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 sequences of 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 and large distortions in match likelihoods of otherwise comparable team pairs. 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 can be adopted to 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 the many factors involved, the perception among customers and employees that they have been fairly treated is an imperative. Recognizing this, the operations literature has begun to explicitly incorporate equity/efficiency trade-offs into assignment optimizations. However, while fairness may be achievable ex post in some instances (for example, in queuing settings with time as a continuous variable within the objective) in others the indivisibility of the assigned objects necessarily leads to substantial inequality in *realized* outcomes. Under

---

Date: February 2022.

Key words and phrases. Constrained assignment; UEFA Champions League; Public draw mechanism. *JEL*: C44, C78, D82.

Authors: Boczoń: Department of Economics, Copenhagen Business School, Porcelænshaven 16A, 2000 Frederiksberg. E-mail: mbo.eco@cbs.dk. 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. Moreover, we are very grateful to the three anonymous referees for their insightful comments and helpful suggestions.

such circumstances, equitable treatment needs to instead be driven by ex-ante fairness, emphasizing the similar chances of good or bad outcomes across the participants. Reflecting this, our paper analyzes a field design for a constrained randomization, where fairness is only achievable in expectation rather than through a specific realization. In this setting more-behavioral requirements for the randomization design become critical: transparency and credibility.

While a designer might endeavor to make a randomized assignment as fair as possible in expectation, a separate issue is ensuring that the randomization is *perceived* and *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 when they can observe the draw, understanding that their peers were at equal jeopardy. In contrast, if the realized assignment comes from a black-box (say through a computer randomization) they may suspect they were unfairly selected, that a manager cherry-picked the outcome. While the design of physical easy-to-follow random draws is trivial in many settings, for assignments with numerous tasks, workers, and/or constraints, designing a transparent randomization procedure becomes considerably more complicated. In this paper, we document a field-proven solution for a complex constrained assignment in a high-stakes sports tournament under huge public scrutiny: 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 mutually disjoint 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 the novelty of the matched teams relative to concurrent national competitions).

The intensity of interest in the UCL means that the tournament is under a magnifying glass: from teams, sponsors, fans, and the media. As an organization, UEFA must appease the various stakeholders, despite their often diametrically opposed interests. The tournament organizers therefore have a clear interest in creating transparent and 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 stakeholders' common interest in maximizing the tournament's entertainment value (all three constraints). In terms of the chosen randomization procedure, while the constraints substantially reduce the number of possible outcomes, there are still thousands of possible assignments, where any realized draw necessarily leaves some teams and their fans ecstatic, and others bereft. UEFA's design objective is therefore to ensure that, modulo the constraints, teams are treated fairly ex ante by the randomization. But more than that, UEFA also needs the draw to be understood as fair, a task that becomes substantially more complicated under the imposed matching constraints.

In response to these design issues, the R16 procedure developed by UEFA follows a hybrid approach, making the parts of the draw that involve randomization as easy to follow as possible, embedding all of the combinatorial complexity in a series of deterministic steps. The R16 matching is formed through a physical draw of teams to be matched from two urns; however, as the draw proceeds, the urns compositions are dynamically adapted by a computer to ensure that all the constraints on the assignment are satisfied. As such, the random component of the draw is not only easy to comprehend (a series of discrete uniform draws) but also credible (each selection is an observed physical draw). In contrast, the draw's computer-assist algorithm, which carries out a number of non-trivial calculations, is effectively a black-box. However, since all of the computer's calculations are deterministic, they can be verified during and/or after the draw by more-sophisticated viewers. Indeed, in the 2021–22 draw, a mistake in implementing the deterministic parts of the procedure was detected, leading to a redo for the entire R16 draw.

Our paper analyzes the properties of this designed randomization, using the tools of market design: theory, estimation, and simulation (Roth, 2002). First, we 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 average 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 draws for the 2004–22 seasons of the UCL, supplemented by an array of complementary simulations. Overall, we show that while marginally better randomizations *are* possible, the tournament's transparency-first procedure under our objective resembles the fairest-possible lottery over the constrained assignments.

Having demonstrated that there exists only minimal scope to improve on the UEFA design within the application's domain (perfect one-to-one matchings under direct exclusions) we analyze two extensions. First, we examine whether substantially better outcomes are possible when slacking the tournament's constraints. This exercise not only helps to 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 fairness possible through an optimal randomization when the applications' hard exclusion constraints are softened. In a second extension, we illustrate how the UEFA randomization procedure can be extended to a more-general many-to-many setting. Within this second extension (a committee randomization), we echo the previous finding that where the imposed constraints exhibit greater slack, the UEFA-like randomization is no longer close to optimal. However, we also show that by selectively imposing *further* constraints, managers might be able to design for fairness within the transparent randomization procedure. While both extensions function as illustrative examples, the discussion opens up a number of potential avenues for future research. In particular, an open research question is over what is/is not achievable when designing randomized assignments built upon easy-to-follow urn draws.

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 expected prize money from the tournament and teams' match likelihoods. Section 5 shows near-optimality of the UEFA procedure. Section 6 outlines the extensions, and finally, Section 7 concludes.[1]

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

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 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.[4] 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 questions 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).[5] 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

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

[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] Also see Rubin et al. (2021) and Pathak et al. (2021) for further work on the assignment of scarce healthcare resources (vaccinations, vaccines, etc.), where fairness concerns are becoming an important component in the design.

[5] 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).

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 there complicated by strategic requirements that agents reveal their preferences to the mechanism, often through a pseudo-market approaches (also see Budish, 2011). In particular, a literature stemming from Bogomolnaia and Moulin (2001) examines: (i) the *random priority* mechanism, analogous to the uniform draw we discuss 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 extends to a much wider array of problems, as long as a separability condition holds for the constraints.[6] While this result primarily serves as a constructive tool in the market-design literature—allowing a transition from an expected assignment assembled by mechanisms 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.[7]

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 (2020), 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, for instance, state-lotteries and high-stakes

[6]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.

[7]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).

gambling games). While having a physical draw facilitates credibility, an easy-to-follow randomization also helps to ensure that fairness is understood by participants.[8] 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 open design question.

3. UEFA's randomization procedure

This section provides the context for our application: Section 3.1 discusses the main features of the tournament and the UEFA's chosen randomization procedure. Section 3.2 theoretically characterizes a generalized version of the draw.

3.1. **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.[9] 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,[10] in our empirical analysis we focus on the 19 seasons during 2004–22.

Since the 2004 season, the UCL consists of a number of pre-tournament qualifying rounds followed by a group and then a knockout stage.[11] 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).[12] 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

[8]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.

[9]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.

[10]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 2022 would indicate the 2021–22 season. For more details regarding the format changes, see Table C.1 in Online Appendix C.

[11]A major redesign of the UCL is planned for the 2025 season, with the group-stage being replaced by a league of 36 teams. However, 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.

[12]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.

