ECAI 2024
U. Endriss et al. (Eds.)

doi:10.3233/FAIA240606

Temporal Fairness in Decision Making Problems

Manuel R. Torres[a,*,1], **Parisa Zehtabi**[a,1], **Michael Cashmore**[a], **Daniele Magazzeni**[a] and **Manuela Veloso**[a]

[a]J.P. Morgan AI Research

Abstract. In this work we consider a new interpretation of fairness in decision making problems. Building upon existing fairness formulations, we focus on how to reason over fairness from a temporal perspective, taking into account the fairness of a history of past decisions. After introducing the concept of temporal fairness, we propose three approaches that incorporate temporal fairness in decision making problems formulated as optimization problems. We present a qualitative evaluation of our approach in four different domains and compare the solutions against a baseline approach that does not consider the temporal aspect of fairness.

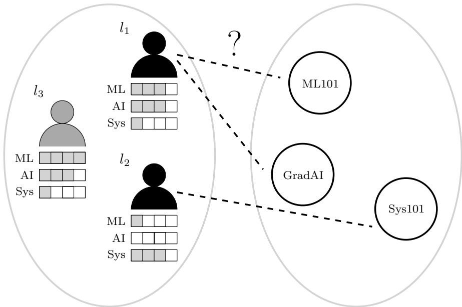

Figure 1: Depicts a *university course assignment*, in which courses are to be assigned to lecturers with different levels of expertise in different areas. The lecturer in gray is on sabbatical leave.

1 Introduction

Automated decision making is an important part of artificial intelligence with a variety of application areas, from scheduling and resource allocation, to robotics and autonomous vehicles. Decision making processes typically aim to optimize an overall benefit or cost. However, as we strive to make our algorithms and agents more intelligent, it is important to ensure that they also account for ethical considerations such as fairness. The need for fair algorithms and agents has been widely studied across different areas, such as robotics [3], healthcare [4], telecommunications [11], and resource allocation [15], among others.

Formulating fairness concerns in different domains can be challenging, and it has been the subject of many studies [21]. In this paper, we take a new angle to considering fairness in decision making processes. We build upon previous fairness formulations, and focus on how to reason about fairness from a temporal perspective, accounting for the fairness of a history of past decisions. We aim to introduce the concept of *"temporal fairness"* into the decision making process, which measures the fairness of solutions throughout time.

As a motivating example consider the scenario depicted in Figure 1 where courses must be assigned to a pool of lecturers (l_1, l_2 and l_3) in semester t. Each lecturer is specialized in different areas and the teaching quality of a course is proportional to the expertise of its lecturer (the gray bars below the lecturers depict their expertise on different topics). Figure 2a depicts the number of courses assigned to each lecturer in the past four semesters. Lecturer l_1 has received a higher teaching load than l_2 over the past four semesters. Regardless of the reasons that have led to the scenario in Figure 2a, the reality is that there has been some *"historical unfairness"*. Figure 2b depicts the cumulative teaching load over time for each of the two available lecturers l_1 and l_2.

A new course allocation must be made for semester t. If an allocation is made that is presently fair, in which both lecturers teach the same number of courses (dashed gray scenario in Figure 2b), an overall *"temporal unfairness"* remains, with the gap in cumulative lecturing load not reducing. In fact, as depicted in Figure 2c, even in the case where an unfair allocation is made for semester t and l_2 lectures all courses, there would still exist a gap in the lecturing load.

In this work, we focus on the problem of decision making while accounting for a historical fairness and considering the impact of future decisions in overall temporal fairness.

We first introduce the definition of an optimization problem that reasons over the trade-off between quality and fairness (Section 2.1). Then we introduce the concept of temporal fairness by including historical fairness into our formulation (Sections 2.2 and 2.3). We then extend this optimization problem to reason over this trade-off while accounting for future predictions and forecasts (Section 2.4). This allows the generation of solutions that may look unfair in the short term, but fairer when analyzed over a longer period of time into the future. We incorporate the notion of temporal fairness via the introduction of a framework for fairness metrics that considers historical solutions.

The main contributions of this paper are: (i) introducing the concept of *temporal fairness* in decision making problems, (ii) a formulation for addressing historical unfairness from past solutions, (iii) a formulation for both addressing historical unfairness and considering future historical fairness, and (iv) a qualitative evaluation on different domains that examines the differences between solutions generated with and without considering the temporal aspect of fairness.

The remainder of the paper is structured as follows. Section 2 introduces the formulations of optimization problems that consider fairness, as described above. We then present the qualitative evaluation of our framework in Section 3. Section 4 discusses relevant related work, and the paper concludes in Section 5 with final remarks and a discussion on avenues for future work.

* Corresponding Author. Email: manuel.torres@jpmchase.com.
[1] Equal contribution.

Courses lectured in past 4 semesters

	$t-4$	$t-3$	$t-2$	$t-1$
l_1	2	1.5	3	2
l_2	1	1.5	0	1

(a)

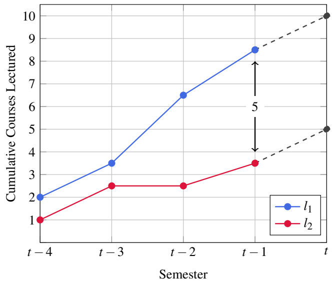

(b)

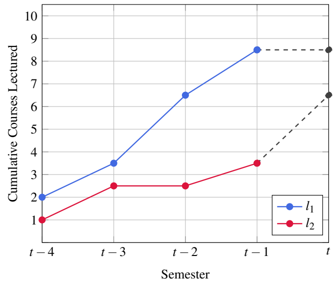

(c)

Figure 2: Depicts a concrete scenario of a university course assignment. Figure 2a depicts the scenario in which l_1 was assigned a higher lecturing load over the past 4 semesters. Figure 2b and 2c depict the cumulative lecturing load of l_1 (blue) and l_2 (red) with possible solutions to the current allocation problem (dashed gray). 2b depicts the case where at timestep t a fair allocation is made with both lecturers given the same load (1.5 courses each). 2c depicts the case where an unfair allocation is made and l_2 is assigned all the lecturing load (3 courses).

