

Fairness Incentives in Response to Unfair Dynamic Pricing

JESSE THIBODEAU, Mila - Quebec AI Institute, HEC Montréal, Québec, Canada

HADI NEKOEI, Mila - Quebec AI Institute, Université de Montréal, Québec, Canada

AFAF TAÏK, Mila - Quebec AI Institute, Université de Montréal, Québec, Canada

JANARTHANAN RAJENDRAN, Dalhousie University, Nova Scotia, Canada

GOLNOOSH FARNADI, Mila - Quebec AI Institute, McGill University, Québec, Canada

The use of dynamic pricing by profit-maximizing firms gives rise to demand fairness concerns, measured by discrepancies in consumer groups' demand responses to a given pricing strategy. Notably, dynamic pricing may result in buyer distributions unreflective of those of the underlying population, which can be problematic in markets where fair representation is socially desirable. To address this, policy makers might leverage tools such as taxation and subsidy to adapt policy mechanisms dependent upon their social objective. In this paper, we explore the potential for AI methods to assist such intervention strategies. To this end, we design a basic simulated economy, wherein we introduce a dynamic social planner (SP) to generate corporate taxation schedules geared to incentivizing firms towards adopting fair pricing behaviours, and to use the collected tax budget to subsidize consumption among underrepresented groups. To cover a range of possible policy scenarios, we formulate our social planner's learning problem as a multi-armed bandit, a contextual bandit and finally as a full reinforcement learning (RL) problem, evaluating welfare outcomes from each case. To alleviate the difficulty in retaining meaningful tax rates that apply to less frequently occurring brackets, we introduce FairReplayBuffer, which ensures that our RL agent samples experiences uniformly across a discretized fairness space. We find that, upon deploying a learned tax and redistribution policy, social welfare improves on that of the fairness-agnostic baseline, and approaches that of the analytically optimal fairness-aware baseline for the multi-armed and contextual bandit settings, and surpassing it by 13.19% in the full RL setting.

CCS Concepts: • **Applied computing** → **Economics**; • **Theory of computation** → **Reinforcement learning**; **Algorithmic mechanism design**;

Additional Key Words and Phrases: Fairness, Contextual bandits, Reinforcement learning, Market simulation

1 INTRODUCTION

Firms equipped with modern compute power and vast consumer data logs may enjoy the economic benefit of engaging in the business practice of dynamic (or personalised) pricing. In doing so, profit-maximizing firms are able to charge potential customers, or customer segments, individualized prices based on estimates of their willingness-to-pay. From an efficiency standpoint, dynamic pricing has been shown to increase firm profitability and sales speeds [25]. However, the welfare implications of the practice are somewhat unclear, and may occasionally yield socially undesirable outcomes as evidenced in the markets for insurance and housing, among others [5, 29]. For instance, while health insurers are known to make extensive use of dynamic pricing, survey data reveal that members of the hispanic population in the U.S. are on average roughly twice as unlikely to have healthcare coverage as members of the black population, who are in turn twice as unlikely as members of the white and asian populations [20]. In addition, substantial price differences between consumers may lead to negative perceptions of fairness [17], which can have their own disparate impact

on market outlooks with respect to buyer participation, thereby propagating existing disparities. In housing, new constructions are often priced such that they are more appealing to privileged groups, thus contributing to growing wealth gaps and gentrification. Other sectors where buyer distributions should reflect those of the underlying population are the financial services and public transportation.

Dynamic pricing is generally employed by firms to assign prices to consumer segments such that the expected profit derived from each is maximized. Thus, a profit-maximizing firm engaging in dynamic pricing rarely bears any weight to its resulting buyer distribution, which, if not reflective of the underlying population, may be considered unfair. In this research, we address the problem of such imbalanced buyer distributions. Specifically, we examine the issue of dynamic pricing under *demand fairness* [6], in markets where buyer distributions should be proportional to that of the underlying population. We begin by showing how profit-maximizing price allocations may fail to satisfy demand fairness. We then employ nonlinear programming (NLP) at the firm level to demonstrate how producers willing to consider fairness in their price allocations may do so with often minimal sacrifice to profitability. We call this case *self-regulation*. Building on this, we acknowledge that in many cases, expecting a profit-maximizer to take it upon itself to consider fairness in their price allocations is unreasonable. Thus, we consider how a benevolent social planner (SP) might leverage available policy tools, notably taxation and subsidy, in order to dynamically incentivize market participation among underrepresented consumer groups, while penalizing unfair firm behaviour. To achieve this, we train an agent using reinforcement learning (RL) methods to generate financial incentives aimed at improving demand fairness, i.e. reducing the distributional gaps in buyer populations. In addition, we allow it to learn a simple redistribution scheme whereby consumers are subsidized based on their group membership. We find that social welfare can be improved using taxation aimed at incentivizing firms to adopt fairer behaviours. With the inclusion of subsidy, we find that firm profits and fairness outcomes both increase on average relative to our positive control benchmark.

Throughout our work, we make use of RL methods, which have a growing presence in the body of literature surrounding dynamic pricing and mechanism design. Specifically, we apply soft actor-critic to a variety of mechanism design tasks, from multi-armed bandits, to contextual bandits, and finally to a full RL problem. In addition, RL has been applied to several domains such as modelling markets [15], efficient energy pricing [16, 18] and insurance pricing [23], where RL was shown to be useful in such scenarios where demand for a specific product is either unobserved or misunderstood

to the point that it is difficult to uncover via traditional analytical methods.

Furthermore, the rise in use cases for AI in industry indicates that RL methods will likely continue to gain popularity for applications to dynamic pricing. Finally, RL methods are scalable and therefore appropriate for the complexity of large consumer bases. Our research proactively addresses the social challenges arising from dynamic pricing, notably via the following contributions:

- Introducing a new framework featuring three distinct policy mechanisms aimed at addressing demand fairness within single-product markets;
- Implementing various fairness incentive mechanisms through a range of economic policy variants, formalized via multi-armed and contextual bandits, as well as a full RL formulation;
- Presenting FairReplayBuffer, a replay buffer for the soft actor-critic (SAC) algorithm, designed specifically for learning fairness taxation;
- Proposing a combined approach of subsidy and taxation to effectively address demand fairness issues and enhance social welfare;
- Conducting a comprehensive evaluation of our proposed framework through a simulation study involving firms each addressing two distinct consumer behaviours ¹.

2 RELATED WORK

Given the interdisciplinary nature of our work, we conduct a literature review consisting of topics from economics, specifically in the subfields of welfare economics and consumer choice theory, as well as a broad overview of fairness applications in sequential decision making tasks.

