```
Algorithm 1: Double DRO policy optimization
Data: Dataset \(\mathcal{D}=\left\{X_{i}, A_{i}, R_{i}\right\}, \pi_{0}, \delta_{c}\) (reward/cost KL radius), \(\delta_{x}\) (context KL radius)
Result: \(\hat{\pi}_{d r o}^{*}, \hat{V}^{*}, \hat{\alpha}_{x}^{*}, \hat{\alpha}_{c}^{*}(n \times k\) values \()\)
// First DRO (rewards) (this step only needs to be done once):
if Binning then
        \(R_{X, A} \leftarrow\) Bin values of \(R_{i}\) by unique \((X, A) \in \mathcal{X} \times \mathcal{A}\) pairs.;
        Solve for \(\hat{Q}_{X, A}\) for every unique \((X, A)\) using Eq. 5 .
else
    // Using function approximation
        \(\{\mathcal{F}\}_{i=1}^{m} \leftarrow\) Fit approximators to evaluate \(\hat{\mathbb{E}}\left[\exp \left(-R / \alpha_{c}\right)\right]\) using the method in Sec. 5 ,
        Solve for \(\hat{Q}_{X_{i}, a}\) for every \(X_{i} \in \mathcal{D}, a \in \mathcal{A}\) using Eq. \(12 / / n \times k\) total \(\hat{Q}_{X_{i}, a}\) values
// Policy Evaluation and Learning
\(\pi \leftarrow\) initial \(\pi\);
while \(\pi\) not converged do
    \(J\left(\alpha_{x}\right)=\alpha_{x} \log \frac{1}{n} \sum_{i} \exp \left(\frac{-\sum_{a \in \mathcal{A}} \pi\left(a \mid X_{i}\right) \hat{Q}_{X_{i}, a}}{\alpha_{x}}\right)+\alpha_{x} \delta_{x} ;\)
    \(\hat{V}_{\pi}=-\min _{\alpha_{x} \geq 0} J\left(\alpha_{x}\right), \hat{\alpha}_{x}^{*}=\operatorname{argmin}_{\alpha_{x} \geq 0} J\left(\alpha_{x}\right) / /\) Policy Evaluation
    Update \((\pi)\) (ex. for differentiable \(\pi\), fix \(\hat{\alpha}_{x}^{*}\) and perform gradient ascent.)
// Note the optimization problem on lines 3, 6, and 10 are all convex
    and can be solved using gradient descent, bisection search, etc.
```


## A Algorithm and Additional Experimental Details

## A. 1 Complete Algorithm

We give the pseudocode of our algorithm in Algorithm 1

## A. 2 Experiment Details

In this section we give additional details for our experiments not included in the main text. Each experiment is run on an internal cluster machine with 4 CPU and 16 GB of memory. A single run of each experiment runs between 10 minutes and 2 hours depending on dataset size.
For all experiments, our $\phi$ vector is defined as:

$$
\begin{equation*}
\phi(X, A)=[X, \text { OneHot }(A)] \tag{14}
\end{equation*}
$$

## A.2.1 Simulated Experiments

As mentioned in the main text, our simulated setting has 4 cgs , where cgs 1 and $3\left(c g_{1,3}\right)$, and cgs 2 and $4,\left(c g_{2,4}\right)$ are the same. In addition to the cgs, there is also 4 other context features, $X_{i}^{f}$ randomly sampled from a uniform distribution. The first 2 features are randomly sampled from Uniform $([0,0.4])$ while the second 2 features are randomly sampled from Uniform $([0.3,0.7])$. There are 4 actions and each $(c g, a)$ pair has a set of true reward generating parameters $\phi_{(1,3), a}^{*}$ and $\phi_{(2,4), a}^{*}$ of length 4 , and the observed reward is $R_{i}=\operatorname{Bernoulli}\left(\phi_{X_{i}, a}^{* T} X_{i}^{f}\right)$. The $\phi^{*}$ 's are:

$$
\phi_{(1,3), a}^{*}=\left\{\begin{array}{ll}
a_{0}: & {[0.5,0.0,0.5,0.0]} \\
a_{1}: & {[1.0,0.0,0.0,0.0]} \\
a_{2}: & {[0.0,0.0,1.0,0.0]} \\
a_{3}: & {[0.1,0.1,0.1,0.1]}
\end{array}, \quad \phi_{(2,4), a}^{*}= \begin{cases}a_{0}: & {[0.5,0.0,0.5,0.0]} \\
a_{1}: & {[0.0,0.0,0.0,1.0]} \\
a_{2}: & {[0.0,1.0,0.0,0.0]} \\
a_{3}: & {[0.1,0.1,0.1,0.1]}\end{cases}\right.
$$

We use a misspecified context for the experiments and the full context used is the $c g$ concatenated with the other features $X_{i}=\left[c g, X_{i}^{f}\right]$.

For the Gaussian experiments, we use the same setup, but instead we first add gaussian noise to the rewards and then clip them to $[0,1]: R_{i}=\operatorname{clip}\left(\phi_{X_{i}, a}^{* T} X_{i}^{f}+\mathcal{N}(0,0.1)\right.$, $\left.\min =0, \max =1\right)$. For the plot in the main text, the training dataset population split is as listed in Table 1 while the test dataset population split between the coarse groups was shifted to $(0.4,0.1,0.4,01)$.

Reported results are averaged over 10 runs where we average over the randomization of the dataset split, the initial policy parameters, and the gradient descent procedure.

## A.2.2 Voting Experiments

The voting dataset of Gerber et al. [2008] is licensed under the Creative Commons Attribution-Non Commercial-No Derivatives 3.0 license. The dataset contains 180002 datapoints, each a voter in a different household across the state of Michigan. The researchers designed one control, and 4 treatment actions that involved mailing the selected individuals a letter ahead of the 2006 Michigan primary election. The actions are:

- Nothing (control): No letter is sent
- Civic: A letter with "Do your civic duty"
- Hawthorne: A letter with "You are being studied"
- Self-History: A letter with the voter's past voting participation record as well as that of the other members of the household. The letter also mentioned a follow-up letter will be sent with the household's updated voting participation after the election
- Neighbors-History: A letter with the voting participation records of the individual, the other members of the individual's household, as well as the neighbors. The letter also mentioned a follow-up letter will be sent after the election with everyone's updated participation, so the individual's participation will be made known among the neighbors.

The data collection policy randomly sampled actions with probability $\frac{5}{9}$ for the control action and $\frac{1}{9}$ for the other actions. We use the binary indicator of voting outcomes as the reward. The neighbors action was the best action for the population, with participants under that action voting at a rate of around $10 \%$ higher than those under the control action. However all actions except the control require effort and cost in the form of printing and mailing a letter. Therefore, for each non-control action, we induce an "effort" cost by randomly flipping $0.09 n_{a}$ of the positive outcomes to 0 , where $n_{a}$ for each action is the number of participants that experienced that action. This lowers the average reward for all non-control action by 0.09 . We choose this cost to induce heterogeneous effects within the dataset.

For the experiment in the main text we take data from one city to use to train a policy deployed on other cities. We choose the train and deployment cities to have a large context shift. For simplicity we approximate a large context shift by choosing one context feature to examine, the 'p2004' feature which is an indicator if the individual voted in the 2004 primary election. We chose the one city with the largest average 'p2004' to be the training city and the 35 cities with the smallest 'p2004' to be the test cities. For computation efficiency and speed, we use a random sample of $25 \%$ of the training data.

We additionally provide a experiment similar to that of the artificial shift in the Scene and Simulated datasets that uses all the data. We set the 'coarse group' using this p2004 features. From the subgroup that did vote in the 2004 primary election, we partition 0.8 of the dataset into the training set and the remaining 0.2 into the test set. Of the subgroup that did not vote, we partition 0.2 of the dataset into the training set and the remaining 0.8 into the test set. This gives the coarse group composition of $(0.75,0.25)$ in the training set and $(0.15,0.85)$ in the test set. While we choose this split, different splits where there is a shift between the training distribution and test distribution also work. We use the features of household size and gender as contexts. We illustrate the more difficult case where some variables, such as the coarse group variable (whether the participants voted in the 2004 primaries), are not available at test time.

We report results averaged across 30 runs.

## A.2.3 Scene Experiments

The Scene classification dataset from the LibSVM library Chang and Lin [2011] is a multiclasssupervised learning dataset. Datsets in the LibSVM library are licensed under the BSD 3-Clause. The
features are processed features of the image and the labels correspond to various landscapes (such as 'mountains', etc). As mentioned in the main text, we use a supervised-to-bandit conversion method where we assume we do not observe the correct labels directly and instead each label is an action and selecting a correct action/label gives positive reward (+1) while a wrong action/label gives no reward (0). We assume we only observe 1 action per datapoint and we randomly sample the actions we observe. We use the true labels to partition the dataset into coarse groups, we place datapoints that contain labels 1,2 , and 3 into one group $\left(c g_{0}\right)$ and the others into another group $\left(c g_{1}\right)$. As in the voting dataset we induce a train test split such that 0.1 of the datapoints in $c g_{0}$ and 0.6 of the datapoints in $c g_{1}$ are in the train dataset with the rest in the test dataset. As in the voting dataset, we assume this feature will not be available at test time.

Due the small dataset size, we use orthonormal random projection, $P \in \mathbb{R}^{d \times d_{p}}$ to project the dataset onto 3 dimensions. Let $d=294$ be the original context feature length and $d_{p}=3$ be the desired projected length. To create our random orthonormal projection matrix of size $\mathbb{R}^{d \times d_{p}}$ we first sample a matrix of size $\mathbb{R}^{d \times d_{p}}$ where each item is sampled from Uniform $([0,1])$. We then take the singular value decomposition and calculate the left and right singular values $U \in \mathbb{R}^{d \times d_{p}}$ and $V \in \mathbb{R}^{d_{p} \times d_{p}}$. We then calculate: $P=U V$.