2 Problem Formulation

We consider decision making problems solved by finding a solution that maximizes an objective function while satisfying a set of constraints. First, we consider a decision making problem where no fairness metric is considered. Such problems can be formulated as an Optimization Problem (OP).

Definition 1. *An **Optimization Problem (OP)** is a tuple $\langle Q, \mathscr{X}, \mathscr{C}\rangle$ where Q is a quality metric, $\mathscr{X}$ is the domain for optimization, and $\mathscr{C}$ is the set of constraints.*

Formally, we define an OP as:

$$\begin{aligned} \max_{x\in\mathscr{X}} \quad & Q(x) \\ \text{s.t.} \quad & \mathscr{C}(x) \end{aligned} \tag{1}$$

In this setting, the goal is to find a solution x^* that maximizes a given *quality metric* Q, while being subject to a set of constraints $\mathscr{C}(x)$. We let variables x denote the optimization variables of the problem. Since this formulation only reasons over the quality metric, it is possible that the optimal solutions may be deemed unfair according to some fairness metric. Moreover, as depicted in Figure 2a, this formulation may lead to a fast accumulation of unfair solutions.

2.1 FOP: Incorporating Fairness

We now incorporate a fairness metric F into the formulation of the optimization problem. A Fair Optimization Problem (FOP) can be defined as:

Definition 2. *A **Fair Optimization Problem (FOP)** is a tuple $\langle Q, F, \mathscr{X}, \mathscr{C}, \beta\rangle$ where F is the fairness metric and $\beta \in \mathbb{R}$ is a parameter than controls the trade-off between quality and fairness. The remaining elements follow the original OP.*

Formally, an FOP can be modelled as:

$$\begin{aligned} \max_{x\in\mathscr{X}} \quad & Q(x) + \beta F(x) \\ \text{s.t.} \quad & \mathscr{C}(x) \end{aligned} \tag{2}$$

In general, we will assume that F returns higher values for fair solutions and lower values for unfair solutions. In practice, it may be convenient for both Q and F to have well-specified ranges, rendering it easier to understand the impact of the parameter β. However, the formulation is general and supports arbitrary quality and fairness metrics. Finally, we note that the specification of the fairness metric F may potentially require the introduction/modification of constraints. In order to keep notation simple, we will continue denoting the set of constraints as before, $\mathscr{C}(x)$.

As an example building upon our previous scenario of the course assignment domain, let us consider a *relative max-min* fairness metric F^{rmm}, which compares the maximum and minimum number of courses lectured by all lecturers, versus the total number of courses lectured during that time. Formally,

$$F^{\text{rmm}}(x) = 1 - \frac{\max_i S_i(x) - \min_j S_j(x)}{S(x)},$$

where $S_i(x)$ is the number of courses lectured by l_i in solution x, and $S(x)$ is the total number of courses lectured. The range of F^{rmm} is $[0,1]$. It is maximized when lecturers get an equal lecturing load, and minimized when one of the lecturers takes the entire load. While incorporating the new fairness metric in the FOP leads to solutions that are fair according to F (or at least fairer, depending on β), there may still exist some historical unfairness that remains from previous allocations. Figures 2a and 2b hinted at this, depicting a scenario where scheduling a fair plan at time step t would have maintained the gap of cumulative courses lectured.

2.2 HFOP: Incorporating Historical Fairness

FOPs assume a fairness metric F that only reasons over the fairness of a solution x. In order to account for existing historical unfairness, it is thus important to reason over the fairness of a solution x in the context of the history of past solutions $H = (x_{t-T}, \ldots, x_{t-1})$, where $x_{t-\Delta}$ is a previous solution from time step $t - \Delta$.

We formalize the notion of such fairness metrics in the following definition.

Definition 3. *A* historical fairness *metric $F_H : \mathscr{X} \to \mathbb{R}$ is a fairness metric for an FOP $\langle Q, F_H, \mathscr{X}, \mathscr{C}, \beta\rangle$ where $H \subseteq \mathscr{X}$ contains solutions satisfying $\mathscr{C}$. We assume $F_\emptyset$ is a fairness metric and write it as F.*

The historical fairness metric F_H can be used to control how fast or slow historical unfairness is compensated. Also, as it is not a limitation for the real-world scenarios we consider in this paper, we assume the time between the historical solutions in H is uniform.

Definition 4. *A **Historical Fair Optimization Problem (HFOP)** is an FOP tuple $\langle Q, F_H, H, \mathscr{X}, \mathscr{C}, \beta\rangle$ where F_H is a historical fairness metric.*

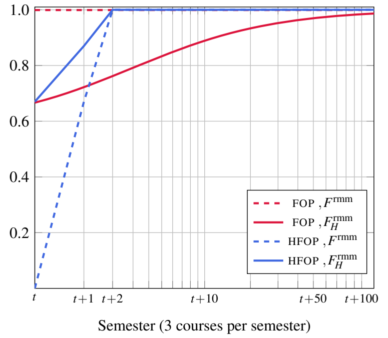

Figure 3: Compares the fairness of the solutions computed using a FOP (red) and HFOP (blue) in a simple scenario with no quality metric Q. Each solution is evaluated with both the *relative max-min* fairness metric F^{rmm}, and its historical variant F_H^{rmm}, shown respectively in dashed and solid lines. The x axis is in log scale.

Formally, we can formulate an HFOP as:

$$\begin{aligned} \max_{x\in\mathscr{X}} \quad & Q(x)+\beta F_H(x) \\ \text{s.t.} \quad & \mathscr{C}(x) \end{aligned} \tag{3}$$

As before, parameter β provides control over the quality/fairness trade-off, with higher values of β leading to a faster compensation of historical unfairness. It is worth highlighting that the optimal solution to an HFOP may actually be an unfair solution from the perspective of a fairness metric F. To see this, let us consider an example.