2.1 Fairness in Dynamic Pricing

The study of welfare within the context of policy design began as a field of economics and gradually found applications in the fields of operations research [9] and computer science [7]. Welfare theory covers a broad variety of subtopics, which include definitions of fairness and how these interact given an incumbent policy. For instance, a recent paper by Cohen et al. [6] formalizes multiple fairness definitions for dynamic pricing in terms of price, demand, consumer surplus, and no-purchase valuation, while proving analytically that no two of these are simultaneously satisfiable. Our analysis uses the demand fairness definition [6, 14] as it most closely measures the disparate impact that dynamic pricing may have on buyer distributions, and which is well motivated by applications in education and healthcare [14]. In addition, Bertsimas et al. [3] define proportional fairness, a fairness criterion ensuring that the relative welfare improvement among one population subgroup exceed the corresponding welfare loss among another. Proportional fairness is useful in illustrating the welfare tradeoff between population subgroups under different policies. While the aforementioned works are related through their use of constrained optimization and

linear programming, Maestre et al. [19] use RL to impose fairness using Jain’s index, which they treat as a measure of fairness in the price allocations between groups. While we gain inspiration for our consumer demand curves from theirs, they assume that firms will be self-regulating, while we introduce a benevolent SP to generate fairness incentives.

2.2 Economics Foundations

In this work, we explore dynamics within single-product markets. Consumer demand behaviours are expressed as purchase probabilities, a formulation commonly applied to evaluate how consumers might react to price fluctuations. The concept of nonlinear consumer preference was first examined in economic theory by Becker [2] and later formalized in seminal work by McFadden [21] and continues to prove foundational in consumer choice modelling [22]. In addition, we examine welfare through the lens of fairness, as did by Fleurbaey [8], who advocates for not only fair outcomes but also fairness in the processes that lead to such outcomes.

2.3 Fairness in Sequential Decision Making

Fairness in sequential decision-making is a critical concern as algorithms increasingly influence societal outcomes, in fields such as healthcare [24], loan approval [12], and recommender systems [26]. Existing studies have established foundational approaches to fairness and highlighted challenges in ensuring equitable algorithmic decisions over time. Joseph et al. [13] introduce fairness constraints in bandit algorithms to prevent long-term disadvantages for individuals or groups. Gillen et al. [10] tackle the implementation of fairness when fairness criteria are undefined, proposing a framework for online learning that adapts to evolving fairness metrics. Yin et al. [27] explore the delayed impacts of fairness-aware algorithms, revealing how short-term equitable decisions can lead to unfair outcomes in the long run.

Within this body of research, the closest application to our work is the AI Economist [28], where agents interact in a simulated economy by exchanging goods and services, while a SP aims to learn a taxation strategy that improves social welfare, defined as the product of equality and productivity. While their work is effective at showcasing emergent behaviours among consumer-workers under incumbent tax regimes, our focus lies on how an SP can impact societal outcomes by influencing firm objectives. In addition, while the AI Economist uses a fixed uniform tax redistribution policy, we formulate the subsidy as an additional component to the learning problem, as do Abebe et al. [1], who explore approaches for welfare-maximizing subsidy allocation under income shocks. However, while they use *min-sum* and *min-max* formulations, we consider average welfare across consumer groups.

3 PROBLEM FORMULATION AND METHODOLOGY

In this section, we first define a simulated economy, outlining the consumer dynamics and possible scenarios in which a social planner might interact with firms. We follow with definitions for our firm and social planner learning problems, and finish with a discussion on the solution methods employed.

¹Our code will be made publicly available upon acceptance of this work.

3.1 Consumer Environment

A firm's objective in using dynamic pricing is to charge the maximal price that a prospective consumer is likely willing to pay based on demand estimates. We design a single-product market consisting of distinct consumer groups, wherein members from each group decide whether to purchase a product at a price determined by the firm. Each consumer group i has a unique purchase probability distribution, expressed as a discrete choice model by

$$\mathbb{P}_i(\text{purchase} = 1 \mid p) = [1 + e^{-(b_i + w_i \times p)}]^{-1}, \quad (1)$$

where p is the price assigned to the good, and parameters b_i and w_i capture different characteristics of consumer profile i regarding their sensitivity to price fluctuations. Once a price assignment has been made by the firm, we obtain each consumer profile's purchase probability from Equation 1. By linearity of expectation, we therefore anticipate that the number of consumers from group i that purchase a product at its given price can be computed by $E[n_i] = N_i \times \mathbb{P}_i(\text{purchase} = 1)$, where N_i is the number of consumers belonging to group i . In our experiments, we approximate these outcomes by simply passing these probabilities into a sequence of N_i Bernoulli trials, such that $n_i \sim \mathcal{B}(N_i, \mathbb{P}_i(\text{purchase} = 1))$, where n_i is a Bernoulli random variable corresponding to the resulting number of consumers from group i who purchased the good. This provides an element of stochasticity to our environment, where purchase outcomes may vary over constant price assignments. For simplicity, we consider the problem of a firm targeting two consumer groups, where one group has a higher sensibility to the price compared to the other one, which makes the consumers of this group underrepresented in the customer base of the firm. The fairness-agnostic firm might aim to maximize its profits with no regard for fairness, while its fairness-aware counterpart might weigh fairness outcomes into its maximization objective. In most real-world applications, we would not expect firms to explicitly self-regulate for fairness, as dynamic pricing is typically used to satisfy classical definitions of utility involving economic productivity (i.e., output, profit-maximization). To mitigate this, we introduce a dynamic social planner that aims to incentivize demand fairness by generating a tax schedule which determines the tax that will be levied on a firm based on their performance in terms of both profitability and fairness. In addition, we allow the social planner to redistribute the tax it collects to consumers as a subsidy to encourage higher market participation among underrepresented groups. Throughout our experiments, we consider demand fairness, and thus are concerned with the gap in purchase probabilities between consumer groups. We formalize our fairness notion as

$$\text{fairness}(p) = 1 - |\mathbb{P}_1(\text{purchase} = 1 \mid p) - \mathbb{P}_2(\text{purchase} = 1 \mid p)|, \quad (2)$$

so that smaller gaps in demand correspond to higher fairness scores. Given this notion of fairness, the social planner's ultimate objective is to maximize social welfare, denoted by

$$\text{swf} = \text{profit} \times \text{fairness},$$

which illustrates the nonlinear tradeoff between firm productivity and fairness outcomes [4, 28]. In Figure 1, we illustrate the economic process of the SP agent generating a tax and subsidy mechanism

based on the current welfare context. Below it, a firm sets prices such as to maximize profits under the SP's generated mechanism. Finally, the environment responds, whereby consumers decide whether to accept or reject firms' price outputs.