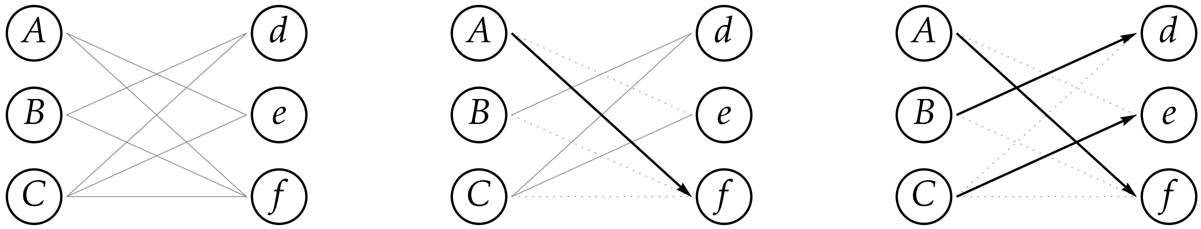

(A) Initial matches (B) Af selected (C) Final matching

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

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.[13]

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.[14] 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 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, and 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

[13]During our period of analysis, ties were broken with the number of goals scored away from home, where a further draw on goals away from home resulted in extra time, and subsequently penalties as the final tiebreaker. Starting from the 2021 season UEFA has abolished the away goals rule.

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

right-hand side, 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 illustrated 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. 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.[15] Second, the draw emphasizes the identity of the match, rather than that of 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 equal. Consequently, 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 identify the maximal set of valid partners), all calculations are entirely deterministic and verifiable.[16]

These three design features transform what could otherwise be a highly esoteric randomization into 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

[15]Unlike many state lotteries, which 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 Sepp Blatter, a former president of the International Federation of Association Football, FIFA, in an interview with Argentine newspaper *La Nacion* on June 13$^{\text{th}}$, 2016).

[16]Speaking to verification, in the 2022 R16 draw a number of implementation errors were made, where some group-stage exclusions were not enforced, leading to a redraw. Importantly, one of the teams (Atletico de Madrid) objected to the initial draw. Whereas the objection could have been over an excluded partner (Liverpool) that was erroneously included in their match draw, it was instead raised over a non-excluded partner (Manchester United) that was not put at equal jeopardy with other draw-eligible teams at this point in the draw. Note that the latter objection requires greater sophistication as it is necessary to verify the precise set of feasible partners before any concern can be raised.

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, 1.6 million views) is likely the most-ardently followed constrained randomization in existence.

3.2. **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 $\psi : 2^{\mathcal{V}} \to \Delta\mathcal{V}$ 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.[17]

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, $\mathbf{v} = (R_1 W_1, \ldots, R_K W_K)$, where the realization of $\psi(\Gamma)$ is given by $\{R_1 W_1, R_2 W_2, \ldots, R_K W_K\} \in \Gamma$.*

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 $\mathbf{v}$.[18]

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 Online Appendix A. □

[17] In Section 6 we generalize this randomization procedure to many-to-many assignments.

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

Proposition 1 indicates that more than $K! \times |\Gamma|$ calculations are required to characterize 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.[19]

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;[20]

(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 Online Appendix A for counter-examples. □

The first two parts of Proposition 2 indicate that the Γ-constrained $\mathcal{R}$-first element-uniform draw is not equivalent to two computationally simpler algorithms, whereas the third part shows that the procedure is asymmetric, in the sense that it *does* matter which side you draw from first.[21]

The above characterizes the randomization procedure used by UEFA to assemble the R16 matching. For the UEFA application, the bipartite constraint is imposed by construction, and the input set of feasible matchings is given by

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

where $H = H_A \cup H_G \subset \mathcal{R} \times \mathcal{W}$ is a set of excluded team pairs, the union of the same-association exclusions H_A and the same-group exclusions H_G. As such, the excluded team-pairs in H vary across seasons depending on the group-level assignment and the composition of teams in the R16.

[19] As such, calculating the entire probability distribution over $\Delta\mathcal{V}$ 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.

[20] 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: $\mathcal{M}(V_t) := \{rw \in \mathcal{R} \times \mathcal{W} | \exists V \in \Gamma \text{ with } rw \in V \wedge V_t \subset V\}$.

[21] 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.

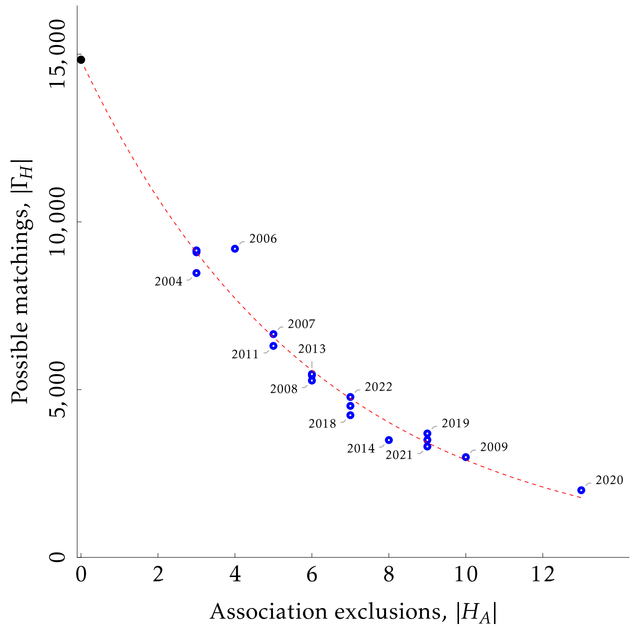

FIGURE 2. Possible matchings against same-nation exclusions
Figure details: Red dashed line indicates a fitted exponential relationship with the intercept constrained to 14,833.

In the absence of the association constraint, the tournament has 14,833 possible R16 matchings, where each same-nation exclusion further reduces the number of valid assignments in Γ_H.[22,23] Across the 19 seasons under consideration, the number of valid assignments ranged from 2,002 in the 2020 season to 6,304 in 2011 to 9,200 in 2006. 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 an exponential function, where each additional same-nation exclusion in H_A decreases the number of possible matchings by 15 percent.

4. Constraint Effects in the UEFA Draw

This section highlights the effects of the matching constraints on teams' tournament outcomes. Section 4.1 discusses the nature of the distortions generated by the tournament's matching constraints. Section 4.2 defines two measures of the effects from the constraints and quantifies them across the 19 UCL seasons under consideration.

[22]See Table C.2 in Online Appendix C for the number of same-nation exclusions generated by each national association between 2004–22.

[23]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.

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)

4.1. **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.[24]

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 H_G. 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.[25]

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.[26] 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 Paris Saint-Germain 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 the French team plays Chelsea is almost three times larger than the probability they play either Basel, Shaktar Donetsk, or Porto (columns 1, 6, and 7).

[24]We calculate all probabilities with a Monte Carlo simulation of size $N = 10^6$, which results in 95 percent confidence intervals for each probability within ±0.001 of the given coefficient (see Proposition 4 in Online Appendix A).

[25]The expected assignment matrices for the R16 draw in the remaining seasons can be found in Online Appendix D.

[26]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.