Building upon the *relative max-min* fairness metric previously discussed, we can now consider its historical variant F_H^{rmm}, where we reason instead over the courses lectured across (H,x)—the concatenation of historical solutions in H with the new solution x.

$$F_H^{\mathrm{rmm}}(x) = 1 - \frac{\max_i S_i((H,x)) - \min_j S_j((H,x))}{S((H,x))},$$

where $S_i((H,x))$ is the number of courses lectured by l_i over all solutions in (H,x) and $S((H,x))$ is the total number of courses lectured.

Let's now see how F^{rmm} and F_H^{rmm} would differ in a concrete scenario, depicted in Table 1. Assume that at time step t we take a solution $x_{(1.5,\,1.5)}$ assigning an equal load of 1.5 courses to each lecturer. While $F^{\mathrm{rmm}}(x_{(1.5,1.5)}) = 1$, we have that $F_H^{\mathrm{rmm}}(x_{(1.5,1.5)}) = 1 - \frac{5}{15} \sim 0.67$. Since the allocation given by $x_{(1.5,1.5)}$ is balanced, the historical max-min gap remains 5 while the total number of courses becomes 15. On the other hand, if we take the solution $x_{(0,3)}$ assigning all 3 courses to lecturer l_2, we would have $F^{\mathrm{rmm}}(x_{(0,3)}) = 0$ and $F_H^{\mathrm{rmm}}(x_{(0,3)}) = 1 - \frac{2}{15} \sim 0.87$ – since all the lecturing load was assigned to l_2, the max-min gap decreases to 2.

Table 1: Compares F and F_H for different solutions, assuming a history H (repeated from Figure 2a). A solution $x_{(m,n)}$ refers to the case where lecturers l_1 and l_2 are assigned m and n courses, respectively. We observe that, under history H, a perfectly fair solution according to F may not be fair according to F_H. The solutions that would be picked under F and F_H are in bold.

	H			
	$t-4$	$t-3$	$t-2$	$t-1$
l_1	2	1.5	3	2
l_2	1	1.5	0	1

x_t	$F(x_t)$	$F_H(x_t)$
$x_{(1.5,1.5)}$	**1.00**	0.67
$x_{(1,2)}$	0.67	0.73
$x_{(0,3)}$	0.00	**0.87**

More generally, it is interesting to compare solutions computed by FOP *vs.* HFOP, and the respective fairness metric F^{rmm} and historical fairness metric F_H^{rmm}. Figure 3 depicts these metrics under a simple course allocation scenario where we assume there exists no quality metric Q and 3 courses per semester. As expected, we observe that the fairness metric F^{rmm} of the solutions computed by FOP is always maximized. In this case, FOP always returns $x_{(1.5,1.5)}$. HFOP, on the other hand, starts by computing unfair solutions $x_{(0,3)}$ and $x_{(0.5,2.5)}$ at time steps t and $t+1$. From $t+2$ onward, HFOP returns the fair solution $x_{(1.5,1.5)}$. These different choices for the solutions have a significant impact on the way the historical unfairness is compensated. Whereas HFOP maximizes the historical fairness in two time steps, we observe that after 10 semesters (or 30 courses) FOP only reaches a value of 0.9. As anticipated in the end of the previous section, we conclude that FOP takes a long time to compensate existing historical unfairness.

2.3 DHFOP: Incorporating Discounted Historical Fairness

We observed in Figure 3 how slowly the historical unfairness would be compensated when following the solutions produced by FOP. In fact, it turns out it would never be fully compensated – since FOP always computes perfectly balanced schedules, the lecturing load gap of 5 would remain unchanged. This behaviour may not fit many domains. It may become especially problematic when considering scenarios with long histories of unfair solutions.

A historical fairness metric F_H allows the specification of an optimization problem HFOP that reasons over remnant historical unfairness. We observed this may lead to solutions that seem unfair at time step t when only considering the current time step (i.e., according to F).

In practice, it makes sense to consider a "forgetting rate phenomenon", where we attribute more importance to recent events than those in a distant past. However, $F_H(x)$ puts equal importance to the fairness of solution x at time step t and all past solutions. In order to model the importance of recent events we propose the discounted historical fairness metric $F_{H,\gamma}$, which discounts past unfairness with a forgetting discount factor γ.

Definition 5. *A **Discounted Historical Fair Optimization Problem (DHFOP)** is a tuple $\langle Q, F_{H,\gamma}, H, \mathscr{X}, \mathscr{C}, \beta\rangle$, where $F_{H,\gamma}$ is a historical fairness metric that reasons over a history of previous solutions $H = (x_{t-T}, \ldots, x_{t-1})$ with the discount factor γ. The remaining elements follow the* HFOP.

Reasoning over a discounted historical fairness metric $F_{H,\gamma}$ allows us to control the importance of the unfairness of past solutions relative to more recent ones. Formally, the optimization problem is:

$$\begin{aligned} \max_{x\in\mathscr{X}} \quad & Q(x)+\beta F_{H,\gamma}(x) \\ \text{s.t.} \quad & \mathscr{C}(x) \end{aligned} \tag{4}$$

There are now two hyper-parameters. The discount factor γ, which sets the importance of the fairness of past solutions, and the parameter β, which controls the quality/fairness trade-off.

Revisiting once more our running example of course assignment and the max-min fairness metric, we could define its discounted historical variant as follows:

$$F_{H,\gamma}^{\mathrm{rmm}}(x) = 1 - \frac{\max_i S_{\gamma,i}((H,x)) - \min_j S_{\gamma,j}((H,x))}{S_\gamma((H,x))},$$

where $S_{\gamma,i}((H,x)) = \sum_{\Delta=0}^{T} \gamma^\Delta S_i(x_{t-\Delta})$ is the discounted number of courses lectured by l_i over all solutions in (H,x) and similarly,

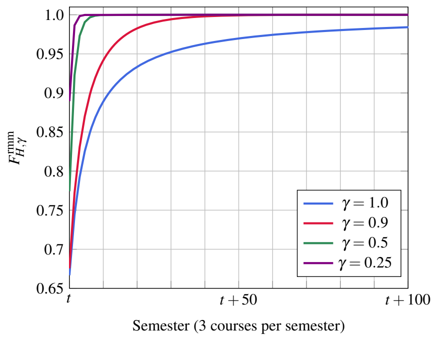

Figure 4: Depicts the discounted historical *relative max-min* fairness metric $F^{\text{rmm}}_{H,\gamma}$ in the setting where from time step $t = 0$ onwards we schedule perfectly balanced loads. The smaller the discount factor γ, the faster the historical fairness is compensated.