3.2 From Policy Preferences to RL Design

In the real world, governments² might adopt varying policy plans dependent upon their maximization objective. In this light, our exploration aims to provide an overview of possible policy scenarios, discussing a range of formulations that an incentive mechanism can take on while associating each to different hypothetical use cases. These include (i) a fixed policy mechanism, (ii) an adaptive policy mechanism conditioned on the current economic environment and (iii) an evolving policy framework that allows for frequent and ongoing changes to existing mechanisms. These are formulated as a multi-armed bandit problem, a contextual bandit problem and finally, a full RL problem. Below, we provide a detailed explanation of each policy mechanism.

3.2.1 Fixed Policy Mechanism: Multi-armed Bandit problem. If a policy planner aims to deploy an incentive mechanism to regulate a new market with uncertain or poorly understood dynamics, they might consider designing a tax and subsidy framework based on simulated data or data from economies with analogous market structures. In this scenario, solving a multi-armed bandit problem would be the most appropriate approach, with policy actions tailored to suit the average firms represented in the data.

3.2.2 Contextual Policy Mechanism: Contextual Multi-armed Bandit Problem. If the government was able to experiment with taxation frameworks that are conditioned on current market dynamics, it would want to create a dynamic policy mechanism. This can be formulated as a contextual bandit problem with an SP capable of generating incentives based on the observable economics context.

3.2.3 Evolving Policy Mechanism: RL Problem. Let us further suppose that an SP is at will to make frequent changes to policy frameworks, under the assumption that firms are equally adaptive and will always compute their best response to incumbent policies. Treating this as a full RL problem, a policy planner may adapt taxation and subsidy to maximize their notion of social welfare over time.

3.3 Firms

The *firm* component of our problem formulation consists of (i) a nonlinear program (NLP) for the fairness-agnostic firm, and (ii) a modified NLP for the fairness-aware firm.

3.3.1 Fairness-agnostic Firm. Having defined a single-product market environment (subsection 3.1), we proceed to formulate the fairness-agnostic firm's problem as the nonlinear program:

$$\begin{aligned} \max_p \quad & \sum_{i=1}^G n_i p \\ \text{s.t.} \quad & 0 < p \leq p_{\max} \end{aligned} \quad (3)$$

where G is the number of consumer groups in the market, and n_i is the number of consumers from group i who purchase and is itself a

²We use *government*, *policy planner* and *social planner* interchangeably




Fig. 1. Firms are capable of efficiently learning profit-maximizing pricing assignments from consumer demand responses. The social planner (SP) learns these implicitly through firms' fairness and profit scores and designs incentive mechanisms that use taxation and subsidy, pushing firms to minimize the gap in demand responses between consumer groups.

function of firm price p , as described in Equation 1. p_{max} may be a value set by firms or governments representing some hypothetical maximum price that can be assigned to the product at hand. Thus, this is a straightforward revenue-maximization problem faced by the firm.

3.3.2 Fairness-aware Firm. In addition, we consider the case where a firm may seek endogenously to maximize some objective which takes the fairness outcomes of their price allocations into account. In similar fashion to Equation 3, this self-regulating firm's problem becomes:

$$\begin{aligned} \max_p \quad & \sum_{i=1}^G n_i p \times \text{fairness} \\ \text{s.t.} \quad & 0 < p \leq p_{max} \end{aligned} \quad (4)$$

where the second term in the product enforces demand fairness by penalizing differences in demand between groups, as expressed in Equation 2.

3.4 Benevolent Social Planner

Given the range of possible preferences for policy makers discussed above, we continue with a mathematical formalism for our dynamic agents, including alternative formulations for our SP agent dependent upon possible policy objectives of government. We first introduce the shared formulations of the SP agent variants, then we outline their differences when introducing the sequential component to the full RL variant.

3.4.1 SP problem formulation. To alleviate the requirement for firms to proactively include fairness considerations in their pricing strategy, we introduce a benevolent SP to incentivize fair firm behaviour by leveraging the policy tools of taxation and subsidization. Specifically, the SP should generate a tax schedule which penalizes the firm based on its *fairness bracket*, and redistributes wealth so as to narrow the distributional gap between market participants. In this scenario, we assume that the SP adopts a revenue-neutral policy; at the end of every tax period, its tax revenue is fully offset by its expenditure, which here takes the form of a consumption subsidy awarded to consumers in proportions determined by their group membership, thereby closing the cycle of wealth in our simulated economy. The SP is formulated as a learning agent that evolves in an environment containing F firms, and each of the firms sets prices for a subset of the population constituted of individuals from both underrepresented and over-represented consumer groups.

For the cases where the SP is a contextual bandit or an RL agent, it observes a context vector from the context space $\mathcal{X}^{SP} = \{f, \phi\}$, where $f \in [0, 1]^F$ denotes a vector of individual firm fairness values, and $\phi \in [0, p_{max}]^F$ denotes a vector of firm profits, normalized per customer, and takes actions from a continuous action space

$$\mathcal{A} = \{(A_1, A_2) \mid A_1 \in [0, 1]^b, A_2 \in [0, 1]\},$$

where A_1 is a vector of size b with elements $\tau \in [0, 1]$ denoting tax rates for each of b tax brackets, and $A_2 \in [0, 1]$, denoting the proportion of the collected tax budget that will be awarded to the underrepresented group as a consumption subsidy in the following period. It follows that the remaining proportion of the tax budget

$1 - A_2$ will be distributed to the relatively overrepresented group. Therefore, to determine the effective price faced by a customer from the underrepresented group 1 and majority group 2, we update p via

$$p_1^t := p^t - \frac{A_2 \times \mathcal{B}}{n_1} \quad ; \quad p_2^t := p^t - \frac{(1 - A_2) \times \mathcal{B}}{n_2}, \quad (5)$$

where $\mathcal{B} = \sum_{i=1}^F \tau_i^{t-1} \times \phi_i^{t-1}$ denotes the total tax revenue generated by F firms in the previous period, and n_1, n_2 denoting the number of consumers belonging to group 1 and 2 respectively.

