## A Appendix

## A. 1 Connection to online learning

In Section 2 we motivated the update (2) as a way to adjust the size of our prediction sets in response to the realized historical miscoverage frequency. Alternatively, one could also derive (2) as an online gradient descent algorithm with respect to the pinball loss. To be more precise let

$$
\beta_{t}:=\sup \left\{\beta: Y_{t} \in \hat{C}_{t}(\beta)\right\}
$$

where we remark that $\hat{C}_{t}\left(\beta_{t}\right)$ can be thought of as the smallest prediction set containing $Y_{t}$. Additionally, define the pinball loss

$$
\rho_{\alpha}(u)=\left\{\begin{array}{l}
\alpha u, u>0 \\
-(1-\alpha) u, u \leq 0
\end{array}\right.
$$

and consider the loss $\ell\left(\alpha_{t}, \beta_{t}\right)=\rho_{\alpha}\left(\beta_{t}-\alpha_{t}\right)$. Then, one directly computes that

$$
\alpha_{t}-\gamma \partial_{\alpha_{t}} \ell\left(\alpha_{t}, \beta_{t}\right)=\alpha_{t}+\gamma\left(\alpha-\mathbb{1}_{\alpha_{t}>\beta_{t}}\right)=\alpha_{t}+\gamma\left(\alpha-\operatorname{err}_{t}\right) .
$$

Because the pinball loss is convex, this gradient descent update falls within a well understood class of algorithms that have been extensively studied in the online learning literature (see e.g. [17]). A standard analysis may then be to bound the regret of $\alpha_{t}$ defined as

$$
\operatorname{Reg}_{T}:=\sum_{t=1}^{T} \ell\left(\alpha_{t}, \beta_{t}\right)-\min _{\beta} \sum_{t=1}^{T} \ell\left(\beta, \beta_{t}\right)
$$

Unfortunately, this notion of regret fails to capture our intuition that $\alpha_{t}$ is adaptively tracking the moving target $\alpha_{t}^{*}$. Thus, we make the connection to online gradient descent only in passing and develop alternative theoretical tools in Section 4

## A. 2 Stock prices

Figure 4 shows daily open prices for the four stocks considered in Section 2.2.


Figure 4: Daily open prices for the four stocks considered in Section 2.2

## A. 3 Trajectories of $\boldsymbol{\alpha}_{\boldsymbol{t}}$

In this section we show the realized trajectories of $\alpha_{t}$ obtained in our experiments from Sections 2.2 and 6 .


BlackBerry


AMD


Fannie Mae


Figure 5: Realized trajectories of $\alpha_{t}$ for predicting stock market volatility as outlined in Section 2.2 using update (2).


Figure 6: Realized trajectories of $\alpha_{t}$ for predicting stock market volatility as outlined in Section 2.2 using update (3) with $w_{s} \propto 0.95^{t-s}$.


Figure 7: Realized trajectory of $\alpha_{t}$ for election night forecasting as outlined in Section6 using update (2).


Figure 8: Realized trajectory of $\alpha_{t}$ for election night forecasting as outlined in Section6using update (3) with $w_{s} \propto 0.95^{t-s}$.

## A. 4 Coverage for additional stocks

Figure 9 shows the local coverage level of adaptive and non-adaptive conformal inference for the prediction of market volatility (see Section 2.2) for 8 additional stocks/indices.


Figure 9: Local coverage frequencies for adaptive conformal (blue), a non-adaptive method that holds $\alpha_{t}=\alpha$ fixed (red), and an i.i.d. Bernoulli(0.1) sequence (grey) for the prediction of market volatility. The coloured dotted lines mark the average coverage obtained across all time points, while the black line indicates the target level of $1-\alpha=0.9$.

## A. 5 Existence of a stationary distribution for $\left(\alpha_{t}, A_{t}\right)$

In this section we give one simple example in which the Markov chain $\left\{\left(\alpha_{t}, A_{t}\right)\right\}_{t \in \mathbb{N}}$ will have a unique stationary distribution. The setting considered here is the same as the one described in Section 4.2.

Let $\alpha_{t}$ be initialized with $\alpha_{1} \in\{\alpha+k \gamma \alpha: k \in \mathbb{Z}\}$. Assume that $\mathcal{A}$ is finite. Let $P$ be the transition matrix of $\left\{A_{t}\right\}_{t \in \mathbb{N}}$ and assume that for all $a_{1}, a_{2} \in \mathcal{A}, P_{a_{1}, a_{2}}=\mathbb{P}\left(A_{t+1}=a_{2} \mid A_{t}=a_{1}\right)>0$. Assume that $\alpha$ satisfies $\alpha^{-1}(1-\alpha) \in \mathbb{N}$. This will hold for common choices of $\alpha$ such as $\alpha=0.1$ and $\alpha=0.05$. Finally, assume that for all $a \in \mathcal{A}$ and all $p \in(0,1), \mathbb{P}\left(S\left(X_{t}, Y_{t}\right) \leq \hat{Q}(p) \mid A=a\right) \in$ $(0,1)$. This will occur for example when $S\left(X_{t}, Y_{t}\right) \mid A_{t}=a$ is supported on $\mathbb{R}$ and $\hat{Q}(\cdot)$ is finite valued for all $p \in(0,1)$.
We claim that in this case $\left\{\left(\alpha_{t}, A_{t}\right)\right\}_{t \in \mathbb{N}}$ has a unique stationary distribution. To prove this it is sufficient to show that this chain is irreducible and has a finite state space. To check that it has a finite state space it is sufficient to show that $\alpha_{t}$ has a finite state space. We claim that with probability one we have that for all $t \in \mathbb{N}, \alpha_{t} \in\{\alpha+k \gamma \alpha: k \in \mathbb{Z}\}$. To prove this we proceed by induction. The base case is given by our choice of $\alpha_{1}$. For the inductive step note that

$$
\begin{aligned}
& \alpha_{t} \in\{\alpha+k \gamma \alpha: k \in \mathbb{Z}\} \\
& \quad \Longrightarrow \alpha_{t+1}=\alpha_{t}+\gamma\left(\alpha-\operatorname{err}_{t}\right)=\left\{\begin{array}{l}
\alpha_{t}+\gamma \alpha, \text { if } \operatorname{err}_{t}=0, \\
\alpha_{t}-\gamma \alpha\left(\alpha^{-1}(1-\alpha)\right), \text { if } \operatorname{err}_{t}=1,
\end{array} \quad \in\{\alpha+k \gamma \alpha: k \in \mathbb{Z}\} .\right.
\end{aligned}
$$

Since $\alpha_{t}$ is also bounded (see Lemma 4.1) this implies that $\alpha_{t}$ has a finite state space.
Finally, the fact that $\left\{\left(\alpha_{t}, A_{t}\right)\right\}_{t \in \mathbb{N}}$ is irreducible follows easily from our assumptions on $\left\{A_{t}\right\}_{t \in \mathbb{N}}$, $\hat{Q}(\cdot)$, and $\mathbb{P}\left(S\left(X_{t}, Y_{t}\right) \leq \hat{Q}(p) \mid A=a\right)$.

