

Chapter 10

Fairness Trade-offs in Sports Timetabling

Dries Goossens*, Xiajie Yi†, David Van Bulck‡

Abstract

Any sports tournament needs a timetable, specifying which teams will meet at what time, and where. However, not all sport timetables are equally fair for the contestants. In this chapter, we discuss three fairness issues, namely consecutive home games, the carry-over effect (which relates to the opponent's previous game), and the number of rest days each team has between consecutive games. Since we typically cannot obtain a timetable that scores well on all these issues, we study how to make a good trade-off. Furthermore, we look at the trade-off between a timetable that is as fair as possible for the league overall, versus a timetable that equitably splits its unfair aspects over the teams. We verify how a number of official timetables from major European football competitions score with respect to fairness criteria. Finally, we generate timetables for an amateur indoor football competition that reconcile overall fairness with an equitable distribution of unfairness over the teams.

*Corresponding author. Ghent University, Faculty of Economics and Business Administration, Department of Business Informatics and Operations Management, Tweeckerkenstraat 2, 9000 Gent, Belgium. Email: dries.goossens@ugent.be

†Ghent University, Faculty of Economics and Business Administration, Department of Business Informatics and Operations Management

‡Ghent University, Faculty of Economics and Business Administration, Department of Business Informatics and Operations Management

1 Introduction

Every sports competition needs a timetable, stating when and where each match of the tournament will be played. In professional sports, the timetable has high economical stakes, because it has an impact on commercial interests and revenues of clubs, broadcasters, sponsors, etc. In amateur sports, practical concerns like venue and team availability are more prominent. However, in all sports, a fair timetable is paramount, in the sense that it should not a priori give an advantage or a disadvantage to a particular contestant. Dealing with fairness, given the many stakeholders and the diversity of their (often conflicting) requirements, makes sports timetabling a very challenging optimization problem (Van Bulck et al., 2020). We explain a number of essential concepts in sports timetabling in Section 2.1.

In this chapter, we focus on so-called *round robin tournaments*. These are tournaments where each team meets each other team a fixed number of times, as opposed to e.g. knock-out tournaments (used in e.g. Grand Slam tennis tournaments), where contestants do not face each opponent. The double round robin tournament, where each team faces each other team twice (typically once in its home venue and once in the venue of the opponent) is a particularly popular format; it is omnipresent e.g. in football (Goossens and Spieksma, 2012b). While this tournament design alone already adds a substantial level of fairness to any timetable designed for it, there are many other fairness issues to deal with. In Section 2.2, we discuss breaks, the carry-over effect, and rest times as our main fairness criteria.

Unfortunately, there ain't no such thing as a free lunch (Friedman, 1975). Indeed, in sports timetabling, even though we prefer a timetable that is fair in every possible way, we are facing a trade-off between various fairness criteria. How much of one fairness criterion do we have to give up in order to make the timetable better with respect to some other fairness issue? Moreover, even if there was only one fairness criterion to deal with, the question still remains how to fairly distribute the undesirable properties of the timetable over the teams. We explore some basic concepts on these fairness trade-offs in Section 2.3.

For academics, we detail on methods to minimize the number of breaks and the carry-over effect in Section 3. We develop a so-called first-carry-over, then-break approach in order to obtain a set of timeta-

This is an author-version, published as Chapter 10 in “Science Meets Sports: 2 when statistics are more than numbers” (Cambridge Scholars Publishing).

bles that provide insight in this trade-off. Furthermore, we study how we can optimize the rest times of a timetable, without overlooking the distribution of undesirable rest times over the teams. To this extent, we develop a mathematical formulation as well as a bi-criteria evolutionary algorithm.

For practitioners, we apply our ideas to two real-life settings in Section 4. We study the official timetables of ten main professional European football leagues, in order to understand how the breaks versus carry-over trade-off is made in practice. Moreover, we create a set of timetables for an amateur indoor football league, which illustrates the price of fairness: by how much do the overall rest times decrease if we wish to balance short rest times over the teams?

2 Preliminaries

2.1 Sports Timetabling

A game (or match) between team i and j , denoted $i - j$, means that team i plays at home, i.e., it uses its own venue (stadium) for that game, against away team j . A round is a set of games, often played on the same weekend, in which every team plays at most one game. A timetable essentially assigns a round to each match.

A timetable is said to be *time-constrained* if it uses the minimum number of rounds required to schedule all the games. In this chapter, we assume that time-constrained competitions have an even number of teams. Hence, each team plays exactly one game in each round in a time-constrained timetable. In contrast, *time-relaxed* timetables utilize (many) more rounds than there are games per team. If a team doesn't play in a round, it is said to have a *bye* in that round. Table 1 gives an example of a time-constrained timetable for eight teams (named **A** to **H**); a time-relaxed timetable for the same tournament is given in Table 2.

Time-constrained timetables are common practice in e.g. professional football leagues in Europe (Goossens and Spieksma, 2012b). However, there are a number of reasons why organizers may opt for a time-relaxed timetable. The main reason is its flexibility to take into account venue or team availability constraints. Besides, it may simply be unattractive or even impractical to play multiple games simultaneously. This is an author-version, published as Chapter 10 in “Science Meets Sports: 3 when statistics are more than numbers” (Cambridge Scholars Publishing).

ously. The most extreme case example is an asynchronous timetable where at most one match takes place in each round (see Suksompong (2016)). Asynchronous tournaments occur when there is only one venue, as is the case in the top-tier national football league of Gibraltar, or when fans need to be able to watch all games live (e.g. the 2012 Premier League Snooker in England).

When the second half of a double round robin timetable is identical to the first half, except that the home advantage is flipped, we say that the timetable is *mirrored*.

Table 1: A time-constrained single round robin timetable with eight teams.

r1	r2	r3	r4	r5	r6	r7
A-H	C-A	A-E	G-A	A-B	D-A	A-F
B-G	H-B	D-B	B-F	C-G	B-C	E-B
F-C	G-D	C-H	E-C	F-D	G-E	C-D
D-E	E-F	F-G	H-D	E-H	H-F	G-H

Table 2: A time-relaxed single round robin timetable with eight teams.

r1	r2	r3	r4	r5	r6	r7	r8	r9	r10
A-H	B-G	C-A	A-E	G-A	E-C	A-B	F-D	D-A	A-F
	F-C	H-B	D-B	B-F	H-D	C-G	E-H	B-C	E-B
	D-E	G-D	C-H					G-E	C-D
		E-F	F-G					H-F	G-H

2.2 Fairness criteria

In this section, we discuss three fairness criteria: breaks, the carry-over effect, and rest times. In our opinion, these criteria are quite universal, in the sense that in most sports, league organizers, teams, and fans would like to see them respected. However, we do not claim that there are no other universal fairness criteria, and there will certainly be relevant sport-specific fairness criteria as well.

2.2.1 Breaks

Determining the venue of games is crucial in terms of the fairness of a timetable. The sequence of home matches ('H') and away matches ('A') played by a single team is called its *home-away pattern* (HAP). Given such a HAP, the occurrence of two consecutive home matches, or two consecutive away matches is called a *break*. Teams can have consecutive breaks, causing them to play three or more home (away) games in a row.

In some cases, away breaks may be beneficial. For instance, to reduce travel costs, a team may prefer to have two (or more) consecutive away games if its stadium is located far from the opponents' venues, and the venues of these opponents are close to each other (provided of course that the teams do not return home after each game). However, in most competitions, breaks - and successive breaks in particular - are avoided as much as possible. Indeed, Forrest and Simmons (2005) observe that scheduling consecutive home games has a negative impact on attendance. Moreover, given the home advantage, a morale boost (blow) after two consecutive home (away) games may have an impact on the outcome of the next game. Hence, breaks are considered unfair.

As an illustration, Table 3a gives the HAPs corresponding to the timetable presented in Table 1. Note that the HAPs of all teams except for team **A** and **H** contain a break. The timetable depicted here contains 6 breaks in total.

2.2.2 The carry-over effect

Any timetable implies an order in which each team meets its opponents. Playing against a strong or a weak opponent has impact on the performance of teams (Briskorn and Knust, 2010). For instance, a team is more likely to be exhausted or demoralized, or to suffer injuries or suspensions from playing against a very strong opponent, which in turn can have a negative impact on this team's performance in its next game. In this way, the opponent of this team in the next game receives an indirect advantage from that strong team. Following this idea, we say that a team i gives a *carry-over effect* to a team j , if some other team t 's game against i is immediately followed by a game of t against j .