4.2. **Quantifying the Association-Constraint Effect.** Below, we further analyze the unequal match chances illustrated in Table 1 by formally quantifying the distortions generated by the association constraint. First, we measure the *monetary effect* of imposing the constraint in terms of expected prize money—though the monetary distortions are not easily interpretable as fairness distortions, since the size of the monetary effect vary with the teams' underlying ability. Then, we quantify the *distortive effect* of the constraint by focusing on the difference in match chances for teams with a common partner, providing a better proxy for the observed inequality, as the association constraint is agnostic as to the teams' identity.

Focusing on the 19 UCL seasons in 2004–22 we find that:

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, the monetary effect; and (ii) creating large inequalities in the match chances for otherwise comparable teams, the distortive effect.*

Evidence for Result 1 (i): In order to measure the association-constraint 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 2004–22. Armed with these estimates and data on the team prizes awarded for reaching each stage of the competition,[27] we simulate the tournament outcomes and calculate the *expected* prize for each team i in each season t, under each realized R16 matching. As the parameter estimation is standard in the literature we relegate the details to Online Appendix B.

To estimate the monetary effect of the association constraint we calculate differences in the expected prize money under the current UEFA draw (input set $\Gamma_{H_A \cup H_G}$) and a counterfactual procedure that drops the association constraint (input set Γ_{H_G}).[28] By construction, teams with a positive association-constraint effect are those benefiting from the constraint, whereas those with a negative value are being disadvantaged. Across all 19 UCL seasons, the association-constraint effect has a standard deviation of 0.3 million euro (it

[27]We use tournament prize amounts from the 2019 season: from a minimum of 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.

[28]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.

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 in the 2014 R16) and a subsidy of 1.2 million euro to Real Madrid in the 2017 season (the 2017 tournament champion).

In Figure 3(A) we illustrate the range in the association-constraint effect (defined as 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. The illustrated relationship indicates that the association-constraint effect increases in the number of same-nation exclusions. In particular, we find that ten association exclusions lead to an expected range of approximately 1.5 million euro.

Evidence for Result 1 (ii): We next quantify the constraint effect on the chances of each match pair and specifically, on the inequality over team's treatment in an ex-ante sense. Our fairness objective measures the average absolute difference in the 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 $\mathbf{A}$, where a_{ik} indicates the probability that ik is selected, we measure the distortion between teams i and j as $\left|a_{ik} - a_{ij}\right|$. Then, taking averages across all possible comparisons, we define our fairness distortion 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:

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

For example, in a hypothetical season with no same-nation exclusions within the association constraint the distortion measure Q has a minimum at zero (a fully symmetric match chance across the seven unconstrained teams) and a maximum at $2/7$ (a degenerate assignment).

In Figure 3(B), we graph the fairness measure Q for the UEFA randomization's expected assignment matrix in each season in 2004–22 on the vertical axis, against the number of same-nation exclusions on the horizontal. 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 4 percentage points. This represents a large relative swing of approximately one-third compared to a one-in-seven match chance were we to drop the association constraint (and conduct a uniform draw over the seven non-group partners).

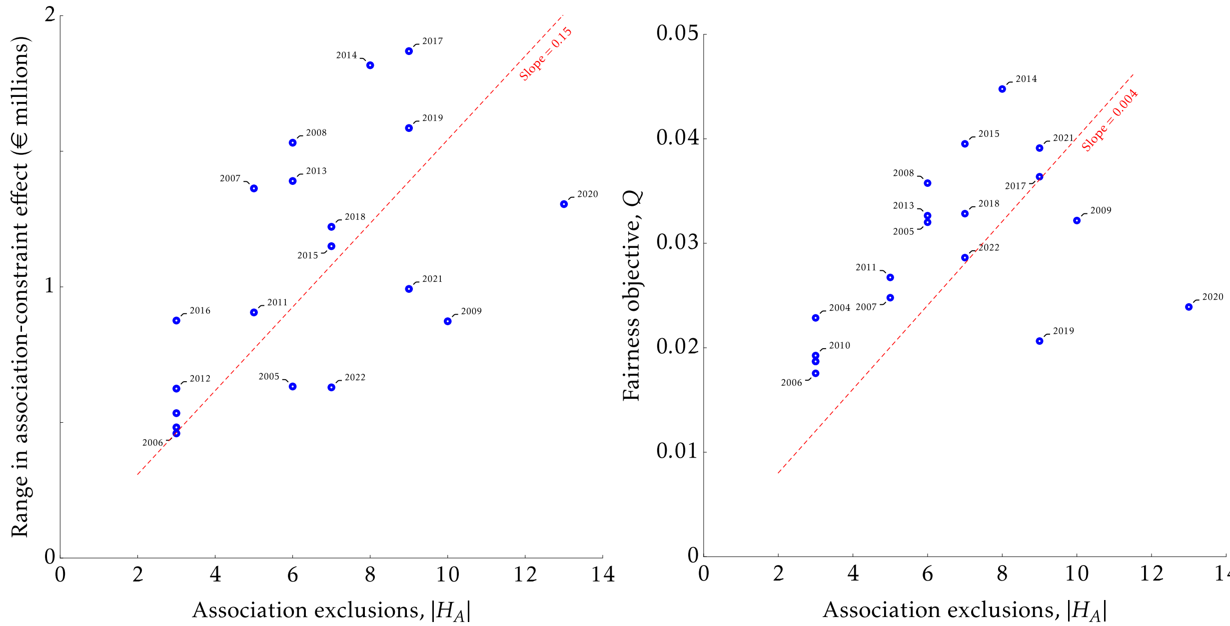

(A) Monetary effect (prize money) (B) Distortive effect (match likelihoods)

FIGURE 3. Effects from imposing the association constraint

Figure details: Red dashed line indicates fitted linear relationship.

While Result 1 points to quantitatively large spillovers from the association constraint, we next show that there is only limited scope to ameliorate these effects through a better randomization procedure.

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. In order to answer this question comprehensively, one would need to optimize over all assignment lotteries, a non-trivial and computationally complex problem, especially given high dimensionality of $\Delta\Gamma_H$. To address this concern we turn to the core result in Budish et al. (2013), which states that as long as the constraint structure can be separated into a bi-hierarchy, any expected assignment matrix satisfying the constraints is implementable as a lottery over constrained assignments.[29]

[29] 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 for team i, and the singleton exclusion ij) or disjoint (for example, the sum-to-one constraint for team k and the singleton 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.

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., it is sufficient to provide any bihierarchy construction over the constraints in H. As such, the UEFA 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, and each of the singleton exclusions in H. (For a formal construction see Online Appendix A.) □

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). Hence, for our specific UEFA application with $K = 8$, Proposition 3 reduces the degrees of freedom by two orders of magnitude. That is, across the 19 UCL seasons between 2004–22 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 2004–22 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 problem:

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

subject to the matching constraints: (i) $\forall ij \in H : a_{ij} = 0$; (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 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 $\hat{\mathbf{A}}_t$ on the horizontal, for each season t between 2004–22. While some improvements are possible across the realized constraints, the gain from a fairness-optimal randomization is marginal. On average, we find that optimization can reduce the fairness distortions by approximately a tenth (which corresponds to the estimated slope in Figure 4 of 0.90).[30,31]

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 randomization, as the draw are most likley assembled over complete matchings, rather than over the component parts.[32,33] Put against the implementation cost of reduced transparency, a reduction in the average match-chance difference from 5 to 4.5 percentage points seems marginal.