## A. 6 Additional information for Section 6

## A.6.1 Dataset description

The county-level demographic characteristics used for prediction were the proportion of the total population that fell into each of the following race categories (either alone or in combination): black or African American, American Indian or Alaska Native, Asian, Native Hawaiian or other Pacific islander. In addition to this, we also used the proportion of the total population that was male, of Hispanic origin, and that fell within each of the age ranges 20-29, 30-44, 45-64, and 65+. Demographic information was obtained from 2019 estimates published by the United States Census Bureau and available at [5]. To supplement these demographic features we also used the median household income and the percentage of individuals with a bachelors degree or higher as covariates. Data on county-level median household incomes was based on 2019 estimates obtained from [7]. The percentage of individuals with a bachelors degree or higher was computed based on data collected during the years 2015-2019 and published at [6]. As an aside, we remark that we used 2019 estimates because this was the most recent year for which data was available.
Vote counts for the 2016 election were obtained from [14], while 2020 election data was taken from [25]. In total, matching covariate and election vote count data were obtained for 3111 counties.

## A.6.2 Detailed prediction algorithm

Algorithm 1 below outlines the core conformal inference method used to predict election results. An R implementation of this algorithm as well as the core method outlined in Section 2.2 can be found at https://github.com/isgibbs/AdaptiveConformal.

## A. 7 Large deviation bounds for the error sequence

In this section we prove Theorem4.1. So, throughout we define the sequences $\left\{\alpha_{t}\right\}_{t \in \mathbb{N}},\left\{\operatorname{err}_{t}\right\}_{t \in \mathbb{N}}$, and $\left\{A_{t}\right\}_{t \in \mathbb{N}}$ as in Section 4.2 and we assume that $\left\{\left(\alpha_{t}, A_{t}\right)\right\}_{t \in \mathbb{N}}$ is a stationary Markov chain from which it follows immediately that $\left\{\left(\alpha_{t}, A_{t}, \mathrm{err}_{t}\right)\right\}_{t \in \mathbb{N}}$ is also stationary. Additionally, we assume that $\hat{Q}(\cdot)$ and $S(\cdot)$ are fixed functions such that $\hat{Q}(\cdot)$ is non-decreasing with $\hat{Q}(x)=-\infty$ for all $x<0$ and $\hat{Q}(x)=\infty$ for all $x>1$. The proof of Theorem 4.1 will rely on the following lemmas.

```
Algorithm 1: CQR method for election night prediction
Data: Observed sequence of county-level votes counts and covariates \(\left\{\left(X_{t}, Y_{t}\right)\right\}_{1 \leq t \leq T}\) and vote
    counts for the democratic candidate in the previous election \(\left\{Y_{t}^{\text {prev }}\right\}_{1 \leq t \leq T}\).
for \(t=1,2, \ldots, T\) do
    Compute the residual \(r_{t}=\left(Y_{t}-Y_{t}^{\mathrm{prev}}\right) / Y_{t}^{\mathrm{prev}}\)
for \(t=501,502, \ldots, T\) do
    // We start making predictions once 500 counties have been observed.
    Randomly split the data \(\left\{\left(X_{l}, r_{l}\right)\right\}_{1 \leq l \leq t-1}\) into a training set \(\mathcal{D}_{\text {train }}\) and a calibration set \(\mathcal{D}_{\text {cal }}\)
        with \(\left|\mathcal{D}_{\text {train }}\right|=\lfloor(t-1) \cdot 0.75\rfloor\);
    Fit a linear quantile regression model \(\hat{q}(x ; p)\) on \(\mathcal{D}_{\text {train }}\);
    for \(\left(X_{l}, r_{l}\right) \in \mathcal{D}_{\text {cal }}\) do
            Compute the conformity score \(S_{l}=\max \left\{\hat{q}\left(X_{l} ; \alpha / 2\right)-r_{l}, r_{l}-\hat{q}\left(X_{l} ; 1-\alpha / 2\right)\right\}\);
        Define the quantile function \(\hat{Q}_{t}(p)=\inf \left\{x:\left(\frac{1}{\left|\mathcal{D}_{\text {cal }}\right|} \sum_{\left(X_{l}, r_{l}\right) \in \mathcal{D}_{\text {cal }}} \mathbb{1}_{S_{l} \leq x}\right) \geq p\right\}\);
    Return the prediction set
        \(\hat{C}_{t}(\alpha):=\left\{y: \max \left\{\hat{q}\left(X_{t} ; \alpha / 2\right)-\frac{y-Y_{t}^{\text {prev }}}{Y_{t}^{\text {prev }}}, \frac{y-Y_{t}^{\text {prev }}}{Y_{t}^{\text {prev }}}-\hat{q}\left(X_{t} ; 1-\alpha / 2\right)\right\} \leq \hat{Q}_{t}(1-\alpha)\right\} ;\)
```

Lemma A. 1 Let $f: \mathbb{R} \rightarrow \mathbb{R}$ and $g: \mathbb{R} \rightarrow \mathbb{R}$ be bounded functions such that either

1. $f$ is non-increasing and $g$ is non-decreasing,
2. or $f$ is non-decreasing and $g$ is non-increasing.