Our reported results are averaged over 30 runs where each run randomizes across the datapoints in the train-test splits, the initial policy parameters, the gradient descent steps during policy optimization, and the random generated random projection matrix.

## A.2.4 Reward Shift Experiment

Our simple environment for reward shift has two unique contexts ( $X_{0}, X_{1}$ ), with $X_{0}=[0,1]$ and $X_{1}=[1,0]$, and 3 actions. We assume the rewards are binary, but with different support for each ( $X$, a) pair. The reward generation process is defined by a tuple of 3 values $\left(r_{0}, r_{1}, p\right)$ with $r_{1} \geq r_{0}$ and $R_{i}=\left(r_{1}-r_{0}\right)$ Bernoulli $(\mathrm{p})+r_{0}$.

$$
X_{0}=\left\{\begin{array}{ll}
a_{0}: & (0.0,1.0, p) \\
a_{1}: & (0.3,0.5, p) \\
a_{2}: & (0.1,0.15, p)
\end{array}, \quad X_{1}= \begin{cases}a_{0}: & (0.0,1.0, p) \\
a_{1}: & (0.0,1.0, p) \\
a_{2}: & (0.1,0.15, p)\end{cases}\right.
$$

We induce a reward shift by setting $p=0.7$ when generating the training data and $p=0.5$ for the test data.

## A. 3 Additional Results: Different dataset splits and hyperparameters

In Figure 2 we give some results from different dataset splits and hyperparameters. In Figure 2 (a) we give the plot similar to plots (b-e) in Figure 1 of the main text for the Simulated Bernoulli experiment. In these plots, we fix the training and test distribution and examine different values of the $\delta / \delta_{x}$ hyperparameter. We see FDRO can learn more robust policies that work well in the test set under context shift. In Figure 2(b) and (c), we show the result of running different dataset splits. Figure 2(b) considers a different training and test context distributions. It considers a shift of group distribution from $(0.63,0.17)$ in the training set to $(0.3,0.7)$ in the test set. Figure 2 (b) considers a different number of groups. As opposed to 2 coarse groups, we now partition the dataset into 3 groups (with items with labels 0 and 1 in $c g_{0}$, the remaining items with labels 2,3 in $c g_{1}$, and all the remaining items in $c g_{2}$ ). We see in both cases FDRO is able to learn more robust policies that work well on the test context distribution compared to baselines. In Figure 2(d) and (e) we test out different hyperparameters to generate plots similar to that of Figure 1a) of the main text. Like in the main text plot, we fix the $\delta_{x}$ and training distribution and generate test datasets with different shifted group distributions. Also, as in the main text, while we fix $\delta_{x}$ for FDRO, we optimize the uncertainty parameter at each test distribution for the baselines. We plot the performance of each algorithm as the performance gain in the test dataset over baseline-IS. We see that for different $\delta_{x}$ hyperparameters, FDRO can still perform better than baselines when there is shift. Setting a lower value of $\delta_{x}$ leads to a smaller decrease in performance when there is very little shift; however, the benefits at large shifts is also smaller.


Figure 2: In these plots, we demonstrate FDRO can perform well even under different hyperparameter and dataset configurations. The dotted line 'Middle' serves as a constant horizontal reference across the two panels in plots (a-d, g). All plots are plotted with 95\% confidence intervals shaded. Reported results and confidence intervals are averaged over 10 runs in all simulated domains and 30 runs in Scene and Voting. Top Row: (a-c) We fix the train and test distributions and plot the performance at different $\delta / \delta_{x}$ hyperparameters in the Scene domain. The plots show results from different configurations than that in the main text, including different dataset splits (a), different number of groups (b), and different number of features for random projection (c). We see FDRO is able to learn more robust policies that does well in the test distribution in all cases. Middle Row: (d-f) We give results for different configurations in the Simulated Bernoulli dataset. In (d) we give the two panel plot with the parameters described in the main text. In (e,f) We vary the shift in the test distribution and plot results as the performance gain over baseline-IS. Bottom Row: We give a plot for the voting dataset with an artifically induced shift.


Figure 3: We give additional results on hyperparameter sensitivity. The dotted line 'Middle' serves as a constant horizontal reference across the two panels in plots (a) we show sensitivity to $\delta_{c}$ by showing performance Simulation Bernoulli at multiple values of $\delta_{c}$ (b) The extracted worst case population of each $c g$ of our FDRO for different $\delta_{x}$. (c) We plot the reward for each action for the joint formulation (baseline-DRO) at different values of the dual parameter $\alpha$. As described in the text, the order between the actions never changes.

## A. 4 Additional Discussion

We provide some additional points of discussion. In Figure 3a we plot performance at different values of $\delta_{c}$ for the Simulated Bernoulli case and show performance to be similar across parameters.

## A.4.1 Examining the Worst Case Context

The worst case distribution for DRO formulation with KL distance bounds is an softmax scaled weighting of the datapoints Hu and Hong, 2013]. with a temperature determined by optimal dual variable $\alpha^{*}$. If we have discrete or easily discretized context features, we can examine the worst case distribution of that feature. We plot the worst case distributions of the $c g$ context feature of FDRO in Figure 3b in the Bernoulli simulated environment. For each $\delta$ value, the end user can examine the estimated performance on the training distribution as well as the worst-case distribution of some of the features to decide if the policy is reasonable.
Formally, from Hu et al. [2018], the worst case distribution is an exponentially/softmax reweighing of the points:

$$
\begin{equation*}
p_{i}=\frac{\exp \left(\left(-\sum_{a \in \mathcal{A}} \hat{\pi}^{*}\left(a \mid X_{i}\right) \hat{Q}_{X, A}\right) / \alpha_{x}^{*}\right)}{\sum_{i} \exp \left(\left(-\sum_{a \in \mathcal{A}} \hat{\pi}^{*}\left(a \mid X_{i}\right) \hat{Q}_{X, A}\right) / \alpha_{x}^{*}\right)} \tag{15}
\end{equation*}
$$

where $\alpha_{x}^{*}$ is the optimal dual variable of $\hat{\pi}^{*}$ in Equation 7. Consider the $j^{t h}$ feature in the context and assume this feature is either discrete, or easily discretized. We can calculate the worst case distribution over values the feature takes on as:

$$
\begin{equation*}
P(X[j]=x)=\sum p_{i} \mathbb{1}\left[X_{i}[j]=x\right] \tag{16}
\end{equation*}
$$

## A.4.2 Utilizing unavailable features:

In two of the settings - voting and scene, we assumed there were context features in the train set that are unavailable at deployment. This situation occurs when those features are unavailable at deployment, or when they are sensitive features that should not be used for fairness considerations. Even without using it for the policy, FDRO is still able to better utilize this information in the first DRO step when calculating the worst case rewards $\hat{Q}_{X, A}$.

## A.4.3 Additional Discussion: Equivalence of Solutions

As mentioned in the main text, we can compare baseline-DRO and FDRO by examining the KLdivergence chain rule decomposition:

$$
D_{k l}\left(p_{0}(R, X) \| p(R, X)\right)=D_{k l}\left(p_{0}(X) \| p(X)\right)+E_{P_{x_{0}}}\left[D_{k l}\left(p_{0}(R \mid X) \| p(R \mid X)\right)\right]
$$

Recall $D_{k l}\left(p_{0}(R, X) \| p(R, X)\right)=\delta, D_{k l}\left(p_{0}(X) \| p(X)\right)=\delta_{x}, D_{k l}\left(p_{0}(R \mid X) \| p(R \mid X)\right)=\delta_{c}$. This informs us that we have the following relation between the baseline-DRO parameter $\delta$ and the FDRO ones $\delta_{x}$ and $\delta_{c}$ :

$$
\delta=\delta_{x}+\sum_{x} P_{x_{0}} \delta_{c}
$$

From this we can see that if we allow $\delta_{c}$ to vary with context $X$ and action $A$, we can set parameters of $\delta_{x}$ and $\delta_{c}(X, A)$ such that the FDRO will find the same solution as baseline-DRO. However this increase in parameter may be difficult for the algorithm designer. Therefore we expect FRDO to be particularly useful when there are independent shifts in the reward and context distributions, which is a common setting- such as shifting locations (for hospitals or social welfare solutions).

We give a detailed example of what the worst case shift in rewards look like for a single action, and how differs between the factored formulation (FDRO) and the joint formulation (baseline-DRO) Figure 4. In this example, we consider a simple two context, discrete reward setting and examine the worst case distributions for a single action. We compare with setting the single joint parameter to $\delta=0.1$ and setting the context shift parameter of the factored formulation to $\delta_{x}=0.1$, while allowing for very little reward shift $\delta_{c}=0$. We can see the resulting worst case joint distribution is different under the two formulations. We can additionally see for the factored formulation to achieve the equivalent worst case distribution as the joint formulation, we need to set different reward shift radius parameters $\left(\delta_{c}\right)$ for context $X_{0}$ and context $X_{1}$.

## A.4.4 Additional Details: Why Joint DRO May Not Work for Binary Rewards

In the main text we mentioned the formulation from prior work [Si et al., 2020] that jointly considers the context and reward distribution shift may lead to a degenerate result of finding the best policy on the training context distribution under binary reward feedback. This empirical observation is similar to a proven phenomena from the supervised learning literature in distributionally robust optimization [Hu et al. 2018] under continuous inputs. The intuition behind this phenomena is that DRO can be seen as weighing the datapoints based on the incurred loss/or incurred cost, where cost are negative rewards. Under binary reward feedback, the lower reward points for each context action pair are all upweighted equally and we empirically observe this leads to a learned policy equivalent to that optimal for the training context distribution. In this section we provide a simple example and additional details around this observation.