$S_\gamma((H,x))$ is the discounted total number of courses lectured:

$$S_\gamma((H,x)) = \sum_{\Delta=0}^{T} \gamma^\Delta S(x_{t-\Delta})$$

We now analyze the behavior of $F^{\text{rmm}}_{H,\gamma}$ for different values of γ. We build upon our course assignment example introduced in the previous section, assuming that at time step t and onward we accept a solution that assigns an equal load to each lecturer (the solution that the FOP would compute). Figure 4 depicts $F^{\text{rmm}}_{H,\gamma}$ for different values of γ. We observe that smaller values of γ lead to a faster compensation of historical fairness. For example, for γ values of 0.25, 0.5, and 0.9, it takes, 2, 5, and 26 semesters, respectively, for $F^{\text{rmm}}_{H,\gamma}$ to reach a value of 0.99.

2.4 MSDHFOP: Historically Fair Planning with Future Forecasts

All problems introduced so far are of a single-shot nature, where the solver is assumed to make a decision for the current time step t. However, single-shot decisions can often result in sub-optimal solutions in complex domains with extended horizons, such as *planning problems* [8]. Reasoning over multiple time steps into the future can allow for more effective solutions given knowledge or predictions about future events.

In the setting of fairness, reasoning over multiple steps into the future may allow for interesting solutions. For example, due to future constraints, a solution that is fair over a given horizon may require initial solutions that seem unfair when analyzed independently. In order to account for both existing historical unfairness and a planning horizon into the future, we let $F_{H,\gamma,\tau}(x_t,\dots,x_{t+T_F})$ denote a fairness metric that considers both a history of solutions H and a sequence of T_F future problems $(x_t,\dots,x_{t+T_F})$, with the past and the future being discounted according to γ and τ.

Definition 6. *A **Multi Step Historical Fair Optimization Problem (MSDHFOP)** is a tuple $\langle Q, F_{H,\gamma,\tau}, H, \mathscr{X}, \mathscr{C}, \beta, \gamma, \tau\rangle$, where $F_{H,\gamma,\tau}$ is a historical fairness metric that reasons over a history of previous solutions $H = (x_{t-T_H},\dots,x_{t-1})$ and a sequence of future solutions $(x_t,\dots,x_{t+T_F})$. γ and τ are discount factors. The remaining elements follow the* DHFOP.

We formulate an MSDHFOP as:

$$\begin{aligned} \max_{x_0,\dots,x_{T_F}} \quad & \textstyle\sum_{t=0}^{T} \tau^t Q(x_t) + \beta F_{H,\gamma,\tau}(x_0,\dots,x_{T_F}) \\ \text{s.t.} \quad & \mathscr{C}(x) \end{aligned} \tag{5}$$

The first term computes the discounted sum of quality of the planned solutions. The second term computes the multi-step historical fairness metric. As before, the discount factor γ sets the importance of past solutions relative to more recent ones in the computation of fairness. Similarly, the discount factor τ discounts future solutions relative to the previous one, impacting both fairness and solution quality. Whereas γ seeks to model the "recency effect" from a fairness perspective (i.e., we tend to attribute more importance to recent events than those in a distant past), τ seeks to model uncertainty in planning into the future (i.e., it is easier to predict states closer in time than those in a distant future).

We can again build upon the discounted historical relative max-min fairness metric, and introduce a variant that also reasons over the next T_F solutions $x_{t:T_F} = (x_t,\dots,x_{t+T_F})$, where we define

$$F^{\text{rmm}}_{H,\gamma,\tau}(x_{t:T_F}) = 1 - \frac{\max_i S_{\gamma,\tau,i}((H,x_{t:T_F})) - \min_j S_{\gamma,j}((H,x_{t:T_F}))}{S_\gamma((H,x_{t:T_F}))},$$

where

$$S_{\gamma,\tau,i}((H,x_{t:T_F})) = \sum_{\Delta=1}^{T_H} \gamma^\Delta S_i(x_{t-\Delta}) + \sum_{\Delta=0}^{T_F} \tau^\Delta S_i(x_{t+\Delta})$$

is the discounted number of courses lectured by l_i, over all solutions in history H and future planned solutions $x_{t:T_F}$. Similarly, $S_{\gamma,\tau}((H,x_{t:T_F}))$ is the discounted total number of courses lectured.

Table 2: Builds upon the results in Table 1, reporting $F_{H,\gamma,\tau}$ for different solutions, assuming a history as depicted in Figure 2a, and $\gamma = \tau = 1$. From Table 1 we observed that a single-shot decision based on F_H would pick solution $x_{(0,3)}$. However, if we are now aware there exists a constraint preventing l_1 from teaching any course at time step $t+1$, then the best decision is to first choose solution $x_{(0.5,2.5)}$.

x_t	x_{t+1}	$F_{H,\gamma,\tau}((x_t,x_{t+1}))$
$x_{(0,3)}$	$x_{(0,3)}$	0.94
$x_{(0.5,2.5)}$	$x_{(0,3)}$	1.00
$x_{(1.5,1.5)}$	$x_{(0,3)}$	0.88

Table 2 depicts an example where planning multiple steps into the future can lead to better solutions. This example builds upon our analysis of Table 1 from which we concluded the optimal solution according to F_H is to assign at time step t all the lecturing load to l_2. However, suppose now we are allowed to plan over a horizon $T = 2$ into the future, and that we are aware of a constraint preventing l_1 from lecturing any courses in the second semester $t+1$. From Table 2 we conclude the best sequence of actions is actually $(x_{(0.5,2.5)}, x_{(0,3)})$. Following HFOP instead would yield to the less rewarding solution $(x_{(0,3)}, x_{(0,3)})$.