Then, for any random variable $Y$

$$
\mathbb{E}[f(Y) g(Y)] \leq \mathbb{E}[f(Y)] \mathbb{E}[g(Y)]
$$

The proof of this result is straightforward and can be found in Section A. 10
Lemma A. 2 For any $\lambda \in \mathbb{R}$ and $t \in \mathbb{N}$ we have that

$$
\begin{equation*}
\mathbb{E}\left[\prod_{s=1}^{t} \exp \left(\lambda\left(\operatorname{err}_{s}-\mathbb{E}\left[\operatorname{err}_{s} \mid A_{s}\right]\right)\right)\right] \leq \exp \left(\lambda^{2} / 2\right) \mathbb{E}\left[\prod_{s=1}^{t-1} \exp \left(\lambda\left(\operatorname{err}_{s}-\mathbb{E}\left[\operatorname{err}_{s} \mid A_{s}\right]\right)\right)\right] \tag{8}
\end{equation*}
$$

Proof: By conditioning on $\alpha_{1}$ and $A_{1}, \ldots, A_{t}$ on both the left and right-hand side of (8) we may view these quantities as fixed. Thus, while for readability we do not denote this conditioning explicitly, the following calculations should be read as conditional on $\alpha_{1}$ and $A_{1}, \ldots, A_{t}$. Then,

$$
\begin{aligned}
\mathbb{E}[ & \left.\prod_{s=1}^{t} \exp \left(\lambda\left(\operatorname{err}_{s}-\mathbb{E}\left[\operatorname{err}_{s} \mid A_{s}\right]\right)\right)\right] \\
& =\mathbb{E}\left[\prod_{s=1}^{t-1} \exp \left(\lambda\left(\operatorname{err}_{s}-\mathbb{E}\left[\operatorname{err}_{s} \mid A_{s}\right]\right)\right) \mathbb{E}\left[\exp \left(\lambda\left(\operatorname{err}_{t}-\mathbb{E}\left[\operatorname{err}_{t} \mid A_{t}\right]\right)\right) \mid \operatorname{err}_{1}, \ldots, \operatorname{err}_{t-1}\right]\right] .
\end{aligned}
$$

Recall that $\alpha_{t}=\alpha_{1}+\gamma \sum_{s=1}^{t-1}\left(\alpha-\operatorname{err}_{s}\right)$ is a deterministic function of $\alpha_{1}$ and $\sum_{s=1}^{t-1} \operatorname{err}_{s}$. So, we may define the functions $f\left(\sum_{s=1}^{t-1} \operatorname{err}_{s}\right)=\prod_{s=1}^{t-1} \exp \left(\lambda\left(\operatorname{err}_{s}-\mathbb{E}\left[\operatorname{err}_{s} \mid A_{s}\right]\right)\right)$ and

$$
\begin{aligned}
g\left(\sum_{s=1}^{t-1} \operatorname{err}_{s}\right): & =\mathbb{E}\left[\exp \left(\lambda\left(\operatorname{err}_{t}-\mathbb{E}\left[\operatorname{err}_{t} \mid A_{t}\right]\right)\right) \mid \operatorname{err}_{1}, \ldots, \operatorname{err}_{t-1}\right] \\
= & P_{A_{t}}\left(S\left(X_{t}, Y_{t}\right) \leq \hat{Q}\left(1-\alpha_{t}\right)\right) \exp \left(-\lambda \mathbb{E}\left[\operatorname{err}_{t} \mid A_{t}\right]\right) \\
& +\left(1-P_{A_{t}}\left(S\left(X_{t}, Y_{t}\right) \leq \hat{Q}\left(1-\alpha_{t}\right)\right)\right) \exp \left(\lambda\left(1-\mathbb{E}\left[\operatorname{err}_{t} \mid A_{t}\right]\right)\right)
\end{aligned}
$$

where we emphasize that on the last line $A_{t}$ and $\alpha_{t}$ should be viewed as fixed quantities. Now, since $\alpha_{t}$ is monotonically decreasing in $\sum_{s=1}^{t-1} \operatorname{err}_{s}$ it should be clear that if $\lambda \geq 0$ then $g$ is non-increasing
(resp. non-decreasing for $\lambda<0$ ) and $f$ is non-decreasing (resp. non-increasing for $\lambda<0$ ). So, by Lemma A.1 we have that

$$
\begin{aligned}
& \mathbb{E}\left[\prod_{s=1}^{t-1} \exp \left(\lambda\left(\operatorname{err}_{s}-\mathbb{E}\left[\operatorname{err}_{s} \mid A_{s}\right]\right)\right) \mathbb{E}\left[\exp \left(\lambda\left(\operatorname{err}_{t}-\mathbb{E}\left[\operatorname{err}_{t} \mid A_{t}\right]\right)\right) \mid \operatorname{err}_{1}, \ldots, \operatorname{err}_{t-1}\right]\right] \\
& \leq \mathbb{E}\left[\prod_{s=1}^{t-1} \exp \left(\lambda\left(\operatorname{err}_{s}-\mathbb{E}\left[\operatorname{err}_{s} \mid A_{s}\right]\right)\right)\right] \mathbb{E}\left[\mathbb{E}\left[\exp \left(\lambda\left(\operatorname{err}_{t}-\mathbb{E}\left[\operatorname{err}_{t} \mid A_{t}\right]\right)\right) \mid \operatorname{err}_{1}, \ldots, \operatorname{err}_{t-1}\right]\right] \\
& =\mathbb{E}\left[\prod_{s=1}^{t-1} \exp \left(\lambda\left(\operatorname{err}_{s}-\mathbb{E}\left[\operatorname{err}_{s} \mid A_{s}\right]\right)\right)\right] \mathbb{E}\left[\exp \left(\lambda\left(\operatorname{err}_{t}-\mathbb{E}\left[\operatorname{err}_{t} \mid A_{t}\right]\right)\right)\right] \\
& \leq \mathbb{E}\left[\prod_{s=1}^{t-1} \exp \left(\lambda\left(\operatorname{err}_{s}-\mathbb{E}\left[\operatorname{err}_{s} \mid A_{s}\right]\right)\right)\right] \exp \left(\lambda^{2} / 2\right)
\end{aligned}
$$

where the last inequality follows by Hoeffding's lemma (see Lemma A.5).
The final result we will need in order to prove Theorem 4.1 is a large deviation bound for Markov chains.

Definition A. 1 Let $\left\{X_{i}\right\}_{i \in \mathbb{N}} \subseteq \mathcal{X}$ be a Markov chain with transition kernel $P$ and stationary distribution $\pi$. Define the inner product space

$$
L^{2}(\mathcal{X}, \pi)=\left\{h: \int_{\mathcal{X}} h(x)^{2} \pi(d x)<\infty\right\}
$$

with inner product

$$
\left\langle h_{1}, h_{2}\right\rangle_{\pi}=\int_{\mathcal{X}} h_{1}(x) h_{2}(x) \pi(d x)
$$

For any $h \in L^{2}(\mathcal{X}, \pi)$, let

$$
L^{2}(\mathcal{X}, \pi) \ni P h:=\int h(y) P(\cdot, d y)
$$

Then, we say that $\left\{X_{i}\right\}_{i \in \mathbb{N}}$ has non-zero absolute spectral gap $1-\eta$ if

$$
\eta:=\sup \left\{\sqrt{\langle P h, P h\rangle_{\pi}}:\langle h, h\rangle_{\pi}=1, \int_{\mathcal{X}} h(x) \pi(d x)=0\right\}<1
$$