3.4.2 Social Welfare Over Time. For the multi-armed and contextual bandit policy scenarios discussed in subsubsection 3.2.1 and subsubsection 3.2.2, we formulate the social planner’s problem as one of setting a tax schedule and redistribution scheme that will remain in effect over a time horizon T . Rewards in these settings are therefore computed only at time T . We further assume that the SP is able to compute fairness values by evaluating the buyer distribution resulting from the firm’s learned pricing strategy. Consequently, the SP’s reward is formulated as

$$\mathcal{R}^{SP} = \frac{1}{F} \sum_{i=1}^F \sum_{j=1}^{G_j} n_i p_j (1 - \tau_j) \times \text{fairness}_j, \quad (6)$$

where F is the number of firms participating in the economy, and G_j is the number of consumer groups addressed by firm j . We note that, with the addition of taxation, the firm’s maximization objective becomes $\max_p \sum_{i=1}^G n_i p \times (1 - \tau)$ because they care about maximizing net profits.

For the RL solutions discussed in subsubsection 3.2.3, the SP’s objective is to maximize discounted reward over the same horizon. In the RL setting, the SP agent’s reward function is denoted

$$\mathcal{R}_{RL}^{SP} = \sum_{t=1}^T \gamma^{t-1} \left(\frac{1}{F} \sum_{j=1}^F \sum_{i=1}^{G_j} n_{i,t} p_{j,t} (1 - \tau_{j,t}) \times \text{fairness}_{j,t} \right),$$

where $\gamma \in [0, 1]$ is a discount factor.

3.5 Solution Methods

In this section, we outline the methods employed. The SP agent learns its policy for choosing good incentive mechanisms given the welfare context using soft actor-critic (SAC), an RL algorithm designed to solve Markov decision processes, where there are temporal transition dynamics. S with emphasis on sample efficiency. SAC draws samples from a *replay buffer* to perform policy gradient updates. Upon experimentation, our SAC agent expressed difficulty in generating taxation frameworks which were consistent in retaining meaningful information from lower brackets, which become less common as firms become fairer, leading us to introduce FairReplayBuffer, which alleviates this concern.

3.5.1 Soft Actor-Critic. In our study, we employed the soft actor-critic (SAC) algorithm [11], an advanced RL technique suited for continuous action spaces. SAC is distinguished by its incorporation of entropy maximization into the reward function, promoting exploration while simultaneously striving for optimal policy development. This approach enhances the algorithm’s sample efficiency and stability, making it particularly suited for environments where

precise, adaptive control is required. While typically used for sequential decision-making over time, SAC is also appropriate for our bandit setup, as its sample efficiency makes it applicable to public policy where experimenting over long periods entails risk.

3.5.2 FairReplayBuffer. A fundamental challenge in generating meaningful incentives for less frequently seen contexts across the fairness space lies in how our SAC agent performs gradient updates from a traditional First-In-First-Out (FIFO) replay buffer, which would mostly consist of common firm behaviours in training. Due to this, a learning agent in our setting would gradually forget any information gained from infrequently observed regions of the observation space, thereby reducing its effectiveness on out-of-sample fairness brackets. To address this, we define a replay buffer that uniformly stores examples from across the fairness space, allowing the learning agent to output meaningful tax actions for each region. The procedure by which we populate a FairReplayBuffer is outlined in algorithm 1. The algorithm, along with ablations, can be found in Appendix C.

4 EXPERIMENTAL DESIGN AND ANALYTICAL BASELINES

In this section, we define our working examples, which are designed to demonstrate the effects that a benevolent SP may have on market dynamics. We begin by defining four firms, each facing varying demand distributions between two consumer groups. Then, we record each firm’s optimal pricing strategy and resulting fairness scores under profit-maximizing objectives, with and without self-regulation for fairness, in the absence of policy intervention. We set $p_{max} = 10$. Finally, we allow the SP to dynamically learn a taxation framework with respect to profit and fairness outcomes which will push firms to adopt prices which promote fairness levels approaching those achieved by fairness-aware firms.

4.1 Multi-firm Setup

We conduct our experiments in an environment where each of four firms addresses two distinct consumer groups. To describe their behaviours via purchase probability distributions, we apply parameters b and w found in Appendix A (some of which are borrowed from Maestre et al. [19]), to Equation 1, in order to obtain the curves depicted in Figure 2. Using the ground-truth purchase probabilities embedded in our consumer environment, we can analytically derive optimal prices for each firm’s objective. For instance, the fairness-agnostic firm’s optimal price allocation is obtained by finding

$$\arg \max_p \mathbb{P}_i(\text{purchase} = 1 \mid p) \times p.$$

4.2 Baselines

Assuming firms are efficient at finding profit-maximizing prices under current market conditions (i.e., demand distributions and incentive schemes), and carrying our working example forward, we would expect fairness-agnostic firms to converge to the prices and fairness values in Table 1a in the absence of policy intervention. By solving for the profit-fairness objective, also in the absence of policy intervention, i.e. self-regulation, we find that firms may change their

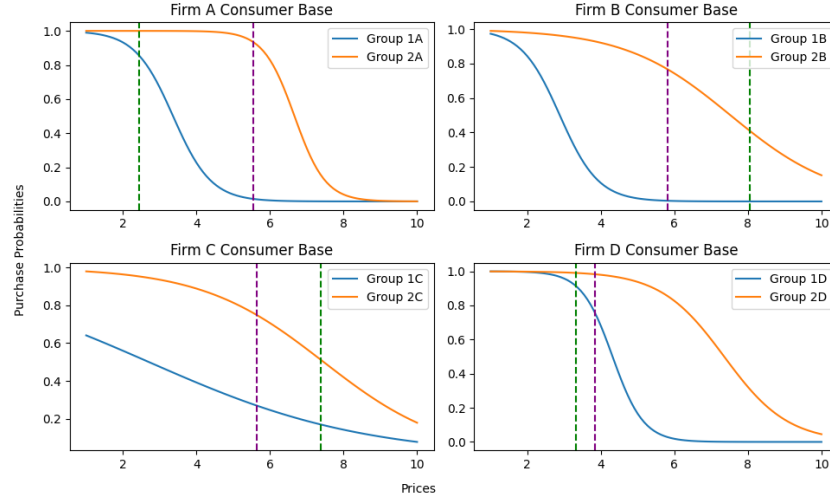


Fig. 2. Consumer profiles: Each firm serves two consumer groups. For each, group 1 may be considered to have lower tolerance to rising prices than group 2. The vertical purple and green lines represent the analytical profit-maximizing price assigned by each firm in the fairness-agnostic and fairness-aware cases respectively, with the resulting vertical gaps between the orange and blue lines illustrating important discrepancies in purchase probabilities between consumer groups under price allocations associated to both behaviours.