3 Experimental Evaluation

3.1 Setup

We evaluate our formulations across multiple domains using different fairness metrics. We start with a technical description of each domain, introducing the decision variables, and the quality and fairness metrics to be used. The machine used to run experiments is an Intel(R) Xeon(R) CPU E3-1585L v5 @ 3.00GHz with 64GB of RAM.

3.1.1 Course Assignment Problem (CAP)

This is the domain that has been used throughout the paper, where a set of lecturers $\mathscr{L}$ is to be assigned to a set of courses $\mathscr{C}$. When dealing with multi-step decision making settings, we may denote the set of courses at time step t as $\mathscr{C}_t$. The expertise of lecturer l in course c is measured by $S : \mathscr{L} \times \mathscr{C} \to \mathbb{R}$, and higher values correspond to higher expertise. Decision variable $x_{l,c} \in \{0, 0.5, 1\}$ indicate the load of lecturer l in teaching course c— a lecturer may not lecture the course at all, or lecture either half a course or the full course.

We consider a quality metric $Q = \frac{1}{Q_{\max}} \sum_{c \in \mathscr{C}_t} \sum_{l \in \mathscr{L}} x_{l,c}\ S(l,c)$, which rewards course assignments with skilled lecturers. $Q_{\max}$ is a normalization constant, denoting the maximum sum of expertise possible—this ensures Q is bounded between 0 and 1. In order to display the generality of our formulation, throughout the experimental evaluation with this domain we may use different fairness metrics.

3.1.2 Vehicle Routing Problem (VRP)

In this domain, given a set of vehicles V, a set of points P that must all be traveled to exactly once, a depot $r \in P$ the vehicles must leave from and return to, and distances between all points $D : P \times P \to \mathbb{R}_+$, determine a route for each vehicle that minimizes the total distance traveled. We consider a standard integer program to model the OP where quality $Q(x)$ is the total distance traveled (see Supplementary Materials in [20] for a full definition). To model fairness, for a given solution x to the integer program, let $U_v(x)$ be the total distance vehicle v travels under x and define $F(x) := \max_{v \in A} U_v(x) - \min_{w \in A} U_w(x)$. This notion of fairness is similar to proportional equality and is used in [13] in a multi-objective version of VRP.

3.1.3 Task Allocation Problem (TAP)

In this domain, given a set of agents A, a set of tasks T, and a cost associated with each agent for each task $C : A \times T \to \mathbb{R}_+$, find an assignment of tasks to agents such that the sum of costs is minimized. We consider a standard integer program to model the OP where quality $Q(x)$ is the sum of costs (see Supplementary Materials in [20] for a full definition). The fairness metric we consider is the classic minimax notion of fairness (see [21] and references therein), where for a given solution x to the integer program, let $U_a(x)$ be the total cost agent a incurs under x and define $F(x) := \max_{a \in A} U_a(x)$.

3.1.4 Nurse Scheduling Problem (NSP)

We consider a version of this classical problem in operations research. In the Supplementary Materials in [20], the authors first formally define the problem and then we show the impact of different histories on the NSP, in particular showcasing the impact of the discount factor in DHFOP.

3.2 Quality vs. Fairness

We start our experimental evaluation with an example depicting the quality *vs.* fairness trade-off, and the impact of the parameter β therein. Let's consider an instance of CAP with 3 lecturers, l_1, l_2, l_3 and 2 courses c_1, c_2. Across all 3 courses, l_1 has high skills ($S = 2$), l_2 has medium skills ($S = 1.5$), and l_3 has low skills ($S = 0$). We use a historical *quadratic max-min gap* fairness metric

$$F_H^{\text{qmmg}}(x) = -\left(\frac{1}{2}\ (\max_i S_i((H,x)) - \min_j S_j((H,x)))\right)^2,$$

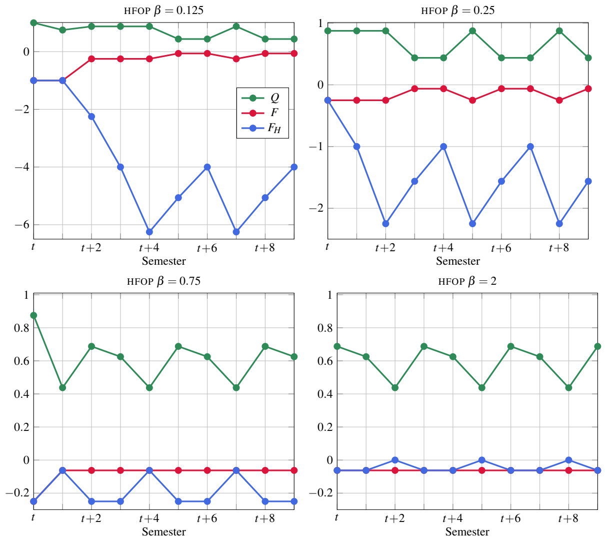

Figure 5: The quality Q, fairness F^{qmmg}, and historical fairness F_H^{qmmg} of the solutions computed by HFOP under different values of β, in the course assignment domain.

where $S_{\gamma,i}$ is defined as previously. The original fairness metric that disregards H follows naturally. F^{qmmg} has range $[-\infty, 0]$ and, when compared to F^{rmm}, should allow for heavier penalization of solutions that increase the lecturing load, due to the quadratic term and the lack of normalization.

We analyzed the solutions computed by HFOP under different values of β, for 10 consecutive semesters, starting with no previous history. Figure 5 compares the quality Q, fairness F^{qmmg}, and historical fairness F_H^{qmmg} of the solutions computed by HFOP under different values of β. Table 3 provides a summary of the results for the different metrics. We observe that, as β increases, HFOP computes solutions with lower quality Q, but higher fairness F and historical fairness F_H. This follows our expectation, since β is the parameter setting the quality *vs.* fairness trade-off. From the figure we also observe that HFOP converges to a pattern of first selecting higher-quality/lower-fairness solutions, and once the historical fairness reaches a certain (low) level, starts selecting lower-quality/higher-fairness solutions to compensate for it. This pattern is even more noticeable for lower values of β, where at earlier time steps the solutions produced tend to be characterized by high quality and low fairness.

