_i}^C = 0$; (iv) Workload: $p_{S_1}^B + p_{S_1}^C = 1$, $\forall i > 1\ 1 \leq p_{S_i}^A + p_{S_i}^B + p_{S_i}^C \leq 2$ and $\forall i\ 0 \leq p_{J_i}^A + p_{J_i}^B + p_{J_i}^C \leq 1$.; (v) Idiosyncratic: $\forall j\ 0 \leq p_{S_2}^j + p_{S_3}^j \leq 1$

[42] The only structural change to the constrained element-uniform draw algorithm is that we only remove committees from consideration in step (iii) once their capacities are exhausted (the three slots).

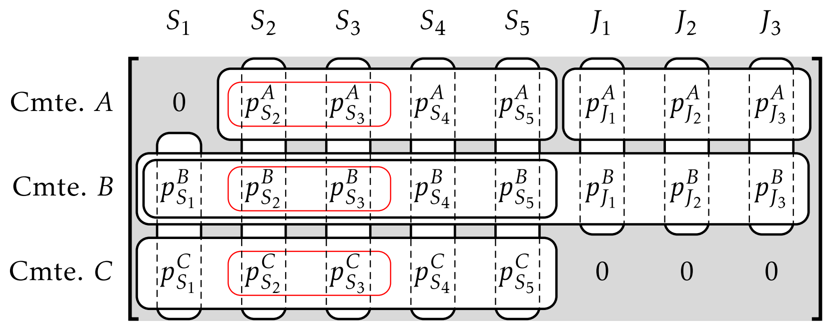

(A) Generic form of the expected assignment

	S_1	S_2	S_3	S_4	S_5	J_1	J_2	J_3
Cmte. A	0	0.45	0.45	0.55	0.55	0.33	0.34	0.33
Cmte. B	0.35	0.41	0.4	0.39	0.39	0.35	0.35	0.36
Cmte. C	0.65	0.47	0.47	0.70	0.70	0	0	0
Total:	1	1.33	1.33	1.64	1.64	0.69	0.69	0.69

(B) Expected assignment under element-uniform randomization

	S_1	S_2	S_3	S_4	S_5	J_1	J_2	J_3
Cmte. A	0	1/2	1/2	1/2	1/2	1/3	1/3	1/3
Cmte. B	0	1/2	1/2	1/2	1/2	1/3	1/3	1/3
Cmte. C	1	1/2	1/2	1/2	1/2	0	0	0
Total:	1	1 1/2	1 1/2	1 1/2	1 1/2	2/3	2/3	2/3

(C) Expected assignment under alternative randomization

FIGURE 7. Committee Examples

can be conducted with a physical draw in front of an audience, in a transparent manner. A computerized aid would still be required to quickly sift through the 576 feasible assignments, though again this component is deterministic.

Inspecting the expected assignment in Figure 7(B) suggests that the department's fairness objectives are largely met with this randomization. However, one legitimate complaint from seniors S_4 and S_5 is that they are negatively affected by the inability of S_2 and S_3 to work together. The two senior faculty with the least constraints are twice as likely to have two assignments than either S_2 or S_3. Does a fairer randomization exist?

The example's constraints satisfy the bihierarchy condition in Budish et al., so the search for fairer randomizations can be conducted entirely over feasible expected assignments.[43] While we could mirror the analysis for the UEFA problem by forming an objective to optimize over, we instead simply point to the feasible expected assignment given in Figure 7(C) that equates expected workloads as well as the chances of each particular assignment (among the comparable faculty).

The developed methodology therefore can come to a distinct conclusion in alternative settings than our field study, where a substantially fairer randomization does exist in this case.[44] However, with the different conclusion, also come different options, where assignment problems have many degrees of freedom. In particular, we note that the designer can also achieve the expected assignment in Figure 7(C) by strengthening the constraint in the element-uniform approach. That is, consider imposing the additional constraints that S_1 serves on committee C with certainty, and that the faculty pairs (S_2, S_3) and (S_4, S_5) are both given exactly three assignments in total. The draw that builds the committee assignment sequentially, uniformly drawing each feasible element in turn has exactly the expected assignment illustrated in Figure 7(C).

Going beyond the specifics of the example, our second extension demonstrates that the simple-to-follow element-wise uniform draw offers a constructive solution in a much wider set of problems. In situations where equity concerns are paramount, such a draw can offer a credible and transparent randomization, providing the various interested parties with a better understanding of their equal treatment. Moreover, when substantially fairer lotteries over the assignments exist, paring down the (typically large) feasible set of assignment with *more* constraints offers a constructive design channel.

[43] From Figure 7(A), the vertical and horizontal blocks can be separated into two distinct hierarchies, where we can put the singleton restrictions for each p_i^j in either. However, plausible examples violating the bihierarchy condition are also easy to conjure. For example, if the faculty subscripts 1 through 4 represent research fields and we imposed a constraint requiring field diversity on the steering committee A.

[44] Moreover, after arriving at this expected assignment, it is also clear that it has a simple implementation, with a randomization over six feasible assignments. A draw could be conducted by rolling a single die here.

8. Conclusion

In many circumstances—particularly those where direct compensation is not possible—managers must create solutions built around fair and equitable treatment. Where outcomes are highly discrete and equity is not possible over a particular realization, designs must necessarily focus on fairness in an expected sense. But this relies on participants' ability to recognize and put faith in their fair treatment by the randomization. For simple settings such as assigning a single task, this can be as achieved with similarly simple means: a physical random draw of names from a hat, conducted in front of the workers being assigned. However, as the complexity of the underlying assignment increases (many tasks, many workers, constraints on the outcomes, etc.) the problem of designing lotteries where participants can perceive their equal treatment becomes much harder.