Consider a simplified version of the Simulated Bernoulli setting described above, however in this setting, let $c g$ be the only context feature. Additionally in this case there are only 2 actions $\left\{a_{0}, a_{1}\right\}$. $a_{0}$ achieves medium reward for all contexts (Bernoulli(0.4)), while $a_{1}$ achieves high reward for $c g_{1,3}$ (Bernoulli(0.5)) and low reward for $c g_{2,4}$ (Bernoulli(0.3)). Consider optimizing for the parameters, $\theta$, of a linear contextual bandit policy of the form $\operatorname{argmax}_{a} \phi(x, a)^{T} \theta$ where $\phi(x, a)$ concatenates the context with a one hot encoding of the action. This representation attempts to fit a linear dependence on context when the relationship is very nonlinear. Due to this misspecification, this ends up being equivalent to optimizing a non-contextual bandit which chooses the single best action. Consider the case where $c g_{1,3}$ comprises a large majority ( 0.8 ) of the training context distribution, $P_{x, \text { train }}$. In this case, action $a_{0}$ is optimal for $P_{x, \text { train }}$. If we are certain only the context distribution will shift, as we expand the uncertainty set around $P_{x, \text { train }}$, we expect best worst-case action to eventually shift to the action that is good for both groups, $a_{1}$. However due to binary rewards, we empirically see the joint formulation upweights the lower rewards for each action equally. We demonstrate this in Figure 3(c), where we examine the re-weighted rewards for each action in the joint formulation, which the joint method uses with weighted importance sampling to optimize the policy. We look at the reward values under different values of the dual variable $\alpha$. Notice that for decreasing $\alpha$, while the re-weighted rewards are lower, the order of the benefit of the rewards does not change. Therefore in this example, there is no value of $\delta$ that will result in a policy different from that optimal for the training context distribution. Our method which decomposes context and reward generation shift mitigates this issue.

## B Proofs

Throughout the proofs, we will define a new variable $C$, or costs, for convenience, with $C=-R$ and $C_{i}=-R_{i}$ for every $i$.

## Simple 2 Context, Discrete Reward Case



Figure 4: A detailed example of what the worst case shift in rewards look like for a single action. We choose a simple two context, discrete reward setting. We see that even under the same amount of KL shift $(\delta=0.1)$, the difference in results under the factored approach and the joint formulations are different. We also give the equivalent parameters for FDRO to be equivalent to the solutions of the joint formulation.

## B. 1 The Convexity and Derivatives of the Policy Evaluation steps

We first prove the convexity of the objectives inside the minimization of equations 6 and 7 and give their derivatives as mentioned in the main text. We first give this intermediate result:
Lemma 1 (Convexity of DRO formulation). Consider a set of $n$ values, $\left\{B_{i}\right\}$ for $i=1, \ldots, n$ and function

$$
f(\alpha)=\alpha \log \left(\sum_{i} \frac{1}{n} \exp \left(\frac{B_{i}}{\alpha}\right)\right)+\alpha \delta
$$

Then $f(\alpha)$ is convex in $\alpha$ for $\alpha>0$

Proof.

$$
\begin{aligned}
f(\alpha) & =\alpha \log \left(\sum_{i} \frac{1}{n} \exp \left(\frac{B_{i}}{\alpha}\right)\right)+\alpha \delta \\
& =\alpha \log \left(\sum_{i} \exp \left(\frac{B_{i}}{\alpha}\right)\right)+\alpha \delta+\alpha \log \left(\frac{1}{n}\right)
\end{aligned}
$$

First recall a classic result from convex optimization that states a function $f(x, y)=y g\left(\frac{x}{y}\right)$ is convex for $y>0$ if an only if $g$ is convex. [Boyd et al., 2004]
Note the expression $g(B)=\log \left(\sum_{i} \exp \left(B_{i}\right)\right)$ takes on the form of a log-sum-exp which is convex. Therefore the function $f(B, \alpha)=\alpha g\left(\frac{B}{\alpha}\right)=\alpha \log \sum_{i}\left(\exp \left(\frac{B_{i}}{\alpha}\right)\right)$ is convex, and consequently convex in $\alpha$ when $\alpha>0$.
The linear combination of convex function is still convex, so we have the desired result.
Now we return to proving the convexity of equations 6 and 7 We first reproduce them here with new intermediate variables $\hat{Z}_{X, A}$ and $\hat{J}\left(\pi, \alpha_{x}\right)$ that represent the objective inside the optimization:

$$
\begin{gather*}
\hat{Z}_{X, A}\left(\alpha_{c}\right)=\alpha_{c} \log \left(\frac{1}{m_{x a}} \sum_{R_{j} \in \mathcal{R}_{X, A}}\left[\exp \left(\frac{-R_{j}}{\alpha_{c}}\right)\right]\right)+\alpha_{c} \delta_{c}  \tag{17}\\
\hat{J}\left(\pi, \alpha_{x}\right)=\alpha_{x} \log \sum_{i} \frac{1}{n} \exp \left(\frac{-\sum_{a \in \mathcal{A}} \pi\left(a \mid X_{i}\right) \hat{Q}_{X_{i}, a}}{\alpha_{x}}\right)+\alpha_{x} \delta_{x} \tag{18}
\end{gather*}
$$

With these variables, we have:

$$
\begin{equation*}
\hat{Q}_{X, A}=-\min _{\alpha_{c} \geq 0}\left\{\hat{Z}_{X, A}\left(\alpha_{c}\right)\right\} \quad \text { and } \hat{V}(\pi)=-\min _{\alpha_{x} \geq 0}\left\{\hat{J}\left(\pi, \alpha_{x}\right)\right\} \tag{19}
\end{equation*}
$$

We can directly apply Lemma 1 to $\hat{Z}_{X, A}\left(\alpha_{c}\right)$ and $\hat{J}\left(\pi, \alpha_{x}\right)$ as defined above with $B_{j}=R_{j}$ and $B_{j}=\sum_{a \in \mathcal{A}} \pi\left(a \mid X_{j}\right) \hat{Q}_{X_{j}, a}$ respectively to obtain both $\hat{Z}_{X, A}\left(\alpha_{c}\right)$ and $\hat{J}\left(\pi, \alpha_{x}\right)$ are convex.
We also give the gradient of $\hat{Z}_{X, A}\left(\alpha_{c}\right)$ and $\hat{J}\left(\pi, \alpha_{x}\right)$ :
Remark 1 (Gradient of DRO formulation). The gradient $\frac{\partial f(\alpha)}{\partial \alpha}$ of $f(\alpha)$ given in Lemma $\square$ is:

$$
\frac{\partial f(\alpha)}{\partial \alpha}=-\frac{\sum_{i} B_{i} \exp \left(\frac{B_{i}}{\alpha}\right)}{\alpha \sum_{i} \exp \left(\frac{B_{i}}{\alpha}\right)}+\log \left(\sum_{i} \exp \left(\frac{B_{i}}{\alpha}\right)\right)+\delta
$$

We now prove the main lemmas and theorems from the main text.

## B. 2 Collection of Useful Results From Prior Work

We first collect results from prior work we will use for proofs of Theorems 1, 2 and 3.
We will utilize Hoeffding's inequality:
Lemma 2 (Hoeffding's Inequality). Let $Z_{1}, \ldots Z_{n}$ be independent bounded random variables with $Z_{i} \in[a, b]$ for all $i$, where $-\infty<a \leq b<\infty$. Then:

$$
\begin{equation*}
\mathbb{P}\left(\frac{1}{n} \sum_{i=1}^{n}\left(Z_{i}-\mathbb{E}\left[Z_{i}\right]\right) \geq t\right) \leq \exp \left(-\frac{2 n t^{2}}{(b-a)^{2}}\right) \tag{20}
\end{equation*}
$$

and

$$
\begin{equation*}
\mathbb{P}\left(\frac{1}{n} \sum_{i=1}^{n}\left(Z_{i}-\mathbb{E}\left[Z_{i}\right]\right) \leq t\right) \leq \exp \left(-\frac{2 n t^{2}}{(b-a)^{2}}\right) \tag{21}
\end{equation*}
$$

for all $t \geq 0$

## Distributionally Robust Optimization with KL divergence [Hu and Hong, 2013]

We also use a result from the DRO literature in stochastic optimization from Hu and Hong [2013]. The distributionally robust optimization formulation from stochastic optimization literature considers optimizing the model variables $\theta \in \Theta$ (for example, parameters of a logisitic regression model) to achieve the best worst-case future performance (ex. performance on the test set). The goal is to use the observed random values $\xi$ (ex. (input, label) pairs of the training data) to optimize a cost function $\mathcal{L}(\theta, \xi)$ (ex.classification loss) when evaluated in the unknown test distribution. The test distribution is assumed to be within an uncertainty set, $\mathbb{P}_{\xi_{0}}$. In this work, we focus on ambiguity sets defined as a bound on the Kullback-Leibler (KL) divergence from the data generation distribution $P_{\xi_{0}}$. The DRO optimization problem is:

$$
\theta^{*}=\arg \min _{\theta \in \Theta} \max _{P_{\xi} \in \mathbb{P}_{\xi}} \mathbb{E}_{P_{\xi}}[\mathcal{L}(\theta, \xi)], \text { where } \mathbb{P}_{\xi}=\left\{P_{\xi}: D_{k l}\left(P_{\xi} \| P_{0}\right) \leq \delta\right\}
$$

The inner maximization over $P_{\xi}$ is difficult to solve directly as it is a constrained optimization problem in probability space. However Hu and Hong [2013] show through strong duality that this can be transformed into a much easier unconstrained convex minimization over a scalar:

Lemma 3 (Theorem 1 from Hu and Hong [2013]:). Assume $\mathcal{L}(\theta, \xi)$ takes on a finite set of values or is bounded almost surely. Then through strong duality, the inner maximization is equivalent to solving the following optimization problem:

$$
\max _{P_{\xi} \in \mathbb{P}_{\xi}} \mathbb{E}_{P_{\xi}}[\mathcal{L}(\theta, \xi)]=\min _{\alpha \geq 0} \alpha \log \mathbb{E}_{P_{0}}[\exp (\mathcal{L}(\theta, \xi) / \alpha)]+\alpha \delta
$$

## Uniform Convergence of Stochastic Optimization [Shalev-Shwartz et al., 2009]

We will also make use of the following uniform convergence result from Shalev-Shwartz et al. [2009] of the optimization of convex objectives.
Consider the following stochastic minimization problem:

$$
\begin{equation*}
w^{*}=\underset{w \in \mathcal{W}}{\operatorname{argmin}} F(w) \tag{22}
\end{equation*}
$$

Where $f(w ; Z)$ denote the loss objective and $F(w)=\mathbb{E}_{Z}[f(w, Z)]$. Define the empirical objective and minimizer as $\hat{F}(w)=\frac{1}{n} \sum_{i=1}^{n} f\left(w ; z_{i}\right)$ and $\hat{w}=\operatorname{argmin}_{w} \hat{F}(w)$.
Then the following result guarantees uniform convergence of $F(w)$ over $w$ :
Lemma 4 (Uniform Convergence of Stochastic Convex Optimization (Theorem 5 of Shalev-Shwartz et al. 2009])). Let the following conditions hold:

1. $f(w, z)$ is convex in $w$
2. $\mathcal{W} \subset \mathbb{R}^{d}$ is bounded by $B$

## 3. $f(w, z)$ is L-Lipschitz w.r.t $w$

Then with probability of at least $1-\delta$, for all $w \in \mathcal{W}$ :

$$
\begin{equation*}
|F(w)-\hat{F}(w)| \leq \mathcal{O}\left(L B \sqrt{\frac{d \log (n) \log (d / \delta)}{n}}\right) \tag{23}
\end{equation*}
$$

We also use Si et al. [2020]'s result on the upper bound on $\alpha$ :
Lemma 5 (Lemma 11 of Si et al. 2020]). Let variable B be bounded between 0 and M, then the optimal $\alpha$ in the objective:

$$
\inf _{\alpha \geq 0} \alpha \log \mathbb{E}_{P_{B}}[\exp (B / \alpha)]+\alpha \delta
$$

is upper bounded: $\alpha^{*} \leq M / \delta$.

## B. 3 Proof of Theorem 1

We first restate Theorem 1 .
Theorem 1 (Strong Duality). The optimization problem in equation 2 is equivalent to solving:

$$
\begin{align*}
& V(\pi)=-\min _{\alpha_{x}>0}\left\{\alpha_{x} \log \mathbb{E}_{X \sim P_{x_{0}}}\left[\exp \left(\frac{\mathbb{E}_{\pi}\left[-Q_{X, A}\right]}{\alpha_{x}}\right)+\alpha_{x} \delta_{x}\right]\right\}  \tag{4}\\
& \text { where } Q_{X, A}=-\min _{\alpha_{c}>0}\left\{\alpha_{c} \log \mathbb{E}_{R \sim P_{r_{0} \mid X, A}}\left[\exp \left(\frac{-R}{\alpha_{c}}\right)\right]+\alpha_{c} \delta_{c}\right\} \tag{5}
\end{align*}
$$

The proof of Theorem 1 directly follows from applying Lemma 3 twice. Notice that because rewards, R , are bounded $\left[0, R_{M}\right]$, the assumption of Lemma 3 is satisfied for $Q_{X, A}$. Similarly, $Q_{X, A}$ are also bounded $\left[0, R_{M}\right]$ so the assumption is also satisfied for $V(\pi)$.

## B. 4 Policy Evaluation and Learning convergence

We first prove the probability of convergence for $\hat{Q}_{X, A}$ for each $(X, A)$ pair separately. We then apply the union bound to bound the probability they hold jointly for all $(X, A)$ pairs. We then use this result along with local Lipschitz continuity to prove the rest of the results.

We start with proving PAC uniform convergence for $\hat{Q}_{X, A}$ :
Lemma 6 (PAC convergence of $Q_{X, A}$ (first DRO layer)). Given a $(X, A)$ pair, when $n \geq \frac{\left(p_{0}(x, a)\right)^{2}}{2 \log (2 / \delta)}$, the following result holds with probability at least $1-\delta$ and all $\alpha_{c} \in\left[\underline{\alpha}_{c} / 2,2 \bar{\alpha}_{c}\right]$ :

$$
\begin{equation*}
\left|Q_{X, A}-\hat{Q}_{X, A}\right| \leq \mathcal{O}\left(c \sqrt{\frac{\log (n) \log (2 / \delta)}{n}}\right) \tag{24}
\end{equation*}
$$

where

$$
\begin{equation*}
c=\frac{32\left(R_{\max }\right)^{3}}{\delta_{c}^{2} \underline{\alpha}_{c}^{2} \exp \left(-R_{\max } / \underline{\alpha}_{c}\right) \underline{p}_{0}(x, a)} \tag{25}
\end{equation*}
$$

and $\underline{p}_{0}(x, a)$ is the lowest probability $(X, A)$ pair: $\underline{p}_{0}(x, a)=\min _{(x, a)} p_{0}(x, a)$, and $\underline{\alpha}_{c}$ is the minimum value for $\alpha_{c}$ and satisfies $\underline{\alpha}_{c}>0$.

Proof We first rewrite the objective in terms the total number of samples, $n$, across all state, action pairs, as opposed to the number of samples $m_{x a}$ for the $(X, A)$ pair:

$$
\begin{align*}
\hat{Q}_{X, A} & =-\min _{\alpha_{c}>0}\left\{\alpha_{c} \log \left(\frac{1}{m_{x a}} \sum_{R_{j} \in \mathcal{R}_{X, A}}\left[\exp \left(\frac{-R_{j}}{\alpha_{c}}\right)\right]\right)+\alpha_{c} \delta_{c}\right\}  \tag{26}\\
& =-\min _{\alpha_{c}>0}\left\{\alpha_{c} \log \left(\frac{1}{p_{0}(x, a) n} \sum_{i=1}^{n}\left[\mathbb{1}_{x a} \exp \left(\frac{-R_{i}}{\alpha_{c}}\right)\right]\right)+\alpha_{c} \delta_{c}\right\} \tag{27}
\end{align*}
$$

Where $\mathbb{1}_{x a}$ is the indicator function that is 1 if both $X_{i}=x$ and $A_{i}=a$, and 0 otherwise. We use the fact $\sum_{R_{j} \in \mathcal{R}_{X, A}}\left[\exp \left(\frac{-R_{j}}{\alpha_{c}}\right)\right]=\sum_{i=1}^{n}\left[\mathbb{1}_{x a} \exp \left(\frac{-R_{i}}{\alpha_{c}}\right)\right]$
Consider the value of $Q_{X, A}$ for a single $(X, A)$ pair. Consider the function $F_{x a}\left(\alpha_{c}\right)=$ $\mathbb{E}\left[\mathbb{1}_{x a} \exp \left(C_{i} / \alpha_{c}\right)\right]$ with $f_{x a}\left(\alpha_{c}, C_{i}\right)=\mathbb{1}_{x a} \exp \left(C_{i} / \alpha_{c}\right)$. Then $f_{x a}\left(\alpha_{c}, C_{i}\right)$ is upper bounded by 1 and lower bounded by 0 . Additionally note that $f_{x a}\left(\alpha_{c}, C_{i}\right)$ is locally L-Lipschitz with respect to $\alpha_{c}$ on $\left[\underline{\alpha}_{c} / 2,2 \bar{\alpha}_{c}\right]$ :

$$
\begin{align*}
\frac{\partial f_{x a}\left(\alpha_{c}\right)}{\alpha_{c}} & =-\frac{\mathbb{1}_{x a} C_{i} \exp \left(C_{i} / \alpha_{c}\right)}{\alpha_{c}^{2}}  \tag{28}\\
\left|\frac{\partial f_{x a}\left(\alpha_{c}\right)}{\alpha_{c}}\right| & \leq \frac{4 R_{\max }}{\underline{\alpha}_{c}^{2}} \tag{29}
\end{align*}
$$

Additionally, from Lemma 5 , we have $\bar{\alpha}_{c}=R_{\max } / \delta_{c}$. Therefore, we can directly apply Lemma 4 to $F_{x a}\left(\alpha_{c}\right)$ and $\hat{F}_{x a}=\frac{1}{n} \sum_{i=1}^{n} f_{x a}\left(\alpha_{c}, C_{i}\right)$ with $d=1, B=2 \bar{\alpha}_{c}=2 R_{\max } / \delta_{c}$ and $L=4 R_{\max } / \underline{\alpha}_{c}^{2}$ and conclude uniform convergence in $\alpha_{c}$ on the range $\left[\underline{\alpha}_{c} / 2,2 \overline{\alpha_{c}}\right]$ : with probability $1-\delta^{\prime}$ :

$$
\begin{equation*}
\sup _{\alpha_{c} \in\left[\underline{\alpha}_{c} / 2,2 \overline{\alpha_{c}}\right]}\left|F_{x a}\left(\alpha_{c}\right)-\hat{F}_{x a}\left(\alpha_{c}\right)\right| \leq \mathcal{O}\left(\frac{8\left(R_{\max }\right)^{2}}{\delta_{c} \underline{\alpha}_{c}^{2}} \sqrt{\frac{\log (n) \log \left(1 / \delta^{\prime}\right)}{n}}\right) \tag{30}
\end{equation*}
$$