3.3 Planning with Future Forecasts

We now evaluate the benefits from a fairness perspective of MSD-HFOP reasoning over multiple steps into the future from a fairness perspective. Consider a simplified instance of CAP with two lecturers l_1 and l_2, and two courses c_1 and c_2. Across all courses, l_1 has high skills ($S = 2$) and l_2 has medium skills ($S = 1$). To showcase the flexibility of our approach to fairness metrics, we now consider another version of the *maximin* fairness metric where utility U_i measures the number of courses taught by lecturer l_i:

$$F^{\text{mm}}(x) = \frac{\min_i U_i(x)}{\max_j U_j(x)}.$$

Assuming no historical solutions, the scheduler is now to plan the course assignments for the next $T = 4$ semesters. There exists a known constraint about the future—l_1 will not be able to take any lecturing load on semesters $t+2$ and $t+3$ due to a sabbatical leave.

Table 3: Summary of the results on the comparison of the quality Q, fairness F^{qmmg}, and historical fairness F_H^{qmmg} of the solutions computed by HFOP under different values of β.

		Q			F^{qmmg}			F_H^{qmmg}		
		max	min	$\mu \pm \sigma$	max	min	$\mu \pm \sigma$	max	min	$\mu \pm \sigma$
	OP	1.0	1.0	1.0 ± 0	-1.0	-1.0	-1.0 ± 0.0	-1.0	-100.0	-38.5 ± 2.42
HFOP	β									
	0.125	1.0	0.44	0.7 ± 0.22	-0.06	-1.0	-0.33 ± 0.35	-1.0	-6.25	-3.89 ± 1.83
	0.25	0.88	0.44	0.66 ± 0.22	-0.06	-0.25	-0.16 ± 0.09	-0.25	-2.25	-1.47 ± 0.63
	0.75	0.88	0.44	0.61 ± 0.13	-0.06	-0.25	-0.08 ± 0.06	-0.06	-0.25	-0.19 ± 0.09
	2	0.69	0.44	0.59 ± 0.11	-0.06	-0.06	-0.06 ± 0.0	-0	-0.06	-0.04 ± 0.03

Table 4 depicts the solutions computed by HFOP and MSDHFOP, in a setting with $\beta = 2$ and discount factors $\gamma = 1, \tau = 1$. Since HFOP plans a single step at a time, it is not able to take advantage of the information on l_1's future constraints. As a result, it schedules the perfectly balanced solution in the two initial steps, and is then forced to schedule the last two time steps as $x_{(0,2)}$. This results in a sequence of solutions leading to an overall lower quality (10 *vs.* 12) and fairness (0.33 *vs.* 1).

Table 4: Comparison between FOP and MSDHFOP, showing the benefits of reasoning over multiple steps into the future. Since HFOP performs single-shot decisions, it ends up ignoring the known constraint that l_1 will not lecture any course in time steps $t+2$ and $t+3$.

	x_t	x_{t+1}	x_{t+2}	x_{t+3}	$\sum_t Q(x_t)$	$F_{H,\gamma,\tau}(x)$
HFOP	$x_{(1,1)}$	$x_{(1,1)}$	$x_{(0,2)}$	$x_{(0,2)}$	10	0.33
MSDHFOP	$x_{(2,0)}$	$x_{(2,0)}$	$x_{(0,2)}$	$x_{(0,2)}$	12	1.0

3.4 Increasing Complexity and Benchmarking

We now examine a more complex problem and show the impact of considering fairness on running time. We first introduce a method for generating random instances and a history of past solutions for VRP. Consider a square integer grid of a fixed size. We deterministically place the depot at the center of the grid and, given a fixed number of points n, choose n of the grid points uniformly at random (not including the depot). For generating history, we generate random instances and solve the problem optimally on these random instances.

For our experiments in this section and the following, we implemented the integer program with the corresponding fairness constraints using the PuLP Python library [17] and used the CBC solver [6] with a standard linearization of $F(x)$ (see, e.g., [21]). All times measured are wall-clock times for the combined model-building and solving times.

In our experiments, we consider 4 vehicles $V = \{V_1, V_2, V_3, V_4\}$ and 12 locations. We generate a history of 5 steps and a single random instance. For each historical instance, we assume V_1 always had the shortest route, V_2 the second shortest, and similarly for V_3 and V_4. Table 5a shows the total distance traveled for each vehicle. We compare the solutions of OP, FOP, and HFOP. We set $\beta = 10$ for all experiments. Table 5b shows the results for each of the 4 vehicles.

In FOP, the notion of fairness considered should encourage solutions where all distances traveled are similar. This, of course, should come at the expense of increasing the overall distance traveled. We see this exact scenario play out when comparing OP and FOP. The distances in the solution for OP are not uniform but attain a total distance of 74.4 and the distances in the solution for FOP are all similar but the total distance traveled is 112.1. When comparing OP and FOP to HFOP, we expect that HFOP should account for the historical unfairness received by vehicle V_4. In fact, we expect and see in the results that the solution to HFOP should give the shortest routes (in order) to V_4, V_3, V_2 and V_1. In terms of total time, FOP and HFOP require roughly 3 times as long to run, thus showing that the cost of incorporating fairness in our framework is not computationally prohibitive for VRP.

Table 5: Results for VRP experiments.

V_1	V_2	V_3	V_4
37.1	44.9	154.4	202.6

(a) Total historical distance traveled.

	V_1	V_2	V_3	V_4	time (s)
OP	4.5	6.3	13.8	49.8	11.6
FOP	26.6	28.3	28.6	28.6	34.2
HFOP	74.5	65.3	6.32	4.47	33.9

(b) Distance traveled for all 4 vehicles on a single instance.