Firm^{agnostic}	A	B	C	D	Avg
f	0.08	0.24	0.52	0.78	0.41
ϕ	2.63	2.24	2.87	3.34	2.77
swf	0.21	0.54	1.49	2.61	1.14

(a) Fairness, profit, and corresponding social welfare scores achieved by fairness-agnostic firms A, B, C and D in the absence of policy intervention.

Firm^{aware}	A	B	C	D	Avg
f	0.85	0.59	0.66	0.92	0.76
ϕ	2.27	1.65	2.51	3.17	2.40
swf	1.93	0.97	1.66	2.92	1.82

(b) Fairness, profit and corresponding social welfare scores achieved by fairness-aware firms A, B, C and D, who self-regulate to .

Table 1. Resulting fairness (f), profit (ϕ) and social welfare (swf) values from analytical optimal price assignments for fairness-agnostic and fairness-aware firms.

prices significantly while sometimes achieving profits comparable to those in the fairness-agnostic case. These welfare-maximizing results for the fairness-aware firm (Table 1b) serve as analytical baselines which can be used as a positive control for comparison with the results from our bandit and RL experiments. Further, referring back to Figure 2 provides an alternative visualization of fairness outcomes, where the difference between the orange and blue curves is the gap in demand between consumer profiles resulting from the firms’ price assignments. We finally note from Table 1a that our initializations for each firm lead them to converge to profit-maximizing prices resulting in a broad coverage of the fairness space. This is a

design choice to demonstrate how optimal decisions from the firms’ standpoint can have unforeseeable fairness implications.

5 EMPIRICAL EVALUATION

We deploy a learning agent in the experimental setting defined in section 4 using our multi-firm setup. While we initialize firms as described, we allow for a degree of stochastic fluctuations in consumer behaviour to occur, reflecting how demand can change subtly in the short-run. To reflect this, we sample purchase probabilities from a normal distribution around means given by the output of the sigmoid demand curves denoted in Equation 1. The curves themselves are parameterized by w_i and b_i , as found in Appendix A. This indicates that the social planner may experience a different reward from having taken identical actions. This, paired with the Bernoulli trials which determine a consumer group’s real purchase outcomes, introduces a degree of noise into our environment, thereby adding complexity to our learning problem. The purchase probability for a consumer group in a given training step is therefore sampled as $\mathbb{P}_{i,j} \sim \mathcal{N}(\mathbb{P}_{i,j}, \sigma^2)$. In our experiments, we set $\sigma = 0.05$, allowing for small fluctuations in demand. We show the best results in terms of swf below. All experiments were run for 20 seeds. Our experiments are designed to answer the following questions:

- (1) What effects do different policy mechanisms, reflected through RL design choices, have on welfare?
- (2) How are individual market participants affected by a dynamic policy mechanism?
- (3) How can RL-generated policies be adopted to assist policy makers and what are the long-term effects of such policy designs?

We train the SP to tax firms based on where they lie along the fairness dimension, which we discretize into 5 brackets, using the

multi-armed bandit, contextual bandit, and full RL formulations (see section 3). Simultaneously, the SP chooses a proportion of the tax budget that it will re-distribute to the underrepresented consumer group in the next timestep. Our results show that, due to efficient wealth redistribution, the introduction of dynamic policies can approach and even improve upon the global welfare obtained from the analytical optimal case where firms are self-regulating.

5.1 RQ1: Welfare Effects of Various RL Methods

Here we evaluate the impact each formulation has on global welfare. Table 2 breaks down the average results for the multi-armed bandit, contextual bandit and full RL settings. First, we note that the RL SP yields the best *swf* results, on-par with the fairness-aware firm baseline, for which results are shown in Table 1b, illustrating the effectiveness of the proposed incentive mechanisms. However, while there is a clear pattern of improvement in *swf* as the problem formulation complexifies, it is worth noting the limitations of each implementation. While a multi-armed bandit SP underperforms relative to other solution methods, it nonetheless significantly improves welfare outcomes compared to profit-maximizing (fairness-agnostic) firms while requiring the least amount of data from the environment. This is suitable, under stationarity assumptions, for new markets with unknown dynamics.

The contextual bandit SP approaches the analytical optimal solution for self-regulating firms. This formulation, while resulting in a more adaptive policy framework, requires more granular environment data in training, which may be difficult to obtain in the real world. Finally, the full RL formulation yields the best results in terms of welfare, achieving improvements relative to the baseline optimum on average. However, it intervenes at higher frequencies, re-defining policy mechanisms at each timestep. In addition to requiring granular training data, this may also reduce its range of applications to highly dynamic or volatile markets with shorter cycles. Globally, our results reveal that it is possible for a dynamic SP to incentive fairness in markets while retaining or even improving economic productivity.

5.2 RQ2: Impact on individual firms

We continue with a discussion on the individual impact to welfare bore by firms. Table 3 denotes the per-firm welfare comparisons between the analytical baselines in the fairness-agnostic and fairness-aware cases, and firms subjected to the RL SP’s learned incentive mechanisms. We note that, while most firms experience improvements in individual welfare, firm C’s is significantly reduced. This is likely due to the almost parallel nature of the demand curves of the two groups, and the fact that the fairer price is higher than the profit-maximizing one. Nonetheless, the policy mechanism applied by the social planner RL agent significantly improves welfare outcomes compared to the fairness-agnostic case, and even manages to surpass those in the (unrealistic) case where firms self-regulate for fairness, with a 13.19% improvement in global welfare. We note further that while global welfare improves under this setting, it is impossible for welfare to improve for every individual firm. Firm C’s welfare loss indicates that there is indeed no such thing as a free lunch.

