

# What Hides behind Unfairness?

## Exploring Dynamics Fairness in Reinforcement Learning

Zhihong Deng<sup>1</sup>, Jing Jiang<sup>1</sup>, Guodong Long<sup>1</sup> and Chengqi Zhang<sup>1</sup>

<sup>1</sup>Australian Artificial Intelligence Institute, University of Technology Sydney  
zhi-hong.deng@student.uts.edu.au, {jing.jiang,guodong.long,chengqi.zhang}@uts.edu.au,

### Abstract

In sequential decision-making problems involving sensitive attributes like race and gender, reinforcement learning (RL) agents must carefully consider long-term fairness while maximizing returns. Recent works have proposed many different types of fairness notions, but how unfairness arises in RL problems remains unclear. In this paper, we address this gap in the literature by investigating the sources of inequality through a causal lens. We first analyse the causal relationships governing the data generation process and decompose the effect of sensitive attributes on long-term well-being into distinct components. We then introduce a novel notion called dynamics fairness, which explicitly captures the inequality stemming from environmental dynamics, distinguishing it from those induced by decision-making or inherited from the past. This notion requires evaluating the expected changes in the next state and the reward induced by changing the value of the sensitive attribute while holding everything else constant. To quantitatively evaluate this counterfactual concept, we derive identification formulas that allow us to obtain reliable estimations from data. Extensive experiments demonstrate the effectiveness of the proposed techniques in explaining, detecting, and reducing inequality in reinforcement learning.

## 1 Introduction

Algorithmic fairness has emerged as an increasingly important research topic in the era of autonomous decision-making systems driven by machine learning [Kasy and Abebe, 2021; Mehrabi *et al.*, 2021; Pessach and Shmueli, 2023]. The widespread application of machine learning algorithms in fields like education [Wolff *et al.*, 2013; Anderson *et al.*, 2019], finance [Dwork *et al.*, 2012; Lee *et al.*, 2014], and law [Angwin *et al.*, ; Tolan *et al.*, 2019] carries the risk of unintentionally perpetuating or introducing biases and discrimination against race, gender, or other types of sensitive attributes. Despite numerous approaches developed to address algorithmic fairness, many focus on supervised

learning or decision-making problems without state transitions [Hardt *et al.*, 2016; Zhang *et al.*, 2017; Liu *et al.*, 2018; Patil *et al.*, 2021]. However, fairness is not static [D’Amour *et al.*, 2020]. In reality, the dynamic nature of the world requires intelligent agents to trace the ongoing interaction with the external environment and adjust their behavior based on environmental feedback. This highlights the need to study long-term fairness in sequential decision-making problems.

Recent works have proposed to formulate and investigate long-term fairness within the framework of Markov Decision Processes (MDPs) [Wen *et al.*, 2021; Chi *et al.*, 2022; Tang *et al.*, 2023], which makes this topic closely related to reinforcement learning. Under this framework, a group fairness problem can be formulated as follows: different demographic groups are represented by a set of MDPs that share the same state and action space and run in parallel, and the agent’s decisions trigger group-level state transitions and reward assignments. In this case, long-term fairness regarding the well-being [Plecko and Bareinboim, 2022] of each group can be measured by the difference in future returns under the agent-specified policy. Some other fairness notions have been proposed as well [Liu *et al.*, 2018; Mouzannar *et al.*, 2019; Zhang *et al.*, 2020; Tang *et al.*, 2023], yet the sources of inequality that lead to unfairness remain unclear.

To highlight the importance of addressing this gap, consider an illustrative example involving two demographic groups, in which the state of each group is represented by a one-dimensional variable, as shown in Figure 1. The first two columns in this figure show the possible outcomes when two common fairness notions are adopted as criteria. The left one corresponds to fair decision, requiring the agent to make consistent decisions for both groups, regardless of sensitive attributes. This criterion is popular because people usually expect decision-makers to provide equal resources and opportunities for different groups instead of treating them separately, but it does not guarantee fair long-term well-being. The middle column, on the other hand, pursues fair state, urging the agent to improve the socio-economic status or qualification profile of the disadvantaged group, so that different groups will reach similar states in the long run. While this notion is very tempting, its implication for long-term well-being is also uncertain, as the disparities would remain if the reward assignment mechanism governing the environment is biased against the disadvantaged group. Both criteria may lead to

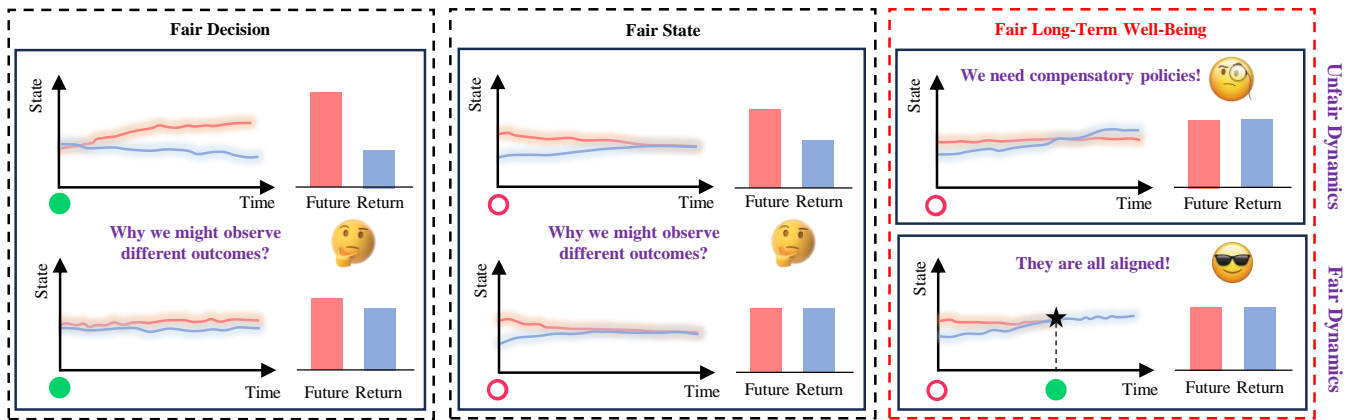


Figure 1: Dynamics fairness can help us gain a better understanding of the outcomes of different fairness criteria. This figure shows a sequential decision-making problem with two demographic groups: **red** represents the advantaged group and **blue** represents the disadvantaged. We present six subfigures, highlighting different fairness criteria (indicated by columns) alongside fair or unfair environmental dynamics (indicated by rows). Each subfigure consists of a line graph showing the evolution of group-wise states over time, accompanied by a histogram illustrating the future returns. The advantaged group may occupy a better initial state (e.g., socio-economic status, qualification profile, expertise level, market competitiveness, etc.), be easier to receive high rewards or reach better states compared to the disadvantaged group. Solid green circles represent consistent decisions for both groups, while empty red circles denote tailored decisions based on sensitive attributes.

different outcomes depending on the environmental dynamics, requiring us to investigate the mechanisms governing the environment, which is overlooked in previous studies.

In this paper, we address the above challenge from a causal perspective. Specifically, we trace the sources of inequality by decomposing the causal effect of sensitive attributes on long-term well-being. We formulate the long-term fairness problem via MDP and introduce a novel notion called *dynamics fairness* to capture the fairness of the mechanisms that govern the environment, measuring the expected changes in the next state and the reward induced by changing the identity of a demographic group (the value of a sensitive attribute) while holding everything else constant. This notion provides a useful tool for analysing the sources of inequality, distinguishing the inequality introduced by the environmental dynamics from those induced by decision-making or inherited from the past. Drawing inspiration from research on mediation analysis in the field of causal inference [Pearl, 2009; Pearl *et al.*, 2016; Pearl, 2022], we propose to characterize dynamics fairness by natural direct effects. Since the computation of these effects involves nested counterfactuals that cannot be observed in the factual world, we further derive identification formulas that allow us to reliably estimate these quantities from data. As depicted by the right column of Figure 1, breaching the fair decision and state criteria is generally inevitable when facing an environment that violates dynamics fairness. In this case, to achieve fair long-term well-being, an agent must adopt compensatory policies for the disadvantaged group. Conversely, when the environment satisfies dynamics fairness, these criteria can be aligned after eliminating the inequality inherited from the past, as indicated by the star.

In summary, our main contributions are as follows:

- We introduce a novel causal fairness notion called dynamics fairness. Previous notions were mainly defined in terms of observable disparities in states, decisions, or

rewards. To the best of our knowledge, dynamics fairness is the first notion that is defined on the underlying mechanisms governing the environment. Thus, it fills in the missing piece in studying long-term fairness in reinforcement learning.

- We further derive a method to quantitatively evaluate dynamics fairness through theoretical analysis, which allows this notion to be applied in practical situations. More importantly, this is a general method that does not rely on making parametric assumptions about the environment, such as linearity.
- We conducted extensive experiments to verify the theoretical results and demonstrate the effectiveness of the proposed methods in explaining, detecting, and reducing inequality in reinforcement learning problems.

## 1.1 Preliminaries

In this paper, we apply causal inference techniques to explore the mechanisms that explain the inequality in reinforcement learning. We adopt the framework of structural causal models (SCMs) [Pearl, 2009; Pearl *et al.*, 2016], which provide a mathematical formalization of causality. SCMs, denoted as  $\mathcal{C} := \langle \mathcal{V}, \mathcal{U}, \mathcal{F}, P(u) \rangle$ , consist of endogenous (observable) variables  $\mathcal{V}$  and background (latent) variables  $\mathcal{U}$ , linked by structural equations  $f_{V_i}$  in  $\mathcal{F}$ . These equations model the causal mechanisms among variables, relating each  $V_i \in \mathcal{V}$  to its parents  $\text{pa}(V_i)$  and associated exogenous variables  $U_{V_i}$ , under a probability distribution  $P(u)$  over  $\mathcal{U}$ . An SCM induces a causal diagram  $\mathcal{G}$  over the nodes  $\mathcal{V}$ , in which a directed edge  $V_i \rightarrow V_j$  exists if  $V_i$  is an argument of  $f_{V_j}$ .