Now define, $G_{x a}\left(\alpha_{c}\right)=\alpha_{c} \log \left(\frac{F_{x a}\left(\alpha_{c}\right)}{p_{0}(x, a)}\right)+\alpha_{c} \delta_{c}$. Note that $Q_{X, A}=\min _{\alpha_{c} \geq 0} G_{x a}\left(\alpha_{c}\right)$

$$
\begin{align*}
\sup _{\alpha_{c}}\left|G_{x a}\left(\alpha_{c}\right)-\hat{G}_{x a}\left(\alpha_{c}\right)\right| & =\sup _{\alpha_{c}}\left|\alpha_{c}\left(\log \left(F\left(\alpha_{c}\right)\right)-\log \left(\hat{F}\left(\alpha_{c}\right)\right)\right)\right|  \tag{31}\\
& \leq 2 \bar{\alpha}_{c} L \sup _{\alpha_{c}}\left|F\left(\alpha_{c}\right)-\hat{F}\left(\alpha_{c}\right)\right|  \tag{32}\\
& \leq \mathcal{O}\left(2 \bar{\alpha}_{c} \frac{8\left(R_{\max }\right)^{2} L}{\delta_{c} \underline{\alpha}_{c}^{2}} \sqrt{\frac{\log (n) \log (\delta / 2)}{n}}\right) \tag{33}
\end{align*}
$$

Where the second line utilizes the fact that $\log$ is locally L-Lipschitz and:

$$
L=\frac{1}{\min \left\{\inf _{\alpha_{c}} \hat{F}\left(\alpha_{c}\right), \inf _{\alpha_{c}} F\left(\alpha_{c}\right)\right\}}
$$

We now derive the quantity $L$. First note the expression $\exp \left(C_{i} / \alpha_{c}\right) \geq \exp \left(-R_{\max } / \underline{\alpha}_{c}\right)$. With this we can write:

$$
\begin{align*}
F_{x a}\left(\alpha_{c}\right) & =p_{0}(x, a) \mathbb{E}\left[\exp \left(C_{i} / \alpha_{c}\right) \mid X_{i}=x, A_{i}=a\right]  \tag{34}\\
& \geq p_{0}(x, a) \exp \left(-R_{\max } / \underline{\alpha}_{c}\right)  \tag{35}\\
& \text { and }  \tag{36}\\
\hat{F}_{x a}\left(\alpha_{c}\right) & =\frac{1}{n} \sum_{i=1}^{n} \mathbb{1}_{x a} \exp \left(C_{i} / \alpha_{c}\right)  \tag{37}\\
& \geq \exp \left(-R_{\max } / \underline{\alpha}_{c}\right) \frac{1}{n} \sum_{i=1}^{n} \mathbb{1}_{x a}  \tag{38}\\
& =\exp \left(-R_{\max } / \underline{\alpha}_{c}\right) \hat{p}_{0}(x, a) \tag{39}
\end{align*}
$$

Where $\hat{p}_{0}(x, a)=\frac{1}{n} \sum_{i=1}^{n} \mathbb{1}_{x a}$ is the empirical probability of observing $(X, A)$ in the dataset. Now we need to lower bound $\hat{p}_{0}(x, a)$. Let us lower bound the probability it is less than half of the true
value $p_{0}(x, a) / 2$ :

$$
\begin{align*}
P\left(\hat{p}_{0}(x, a) \leq \frac{p_{0}(x, a)}{2}\right) & =P\left(-\hat{p}_{0}(x, a) \geq-\frac{p_{0}(x, a)}{2}\right)  \tag{40}\\
& =P\left(p_{0}(x, a)-\hat{p}_{0}(x, a) \geq p_{0}(x, a)-\frac{p_{0}(x, a)}{2}\right)  \tag{41}\\
& =P\left(p_{0}(x, a)-\hat{p}_{0}(x, a) \geq \frac{p_{0}(x, a)}{2}\right)  \tag{42}\\
& \stackrel{(1)}{\leq} \exp \left(-2 n\left(\frac{p_{0}(x, a)}{2}\right)^{2}\right)  \tag{43}\\
& =\exp \left(-\frac{n p_{0}(x, a)^{2}}{2}\right) \tag{44}
\end{align*}
$$

Where step (1) applies Hoeffding inequality on the binary random variables $\mathbb{1}_{x a}$ and with $t=\frac{p_{0}(x, a)}{2}$ We can now calculate the minimum number of timesteps such that $P\left(\hat{p}_{0}(x, a) \leq \frac{p_{0}(x, a)}{2}\right) \leq \delta^{\prime}$ :

$$
\begin{equation*}
\exp \left(-\frac{p_{0}(x, a)^{2}}{2 n}\right) \leq \delta^{\prime} \xrightarrow[\text { Solve for } \mathrm{n}]{ } n \geq \frac{2 \log \left(1 / \delta^{\prime}\right)}{p_{0}(x, a)^{2}} \tag{45}
\end{equation*}
$$

Therefore for $n \geq \frac{2 \log \left(1 / \delta^{\prime}\right)}{p_{0}(x, a)^{2}}$, with probability $1-\delta$ we have $\hat{p}_{0}(x, a) \geq \frac{p_{0}(x, a)}{2}$. We therefore have:

$$
\begin{equation*}
L=\frac{2}{\exp \left(\frac{-R_{\max }}{\underline{\alpha}_{c}}\right) p_{0}(x, a)} \geq \frac{2}{\exp \left(\frac{-R_{\max }}{\underline{\alpha}_{c}}\right) \underline{p}_{0}(x, a)} \tag{46}
\end{equation*}
$$

We can now put everything together for the constant from equation 33

$$
\begin{equation*}
c=2 \bar{\alpha}_{c} \frac{8\left(R_{\max }\right)^{2} L}{\delta_{c} \underline{\alpha}_{c}^{2}}=\frac{32\left(R_{\max }\right)^{3}}{\delta_{c}^{2} \underline{\alpha}_{c}^{2} \exp \left(\frac{-R_{\max }}{\underline{\underline{\alpha}}_{c}}\right) \underline{p}_{0}(x, a)} \tag{47}
\end{equation*}
$$

Let $\hat{\alpha}_{c}=\operatorname{argmin}_{\alpha_{c}} \hat{G}_{x a}\left(\alpha_{c}\right)$. Note that $\hat{Q}_{X, A}=\hat{G}_{x a}\left(\hat{\alpha}_{c}\right)$. Let $\alpha_{c}^{*}=\operatorname{argmin}_{\alpha_{c}} G_{x a}\left(\alpha_{c}\right)$. Note that $Q_{X, A}=G_{x a}\left(\alpha_{c}^{*}\right)$.
It then follows:

$$
\begin{align*}
G_{x a}\left(\hat{\alpha}_{c}\right)-\hat{G}_{x a}\left(\hat{\alpha}_{c}\right) & \leq \mathcal{O}\left(c \sqrt{\frac{\log (n) \log \left(1 / \delta^{\prime}\right)}{n}}\right)  \tag{48}\\
G_{x a}\left(\hat{\alpha}_{c}\right)-\hat{G}_{x a}\left(\hat{\alpha}_{c}\right)+G_{x a}\left(\alpha_{c}^{*}\right) & \leq G_{x a}\left(\alpha_{c}^{*}\right)+\mathcal{O}\left(c \sqrt{\frac{\log (n) \log \left(1 / \delta^{\prime}\right)}{n}}\right)  \tag{49}\\
G_{x a}\left(\alpha_{c}^{*}\right)-\hat{G}_{x a}\left(\hat{\alpha}_{c}\right) & \leq G_{x a}\left(\alpha_{c}^{*}\right)-G_{x a}\left(\hat{\alpha}_{c}\right)+\mathcal{O}\left(c \sqrt{\frac{\log (n) \log \left(1 / \delta^{\prime}\right)}{n}}\right)  \tag{50}\\
Q_{X, A}-\hat{Q}_{X, A} & \leq \underbrace{G_{x a}\left(\alpha_{c}^{*}\right)-G_{x a}\left(\hat{\alpha}_{c}\right)}_{\leq 0}+\mathcal{O}\left(c \sqrt{\frac{\log (n) \log \left(1 / \delta^{\prime}\right)}{n}}\right)  \tag{51}\\
Q_{X, A}-\hat{Q}_{X, A} & \leq \mathcal{O}\left(c \sqrt{\frac{\log (n) \log \left(1 / \delta^{\prime}\right)}{n}}\right) \tag{52}
\end{align*}
$$

Similarly we can start with $G_{x a}\left(\alpha_{c}^{*}\right)-\hat{G}_{x a}\left(\alpha_{c}^{*}\right) \leq \mathcal{O}\left(c \sqrt{\frac{\log (n) \log \left(1 / \delta^{\prime}\right)}{n}}\right)$ and prove the other side $\hat{Q}_{X, A}-Q_{X, A} \leq \mathcal{O}\left(c \sqrt{\frac{\log (n) \log \left(1 / \delta^{\prime}\right)}{n}}\right)$.

By setting $\delta^{\prime}=\delta / 2$, and taking the union bound over the two sides, the desired lemma result follows.