Clearly, carry-over effects cannot be avoided as – except in the
This is an author-version, published as Chapter 10 in “Science Meets Sports: 5
when statistics are more than numbers” (Cambridge Scholars Publishing).

first round – teams always have a previous opponent. It is however considered unfair if some team predominantly gives a carry-over effect to the same team. Indeed, Goossens and Spieksma (2012a) mention examples from football in Norway and Belgium, where the carry-over effect was held responsible for determining the league champion and the relegated team respectively (they could however not find any statistical evidence in favor of this claim).

The extent to which carry-over effects are balanced is measured by the so-called *carry-over effects value* (Russell, 1980). The carry-over effects value is defined as $\sum_{i,j} c_{ij}^2$, where c_{ij} corresponds to the number of times that team i gives a carry-over effect to team j in a tournament, i.e., the number of times that team j plays against the opponent of team i in the previous round. Note that according to Russell’s definition, carry-over effects from the last round to the first round are also counted. An illustration of the carry-over effects (c_{ij}) of the timetable in Table 1 is given in Table 3b; in this case, the carry-over effects value equals 196.

Table 3: Illustration of HAPs and c_{ij} values of the timetable in Table 1.

(a) HAPs								(b) c_{ij} values								
	r1	r2	r3	r4	r5	r6	r7		A	B	C	D	E	F	G	H
A	H	A	H	A	H	A	H	A	0	1	5	0	0	0	0	1
B	H	A	A	H	A	H	A	B	0	0	1	5	0	0	0	1
C	A	H	H	A	H	A	H	C	0	0	0	1	5	0	0	1
D	H	A	H	A	A	H	A	D	0	0	0	0	1	5	0	1
E	A	H	A	H	H	A	H	E	0	0	0	0	0	1	5	1
F	H	A	H	A	H	A	A	F	5	0	0	0	0	0	1	1
G	A	H	A	H	A	H	H	G	1	5	0	0	0	0	0	1
H	A	H	A	H	A	H	A	H	1	1	1	1	1	1	1	0

2.2.3 Rest times

Since time-relaxed timetables contain (many) more rounds than games per team, the rest time between consecutive games of a team can vary substantially and a team’s timetable may therefore contain congested periods. From a fairness perspective, congested periods are problematic as they can lead to player injuries (Bengtsson et al., 2013).

For example, when Manchester City had to play four games in 11 days during the Christmas and New Year period in the 2018-2019 Premier League season, their coach Pep Guardiola said “the Premier League fixture list is a disaster for player well being”. West Ham’s head of medical services Gary Lewin added: “I don’t think it is particularly fair – physically it is not a level playing field for all clubs, as some are playing every few days, some have a longer break between games and some have a week off before playing twice in three days”¹.

Suksompong (2016) proposes to measure timetable congestion by the so-called *guaranteed rest time* (GRT), i.e., the number of rounds without a game that any team will at least have between two consecutive games. However, the problem with the GRT is that it only considers the worst-case rest time. Hence, in this chapter, we use the *aggregated rest time penalty* (ARTP) as a fairness measure. This measure penalizes a timetable with a value p_r each time a team has only r rounds between two consecutive games. The idea is of course that this penalty increases as there are fewer rounds in between; on the other hand, when sufficient rounds are in between such that the team is fully recovered, no penalty is incurred (Van Bulck et al., 2019). The ARTP of a timetable is then simply the sum of these penalties.

Table 4: Calculation of the rest time penalties for the timetable in Table 2.

	r1	r2	r3	r4	r5	r6	r7	r8	r9	r10
A			p_1	p_0	p_0		p_1		p_1	p_0
B			p_0	p_0	p_0		p_1		p_1	p_0
C			p_0	p_0		p_1	p_0		p_1	p_0
D			p_0	p_0		p_1		p_1	p_0	p_0
E			p_0	p_0		p_1		p_1	p_0	p_0
F			p_0	p_0	p_0			p_2	p_0	p_0
G			p_0	p_0	p_0		p_1		p_1	p_0
H			p_1	p_0		p_1		p_1	p_0	p_0

As an illustration, Table 4 gives the calculation of the rest time penalties for the timetable depicted in Table 2. If we assume $p_0 = 4$, $p_1 = 2$ and $p_2 = 1$, we have ARTP = 157. Although the GRT of any timetable with eight teams and 10 rounds is always 0, there are

¹CNN Sports, <https://edition.cnn.com/2018/01/02/football/pep-guardiola-english-premier-league-fixture-scheduling/index.html>

several opportunities to optimize the ARTP.

2.3 Fairness trade-offs

Unfortunately, a single timetable that gives the best performance for each fairness criterion does not exist. Consequently, league organizers are facing a trade-off between fairness criteria. One way to organize this trade-off is through *Pareto-efficiency*. A timetable is called Pareto-efficient if it is impossible to improve one fairness criterion without deteriorating at least one other fairness criterion. This concept is named after the Italian economist and engineer Vilfredo Pareto (1848–1923), who applied it in his income distribution studies. Hence, it would be wise to pick a timetable from the set of Pareto-efficient timetables, known as the *Pareto-front*. The choice for a specific timetable on the Pareto-front depends on the relative importance the league organizer attaches to the fairness criteria.

Even if the league organizer manages to strike a balance between multiple fairness criteria, his struggle is not over. Indeed, the previously described approach only considers the overall presence of unfairness, it does not consider how this unfairness is distributed over the teams. Even a timetable with very few undesired properties may be conceived as unfair when only a small number of teams carry the majority of this burden. We assume that each undesired property of a timetable can be linked with a penalty, which can in turn be split over the teams that suffer from it. Hence, we need to reconcile *criterion efficiency*, i.e. minimizing the total penalty of the timetable, with *distribution equity*, namely making sure that the distribution of the penalties is well balanced over the teams.

We aim to generate equitably-efficient timetables. Practically, a timetable is equitably-efficient if it is Pareto-efficient, meaning that we cannot improve a team's timetable without deteriorating the timetable of at least one other team, and it complies with the Pigou-Dalton principle of fairness, meaning that we cannot shift a penalty from a worse-off team to a better-off team without increasing the total penalties (Ogryczak, 1997). Equitable-efficiency can be seen as a refinement of Pareto-efficiency: any equitably-efficient timetable is Pareto-efficient, but not all Pareto-efficient timetables are equitably-efficient.

3 For academics

3.1 Trade-offs between two fairness criteria

In time-constrained timetables for round robin tournaments with an even number of teams, each team will play on each round, and consequently, their rest times will be the same. Hence, in these competitions, we focus on minimizing breaks (Section 3.1.1) and balancing the carry-over effect (Section 3.1.2). We survey what is known in the literature on these fairness criteria, as well as their trade-off, before we develop our own approach in Section 3.1.3.

3.1.1 Minimizing the number of breaks

Many of the theoretical results and algorithms in sports timetabling are based on graph theory. For instance, De Werra (1980) uses the complete graph K_{2n} on $2n$ nodes for constructing single round robin tournaments, with the nodes corresponding with the teams and each edge with a game between the teams of the nodes it connects. A time-constrained timetable can then be seen as a one-factorization of K_{2n} , i.e., a partitioning into edge-disjoint one-factors F_i with $i = 1, \dots, 2n - 1$. A one-factor is a set of edges such that each node in the graph is incident to exactly one of these edges. Each one-factor corresponds to a round and represents n matches.

One particular one-factorization results in so-called *canonical* timetables, which are highly popular in sport timetabling (Goossens and Spieksma, 2012b). This one-factorization has its one-factors F_i for $i = 1, \dots, 2n - 1$ defined as

$$F_i = \{(2n, i)\} \cup \{([i + k]^+, [i - k]^-) : k = 1, \dots, n - 1\},$$

where $[x]^+ = x$ if $x \leq 2n - 1$ and $[x]^+ = x - 2n + 1$ otherwise, while $[x]^- = x$ if $x \geq 1$ and $[x]^- = x + 2n - 1$ otherwise. Figure 1 illustrates the canonical one-factorization for a single round robin tournament with 6 teams.

If the league organizer can determine which match is played in which round, the minimal number of breaks for a single-round robin tournament with $2n$ teams is $2n - 2$, with $2n - 2$ teams having 1 break and 2 teams without breaks (De Werra, 1981). Moreover, De

Werra (1981) shows that this can always be achieved with a canonical timetable. For a double-round robin tournament, a timetable with $2n - 4$ breaks can easily be constructed from a single-round robin timetable with a minimal number of breaks; if we want a mirrored double-round robin timetable, the minimal number of breaks is $3n - 6$ (De Werra, 1981). However, if there is no need for a timetable that consists of consecutive single-round robin tournaments, we can limit the number of breaks to $n - 2$, even if all teams meet each other team more than twice (Goossens and Spieksma, 2012b).