Theorem A. 1 (Theorem 1 in [19]) Let $\left\{X_{i}\right\}_{i \in \mathbb{N}} \subseteq \mathcal{X}$ be a stationary Markov chain with invariant distribution $\pi$ and non-zero absolute spectral gap $1-\eta>0$. Let $f_{i}: \mathcal{X} \rightarrow[-C, C]$ be a sequence of functions with $\pi\left(f_{i}\right)=0$ and define $\sigma^{2}:=\sum_{i=1}^{n} \pi\left(f_{i}^{2}\right) / n$. Then, for $\forall \epsilon>0$,

$$
\mathbb{P}_{\pi}\left(\frac{1}{n} \sum_{i=1}^{n} f_{i}\left(X_{i}\right) \geq \epsilon\right) \leq \exp \left(-\frac{n(1-\eta) \epsilon^{2} / 2}{(1+\eta) \sigma^{2}+5 C \epsilon}\right)
$$

Finally, we are ready to prove Theorem 4.1.
Proof: [Proof of Theorem4.1] Write

$$
\begin{align*}
& \mathbb{P}\left(\left|\frac{1}{T} \sum_{t=1}^{T} \operatorname{err}_{t}-\alpha\right|>\epsilon\right)  \tag{9}\\
& \leq \mathbb{P}\left(\left|\frac{1}{T} \sum_{t=1}^{T} \operatorname{err}_{t}-\mathbb{E}\left[\operatorname{err}_{t} \mid A_{t}\right]\right|>\epsilon / 2\right)+\mathbb{P}\left(\left|\frac{1}{T} \sum_{t=1}^{T} \mathbb{E}\left[\operatorname{err}_{t} \mid A_{t}\right]-\alpha\right|>\epsilon / 2\right) \tag{10}
\end{align*}
$$

By applying lemma A. 2 inductively we have that for all $\lambda>0$,

$$
\begin{aligned}
\mathbb{P}\left(\frac{1}{T} \sum_{t=1}^{T} \operatorname{err}_{t}-\mathbb{E}\left[\operatorname{err}_{t} \mid A_{t}\right]>\epsilon / 2\right) & \leq \exp (-T \lambda \epsilon / 2) \mathbb{E}\left[\prod_{t=1}^{T} \exp \left(\lambda\left(\operatorname{err}_{t}-\mathbb{E}\left[\operatorname{err}_{t} \mid A_{t}\right]\right)\right)\right] \\
& \leq \exp (-T \lambda \epsilon / 2) \exp \left(T \lambda^{2} / 2\right)
\end{aligned}
$$

with an identical bound on the left tail. Choosing $\lambda=\epsilon / 2$ gives the bound

$$
\mathbb{P}\left(\left|\frac{1}{T} \sum_{t=1}^{T} \operatorname{err}_{t}-\mathbb{E}\left[\operatorname{err}_{t} \mid A_{t}\right]\right|>\epsilon / 2\right) \leq 2 \exp \left(-T \epsilon^{2} / 8\right)
$$

On the other hand, the second term in can be bounded directly using Theorem A. 1.

## A. 8 Approximate marginal coverage

In this section we prove Theorem4.2.
Proof: [Proof of Theorem4.2] Our proof follows similar steps to those presented in previous works on stochastic gradient descent under distribution shift [9, 35]. First, note that

$$
\left(\alpha_{t+1}-\alpha_{A_{t}}^{*}\right)^{2}=\left(\alpha_{t}-\alpha_{A_{t}}^{*}\right)^{2}+2 \gamma\left(\alpha-\operatorname{err}_{t}\right)\left(\alpha_{t}-\alpha_{A_{t}}^{*}\right)+\gamma^{2}\left(\alpha-\operatorname{err}_{t}\right)^{2}
$$

Now recalling that $M\left(\alpha_{A_{t}}^{*} \mid A_{t}\right)=\alpha$, we find that

$$
\begin{aligned}
-\mathbb{E}\left[\left(\alpha-\operatorname{err}_{t}\right)\left(\alpha_{t}-\alpha_{A_{t}}^{*}\right)\right] & =\mathbb{E}\left[\mathbb{E}\left[\left(\operatorname{err}_{t}-\alpha\right)\left(\alpha_{t}-\alpha_{A_{t}}^{*}\right) \mid A_{t}, \alpha_{t}\right]\right] \\
& =\mathbb{E}\left[\left(M\left(\alpha_{t} \mid A_{t}\right)-M\left(\alpha_{A_{t}}^{*} \mid A_{t}\right)\right)\left(\alpha_{t}-\alpha_{A_{t}}^{*}\right)\right] \\
& \geq \frac{1}{L} \mathbb{E}\left[\left(M\left(\alpha_{t} \mid A_{t}\right)-M\left(\alpha_{A_{t}}^{*} \mid A_{t}\right)\right)^{2}\right] \\
& =\frac{1}{L} \mathbb{E}\left[\left(M\left(\alpha_{t} \mid A_{t}\right)-\alpha\right)^{2}\right]
\end{aligned}
$$

Thus it follows that

$$
\begin{aligned}
& 2 \gamma L^{-1} \sum_{t=1}^{T} \mathbb{E}\left[\left(M\left(\alpha_{t} \mid A_{t}\right)-\alpha\right)^{2}\right] \leq \sum_{t=1}^{T} \mathbb{E}\left[\left(\alpha_{t}-\alpha_{A_{t}}^{*}\right)^{2}-\left(\alpha_{t+1}-\alpha_{A_{t}}^{*}\right)^{2}+\gamma^{2}\left(\alpha-\operatorname{err}_{t}\right)^{2}\right] \\
& \leq \sum_{t=1}^{T} \mathbb{E}\left[\left(\alpha_{t+1}-\alpha_{A_{t+1}}^{*}\right)^{2}-\left(\alpha_{t+1}-\alpha_{A_{t}}^{*}\right)^{2}\right]+\mathbb{E}\left[\left(\alpha_{1}-\alpha_{A_{1}}^{*}\right)^{2}\right]+\gamma^{2} T \\
& \leq \sum_{t=1}^{T} \mathbb{E}\left[2 \alpha_{t+1}\left(\alpha_{A_{t}}^{*}-\alpha_{A_{t+1}}^{*}\right)\right]+\left(\alpha_{A_{T+1}}^{*}\right)^{2}+\mathbb{E}\left[\left(\alpha_{1}-\alpha_{A_{1}}^{*}\right)^{2}\right]+\gamma^{2} T \\
& \leq \sum_{t=1}^{T} 2(1+\gamma) \mathbb{E}\left[\left|\alpha_{A_{t+1}}^{*}-\alpha_{A_{t}}^{*}\right|\right]+\left(\alpha_{A_{T+1}}^{*}\right)^{2}+\mathbb{E}\left[\left(\alpha_{1}-\alpha_{A_{1}}^{*}\right)^{2}\right]+\gamma^{2} T
\end{aligned}
$$