Summarizing the section, the thrust of our analysis has shown that despite substantial distortions generated by the constraints, the transparent element-uniform approach to randomizing the R16 assignments between 2004–22 is very close to optimal in fairness terms. Now, we extend this result outward and show that near-optimality of the UEFA randomization continues to hold for a variety of simulated alternatives that differ with respect to the scale of the matching problem, the number of exclusion constraints, and the degree of constraint dispersion across the expected assignment matrix. Hence, we find that for settings where a manager attempts to generate a fair matching between two groups under a series of exclusion constraints, the element uniform procedure is not only transparent to outside observers, it also comes very close to a first-best procedure in fairness terms.

5.2. **Examining Simulated Draws.** Above we demonstrate that the UEFA procedure is close to a constrained-best for all UCL seasons across 2004–22. We now augment that result by demonstrating that the same property holds under an array of simulated alternatives:

[30] Similar qualitative results hold when we repeat the analysis over the following alternative objectives: minimizing the squared 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 focus on the pairwise distortion measure Q, which has the benefit of a simpler interpretation.

[31] To add context to our conclusion that the UEFA procedure comes very close to the constrained best, we analyze the performance of a reasonable modification of the current UEFA procedure. Instead of uniformly drawing a team at step k(i), we fix the ordering from $\mathcal{R}$—matching the most-constrained teams first—but continue to uniformly draw feasible match partners at step k(ii). In Figure C.1 in Online Appendix C we illustrate that this simple modification aimed at matching the most-constrained teams first yields a loss in efficiency four times larger than for the current UEFA algorithm.

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

[33] 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.2–C.4 in Online Appendix C).

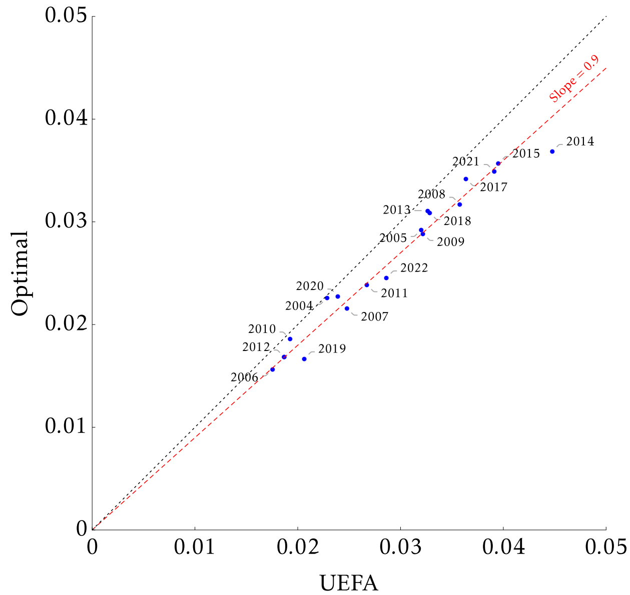

FIGURE 4. Comparison of fairness distortion measure Q: optimal vs. current UEFA procedure

Figure details: Red dashed line indicates fitted linear relationship.

Result 3. *The element-uniform randomization for a perfect one-to-one matching with direct exclusions H continues to be close to a constrained-best as we shift: (i) the number of constraints $|H|$, (ii) the likely location of the constraints within the expected assignment matrix, and (iii) the underlying dimension of the assignment problem K.*

Evidence for Result 3 (i)-(ii): 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 from using the element-uniform draw as opposed to a fairness-optimized randomization via

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

where $\overline{Q}_{\emptyset} = {}^{2}/_{K}$ denotes the maximum value of Q for a degenerate assignment under neither group nor association constraint.

In our simulations, we consider constraint structures where the bipartite and group constraints are fixed, but where we vary the number and arrangement of association exclusions within the association constraint. In particular, we vary: (i) the number of association exclusions, from $|H_A| = 5$ to $|H_A| = 20$ in unit increments; and (ii) the distribution

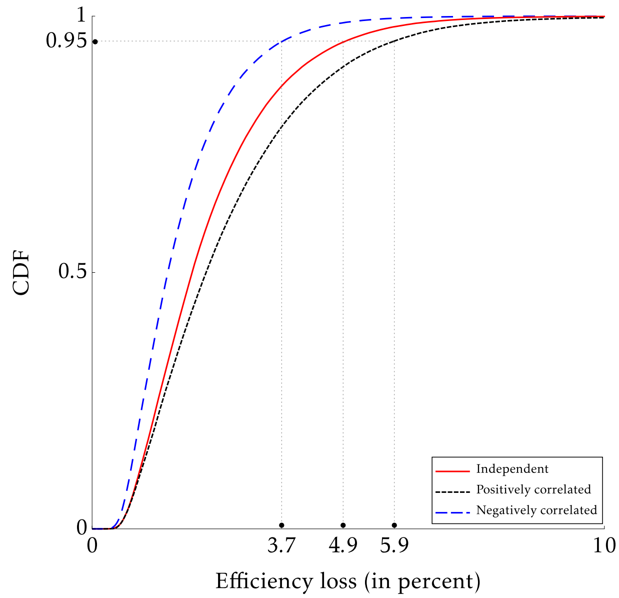

FIGURE 5. Efficiency loss CDFs across simulations

of these exclusions across the assignment matrix (i.e. the relative chances of multiple exclusions in the same row or column).[34]

Figure 5 illustrates the empirical CDF for the efficiency loss ϕ at $K = 8$. The results are pooled across values of $|H_A|$—as we do not observe any relative differences in ϕ across $|H_A|$ (per Figure 4 the effect is proportional)—but broken out by the correlation of the constraints within the matrix, where the figure indicates a clear stochastic dominance relationship. The simulations suggest that the estimated efficiency loss of the element-uniform draw is largest when exclusions are more likely to fall in the same row or column—the case for the UCL R16 assignment problem given that many of the constraints stem from a relatively small number of associations—rather than spread more evenly across the expected assignment matrix. However, while the inefficiency of the element-uniform randomization does increase when the constraint locations are interrelated, the quantitative level of the effect remains small. Superimposed on Figure 5 we

[34] For each exclusion-number parameter $|H_A|$ we generate 30,000 constraint structures: one-third using conditional independence across each sequentially drawn exclusion; another third under a positive correlation, making sequentially drawn exclusions in the same row or column more likely; and the final third under negative correlation, with subsequently drawn exclusions in the same row or column less likely. For illustration, consider an eight-by-eight assignment problem with $|H_A| = 5$. *Conditional independence*: In order to generate a structure comprised of five conditionally independent constraints, we sequentially draw five pairs ij without replacement from the set of non-group pairs $\mathcal{U} = \{(ij) : i, j = 1, \ldots, K, i \neq j\}$. *Positive (negative) dependence*: We sequentially draw five positively (negatively) correlated pairs, but where conditional on a draw (i^*j^*) we assign higher (smaller) sampling weights to pairs in the same row or column, making it more (less) likely that (i^*j) or (ji^*) is chosen. For full details (and precise sampling weights) see Online Appendix E.