3.5 Larger Scale Experimentation

In this section, we show that our framework can be applied on a larger scale than considered in the previous sections. We introduce a method for generating random instances for TAP. We create instances with $|A| = |T| = 40$. For each agent $a \in A$, one task is chosen uniformly at random to have cost 5, three tasks are chosen uniformly at random to have cost 20, and the rest of the tasks have cost 30. We sometimes deterministically enforce that an agent a does not have a task of cost 5 and this task is replaced with a cost 30 task, in which case we say agent a is *constrained*.

We refer to the following setup as a single *run* and we average our results over 10 runs. Sample 8 agents uniformly at random from A and denote this subset as C. Produce 3 random instances according to our random instance generation given above. Then produce 3 more random instances where all agents in C are constrained. These 6 instances are the future instances. To generate history, we run the OP on each of these 6 instances. Sort the agents according to total cost. In this order, the last 4 agents not in C are assigned a historical cost of 180, call these agents W. Amongst the remaining agents, the first 24 agents are assigned a historical cost of 30. The remaining 12 agents are assigned a historical cost of 120. More justification for random instance generation and history generation is provided in the Supplementary Materials in [20]. All instances are run with $\beta = 10$.

We first evaluate the maximum cost assigned to any agent in OP and FOP. Table 6 shows the number of times per run the maximum cost is 30. We expect the number to be much larger in OP compared to FOP, which is confirmed in Table 6. This is at the expense of incurring a larger total cost, which is expected since FOP is also prioritizing minimizing the maximum cost and not just the total cost.

Table 6: Evaluating OP, FOP, HFOP for TAP experiments. (First column) Measures the number of times the maximum cost is 30—the largest value is 6 as there are 6 instances per run. (Last two columns) For each run, we compute the total cost for each agent in W and average it by $|W| = 4$, similarly for $A \setminus W$. All values averaged over 10 runs.

	max cost=30	avg sum of costs	avg cost of W	avg cost of $A \setminus W$
OP	5.9 ± 0.3	470.8 ± 12.1	97.6 ± 7.8	67.6 ± 1.7
FOP	2.4 ± 0.8	484.0 ± 13.2	86.7 ± 9.3	71.0 ± 2.2
HFOP	6.0 ± 0.0	478.2 ± 12.5	50.7 ± 7.7	74.0 ± 1.9

Table 7: Evaluating the impact of the constrained agents on the TAP experiments. For each run, we compute the total cost for each agent in C and average it by $|C| = 8$. We compare the results over the first 3 instances and last 3 instances of each run. The times for each run are averaged over the 6 instances. All values averaged over 10 runs.

	avg cost of C (first 3)	avg cost of C (last 3)	time (s)
OP	32.5 ± 4.4	67.8 ± 1.7	0.5 ± 0.1
FOP	32.3 ± 3.2	62.1 ± 2.1	54.4 ± 42.5
HFOP	36.6 ± 4.5	67.3 ± 1.9	0.6 ± 0.1
MSDHFOP	23.3 ± 3.7	64 ± 0.9	4.2 ± 1.2

We next evaluate the impact of history in OP, FOP, and HFOP. Recall that for each run, W is the set of agents who received the largest historical cost. OP and FOP do not consider history and therefore will not necessarily prioritize the agents in W. We see this exact behavior in Table 6. Furthermore, even as OP continues to not prioritize the agents in W, which is how W is defined, the average cost of agents not in W for OP is still less than for FOP or HFOP. This is expected, at least for HFOP, as HFOP prioritizes agents in W.

We now examine the outcomes of the constrained agents C, who are all constrained in the last 3 instances of each run. Table 7 shows that the constrained agents in OP, FOP, and HFOP all have a similar average over both the first and second 3 instances, which is expected as these agents should not necessarily receive special treatment in any of these frameworks. However, MSDHFOP reasons about the future, and therefore we expect it to adjust for the fact that the agents in C are constrained over the last 3 instances. We see this behavior in Table 7. Note that we set $\gamma = \tau = 0.75$.

We also report the running times in Table 7. The running times of HFOP and MSDHFOP are about the same as OP, especially considering MSDHFOP runs all 6 future instances at once. FOP, however, has a large average running time, but the median running time is only 1.6 s—some instances require large amounts of time but are not common. One hypothesis is that the solver we use takes a lot of time when trying to minimize the max cost when there are multiple agents that can achieve the max cost. In HFOP on the first future instance, for example, the historical imbalance ensures that only the 4 agents in W can achieve the max cost, which may reduce the set of candidate optimal solutions considerably.

4 Related Work

In recent times, a significant amount of research has been dedicated to fairness in AI, with a focus on predictive models and algorithmic fairness [14]. In machine learning, in particular, the topic of long-term fairness has been the subject of much attention [5, 7, 10]. The long-term consideration of fairness is relevant to the MSDHFOP formulation, where we consider both future and past history.

Other lines of research look into the connections of algorithmic fairness and ethical decision making in the context of sequential decision making and planning [18]. Nashed et al. explore how each of these settings has articulated its normative concerns, the viability of different techniques for these different settings, and how ideas from one may be useful for the other.

Motivated by computational resource allocation problems, there exists a vast literature on the topic of fairness in real-time scheduling. Examples include the fair scheduling of periodically arriving tasks with deadlines [2], or more generally, the problem of scheduling tasks to long lived processes while taking into account the benefit/cost to each process [1]. While these works look at fairness from a temporal perspective—seeking to ensure a fair load to the different processes—they do not consider possible historical unfairness due to previous solutions and tend to focus on a specific fairness metric.

In the areas of decision-making and planning, several works have focused on different ways to mathematically formulate fairness metrics. Recent work surveys various schemes that have been proposed for formulating ethics-related criteria, including those that integrate efficiency and fairness concerns [21]. They emphasize the challenges of having a single definition of fairness, as different definitions are appropriate for different contexts. Additionally, different fairness models are grouped into clusters, each representing a different type of fairness principle, to facilitate comparisons and help identify the most suitable model for practical applications. While the fairness metrics introduced did not consider fairness from a temporal perspective where a history of past solutions exists, they can be adapted and used as part of all our formulations.