5.3 RQ3: Deployment and performance over time

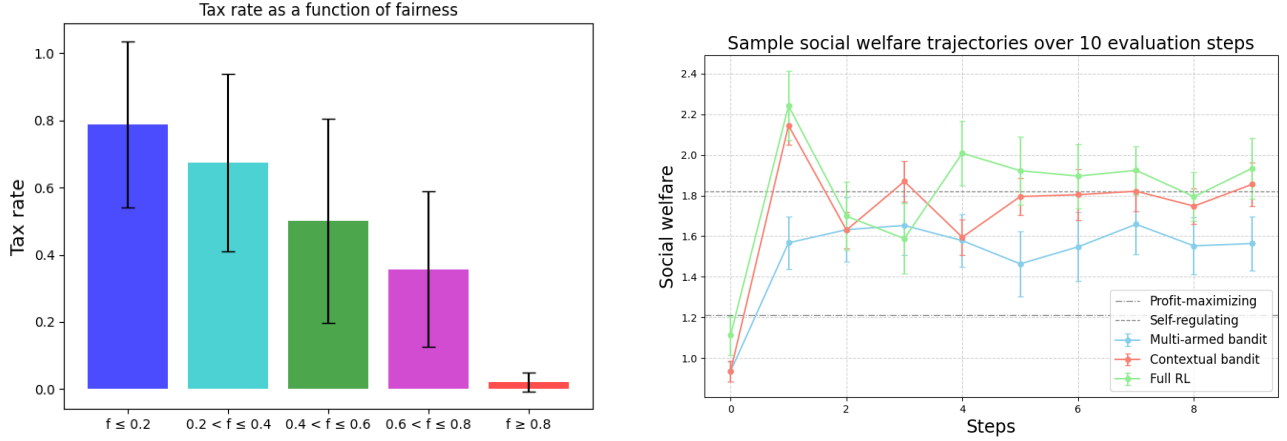
As a high social impact application, it is necessary to study the usability and long-term impact of our framework. This is particularly relevant in the case that a similar solution is used to inform and assist policy makers in mechanism design. We note in Figure 3 the emergence of a clear pattern between fairness and taxation, in line to a large degree with human intuition: firms demonstrating fairer behaviours are incentivized by lower tax rates, and this trend persists for each problem formulation (Appendix C). In addition, each formulation of the SP learns to give the majority of the subsidy budget to the underrepresented consumer group (noted in Table 2). For example, a proportion $S = 0.66$ means that for every unit of tax budget the social planner collects, it redistributes 0.66 units to *group 1*, and 0.34 units to *group 2*, thereby reducing their effective price and increasing their market participation. In our evaluation of welfare over time, the full RL SP outperforms both multi-armed and contextual bandit SPs in most timesteps, as shown in Figure 3.

6 LIMITATIONS AND SOCIAL IMPACT

There exist opposing schools of thought regarding fiscal policy and policy intervention in general. While one would advocate for intervention for the sake of equality or social progressivity, another would claim that any form of intervention would disrupt the economic processes by which consumption and production are governed, and that any deviation from these processes is inefficient. In our simplified setting, we build an economic subsystem to illustrate the tradeoffs between policy intervention and its absence. Profit-maximizing firms are not likely to be intrinsically motivated to behave fairly with respect to their buyer distributions, especially when said behaviour entails a sacrifice in profitability. We illustrate this in Table 1a, where profit-maximizing prices may have unforeseeable fairness outcomes dependent on distributions in demand. To mitigate this, we introduce a third-party social planner whose objective is to generate incentives which nudge firms into considering social welfare. Our experiments ultimately show that net improvements in social welfare are achievable with tailored taxation and redistribution schemes learned through a range of RL methods, each with their own applications. We note further that the adjusted prices paid by consumers (post-subsidy) are effective at increasing market participation.

Given that taxation and subsidy have important social implications, deploying AI-based strategies might not be easily accepted by decision-makers or the general public due to interpretability and accountability concerns. While we were able to generate seemingly interpretable tax schedules using RL, the chosen tax brackets themselves might not be easily accepted or interpreted, especially since they are linked with a penalty. Furthermore, while the resulting pattern of tax schedules is consistent across random seeds, the values for each tax bracket varies. Deployment of such a solution therefore requires additional considerations for robustness against distribution shifts and possible poisoning attacks. Additionally, dynamic pricing is becoming more algorithmically driven, thus requiring more studies into how these algorithms might adapt in response to our incentive strategies.

Firms	Multi-armed bandit ($S = 0.63$)					Contextual bandit ($S = 0.65$)					Full RL ($S = 0.66$)				
	A	B	C	D	Avg	A	B	C	D	Avg	A	B	C	D	Avg
f	0.72	0.66	0.58	0.93	0.72 ± 0.04	0.87	0.80	0.58	0.94	0.80 ± 0.02	0.84	0.79	0.58	0.94	0.79 ± 0.03
ϕ	2.39	1.69	1.67	2.96	2.25 ± 0.07	2.35	1.89	1.70	2.97	2.23 ± 0.07	2.37	2.26	2.00	3.43	2.51 ± 0.11
swf	1.82	0.93	0.89	3.00	1.64 ± 0.14	2.04	1.51	0.98	2.78	1.78 ± 0.09	1.99	1.78	1.15	3.23	2.06 ± 0.15

Table 2. Aggregate Results per experimental design. For the full RL, $\gamma = 0.99$. Standard errors are reported only for averages to reduce clutter.Fig. 3. *Left*: Tax actions taken by the RL social planner. Reported policy mechanisms record SP actions averaged over 20 seeds. *Right*: Social welfare trajectories during evaluation for the SP’s learned policy frameworks from multi-armed bandit, contextual bandit, and RL formulations.

Firm	A	B	C	D	\overline{swf}
$swf_{agnostic}$	0.21	0.54	1.47	2.61	1.21
swf_{aware}	1.93	0.77	1.66	2.92	1.82
swf_{SP}	1.99	1.78	1.15	3.23	2.06
$\overline{swf} \Delta_{SP}^{aware}$	3.11%	131.17%	-30.72%	10.62%	13.19%

Table 3. Summary of the welfare effects of the social planner’s learned taxation and redistribution scheme compared to analytical baselines. The bottom row denotes the percentage change in welfare between the fairness-aware baseline and the social planner’s generated *fairness tax*.

Our research has several simplifying assumptions, making it difficult to utilize this framework in real-world scenarios. For instance, considering only demand fairness might lead to counter-intuitive positive benchmarks, as seen in our examples where analytically-optimal prices are, in some cases, higher than their fairness-agnostic counterpart. Further, it is difficult to conceive of a mechanism for assessing fairness on a per-firm basis in the real world. This would presumably necessitate policy makers to access broader market data, as well as consumer-level data from individual firms, from which they may compute distributions which would be necessary for a fairness-based tax to be enforced. Additionally, our experiments consider a single-product market with no competitive dynamics between firms. While our setting illustrates a real concern for the disparate impacts of dynamic pricing on consumer distributions, introducing consumer choice would increase realism, but would

perhaps make achieving improvements in welfare more challenging. Lastly, we focused on demand fairness given how it reflects accessibility to essential goods and services in critical domains such as healthcare, but further research is necessary to explore how our framework can be applied to alternative fairness notions such as price and consumer surplus fairness, which expand to other applications such as loan approval.