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).

Evidence for Result 3 (iii): In addition to checking for near-optimality of the UEFA procedure at $K = 8$, we conduct simulations for $K = 6$ and $K = 7$. Again, using randomly generated constraint structures (where $|H_A|$ varies from $|H_A| = 5$ to $|H_A| = 20$, and where exclusion locations are sequentially independent draws) we find that the UEFA assignment rule continues to be close to a constrained-best.[35] Using a linear regression model to examine how ϕ responds to changes in K we find that the efficiency loss decreases by approximately 1.35 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.[36]

6. Discussion: Beyond the UEFA Application

Above we have analyzed the properties of the dynamic element-uniform randomization procedure used within the UEFA tournament setting, a matching environment where (A) the sought matching is a perfect one-to-one assignment; (B) the constraints on the process are direct exclusions; and (C) the constraints satisfy the bihierarchy condition from the main result in Budish et al. (2013). In this section we examine how managers might adapt the UEFA randomization together with the illustrated methods of analysis to more-general settings by relaxing properties (A) through (C) .

We begin our analysis by generalizing the element-uniform randomization procedure to a many-to-many assignment problem under an arbitrary constraint structure. Next, we outline a first extension, where we soften the direct exclusion constraints used in the UEFA application (dropping properties B and C). By allowing greater flexibility in the constraints we find that the element-uniform randomization is no longer close to optimal. To show that this is not necessarily driven by failure of the bihierarchy condition, we then outline a second extension for a many-to-many assignment that allows the constraint structure to be softer while still satisfying the bihierarchy condition (dropping properties A and B, but keeping C). Again, we find that alternative randomizations can substantially improve fairness, concluding that near-optimality of the UEFA implementation seems to hang upon property B, that the imposed constraints are direct exclusions. While both extensions suggest a more negative result for the element-uniform randomization—as softer constraints temper managers' ability to use an element-uniform randomization as

[35] See Online Appendix E for more details and comparisons at differing K under the independently drawn same-nation exclusions.

[36] See Table C.3 in Online Appendix C for the estimation results.

an out-of-the-box design—the analysis continues to demonstrate the usefulness of the techniques and methods developed. Moreover, it suggests that if problems with the element-uniform procedure stem from too many degrees of freedom in the constraints, there exists an easy and fruitful solution: adding further constraints.

6.1. **Theory for the many-to-many element-uniform draw.** In the generalized many-to-many randomization procedure we retain a bipartite vertex structure, where a set of winners $\mathcal{W} = \{w_1, \ldots, w_L\}$ must be assigned to a set of runners-up $\mathcal{R} = \{r_1, \ldots, r_K\}$. However, the universe of possible assignments is now far less constrained and given by $\mathcal{V} = 2^{\mathcal{R}\times\mathcal{W}}$, where each entry is a subset of the possible undirected edges between $\mathcal{R}$ and $\mathcal{W}$. Moreover, note that rather than the one-to-one restriction in the UEFA application, where each entry in $\mathcal{R}\cup\mathcal{W}$ acts as a vertex for exactly one edge (a perfect one-to-one match), the generalization allows each vertex to be a part of many separate edges.

Mirroring the definitions in Budish et al. (2013), constraints on the assignment are specified via a collection of edge sets $\mathcal{H} = \{H_1, \ldots, H_N\}$, where each $H_j \subseteq \mathcal{R}\times\mathcal{W}$ has a corresponding lower and upper-bound on the permissible edge-count in any assignment. Feasibility of an assignment $V \in \mathcal{V}$ under the constraint structure $(\mathcal{H}, \mathbf{q}) := \left(H, \left[\underline{q}_H, \overline{q}_H\right]\right)_{H\in\mathcal{H}}$ therefore requires that $\underline{q}_H \leq |H \cap V| \leq \overline{q}_H$ for all $H \in \mathcal{H}$ and so, the feasible assignment set is given by

$$\Gamma(\mathcal{H}, \mathbf{q}) = \left\{ V \in \mathcal{V} \middle| \underline{q}_H \leq |H \cap V| \leq \overline{q}_H \text{ for all } H \in \mathcal{H} \right\}.$$

For any input set of feasible assignments, the transparent element-uniform randomization can be generalized to the many-to-many setting as follows:

Algorithm (Many-to-many Γ-constrained element-uniform draw)**.** *Given an input set of admissible matchings $\Gamma \subseteq \mathcal{V}$, the algorithm selects a matching $\psi(\Gamma)$ in a finite number of steps ($K \leq |\cup_{V\in\mathcal{V}} V|$), where at each step k an edge $r_k w_k \in \mathcal{R}\times\mathcal{W}$ is selected via two sequential uniform draws.*

Step 0 *(Initialization): Set $V_0 = \emptyset$ and $\Gamma_0 = \Gamma$.*

Step k *(draw selected edge $k = 1, 2, 3, \ldots$):*

- *(i) Select r_k: uniform draw over feasible set $\{r \in \mathcal{R} \,|\, \exists V \in \Gamma_{k-1}, w \in \mathcal{W} \text{ s.t. } rw \in V \setminus V_{k-1}\}$;*
- *(ii) Select w_k: uniform draw over feasible set $\{w \in \mathcal{W} \,|\, \exists V \in \Gamma_{k-1} \text{ s.t. } r_k w \in V \setminus V_{k-1}\}$;*
- *(iii) Update the set of edges $V_k := V_{k-1} \cup \{r_k w_k\}$, and pare-down the set of feasible assignments to $\Gamma_k := \{V \in \Gamma_{k-1} \,|\, V_k \subseteq V\}$.*
- *(iv) Termination check:*
  - *If $|\Gamma_k| = 1$, stop and output the unique entry $V \in \Gamma_k$ as the final realization;*
  - *Else continue to step $k+1$.*

The above randomization nests the element-uniform draw used in the UEFA application,[37] maintaining transparency in each randomization, steps k(i) and k(ii), through simple uniform draws from a comparatively small set of options.[38]