Figure 1: Illustration of the canonical schedule for a single round robin tournament with 6 teams (based on Januario et al. (2016)).

If the opponents are fixed for each round, and the league organizer can only determine the home advantage, finding a timetable with a minimal number of breaks is known as the *break-minimization problem*. This problem has been tackled using e.g. constraint programming

(Régim, 2001), integer programming (Trick, 2000) and semidefinite programming (Miyashiro and Matsui, 2006).

3.1.2 Balancing the carry-over effect

Since its introduction by Russell (1980), there have been several attempts to find timetables with minimal carry-over effects value (or in other words, to balance the carry-over effects as good as possible over the teams). The lowest carry-over effects value that we may hope for in a single-round robin tournament with $2n$ teams is $2n(2n - 1)$. This is the case when each team gives a carry-over effect to each other team (except itself) exactly once. A timetable that achieves this is called a *balanced* timetable.

Russell (1980) presents an algorithm that results in a balanced timetable when the number of teams is a power of 2. Anderson (1999) found balanced timetables for 20 and 22 teams, and improved several of Russell's results. In fact, despite various approaches (Henz et al., 2004, Kidd, 2010, Miyashiro and Matsui, 2006, Trick, 2000), only Guedes and Ribeiro (2011) were able to further improve one of Anderson's results (namely, for 12 teams). Trick (2000) proved that 60 is the optimal carry-over effects value for 6 teams; Lambrechts et al. (2017) later showed that all timetables for 6 teams have this carry-over effects value. Table 5 summarizes the results.

3.1.3 Breaks versus carry-over effects

The best-known timetables with respect to the carry-over effects value do not specify the home advantage, and hence make no claims about the number of breaks. On the other hand, the canonical timetables, which allow a minimal number of breaks, have been shown to be the worst timetables with respect to balancing the carry-over effects (Lambrechts et al., 2017). For example, the timetable in Table 1 is canonical; its carry-over effects value amounts to 196 (see Table 3b).

The trade-off between the carry-over effect and breaks has been studied before in the literature. Çavdaroğlu and Atan (2019) start from the canonical timetable, and apply a round swapping procedure in order to reduce its carry-over effects value. Günneç and Demir (2019) start from a timetable with at most one break per team, and then swap rounds using a tabu-search algorithm in order to obtain a This is an author-version, published as Chapter 10 in "Science Meets Sports: 11 when statistics are more than numbers" (Cambridge Scholars Publishing).

Table 5: Carry-over effects values found in the literature for various league sizes (proven optimal values are bolded).

$2n$	$2n(2n-1)/$ Best found	Russell	Anderson	Trick	Henz et al.	Miyashiro et al.	Guedes & Ribeiro	Kidd
4	12 / 12	12	-	-	-	-	12	12
6	30 / 60	60	-	60	60	-	60	60
8	56 / 56	56	56	-	-	-	56	56
10	90 / 108	138	108	122	128	108	108	108
12	132 / 160	196	176	-	188	176	160	176
14	182 / 234	260	234	-	-	254	254	234
16	240 / 240	240	-	-	-	240	-	240
18	306 / 340	428	340	-	-	400	-	340
20	380 / 380	520	380	-	-	488	-	380
22	462 / 462	-	462	-	-	-	-	462
24	552 / 598	-	644	-	-	-	-	598

small carry-over effects value. One can see these contributions as *first-break, then carry-over* approaches. Motivated by this idea, and taking into account that carry-over effects are only related to the opponents, here, we develop a reversed approach: minimizing the carry-over effects value first, and then optimizing the number of breaks, or in other words *first-carry-over, then-break* (FCTB).

We first match the opponent pairs to rounds in order to minimize the carry-over effects value. Let T and R be the set of teams and rounds respectively, with $|T|$ even, and $|R| = |T| - 1$ (i.e., time-constrained single round robin). We say that x_{ijr} equals 1 if teams i and j ($i, j \in T, i \neq j$) play against each other in round $r \in R$, and 0 otherwise. Next, we say that c_{ijr} equals 1 if team $i \in T$ gives a carry-over effect to team $j \in T$ in round $r \in R$, and 0 otherwise. As a result, the number of carry-over effects c_{ij} that team i gives to team j is the sum of c_{ijr} over $r \in R$. Thus, we obtain the following formulation (note that the carry-over effects from the last round to the first round are also considered):

$$\min \sum_{i \in T} \sum_{j \in T} c_{i,j}^2 \quad (1)$$

subject to

$$x_{i,i,r} = 0 \quad \forall i \in T, r \in R \quad (2)$$

$$x_{i,j,r} = x_{j,i,r} \quad \forall i, j \in T, r \in R \quad (3)$$

$$\sum_{r \in R} x_{i,j,r} = 1 \quad \forall i, j \in T, i \neq j \quad (4)$$

$$\sum_{j \in T} x_{i,j,r} = 1 \quad \forall i \in T, r \in R \quad (5)$$

$$x_{i,l,r} + x_{j,l,r+1} - 1 \leq c_{i,j,r} \quad \forall i, j, l \in T, r \in R \setminus \{|R|\} \quad (6)$$

$$x_{i,l,|R|} + x_{j,l,1} - 1 \leq c_{i,j,|R|} \quad \forall i, j, l \in T \quad (7)$$

$$\sum_{r \in R} c_{i,j,r} = c_{i,j} \quad \forall i, j \in T, i \neq j \quad (8)$$

$$x_{i,j,r} \in \{0, 1\} \quad \forall i, j \in T, r \in R \quad (9)$$

$$c_{i,j,r} \in \{0, 1\} \quad \forall i, j \in T, i \neq j, r \in R \quad (10)$$

$$c_{i,j} \geq 0 \quad \forall i, j \in T, i \neq j \quad (11)$$

Constraints (2) state that a team cannot play against itself, and constraints (3) enforce that if team i meets team j on some round, that j then also meets i on that round. Constraints (4) regulate that each pair of teams meets once during the tournament and constraints (5) imply that each team plays exactly one game per round. Constraints (6) and (7) calculate the carry-over effects from team i to j over the tournament, including the effects passed from the last round to the first round. The relations between c_{ijk} and c_{ij} are shown in constraints (8).

After obtaining the assignment of opponent pairs to rounds with minimized carry-over effects from the formulation above (however, without the home advantage), the number of breaks is to be minimized. We use the following two decision variables: h_{ir} is 1 if team $i \in T$ plays home in round $r \in R$, and 0 otherwise, and b_{ir} equals 1 if team $i \in T$ has a break in round $r \in R$, and 0 otherwise. Note that x_{ijr} is known from the outcome of the before-mentioned formulation.

$$\min \sum_{i \in T} \sum_{r \in R} b_{i,r} \quad (12)$$

subject to

$$2 - h_{i,r} - h_{j,r} \geq x_{i,j,r} \quad \forall i, j \in T, r \in R \quad (13)$$

$$h_{i,r} + h_{j,r} \geq x_{i,j,r} \quad \forall i, j \in T, r \in R \quad (14)$$

$$h_{i,r} + h_{i,r+1} - 1 \leq b_{i,r+1} \quad \forall i \in T, r \in R \setminus \{|R|\} \quad (15)$$

$$1 - h_{i,r} - h_{i,r+1} \leq b_{i,r+1} \quad \forall i \in T, r \in R \setminus \{|R|\} \quad (16)$$

$$\sum_{r \in R} h_{i,r} \leq |T|/2 \quad \forall i \in T \quad (17)$$

$$\sum_{r \in R} h_{i,r} \geq |T|/2 - 1 \quad \forall i \in T \quad (18)$$

$$h_{i,r}, b_{i,r} \in \{0, 1\} \quad \forall i \in T, r \in R \quad (19)$$

According to constraints (13) two teams who play against each other cannot both play home simultaneously; nor can they both have an away game on that round (constraints (14)). Constraints (15) and (16) keep track of the home and away breaks respectively. Additionally, the number of home and away games for each team should also be balanced, as enforced by constraints (17) and (18).