In this paper, we outline a field solution developed for a random constrained assignment under huge public scrutiny: the draw of competing teams in a sports tournament. The developed randomization is both transparent (in terms of being publicly conducted with a series of simple step) and credible (in the sense of being truly random, where the designer cannot be accused of cherry-picking the realization). At each step, simple uniform draws are used to generate each element of the aggregate assignment, where a computer is used to deterministically enforce the imposed constraints. Our paper demonstrates that the imposed constraints have a substantial effect, both monetarily over expected prizes, and in distorting the fair treatment of otherwise comparable teams. Normatively though, looking across all possible lotteries for the constrained assignments, we show that the chosen procedure comes very close to achieving the fairest possible outcome. Not only is the randomization transparent to the various stakeholders, at least for the one-to-one matching case in the application, it is close to optimal.

The field-proven procedure we document is a dynamic variant of the random-priority mechanism (Bogomolnaia and Moulin, 2001), though here without strategic choice by the selected teams. The randomization provides a positive construction with the potential for application across a number of alternative settings, where satisfaction of any constraints is built directly into the procedure. Moreover, our analysis of the tournament across the past sixteen seasons (as well as across a large number of simulations) suggests that the transparent randomization is close to optimal in this field setting.

In two extension, we show that this near-optimality can fail in alternative settings, where better designed solutions are possible. However, the developed methodology makes possible both the detection of these alternatives and the potential for constructive alternatives. While the documented procedure offers a simple construction, if the constraints can be decomposed á la Budish et al. (2013) a computationally-tractable channel exists for normative assessment. Even within the simple dynamic draw, the large number of degrees of freedom that can make the problem intractable also offers an out when designing alternative randomizations, where imposing *further* constraints on the process can be used as a tool to enhance fairness.

References

Akbarpour, Mohammad and Afshin Nikzad, "Approximate random allocation mechanisms," *Review of Economic Studies*, 2020, *87* (6), 2473–2510.

——— **and Shengwu Li**, "Credible Auctions: A Trilemma," *Econometrica*, forthcoming.

Anbarci, Nejat, Ching-Jen Sun, and M Utku Ünver, "Designing fair tiebreak mechanisms: The case of FIFA penalty shootouts," January 2015. Boston College working paper.

Apesteguia, Jose and Ignacio Palacios-Huerta, "Psychological pressure in competitive environments: Evidence from a randomized natural experiment," *American Economic Review*, 2010, *100* (5), 2548–64.

Baccara, Mariagiovanna, Ayşe İmrohoroğlu, Alistair J Wilson, and Leeat Yariv, "A field study on matching with network externalities," *American Economic Review*, 2012, *102* (5), 1773–804.

Bertsimas, Dimitris and Shubham Gupta, "Fairness and collaboration in network air traffic flow management: an optimization approach," *Transportation Science*, 2016, *50* (1), 57–76.

———**, Vivek F Farias, and Nikolaos Trichakis**, "The price of fairness," *Operations research*, 2011, *59* (1), 17–31.

———**,** ———**, and** ———, "On the efficiency-fairness trade-off," *Management Science*, 2012, *58* (12), 2234–2250.

———**,** ———**, and** ———, "Fairness, efficiency, and flexibility in organ allocation for kidney transplantation," *Operations Research*, 2013, *61* (1), 73–87.

Bhaskar, V, "Rational adversaries? evidence from randomised trials in one day cricket," *Economic Journal*, 2008, *119* (534), 1–23.

Bó, Inácio and Li Chen, "Designing heaven's will: Lessons in market design from the Chinese Imperial Civil Servants Match," *Working Paper*, 2019.

Bogomolnaia, Anna and Hervé Moulin, "A new solution to the random assignment problem," *Journal of Economic theory*, 2001, *100* (2), 295–328.

Budish, Eric, "The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes," *Journal of Political Economy*, 2011, *119* (6), 1061–1103.

——— **and Estelle Cantillon**, "The multi-unit assignment problem: Theory and evidence from course allocation at Harvard," *American Economic Review*, 2012, *102* (5), 2237–71.

———**, Yeon-Koo Che, Fuhito Kojima, and Paul Milgrom**, "Designing Random Allocation Mechanisms: Theory and Applications," *American Economic Review*, April 2013, *103* (2), 585–623.

Burkard, Rainer, Mauro Dell'Amico, and Silvano Martello, *Assignment problems: revised reprint*, SIAM, 2012.

Chiappori, Pierre-André, Steven Levitt, and Timothy Groseclose, "Testing mixed-strategy equilibria when players are heterogeneous: The case of penalty kicks in soccer," *American Economic Review*, 2002, *92* (4), 1138–51.

Dagaev, Dmitry and Konstantin Sonin, "Winning by losing: Incentive incompatibility in multiple qualifiers," *Journal of Sports Economics*, 2018, *19* (8), 1122–46.

Dixon, Mark J and Stuart G Coles, "Modelling association football scores and inefficiencies in the football betting market," *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 1997, *46* (2), 265–80.

Foellmi, Reto, Stefan Legge, and Lukas Schmid, "Do Professionals Get It Right? Limited Attention and Risk-taking Behaviour," *Economic Journal*, 2016, *126* (592), 724–55.

Fréchette, Guillaume R, Alvin E Roth, and M Utku Ünver, "Unraveling yields inefficient matchings: evidence from post-season college football bowls," *RAND Journal of Economics*, 2007, *38* (4), 967–82.

Guyon, Julien, "Rethinking the FIFA World Cup Final Draw," *Journal of Quantitative Analysis in Sports*, 2015, *11* (3), 169–82.