In a profit-driven market, our work shows that merely relying on self-regulation is not sufficient, thus requiring interventions from policy makers to ensure equity and equal access to multiple goods and services. While we do not present our framework as the ultimate solution, we believe it can guide future research and inspire exploration at the intersection of AI and society, particularly in the design of dynamic incentive mechanisms, especially in applications where dynamic pricing has high social and economic impacts.

7 CONCLUSION

When used for fairness-agnostic profit-maximization, dynamic pricing can have harmful consequences with respect to equal access. We demonstrate that, by leveraging the policy tools of taxation and subsidy, a dynamic social planner may create global efficiencies which would have otherwise been unattainable, even with locally optimal firm behaviour by self-regulation. This not only reveals that a dynamic third party may be trained to improve fairness outcomes in a broad sense, but may also generate policies which are tailored to specific markets and consumer dynamics, and that these policies may be deployed in ways that improve upon free-market outcomes. While we believe these findings to be relevant for monopolistic

markets or geographically-delimited markets with little potential for consumer displacement, we see great value in extending this work to competitive markets, where consumers have more freedom in choosing among different products. Additionally, exploring more realistic and challenging environments with imbalanced consumer distributions is necessary to further understand and analyze the financial and social implications of such frameworks. Finally, extending the framework to use various fairness definitions would be needed to ensure its compatibility with other application domains.

ACKNOWLEDGMENTS

Funding support for project activities has been partially provided by Canada CIFAR AI Chair, Google award, MEI award, and the Natural Sciences and Engineering Research Council of Canada (NSERC). We also express our gratitude to Compute Canada for their support in providing facilities for our evaluations. Finally, we thank Nobuhiro Kiyotaki for the valuable insight he provided during the early ideation stages of this research.

REFERENCES

- [1] Rediet Abebe, Jon Kleinberg, and S Matthew Weinberg. 2020. Subsidy allocations in the presence of income shocks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 7032–7039.
- [2] Gary S Becker. 1962. Irrational behavior and economic theory. *Journal of political economy* 70, 1 (1962), 1–13.
- [3] Dimitris Bertsimas, Vivek F Farias, and Nikolaos Trichakis. 2011. The price of fairness. *Operations research* 59, 1 (2011), 17–31.
- [4] Dimitris Bertsimas, Vivek F Farias, and Nikolaos Trichakis. 2012. On the efficiency-fairness trade-off. *Management Science* 58, 12 (2012), 2234–2250.
- [5] Jose M Betancourt, Ali Hortaçsu, Aniko Oery, and Kevin R Williams. 2022. *Dynamic price competition: Theory and evidence from airline markets*. Technical Report. National Bureau of Economic Research.
- [6] Maxime C Cohen, Adam N Elmachtoub, and Xiao Lei. 2022. Price discrimination with fairness constraints. *Management Science* 68, 12 (2022), 8536–8552.
- [7] Shantanu Das, Swapnil Dhamal, Ganesh Ghalme, Shweta Jain, and Sujit Gujar. 2022. Individual fairness in feature-based pricing for monopoly markets. In *Uncertainty in Artificial Intelligence*. PMLR, 486–495.
- [8] Marc Fleurbaey. 2008. *Fairness, responsibility, and welfare*. OUP Oxford.
- [9] Guillermo Gallego, Huseyin Topaloglu, et al. 2019. *Revenue management and pricing analytics*. Vol. 209. Springer.
- [10] Stephen Gillen, Christopher Jung, Michael Kearns, and Aaron Roth. 2018. Online learning with an unknown fairness metric. *Advances in neural information processing systems* 31 (2018).
- [11] Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, and Sergey Levine. 2018. Soft Actor-Critic Algorithms and Applications. <https://doi.org/10.48550/ARXIV.1812.05905>
- [12] Yaowei Hu and Lu Zhang. 2022. Achieving long-term fairness in sequential decision making. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 9549–9557.
- [13] Matthew Joseph, Michael Kearns, Jamie H Morgenstern, and Aaron Roth. 2016. Fairness in learning: Classic and contextual bandits. *Advances in neural information processing systems* 29 (2016).
- [14] Nathan Kallus and Angela Zhou. 2021. Fairness, welfare, and equity in personalized pricing. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*. 296–314.
- [15] Alexander Kastius and Rainer Schlosser. 2021. Dynamic pricing under competition using reinforcement learning. *Journal of Revenue and Pricing Management* (2021), 1–14.
- [16] Byung-Gook Kim, Yu Zhang, Mihaela Van Der Schaar, and Jang-Won Lee. 2015. Dynamic pricing and energy consumption scheduling with reinforcement learning. *IEEE Transactions on smart grid* 7, 5 (2015), 2187–2198.
- [17] Simon Lee, Abdou Illia, and Assion Lawson-Body. 2011. Perceived price fairness of dynamic pricing. *Industrial Management & Data Systems* (2011).
- [18] Renzhi Lu, Seung Ho Hong, and Xiongfeng Zhang. 2018. A dynamic pricing demand response algorithm for smart grid: Reinforcement learning approach. *Applied energy* 220 (2018), 220–230.
- [19] Roberto Maestre, Juan Duque, Alberto Rubio, and Juan Arévalo. 2019. Reinforcement learning for fair dynamic pricing. In *Intelligent Systems and Applications: Proceedings of the 2018 Intelligent Systems Conference (IntelliSys) Volume 1*. Springer, 120–135.
- [20] Michael E. Martinez. 2022. QuickStats: Percentage of Uninsured Adults Aged 18–64 Years, by Race and Selected Hispanic Origin Subgroup – National Health Interview Survey, United States, 2020. *MMWR Morb Mortal Wkly Rep* 2022 71, 834 (2022). <https://doi.org/10.15585/mmwr.mm7125a3>
- [21] Daniel McFadden. 1972. Conditional logit analysis of qualitative choice behavior. (1972).
- [22] Hybrid Choice Models. 2002. Progress and Challenges. *Marketing Letters* 13, 3 (2002), 163–175.
- [23] Marlon Nieuwenhuis, Antony SR Manstead, and Matthew J Easterbrook. 2019. Accounting for unequal access to higher education: The role of social identity factors. *Group Processes & Intergroup Relations* 22, 3 (2019), 371–389.
- [24] Alvin Rajkomar, Michaela Hardt, Michael D Howell, Greg Corrado, and Marshall H Chin. 2018. Ensuring fairness in machine learning to advance health equity. *Annals of internal medicine* 169, 12 (2018), 866–872.
- [25] Rainer Schlosser and Martin Boissier. 2018. Dynamic pricing under competition on online marketplaces: A data-driven approach. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*. 705–714.
- [26] Maria Stratigi, Jyrki Nummenmaa, Evangelia Pitoura, and Kostas Stefanidis. 2020. Fair sequential group recommendations. In *Proceedings of the 35th Annual ACM Symposium on Applied Computing (Brno, Czech Republic) (SAC ’20)*. Association for Computing Machinery, New York, NY, USA, 1443–1452. <https://doi.org/10.1145/3341105.3375766>
- [27] Tongxin Yin, Reilly Raab, Mingyan Liu, and Yang Liu. 2024. Long-term fairness with unknown dynamics. *Advances in Neural Information Processing Systems* 36 (2024).
- [28] Stephan Zheng, Alexander Trott, Sunil Srinivasa, Nikhil Naik, Melvin Gruesbeck, David C Parkes, and Richard Socher. 2020. The ai economist: Improving equality and productivity with ai-driven tax policies. *arXiv preprint arXiv:2004.13332* (2020).
- [29] Yushu Zhu, Yue Yuan, Jiaxin Gu, and Qiang Fu. 2023. Neoliberalization and inequality: disparities in access to affordable housing in urban Canada 1981–2016. *Housing Studies* 38, 10 (2023), 1860–1887.