In our FCTB approach, the idea is to first solve formulation (1)-(11), and then use its outcome as the input for formulation (12)-(19). The models are solved using IBM Ilog Cplex 12.71 on a MacOS 10.13.6 system with 8 GB RAM and an Intel Core i5 processor. For practical reasons, we bound the computation time to 2 hours for solving formulation (1)-(11). For up to 8 teams, this suffices for finding an optimal solution. For more teams, we check whether further improvements are possible by randomly switching the order of the rounds. As an alternative approach, we skip formulation (1)-(11), and instead use the best-known timetables with respect to the carry-over effects value published by Kidd (2010) as input for formulation (12)-(19). We refer to this approach as FCTB-best-known.

The results for the canonical timetable, from the literature, and from both our approaches for 4 to 24 teams are shown in Table 6. For 4 and 6 teams, FCTB obtains the same results as the canonical timetables, and both the carry-over effects value and the number of breaks are optimal. In comparison with the results from Çavdaroğlu and Atan (2019), the excellent carry-over effects values obtained by FCTB-best-known comes at the cost of a slightly higher number of breaks (except for 16 teams). In most cases, we have better solutions on both carry-over effects values and the number of breaks than the result from Günneç and Demir (2019). The timetables that are on the Pareto-front are indicated in bold. Note that the number of breaks per team is not bounded in our approach, but we enforce a balance in the number of home and away games played by every team during the tournament, which is not considered in the literature.

Table 6: Trade-offs between breaks and carry-over effects value.

Number of teams		4	6	8	10	12	14	16	18	20	22	24
Canonical timetable	breaks	2	4	6	8	10	12	14	16	18	20	22
	carry-over	12	60	196	468	924	1612	2580	3876	5548	7644	10212
FCTB	breaks	2	4	8	12	18	24	34	40	-	-	-
	carry-over	12	60	56	138	206	292	400	508	-	-	-
FCTB-best-known	break	2	4	8	14	22	28	24	48	50	58	64
	carry-over	12	60	56	108	176 ²	234	240	340	380	462	598
Çavdaroğlu & Atan	breaks	-	-	-	12	18	26	32	42	-	-	-
	carry-over	-	-	-	136	192	254	330	406	-	-	-
Gümneç & Demir ¹	breaks	-	-	-	12	16	24	34	40	-	-	-
	carry-over	-	-	-	144	212	302	396	556	-	-	-

¹ These results are from the setting with unbounded breaks.

² The best carry-over effects value for 12 teams is 160 (Guedes and Ribeiro, 2011), however, we were not able to retrieve the corresponding timetable from the literature.

3.2 Trade-offs between efficiency and equity

There are typically no timetables that make every team happy. Indeed, for each timetable $D \in \mathcal{D}$, team $1 \leq i \leq n$ has a non-negative aversion $f_i(D) : \mathcal{D} \rightarrow \mathbb{R}^+$. This aversion is based on one or more properties of the timetable (e.g. a break) that is undesirable for team i . We assume that the league organizer has perfect knowledge of the aversion of each team, and is in control to select a timetable D from \mathcal{D} .

As it is in the league organizer’s interest to please as many teams as possible, each function f_i needs to be minimized. A traditional approach is to use a *minisum* objective $\sum_{i=1}^n f_i(D)$, which minimizes the total (or equivalently, the mean) aversion to the timetable. The main flaw of this approach is that only the aggregate of the undesirable property is considered (i.e., *criterion efficiency*), and not the distribution of the property over the teams (i.e., *distribution equity*).

In recent years, the literature has come up with various inequity-averse optimization techniques that incorporate distribution equity alongside criterion efficiency (see Karsu and Morton (2015) for an overview). Of particular interest is the concept of *equitable-efficiency* (Ogryczak, 1997). To explain this concept, we need some additional terminology. Aversion vector $\vec{x}_D = (f_1(D), f_2(D), \dots, f_n(D))$ collects the aversion of all teams. We assume that \vec{x}_D is sorted in non-increasing order meaning that $f_1(D) \geq f_2(D) \geq \dots \geq f_n(D)$. As the league organizer strives for maximal efficiency, less aversion for

one team is always better: \vec{x}_D *rationaly dominates* $\vec{x}_{D'}$ if the aversion for at least one team is smaller whereas no other team is worse off. We call \vec{x}_D *Pareto-efficient* in \mathcal{D} if and only if there exists no timetable $D' \in \mathcal{D}$ such that $\vec{x}_{D'}$ rationally dominates \vec{x}_D . The image of all non-dominated solutions in the objective space is called the *Pareto-frontier*.

In order to fairly distribute the total aversion over the teams, the concept of equitable-efficiency is built around two fundamental axioms. The first axiom prescribes *anonymous identities*, i.e., the identities of the teams are not important to analyze the equity of the aversion vector. Although this seems reasonable, the practice of sports timetabling may be different in the sense that some teams can be considered *more equal* than others. The second axiom is known as the *Pigou-Dalton principle of transfers*: any transfer from a worse-off team to a better-off team, *ceteris paribus*, should always result in a more preferable aversion vector. For a formal definition of the axioms, we refer to Ogryczak (2000).

Vector \vec{x}_D *equitably dominates* vector $\vec{x}_{D'}$ if we can produce \vec{x}_D from $\vec{x}_{D'}$ after a finite sequence of index permutations and at least one aversion transfer from a worse-off team to a better-off team (also called a Robin Hood operation) or a decrease of the aversion of a particular team. We call \vec{x}_D *equitably-efficient* if and only if there exists no timetable $D' \in \mathcal{D}$ such that $\vec{x}_{D'}$ equitably dominates \vec{x}_D . As both axioms do not conflict with Pareto-efficiency, an equitably-efficient timetable is also Pareto-efficient but not the other way around (Karsu and Morton, 2015).

The remainder of this section is as follows. First, Section 3.2.1 describes a typical time-relaxed sports timetabling problem in which venue and player availability constraints need to be taken into account. Next, Section 3.2.2 proposes an integer programming (IP) model to generate equitably-efficient timetables with regard to teams' rest periods. This IP model, however, requires a considerable amount of computational resources. Section 3.2.3 therefore presents a multi-objective evolutionary algorithm capable of generating a rich set of equitably-efficient solutions in a single run.

3.2.1 Availability constraints in time-relaxed timetabling

In the time-Relaxed Availability Constrained double Round-Robin Tournament problem (RAC-2RRT), the input consists of a set of rounds R , and a set of teams T where $n = |T|$. Each team $i \in T$ also provides a venue availability set $H_i \subseteq R$ containing all rounds during which i 's venue is available, and a player availability set A_i containing all rounds during which i 's players are available. Since a team can only play (at home or away) when its players are available, we assume without loss of generality that $H_i \subseteq A_i$ for each $i \in T$. This makes that a team can play at home on all rounds in H_i , and that it can play away on all rounds in A_i . Finally, a parameter τ is given, defining the total number of rounds after which we assume a team is fully recovered from its previous game (e.g. 5 days).

RAC-2RRT consists of finding a feasible timetable, that is an assignment of games to rounds such that:

- (C1) each team plays exactly one home game against every other team,
- (C2) the venue availability H_i with $i \in T$ is respected (i.e., no game i - j is planned on a round $r \notin H_i$),
- (C3) the player availability A_i with $i \in T$ is respected (i.e., no game i - j or j - i is planned on a round $r \notin A_i$),
- (C4) each team plays at most one game per round $r \in R$, and
- (C5) each team plays at most two games per $\tau + 1$ rounds.

Van Bulck et al. (2019) propose IP formulation (20)-(24) to solve RAC-2RRT. In this model, the variable $z_{i,j,r}$ is 1 if team $i \in T$ and team $j \in T \setminus \{i\}$ meet at the venue of i on round $r \in R$. Constraints (20) ensure that each team plays exactly one home game against every other team (C1). Constraints (21) require a team to play at most one game per round (C4), whereas constraints (22) enforce that a team plays at most two games per $\tau + 1$ rounds (C5). Constraints (23) reduce the number of variables in the system by explicitly stating that two teams can only meet when the venue of the home team and the players of the away team are simultaneously available (C2), (C3); in practice these variables are not created. Finally, constraints (24) are the binary constraints on the z -variables.