______, "What a fairer 24 team UEFA Euro could look like," *Journal of Sports Analytics*, 2018. in press.

Hylland, Aanund and Richard Zeckhauser, "The efficient allocation of individuals to positions," *Journal of Political economy*, 1979, *87* (2), 293–314.

Karlis, Dimitris and Ioannis Ntzoufras, "Analysis of sports data by using bivariate Poisson models," *The Statistician*, 2003, *52* (3), 381–93.

Koopmans, Tjalling C and Martin Beckmann, "Assignment problems and the location of economic activities," *Econometrica*, 1957, pp. 53–76.

Kuhn, Harold W, "The Hungarian method for the assignment problem," *Naval research logistics quarterly*, 1955, *2* (1-2), 83–97.

Maher, Michael J, "Modelling association football scores," *Statistica Neerlandica*, 1982, *36* (3), 109–18.

Neumann, John Von, "A Certain Zero-sum Two-person Game Equivalent to the Optimal Assignment Problem," in H.W. Kuhn and A.W Tucker, eds., *Contributions to the Theory of Games*, Vol. II of *Annals of Mathematics Studies*, Princeton University Press, 1953, pp. 5–12.

Orden, Alex, "The transhipment problem," *Management Science*, 1956, *2* (3), 276–285.

Palacios-Huerta, Ignacio, "Professionals play minimax," *Review of Economic Studies*, 2003, *70* (2), 395–415.

Pope, Devin G and Maurice E Schweitzer, "Is Tiger Woods loss averse? Persistent bias in the face of experience, competition, and high stakes," *American Economic Review*, 2011, *101* (1), 129–57.

Prendergast, Candice, "The Provision of Incentives in Firms," *Journal of Economic Literature*, 1999, *37* (1), 7–63.

Radunovic, Bozidar and Jean-Yves Le Boudec, "Rate performance objectives of multihop wireless networks," *IEEE transactions on Mobile computing*, 2004, *3* (4), 334–349.

Ribeiro, Celso C., "Sports scheduling: Problems and applications," *International Transactions in Operational Research*, 2013, *19* (1-2), 201–26.

Roth, Alvin E, "The economist as engineer: Game theory, experimentation, and computation as tools for design economics," *Econometrica*, 2002, *70* (4), 1341–78.

Scarf, Philip and Muhammad M. Yusof, "A numerical study of tournament structure and seeding policy for the soccer World Cup Finals," *Statistica Neerlandica*, 2011, *65* (1), 43–57.

———, ———, **and Mark Bilbao**, "A numerical study of designs for sporting contests," *European Journal of Operational Research*, 2009, *198* (1), 190–8.

Shreedhar, Madhavapeddi and George Varghese, "Efficient fair queuing using deficit round-robin," *IEEE/ACM Transactions on Networking*, 1996, *4* (3), 375–385.

Vong, Allen, "Strategic manipulation in tournament games," *Games and Economic Behavior*, 2017, *102*, 562–7.

Vossen, Thomas, Michael Ball, Robert Hoffman, and Michael Wambsganss, "A general approach to equity in traffic flow management and its application to mitigating exemption bias in ground delay programs," *Air Traffic Control Quarterly*, 2003, *11* (4), 277–292.

Walker, Mark and John Wooders, "Minimax play at Wimbledon," *American Economic Review*, 2001, *91* (5), 1521–538.

A.1. **Proof of Proposition 1.** For any matching $V \in \Gamma$, each of the $K!$ possible permutations of V has strictly positive probability, so $\Pr\{V\} = \sum_{\mathbf{v}\in\mathcal{P}(V)} \Pr\{\mathbf{v}\}$, which can be rewritten using the chain rule as

$$\Pr\{V\} = \sum_{\mathbf{v}\in\mathcal{P}(V)} \prod_{k=1}^{K} \left(\Pr\{R_k|\mathbf{v}_{k-1}\} \cdot \Pr\{W_k|R_k, \mathbf{v}_{k-1}\}\right),$$

where $\mathbf{v}_{k-1}$ denotes the matches selected in steps 1 through $k-1$. Since the randomization is uniform at each step, $\Pr\{R_k|\mathbf{v}_{k-1}\}$ simplifies to $\frac{1}{K-k+1}$ and $\Pr\{W_k|R_k, \mathbf{v}_{k-1}\} = \frac{1}{\mathcal{W}_k(\mathbf{v})}$.

A.2. **Proof of Proposition 2.** The constrained $\mathcal{R}$-first element-uniform draw is distinct from: (i) a uniform draw over Γ; (ii) A sequential uniform draw of Γ-feasible team pairs; and (iii) The same draw where we switch the labeling of $\mathcal{R}$ and $\mathcal{W}$ (the Γ-constrained element-uniform draw where we draw from $\mathcal{W}$ first).