Note about $\underline{\alpha}_{c}$ : In this work we only consider cases where $\alpha_{c}>0 . \alpha_{c}$ is strictly bounded away from 0 by some $\underline{\alpha}_{c}>0$. This occurs either when rewards have continuous support, or, in the case of discrete reward support, when $\delta_{c}$ is small enough such that the uncertainty set does not include the distribution with all mass at a single support value.
Since we are considering the case of discrete, finite context and action sets, we can now take a union bound over all $(X, A)$ pairs and achieve joint uniform convergence over all $(X, A)$ pairs to arrive at the following result:

Lemma 7 (Uniform convergence over all $(X, A)$ ). When $n \geq \frac{(2 \log (2|\mathcal{X}| \mathcal{A} \mid / \delta)}{\left(\underline{p}_{0}(x, a)\right)^{2}}$, the following holds with probability of at least $1-\delta$ :

$$
\begin{equation*}
\max _{(X, A)}\left|Q_{X, A}-\hat{Q}_{X, A}\right| \leq \mathcal{O}\left(c \sqrt{\frac{\log (n) \log (2|\mathcal{X}||\mathcal{A}| / \delta)}{n}}\right) \tag{53}
\end{equation*}
$$

Where

$$
\begin{equation*}
c=\frac{32\left(R_{\max }\right)^{3}}{\delta_{c}^{2} \underline{\alpha}_{c}^{2} \exp \left(\frac{-R_{\max }}{\underline{\underline{\alpha}}_{c}}\right) \underline{p}_{0}(x, a)} \tag{54}
\end{equation*}
$$

With this, we continue the proof, recall our goal is to optimize the distributionally robust objective:

$$
\begin{equation*}
\hat{V}=\min _{\pi, \alpha_{x} \geq 0}\left\{\hat{J}\left(\pi, \alpha_{x}\right)\right\} \quad \text { and } \quad V(\pi)=\min _{\pi, \alpha_{x} \geq 0}\left\{J\left(\pi, \alpha_{x}\right)\right\} \tag{55}
\end{equation*}
$$

Where

$$
\begin{align*}
& \hat{J}\left(\pi, \alpha_{x}\right)=\alpha_{x} \log \mathbb{E}\left[\exp \left(\frac{-\hat{Q}_{X, \pi}^{a v g}}{\alpha_{x}}\right)\right]+\alpha_{x} \delta_{x}  \tag{56}\\
& J\left(\pi, \alpha_{x}\right)=\alpha_{x} \log \mathbb{E}\left[\exp \left(\frac{-Q_{X, \pi}^{a v g}}{\alpha_{x}}\right)\right]+\alpha_{x} \delta_{x} \tag{57}
\end{align*}
$$

And for a fixed policy $\pi$ and a context $X$

$$
\begin{equation*}
\hat{Q}_{X, \pi}^{a v g}=\sum_{a} \pi(a \mid X) \hat{Q}_{X, a} \quad \text { and } \quad Q_{X, \pi}^{a v g}=\sum_{a} \pi(a \mid X) Q_{X, a} \tag{58}
\end{equation*}
$$

We will next show the following result which gives uniform convergence of $J\left(\pi, \alpha_{x}\right)$ in $\pi$ and $\alpha_{x}$.
Lemma 8 (Uniform convergence of optimization objective). With probability of at least $1-\delta$ and with $n \geq \frac{(2 \log (2|\mathcal{X}| \mathcal{A} \mid / \delta)}{\left(\underline{p}_{0}(x, a)\right)^{2}}$ :

$$
\begin{equation*}
\sup _{\pi, \alpha_{x}}\left|\hat{J}\left(\pi, \alpha_{x}\right)-J\left(\pi, \alpha_{x}\right)\right| \leq \mathcal{O}\left(c \sqrt{\frac{\log (n) \log \left(\frac{2|\mathcal{X}||\mathcal{A}|}{\delta}\right)}{n}}\right) \tag{59}
\end{equation*}
$$

Proof: For all steps, assume $n \geq \frac{(2 \log (2|\mathcal{X}| \mathcal{A} \mid / \delta)}{\left(\underline{p}_{0}(x, a)\right)^{2}}$. Because $\sum_{a} \pi(a \mid X)=1, Q_{X, \pi}^{a v g}$ is a weighted average of $Q_{X, A}$.

$$
\begin{equation*}
\left|\hat{Q}_{X, \pi}^{a v g}-Q_{X, \pi}^{a v g}\right| \leq \max \left|Q_{X, A}-\hat{Q}_{X, A}\right| \leq \mathcal{O}\left(c \sqrt{\frac{\log (n) \log \left(\frac{2|\mathcal{X}||\mathcal{A}|}{\delta}\right)}{n}}\right) \tag{60}
\end{equation*}
$$

Additionally because Lemma 7 holds for simultaneously all $(X, A)$ pairs, and every $\hat{Q}_{i, \pi}^{a v g}$ depends on the same set of $Q_{X, A}$, we have the above result holds with probability with $1-\delta$ for all $i$. Formally, with probability $1-\delta$ :

$$
\begin{equation*}
\max _{X}\left|\hat{Q}_{X, \pi}^{a v g}-Q_{X, \pi}^{a v g}\right| \leq \max \left|Q_{X, A}-\hat{Q}_{X, A}\right| \leq \mathcal{O}\left(c \sqrt{\frac{\log (n) \log \left(\frac{2|\mathcal{X} \| \mathcal{A}|}{\delta}\right)}{n}}\right) \tag{61}
\end{equation*}
$$

Then we have:

$$
\begin{align*}
\sup \left|\hat{J}\left(\pi, \alpha_{x}\right)-J\left(\pi, \alpha_{x}\right)\right| & =\alpha_{x} \sup \left|\log \mathbb{E}\left[\exp \left(\frac{-\hat{Q}_{X, \pi}^{a v g}}{\alpha_{x}}\right)\right]-\log \mathbb{E}\left[\exp \left(\frac{-Q_{X, \pi}^{a v g}}{\alpha_{x}}\right)\right]\right|  \tag{62}\\
& =\alpha_{x} \sup \left|\left(\log \mathbb{E}\left[\frac{\exp \left(\frac{-\hat{Q}_{i, \pi}^{a v g}}{\alpha_{x}}\right)}{\exp \left(\frac{-Q_{i, \pi}^{a v g}}{\alpha_{x}}\right)}\right]\right)\right|  \tag{63}\\
& =\alpha_{x} \sup \left|\left(\log \mathbb{E}\left[\exp \left(\frac{\left.Q_{i, \pi}^{a v g}-\hat{Q}_{i, \pi}^{a v g}\right)}{\alpha_{x}}\right)\right]\right)\right|  \tag{64}\\
& \leq \alpha_{x}\left(\log \mathbb{E}\left[\exp \left(\frac{\sup _{i}\left|Q_{i, \pi}^{a v g}-\hat{Q}_{i, \pi}^{a v g}\right|}{\alpha_{x}}\right)\right]\right)  \tag{65}\\
& \leq \alpha_{x}\left(\operatorname { l o g } \left(\exp \left(\frac{\left.\left.\mathcal{O}\left(c \sqrt{\left.\frac{\log (n) \log \left(\frac{2 \mid \mathcal{X I | I \mathcal { A } |}}{\delta}\right)}{n}\right)} \alpha_{x}\right)\right)\right)}{n}\right)\right.\right.  \tag{66}\\
& =\mathcal{O}\left(c \sqrt{\left.\frac{\log (n) \log \left(\frac{2|\mathcal{X} \||\mathcal{A}|}{\delta}\right)}{n}\right)}\right) \tag{67}
\end{align*}
$$

Which is the desired result.
We can now show convergence in policy evaluation. We first rewrite the objectives, $\hat{V}(\pi)$ and $\hat{V}(\pi)$, using $J\left(\pi, \alpha_{x}\right)$ and $\hat{J}\left(\pi, \alpha_{x}\right)$. Given a fixed policy $\pi$ :

$$
\begin{equation*}
\hat{V}(\pi)=\min _{\alpha_{x} \geq 0}\left\{\hat{J}\left(\pi, \alpha_{x}\right)\right\} \quad \text { and } \quad V(\pi)=\min _{\alpha_{x} \geq 0}\left\{J\left(\pi, \alpha_{x}\right)\right\} \tag{68}
\end{equation*}
$$

We now restate the theorem:
Theorem 2 (Convergence of policy evaluation). For $n \geq \frac{(2 \log (2|\mathcal{X}| \mathcal{A} \mid / \delta)}{\left(\underline{p}_{0}(x, a)\right)^{2}}$, the following holds for any $\pi$ with probability of at least $1-\delta$ :

$$
\begin{equation*}
|V(\pi)-\hat{V}(\pi)| \leq \mathcal{O}\left(c \sqrt{\frac{\log (n) \log \left(\frac{2|\mathcal{X} \| \mathcal{A}|}{\delta}\right)}{n}}\right) \quad \text { where } c=\frac{32\left(R_{\max }\right)^{3}}{\delta_{c}^{2} \underline{\alpha}_{c}^{2} \exp \left(\frac{-R_{\max }}{\underline{\alpha}_{c}}\right) \underline{p}_{0}(x, a)} \tag{8}
\end{equation*}
$$

and $\underline{p}_{0}(x, a)=\min _{(x, a)} p_{0}(x, a) \geq \epsilon_{\pi} \epsilon_{X}$ is the minimum over probability of occurrence of $(X, A)$ pairs, and $\underline{\alpha}_{c}$ is the minimum value for $\alpha_{c}$.