This is an author-version, published as Chapter 10 in "Science Meets Sports: 17 when statistics are more than numbers" (Cambridge Scholars Publishing).

$$\sum_{r \in H_i \cap A_j} z_{i,j,r} = 1 \quad \forall i, j \in T : i \neq j \quad (20)$$

$$\sum_{j \in T \setminus \{i\}} (z_{i,j,r} + z_{j,i,r}) \leq 1 \quad \forall i \in T, r \in A_i \quad (21)$$

$$\sum_{j \in T \setminus \{i\}} \sum_{k=r}^{r+\tau} (z_{i,j,k} + z_{j,i,k}) \leq 2 \quad \forall i \in T, r \in A_i \quad (22)$$

$$z_{i,j,r} = 0 \quad \forall i, j \in T : i \neq j, r \in R \setminus \{H_i \cap A_j\} \quad (23)$$

$$z_{i,j,r} \in \{0, 1\} \quad \forall i, j \in T : i \neq j, r \in H_i \cap A_j \quad (24)$$

The ARTP penalizes the timetable with a positive value p_r each time a team has only $r < \tau$ rounds between consecutive games, with $p_r \leq p_{r-1}$. We measure the aversion of team i by summing over the rest time penalties related to all games of team i and refer to this sum with ARTP_i . To minimize the ARTP, the minisum model uses an auxiliary variable $y_{i,r,t}$ which is 1 if team i plays a game on round r followed by its next game on round t , and 0 otherwise. Constraints (25) regulate the value of the $y_{i,r,t}$ variables by considering the number of rounds between two consecutive games of the same team. We note that it follows from (C5) that the games are consecutive if team i plays on round r and t and $|t - r| \leq \tau$. Hence, we can strengthen the formulation by dropping the negative summation term of Equation (25). Constraints (26) model the aversion of team i by setting f_i equal to ARTP_i . Finally, constraints (27) state that the y -variables are binary.

Minisum model

$$\min \sum_{i \in T} f_i$$

subject to

$$(20) - (24)$$

$$\sum_{j \in T \setminus \{i\}} \left(z_{i,j,r} + z_{j,i,r} + z_{i,j,t} + z_{j,i,t} - \sum_{k=r+1}^{t-1} (z_{i,j,k} + z_{j,i,k}) \right) - 1 \leq y_{i,r,t} \quad \forall i \in T, r, t \in A_i : r < t, t - r \leq \tau \quad (25)$$

$$f_i = \sum_{r \in A_i} \sum_{t=r+1}^{r+\tau} p_{(t-r-1)} y_{i,r,t} \quad \forall i \in T \quad (26)$$

$$y_{i,r,t} \in \{0, 1\} \quad \forall i \in T, r, t \in A_i : r < t, t - r \leq \tau \quad (27)$$

This is an author-version, published as Chapter 10 in “Science Meets Sports: 18 when statistics are more than numbers” (Cambridge Scholars Publishing).

3.2.2 Balancing rest times over teams with IP

Ogryczak (2000) shows that equitable dominance between timetables D and D' can be identified by comparing their cumulative ordered aversion vectors. To this purpose, let $\vec{\Theta}_D = (\theta_1(D), \theta_2(D), \dots, \theta_n(D))$ be the cumulative ordered aversion vector where $\theta_i(D) = \sum_{j=1}^i f_j(D)$ for $i \in T$ (recall that $f_i(D) \leq f_{i+1}(D)$). Timetable D equitably dominates D' if and only if $\theta_i(D) \leq \theta_i(D')$ for all $i \in T$ with at least one strict inequality (Ogryczak, 2000). As an example, timetable D with $\vec{x}_D = (15, 10, 5)$ dominates D' with $\vec{x}_{D'} = (20, 10, 0)$ as we have $\vec{\Theta}_D = (15, 25, 30)$ and $\vec{\Theta}_{D'} = (20, 30, 30)$.

An important consequence is that all equitably-efficient timetables may be generated by enumerating all Pareto-efficient solutions with respect to $\min_{D \in \mathcal{D}} \vec{\Theta}_D$. To this purpose, Kostreva et al. (2004) propose constraints (28)-(30). Constraints (28) model auxiliary variables $d_{i,j}^+$ representing the upside deviation of the aversion of team i from the value of the unrestricted variable t_i . These auxiliary variables are then used in constraints (29) to model $\theta_i(D)$. For a correctness proof, we refer to Kostreva et al. (2004) and Ogryczak et al. (2008). Applying these constraints to our problem leads to the following IP-EQ model.

IP-EQ model

$$\min \vec{\Theta} = (\theta_1, \theta_2, \dots, \theta_n)$$

subject to

$$(20) - (27)$$

$$t_i + d_{i,j}^+ \geq f_i \quad \forall i, j \in T \quad (28)$$

$$\theta_i = it_i + \sum_{j \in T} d_{i,j}^+ \quad \forall i \in T \quad (29)$$

$$d_{i,j}^+ \geq 0 \quad \forall i, j \in T \quad (30)$$

There are several approaches to generate Pareto-efficient solutions for the IP-EQ model. One approach is to use an aggregation function which maps $\vec{\Theta}_D$ into a single objective to be optimized. Kostreva et al. (2004) characterize aggregation functions that result in an equitably-efficient solution. One example of such aggregation function is $\sum_{i \in T} w_i \theta_i(D)$

with $w_i = (n + (n - 2i + 1)\lambda)/n^2$ and λ a trade-off parameter in the range $]0, 1]$ (see Ogryczak (2000)). We will use this aggregation function in the computational results of Section 4.2, where we solve the IP-EQ model with Gurobi Optimizer 7.5.2 using 13 threads and 8 GB of RAM, and a time limit of 3 hours.

Another approach is to treat the objectives hierarchically. Rawl’s difference principle states that “social and economic inequalities are to be arranged so that they are to be of greatest benefit to the least-advantaged members of society” (Rawls and Kelly, 2001) and therefore suggests a *minimax* approach which minimizes $\theta_1(D)$ first. The minimax approach, however, only concerns the aversion of the worst-off team (i.e., $\theta_1(D)$) and thereby misses remaining optimization opportunities. Ogryczak (1997) proves that equitably-efficient solutions can be generated by the *lexicographic minimax* approach which optimizes $\theta_i(D)$ in increasing order of i , without allowing to deteriorate previously optimized objectives. Utilitarian philosophers, on the other hand, would advocate the *minisum* approach which corresponds to minimization of $\theta_n(D)$. An equitably-efficient solution can be generated by the *lexicographic minisum* approach which optimizes $\theta_i(D)$ in decreasing order of i .

3.2.3 A bi-criteria evolutionary algorithm

The set of Pareto-efficient solutions for $\vec{\Theta}_D$ defines the entire set of equitably-efficient solutions. Unfortunately, direct optimization of $\vec{\Theta}_D$ implies the use of an aggregation function or requires that the objectives are optimized hierarchically. As it is daunting, if not impossible, to define the preference of the league organizer prior with regard to the efficiency-equity trade-off, this section proposes a bi-criteria evolutionary algorithm generating a rich set of equitably-efficient compromise solutions. This approach is not only appealing from a computational point of view but also allows to visualize the trade-off in the efficiency-equity image space.

In particular, the bi-criteria evolutionary algorithm takes into account criterion efficiency via mean aversion $\mu_D = \sum_{i \in T} f_i(D)/n$ and distribution equity via the mean absolute difference MD_D :

$$\text{MD}_D = \frac{1}{2n^2} \sum_{i,j \in T} |f_i(D) - f_j(D)|. \quad (31)$$

Intuitively, MD_D expresses the expected difference in the aversion of two uniformly chosen teams.

Noteworthy, existing timetabling literature measures equity via Jain’s fairness index $J_D = (\sum_{i=1}^n f_i(D))^2 / (n \sum_{i=1}^n f_i(D)^2)$ (see e.g. Mühlenhaller and Wanka (2016) and Muklason et al. (2017)) or Gini’s coefficient $G_D = MD_D / \mu_D$. Unfortunately, we cannot use Jain’s index and Gini’s coefficient as they do not comply with the concept of equitable-efficiency (see (Ogryczak, 2000)).

In order to minimize the ARTP, Van Bulck et al. (2019) propose a tabu-search based algorithm. The key component in this algorithm consists of scheduling (or rescheduling) all home games of one team, which is modeled as a transportation problem. To fairly distribute rest times over teams, we propose a multi-objective genetic algorithm, which includes a local search operator based on the algorithm of Van Bulck et al. (2019). Genetic algorithms which involve a local search operator are known as memetic algorithms; we refer to Jazzkiewicz et al. (2012) for an introduction on memetic algorithms. To cope with multiple objectives, our algorithm uses several ideas proposed by Ishibuchi and Murata (1998). The remainder of this section briefly summarizes the main components of our algorithm. A general overview of the algorithmic flow is depicted in Figure 2.

Solution representation The algorithm encodes a double round-robin timetable via a matrix where each cell (i, j) carries round $r_{i,j}$ on which home team $i \in T$ plays against away team $j \in T \setminus \{i\}$. We allow the algorithm not to plan a game by leaving the corresponding cell blank, but this results in a high cost P during the fitness evaluation.