A DEMAND DISTRIBUTION PARAMETERS

In our experiments, we designed four firms A, B, C and D who each address a customer base whose demand distributions depend on parameters sampled in the table below.

Params	Firm A		Firm B		Firm C		Firm D	
	$g1$	$g2$	$g1$	$g2$	$g1$	$g2$	$g1$	$g2$
w	-1.926	-2.369	-1.9	-0.695	-0.340	-0.600	-2.369	-1.1526
b	6.4757	15.7900	5.4757	5.229	0.9195	4.4757	10.2290	8.4757

* In stochastic experiments, values are sampled normally around table values with a standard deviation of $\sigma = 0.1$ every 50,000 steps, and $\sigma = 0.0125$ in intermediate steps.

B ADDITIONAL EXPERIMENTS

In addition to the experimental results reported in the main section, we also ran experiments over a variety of reward functions without the use of a subsidy. For instance, including gross profit (pre-tax) into the social planner's reward function made it less concerned for the after-tax income of firms, and more with the overall size of the economy measured by the sum of net-profit and a tax budget. This meant that firms were taxed at very high rates, and while fairness was improved, the profit values reported do not reflect the welfare of individual firms.

Firms	Multi-armed bandit					Contextual bandit					Full RL				
	A	B	C	D	Avg	A	B	C	D	Avg	A	B	C	D	Avg
f	0.76	0.32	0.57	0.76	0.60	0.79	0.33	0.56	0.81	0.62	0.76	0.46	0.62	0.83	0.67
ϕ	2.39	2.16	2.71	3.22	2.62	2.37	2.16	2.69	3.33	2.63	2.39	1.97	2.52	3.23	2.53
swf	1.82	0.69	1.54	2.44	1.57	1.87	0.71	1.51	2.69	1.63	1.81	0.91	1.56	2.68	1.70

Table 4. Aggregate Results w/ Gross Profit

To evaluate firm welfare in the absence of subsidy, the social planner considers net profit in its reward. In this setting, it faces the tradeoff between welfare and taxation, but the tax budget is not reinjected into the economy, making it difficult to draw comparisons with the analytical baselines, where firms are self-regulating and do not pay taxes. Nonetheless, fairness improvements are still achievable.

Firms	Multi-armed bandit					Contextual bandit					Full RL				
	A	B	C	D	Avg	A	B	C	D	Avg	A	B	C	D	Avg
f	0.79	0.33	0.58	0.69	0.60	0.62	0.36	0.57	0.78	0.58	0.73	0.27	0.55	0.77	0.58
ϕ	1.98	1.19	1.35	2.19	1.43	2.07	1.42	2.44	2.91	2.21	2.30	1.52	2.26	2.89	2.24
swf	1.56	0.40	0.79	1.52	0.86	1.28	0.51	1.39	2.27	1.28	1.68	0.41	1.25	2.23	1.30

Table 5. Aggregate Results w/ Net Profit.

C FAIR REPLAY BUFFER

Here, we include details regarding the FairReplayBuffer, including the algorithm along with ablations highlighting the differences between the tax schedules learned by the SP agent when using FIFO, and FairReplayBuffer.

C.1 Algorithm

Algorithm 1: FairReplayBuffer	
Input: buffer capacity $ \mathcal{D} $, obs , $action$, $reward$, $done$, $infos$, brackets \mathcal{I} , batch size $ \mathcal{B} $	
Output: A replay buffer ensuring uniform distribution of experiences across the fairness-profit space	
1	Buffer initialization, $\mathcal{D} \leftarrow []$
2	Storage per bracket initialization, $\mathcal{S} \leftarrow \{\}$
	/* Procedure to add experiences */
3	Procedure ADD(obs , $action$, $reward$, $done$, $infos$):
4	Determine i based on obs
5	Append(obs , $action$, $reward$, $done$, $infos$) to \mathcal{D}
6	Append index of new experience to $\mathcal{S}[i]$
	/* Procedure to sample experiences */
7	Procedure SAMPLE($ \mathcal{B} $):
8	Sampled batch initialization, $\mathcal{B} \leftarrow []$
9	$ b \leftarrow \mathcal{B} / \mathcal{I} $
10	for i in \mathcal{I} do
11	$b \leftarrow \text{Sample}(o, a, r, d, info)$ from $\mathcal{S}[i]$;
12	Add b to \mathcal{B}
13	Shuffle \mathcal{B}
14	return \mathcal{B}

C.2 Ablations

We ablate the FairReplayBuffer by comparing it with trials run with a first-in-first-out (FIFO) buffer. The intention behind designing this replay buffer was to generate a tax schedule with intuitive patterns, achieved by ensuring that the learning agent retains information from less frequently-observed fairness brackets.






Fig. 4. FairReplayBuffer vs. FIFO for the RL setting




(a) C-MAB, FairReplayBuffer




(b) C-MAB, FIFO

Fig. 5. FairReplayBuffer vs. FIFO for the C-MAB setting



(a) MAB, FairReplayBuffer



(b) MAB, FIFO

Fig. 6. FairReplayBuffer vs. FIFO for the MAB setting