A.2.1. *Part (i).* Consider a problem of matching $\mathcal{R} = \{a,b,c,d\}$ to $\mathcal{W} = \{e,f,g,h\}$ under the constraints $H = \{ae, bf, cg, dh, ah, bg, dg\}$.[45] The three resulting feasible matchings are given by $V_1 = \{ag, be, ch, df\}$, $V_2 = \{ag, bh, ce, df\}$, and $V_3 = \{ag, bh, cf, de\}$. Under a uniform draw over $\Gamma = \{V_1, V_2, V_3\}$, the probability of V_1 is equal to $\frac{1}{3}$. However, under the constrained $\mathcal{R}$-first element-uniform draw, (and knowing that a is degenerate) the probability of V_1 is given by:

$$\Pr\{V_1\} = \Pr\left\{be \in V^\star\right\} = \sum_{x\in\mathcal{W}} \Pr\{w_1 = x\} \cdot \Pr\left\{be \in V^\star | w_1 = x\right\} = \tfrac{13}{36}.$$

Hence, the two procedures are distinct.[46]

A.2.2. *Part (ii).* Consider a problem of matching $\mathcal{W} = \{a,b,c,d\}$ to $\mathcal{R} = \{e,f,g,h\}$ under the constraints $H = \{bf, cg, ch, bg, dh\}$. Among the five resulting feasible matchings, only $V_1 = \{af, bh, ce, dg\}$ contains the match af. Under a sequential uniform draw of Γ-feasible team pairs $\Pr\{V_1\} = \Pr\{af \in V^\star\} = \frac{59}{308}$; However, under the constrained $\mathcal{R}$-first element-uniform draw $\Pr\{V_1\} = \Pr\{af \in V^\star\} = \frac{55}{288}$. Hence, the two procedures are distinct.

[45]The proof requires at least a 4×4 market, as a 3×3 is degenerate with the standard symmetric group constraints and a single asymmetric same-nation exclusion. The proof can obviously be extended to any $n\times n$ market by making the remaining $(n-3)$ match partners unique through exclusions.

[46]The proof becomes more cumbersome but still goes through if we removed the dg exclusion that forces ag to be degenerate.

A.2.3. *Part (iii).* Consider a problem of matching $\mathcal{R} = \{a, b, c, d\}$ to $\mathcal{W} = \{e, f, g, h\}$ under the constraints $H = \{bf, cg, ch, bg, dh\}$. Among the five resulting feasible matchings, only $V_1 = \{af, bh, ce, dg\}$ contains the match af.[47] Under the constrained $\mathcal{W}$-first element-uniform draw $\Pr\{V_1\} = \Pr\{af \in V^\star\} = \frac{161}{864}$; However, under the constrained $\mathcal{R}$-first element-uniform draw $\Pr\{V_1\} = \Pr\{af \in V^\star\} = \frac{55}{288}$.[48] Hence, the two procedures are distinct.

A.3. **Proof of Proposition 3.** First, notice that any assignment V can be rewritten as a matrix $\mathbf{X}(V) \in \{0,1\}^{K \times K}$ with a generic entry $x_{ij}(V) = \mathbf{1}\left\{r_i w_j \in V\right\}$ indicating whether or not runner-up r_i is matched to winner w_j. Since V represents a perfect matching between $\mathcal{R}$ and $\mathcal{W}$, $\mathbf{X}(V)$ is a rook-matrix where each row and column have exactly one unit-valued entry with all other entries equal to zero.

Second, for any random draw over Γ_H, the expected assignment matrix is defined as $\mathbf{A} := \mathbb{E}\mathbf{X}(V) = \sum_{V \in \Gamma_H} \Pr\{V\} \cdot \mathbf{X}(V)$, with the generic entry a_{ij} representing the probability of the $r_i w_j$ match.

Next, recall that an expected assignment matrix $\mathbf{A}$ in our setting satisfies the matching constraints if:

(i) Each entry a_{ij} can be interpreted as the probability of $r_i w_j$ being part of V ($\forall i, j : 0 \leq a_{ij} \leq 1$);

(ii) Excluded entries have zero probability ($\forall r_i w_j \in H : a_{ij} = 0$);

(iii) Each row and column can be interpreted as the marginal probability distribution for the respective team ($\forall i, j : \sum_{k=1}^{K} a_{kj} = \sum_{k=1}^{K} a_{ik} = 1$).

As such, the matching constraints can be grouped into two distinct sets: (i) the union of the singleton constraints and the K row constraints; and (ii) the K column constraints and the result follows.[49] Consequently, the matching constraints satisfy the bihierarchy condition in Budish et al. (2013, Theorem 1) and the result follows.

A.4. **Proof of Proposition 4.** Set $J = |\Gamma_H| < K!$ and for each entry $V_j \in \Gamma_H$ set $\Gamma_j = \left\{V_j\right\}$. By Proposition 3, there exists a probability p_j of selecting each admissible matching V_j

[47] A similar counterexample can be constructed with the standard group restriction enforced, but would require a 5×5 market; We omit it for tractability and instead, focus on a 4×4 sub-market.

[48] The four other matchings are: $V_2 = \{ae, bh, cf, dg\}$; $V_3 = \{ag, bh, ce, df\}$; $V_4 = \{ag, bh, cf, de\}$; $V_5 = \{ah, be, cf, dg\}$.

[49] The quotas for each element a_{ij} are a min and a max of zero for the excluded singletons; a min of zero and a max of one for the non-excluded singletons; a min and a max of one for the row sum; and a min and a max of one for the column sum.

that induces and expected assignment matrix satisfying the matching constraints. The existence follows trivially by setting $\Pr\left\{\Gamma_j\right\} = p_j$.

A.5. **Draw procedure of C correlated constraints.** Let $\mathcal{U} = \{(ij) : i, j = 1,...,K, i \neq j\}$, be a set of all off-diagonal elements of $K \times K$ matrix. Let $\mathcal{B}$ denote a set of sampling weights of elements in $\mathcal{U}$. Let $\mathcal{H}_A$ be the set of C drawn constraints Let m_i for all $i = 1,...,K$ denote the cumulative number of times element $(i\cdot)$ has been drawn. Let m_j for all $j = 1,...,K$ denote the cumulative number of times element $(\cdot j)$ has been drawn.