Fitness of individuals For each candidate timetable D , we count the total number of unscheduled games u_D and assign fitness $Pu_D + \mu_D + \lambda MD_D$. Ogryczak (2000) proves that any timetable which is optimal for this fitness function is equitably-efficient when $0 < \lambda \leq 1$. At the beginning of each generation, the algorithm varies the search direction by uniformly choosing the trade-off parameter λ in the range $]0, 1]$. This facilitates finding a diverse set of equitably-efficient solutions (see Ishibuchi and Murata (1998)).

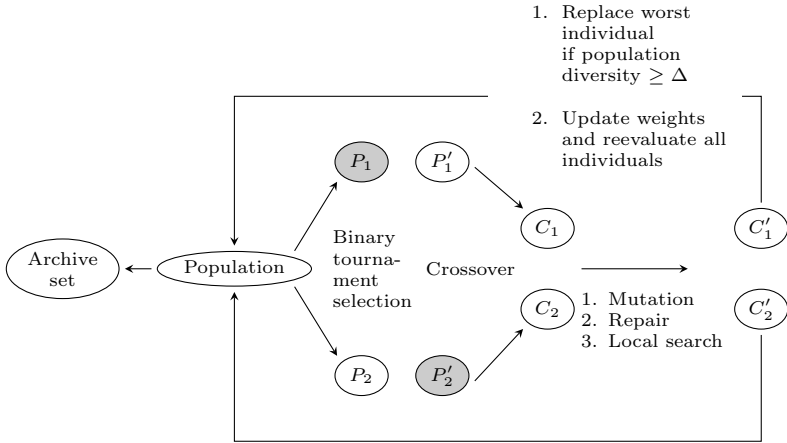


Figure 2: Illustration of the memetic algorithm with population management (MA|PM). Grey ellipses represent the parents selected by the binary tournament operator.

Recombination and mutation To improve solutions and to avoid getting trapped in local optima, our genetic algorithm varies candidate solutions via crossover and mutation. The crossover operator takes as input two parent solutions and generates a new offspring solution as follows. First, it draws a random number t between 1 and $n-1$. Next, it copies all home game assignments of the first t teams from the first parent, and all other home game assignments from the second parent. Since each game features one home team and one away team, this makes that all games are scheduled. We also consider a variant of this operator which copies away game assignments. After crossover, each offspring solution undergoes mutation. The algorithm randomly decides for each cell (i, j) independently whether cell value $r_{i,j}$ is to be mutated, in which case $r_{i,j}$ is replaced by a uniformly chosen value from $\{H_i \cap A_j\} \setminus \{r_{i,j}\}$.

Repair and local search As new offspring solutions may violate constraints (C4) and (C5), we fully repair a timetable for all teams in a randomly chosen order by solving a transportation network which schedules or reschedules all home games of a chosen team. Van Bulck

et al. (2019) explain how this transportation network can be adapted to optimize the ARTP. The fitness function used in our version of the algorithm, however, varies the value of λ to take into account distribution equity alongside criterion efficiency. Therefore, the fitness of a solution may deteriorate after solving the transportation network. We only accept the new solution if it improves the fitness value.

Population management Genetic algorithms need a parent and survivor selection scheme to guide the evolution of the population. We create the initial population by repeatedly solving the IP model presented in Equations (20)-(24). At the beginning of each iteration, our algorithm uniformly chooses $0 < \lambda \leq 1$ and re-evaluates all individuals in the current population. Next, a binary tournament operator selects two parent solutions based on the re-evaluated fitness values. With a probability of p_c the two parents mate in which case two offsprings are generated by crossover. In the other case, the two offsprings are identical to the two parents. With a probability of p_m , the offspring solutions undergo mutation. After recombination and mutation, the local improvement heuristic is applied. Subsequently, the algorithm calculates for each offspring and each member of the population a dissimilarity distance expressed in terms of the percentage of different game to round assignments. If the distance between the offspring and each member of the population is greater than the diversity parameter Δ , the offspring replaces the worst solution in the current population (for more details, see Sörensen and Sevaux (2006)).

Apart from the current population, our version of the memetic algorithm stores an archive set of equitably-efficient solutions. If an offspring solution is not equitably dominated by any other solution in the archive, the offspring solution is added to the archive from which we subsequently remove newly equitably dominated solutions. The use of this archive set makes that no equitably-efficient solutions are lost (for more information, see Ishibuchi and Murata (1998)). Ideally, solutions in the archive set are close to the Pareto-efficient frontier and cover the entire mean-equity Pareto front.

Computational setting The memetic algorithm is implemented in in C++, compiled with g++ 4.8.5 using optimization flag -O3, and parallelized with OpenMP. The parameters of the algorithm were tuned with *irace* (López-Ibáñez et al., 2016) using the hypervolume indicator. This is an author-version, published as Chapter 10 in “Science Meets Sports: 23 when statistics are more than numbers” (Cambridge Scholars Publishing).

tor. To solve the transportation problems, we use an $\mathcal{O}(n^3)$ implementation of the Kuhn-Munkres algorithm. We run the genetic algorithm with a time-limit of 10 minutes, 12 cores, and 2 GB of RAM on a CentOS 7.4 GNU/Linux based system with an Intel E5-2680 processor, running at 2.5 GHz.

4 For practitioners

4.1 Trade-offs between two fairness criteria

Every year, football fans and clubs eagerly await the announcement of the official timetable for the new season by the league organizers. This timetable is the result of a complex planning process, and tries to take into account wishes and requirements from the police, broadcasters, sponsors, etc. Since it is rarely revealed exactly what these requirements are, the media focus mostly on the fairness of the timetable.

In this section, we discuss the number of breaks and the carry-over effect present in the official timetables of 10 main European football leagues: Belgium, England, France, Germany, Italy, the Netherlands, Portugal, Russia, Spain, and Ukraine (for other empirical studies on football timetables, we refer to Goossens and Spieksma (2012b) and Yi et al. (2019)). We studied all timetables from season 2009–2010 up to season 2018–2019. Table 7 lists the average carry-over effects value per team and average number of breaks per team for each season.

Looking at breaks, England has a strikingly high number, with the Netherlands coming second. It is safe to say that the timetabling process in these countries leaves a lot of room for improvement. In most other countries, the number of breaks is low and quite stable, although Russia has had a few aberrant seasons, and France seems to have abandoned their timetables with two breaks per team since season 2014–2015.

Except for the last season, Spain displays the largest carry-over effects value among the 10 leagues. The main reason is that it uses a so-called canonical timetable (see Section 3.1.3). The canonical timetable has been (and to some extent still is) very popular in sports timetabling (Goossens and Spieksma, 2012b), however, it is also the worst possible timetable with respect to balancing carry-over effects. This is an author-version, published as Chapter 10 in “Science Meets Sports: 24 when statistics are more than numbers” (Cambridge Scholars Publishing).

Table 7: Overview of the carry-over effects value / breaks for 10 main European football leagues for seasons 2009 – 2010 till 2018 – 2019

Year	Belgium	England	France	Germany	Italy	Netherlands	Portugal	Russia	Spain	Ukraine
09/10	2383/42	2436/172	4548/40	3542/48	2974/66	1966/114	2307/42	8345/40	20657/54	9015/50
10/11	2441/42	2892/124	4730/40	5154/48	3306/64	1778/112	2231/42	8491/32	21044/54	9015/44
11/12	2571/42	2574/126	4740/40	4460/48	3144/66	1906/124	2287/42	9591/50	21044/54	9015/60
12/13	2663/42	2034/130	4680/40	4202/48	3124/62	1940/106	2307/42	8491/40	21044/54	9015/54
13/14	2265/66	2764/126	4536/40	3622/48	3508/64	1930/116	3962/42	1431/112	21044/54	9015/42
14/15	2423/44	2764/116	4536/48	4342/48	3508/62	1930/116	3962/48	1431/92	21044/54	5517/52
15/16	2201/44	2594/138	2258/84	4050/48	3232/64	1806/110	3842/48	8491/32	21044/54	5477/36
16/17	1545/52	2696/154	2510/74	4046/48	3124/64	2076/110	3854/48	4859/78	21044/54	3075/32
17/18	1477/48	2540/152	2416/66	3874/48	3160/72	2142/88	4058/48	8491/44	21044/58	3075/32
18/19	1513/48	2642/132	2498/70	3734/48	3268/64	2132/134	4330/48	8491/60	2598/72	3075/36

(Lambrechts et al., 2017). The canonical timetable has also been used in Ukraine (all seasons) and Russia (7 of the 10 seasons).