Proof For all steps, assume $n \geq \frac{(2 \log (2|\mathcal{X}| \mathcal{A} \mid / \delta)}{\left(\underline{p}_{0}(x, a)\right)^{2}}$. Define:

$$
\begin{equation*}
\hat{\alpha}_{x}=\underset{\alpha_{x} \geq 0}{\operatorname{argmin}}\left\{\hat{J}\left(\pi, \alpha_{x}\right)\right\} \quad \text { and } \quad \alpha_{x}^{*}=\underset{\alpha_{x} \geq 0}{\operatorname{argmin}}\left\{J\left(\pi, \alpha_{x}\right)\right\} \tag{69}
\end{equation*}
$$

It then follows:

$$
\left.\begin{array}{rl}
J\left(\pi, \hat{\alpha}_{x}\right)-\hat{J}\left(\pi, \hat{\alpha}_{x}\right) & \leq \mathcal{O}\left(c \sqrt{\left.\frac{\log (n) \log \left(\frac{2|\mathcal{X}||\mathcal{A}|}{\delta}\right)}{n}\right)}\right. \\
J\left(\pi, \hat{\alpha}_{x}\right)-\hat{J}\left(\pi, \hat{\alpha}_{x}\right)+J\left(\pi, \alpha_{x}^{*}\right) & \leq J\left(\pi, \alpha_{x}^{*}\right)+\mathcal{O}\left(c \sqrt{\left.\frac{\log (n) \log \left(\frac{2|\mathcal{X}||\mathcal{A}|}{\delta}\right)}{n}\right)}\right) \\
J\left(\pi, \alpha_{x}^{*}\right)-\hat{J}\left(\pi, \hat{\alpha}_{x}\right) & \leq J\left(\pi, \alpha_{x}^{*}\right)-J\left(\pi, \hat{\alpha}_{x}\right)+\mathcal{O}\left(c \sqrt{\frac{\log (n) \log \left(\frac{2|\mathcal{X} \| \mathcal{A}|}{\delta}\right)}{n}}\right) \\
V(\pi)-\hat{V}(\pi) & \leq \underbrace{J\left(\pi, \alpha_{x}^{*}\right)-J\left(\pi, \hat{\alpha}_{x}\right)+\mathcal{O}\left(c \sqrt{\frac{\log (n) \log \left(\frac{2|\mathcal{X}||\mathcal{A}|}{\delta}\right)}{n}}\right)} \\
& \leq 0 \text { by definition of } \alpha_{x}^{*} \tag{74}
\end{array}\right)
$$

We can also show the other side:

$$
\begin{align*}
& \hat{J}\left(\pi, \alpha *_{x}\right)-J\left(\pi, \alpha_{x}^{*}\right) \leq \mathcal{O}\left(c \sqrt{\frac{\log (n) \log \left(\frac{2|\mathcal{X}||\mathcal{A}|}{\delta}\right)}{n}}\right)  \tag{75}\\
& \hat{J}\left(\pi, \alpha *_{x}\right)-J\left(\pi, \alpha_{x}^{*}\right)+\hat{J}\left(\pi, \hat{\alpha}_{x}\right) \leq \hat{J}\left(\pi, \hat{\alpha}_{x}\right)+\mathcal{O}\left(c \sqrt{\frac{\log (n) \log \left(\frac{2|\mathcal{X}||\mathcal{A}|}{\delta}\right)}{n}}\right)  \tag{76}\\
& \hat{J}\left(\pi, \hat{\alpha}_{x}\right)-J\left(\pi, \alpha_{x}^{*}\right) \leq \hat{J}\left(\pi, \hat{\alpha}_{x}\right)-\hat{J}\left(\pi, \alpha *_{x}\right)+\mathcal{O}\left(c \sqrt{\frac{\log (n) \log \left(\frac{2|\mathcal{X}||\mathcal{A}|}{\delta}\right)}{n}}\right)  \tag{77}\\
& \hat{V}(\pi)-V(\pi) \leq \underbrace{\hat{J}\left(\pi, \hat{\alpha}_{x}\right)-\hat{J}\left(\pi, \alpha *_{x}\right)}_{\leq 0 \text { by definition of } \hat{\alpha}_{x}}+\mathcal{O}\left(c \sqrt{\frac{\log (n) \log \left(\frac{2|\mathcal{X}||\mathcal{A}|}{\delta}\right)}{n}}\right)  \tag{78}\\
& \hat{V}(\pi)-V(\pi) \leq \mathcal{O}\left(c \sqrt{\frac{\log (n) \log \left(\frac{2|\mathcal{X}||\mathcal{A}|}{\delta}\right)}{n}}\right) \tag{79}
\end{align*}
$$

Combining these two shows the result gives the desired result.
Now lastly consider the case of policy learning. Define

$$
\begin{equation*}
V^{*}=\min _{\pi, \alpha_{x}} J\left(\pi, \alpha_{x}\right) \quad \text { and } \quad \hat{V}^{*}=\min _{\pi, \alpha_{x}} \hat{J}\left(\pi, \alpha_{x}\right) \tag{80}
\end{equation*}
$$

Additionally define optimal parameters:

$$
\begin{equation*}
\left(\pi^{*}, \alpha_{x}^{*}\right)=\underset{\pi, \alpha_{x}}{\operatorname{argmin}} J\left(\pi, \alpha_{x}\right) \quad \text { and } \quad\left(\hat{\pi}, \hat{\alpha}_{x}\right)=\underset{\pi, \alpha_{x}}{\operatorname{argmin}} \hat{J}\left(\pi, \alpha_{x}\right) \tag{81}
\end{equation*}
$$

Finally, recall the the set of $\epsilon$-optimal policies $\Pi_{\epsilon}^{*}=\left\{\pi: V(\pi) \geq V^{*}-\epsilon\right\}$.
We first restate the theorem:

Theorem 3 (Convergence of Policy Learning). For $n \geq \frac{(2 \log (2|\mathcal{X}| \mathcal{A} \mid / \delta)}{\left(\underline{p}_{0}(x, a)\right)^{2}}$, the following holds with probability of at least $1-\delta$ :

$$
\begin{equation*}
\left|\hat{V}^{*}-V^{*}\right| \leq \mathcal{O}\left(c \sqrt{\frac{\log (n) \log \left(\frac{2|\mathcal{X}||\mathcal{A}|}{\delta}\right)}{n}}\right) \tag{9}
\end{equation*}
$$

Equivalently this says

$$
\begin{equation*}
P\left(\hat{\pi} \in \Pi_{\epsilon}^{*}\right)>1-\mathcal{O}\left(2|\mathcal{X} \| \mathcal{A}| \exp \left(-\frac{n \epsilon^{2}}{\log (n) c^{2}}\right)\right) \tag{10}
\end{equation*}
$$

Proof: This proof is similar to the proof of policy evaluation of the previous theorem. The first side of the inequality

$$
\begin{align*}
& J\left(\hat{\pi}, \hat{\alpha}_{x}\right)-\hat{J}\left(\hat{\pi}, \hat{\alpha}_{x}\right) \leq \mathcal{O}\left(c \sqrt{\left.\frac{\log (n) \log \left(\frac{2|\mathcal{X}||\mathcal{A}|}{\delta}\right)}{n}\right)}\right.  \tag{82}\\
& J\left(\hat{\pi}, \hat{\alpha}_{x}\right)-\hat{J}\left(\hat{\pi}, \hat{\alpha}_{x}\right)+J\left(\pi^{*}, \alpha_{x}^{*}\right) \leq J\left(\pi *^{\prime} \alpha_{x}^{*}\right)+\mathcal{O}\left(c \sqrt{\left.\frac{\log (n) \log \left(\frac{2|\mathcal{X}||\mathcal{A}|}{\delta}\right)}{n}\right)}\right.  \tag{83}\\
& J\left(\pi^{*}, \alpha_{x}^{*}\right)-\hat{J}\left(\hat{\pi}, \hat{\alpha}_{x}\right) \leq J\left(\pi^{*}, \alpha_{x}^{*}\right)-J\left(\hat{\pi}, \hat{\alpha}_{x}\right)+\mathcal{O}\left(c \sqrt{\frac{\log (n) \log \left(\frac{2|\mathcal{X}||\mathcal{A}|}{\delta}\right)}{n}}\right)  \tag{84}\\
& V^{*}-\hat{V}^{*} \leq \underbrace{J\left(\pi^{*}, \alpha_{x}^{*}\right)-J\left(\hat{\pi}, \hat{\alpha}_{x}\right)+\mathcal{O}} c \sqrt{\left.\frac{\log (n) \log \left(\frac{2|\mathcal{X} \||\mathcal{A}|}{\delta}\right)}{n}\right)}  \tag{85}\\
& V^{*}-\hat{V}^{*} \leq \mathcal{O}\left(c \sqrt{\left.\frac{\log (n) \log \left(\frac{2 \mid \mathcal{X} \| \operatorname{lition} \text { of }\left(\pi^{*}, \alpha_{x}^{*}\right)}{\delta}\right)}{n}\right)}\right. \tag{86}
\end{align*}
$$

The other side of the inequality:

$$
\begin{align*}
& \hat{J}\left(\pi^{*}, \alpha *_{x}\right)-J\left(\pi^{*}, \alpha_{x}^{*}\right) \leq \mathcal{O}\left(c \sqrt{\left.\frac{\log (n) \log \left(\frac{2|\mathcal{X}||\mathcal{A}|}{\delta}\right)}{n}\right)}\right.  \tag{87}\\
& \hat{J}\left(\pi^{*}, \alpha *_{x}\right)-J\left(\pi^{*}, \alpha_{x}^{*}\right)+\hat{J}\left(\hat{\pi}, \hat{\alpha}_{x}\right) \leq \hat{J}\left(\hat{\pi}, \hat{\alpha}_{x}\right)+\mathcal{O}\left(c \sqrt{\left.\frac{\log (n) \log \left(\frac{2|\mathcal{X}||\mathcal{A}|}{\delta}\right)}{n}\right)}\right.  \tag{88}\\
& \hat{J}\left(\hat{\pi}, \hat{\alpha}_{x}\right)-J\left(\pi^{*}, \alpha_{x}^{*}\right) \leq \hat{J}\left(\hat{\pi}, \hat{\alpha}_{x}\right)-\hat{J}\left(\pi^{*}, \alpha *_{x}\right)+\mathcal{O}\left(c \sqrt{\left.\frac{\log (n) \log \left(\frac{2|\mathcal{X}||\mathcal{A}|}{\delta}\right)}{n}\right)}\right. \\
& \hat{V}^{*}-V^{*}\leq \underbrace{\leq 0 \text { by definition of }\left(\hat{\pi}, \hat{\alpha}_{x}\right)}, \hat{\alpha}_{x})-\hat{J}\left(\pi^{*}, \alpha *_{x}\right)+\mathcal{O} \\
&\left.c \sqrt{\frac{\log (n) \log \left(\frac{2|\mathcal{X}||\mathcal{A}|}{\delta}\right)}{n}}\right)  \tag{90}\\
& \hat{V}^{*}-V^{*} \leq \mathcal{O}\left(c \sqrt{\frac{\log (n) \log \left(\frac{2|\mathcal{X} \||\mathcal{A}|}{\delta}\right)}{n}}\right) \tag{91}
\end{align*}
$$

Combining these two shows the result gives the desired result.