where the last inequality follows from Lemma 4.1. So, re-arranging we get the inequality

$$
\begin{aligned}
& \frac{1}{T} \sum_{t=1}^{T} \mathbb{E}\left[\left(M\left(\alpha_{t} \mid A_{t}\right)-\alpha\right)^{2}\right] \\
& \leq \frac{L}{2 T \gamma}\left(\sum_{t=1}^{T} 2(1+\gamma) \mathbb{E}\left[\left|\alpha_{A_{t+1}}^{*}-\alpha_{A_{t}}^{*}\right|\right]+\left(\alpha_{A_{T+1}}^{*}\right)^{2}+\mathbb{E}\left[\left(\alpha_{1}-\alpha_{A_{1}}^{*}\right)^{2}\right]+\gamma^{2} T\right)
\end{aligned}
$$

Finally, since $\left\{\left(A_{t}, \alpha_{t}\right)\right\}_{t \in \mathbb{N}}$ is stationary we may let $T \rightarrow \infty$ to get that

$$
\mathbb{E}\left[\left(M\left(\alpha_{t} \mid A_{t}\right)-\alpha\right)^{2}\right] \leq \frac{L(1+\gamma)}{\gamma} \mathbb{E}\left[\left|\alpha_{A_{t+1}}^{*}-\alpha_{A_{t}}^{*}\right|\right]+\frac{L \gamma}{2}
$$

as claimed.

## A. 9 Bounds on $B$ and $\sigma_{B}^{2}$

In this section we bound the constants $B$ and $\sigma_{B}^{2}$ appearing in the statement of Theorem4.1 Let

$$
\begin{aligned}
\epsilon_{1} & =\sup _{k \in\{0,1,2, \ldots\}} \sup _{a \in \mathcal{A}} \mathbb{E}\left[\mid \alpha_{A_{T-k}}^{*}-\alpha_{A_{T-k-1}}^{*} \| A_{T}=a\right] \\
\text { and } \epsilon_{2} & =\sup _{k \in\{0,1,2, \ldots\}} \sup _{a \in \mathcal{A}} \mathbb{E}\left[\left(\alpha_{A_{T-k}}^{*}-\alpha_{A_{T-k-1}}^{*}\right)^{2} \mid A_{T}=a\right] .
\end{aligned}
$$

Then, our main result is Lemma A. 4 which shows that

$$
\begin{equation*}
B \leq C\left(\gamma+\gamma^{-1}\left(\epsilon_{1}+\epsilon_{2}\right)\right) \quad \text { and } \quad \sigma_{B}^{2} \leq B^{2} \tag{11}
\end{equation*}
$$

where the constant $C$ depends on how close $M(\cdot \mid a)$ is to the ideal linear function $M(p \mid a)=p$.

Plugging 11 into Theorem 4.1 gives a concentration inequality for $\left|T^{-1} \sum_{t=1}^{T} \operatorname{err}_{t}-\alpha\right|$. In particular, suppose we use an optimal stepsize of $\gamma \propto \sqrt{\epsilon_{1}}$. Then, combining (11) with Theorem4.1 roughly tells us that

$$
\begin{equation*}
\left|\frac{1}{T} \sum_{t=1}^{T} \operatorname{err}_{t}-\alpha\right| \leq O\left(\max \left\{\frac{1}{\sqrt{T}}, \frac{\sqrt{\epsilon_{1}}}{\sqrt{T(1-\eta)}}, \frac{\sqrt{\epsilon_{1}}}{T(1-\eta)}\right\}\right) . \tag{12}
\end{equation*}
$$

As a comparison it may be instructive to note that the more naive bound given in Proposition 4.1 can be written as

$$
\begin{equation*}
\left|\frac{1}{T} \sum_{t=1}^{T} \operatorname{err}_{t}-\alpha\right| \leq O\left(\frac{1}{T \gamma}\right)=O\left(\frac{1}{T \sqrt{\epsilon_{1}}}\right) \tag{13}
\end{equation*}
$$

A sharp reader may notice that the naive bound given in (13) actually goes to 0 faster in $T$ than the HMM-based bound shown in (12). While this is true, the bound (13) has the highly undesirable property of increasing in $1 / \sqrt{\epsilon_{1}}$, i.e. the bound increases as the size of the distribution shift decreases. On the other hand, the HMM-based bound has the more intuitive property of decreasing with the distribution shift.
The only remaining issue is to determine the size of $\sqrt{\epsilon_{1} /(1-\eta)}$. To provide some insight into this quantity note that there are two main regimes in which we expect $\left|T^{-1} \sum_{t=1}^{T} \operatorname{err}_{t}-\alpha\right|$ to be small:

1. Environments in which the state $A_{t}$ changes frequently, but $\left|\alpha_{A_{t+1}}^{*}-\alpha_{A_{t}}^{*}\right|$ is always small. In this case it is reasonable to expect $1-\eta$ to not be too small and so we anticipate that 12 will give a reasonable bound.
2. Environments in which the state changes very infrequently. In this case $\left\{A_{t}\right\}_{t \in \mathbb{N}}$ will mix slowly and so we expect $1-\eta$ to be quite small. Additionally, we also have that $\alpha_{A_{t+1}}^{*}=\alpha_{A_{t}}^{*}$ a large proportion of the time and thus $\epsilon_{1}$ will also be small. As a result, it is not immediately clear what 12 tells us about $\left|T^{-1} \sum_{t=1}^{T} \operatorname{err}_{t}-\alpha\right|$. Below we give one simple example that demonstrates that (12) can also be a reasonable bound in this instance.