It should be noted that Table 7 sketches a somewhat distorted picture, since the number of teams in a league has an impact on the minimal number of breaks and the minimal carry-over effects value that can be attained. The leagues indeed do not feature the same number of teams: England, France, Italy and Spain have 20 teams, the Netherlands and Germany have 18 teams, and Belgium and Russia have 16 teams. In Portugal, the number of teams increased from 16 to 18 in season 2014 – 2015, while in Ukraine, the number of teams dropped from 16 to 14 in that same season, and even to 12 in season 2016 – 2017. Hence, in order to allow a proper comparison, Figure 3 shows the average number of breaks per team on its horizontal axis. While this makes sense for breaks, this is much less the case for the average carry-over effect per team, since the minimal carry-over effects value increases non-linearly with the number of teams. Hence, in order to allow a proper comparison, Figure 3 presents how close the carry-over effects value is to the maximal carry-over effects value for a league of that size (corresponding to an indexed value of 1) on its vertical axis. Since the maximal carry-over effects value is not known for double round robin tournaments, and since many timetables apply mirroring, we have focused on the first half of the season only.

The trade-off between balancing carry-over effects and minimizing the total number of breaks is now apparent in Figure 3. The figure clearly shows the Netherlands and England on the one hand, who totally neglect the number of breaks, and Spain, Russia and Ukraine on the other hand, who ignore the carry-over effects value. The Pareto-front, which corresponds to the best trade-off one can make, is also indicated in Figure 3. The Pareto-front is based on timetables from Italy, France and Belgium. That means that for these seasons, no one does better with respect to breaks *and* carry-over. Hence, although they make different choices with respect to the trade-off, these three countries are ‘getting the most value for their money’. Portugal and Germany are not far away in most seasons, but leagues like England could get a much lower number of breaks for the same carry-over effects value. Likewise, Spain could drastically improve its balance of carry-over effects, without incurring more breaks. It is interesting to see that while most leagues have their seasons in the same area of the plot, the league organizers in Russia make very different choices

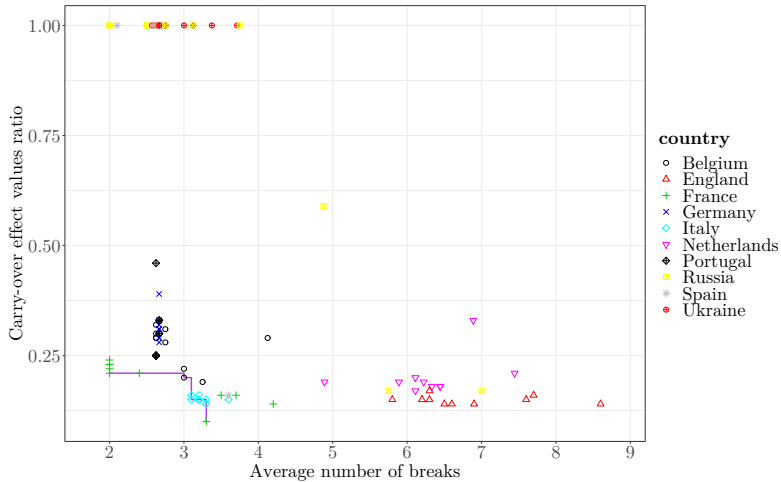


Figure 3: Trade-off between balancing carry-over effects and minimizing breaks (each entry corresponds to one season in one league).

from one season to the next, i.e., sometimes opting for a very high carry-over effects value and small number of breaks, and sometimes vice versa.

4.2 Trade-offs between efficiency and equity

The ‘Liefhebbers Zaalvoetbal Cup’ (LZV Cup) is a non-professional indoor football league founded in 2002. This league currently involves 548 teams, grouped divisions in 20 regions in Flanders (Belgium). In a division, each team plays each other team once at home and once away in a double round-robin tournament. The league aims to attract teams that consist of friends, is open to all ages, and considers fair play of utmost importance. The games are played without referees, since, according to the organizers, “referees are expensive, make mistakes, and invite players to explore the borders of sportsmanship”². In this context, it makes sense that also their timetables display a

²see www.lzvcup.be [in Dutch]

high level of fairness.

In this chapter, we consider 3 divisions of the LZV Cup: two of these divisions have 15 teams, the other has 14 teams³. The season starts on September 1st and ends on May 31st, which corresponds to 273 rounds. A time-relaxed timetable is required in which no team has more than two games in a period of τ rounds or less. For the divisions with 15 teams, $\tau = 8$ whereas in the division with 14 teams $\tau = 9$. We use $p_r = (\tau - r)^2$ to denote the penalty incurred for every pair of consecutive games played by a team within a period of $r < \tau$ rounds. We refer to the total sum of penalties incurred by team i as the aggregated rest time penalty of i (ARTP_i); the sum of the ARTP_i values over all teams results in the ARTP score of the timetable.

In order to evaluate the fairness of a timetable, we use two decision criteria. First, we use the mean aggregated rest time penalty $\mu_D = \text{ARTP}/n$ to measure the criterion efficiency of timetable D . Second, we use the mean absolute difference $\text{MD}_D = \frac{1}{2n^2} \sum_{i=1}^n \sum_{j=1}^n |\text{ARTP}_i - \text{ARTP}_j|$ to measure the equity of timetable D with regard to the distribution of the rest time penalties over the teams. Intuitively, MD_D measures the expected difference in the aggregated rest time penalty of two randomly chosen teams.

The league organizers are confronted with the price of equity: by how much does the mean aggregated rest time penalty increase when we additionally consider the distribution of the penalties over the teams? After generating schedules manually for a number of years, the league organizers have moved to an integer programming formulation (based on constraints (20)-(24) described in Section 3.2.1). However, this approach only gives them one single solution per division, and no insight on the trade-off between equity and efficiency. In this chapter, we have developed an IP formulation IP-EQ (Section 3.2.2) and an evolutionary algorithm (Section 3.2.3) in order to produce alternative timetables, and to shed light on the price of equity.

Figure 4 shows the Pareto-front with regard to criterion efficiency (μ_D) and distribution equity (MD_D) for each of the three divisions. Besides, we draw a straight trade-off line through the solution with

³More details on these divisions can be found in Van Bulck et al. (2019), in which divisions 1 to 3 are referred to as benchmark instance #2, #3 and #4, respectively

the best efficiency score and set the slope such that a 1% increase in equity yields a 1% increase in efficiency. Mühenthaler and Wanka (2016) argue that picking a timetable below this trade-off line might be an attractive option in real-world applications as the increase in fairness is larger than the decrease in efficiency.

When analyzing Figure 4, we see that the official timetable is far above the Pareto frontier: more efficient and equitable timetables exist. The genetic algorithm is able to produce a rich set of equitably-efficient solutions which are well separated in the efficiency-equity (μ_D - MD_D) image space. This is remarkable as the genetic algorithm produces all its solutions in a single run with 10 minutes of computation time, whereas the IP-formulations generate only one solution after 3 hours of computation time.

Figure 4 also shows that the price of fairness can be quite different between divisions (compare for example a 1% reduction of MD_D in division 1 and 2). This motivates our approach to provide the league organizer with a rich set of equitably-efficient timetables. From this set, the league organizer may select a final timetable by e.g. further analyzing the rest times distribution.

As an example, Figure 5 shows the rest time penalty $ARTP_i$ for each of the 15 teams in division 2 for various timetables, where team 1 is the worst-off team in that timetable with respect to rest times, and team 15 has the most favorable rest times. The figure compares the timetable with the lowest maximal $ARTP_i$, the most equitable timetable, the most efficient timetable, and the timetable that was actually used in the LZV Cup competition. It would not be hard for the league organizers to convince the teams of the added value of our timetabling algorithms. Indeed, when comparing the most efficient timetable with the official one, we see that every team is better off. Unfortunately, the most efficient timetable does not take into account the distribution of the rest times over the teams. The timetable with the lowest maximal $ARTP_i$ illustrates that the rest times of the worst-off team can be improved considerably, although this happens at the expense of nearly every other team when compared with the most efficient timetable. The most equitable timetable may work as a compromise solution as, except for team 2, mainly the teams with low $ARTP_i$ penalties pay the price to improve the situation for the worst-off team.