SCMs provide a formal language for articulating counterfactual statements, allowing researchers to examine hypothetical scenarios by evaluating potential outcomes under varied situations. Specifically, researchers use  $Y_u(x)$  to denote the potential outcome of a variable  $Y$  obtained by solving

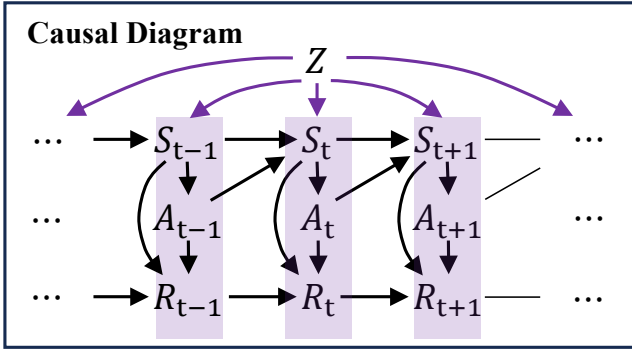


Figure 2: A causal diagram for sequential decision-making problems with a sensitive attribute  $Z$ . For clarity, the edges from the sensitive attribute to states, decisions, and rewards are represented by the connection between purple arrows and shaded areas.

the equations in  $\mathcal{F}$  with  $X = x$  for an individual  $U = u$ . In causal inference, many research efforts focus on predicting causal effects such as the total effect (TE)  $\text{TE}_{x,x^\theta}(Y) := E[Y(x^\theta)] - E[Y(x)]$ , where  $X$  is the treatment variable and  $Y$  is the outcome variable. TE measures the difference between the potential outcomes of a variable under different treatments for the whole population (rather than for a specific individual  $u$ ). In this paper, we are interested in the TE of the sensitive attribute on the future return, i.e., how a change in identity would affect a demographic group’s long-term well-being if all other factors were held constant?<sup>1</sup> The following section explores decomposing this effect into distinct components, examining the sources of inequality in RL problems, particularly focusing on the role of environmental dynamics.

## 2 Analysing inequality in reinforcement learning

In this section, we aim to explore and explain the inequality in reinforcement learning through a causal lens. We first present the problem formulation, employing causal modeling to analyse environmental dynamics in sequential decision-making scenarios. Next, we conduct an in-depth analysis of inequality and investigate its sources through the decomposition of inequality. Furthermore, we introduce dynamics fairness and propose an approach for evaluating this counterfactual concept quantitatively. Lastly, we present a model-based reinforcement learning algorithm that incorporates dynamics fairness into the planning process.

### 2.1 Problem Formulation

**Problem setup.** In reinforcement learning, sequential decision-making problems under uncertainty are often modeled as an MDP  $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, R, \mu, \gamma)$ .  $\mathcal{S}$  and  $\mathcal{A}$  denote the state and the action space respectively;  $P : \mathcal{S} \times \mathcal{A} \rightarrow \Delta_{\mathcal{S}}$  is the transition kernel, where  $P(s^\theta | s, a)$  is the probability of transitioning to state  $s^\theta$  after taking action  $a$  in state

<sup>1</sup>While in reality manipulating one’s identity is generally infeasible, one can still obtain reliable estimates of the causal effect of interest under appropriate structural assumptions [Pearl, 2009].

$s$ ;  $R : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$  specifies the immediate reward received after taking action  $a$  in state  $s$ ;  $\mu \in \Delta_{\mathcal{S}}$  stands for the initial state distribution. At step  $t$ , the goal of an RL agent is to maximize the future return  $G_t = R_t + \gamma R_{t+1} + \gamma^2 R_{t+2} + \dots = \sum_{k=0}^{\infty} \gamma^k R_{t+k}$ , where  $\gamma \in [0, 1)$  is the discount factor. In model-based RL, a common practice is to approximate the real environment using a parametric model  $f_\theta$ , which is employed to generate additional training data for policy learning or predict future transitions and rewards for planning. To model the environmental dynamics of multiple demographic groups, we consider  $N$  MDPs, all sharing the same state and action space, with each MDP representing a distinct demographic group. We abstract the group identity as a variable  $Z$ , referred to as the sensitive attribute. Then we can characterize the set of all  $N$  MDPs through an augmented dynamics model  $f_\theta$  that takes  $Z$  as an additional input. This is a common strategy used to address the dynamics generalization problem [Lee *et al.*, 2020; Guo *et al.*, 2022], enabling agents to learn the environment efficiently and accurately.

**SCM representation of the dynamics model.** Based on the above formulation, we can cast the augmented dynamics model into an SCM  $\mathcal{C}$  using auto-regressive uniformization (See Lemma 2 in [Buesing *et al.*, 2019]). For example, let’s consider the conditional distribution  $P(R|S, A, Z)$  and denote its cumulative distribution function as  $F(R|S, A, Z)$ . We can construct a new variable  $\bar{R} := F^{-1}(U_R|S, A, Z)$  where  $U_R$  is a uniform random variable sampled from  $[0, 1]$ . Since the distribution of  $\bar{R}$  matches that of  $R$ , by defining  $f_R := F^{-1}(U_{S^\theta}|S, A, Z)$ , we obtain the structural equation for  $R$ . A similar approach can be applied to other variables as well. The causal diagram  $\mathcal{G}$  induced by the resulting SCM  $\mathcal{C}$  is illustrated in Figure 2. Note that the impact of the sensitive attribute may cover various facets of the system, including state transitions, decision-making, and reward assignments. Therefore, inequality may stem from multiple sources following different causal paths in  $\mathcal{G}$ .

### 2.2 Decomposition of Inequality

In this work, we focus on the inequality in long-term well-being, a causal quantity that is captured by the difference in future return  $G_t$  when setting the sensitive attribute  $Z = z_1$  compared to  $Z = z_0$ . Using the causal language, we can now formally define this quantity as follows.

**Definition 1** (Well-being Gap). At step  $t$ , the well-being gap between the two demographic groups distinguished by the sensitive attribute  $Z$  is defined as:

$$\text{TE}_{z_0, z_1}(G_t) := E[G_t(z_1)] - E[G_t(z_0)].$$

This quantity measures the total effect of the sensitive attribute  $Z$  on the future return  $G_t$ . Since  $G_t$  is a linear combination of the reward signals, the well-being gap satisfies the following lemma:

**Lemma 1.** The well-being gap  $\text{TE}_{z_0, z_1}(G_t)$  can be decomposed into the sum of reward gaps at each time step:

$$\text{TE}_{z_0, z_1}(G_t) = \sum_{k=0}^{\infty} \gamma^k \text{TE}_{z_0, z_1}(R_{t+k}),$$

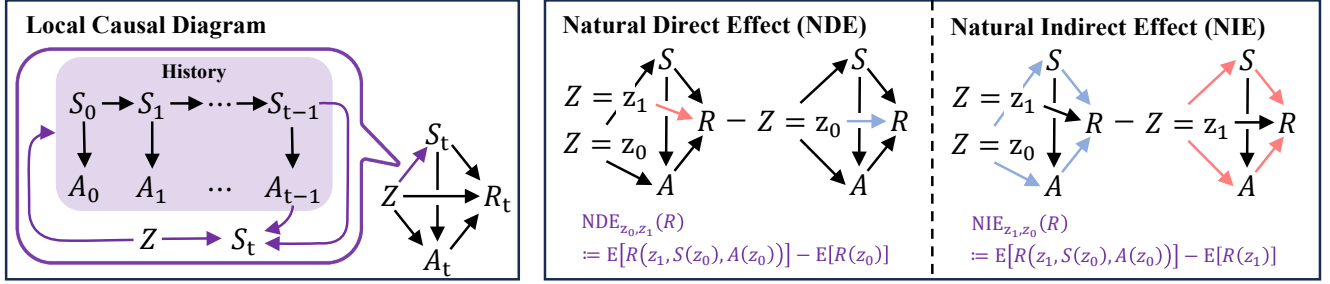


Figure 3: (Left) A local causal diagram in which the directed edge  $Z \rightarrow S_t$  carries the direct effect of  $Z$  on  $S_t$  and the indirect effect mediated by past states and actions (the history). We can combine them into one when focusing on analysing the effect of  $Z$  on  $R_t$  (or  $S_{t+1}$ ). (Right) A graphical representation of natural direct and indirect effects, where the contrast between the two quantities is highlighted in red and blue.

where  $\text{TE}_{z_0, z_1}(R_t) := E[R_t(z_1)] - E[R_t(z_0)]$  is the reward gap at step  $t$ .

Lemma 1 indicates that we can focus on the reward gap at each time step to examine the sources of inequality in long-term well-being. In order to analyse the reward gap  $\text{TE}_{z_0, z_1}(R_t)$  at step  $t$ , we present the local causal diagram in Figure 3 (Left). From the diagram, we can see that  $\text{TE}_{z_0, z_1}(R_t)$  is affected by the sensitive attribute  $Z$  through two types of causal paths: the direct path  $Z \rightarrow R_t$  and the indirect paths mediated by the current state  $S_t$  or the decision  $A_t$ . Notably,  $Z \rightarrow S_t$  carries the inequality inherited from historical factors as well as the one directly caused by  $Z$ .

Next, we introduce two counterfactual quantities that are useful for a refined decomposition of the reward gap. We omit the subscript  $t$  when the context is clear.

**Definition 2** (Natural Direct and Indirect Effect). The natural direct and indirect effects of sensitive attribute  $Z$  on reward  $R$  are defined, respectively, as:

$$\text{NDE}_{z_0, z_1}(R) := E[R(z_1, S(z_0), A(z_0))] - E[R(z_0)].$$

$$\text{NIE}_{z_1, z_0}(R) := E[R(z_1, S(z_0), A(z_0))] - E[R(z_1)].$$

Figure 3 (Right) illustrates these two quantities, both involving nested counterfactuals. For instance, in  $\text{NDE}_{z_0, z_1}(R)$ , the first factor involves computing the expected reward when  $S$  and  $A$  are set to the values they would have naturally taken under  $Z = z_0$  while  $Z$  is set to  $z_1$ . Perhaps surprisingly, despite being not realizable in the factual world, these quantities are of great interest and are widely used in many practical applications. In discrimination cases, by keeping the state  $S$  and decision  $A$  to the level they would have naturally attained under  $Z = z_0$ ,  $\text{NDE}_{z_0, z_1}(R)$  measures the advantage that the individuals in the disadvantaged group  $Z = z_0$  would have gained if there were no direct discrimination within the reward assignment process. Conversely,  $\text{NIE}_{z_1, z_0}(R)$  measures the advantage that the individuals in the disadvantaged group  $Z = z_0$  lose due to indirect discrimination.