Example A. 1 Let $\left\{A_{t}\right\}_{t \in \mathbb{N}}$ be the Markov chain with states $\{1, \ldots, n\}$ and transition matrix

$$
P=\left(p-\frac{1-p}{n-1}\right) I+\frac{1-p}{n-1} 11^{T}
$$

where $p \in[0,1]$ is taken to be very close to 1 . Let $\Delta:=\max _{i \neq j}\left|\alpha_{i}^{*}-\alpha_{j}^{*}\right|$. Then, we have that $\epsilon_{1}, \epsilon_{2} \leq \Delta(1-p)$. Moreover, note that this chain has spectral gap $1-\eta \cong 1-p$. Thus, 12, simplifies to

$$
\begin{equation*}
\left|\frac{1}{T} \sum_{t=1}^{T} e r r_{t}-\alpha\right| \leq O\left(\max \left\{\frac{1}{\sqrt{T}}, \frac{\sqrt{\Delta}}{T \sqrt{1-p}}\right\}\right) . \tag{14}
\end{equation*}
$$

In particular, for $T>\Delta /(1-p)$ we find that the error sequence concentrates at a rate of $O(1 / \sqrt{T})$, which is consistent with the behaviour of an i.i.d. Bernoulli sequence. Finally, to understand this restriction on $T$ note that given a starting state $j \in\{1, \ldots, n\}$ we expect $\alpha_{t}$ to contract towards $\alpha_{j}^{*}$ at a rate of $(1-\gamma)$ and therefore to have that

$$
\frac{1}{T} \sum_{t=1}^{T}\left|\alpha_{t}-\alpha_{j}^{*}\right| \propto \frac{1}{T} \sum_{t=1}^{T}(1-\gamma)^{t-1}\left|\alpha_{1}-\alpha_{j}^{*}\right| \leq \frac{\Delta}{T \gamma}=\frac{\sqrt{\Delta}}{T \sqrt{1-p}}
$$

where here we assumed that $\alpha_{1} \in\left\{\alpha_{1}^{*}, \ldots, \alpha_{n}^{*}\right\}$. Thus, the second term in the maximum in (14) can be seen as accounting for the rate of covergence of $\alpha_{t}$ to $\alpha_{j}^{*}$ during the time that the Markov chain is in state $j$ and given that $\alpha_{1}$ starts at some $\alpha_{i}^{*}, 1 \leq i \leq n$.

We now derive (11).

Lemma A. 3 Assume that $\exists 0<c<1 /(2 \gamma)$ such that for all $a \in \mathcal{A}$ and all $p \in[-\gamma, 1+\gamma]$,

$$
\left|M(p \mid a)-M\left(\alpha_{a}^{*} \mid a\right)\right| \geq c\left|p-\alpha_{a}^{*}\right|
$$

Then, for all $a \in \mathcal{A}, k \in\{0,1,2, \ldots\}$, and $T \in \mathbb{N}$

$$
\begin{aligned}
& \mathbb{E}\left[\left(\alpha_{1}-\alpha_{A_{1}}^{*}\right)^{2} \mid A_{1+k}=a\right] \\
& \leq(1-2 c \gamma)^{T-1} \mathbb{E}\left[\left(\alpha_{1}-\alpha_{A_{1}}^{*}\right)^{2} \mid A_{T+k}=a\right] \\
& +\sum_{t=2}^{T}(1-2 c \gamma)^{T-t}\left(\gamma^{2}+2(1+\gamma) \mathbb{E}\left[\left|\alpha_{A_{t-1}}^{*}-\alpha_{A_{t}}^{*}\right| \mid A_{T+k}=a\right]+\mathbb{E}\left[\left(\alpha_{A_{t-1}}^{*}-\alpha_{A_{t}}^{*}\right)^{2} \mid A_{T+k}=a\right]\right)
\end{aligned}
$$

Furthermore, if we assume that $\forall a \in \mathcal{A}$ and $k \in\{0,1,2, \ldots\}$,

$$
\mathbb{E}\left[\mid \alpha_{A_{t-k}}^{*}-\alpha_{A_{t-k-1}}^{*} \| A_{t}=a\right] \leq \epsilon_{1} \text { and } \mathbb{E}\left[\left(\alpha_{A_{t-k}}^{*}-\alpha_{A_{t-k-1}}^{*}\right)^{2} \mid A_{t}=a\right] \leq \epsilon_{2}
$$

then we find that

$$
\mathbb{E}\left[\left(\alpha_{1}-\alpha_{A_{1}}^{*}\right)^{2} \mid A_{1+k}=a\right] \leq \frac{1}{2 c \gamma}\left(\gamma^{2}+2(1+\gamma) \epsilon_{1}+\epsilon_{2}\right)
$$

Proof: Fix any $T \in \mathbb{N}$. Since $\left\{\left(\alpha_{t}, A_{t}\right)\right\}_{t \in \mathbb{N}}$ is stationary we have that

$$
\mathbb{E}\left[\left(\alpha_{1}-\alpha_{A_{1}}^{*}\right)^{2} \mid A_{1+k}\right]=\mathbb{E}\left[\left(\alpha_{T}-\alpha_{A_{T}}^{*}\right)^{2} \mid A_{T+k}\right] .
$$

Now note that

$$
\begin{aligned}
\mathbb{E}\left[\left(\alpha_{T}-\alpha_{A_{T}}^{*}\right)^{2} \mid A_{T+k}\right]= & \mathbb{E}\left[\left(\alpha_{T}-\alpha_{A_{T-1}}^{*}\right)^{2} \mid A_{T+k}\right]+2 \mathbb{E}\left[\left(\alpha_{T}-\alpha_{A_{T-1}}^{*}\right)\left(\alpha_{A_{T-1}}^{*}-\alpha_{A_{T}}^{*}\right) \mid A_{T+k}\right] \\
& +\mathbb{E}\left[\left(\alpha_{A_{T-1}}^{*}-\alpha_{A_{T}}^{*}\right)^{2} \mid A_{T+k}\right] \\
\leq & \mathbb{E}\left[\left(\alpha_{T}-\alpha_{A_{T-1}}^{*}\right)^{2} \mid A_{T+k}\right]+2(1+\gamma) \mathbb{E}\left[\left|\alpha_{A_{T-1}}^{*}-\alpha_{A_{T}}^{*}\right| \mid A_{T+k}\right] \\
& +\mathbb{E}\left[\left(\alpha_{A_{T-1}}^{*}-\alpha_{A_{T}}^{*}\right)^{2} \mid A_{T+k}\right]
\end{aligned}
$$

where on the last line we have applied Lemma 4.1. The first term above can be bounded as

$$
\begin{aligned}
& \mathbb{E}\left[\left(\alpha_{T}-\alpha_{A_{T-1}}^{*}\right)^{2} \mid A_{T+k}\right] \\
& \leq \mathbb{E}\left[\left(\alpha_{T-1}-\alpha_{A_{T-1}}^{*}\right)^{2} \mid A_{T+k}\right]+2 \mathbb{E}\left[\gamma\left(\alpha-\operatorname{err}_{T-1}\right)\left(\alpha_{T-1}-\alpha_{A_{T-1}}^{*}\right) \mid A_{T+k}\right]+\gamma^{2},
\end{aligned}
$$