## References

A. Agarwal, D. Hsu, S. Kale, J. Langford, L. Li, and R. Schapire. Taming the monster: A fast and simple algorithm for contextual bandits. In International Conference on Machine Learning, pages 1638-1646. PMLR, 2014.
D. Bertsimas and M. Sim. The price of robustness. Operations research, 52(1):35-53, 2004.
S. Boyd, S. P. Boyd, and L. Vandenberghe. Convex optimization. Cambridge university press, 2004.
Y. Chandak, G. Theocharous, S. Shankar, M. White, S. Mahadevan, and P. Thomas. Optimizing for the future in non-stationary mdps. In International Conference on Machine Learning, pages 1414-1425. PMLR, 2020.
Y. Chandak, S. Niekum, B. da Silva, E. Learned-Miller, E. Brunskill, and P. S. Thomas. Universal off-policy evaluation. Advances in Neural Information Processing Systems, 34, 2021.
C.-C. Chang and C.-J. Lin. Libsvm: a library for support vector machines. ACM transactions on intelligent systems and technology (TIST), 2(3):1-27, 2011.
W. Chu, L. Li, L. Reyzin, and R. Schapire. Contextual bandits with linear payoff functions. In Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics, pages 208-214. JMLR Workshop and Conference Proceedings, 2011.
J. Duchi and H. Namkoong. Learning models with uniform performance via distributionally robust optimization. arXiv preprint arXiv:1810.08750, 2018.
J. Duchi, P. Glynn, and H. Namkoong. Statistics of robust optimization: A generalized empirical likelihood approach. arXiv preprint arXiv:1610.03425, 2016.
J. C. Duchi, T. Hashimoto, and H. Namkoong. Distributionally robust losses against mixture covariate shifts. Under review, 2, 2019.
L. Faury, U. Tanielian, E. Dohmatob, E. Smirnova, and F. Vasile. Distributionally robust counterfactual risk minimization. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 34, pages 3850-3857, 2020.
S. G. Finlayson, A. Subbaswamy, K. Singh, J. Bowers, A. Kupke, J. Zittrain, I. S. Kohane, and S. Saria. The clinician and dataset shift in artificial intelligence. The New England journal of medicine, 385(3):283, 2021.
C. Gelada and M. G. Bellemare. Off-policy deep reinforcement learning by bootstrapping the covariate shift. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 33, pages 3647-3655, 2019.
A. S. Gerber, D. P. Green, and C. W. Larimer. Social pressure and voter turnout: Evidence from a large-scale field experiment. American political Science review, 102(1):33-48, 2008.
S. Giguere, B. Metevier, Y. Brun, P. S. Thomas, S. Niekum, and B. C. da Silva. Fairness guarantees under demographic shift. In International Conference on Learning Representations, 2021.
T. Hashimoto, M. Srivastava, H. Namkoong, and P. Liang. Fairness without demographics in repeated loss minimization. In International Conference on Machine Learning, pages 1929-1938. PMLR, 2018.
T. Hatt, D. Tschernutter, and S. Feuerriegel. Generalizing off-policy learning under sample selection bias. arXiv preprint arXiv:2112.01387, 2021.
W. Hu, G. Niu, I. Sato, and M. Sugiyama. Does distributionally robust supervised learning give robust classifiers? In International Conference on Machine Learning, pages 2029-2037. PMLR, 2018.
Z. Hu and L. J. Hong. Kullback-leibler divergence constrained distributionally robust optimization. Available at Optimization Online, 2013.
E. Imani, E. Graves, and M. White. An off-policy policy gradient theorem using emphatic weightings. Advances in Neural Information Processing Systems, 31, 2018.
N. Jiang and J. Huang. Minimax confidence interval for off-policy evaluation and policy optimization. arXiv preprint arXiv:2002.02081, 2020.
N. Jiang and L. Li. Doubly robust off-policy value evaluation for reinforcement learning. In International Conference on Machine Learning, pages 652-661. PMLR, 2016.
N. Kallus and M. Uehara. Double reinforcement learning for efficient off-policy evaluation in markov decision processes. Journal of Machine Learning Research, 21(167):1-63, 2020a.
N. Kallus and M. Uehara. Statistically efficient off-policy policy gradients. In International Conference on Machine Learning, pages 5089-5100. PMLR, 2020 b.
N. Kallus, X. Mao, K. Wang, and Z. Zhou. Doubly robust distributionally robust off-policy evaluation and learning. arXiv preprint arXiv:2202.09667, 2022.
M. Kato, M. Uehara, and S. Yasui. Off-policy evaluation and learning for external validity under a covariate shift. arXiv preprint arXiv:2002.11642, 2020.
E. S. Kim, R. S. Herbst, I. I. Wistuba, J. J. Lee, G. R. Blumenschein, A. Tsao, D. J. Stewart, M. E. Hicks, J. Erasmus, S. Gupta, et al. The battle trial: personalizing therapy for lung cancer. Cancer discovery, 1(1):44-53, 2011.
P. W. Koh, S. Sagawa, H. Marklund, S. M. Xie, M. Zhang, A. Balsubramani, W. Hu, M. Yasunaga, R. L. Phillips, I. Gao, et al. Wilds: A benchmark of in-the-wild distribution shifts. In International Conference on Machine Learning, pages 5637-5664. PMLR, 2021.
Q. Liu, L. Li, Z. Tang, and D. Zhou. Breaking the curse of horizon: Infinite-horizon off-policy estimation. arXiv preprint arXiv:1810.12429, 2018.
Y. Liu, A. Swaminathan, A. Agarwal, and E. Brunskill. Off-policy policy gradient with stationary distribution correction. In Uncertainty in Artificial Intelligence, pages 1180-1190. PMLR, 2020.
P. Michel, T. Hashimoto, and G. Neubig. Distributionally robust models with parametric likelihood ratios. arXiv preprint arXiv:2204.06340, 2022.
W. Mo, Z. Qi, and Y. Liu. Learning optimal distributionally robust individualized treatment rules. Journal of the American Statistical Association, 116(534):659-674, 2021.
H. Namkoong and J. C. Duchi. Stochastic gradient methods for distributionally robust optimization with f-divergences. Advances in neural information processing systems, 29, 2016.
Y. Oren, S. Sagawa, T. B. Hashimoto, and P. Liang. Distributionally robust language modeling. arXiv preprint arXiv:1909.02060, 2019.
D. Precup. Eligibility traces for off-policy policy evaluation. Computer Science Department Faculty Publication Series, page 80, 2000.
C. Qin and D. Russo. On adaptivity and confounding in contextual bandit experiments. In NeurIPS 2021 Workshop on Distribution Shifts: Connecting Methods and Applications, 2021.
H. Rahimian and S. Mehrotra. Distributionally robust optimization: A review. arXiv preprint arXiv:1908.05659, 2019.
S. Sagawa, P. W. Koh, T. B. Hashimoto, and P. Liang. Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. arXiv preprint arXiv:1911.08731, 2019.
O. Sakhi, L. Faury, and F. Vasile. Improving offline contextual bandits with distributional robustness. arXiv preprint arXiv:2011.06835, 2020.
S. Shalev-Shwartz, O. Shamir, N. Srebro, and K. Sridharan. Stochastic convex optimization. In COLT, volume 2, page 5, 2009.
A. Shapiro. Distributionally robust stochastic programming. SIAM Journal on Optimization, 27(4): 2258-2275, 2017.
N. Si, F. Zhang, Z. Zhou, and J. Blanchet. Distributional robust batch contextual bandits. arXiv preprint arXiv:2006.05630, 2020.
H. Singh, S. Joshi, F. Doshi-Velez, and H. Lakkaraju. Learning under adversarial and interventional shifts. arXiv preprint arXiv:2103.15933, 2021.
A. Strehl, J. Langford, L. Li, and S. M. Kakade. Learning from logged implicit exploration data. Advances in neural information processing systems, 23, 2010.
A. Subbaswamy, P. Schulam, and S. Saria. Preventing failures due to dataset shift: Learning predictive models that transport. In The 22nd International Conference on Artificial Intelligence and Statistics, pages 3118-3127. PMLR, 2019.
A. Swaminathan and T. Joachims. Batch learning from logged bandit feedback through counterfactual risk minimization. The Journal of Machine Learning Research, 16(1):1731-1755, 2015a.
A. Swaminathan and T. Joachims. The self-normalized estimator for counterfactual learning. advances in neural information processing systems, 28, 2015b.
P. Thomas and E. Brunskill. Data-efficient off-policy policy evaluation for reinforcement learning. In International Conference on Machine Learning, pages 2139-2148, 2016 a.
P. S. Thomas and E. Brunskill. Magical policy search: Data efficient reinforcement learning with guarantees of global optimality. In European workshop on reinforcement learning, 2016b.
P. S. Thomas, G. Theocharous, M. Ghavamzadeh, I. Durugkar, and E. Brunskill. Predictive off-policy policy evaluation for nonstationary decision problems, with applications to digital marketing. In Twenty-Ninth IAAI Conference, 2017.
A. Xie, J. Harrison, and C. Finn. Deep reinforcement learning amidst lifelong non-stationarity. arXiv preprint arXiv:2006.10701, 2020.
M. Yang, O. Nachum, B. Dai, L. Li, and D. Schuurmans. Off-policy evaluation via the regularized lagrangian. arXiv preprint arXiv:2007.03438, 2020.