Based on Lemma 1 and the above definitions, we can now present a quantitative decomposition of the well-being gap:

**Theorem 1** (Causal Decomposition of Well-being Gap). The well-being gap  $\text{TE}_{z_0, z_1}(G_t)$  can be decomposed as follows:

$$\text{TE}_{z_0, z_1}(G_t) = \sum_{k=0}^{\gamma} \gamma^k (\text{NDE}_{z_0, z_1}(R) - \text{NIE}_{z_1, z_0}(R)).$$

Theorem 1 offers a general decomposition of the well-being gap, quantitatively explaining how it is accumulated over time, without assuming any specific functional form of the structural equations in  $\mathcal{C}$ . The following lemma further reveals the qualitative relationship between the natural direct and indirect effects and the causal paths.

**Lemma 2.** Consider a set of MDPs described by the augmented dynamics model  $f_\theta$ , along with the induced causal diagram  $\mathcal{G}$ , the following statements hold: If  $\text{NDE}_{z_0, z_1}(R) \neq 0$ , then there exists a direct path from  $Z$  to  $R$ ; Similarly, if  $\text{NIE}_{z_1, z_0}(R) \neq 0$ , then there exist indirect paths connecting  $Z$  and  $R$ .

That is to say, the natural direct and indirect effects provide hints for structural information about the causal diagram, thus helping us identify the causal paths that are responsible for the gap. Summing up, the analysis presented herein systematically examines the sources of inequality through a causal lens. By introducing the natural direct and indirect effects, we provide both quantitative and qualitative explanations of the well-being gap, offering insights into the intricacies of inequality in reinforcement learning problems.

### 2.3 Dynamics Fairness

Beyond merely explaining the sources of inequality, in practice, we also seek clarity regarding the responsibility for such inequality. Theorem 1 and Lemma 2 demonstrate that the inequality accumulated at each step can be attributed to two types of causal paths: the direct path ( $Z \rightarrow R$ ) and the indirect paths ( $Z \rightarrow S \rightarrow R$ ,  $Z \rightarrow A \rightarrow R$ , and  $Z \rightarrow S \rightarrow A \rightarrow R$ ). However, the inequality introduced by the environmental dynamics and the decisions is still coupled together. In other words,  $\text{NIE}_{z_0, z_1}(R) \neq 0$  can not tell us how to determine responsibility between the two. Before diving into further details, we first formally define the notion of dynamics fairness.

**Definition 3** (Dynamics Fairness). The dynamics is fair regarding the sensitive attribute  $Z$  if there are no direct paths from  $Z$  to either the reward  $R$  or the next state  $S^0$  in the causal diagram  $\mathcal{G}$  induced by the augmented dynamics model  $f_\theta$ .

Definition 3 provides a structural condition for the environment to satisfy dynamics fairness. Banning the direct paths from the sensitive attribute to the reward and the next state

ensures inequality is not introduced by the environmental dynamics. As a result, there are only two possible sources of inequality: historical inequality, i.e., the disparity in the current state of different demographic groups, and the inequality introduced by the decisions, i.e., disparate treatment in the decision-making process. Specifically, the historical inequality can be estimated by the observed disparity in the current state  $\mathbb{E}[S|Z = z_1] - \mathbb{E}[S|Z = z_0]$ , which is easy to measure. If the environment satisfies dynamics fairness and there is no historical inequality, then the well-being gap is completely attributed to the decisions. In this way, we can solve the problem of attributing responsibility for inequality. Building upon Lemma 2, we establish the following theorem:

**Theorem 2** (Criterion for Dynamics Fairness Violation). The environment is considered to violate dynamics fairness if either  $\text{NDE}_{z_0, z_1}(R) \neq 0$  or  $\text{NDE}_{z_0, z_1}(S^\theta) \neq 0$  holds.

Theorem 2 offers a quantitative criterion for dynamics fairness. By checking the natural direct effects of  $Z$  on reward  $R$  and next state  $S^\theta$ , we can determine whether the mechanisms governing the environmental dynamics would introduce inequality during the sequential decision-making process, thus contributing to the well-being gap in the long run. More importantly, it also provides mathematical justification for compensatory policies. If the environment violates dynamics fairness, then disparate treatment is necessary to compensate the disadvantaged group when aiming to achieve equality in long-term well-being. On the other hand, if the environment is fair, then disparate treatment should be limited to a minimum level, not only for the sake of fair decision-making but also to avoid widening the well-being gap.

As mentioned earlier, the computation of  $\text{NDE}_{z_0, z_1}(\cdot)$  involves nested counterfactuals, which require data that we cannot obtain in the factual world. Fortunately, under the structural assumptions implied by the augmented dynamics model, we can derive identification equations for these effects.

**Theorem 3** (Identification of Dynamics Fairness). Under the structural assumptions implied by the augmented dynamics model depicted by the local causal diagram, the natural direct effects of  $Z$  on  $R$  and  $S^\theta$  can be identified as follows:

$$\begin{aligned} \text{NDE}_{z_0, z_1}(R) &= \sum_{s, a} (\mathbb{E}[R|Z = z_1, S = s, A = a] \\ &\quad - \mathbb{E}[R|Z = z_0, S = s, A = a])P(S = s, A = a|Z = z_0). \\ \text{NDE}_{z_0, z_1}(S^\theta) &= \sum_{s, a} (\mathbb{E}[S^\theta|Z = z_1, S = s, A = a] \\ &\quad - \mathbb{E}[S^\theta|Z = z_0, S = s, A = a])P(S = s, A = a|Z = z_0). \end{aligned}$$

In the above equations, all quantities on the right-hand side are expressed using conditional expectations or probabilities rather than counterfactuals, and they can be estimated directly from data using standard statistical methods. As a result, Theorem 3 provides a practical way to determine whether the environment satisfies dynamics fairness. Proofs are included in Appendix A in the supplementary material.

## 2.4 Algorithm for Achieving Long-Term Fairness

So far, we have examined the sources of inequality and the associated responsibilities, with a particular focus on the in-

fluence of environmental dynamics. The presented results not only demystify the intricate relationships behind unfairness in RL problems but also provide valuable guidance for designing RL algorithms to tackle inequality. Based on the above analyses, we propose a model-based RL algorithm named *InsightFair* that incorporates dynamics fairness into the planning process. Specifically, the algorithm iteratively collects data from the environment and trains an ensemble augmented dynamics model  $f_\theta$ , which is later used to simulate the environment and evaluate the action sequences for planning. In each epoch, the agent first evaluates dynamics fairness following Theorem 3 using  $f_\theta$  and the collected data. If the environment violates dynamics fairness, then disparate treatment is considered necessary in the planning process. Otherwise, the agent samples consistent action sequences for all groups when the observed disparity in the current state approaches zero. The optimization objective is regularized by a constraint on the estimated well-being gap to induce a fair future. The detailed algorithm is presented in Appendix B.

## 3 Experiments

In this section, we empirically validate the theoretical results of our work through extensive experiments. Specifically, the results presented here aim to answer the following questions: (1) Does the proposed decomposition offer an accurate quantitative explanation for inequality? (2) Can the proposed identification formula reliably detect cases where the environment violates dynamics fairness? (3) Does the proposed algorithm, *InsightFair*, work effectively with the aid of dynamics fairness? To be self-contained, we include additional details and results in Appendix C in the supplementary material.

### 3.1 Explaining Inequality

To answer the first question, we consider a simple non-linear model for the reward function, i.e.,  $R = 1$  if  $w_0 + w_1 z + w_2 s + w_3 a + u_R \geq 0$  and  $R = 0$  otherwise, where  $w_0, w_1, w_2, w_3$  are the model parameters and  $U_R \sim \text{logistic}(0, 1)$  is the background variable. Since the reward function is fixed within an episode and the return is a linear combination of the rewards, we can examine the well-being gap by analysing the decomposition of the reward gap  $\text{TE}_{z_0, z_1}(R)$ .

Figure 4 illustrates the comparison between the total effect  $\text{TE}_{z_0, z_1}(R)$  and the natural direct and indirect effects  $\text{NDE}_{z_0, z_1}(R)$  and  $\text{NIE}_{z_1, z_0}(R)$ , respectively. The values are computed analytically under two sets of model parameters  $w_1 = w_2 = w_3 = 0.1$  and  $w_1 = w_2 = w_3 = 3.0$  along with  $w_0$  varying in the range  $[-2.5, 2.5]$ . First, we can see that the sum of  $\text{NDE}_{z_0, z_1}(R)$  and  $-\text{NIE}_{z_1, z_0}(R)$  accurately explains the reward gap  $\text{TE}_{z_0, z_1}(R)$  in both settings. Second, by varying  $w_0$ , the proportion of the direct and indirect effects changes accordingly. In the right figure, the dominance of the direct effect suggests that the inequality is mainly introduced by the environmental dynamics instead of other factors.