Specifically, in the UEFA application the input set is generated under two types of constraint: (i) Perfect one-to-one assignment, where for each $w \in \mathcal{W}$ a column constraint implies a unit-assignment constraint $\underline{q}_{H_w} = \overline{q}_{H_w} = 1$ for the set $H_w = \mathcal{R} \times \{w\}$, with a similar constraint for each $r \in \mathcal{R}$ for a unit assignment on the row $H_r = \{r\} \times \mathcal{R}$. (ii) Direct exclusions over the group and association constraints, where each edge $rw \in H_A \cup H_G$ is directly excluded (with $\underline{q}_{H_A} = \overline{q}_{H_A} = \underline{q}_{H_G} = \overline{q}_{H_G} = 0$).[39] The full constraint structure for the UEFA application is therefore given by:

$$\mathcal{H}_{\text{R16}} = \left\{H_{r_1}, \ldots, H_{r_K}, H_{w_1}, \ldots, H_{w_K}, H_G, H_A\right\},$$

with corresponding lower/upper quota constraints of

$$\mathbf{q}_{\text{R16}} = \left([1,1]_{H_{r_1}}, \ldots, [1,1]_{H_{r_k}}, [1,1]_{H_{w_1}}, \ldots, [1,1]_{H_{w_k}}, [0,0]_{H_G}, [0,0]_{H_A}\right).$$

The constraints in the UEFA application are hard—that is, they allow for no flexibility, imposing equality conditions over the lower and upper quotas ($\underline{q}_H = \overline{q}_H$ for all $H \in \mathcal{H}$). In the two extensions we consider we focus on environments with a "softer" constraint structure, where $\underline{q}_H < \overline{q}_H$ for some $H \in \mathcal{H}$.

6.2. **Weakening the constraints within the UEFA application.** Our paper's main analysis shows that while the association constraint imposed by UEFA generates substantial fairness distortions, the scope for reducing them through a better randomization design is limited. That is, most of the inequality in treatment is an unavoidable consequence of the imposed constraints, and cannot be designed away. Hence, a natural line of inquiry for a designer is to quantify the potential gains in fairness terms from weakening the imposed constraints. This creates an additional motive for our first softer constraint extension.

In order to assess the effect from weakening the association constraint, while maintaining the element-uniform randomization (and its transparency), we allow for *at most one* same-association pair. Formally, this relaxation of the association constraint implies the weakened constraint structure $\left(\mathcal{H}_{\text{R16}}, \mathbf{q}'_{\text{R16}}\right)$ in which the modified quota $\mathbf{q}'_{\text{R16}}$ is identical

[37]The changes here are to the update steps and finalization rules. Note that this definition assumes that some of the upper-bound quotas have binding effects on the input set Γ, as the termination check looks for a maximal-edge assignment under the constraints. If the bipartite graph $\overline{V}$ with all possible edges is feasible, it represents the final outcome from the procedure with certainty.

[38]The urn composition calculations in step k(iii) are more complicated, though per the UEFA draw this step is still deterministic, and thus verifiable.

[39]The definition of $\mathcal{V}$ imposes a symmetric implicit constraint on every possible singleton edge $\{rw\}$ a quota of $\underline{q}_{\{rw\}} = 0$ and $\overline{q}_{\{rw\}} = 1$.

to the original $\mathbf{q}_{\mathrm{R16}}$ except for the upper-bound on the association set H_A, which we now set to $\overline{q}'_{H_A} = 1$:

$$\mathbf{q}'_{\mathrm{R16}} = \Big([1,1]_{H_{r_1}}, \ldots, [1,1]_{H_{r_k}}, [1,1]_{H_{w_1}}, \ldots, [1,1]_{H_{w_k}}, [0,0]_{H_G}, [0,1]_{H_A}\Big).$$

Weakening the association constraint in this way strictly expands the input set of assignments with $\Gamma(\mathcal{H}_{\mathrm{R16}}, \mathbf{q}_{\mathrm{R16}}) \subset \Gamma(\mathcal{H}_{\mathrm{R16}}, \mathbf{q}'_{\mathrm{R16}})$ whenever H_A has one or more excluded pairs.

Fixing the element-uniform draw, the analysis of the effects from weakening the association constraint is straightforward: we can simply compare the outcomes under the algorithm with the two different feasible assignment inputs. However, while the comparison within the element-uniform is simple, understanding the effect from weakening the constraints on an *optimal* randomization over $\Delta\Gamma(\mathcal{H}_{\mathrm{R16}}, \mathbf{q}'_{\mathrm{R16}})$ raises a non-trivial concern. While it is easy to optimize over expected-assignment matrices satisfying the relaxed constraint,[40] the problem now violates the bihierarchy condition. From Theorem 2 in Budish et al. (2013) we know that an expected assignment matrix satisfying the constraints exists for this problem that is *not* implementable as a randomization over the feasible assignments. As such, an optimized expected assignment matrix $\mathbf{A}^\star$ only provides a lower-bound $Q(\mathbf{A}^\star)$ on what is possible when optimizing over $\Delta\Gamma(\mathcal{H}_{\mathrm{R16}}, \mathbf{q}'_{\mathrm{R16}})$. Fortunately, as we discuss below, for the UEFA setting the process is still informative, as the bound is attained for the cases where the optimal expected assignment $\mathbf{A}^\star$ indicates a boundary solution.

Analyzing the UEFA application with a softer association constraint we find that:

Result 4. *Softening the association constraint to allow at most one same-association match within the element-uniform randomization generates a quantitatively large* reduction *in the fairness distortions. However, the element-uniform procedure is no longer close to a constrained-best in relative terms. In particular, in cases with 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 our fairness distortion measure for the element-uniform draw under the expanded set of feasible assignments $\Gamma(\mathcal{H}_{\mathrm{R16}}, \mathbf{q}'_{\mathrm{R16}})$ on the vertical axis, against the UEFA procedure under the actual constraints $\Gamma(\mathcal{H}_{\mathrm{R16}}, \mathbf{q}_{\mathrm{R16}})$ on the horizontal.[41] The illustrated results indicate that weakening the constraint to allow a single same-association match in the R16 decreases the fairness distortions by 70

[40] An expected assignment $A \in [0,1]^{K\times L}$ satisfies the constraint structure $(\mathcal{H}, \mathbf{q})$ if $\underline{q}_H \leq \sum_{rw\in H.} a_{rw} \leq \overline{q}_H$ for all $H \in \mathcal{H}$.