Algorithm. ***Initialization***: *Set* $H_{A,0} = \emptyset$, $\mathcal{B} = \mathcal{B}_0 = 1$, $m_{i,0} = m_{j,0} = 0$.

Step-*c: (for* $c = 1$ *to* C*)*

(i) Chose $(i_c j_c)$ *through a weighted draw over* $\mathcal{U}$ *with weights in* $\mathcal{B}_c$.

(ii) Define the set of currently drawn constraints: $H_{A,c} = H_{A,c-1} \cup (i_c j_c)$.

(iii) Update the values of multiplicity factors by the latest draw:
$\forall i = 1,...,K : m_{i,c} = m_{i,c-1} + \mathbf{1}_{i==i_c}$
$\forall j = 1,...,K : m_{j,c} = m_{j,c-1} + \mathbf{1}_{j==j_c}$

(iv) Define the set of current sampling weights: $\mathcal{B}_c = \{b_{ij,c}\}$ *where*
$\forall (ij) \in H_{A,c} : b_{ij,c} = 0$
$\forall (ij) \in \mathcal{U} / H_{A,c} : b_{ij,c} = b^*_{max\{m_{i,c}, mj,c\}}$

For positively correlated constraints we set $\left[b^*_1, b^*_2, b^*_3, b^*_4, b^*_5\right] = [3, 6, 9, 12, 15]$. *For negatively correlated constraints we set* $\left[b^*_1, b^*_2, b^*_3, b^*_4, b^*_5\right] = [1/3, 1/6, 1/9, 1/12, 1/15]$.

A.6. **Additional Results.**

A.6.1. *Simulation Errors.*

Proposition 4. *Simulating the mechanism* $N = 10^6$ *times leads to 95 percent confidence intervals smaller than ±0.001.*

Proof. Assignments are independent draws from a fixed distribution with a probability of selecting assignment V given by $f(V)$. The probability that the particular match ab is selected is given by $p_{ab} = \sum_{V \in M(ab)} f(V)$ where $M_{ab} := \left\{V \in \Gamma \,\middle|\, ab \in \mu\right\}$ is the set of matchings which include ab. We simulate the vector 8^2-vector $\hat{\boldsymbol{p}}$ where each element in $\hat{p}_{ab}$ is calculated from the N independent simulation assignments $\left(\hat{V}_i\right)_{i=1}^{N}$

$$\hat{p}_{ab} := \frac{1}{N}\sum_{i=1}^{N}\mathbf{1}\left\{ab \in \hat{V}_i\right\}.$$

The vector $\hat{\boldsymbol{p}}$ has the obvious property that $\mathbb{E}(\hat{\boldsymbol{p}}) = \boldsymbol{p}$. We can use the central-limit theorem to show that $\sqrt{n}(\hat{\boldsymbol{p}}-\boldsymbol{p}) \xrightarrow{D} \mathcal{N}_{64}(\mathbf{0},\boldsymbol{\Omega})$ where the variance-covariance matrix $\boldsymbol{\Omega}$ has a generic element given by:

$$\omega_{ab,cd} = \Pr\{ab \wedge cd\} - \Pr\{ab\}\Pr\{cd\},$$

which can be estimated by

$$\hat{\omega}_{ab,cd} = \frac{1}{N}\sum_{i=1}^{N}\mathbf{1}\left\{ab, cd \in \hat{V}_i\right\} - \hat{p}_{ab}\hat{p}_{cd}.$$

However, given a simulation-size of $N = 10^6$, a conservative estimates (as $\omega_{ab,ab} \leq \frac{1}{4}$) for the 95 percent confidence interval for each probability p_{ab} is given by $\hat{p}_{ab} \pm \frac{1.96}{2000} \approx \hat{p}_{ab} \pm 0.001$. □

Appendix B. Data and Estimation of Game-Outcome: For Online Publication Model

In order to account for variation in teams' ability while examining potential effects driven by the tournament's constraints, we estimate a commonly used structural model for football-game outcomes: the bivariate Poisson (Maher, 1982; Dixon and Coles, 1997).

B.1. **Model.** Let S_i and S_j be the random variables indicating the number of goals scored by home-team i and guest-team j in a given game. In a bivariate Poisson model with parameters $(\lambda_1, \lambda_2, \lambda_3)$ the realized scoreline $\left(s_i, s_j\right)$ has a joint probability distribution given by

$$\Pr_{\left(S_i,S_j\right)}\left(s_i, s_j\right) = \exp\{-(\lambda_1+\lambda_2+\lambda_3)\}\frac{\lambda_1^{s_i}}{s_i!}\frac{\lambda_2^{s_j}}{s_j!}\sum_{k=0}^{\min\left(s_i,s_j\right)}\binom{s_i}{k}\binom{s_j}{k}k!\left(\frac{\lambda_3}{\lambda_1\lambda_2}\right)^k,$$

where $\mathbb{E}[S_i] = \lambda_1 + \lambda_3$, $\mathbb{E}\left[S_j\right] = \lambda_2 + \lambda_3$ and $\mathbf{Cov}\left(S_i, S_j\right) = \lambda_3$.

In our specification, we follow Karlis and Ntzoufras (2003) and assume that $\ln\lambda_1 = \mu^t + \eta^t + \alpha_i^t - \delta_j^t$, $\ln\lambda_2 = \mu^t + \alpha_i^t - \delta_j^t$, and $\lambda_3 = \rho^t$., where α_k^t and δ_k^t measure the idiosyncratic offensive and defensive abilities for team k in season t (larger values indicating greater ability), μ^t denotes the season-specific constant, and η^t is the season-specific home-advantage parameter.