where we additionally have that

$$
\begin{aligned}
& \mathbb{E}\left[\left(\alpha-\operatorname{err}_{T}\right)\left(\alpha_{T-1}-\alpha_{A_{T-1}}^{*}\right) \mid A_{T+k}\right] \\
& =\mathbb{E}\left[\left(\alpha-\mathbb{E}\left[\operatorname{err}_{T-1} \mid A_{T-1}, \alpha_{T-1}, A_{T+k}\right]\right)\left(\alpha_{T-1}-\alpha_{A_{T-1}}^{*}\right) \mid A_{T+k}\right] \\
& =\mathbb{E}\left[\left(M\left(\alpha_{A_{T-1}}^{*} \mid A_{T_{1}}\right)-M\left(\alpha_{T-1} \mid A_{T-1}\right)\right)\left(\alpha_{T-1}-\alpha_{A_{T-1}}^{*}\right) \mid A_{T+k}\right] \\
& \leq-c \mathbb{E}\left[\left(\alpha_{T-1}-\alpha_{A_{T-1}}^{*}\right)^{2} \mid A_{T+k}\right] .
\end{aligned}
$$

Whence,

$$
\mathbb{E}\left[\left(\alpha_{T}-\alpha_{A_{T-1}}^{*}\right)^{2} \mid A_{T+k}\right] \leq(1-2 c \gamma) \mathbb{E}\left[\left(\alpha_{T-1}-\alpha_{A_{T-1}}^{*}\right)^{2} \mid A_{T+k}\right]+\gamma^{2}
$$

and plugging this into our first inequality yields

$$
\begin{aligned}
\mathbb{E}\left[\left(\alpha_{T}-\alpha_{A_{T}}^{*}\right)^{2} \mid A_{T+k}\right] \leq & (1-2 c \gamma) \mathbb{E}\left[\left(\alpha_{T-1}-\alpha_{A_{T-1}}^{*}\right)^{2} \mid A_{T+k}\right]+\gamma^{2} \\
& +2(1+\gamma) \mathbb{E}\left[\left|\alpha_{A_{T-1}}^{*}-\alpha_{A_{T}}^{*}\right| \mid A_{T+k}\right]+\mathbb{E}\left[\left(\alpha_{A_{T-1}}^{*}-\alpha_{A_{T}}^{*}\right)^{2} \mid A_{T+k}\right]
\end{aligned}
$$

Repeating this argument inductively gives

$$
\begin{aligned}
& \mathbb{E}\left[\left(\alpha_{T}-\alpha_{A_{T}}^{*}\right)^{2} \mid A_{T+k}=a\right] \\
& \leq(1-2 c \gamma)^{T-1} \mathbb{E}\left[\left(\alpha_{1}-\alpha_{A_{1}}^{*}\right)^{2} \mid A_{T+k}=a\right] \\
& +\sum_{t=2}^{T}(1-2 c \gamma)^{T-t}\left(\gamma^{2}+2(1+\gamma) \mathbb{E}\left[\left|\alpha_{A_{t-1}}^{*}-\alpha_{A_{t}}^{*}\right| \mid A_{T+k}=a\right]+\mathbb{E}\left[\left(\alpha_{A_{t-1}}^{*}-\alpha_{A_{t}}^{*}\right)^{2} \mid A_{T+k}=a\right]\right) .
\end{aligned}
$$

The final part of the lemma follows by sending $T \rightarrow \infty$.

Lemma A. 4 Assume that for all $a \in \mathcal{A}$ and $p \in[-\gamma, 1+\gamma], M(\cdot \mid a)$ admits the second order Taylor expansion

$$
M(p \mid a)-M\left(\alpha_{a}^{*} \mid a\right)=C_{a}^{1}\left(p-\alpha_{a}^{*}\right)+C_{p, a}^{2}\left(p-\alpha_{a}^{*}\right)^{2}
$$

where $0<c_{1} \leq C_{a}^{1} \leq C_{1}<1 / \gamma$ and $\left|C_{p, a}^{2}\right| \leq C_{2}$. Then, for all $T \in \mathbb{N}$ and $a \in \mathcal{A}$,

$$
\left|\mathbb{E}\left[\operatorname{err}_{1} \mid A_{1}=a\right]-\alpha\right| \leq C_{1}\left(1-\gamma c_{1}\right)^{T-1} \mathbb{E}\left[\left|\alpha_{1}-\alpha_{A_{1}}^{*}\right| \mid A_{T}=a\right]
$$

$$
\begin{aligned}
& +\sum_{t=1}^{T-1} C_{1} C_{2} \gamma\left(1-c_{1} \gamma\right)^{T-t-1} \mathbb{E}\left[\left(\alpha_{t}-\alpha_{A_{t}^{*}}\right)^{2} \mid A_{T}=a\right] \\
& +\sum_{t=1}^{T-1} C_{1}\left(1-c_{1} \gamma\right)^{T-t-1} \mathbb{E}\left[\left|\alpha_{A_{t+1}}^{*}-\alpha_{A_{t}}^{*}\right| \mid A_{T}\right]+C_{2} \mathbb{E}\left[\left(\alpha_{T}-\alpha_{A_{T}}^{*}\right)^{2} \mid A_{T}=a\right]
\end{aligned}
$$

Furthermore, suppose the assumptions of Lemma A.3 hold and that $\forall a \in \mathcal{A}$ and $k \in\{0,1,2, \ldots\}$,

$$
\mathbb{E}\left[\mid \alpha_{A_{T-k}}^{*}-\alpha_{A_{T-k-1}}^{*} \| A_{T}=a\right] \leq \epsilon_{1} \text { and } \mathbb{E}\left[\left(\alpha_{A_{T-k}}^{*}-\alpha_{A_{T-k-1}}^{*}\right)^{2} \mid A_{T}=a\right] \leq \epsilon_{2}
$$

Then $\forall a \in \mathcal{A}$,

$$
\left|\mathbb{E}\left[\operatorname{err}_{1} \mid A_{1}=a\right]-\alpha\right| \leq\left(C_{1} C_{2} \frac{1}{c_{1}}+C_{2}\right) \frac{1}{2 c \gamma}\left(\gamma^{2}+2(1+\gamma) \epsilon_{1}+\epsilon_{2}\right)+\frac{C_{1}}{c_{1} \gamma} \epsilon_{1}
$$