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

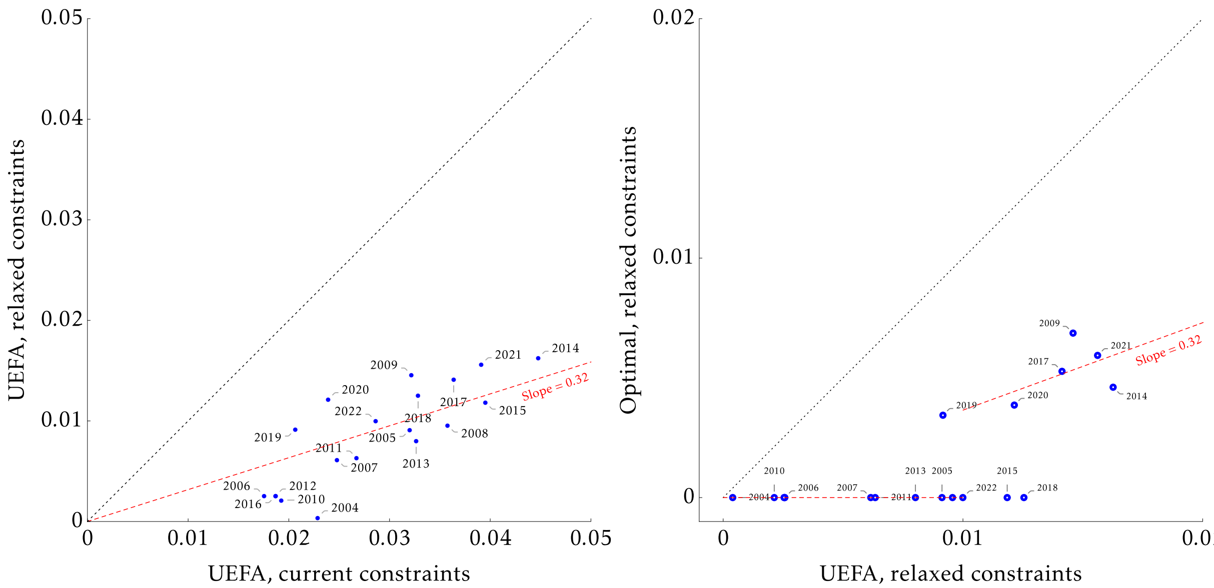

(A) Effect of weakened constraint (element-uniform)

(B) Optimal versus element-uniform (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 $|H_A| \leq 7$.

percent. This is a sizable reduction, especially when compared to the 10 percent reduction from optimal randomization over the original feasible set $\Gamma(\mathcal{H}_{\text{R16}}, \mathbf{q}_{\text{R16}})$.

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 distortion measure $Q(\mathbf{A})$ for an optimal expected assignment satisfying the weakened constraint on the vertical axis, against the element-uniform procedure under the weakened constraint on the horizontal.[42] While both plots indicate substantially greater fairness when the association constraint is weakened, in many years the optimal expected assignment matrix $\mathbf{A}_t^\star$ obtains a first-best outcome with $Q(\mathbf{A}_t^\star) = 0$ while the element-uniform still exhibits distortion. Fairness-optimized expected assignments satisfying the relaxed constraints achieve *perfect* equity in thirteen of the nineteen seasons examined, where in the remaining six, it reduces the admittedly smaller distortions produced by the element-uniform procedure by approximately 70 percent.

[42]Readers' attention is drawn to the different scales for the axes in Figures 6(A) and 6(B).

A concern over the optimal expected assignment matrices is that universal implementability no longer holds as the weakened constraint violates the bihierarchy condition.[43] As such, the optimization results in Figure 6(B) can only be viewed as lower bounds on the fairness distortions possible given a lottery over $\Gamma(\mathcal{H}_{\mathrm{R16}}, \mathbf{q}'_{\mathrm{R16}})$. As bounds go, those coinciding with a boundary on the range of the objective function are naturally of questionable content. However, violation of the bihierarchy condition does not imply that any particular expected assignment is not implementable, just that such is not guaranteed: Focusing on the seasons with a boundary solution for the optimal expected assignment in Figure 6(B) we construct an implementation resulting in a perfectly fair randomization, a one-in-seven chance of matching with each of the non-group partners. Particualrly, we create a partition of the 56 feasible (non-group) pairs into seven disjoint matchings in $\Gamma(\mathcal{H}_{\mathrm{R16}}, \mathbf{q}'_{\mathrm{R16}})$, and then run an equal-chance lottery over the seven disjoint matchings, implementing the optimal expected assignment in Figure 6(B). Such a construction is possible for each of the boundary solutions, and so a perfectly fair optimal randomization *is* possible, demonstrating a large relative improvement over the element-uniform procedure under the weakened constraints.

6.2.1. *Many-to-many example.* Above we have demonstrated that the element-uniform procedure is no longer close to optimal for the one-to-one UEFA application with softer constraints that violate the bihierarchy condition. In this section we show that:

Result 5. *Violation of the bihierarchy condition on the constraint structure does not drive the element-uniform's failure; instead, it is driven by the procedure's incapacity to optimize over the additional degrees of freedom provided by the softer constraint.*

Evidence for Result 5.: We arrive at Result 5 by constructing an example of a many-to-many assignment under a softer constraint structure that *does* satisfy the bihierarchy condition. To frame our many-to-many problem we consider a non-trivial randomization design for an assignment of faculty members (a set of workers $\mathcal{W}$) to committees (a set of roles $\mathcal{R}$). Our example involves eight faculty members: five seniors (S_1 through S_5, where S_1 is the department's chair), and three identical juniors (J_1 through J_3). These eight workers are to be matched to three committees: Cmte. X, Cmte. Y, and Cmte. Z. Each committee needs to be composed of exactly three faculty, where we assume the department's objectives for the randomization are to have: (i) a fair expected division of workload across faculty of the same rank and (ii) equitable chances of assignment to each committee within the randomization.

[43] This is true for all seasons except for 2004, where all of the same-nation restrictions are imposed on a single team. In this case while the bihierarchy condition *is* satisfied, the element-uniform procedure also produces the first-best result as weakening the constraint effectively removes the association constraint.

Complicating the assignment, the department wishes to impose a series of plausible constraints: (i) Steering committee X requires exactly two seniors and one junior, and cannot include the department's chair S_1. (ii) Hiring committee Y requires one senior to chair it. (iii) Tenure committee Z requires all three members to be seniors. Beyond the committee composition constraints, the department imposes constraints aimed at minimizing ex-post differences in workload: (iv) each junior can serve on at most one committee, where each senior must serve on at least one and at most two committees (where we further restrict the department's 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 on the same committee. This problem's expected assignment matrix is illustrated in Figure 7(A), where each entry p^r_w denotes the probability of faculty w being assigned to committee r, and each block of entries represents a constraint set H.[44,45]

The example assignment problem described above and illustrated in Figure 7(A) induces a constraint structure $(\mathcal{H}, \mathbf{q})$ and therefore, a set of feasible assignments $\Gamma(\mathcal{H}, \mathbf{q})$. Using $\Gamma(\mathcal{H}, \mathbf{q})$ as the input we can randomize the committee assignments using the generalized element-uniform randomization characterized in Section 6.1. In Figure 7(B) we illustrate the resulting expected assignment matrix, where the assignment is assembled in sequence, selecting a committee uniformly from those with slots left to fill, and then choosing uniformly among the feasible faculty for that committee slot (given the prior selections and constraints). As in our field application, this randomization can be conducted with a simple physical draw in a public and transparent manner. A computerized aid would still be required to quickly sift through the 576 feasible assignments to determine the urn compositions at each step, though per the UEFA draw this component is deterministic and fully verifiable.

[44] The constraint structure $(\mathcal{H}, \mathbf{q})$ can be written over the following set/quota pairs $\left(H, \left[\underline{q}_H, \overline{q}_H\right]\right)$:

(i) Faculty assignment quotas, the vertical sets in Figure 6 (A), are: the chair S_1 $(\{YS_1, ZS_1\}, [1,1])$; the four other seniors $S \in \{S_2, S_3, S_4, S_5\}$ with $(\{XS, YS, ZS\}, [1,2])$ for each; the three juniors $J \in \{J_1, J_2, J_3\}$ with $(\{XJ, YJ\}, [0,1])$ for each.