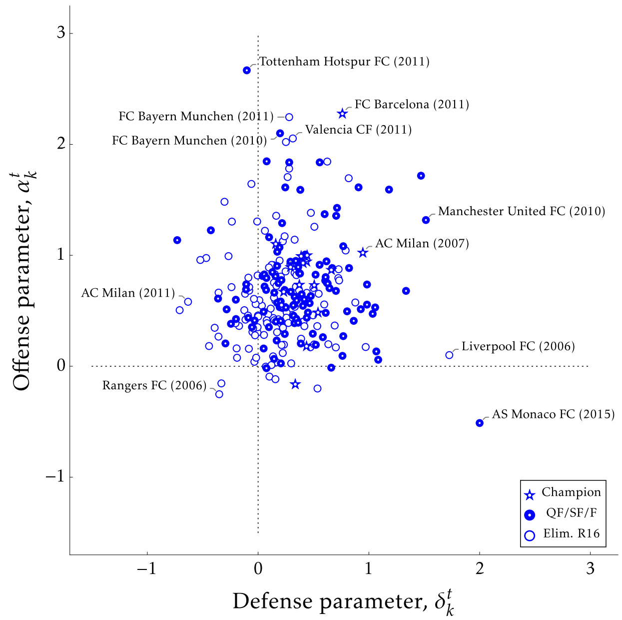

FIGURE B.1. Estimated offense and defense parameters for the R16 teams

Figure details: Figure excludes estimated parameter pairs for teams that failed to reach the R16.

B.2. **Estimation.** We estimate the above model via constrained maximum likelihood separately for each season t between 2004 and 2019. For scale identification we impose two sum-to-zero constraints in each season, forcing $\sum_k \alpha_k^t = \sum_k \delta_k^t = 0$. For estimation in season t we rely on game-level data from the group stage in season t and the group and knock-out stages (except for the final game which is played on a neutral soil) in seasons $t-1$ and $t-2$.[50]

In Figure B.1 we graph the estimated parameters (defense on the horizontal axis, offense on the vertical) for the subset of teams reaching the R16 in the 2004–19 seasons, dropping those that fail to get past the group stage. The strongest teams have large positive values for both the offense and defense parameters; see for example Manchester United in 2010 and FC Barcelona in 2011. Conversely, low-performing teams have either a negative value for the offense parameter (AS Monaco in 2015), the defense parameter (AC Milan

[50]This results in a total of 408 game-level observations used in the estimation for the 2004 season, 376 observations for the 2005 season, and 348 observations for each season between 2006 and 2019. The differences in the number of observations across years result from a change to the tournament design in the 2004 season, where a second group stage feeding into the quarter-finals was replaced by the R16. Table C.5 in Online Appendix C provides summary statistics for the UCL game-level outcomes in all seasons between 2002 and 2019.

Table B.1. Summary statistics for estimated parameters (R16 teams)

Stage	**Offense parameter, α_k^t**				**Defense parameter, δ_k^t**			
	Mean	Med.	Min	Max	Mean	Med.	Min	Max
Elim. in R16	0.62	0.51	-0.20	2.25	0.27	0.17	-0.58	10.99
Elim. in QF	0.69	0.64	-0.51	2.67	0.31	0.26	-0.73	2.00
Reach SF	0.78	0.74	-0.43	2.28	0.66	0.44	-0.43	11.52

in 2011), or both (Rangers in 2006). The large mass of teams are low-to-medium-strength with small but positive values over both the offense and defense parameters.[51]

Complementing the figure, Table B.1 presents summary statistics for the estimated offense and defense parameters broken out by the realized stage reached within the tournament. We find that the eight teams that progress to the quarter-finals are stronger both offensively and defensively than those eliminated in the R16. Similarly, the four teams advancing to the semi-finals have better offensive and defensive performance relative to those knocked out in the quarter-finals. A two-sample Kolmogorov-Smirnov test confirms the pattern, indicating that the empirical distributions of the offense and defense parameters in the R16 and the semi-finals are statistically different ($p = 0.004$ for defense, $p < 0.001$ for offense).

[51]Table C.6 in Online Appendix C provides estimates for the constant term μ^t, the home-advantage parameter η^t, and the correlation coefficient between the number of goals scored by opposing teams ρ^t in the seasons 2004–19.

Table C.1. Format of the post-qualifying stages of the UCL between the 1956 and 2019 seasons

Season(s)	1st knockout phase		Group stage		2nd knockout phase			
	K1	K2	G1	G2	R16	QF	SF	F
1956-1966	16					8	4	2
1967	30	16				8	4	2
1968–1991	32	16				8	4	2
1992–1993	32	16	8					2
1994	32	16	8				4	2
1995–1997			16			8	4	2
1998–1999			24			8	4	2
2000–2003			32	16		8	4	2
2004–2019			32		16	8	4	2

K1 and K2 denote the number of teams competing in the 1st and 2nd knock-out round; G1 and G2 in the 1st and the 2nd round of the group stage; R16 in the R16, QF in the quarter-final, SF in the semi-final, and F in the final game

Table C.2. Number of teams from each association participating in the UCL R16 by season