Proof: Fix any $T \in \mathbb{N}$. Since $\left\{\left(\mathrm{err}_{t}, A_{t}\right)\right\}_{t \in \mathbb{N}}$ is stationary we have that

$$
\mathbb{E}\left[\operatorname{err}_{1} \mid A_{1}=a\right]=\mathbb{E}\left[\operatorname{err}_{T} \mid A_{T}=a\right]
$$

Then, by Taylor expanding $M\left(\cdot \mid A_{T}\right)$ we find that

$$
\begin{aligned}
& \left|\mathbb{E}\left[\operatorname{err}_{T} \mid A_{T}=a\right]-\alpha\right|=\left|\mathbb{E}\left[M\left(\alpha_{T} \mid A_{T}\right)-M\left(\alpha_{A_{T}}^{*} \mid A_{T}\right) \mid A_{T}\right]\right| \\
& \leq\left|\mathbb{E}\left[C_{A_{T}}^{1}\left(\alpha_{T}-\alpha_{A_{T}}^{*}\right) \mid A_{T}\right]\right|+\mathbb{E}\left[C_{\alpha_{T}, A_{T}}^{2}\left(\alpha_{T}-\alpha_{A_{T}}^{*}\right)^{2} \mid A_{T}\right] \\
& \leq\left|\mathbb{E}\left[C_{A_{T}}^{1}\left(\alpha_{T}-\alpha_{A_{T}}^{*}\right) \mid A_{T}\right]\right|+C_{2} \mathbb{E}\left[\left(\alpha_{T}-\alpha_{A_{T}}^{*}\right)^{2} \mid A_{T}\right]
\end{aligned}
$$

The first term above can be further bounded as

$$
\begin{aligned}
& \left|\mathbb{E}\left[C_{A_{T}}^{1}\left(\alpha_{T}-\alpha_{A_{T}}^{*}\right) \mid A_{T}\right]\right| \\
& =\left|\mathbb{E}\left[C_{A_{T}}^{1}\left(\alpha_{T-1}+\gamma\left(\alpha-\operatorname{err}_{T-1}\right)-\alpha_{A_{T-1}}^{*}\right) \mid A_{T}\right]+\mathbb{E}\left[C_{A_{T}}^{1}\left(\alpha_{A_{T-1}}^{*}-\alpha_{A_{T}}^{*}\right) \mid A_{T}\right]\right| \\
& \leq\left|\mathbb{E}\left[C_{A_{T}}^{1}\left(\alpha_{T-1}-\alpha_{A_{T-1}}^{*}\right)\right]+\gamma \mathbb{E}\left[C_{A_{T}}^{1}\left(M\left(\alpha_{A_{T-1}}^{*} \mid A_{T-1}\right)-M\left(\alpha_{T-1} \mid A_{T-1}\right)\right) \mid A_{T}\right]\right| \\
& \quad \quad+C_{1} \mathbb{E}\left[\left|\alpha_{A_{T-1}}^{*}-\alpha_{A_{T}}^{*}\right| \mid A_{T}\right] \\
& =\left|\mathbb{E}\left[C_{A_{T}}^{1}\left(1-\gamma C_{A_{T-1}}^{1}\right)\left(\alpha_{T-1}-\alpha_{A_{T-1}}^{*}\right) \mid A_{T}\right]\right|+C_{1} C_{2} \gamma \mathbb{E}\left[\left(\alpha_{T-1}-\alpha_{A_{T-1}}^{*}\right)^{2} \mid A_{T}\right] \\
& \quad \quad+C_{1} \mathbb{E}\left[\left|\alpha_{A_{T-1}}^{*}-\alpha_{A_{T}}^{*}\right| \mid A_{T}\right] .
\end{aligned}
$$

The desired result follows by repeating this process inductively. Finally, the last part of the Lemma follows by sending $T \rightarrow \infty$ and applying the result of Lemma A. 3

As a final aside we remark that in the main text we claimed that in the ideal case where $M(p \mid a)=p$ for all $p \in[0,1]$ this bound can be replaced by

$$
\left|\mathbb{E}\left[\operatorname{err}_{1} \mid A_{1}=a\right]-\alpha\right| \leq 2\left(\gamma+\gamma^{-1} \epsilon_{1}\right)
$$

This can be justified by using the fact that in this case we have that for all $p \in[-\gamma, 1+\gamma]$

$$
M(p \mid a)-M\left(\alpha_{a}^{*} \mid a\right)=\left(p-\alpha_{a}^{*}\right)+C_{p, a}^{2}
$$

with $\left|C_{p, a}^{2}\right| \leq \gamma$. The desired result then follows by repeating the argument of Lemma A. 4

## A. 10 Technical lemmas

Proof: [Proof of Lemma A.1]] We assume without loss of generality that $f$ is non-decreasing and $g$ is non-increasing as otherwise one can simply multiply both $f$ and $g$ by -1 .
Let $g^{U}:=\sup \left\{g(y): f(y)>\mathbb{E}[f(Y)]\right.$ and $g^{L}:=\inf \{g(y): f(y) \leq \mathbb{E}[f(Y)]$. By the monotonicity of $f$ and $g$ we clearly have that $g^{L} \geq g^{U}$. Therefore,

$$
\begin{aligned}
& \mathbb{E}[f(Y) g(Y)]-\mathbb{E}[f(Y)] \mathbb{E}[g(Y)]=\mathbb{E}[(f(Y)-\mathbb{E}[f(Y)]) g(Y)] \\
& \leq \mathbb{E}\left[(f(Y)-\mathbb{E}[f(Y)]) g^{L} \mathbb{1}_{f(Y) \leq \mathbb{E}[f(Y)]}\right]+\mathbb{E}\left[(f(Y)-\mathbb{E}[f(Y)]) g^{U} \mathbb{1}_{f(Y)>\mathbb{E}[f(Y)]}\right] \\
& \leq \mathbb{E}\left[(f(Y)-\mathbb{E}[f(Y)]) g^{L} \mathbb{1}_{f(Y) \leq \mathbb{E}[f(Y)]}\right]+\mathbb{E}\left[(f(Y)-\mathbb{E}[f(Y)]) g^{L} \mathbb{1}_{f(Y)>\mathbb{E}[f(Y)]}\right] \\
& =0
\end{aligned}
$$

as desired.
Lemma A. 5 [Hoeffding's Lemma [18]] Let $X$ be a mean 0 random variable such that $X \in[a, b]$ almost surely. Then, for all $\lambda \in \mathbb{R}$

$$
\mathbb{E}[\exp (\lambda X)] \leq \exp \left(\lambda^{2} \frac{(b-a)^{2}}{8}\right)
$$