### 3.2 Detecting Violation of Dynamics Fairness

In this part, we empirically verify the effectiveness of the proposed identification formula by customizing the ML-fairness-gym suite [D’Amour *et al.*, 2020] for group fairness problems. Specifically, we use the Allocation environment that

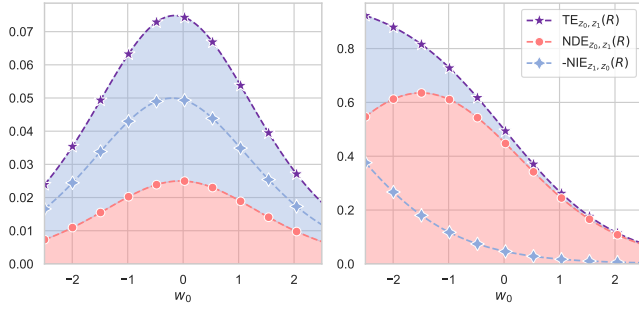


Figure 4: Visualization of total effect, natural direct effect, and natural indirect effect under different model parameters. The purple curve represents  $TE_{z_0, z_1}(R)$ , while the red and blue curves depict  $NDE_{z_0, z_1}(R)$  and  $NIE_{z_1, z_0}(R)$ , respectively. The two figures illustrate scenarios where either the indirect effect dominates (top) or the direct effect dominates (bottom), while the shaded areas show that the sum of  $NDE_{z_0, z_1}(R)$  and  $NIE_{z_1, z_0}(R)$  matches the total effect in both scenarios, as guaranteed by Theorem 1.

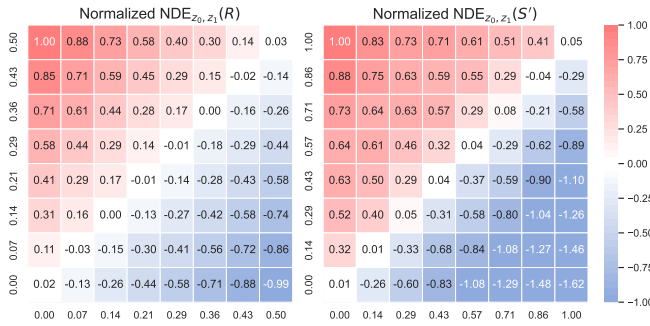


Figure 5: Results of evaluating dynamics fairness using the proposed identification formula. Each tile represents the estimated natural direct effects,  $NDE_{z_0, z_1}(R)$  or  $NDE_{z_0, z_1}(S^\theta)$ , under specific parameter configurations indicated by the row and column. Lighter colors signify values closer to zero (satisfying dynamics fairness), while darker colors represent larger absolute values (violating dynamics fairness). Red denotes the second demographic group is advantaged, while blue denotes disadvantaged.

contains two demographic groups. At each step, some incidents like rat infestations occur in both groups, depending on their current state. The agent must allocate resources effectively to minimize both missed incidents and resources allocated. Groups do not interact with each other. By manipulating the reward function and the transition dynamics, we can generate different environments to simulate scenarios where the environment satisfies or violates dynamics fairness.

We conduct two sets of experiments to evaluate  $NDE_{z_0, z_1}(R)$  and  $NDE_{z_0, z_1}(S^\theta)$ , respectively. In the first set, we vary the reward function while keeping the transition dynamics fixed. Two parameters  $\beta_0$  and  $\beta_1$  are used to control the relative advantage of the second demographic group. When  $\beta_1 > \beta_0$ , the second group is advantaged, and vice versa. The environment is considered fair when  $\beta_1 = \beta_0$ , i.e., the relative advantage reduces to zero. The second set of experiments is conducted similarly, except that we fix the reward function and vary the parameters controlling the tran-

sition dynamics to determine the relative advantage. For both sets, we conduct a total of 64 experiments by varying the parameters in a fixed range to generate a series of environments and estimating the natural direct effects using the proposed identification formula. The results are shown in Figure 5. From this figure, it is clear that the natural direct effects approach zero when the environment satisfies dynamics fairness (as colored in white on the diagonal). On the other hand, when the environment violates dynamics fairness, the natural direct effects vary with the magnitude of the relative advantage, as indicated by the darker colors in the heatmaps. These results empirically support that the proposed approach can reliably detect violations of dynamics fairness.

### 3.3 Promoting Long-Term Fairness

To answer the third question, we evaluate the performance of the proposed algorithm by comparing it with the base algorithm, PETS [Chua *et al.*, 2018], which is a popular model-based RL algorithm but does not consider fairness. In addition, we consider two variants of PETS. The first variant, Fair-A, requires the agent to satisfy the fair decision criterion, sampling consistent action sequences for all groups. The second variant, Fair-S, incorporates a constraint regarding state disparities into the optimization objective, aiming to achieve equality in the state for all groups. Existing studies [Wen and Topcu, 2018; Liu *et al.*, 2021] have shown that model-based algorithms can effectively satisfy constraints while maintaining good task performance. Note that our focus in this paper is orthogonal to these works. The experiments presented herein aim to examine the effectiveness of the proposed fairness notion when incorporated into a model-based RL algorithm, particularly in promoting fair long-term well-being when the environment violates dynamics fairness and in aligning different criteria in a fair environment.

The learning curves in Figure 6 (Left) illustrate the average return and the well-being gap for all algorithms in the Allocation environment. The results are averaged over five runs with different random seeds. InsightFair exhibits a slightly lower task performance compared to PETS and Fair-A, while considerably reducing the well-being gap, indicating its effectiveness in promoting long-term fairness when the environment violates dynamics fairness. Fair-S also reduces the well-being gap but at the cost of task performance, as it sacrifices the utility of the advantaged group to achieve equality in the state. In the next experiment shown in Figure 6 (Right), we consider a fair environment where the relative advantage is zero but the advantaged group begins with a better initial state. We visualize the changes in state and decision gap (the absolute difference in actions) between the two groups. InsightFair makes consistent decisions for both groups when state disparities approach zero, resulting in zero decision gap in most steps, while PETS demonstrates disparate treatment throughout the episode. These results suggest that InsightFair effectively aligns the fair state, action, and well-being criteria in a fair environment.

## 4 Related Work

**Long-term fairness.** Consideration of fairness in machine learning systems has evolved beyond static settings to a dy-

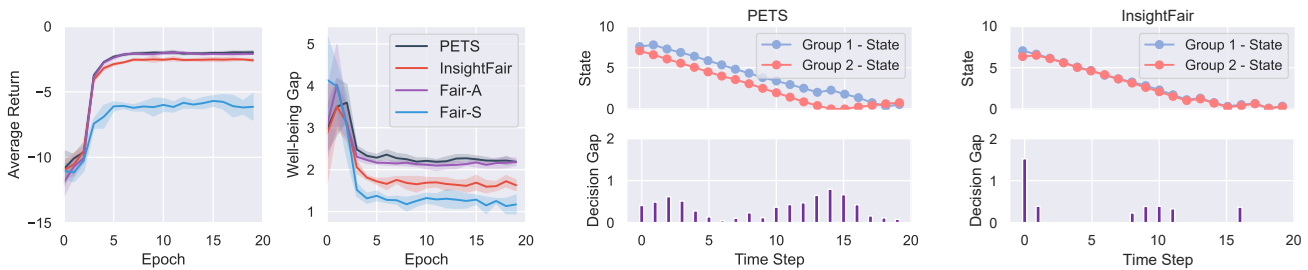


Figure 6: (Left) Learning curves depicting return and the well-being gap. (Right) Changes in state and decision gaps within an episode.

dynamic perspective [D’Amour *et al.*, 2020], i.e., autonomous agents may interact with the environment over time and have a long-term impact on fairness. Notably, there is growing evidence [Liu *et al.*, 2018; Zhang *et al.*, 2020] supporting concerns about the application of static fairness measures, as they may jeopardize long-term fairness. In recent years, there has been a considerable amount of work devoted to developing algorithms that promote long-term fairness [Mouzannar *et al.*, 2019; Wen *et al.*, 2021; Puranik *et al.*, 2022; Chi *et al.*, 2022; Tan *et al.*, 2022; Yu *et al.*, 2022; Xu *et al.*, 2023]. [Wen *et al.*, 2021] investigated the long-term fairness problem under scenarios in which individuals receive decisions repeatedly (e.g., a lending problem in which an individual applies for a loan repeatedly and the agent needs to decide whether or not to approve the applications). Specifically, they define fairness as the gap between the expected cumulative rewards of individuals in different groups. The proposed algorithm relies on a known model and can be reduced to solving a linear programming problem in the finite-state MDPs setting. [Chi *et al.*, 2022] adopt a similar definition of fairness. Their algorithm employs integral probability metrics to align the state visitation distributions among groups, thereby mitigating return disparity. More recently, [Yu *et al.*, 2022] proposed a policy optimization algorithm that incorporates fairness constraints into the advantage function. This enables an agent to efficiently explore the environment and learn a policy that strikes the balance between utility and fairness. However, [Xu *et al.*, 2023] noted that assuming the fairness metric as part of the state is generally impractical because long-term fairness necessitates the agent to focus on the future rather than the past. Their work studied supply-demand MDPs, a variant of MDPs that contain two time-varying variables, supply and demand, for each group depending on states and actions. The authors suggested using the ratio between the cumulative group supply and demand to measure long-term well-being and proposed a policy optimization algorithm that can effectively reduce the discrepancy in this ratio. In this paper, we systematically examine long-term fairness for demographic groups characterized by a set of different MDPs. Our work features causal modeling of the agent-environment interaction and aims to identify the sources of inequality, thereby providing a new perspective on studying the long-term fairness problem.

**Causal fairness analysis.** Causal fairness analysis, an emerging trend in fairness research, leverages techniques and theories from the causal inference community to ad-

vance the exploration of fairness problems in machine learning [Creager *et al.*, 2020; Plecko and Bareinboim, 2022; Deng *et al.*, 2023]. This methodology not only rigorously defines fairness but also provides effective and reliable solutions to ensure it.” For example, [Kusner *et al.*, 2017] proposed counterfactual fairness for fair prediction problems, requiring the distribution over possible predictions of an individual to remain unchanged in a world where its sensitive attributes had been changed. In fact, counterfactuals provide a versatile framework for formulating various fairness notions of different granularity. [Zhang and Bareinboim, 2018] introduced different types of counterfactual effects for the fair decision problem. These effects offer fairness notions at the group level, helping researchers interpret various discrimination mechanisms. Furthermore, [Nabi *et al.*, 2019] and [Wu *et al.*, 2019] demonstrated how causal graphs can assist researchers in defining fine-grained fairness associated with specific causal paths. For instance, using sensitive attributes or their proxies is deemed reasonable in some cases due to business necessity [Plecko and Bareinboim, 2022] but may lead to discrimination in other cases. Recent research has extended the study of causal fairness to contextual bandits and outcome control problems [Huang *et al.*, 2022; Plecko and Bareinboim, 2023]. Unlike previous studies, our paper delves into the fairness problem in the sequential decision-making setting, introducing a novel causal fairness notion, dynamics fairness, which captures the fairness of the mechanisms governing the environment.