Season	TOP5						BEL	CYP	CZE	DEN	GRE	NED	POR	RUS	SCO	SUI	TUR	UKR	Total
	ENG	ESP	FRA	GER	ITA	Total													
2004	3	4	2	2	2	13			1				1	1					16
2005	4	2	2	3	3	14						1	1						16
2006	3	3	1	1	3	11					1	2	1		1				16
2007	4	3	2	1	3	13						1	1		1				16
2008	4	3	1	1	3	12					1		1		1		1		16
2009	4	4	1	1	3	13					1		2						16
2010	3	3	2	2	3	13					1		1	1					16
2011	4	3	2	2	3	14				1								1	16
2012	2	2	2	2	3	11		1					1	2		1			16
2013	2	4	1	3	2	12							1		1		1	1	16
2014	4	3	1	4	1	13					1			1			1		16
2015	3	3	2	4	1	13							1			1		1	16
2016	3	3	1	2	2	11	1					1	1	1				1	16
2017	3	4	2	3	2	14							2						16
2018	5	3	1	1	2	12							1			1	1	1	16
2019	4	3	2	3	2	14						1	1						16
Mean	3.4	3.1	1.6	2.2	2	13													

BEL indicates Belgium, CYP Cyprus, CZE Czech Republic, DEN Denmark, ENG England, ESP Spain, FRA France, GER Germany, NED the Netherlands, ITA Italy, POR Portugal, RUS Russia, SCO Scotland, SUI Switzerland, TUR Turkey, UKR Ukraine. TOP5 denotes English, Spanish, French, German, and Italian associations together

Table C.3. Number of same-nation exclusions generated by each national association in the UCL R16 by season

Season	**TOP5**						**POR**	**RUS**	**Total**
	ENG	ESP	FRA	GER	ITA	Total			
2004		3				3			3
2005	4			2		6			6
2006	1	2				3			3
2007		2	1		2	5			5
2008	4				2	6			6
2009	4	3			2	9	1		10
2010			1		2	3			3
2011	3	2				5			5
2012				1	2	3			3
2013	1	4			1	6			6
2014	4			4		8			8
2015	2		1	4		7			7
2016	2					2		(1)	3
2017	2	4	1	2		9			9
2018	4	2			1	7			7
2019	3	2	1	2	1	9			9
Mean	2.8	2.8	1.0	2.5	1.7	5.5			5.6

(1) indicates a constraint generated by FC Zenit (RUS) and FC Dynamo Kyiv (UKR)

Table C.4. Same-nation exclusions in the UCL R16 by season

Season	**TOP5**					**POR**	**RUS**
	ENG	ESP	FRA	GER	ITA		
2004		1|3					
2005	2|2			1|2			
2006	1|1	2|1					
2007		1|2	1|1		1|2		
2008	2|2				2|1		
2009	2|2	1|3			2|1	1|1	
2010			1|1		2|1		
2011	3|1	2|1					
2012				1|1	1|2		
2013	1|1	2|2			1|1		
2014	2|2			2|2			
2015	1|2		1|1	2|2			
2016	2|1						1|1
2017	2|1	2|2	1|1	1|2			
2018	4|1	1|2			1|1		
2019	3|1	1|2	1|1	1|2	1|1		

In a $(m|n)$-pair m indicates the number of seeded (group stage winners) teams and n the number of unseeded (group stage runners-up) teams. $m \times n$ is the total number of exclusions generated by a given association

Table C.5. Summary statistics for the number of goals scored in the UCL

Season	# Games	Average		Std. Dev.	
		Home	Away	Home	Away
2002	156	1.69	0.95	1.18	0.90
2003	156	1.58	1.19	1.36	1.07
2004	124	1.52	0.94	1.37	0.99
2005	124	1.69	0.97	1.46	1.07
2006	124	1.39	0.86	1.31	0.97
2007	124	1.47	1.00	1.29	1.05
2008	124	1.57	1.07	1.42	1.02
2009	124	1.45	1.19	1.34	1.28
2010	124	1.42	1.15	1.23	1.13
2011	124	1.64	1.19	1.44	1.27
2012	124	1.68	1.09	1.54	1.13
2013	124	1.63	1.31	1.30	1.10
2014	124	1.62	1.26	1.35	1.33
2015	124	1.70	1.19	1.60	1.33
2016	124	1.68	1.12	1.42	1.13
2017	124	1.84	1.19	1.68	1.23
2018	124	1.77	1.43	1.53	1.40
2019	124	1.71	1.23	1.43	1.20

From the group stage onward except for the final game played on a neutral ground

Table C.6. Estimated bivariate Poisson model coefficients by season

Season	μ^t	η^t	ρ^t
2004	-0.25	0.42	-2.06
2005	-0.47	0.46	-2.00
2006	-0.71	0.57	-1.88
2007	-0.86	0.55	-1.99
2008	-0.49	0.46	-2.02
2009	-0.85	0.37	-1.71
2010	-1.34	0.29	-1.37
2011	-1.51	0.31	-1.33
2012	-0.69	0.27	-10.34
2013	-0.25	0.34	-2.23
2014	-0.05	0.31	-13.76
2015	-0.04	0.27	-16.38
2016	-0.09	0.34	-15.78
2017	-0.07	0.39	-14.96
2018	-0.05	0.37	-16.59
2019	-0.22	0.32	-16.06
Mean	-0.50	0.38	-7.53

μ^t denotes the constant term, η^t the home-effect parameter, and ρ the correlation coefficient between the number of goals scored by the two opposing teams in the season between 2004 and 2019

Table C.7. Linear regression for simulation results

Regressor	**Estimate**	**Std.Err**
Constant	2.25E-01	6.48E-04
Assignment's dimension	-3.16E-02	1.01E-04
Number of constraints	3.65E-03	1.38E-05
Number of valid assignments	4.52E-06	3.43E-08
Independent constraints=Yes	-5.07E-03	6.65E-05
Negatively-correlated constraints=Yes	-1.15E-02	6.87E-05

The dependent variable is the efficiency loss ϕ. The number of observations is equal to 540,000.