There has also been a growing interest in fairness in multi-agent decision making, planning [19], and reinforcement learning [12, 9]. Recent work focuses on fairness in long-term decision making problems, introducing a new voting formalism that takes the history of previous decisions into account [16]. While the concept of considering history is similar to the definition of HFOP, our formalism considers a centralized decision making process.

5 Conclusion and Future Work

In this work we took a new angle to considering fairness in decision making processes. Building upon previous fairness formulations, we focused on how to reason about fairness from a temporal perspective, especially when there exists a history of past decisions that may have been potentially unfair. In this setting, we proposed to reason over the concept of *"temporal fairness"* in decision making processes.

Starting from a general decision making problem—OP—we incrementally built our approach to reason over temporal fairness, accounting for both past solutions and predictions about the future. With the introduction of a fairness metric in the objective, FOP extends OP by reasoning over the quality/fairness trade-off of a solution. To reason over historical unfairness, we propose HFOP, where the fairness metric takes into account a history of previous solutions. A discounted version DHFOP is then proposed to allow us to model the importance of more recent events. Finally, the MSDHFOP formulation is extended to reason over both historical and future solutions. In the experimental evaluation we assess our approach across different domains and show, in particular, how our approach is compatible with different fairness metrics.

As directions for future work, we envision exploring scenarios where different fairness metrics are used across time (in the past and future) and reasoning over multiple concurrent fairness metrics.

6 Disclaimer

This paper was prepared for informational purposes in part by the Artificial Intelligence Research group of JPMorgan Chase & Co. and its affiliates ("JP Morgan"), and is not a product of the Research Department of JP Morgan. JP Morgan makes no representation and warranty whatsoever and disclaims all liability, for the completeness, accuracy or reliability of the information contained herein. This document is not intended as investment research or investment advice, or a recommendation, offer or solicitation for the purchase or sale of any security, financial instrument, financial product or service, or to be used in any way for evaluating the merits of participating in any transaction, and shall not constitute a solicitation under any jurisdiction or to any person, if such solicitation under such jurisdiction or to such person would be unlawful.

References

[1] M. Ajtai, J. Aspnes, M. Naor, Y. Rabani, L. J. Schulman, and O. Waarts. Fairness in scheduling. *Journal of Algorithms*, 29(2):306–357, 1998.

[2] S. K. Baruah. Fairness in periodic real-time scheduling. In *Proceedings 16th IEEE Real-Time Systems Symposium*, pages 200–209. IEEE, 1995.

[3] M. Brandao, M. Jirotka, H. Webb, and P. Luff. Fair navigation planning: a resource for characterizing and designing fairness in mobile robots. *Artificial Intelligence*, 282:103259, 2020.

[4] I. Y. Chen, E. Pierson, S. Rose, S. Joshi, K. Ferryman, and M. Ghassemi. Ethical machine learning in healthcare. *Annual review of biomedical data science*, 4:123–144, 2021.

[5] A. D'Amour, H. Srinivasan, J. Atwood, P. Baljekar, D. Sculley, and Y. Halpern. Fairness is not static: deeper understanding of long term fairness via simulation studies. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 525–534, 2020.

[6] J. Forrest and R. Lougee-Heimer. Cbc user guide. In *Emerging theory, methods, and applications*, pages 257–277. INFORMS, 2005.

[7] Y. Ge, S. Liu, R. Gao, Y. Xian, Y. Li, X. Zhao, C. Pei, F. Sun, J. Ge, W. Ou, et al. Towards long-term fairness in recommendation. In *Proceedings of the 14th ACM international conference on web search and data mining*, pages 445–453, 2021.

[8] M. Ghallab, D. Nau, and P. Traverso. *Automated planning and acting*. Cambridge University Press, 2016.

[9] N. A. Grupen, B. Selman, and D. D. Lee. Cooperative multi-agent fairness and equivariant policies. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 9350–9359, 2022.

[10] Y. Hu and L. Zhang. Achieving long-term fairness in sequential decision making. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 9549–9557, 2022.

[11] S. Huaizhou, R. V. Prasad, E. Onur, and I. Niemegeers. Fairness in wireless networks: Issues, measures and challenges. *IEEE Communications Surveys & Tutorials*, 16(1):5–24, 2013.

[12] J. Jiang and Z. Lu. Learning fairness in multi-agent systems. *Advances in Neural Information Processing Systems*, 32, 2019.

[13] N. Jozefowiez, F. Semet, and E.-G. Talbi. An evolutionary algorithm for the vehicle routing problem with route balancing. *European Journal of Operational Research*, 195(3):761–769, 2009.

[14] J. Kleinberg, J. Ludwig, S. Mullainathan, and A. Rambachan. Algorithmic fairness. In *Aea papers and proceedings*, volume 108, pages 22–27, 2018.

[15] A. Kumar and J. Kleinberg. Fairness measures for resource allocation. In *Proceedings 41st annual symposium on foundations of computer science*, pages 75–85. IEEE, 2000.

[16] M. Lackner. Perpetual voting: Fairness in long-term decision making. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 2103–2110, 2020.

[17] S. Mitchell, M. OSullivan, and I. Dunning. Pulp: a linear programming toolkit for python. *The University of Auckland, Auckland, New Zealand*, 65, 2011.

[18] S. B. Nashed, J. Svegliato, and S. L. Blodgett. Fairness and sequential decision making: Limits, lessons, and opportunities. *arXiv preprint arXiv:2301.05753*, 2023.

[19] A. Pozanco and D. Borrajo. Fairness in multi-agent planning. *arXiv preprint arXiv:2212.00506*, 2022.

[20] M. R. Torres, P. Zehtabi, M. Cashmore, D. Magazzeni, and M. Veloso. Temporal fairness in decision making problems, 2024. URL https://arxiv.org/abs/2408.13208.

[21] V. Xinying Chen and J. Hooker. A guide to formulating fairness in an optimization model. *Annals of Operations Research*, pages 1–39, 2023.