## 5 Conclusion

In this paper, we systematically studied inequality in RL problems through a causal lens, aiming to demystify the intricate relationships behind unfairness. We introduced a causal model to formalize the agent-environment interaction and provided a quantitative decomposition of the well-being gap, which explains how inequality is accumulated over time. Furthermore, we proposed dynamics fairness, a novel causal fairness notion that, for the first time, captures the fairness of the underlying mechanisms governing the environment. To measure this counterfactual quantity, we proposed identification formulas that do not rely on specific functional forms or a known model. Extensive experiments were conducted to verify the theoretical results and demonstrate the effectiveness of the proposed approaches. In the future, we plan to extend our work to more complicated scenarios, such as partially observable and non-stationary environments.

## References

- [Anderson *et al.*, 2019] Henry Anderson, Afshan Boodhwani, and Ryan S. Baker. Assessing the Fairness of Graduation Predictions. In *Proceeding of the 12th International Conference on Educational Data Mining*, pages 488–491, 2019.
- [Angwin *et al.*, ] Julea Angwin, Larson Jeff, Mattu Surya, and Kirchner Lauren. Machine Bias: There’s software used across the country to predict future criminals. And it’s biased against blacks. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>.
- [Buesing *et al.*, 2019] Lars Buesing, Theophane Weber, Yori Zwols, Nicolas Heess, Sebastien Racaniere, Arthur Guez, and Jean-Baptiste Lespiau. Woulda, Coulda, Shoulda: Counterfactually-Guided Policy Search. In *International Conference on Learning Representations*, February 2019.
- [Chi *et al.*, 2022] Jianfeng Chi, Jian Shen, Xinyi Dai, Weinan Zhang, Yuan Tian, and Han Zhao. Towards Return Parity in Markov Decision Processes. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, pages 1161–1178. PMLR, May 2022.
- [Chua *et al.*, 2018] Kurtland Chua, Roberto Calandra, Rowan McAllister, and Sergey Levine. Deep Reinforcement Learning in a Handful of Trials using Probabilistic Dynamics Models. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- [Creager *et al.*, 2020] Elliot Creager, David Madras, Toniann Pitassi, and Richard Zemel. Causal modeling for fairness in dynamical systems. In *International Conference on Machine Learning*, pages 2185–2195. PMLR, 2020.
- [D’Amour *et al.*, 2020] Alexander D’Amour, Hansa Sriniwasan, James Atwood, Pallavi Baljekar, David Sculley, and Yoni 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.
- [Deng *et al.*, 2023] Zhihong Deng, Jing Jiang, Guodong Long, and Chengqi Zhang. Causal Reinforcement Learning: A Survey. *Transactions on Machine Learning Research*, July 2023.
- [Dwork *et al.*, 2012] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, ITCS ’12, pages 214–226, New York, NY, USA, January 2012. Association for Computing Machinery.
- [Guo *et al.*, 2022] Jiaxian Guo, Mingming Gong, and Dacheng Tao. A Relational Intervention Approach for Unsupervised Dynamics Generalization in Model-Based Reinforcement Learning. In *International Conference on Learning Representations*, March 2022.
- [Hardt *et al.*, 2016] Moritz Hardt, Eric Price, Eric Price, and Nati Srebro. Equality of Opportunity in Supervised Learning. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.
- [Huang *et al.*, 2022] Wen Huang, Kevin Labille, Xintao Wu, Dongwon Lee, and Neil Heffernan. Achieving User-Side Fairness in Contextual Bandits. *Human-Centric Intelligent Systems*, 2(3):81–94, December 2022.
- [Kasy and Abebe, 2021] Maximilian Kasy and Rediet Abebe. Fairness, Equality, and Power in Algorithmic Decision-Making. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT ’21, pages 576–586, New York, NY, USA, March 2021. Association for Computing Machinery.
- [Kusner *et al.*, 2017] Matt J. Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. Counterfactual fairness. *Advances in neural information processing systems*, 30, 2017.
- [Lee *et al.*, 2014] Eric L. Lee, Jing-Kai Lou, Wei-Ming Chen, Yen-Chi Chen, Shou-De Lin, Yen-Sheng Chiang, and Kuan-Ta Chen. Fairness-Aware Loan Recommendation for Microfinance Services. In *Proceedings of the 2014 International Conference on Social Computing, Social-Com ’14*, pages 1–4, New York, NY, USA, August 2014. Association for Computing Machinery.
- [Lee *et al.*, 2020] Kimin Lee, Younggyo Seo, Seunghyun Lee, Honglak Lee, and Jinwoo Shin. Context-aware Dynamics Model for Generalization in Model-Based Reinforcement Learning. In *Proceedings of the 37th International Conference on Machine Learning*, pages 5757–5766. PMLR, November 2020.
- [Liu *et al.*, 2018] Lydia T. Liu, Sarah Dean, Esther Rolf, Max Simchowitz, and Moritz Hardt. Delayed impact of fair machine learning. In *International Conference on Machine Learning*, pages 3150–3158. PMLR, 2018.
- [Liu *et al.*, 2021] Zuxin Liu, Hongyi Zhou, Baiming Chen, Sicheng Zhong, Martial Hebert, and Ding Zhao. Constrained Model-based Reinforcement Learning with Robust Cross-Entropy Method, March 2021.
- [MacKinnon *et al.*, 2007] D.P. MacKinnon, C.M. Lockwood, C.H. Brown, W. Wang, and J.M. Hoffman. The intermediate endpoint effect in logistic and probit regression. *Clinical Trials*, 4(5):499–513, October 2007.
- [Mehrabi *et al.*, 2021] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A Survey on Bias and Fairness in Machine Learning. *ACM Comput. Surv.*, 54(6):115:1–115:35, July 2021.
- [Mouzannar *et al.*, 2019] Hussein Mouzannar, Mesrob I. Ohannessian, and Nathan Srebro. From Fair Decision Making To Social Equality. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, FAT\* ’19, pages 359–368, New York, NY, USA, January 2019. Association for Computing Machinery.
- [Nabi *et al.*, 2019] Razieh Nabi, Daniel Malinsky, and Ilya Shpitser. Learning Optimal Fair Policies. In *Proceedings of the 36th International Conference on Machine Learning*, pages 4674–4682. PMLR, May 2019.



- [Patil *et al.*, 2021] Vishakha Patil, Ganesh Ghalme, Vineet Nair, and Y. Narahari. Achieving fairness in the stochastic multi-armed bandit problem. *The Journal of Machine Learning Research*, 22(1):174:7885–174:7915, January 2021.
- [Pearl *et al.*, 2016] Judea Pearl, Madelyn Glymour, and Nicholas P. Jewell. *Causal Inference in Statistics: A Primer*. John Wiley & Sons, 2016.
- [Pearl, 2009] Judea Pearl. *Causality*. Cambridge university press, 2009.
- [Pearl, 2022] Judea Pearl. Direct and Indirect Effects. In Hector Geffner, Rina Dechter, and Joseph Y. Halpern, editors, *Probabilistic and Causal Inference*, pages 373–392. ACM, New York, NY, USA, 1 edition, February 2022.
- [Pessach and Shmueli, 2023] Dana Pessach and Erez Shmueli. Algorithmic Fairness. In Lior Rokach, Oded Maimon, and Erez Shmueli, editors, *Machine Learning for Data Science Handbook: Data Mining and Knowledge Discovery Handbook*, pages 867–886. Springer International Publishing, Cham, 2023.
- [Plecko and Bareinboim, 2022] Drago Plecko and Elias Bareinboim. Causal Fairness Analysis, July 2022.
- [Plecko and Bareinboim, 2023] Drago Plecko and Elias Bareinboim. Causal Fairness for Outcome Control, June 2023.
- [Puranik *et al.*, 2022] Bhagyashree Puranik, Upamanyu Madhow, and Ramtin Pedarsani. A Dynamic Decision-Making Framework Promoting Long-Term Fairness. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '22, pages 547–556, New York, NY, USA, July 2022. Association for Computing Machinery.
- [Tan *et al.*, 2022] Xiaoqing Tan, Zhengling Qi, Christopher Warren Seymour, and Lu Tang. RISE: Robust Individualized Decision Learning with Sensitive Variables. In *Advances in Neural Information Processing Systems*, October 2022.
- [Tang *et al.*, 2023] Zeyu Tang, Yatong Chen, Yang Liu, and Kun Zhang. Tier Balancing: Towards Dynamic Fairness over Underlying Causal Factors. In *The Eleventh International Conference on Learning Representations*, February 2023.
- [Tolan *et al.*, 2019] Songül Tolan, Marius Miron, Emilia Gómez, and Carlos Castillo. Why Machine Learning May Lead to Unfairness: Evidence from Risk Assessment for Juvenile Justice in Catalonia. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law*, ICAIL '19, pages 83–92, New York, NY, USA, June 2019. Association for Computing Machinery.
- [Wen and Topcu, 2018] Min Wen and Ufuk Topcu. Constrained Cross-Entropy Method for Safe Reinforcement Learning. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- [Wen *et al.*, 2021] Min Wen, Osbert Bastani, and Ufuk Topcu. Algorithms for fairness in sequential decision making. In *International Conference on Artificial Intelligence and Statistics*, pages 1144–1152. PMLR, 2021.
- [Wolff *et al.*, 2013] Annika Wolff, Zdenek Zdrahal, Andriy Nikolov, and Michal Pantucek. Improving retention: Predicting at-risk students by analysing clicking behaviour in a virtual learning environment. In *Proceedings of the Third International Conference on Learning Analytics and Knowledge*, LAK '13, pages 145–149, New York, NY, USA, April 2013. Association for Computing Machinery.
- [Wu *et al.*, 2019] Yongkai Wu, Lu Zhang, Xintao Wu, and Hanghang Tong. PC-Fairness: A Unified Framework for Measuring Causality-based Fairness. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [Xu *et al.*, 2023] Yuancheng Xu, Chenghao Deng, Yanchao Sun, Ruijie Zheng, Xiyao Wang, Jieyu Zhao, and Furong Huang. Equal Long-term Benefit Rate: Adapting Static Fairness Notions to Sequential Decision Making, September 2023.
- [Yu *et al.*, 2022] Eric Yang Yu, Zhizhen Qin, Min Kyung Lee, and Sicun Gao. Policy Optimization with Advantage Regularization for Long-Term Fairness in Decision Systems. In *Advances in Neural Information Processing Systems*, October 2022.
- [Zhang and Bareinboim, 2018] Junzhe Zhang and Elias Bareinboim. Fairness in Decision-Making — The Causal Explanation Formula. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, April 2018.
- [Zhang *et al.*, 2017] Lu Zhang, Yongkai Wu, and Xintao Wu. A causal framework for discovering and removing direct and indirect discrimination. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 3929–3935, 2017.
- [Zhang *et al.*, 2020] Xueru Zhang, Ruibo Tu, Yang Liu, Mingyan Liu, Hedvig Kjellstrom, Kun Zhang, and Cheng Zhang. How do fair decisions fare in long-term qualification? In *Advances in Neural Information Processing Systems*, volume 33, pages 18457–18469. Curran Associates, Inc., 2020.