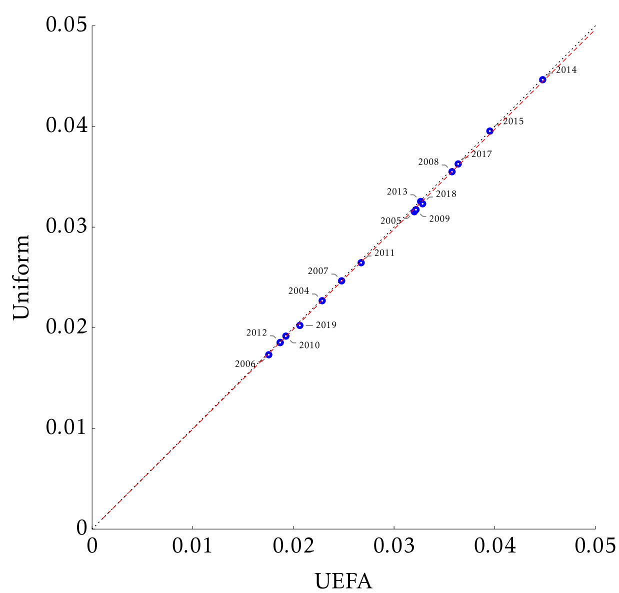

Figure C.1. Comparison of fairness distortion measure Q: Uniform draw over Γ versus current UEFA procedure

Figure details: Red dashed line indicates fitted linear relationship.

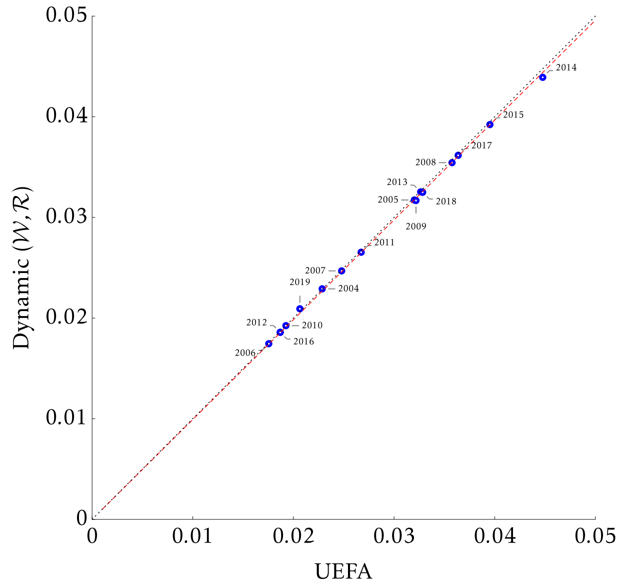

Figure C.2. Comparison of fairness distortion measure Q: Sequential uniform draw of Γ-feasible team pairs versus current UEFA procedure

Figure details: Red dashed line indicates fitted linear relationship.

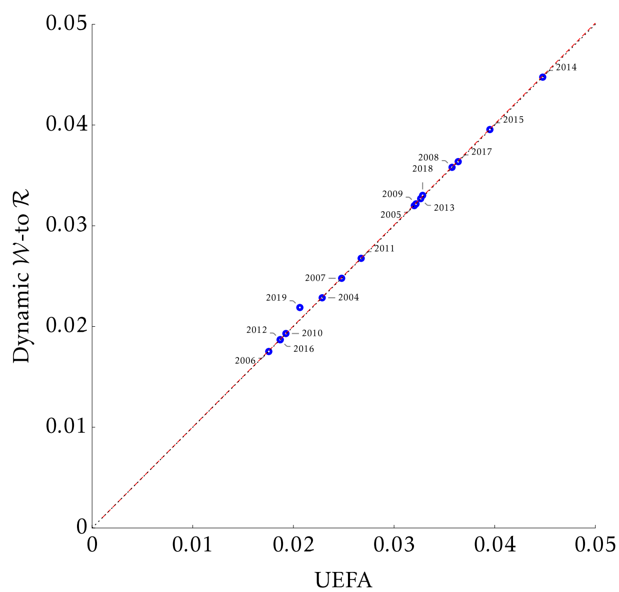

Figure C.3. Comparison of fairness distortion measure Q: Constrained $\mathcal{W}$-to-$\mathcal{R}$ dynamic draw versus current UEFA procedure

Figure details: Red dashed line indicates fitted linear relationship.

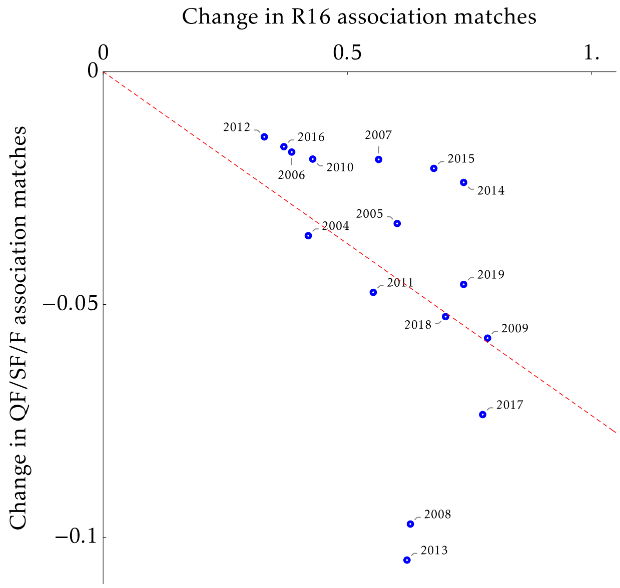

Figure C.4. Change in same-nation match-ups in QF/SF/F against the R16 after relaxing the association constraint

Figure details: Red dashed line indicates fitted linear relationship.