(ii) Committee composition constraints, horizontal black-bordered sets in the figure, are: Cmte. X with 2 seniors $(\{XS_2, XS_3, XS_4, XS_5\}, [2,2])$ and 1 junior $(\{XJ_1, XJ_2, XJ_3\}, [1,1])$; Cmte. Y with at least one senior $(\{YS_1, YS_2, YS_3, YS_4, YS_5\}, [1,3])$ and three total $(\{YS_1, YS_2, YS_3, YS_4, YS_5, YJ_1, YJ_2, YJ_3\}, [3,3])$; and Cmte. Z with three seniors $(\{ZS_1, ZS_2, ZS_3, ZS_4, ZS_5\}, [3,3])$. The committee composition constraints also imply four singleton exclusions $rw \in \{XS_1, ZJ_1, ZJ_2, ZJ_3\}$ with $(\{rw\}, [0,0])$ for each.

(iii) The minimizing acrimony constraints, horizontal red-bordered sets, for each committee $R \in \{X, Y, Z\}$ we have $(\{RS_2, RS_3\}, [0,1])$;

An expected assignment satisfying the constraints therefore requires $p^r_w \in [0,1]$ for all possible edges rw, and $\sum_{wr \in H} p^r_w \in \left[\underline{q}_H, \overline{q}_H\right]$ for each $H \in \mathcal{H}$.

[45] Operationally, while a simpler example can be constructed, our intention with the committee assignment example is to provide a plausible assignment problem that is easy to describe in words but also provides a non-trivial example for the combinatorics involved.

$$
\begin{array}{r}
\\
\text{Cmte. } X \\
\text{Cmte. } Y \\
\text{Cmte. } Z
\end{array}
\begin{array}{c}
\begin{array}{cccccccc}
S_1 & S_2 & S_3 & S_4 & S_5 & J_1 & J_2 & J_3
\end{array}
\\
\begin{bmatrix}
0 & p^X_{S_2} & p^X_{S_3} & p^X_{S_4} & p^X_{S_5} & p^X_{J_1} & p^X_{J_2} & p^X_{J_3} \\
p^Y_{S_1} & p^Y_{S_2} & p^Y_{S_3} & p^Y_{S_4} & p^Y_{S_5} & p^Y_{J_1} & p^Y_{J_2} & p^Y_{J_3} \\
p^Z_{S_1} & p^Z_{S_2} & p^Z_{S_3} & p^Z_{S_4} & p^Z_{S_5} & 0 & 0 & 0
\end{bmatrix}
\end{array}
$$

(A) Generic form of the expected assignment

	S_1	S_2	S_3	S_4	S_5	J_1	J_2	J_3
Cmte. X	0	0.45	0.45	0.55	0.55	0.33	0.34	0.33
Cmte. Y	0.35	0.41	0.40	0.39	0.39	0.35	0.35	0.36
Cmte. Z	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. X	0	1/2	1/2	1/2	1/2	1/3	1/3	1/3
Cmte. Y	0	1/2	1/2	1/2	1/2	1/3	1/3	1/3
Cmte. Z	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) A feasible expected assignment

Figure 7. Committee examples

The expected assignment in Figure 7(B) suggests that the department's fairness objectives are largely met with the many-to-many element-uniform draw. However, one legitimate complaint from inspecting the expected assignment is that seniors S_4 and S_5 are negatively affected by the constraint imposed to minimize acrimony between seniors

S_2 and S_3. Specifically, S_4 and S_5, the senior faculty with the fewest constraints, are almost twice as likely to have two committee assignments as either S_2 or S_3. This begs the question: does a fairer randomization exist?

By construction, the example's constraints are chosen to satisfy the bihierarchy condition from Budish et al. As such, the search for fairer randomizations can be conducted over the feasible expected assignments given a well-defined objective to optimize over.[46] However, to make our point it is sufficient to point to a particular feasible expected assignment that equates both the expected workloads and the chances of particular assignment among the comparable faculty: the feasible expected assignment illustrated in Figure 7(C).

Consequently, our many-to-many element-uniform draw under softer constraints does not guarantee the fairest possible outcome, even though the constraint structure satisfies the bihierarchy condition. The reasoning for this failure is that the element-uniform randomization procedure does not have a channel through which to optimize over the softer constraints. However, this points to the possibility of an easy solution for the designer to arrive at a fairer assignment within the element-uniform paradigm: adding further constraints to remove the degrees of freedom. In particular, the expected assignment illustrated in Figure 7(C) *is* implementable under the element-uniform randomization under a stricter set of constraints. For example, if we impose the additional constraints that S_1 serves on committee Z with certainty and that the faculty pairs (S_2, S_3) and (S_4, S_5) are both given exactly three assignments in total on top of the prior constraints, the expected assignment under the element-uniform procedure is now exactly the expected assignment shown in Figure 7(C). As such, if transparency within the randomization is important, the element-uniform procedure offers a practical constructive framework to design within. Hence, with many additional degrees of freedom available, a manager can attempt to design better randomizations by imposing further constraints.

Going beyond the specifics of our two extensions, this discussion demonstrates that a generalized version of the simple-to-follow element-uniform draw used by UEFA offers a constructive solution to randomized assignment problems with non-trivial constraints. In situations where equity concerns are paramount, the draw offers a credible and transparent randomization, providing the various interested parties with a better understanding of their equal treatment. While near-optimality of the element-uniform draw is not guaranteed under softer constraints, the tools showcased here offer a number of paths forward for designing substantially fairer outcomes.

[46] 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_a^b in either.

7. 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, 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 steps) 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-assist is used to deterministically enforce the imposed constraints.

We demonstrate 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 under direct exclusions (per 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.

In two extensions, we show that near-optimality of the transparent randomization can fail in alternative settings, where better designed solutions are possible. However, the developed methodology makes possible both the detection of these alternative designs, and constructive alternatives that retain transparency. 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 procedure, 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*, 2020, *88* (2), 425–467.

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.

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 Volume II*, number 28. In ‘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.

Pathak, Parag A, Tayfun Sönmez, M Utku Ünver, and M Bumin Yenmez, “Fair allocation of vaccines, ventilators and antiviral treatments: leaving no ethical value behind in health care rationing,” in “Proceedings of the 22nd ACM Conference on Economics and Computation” 2021, pp. 785–786.

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.

Rubin, Emily, Scott L Dryden-Peterson, Sarah P Hammond, Inga Lennes, Alyssa R Letourneau, Parag Pathak, Tayfun Sonmez, and M Utku Ünver, “A novel approach to equitable distribution of scarce therapeutics: institutional experience implementing a reserve system for allocation of Covid-19 monoclonal antibodies,” *Chest*, 2021, *160* (6), 2324–2331.

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.