## A Proofs

**Proof of Lemma 1.** The well-being gap  $\text{TE}_{z_0, z_1}(G_t)$  is defined as the difference between two expectations:

$$\text{TE}_{z_0, z_1}(G_t) := \mathbb{E}[G_t(z_1)] - \mathbb{E}[G_t(z_0)],$$

in which  $\mathbb{E}[G_t(z)]$  denote the expected value of potential outcomes of return  $G_t$  when the sensitive attribute  $Z$  is set to  $z$ . We can expand  $G_t$  as the sum of discounted rewards. Since the expected value operator  $\mathbb{E}[\cdot]$  is linear, we have:

$$\begin{aligned} \mathbb{E}[G_t(z)] &= \mathbb{E}[R_t(z) + \gamma R_{t+1}(z) + \gamma^2 R_{t+2}(z) + \dots] \\ &= \mathbb{E}[R_t(z)] + \gamma \mathbb{E}[R_{t+1}(z)] + \gamma^2 \mathbb{E}[R_{t+2}(z)] + \dots \\ &= \sum_{k=0}^{\infty} \gamma^k \mathbb{E}[R_{t+k}(z)]. \end{aligned}$$

As a result, the well-being gap can be decomposed as the discounted sum of reward gaps at each time step:

$$\begin{aligned} \text{TE}_{z_0, z_1}(G_t) &= \sum_{k=0}^{\infty} \gamma^k \mathbb{E}[R_{t+k}(z_1)] - \sum_{k=0}^{\infty} \gamma^k \mathbb{E}[R_{t+k}(z_0)] \\ &= (\mathbb{E}[R_t(z_1)] - \mathbb{E}[R_t(z_0)]) \\ &\quad + \gamma (\mathbb{E}[R_{t+1}(z_1)] - \mathbb{E}[R_{t+1}(z_0)]) + \dots \\ &= \sum_{k=0}^{\infty} \gamma^k \text{TE}_{z_0, z_1}(R_{t+k}). \end{aligned}$$

□

**Proof of Theorem 1.** Based on Lemma 1, to prove this theorem, it suffices to show that the reward gap  $\text{TE}_{z_0, z_1}(R_t)$  can be decomposed as the sum of  $\text{NDE}_{z_0, z_1}(R)$  and  $-\text{NIE}_{z_1, z_0}(R)$ , which can be derived straightforwardly from the definitions of the natural direct and indirect effects:

$$\begin{aligned} \text{TE}_{z_0, z_1}(R_t) &= \mathbb{E}[R_t(z_1)] - \mathbb{E}[R_t(z_0)] \\ &= (\mathbb{E}[R_t(z_1, S(z_0), A(z_0))] - \mathbb{E}[R_t(z_0)]) \\ &\quad + (\mathbb{E}[R_t(z_1)] - \mathbb{E}[R_t(z_1, S(z_0), A(z_0))]) \\ &= \text{NDE}_{z_0, z_1}(R_t) - \text{NIE}_{z_1, z_0}(R_t). \end{aligned}$$

Since the reward function is time-invariant, we have:

$$\text{TE}_{z_0, z_1}(G_t) = \sum_{k=0}^{\infty} \gamma^k (\text{NDE}_{z_0, z_1}(R) - \text{NIE}_{z_1, z_0}(R)).$$

□

**Proof of Lemma 2.** Consider the contrapositive of the original statement: If there is no direct causal path from  $Z$  to  $R$  and  $S^\theta$  in a causal diagram  $\mathcal{G}$  induced by the augmented dynamics model  $f_\theta$ , then  $\text{NDE}_{z_0, z_1}(R) = 0$  and  $\text{NDE}_{z_0, z_1}(S^\theta) = 0$ ; Similarly, if there are no indirect causal paths from  $Z$  to  $R$ , then  $\text{NIE}_{z_1, z_0}(R) = 0$ . We can divide the proof of this new statement into two parts.

First, following the consistency rule of counterfactuals [Pearl, 2009, Corollary 7.3.2], we have  $R(z_0) = R(z_0, S(z_0), A(z_0))$ . Since there is no direct path from  $Z$  to  $R$ , we have  $R(z, s, a) = R(s, a)$  for all  $z, s, a$  according to

the exclusion restriction rule of structure-based counterfactuals [Pearl, 2009, Section 7.3.2]. As a result, we can now rewrite the natural indirect effect of  $Z$  on  $R$  as follows:

$$\begin{aligned} \text{NDE}_{z_0, z_1}(R) &= \mathbb{E}[R(z_1, S(z_0), A(z_0))] - \mathbb{E}[R(z_0)] \\ &= \mathbb{E}[R(z_1, S(z_0), A(z_0))] - \mathbb{E}[R(z_0, S(z_0), A(z_0))] \\ &= \mathbb{E}[R(S(z_0), A(z_0)) - R(S(z_0), A(z_0))] = 0. \end{aligned}$$

Similarly, when there are no indirect paths from  $Z$  to  $R$ , or more specifically, when there is no path from  $Z$  to  $R$  that passes through  $S$  or  $A$  in  $\mathcal{G}$ , we have  $R(z, s, a) = R(z)$  for all  $z, s, a$ . As a result, the natural direct effect of  $Z$  on  $R$  can be rewritten as follows:

$$\begin{aligned} \text{NIE}_{z_1, z_0}(R) &= \mathbb{E}[R(z_1)] - \mathbb{E}[R(z_1, S(z_0), A(z_0))] \\ &= \mathbb{E}[R(z_1)] - \mathbb{E}[R(z_1)] = 0. \end{aligned}$$

Summing up, we show that the contrapositive statement is true, and by equivalence, the original statement is proved. □

**Proof of Theorem 2.** Definition 3 states that, to satisfy dynamics fairness, the sensitive attribute  $Z$  should not causally influence the reward  $R$  or the next state  $S^\theta$  through direct causal paths in the causal diagram  $\mathcal{G}$  induced by the augmented dynamics model  $f_\theta$ . Lemma 2 shows that  $\text{NDE}_{z_0, z_1}(R)$  can be used to detect the direct causal paths from  $Z$  to  $R$ . We can prove that  $Z \rightarrow S^\theta$  exists if  $\text{NDE}_{z_0, z_1}(S^\theta) \neq 0$  in a similar way. As a result, if either  $\text{NDE}_{z_0, z_1}(R) \neq 0$  or  $\text{NDE}_{z_0, z_1}(S^\theta) \neq 0$  holds, then the environment violates dynamics fairness. In other words, we can detect the violation of dynamics fairness by checking the natural direct effects of  $Z$  on  $R$  and  $S^\theta$ . □

**Proof of Theorem 3.** Here we focus on the identification of  $\text{NDE}_{z_0, z_1}(R)$ , as the identification of  $\text{NDE}_{z_0, z_1}(S^\theta)$  can be proved similarly. We first rewrite the natural direct effect of  $Z$  on  $R$  as follows:

$$\begin{aligned} \text{NDE}_{z_0, z_1}(R) &= \mathbb{E}[R(z_1, S(z_0), A(z_0))] - \mathbb{E}[R(z_0)] \\ &= \sum_{s, a} (\mathbb{E}[R(z_1, s, a) \mid S(z_0) = s, A(z_0) = a] \\ &\quad - \mathbb{E}[R(z_0, s, a) \mid S(z_0) = s, A(z_0) = a]) P(S(z_0) = s, A(z_0) = a) \\ &= \sum_{s, a} (\mathbb{E}[R(z_1, s, a)] - \mathbb{E}[R(z_0, s, a)]) P(S(z_0) = s, A(z_0) = a). \end{aligned}$$

The last step holds because there are no paths between  $R$  and  $Z, S, A$  in  $\mathcal{G}$  containing only background variables. In this case, we have  $R(z, w) \perp\!\!\!\perp W(z^\theta)$  for any  $z, z^\theta, w$  according to the independence restriction rule of structure-based counterfactuals [Pearl, 2009, Section 7.3.2]. The resulting equality is also known as the experimental identification of natural direct effects [Pearl, 2022]. Since none of the quantities on the right-hand side involve nested counterfactuals, it can be estimated using experimental data.