Figure 4: Pareto frontier for division 1 (top), 2 (middle), and 3 (bottom). Red circles represent the solutions found by the genetic algorithm, green triangles the solutions found with IP-EQ, and blue squares the official solutions. The Pareto frontier is indicated by the full line in purple, whereas the 1% trade-off line is indicated by the dashed line in gray.




Figure 5: Aggregated rest time per team for the timetable with the lowest maximal ARTP (‘MiniMax’, red), most equitable timetable (‘MiniMD’, green), most efficient timetable (‘MiniSum’, blue), and the official timetable (purple). Teams are sorted in non-decreasing order of $ARTP_i$, which makes that the worst-off team (i.e., team 1), may be a different team, depending on the timetable.

5 Conclusion

In round robin competitions, it is often said that teams should not complain about the timetable, because “at the end of the day you have to play against everyone”⁴. This chapter has demonstrated that there is more to fairness in sports timetabling than that, by discussing breaks, the carry-over effect, and rest times.

When confronted with multiple fairness issues, it may be difficult to get the best of both worlds. For instance, a better balance of the carry-over effects usually goes at the expense of a higher number of breaks. We developed a new method that generates timetables that offer a better trade-off than those currently available in the literature. We also looked at how the big European football competitions deal with this trade-off, and found that France, Italy and Belgium are making sensible choices; other leagues leave plenty of room for

⁴This was for instance said by Juan Mata (Manchester United), after learning that his team had to open the 2019-2020 season with a challenging game against Chelsea, see www.manutd.com/en/news/detail/juan-mata-previews-manutd-v-chelsea-to-start-new-premier-league-season.

improvement.

Even a timetable with few unfair properties may be ill received when only a small number of teams carry the majority of this burden. In contrast to classic approaches which only minimize the total unfairness, we showed how to equitably distribute the unfair properties over the teams. In particular, we demonstrated the importance of balancing rest times between consecutive games in a real-life indoor football competition. We found that the distribution of rest times over the teams can be improved, although this results in a slightly shorter average rest time.

In conclusion, we want to stress the importance for league planners to work with a timetabling method that offers multiple and diverse timetables that are close to the Pareto front. Only in this way, the league planners can make a thought-out choice on fairness trade-offs in sports timetabling.

Acknowledgement

The authors wish to thank prof. Mario Guajardo for discussing fairness issues in sports, as well as the participants of the Fairness in Sports workshops (April 12th 2018, Ghent, Belgium and June 5th 2018, London, UK) who indirectly contributed to this chapter.

References

- Anderson, I. (1999). Balancing carry-over effects in tournaments. In Holroyd, F., Quinn, K., Rowley, C., and Webb, B., editors, *Combinatorial Designs and Their Applications*, pages 1–16, Boca Raton. CRC Research Notes in Mathematics. Chapman & Hall.
- Bengtsson, H., Ekstrand, J., and Hägglund, M. (2013). Muscle injury rates in professional football increase with fixture congestion: an 11-year follow-up of the UEFA Champions League injury study. *British Journal of Sports Medicine*, 47:743–747.
- Briskorn, D. and Knust, S. (2010). Constructing fair sports league schedules with regard to strength groups. *Discrete Applied Mathematics*, 158:123–135.

- Çavdaroğlu, B. and Atan, T. (2019). Integrated break and carryover minimization problem in round robin tournament. In *MathSport International 2019 Conference Proceedings*, pages 25–32.
- De Werra, D. (1980). Geography, games, and graphs. *Discrete Applied Mathematics*, 2:327–337.
- De Werra, D. (1981). Scheduling in sports. *Studies on Graphs and Discrete Programming*, 11:381–395.
- Forrest, D. and Simmons, R. (2005). New issues in attendance demand: The case of the English football league. *Journal of Sports Economics*, 7(3):247–266.
- Friedman, M. (1975). *There’s No Such Thing As a Free Lunch*. Open Court Publishing Company (US).
- Goossens, D. R. and Spieksma, F. C. R. (2012a). The carryover effect does not influence football results. *Journal of Sports Economics*, 13(3):288–305.
- Goossens, D. R. and Spieksma, F. C. R. (2012b). Soccer schedules in Europe: an overview. *Journal of Scheduling*, 15(5):641–651.
- Guedes, A. C. and Ribeiro, C. C. (2011). A heuristic for minimizing weighted carry-over effects in round robin tournaments. *Journal of Scheduling*, 14(6):655–667.
- Günneç, D. and Demir, E. (2019). Fair-fixtute: Minimizing carry-over effects in football leagues. *Journal of Industrial & Management Optimization*, 15(4):1565–1577.
- Henz, M., Müller, T., and Thiel, S. (2004). Global constraints for round robin tournament scheduling. *European Journal of Operational Research*, 153(1):92–101.
- Ishibuchi, H. and Murata, T. (1998). A multi-objective genetic local search algorithm and its application to flowshop scheduling. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 28:392–403.
- Januario, T., Urrutia, S., Celso, C. R., and de Werra, D. (2016). Edge coloring: A natural model for sports scheduling. *Eur. J. Oper. Res.*, 254:1 – 8.

- Jaszkiewicz, A., Ishibuchi, H., and Zhang, Q. (2012). Multiobjective memetic algorithms. In *Handbook of Memetic Algorithms*, pages 201–217. Springer.
- Karsu, O. and Morton, A. (2015). Inequity averse optimization in operational research. *European Journal of Operational Research*, 245:343–359.
- Kidd, M. (2010). A tabu-search for minimising the carry-over effects value of a round-robin tournament. *ORiON*, 26(2):125–141.
- Kostreva, M. M., Ogryczak, W., and Wierzbicki, A. (2004). Equitable aggregations and multiple criteria analysis. *European Journal of Operational Research*, 158:362–377.
- Lambrechts, E., Ficker, A. M. C., Goossens, D. R., and Spieksma, F. C. R. (2017). Round-robin tournaments generated by the circle method have maximum carry-over. *Mathematical Programming*, 172:270–302.
- López-Ibáñez, M., Dubois-Lacoste, J., Cáceres, L. P., Birattari, M., and Stützle, T. (2016). The irace package: Iterated racing for automatic algorithm configuration. *Operations Research Perspectives*, 3:43–58.
- Miyashiro, R. and Matsui, T. (2006). Minimizing the carry-over effects value in a round-robin tournament. In *Proceedings of the 6th International Conference on the Practice and Theory of Automated Timetabling*, pages 460–463.
- Mühlenthaler, M. and Wanka, R. (2016). Fairness in academic course timetabling. *Annals of Operations Research*, 239:171–188.
- Muklason, A., Parkes, A. J., Özcan, E., McCollum, B., and McMullan, P. (2017). Fairness in examination timetabling: Student preferences and extended formulations. *Applied Soft Computing*, 55:302–318.
- Ogryczak, W. (1997). On the lexicographic minimax approach to location problems. *European Journal of Operational Research*, 100:566–585.
- Ogryczak, W. (2000). Inequality measures and equitable approaches to location problems. *European Journal of Operational Research*, 122:374–391.

- Ogryczak, W., Wierzbicki, A., and Milewski, M. (2008). A multi-criteria approach to fair and efficient bandwidth allocation. *Omega*, 36:451–463.
- Rawls, J. and Kelly, E. (2001). *Justice as fairness: A restatement*. Harvard University Press, Cambridge (US).
- Régin, J.-C. (2001). Minimization of the number of breaks in sports scheduling problems using constraint programming. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 57:115–130.
- Russell, K. (1980). Balancing carry-over effects in round robin tournaments. *Biometrika*, 67(1):127–131.
- Sörensen, K. and Sevaux, M. (2006). MA|PM: memetic algorithms with population management. *Computers & Operations Research*, 33:1214–1225.
- Suksompong, W. (2016). Scheduling asynchronous round-robin tournaments. *Operations Research Letters*, 44:96–100.
- Trick, M. A. (2000). A schedule-then-break approach to sports timetabling. In *Proceedings of the 3rd International Conference on the Practice and Theory of Automated Timetabling*, pages 242–253.
- Van Bulck, D., Goossens, D., Schönberger, J., and Guajardo, M. (2020). RobinX: a three-field classification and unified data format for round-robin sports timetabling. *European Journal of Operational Research*, 270:568–580.
- Van Bulck, D., Goossens, D. R., and Spieksma, F. C. R. (2019). Scheduling a non-professional indoor football league: a tabu search based approach. *Annals of Operations Research*, 275:715–730.
- Yi, X., Goossens, D., and Nobibon, F. T. (2019). Proactive and reactive strategies for football league timetabling. *European Journal of Operational Research*. In press, DOI:10.1016/j.ejor.2019.09.038.