Furthermore, based on the rule 2 of do-calculus [Pearl, 2009, Section 3.4.2],  $\mathbb{E}[R(z, s, a)] = \mathbb{E}[R|z, s, a]$  holds for all  $z, s, a$  and  $P(S(z_0) = s, A(z_0) = a)$  holds for all  $s, a$ . As a result, we can derive the identification equation for

$\text{NDE}_{z_0, z_1}(R)$  as follows:

$$\text{NDE}_{z_0, z_1}(R) = \sum_{s, a} (\mathbb{E}[R|Z = z_1, S = s, A = a] - \mathbb{E}[R|Z = z_0, S = s, A = a]) P(S = s, A = a | Z = z_0),$$

which provides a principled way to estimate  $\text{NDE}_{z_0, z_1}(R)$  using standard statistical methods, hence a tool for detecting violations of dynamics fairness.  $\square$

## B Algorithm for Achieving Long-Term Fairness

Algorithm 1 presents the detailed procedure of the proposed algorithm, InsightFair, which explicitly incorporates dynamics fairness into the planning process, aiming to promote long-term fairness. Specifically, we adopt ensembles of probabilistic dynamics models to capture uncertainty in the environment, along with a planning procedure based on trajectory sampling, similar to PETS [Chua *et al.*, 2018]. However, in each epoch, the agent is required to evaluate dynamics fairness DF using the proposed identification formula in Theorem 3 after training the augmented dynamics model  $f_\theta$  on the collected data. Then it passes the value of DF to the planner.

To select the best action  $a_t$  for each group in the current state  $s_t$ , the planner first checks whether the observed disparity  $d(s_t)$  in  $s_t$  is smaller than a predefined threshold  $\epsilon$ , which can be determined by domain experts for different applications. If  $d(s_t)$  is sufficiently small and  $\text{DF} = 1$ , then the agent samples consistent actions for all groups. Otherwise, the agent samples group-specific action sequences for each group. Then the agent evaluates the samples using the augmented dynamics model  $f_\theta$  and the estimated well-being gap  $\text{TE}_{z_0, z_1}(G_{t:t+H})$  for the next  $H$  steps. Instead of selecting candidate actions with the highest estimated return, to promote long-term fair well-being, the agent sorts the samples according to return regularized by the estimated well-being gap. The agent then updates the action distributions  $\text{CEM}(\cdot)$  and return the first action  $a_t$  in the optimal sample.

## C Additional Experimental Details and Results

In this section, we provide additional details and results of the experiments presented in Section 3.

### C.1 Experimental Details of Section 3.1

In this experiment, to quantitatively examine the decomposition formula in Theorem 1, we consider a simple non-linear model similar to the one studied in [MacKinnon *et al.*, 2007; Zhang and Bareinboim, 2018] but with no confounders between the treatment variable (the sensitive variable  $Z$  in our case) and the outcome variable (the reward  $R$  in our case). Specifically, the model with the following form:

$$R = \begin{cases} 1, & \text{if } w_0 + w_1 z + w_2 s + w_3 a + u_R \geq 0 \\ 0, & \text{otherwise} \end{cases}$$

where  $w_0, w_1, w_2, w_3$  are the model parameters and  $U_R \sim \text{logistic}(0, 1)$  is the background variable. In this case,  $P(R =$

---

### Algorithm 1 InsightFair: Algorithm for Achieving Equality in Long-Term Well-being

---

```

// Planning with Dynamics Fairness
1: procedure PLAN(augmented dynamics model  $f_\theta$ , initial state  $s_t$ , action distributions  $\text{CEM}(\cdot)$ , dynamics fairness DF)
2:   Compute the disparity in the current state:
        $d(s_t) = j_{s_t, z_0} \quad s_{t, z_1} j.$ 
3:   for  $k \geq \bar{f}_1, \dots, Kg$  do ▷ For each iteration
4:     if  $1(\text{DF})$  and  $d(s_t) \in \epsilon$  then
5:       Draw  $N$  samples from the same action distribution:
        $\hat{f}_{a_{t:t+H}} \quad \text{CEM}(\cdot) g_{i=1}^N$ 
6:     else
7:       Draw  $N$  samples from group-specific action distributions:  $\hat{f}_{a_{t:t+H}} \quad \text{CEM}_z(\cdot) g_{i=1}^N$  for each group  $Z = z.$ 
8:     end if
9:     Evaluate the samples using the augmented dynamics model  $f_\theta$  and the estimated well-being gap  $\text{TE}_{z_0, z_1}(G_{t:t+H}).$ 
10:    Sort the samples according to their regularized return:
        $J(a_{t:t+H}) = \sum_{t^0=t}^{t+H} \gamma^{t-t^0} {}^t R_{t^0} \quad \lambda \text{TE}_{z_0, z_1}(G_{t^0:t^0+H}).$ 
11:    Update the action distributions  $\text{CEM}(\cdot)$  to maximize the likelihood of the top  $M$  samples in  $\hat{f}_{a_{t:t+H}} g_{i=1}^N.$ 
12:  end for
13:  return the first action  $a_t$  in the optimal sample.
14: end procedure

// Main Algorithm
15: procedure LEARN(initial dynamics model  $f_{\theta_0}$ , initial data set  $D$  collected by a random policy)
16:  for  $e = 1, \dots, E$  do ▷ For each epoch
17:    Train the augmented dynamics model  $f_\theta$  on  $D.$ 
18:    Evaluate the dynamics fairness of  $f_\theta$  using Theorem 3.
19:    Set  $\text{DF} = 0$  if the environment satisfies dynamics fairness; otherwise, set  $\text{DF} = 1.$ 
20:    for  $t = 0, \dots, \text{Episode Length}$  do ▷ For each episode
21:      Observe the current state  $s_t$  from the environment.
22:      Select actions by  $a_t \quad \text{PLAN}(f_\theta, s_t, \text{CEM}(\cdot))$ 
23:      Execute the action  $a_t$  and observe the next state  $s_{t+1}$  and the reward  $r_t.$ 
24:      Update the data set  $D$  with the new data point.
25:    end for
26:  end for
27: end procedure

```

---

1 |  $Z = z, S = s, A = a$ ) can be written as:

$$\begin{aligned} & P(R = 1 | Z = z, S = s, A = a) \\ &= P(-w_0 > w_1 z + w_2 s + w_3 a + U_R) \\ &= (1 + \exp(-w_0 - w_1 z - w_2 s - w_3 a))^{-1} \\ &= L(w_0 + w_1 z + w_2 s + w_3 a), \end{aligned}$$

where  $L(\cdot)$  is the logistic function. Assuming that the background variables  $U_S \sim \mathcal{N}(0, \sigma_S^2)$  and  $U_A \sim \mathcal{N}(0, \sigma_A^2)$ , where  $\sigma_S \ll 1, \sigma_A \ll 1$ , we can now derive the analytical form of the natural direct and indirect effects of  $Z$  on  $R$  as follows:

$$\begin{aligned} \text{NDE}_{z_0, z_1}(R) &= \int_1^1 \int_1^1 P(R = 1 | Z = z_1, S = s, A = a) \\ &\quad - P(R = 1 | Z = z_0, S = s, A = a) P(S = s, a = a | Z = z_0) ds da \\ &= L(w_0 + w_1) - L(w_0) + 0(\sigma_S^2 + \sigma_A^2) \end{aligned}$$

$$\begin{aligned} \text{NIE}_{z_1, z_0}(R) &= \int_1^7 \int_1^7 P(R = 1 \mid Z = z_1, S = s, A = a) \\ &\quad (P(S = s, a = a \mid Z = z_0) - P(S = s, a = a \mid Z = z_1)) ds da \\ &= L(w_0 + w_1) - L(w_0 + w_1 + w_2 + w_3) + 0(\sigma_S^2 + \sigma_A^2). \end{aligned}$$

As a result, we can compute the exact value of both  $\text{NDE}_{z_0, z_1}(R)$  and  $\text{NIE}_{z_1, z_0}(R)$  using these equations. To simulate the effects in different environments, we fix the model parameters and only vary  $w_0$  to generate different environments. The obtained results are visualized in Figure 4.

## C.2 Experimental Setups of Section 3.2 and Section 3.3

In these experiments, we customize the ML-fairness-gym suite [D’Amour *et al.*, 2020] for group fairness problems. We conduct experiments on two environments, *Allocation-v0* and *Lending-v0*. Both environments are developed based on Gymnasium<sup>2</sup>, which is a widely accepted python library for reinforcement learning. The studied environments use a reward vector instead of a scalar to pass group-specific rewards to the agent, with most of the other APIs remaining the same as Gymnasium. The detailed descriptions are as follows:

*Allocation-v0*: In this environment, groups are characterized by different sites, each with a time-variant incident rates that determine the number of incidents (such as rat infestations) occurring at the site. The agent is tasked with allocating resources to these sites to minimize the total number of missed incidents.

- **Observation space:**  $s_t = [s_{t, z_0}, s_{t, z_1}]$ , where  $s_{t, z}$  is a scalar indicating the incident rate (mean of a Gaussian distribution) of group  $z$  at time step  $t$ .
- **Action space:**  $a_t = [a_{t, z_0}, a_{t, z_1}]$ , where  $a_{t, z}$  is a two-dimensional vector indicating the allocation of resources to group  $z$  at time step  $t$ .
- **Reward function:** For each group  $z$ , the reward  $r_{t, z}$  is defined as the negative number of missed incidents minus the cost of allocating resources at time step  $t$ . Note that, the number of incidents occurred at each site is sampled from a group-specific distribution and the agent is not aware of its value. Therefore, the agent needs to take uncertainty into account when allocating resources.
- **Transition dynamics:** For each group  $z$ , if more than half of the incidents are solved, its incident rate will decrease by a fixed amount; otherwise, its incident rate will increase.

*Lending-v0*: In this environment, each group has a dynamic credit score distribution, which determines the probability of repaying a loan. The agent is tasked with setting the loan threshold for each group to maximize the benefit of the bank.

- **Observation space:**  $s_t = [s_{t, z_0}, s_{t, z_1}]$ , where  $s_{t, z}$  is a five-dimensional vector indicating the credit score distribution (probability masses in a categorical distribution) of group  $z$  at time step  $t$ .

Table 1: Parameter settings for each environment and experiment.

Environment	Allocation-v0	Lending-v0
Experiments in Section 3.2		
$s_{0, z_0}$	6.0	(0.0, 0.2, 0.3, 0.3, 0.2)
$s_{0, z_1}$	6.0	(0.0, 0.2, 0.3, 0.3, 0.2)
$\alpha$	[0, 1.0]	[0, 0.2]
$\beta$	[0, 0.5]	[0, 0.5]
Experiments in Section 3.3 (unfair dynamics)		
$s_{0, z_0}$	6.0	(0.0, 0.2, 0.3, 0.3, 0.2)
$s_{0, z_1}$	6.0	(0.0, 0.2, 0.3, 0.3, 0.2)
$\alpha$	0.05	0.01
$\beta$	0.05	0.05
Experiments in Section 3.3 (fair dynamics)		
$s_{0, z_0}$	6.2	(0.0, 0.2, 0.35, 0.25, 0.2)
$s_{0, z_1}$	6.0	(0.0, 0.2, 0.25, 0.35, 0.2)
$\alpha$	0.0	0.0
$\beta$	0.0	0.0

- **Action space:**  $a_t = [a_{t, z_0}, a_{t, z_1}]$ , where  $a_{t, z}$  is a two-dimensional vector deciding the loan threshold for group  $z$  at time step  $t$ .
- **Reward function:** For each group  $z$ , the reward  $r_{t, z}$  is defined as the interest earned from the loans minus the cost of the loan defaults at time step  $t$ . Note that, in this environment, individuals with different credit scores are sampled from their group-specific distributions and this score determines whether or not they would repay the loan if granted. The agent is not aware of the true credit score of each individual and needs to take uncertainty into account when setting the loan threshold.
- **Transition dynamics:** For each group  $z$ , if an individual repays the loan, a fixed amount of probability mass will be shifted from its current credit score to a higher credit score; otherwise, the probability mass will be shifted to a lower credit score.

To evaluate the effectiveness the proposed techniques, we introduce two sets of parameters,  $[\alpha_1, \alpha_2]$  and  $[\beta_1, \beta_2]$ , to control the relative advantage of the groups in reward assignment and state transition. Along with the initial state of each group, these parameters can be used to generate various environments in which different groups take a dominant position in turn (when the environment is unfair) or behave similarly (when the environment is fair). The parameter settings for each environment and experiment are summarized in Table 1.

For the experiments in Section 3.2, parameters are provided in the form of intervals. We vary the parameters,  $[\alpha_1, \alpha_2]$  and  $[\beta_1, \beta_2]$ , in these fixed intervals to generate a series of environments with different natural direct effects. For example, when  $\text{NDE}_{z_0, z_1}(R)$  is being investigated, we fix  $\alpha_1 = \alpha_2$  and vary  $\beta_1$  and  $\beta_2$  in the interval  $[0, 0.5]$  so that the relative advantage in state transition is zero while the relative advantage in reward assignment varies based on  $\beta_2 - \beta_1$ . We divide the interval  $[0, 0.5]$  into 8 equal parts so there are  $8 \times 8 = 64$  environments in total for each experiment.

## C.3 Additional Experimental Results

<sup>2</sup><https://gymnasium.farama.org/>

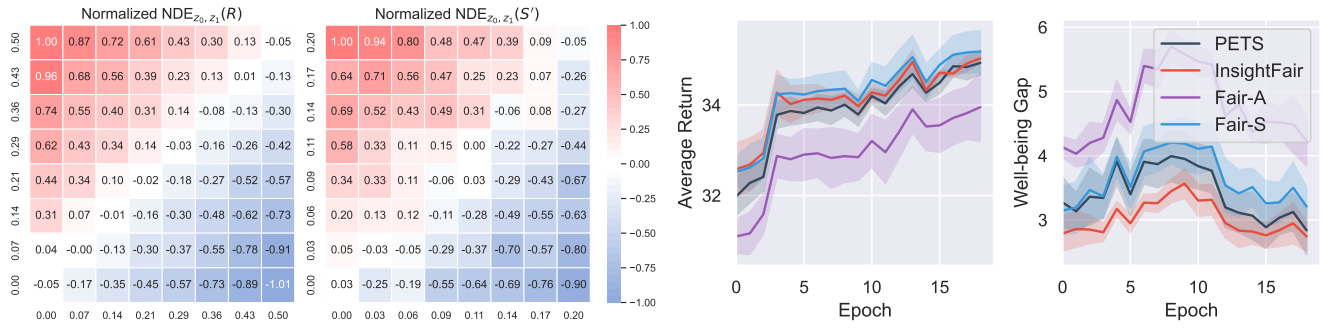


Figure 7: (Left) Results of evaluating dynamics fairness in the *Lending-v0* environment. Each tile represents the estimated natural direct effects,  $NDE_{z_0, z_1}(R)$  or  $NDE_{z_0, z_1}(S^0)$ , under specific parameter configurations indicated by the row and column. The proposed identification formula is capable of detecting the violation of dynamics fairness. (Right) Learning curves depicting the average return and the well-being gap in the *Lending-v0* environment (unfair dynamics). In this environment, InsightFair achieves the smallest well-being gap with a comparable return. Fair-A performs the worst since it cannot flexibly handle the different dynamics of the two groups.

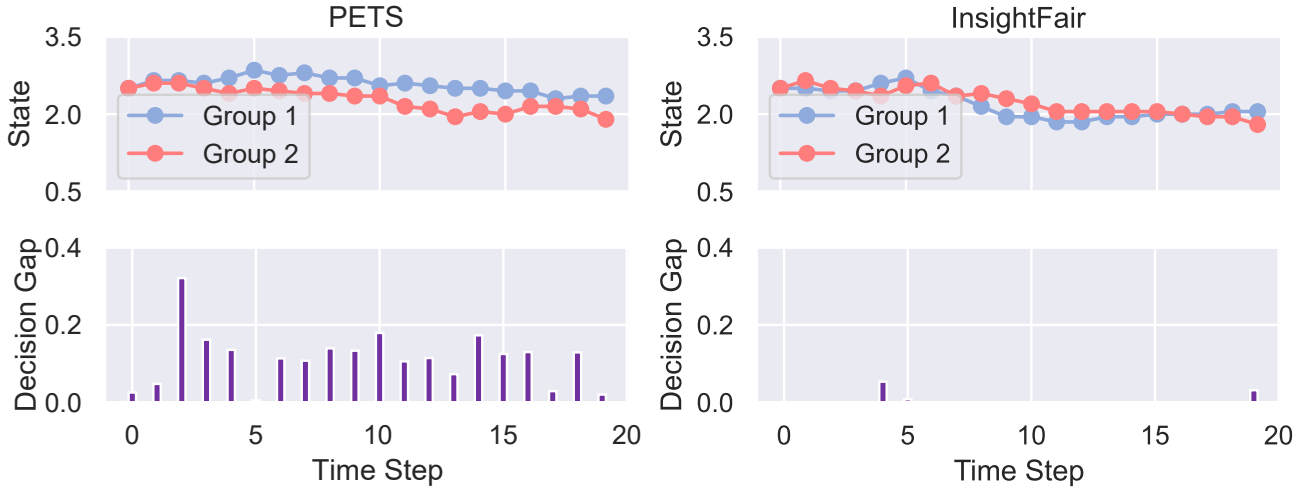


Figure 8: Changes in state and decision gaps within an episode in the *Lending-v0* environment (fair dynamics). Red denotes the advantaged group, which occupies a slightly better initial state as indicated in Table 1, and blue denotes the disadvantaged group. InsightFair has zero decision gaps when the states (measured by the weighted sum of credit scores and its probability mass) are sufficiently close, while PETS shows large decision gaps even when the states are close.

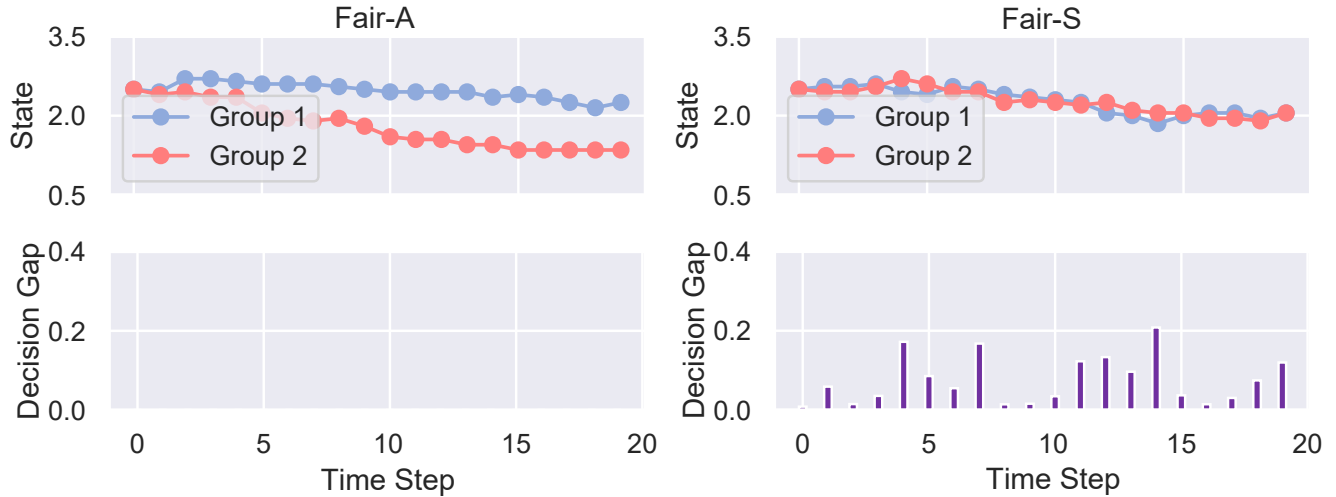


Figure 9: Changes in state and decision gaps within an episode in the *Lending-v0* environment (fair dynamics). The left figure shows Fair-A exhibits a zero decision gap but the state disparity may increase due to the difference in initial states and the aleatoric uncertainty in the dynamics. Fair-S tends to bring the states closer but the decision gap exists throughout the episode.