A Full Proofs of Presented Results

In this appendix, we present full proofs of all results which are not already complete in the main text.

Proof of Theorem 4 For any $i \in [d]$ and $x \in \mathcal{Z}^d$ define

$$M^{(i)} = \mathbb{E}_{X\sim\mu}[f(X)|\mathcal{B}^c, X^{[i]} = x^{[i]}] - \mathbb{E}_{X\sim\mu}[f(X)|\mathcal{B}^c, X^{[i-1]} = x^{[i-1]}] \,. \tag{26}$$

with the edge case

$$M^{(1)} = \mathbb{E}_{X\sim\mu}[f(X)|\mathcal{B}^c, X_1 = x_1] - \mathbb{E}_{X\sim\mu}[f(X)|\mathcal{B}^c] \tag{27}$$

Due to $\mathbb{E}_{X\sim\mu}[f(X)|\mathcal{B}^c, X = x] = f(x)$ for $x \in \mathcal{B}^c$ we have

$$f - \mathbb{E}_{X\sim\mu}[f(X)|\mathcal{B}^c] = \sum_{i=1}^{d} M^{(i)} \tag{28}$$

Since the conditions $X^{[i]} = x^{[i]}$ generate a nested sequence of σ-algebras, the quantities $K^{(i+1)}f(x) = \mathbb{E}_\mu[f(X)|\mathcal{B}^c, X^{[i]} = x^{[i]}]$ are a Doob martingale and (26) is a martingale difference sequence. In order to bound the moment generating function of f, we will bound every $M^{(i)}$ from above and below and apply the Azuma-Hoeffding theorem [34, Theorem 4.1]. We have

$$\begin{aligned}
M^{(i)} &= \mathbb{E}_\mu[f(X)|\mathcal{B}^c, X^{[i]} = x^{[i]}] - \mathbb{E}_\mu[f(X)|\mathcal{B}^c, X^{[i-1]} = x^{[i-1]}] && \text{(29a)}\\
&= \mathbb{E}_\mu[f(X)|\mathcal{B}^c, X^{[i]} = x^{[i]}] - \mathbb{E}_\mu[\mathbb{E}_\mu[f(X)|\mathcal{B}^c, X^{[i-1]} = x^{[i-1]}, X_i]|\mathcal{B}^c, X^{[i-1]} = x^{[i-1]}] && \text{(29b)}\\
&= \int f(x^{[i]}y^{(i,d]})\mu(dy^{(i,d]}|x^{[i]}, \mathcal{B}^c)\\
&\quad - \int \Big(\int f(x^{[i]}u^{(i,d]})\mu(du^{(i,d]}|x^{[i-1]}, y_i, \mathcal{B}^c)\Big)\mu(dy^{[i,d]}|x^{[i-1]}, \mathcal{B}^c) && \text{(29c)}
\end{aligned}$$

by the tower property of conditional expectations. Because $\mu(dy^{[i,d]}|x^{[i-1]}, \mathcal{B}^c)$ is a probability measure, it holds

$$\int f(x^{[i]}y^{(i,d]})\mu(dy^{(i,d]}|x^{[i]}, \mathcal{B}^c) = \int \Big(\int f(x^{[i-1]}x_i u^{(i,d]})\mu(du^{(i,d]}|x^{[i]}, \mathcal{B}^c)\Big)\mu(dy^{[i,d]}|x^{[i-1]}, \mathcal{B}^c) \tag{30}$$

and we find

$$\begin{aligned}
M^{(i)} = \int \mu(dy^{[i,d]}|x^{[i-1]}, \mathcal{B}^c)\Big(&\int f(x^{[i-1]}x_i u^{(i,d]})\mu(du^{(i,d]}|x^{[i]}, \mathcal{B}^c)\\
&- \int f(x^{[i]}u^{(i,d]})\mu(du^{(i,d]}|x^{[i-1]}, y_i, \mathcal{B}^c)\Big)
\end{aligned} \tag{31}$$

Now bound $A^{(i)} \le M^{(i)} \le B^{(i)}$ almost surely with

$$\begin{aligned}
A^{(i)} = \int \mu(dy^{[i,d]}|x^{[i-1]}, \mathcal{B}^c) \inf_{x_i \in \mathcal{B}_i^c(x^{[i-1]})} \Big(&\int f(x^{[i-1]}x_i u^{(i,d]})\mu(du^{(i,d]}|x^{[i]}, \mathcal{B}^c)\\
&- \int f(x^{[i]}u^{(i,d]})\mu(du^{(i,d]}|x^{[i-1]}, y_i, \mathcal{B}^c)\Big) && \text{(32a)}\\
B^{(i)} = \int \mu(dy^{[i,d]}|x^{[i-1]}, \mathcal{B}^c) \sup_{x_i \in \mathcal{B}_i^c(x^{[i-1]})} \Big(&\int f(x^{[i-1]}x_i u^{(i,d]})\mu(du^{(i,d]}|x^{[i]}, \mathcal{B}^c)\\
&- \int f(x^{[i]}u^{(i,d]})\mu(du^{(i,d]}|x^{[i-1]}, y_i, \mathcal{B}^c)\Big) && \text{(32b)}
\end{aligned}$$

where $\mathcal{B}_i^c(x^{[i-1]})$ contains all $x_i \in \mathcal{Z}$ such that there exist $x^{(i,d]} \in \mathcal{Z}^{d-i}$ with $(x^{[i-1]}, x_i, x^{(i,d]}) \in \mathcal{B}^c$. Because every realization of a random variable conditioned on $\mathcal{B}^c$ is in the set of good inputs, the difference $\|B^{(i)} - A^{(i)}\|_\infty$ can be written as

$$\sup_{x,z \in \mathcal{B}^c, x^{[d]\setminus\{i\}} = z^{[d]\setminus\{i\}}} \int f(x^{[i]}u^{(i,d]})\mu(du^{(i,d]}|x^{[i]}, \mathcal{B}^c) - \int f(z^{[i]}u^{(i,d]})\mu(du^{(i,d]}|z^{[i]}, \mathcal{B}^c) \tag{33}$$

By seeing this expression in terms of oscillation of the kernel action $K^{(i+1)}f$, we find

$$\|B^{(i)} - A^{(i)}\|_\infty \leq \|\rho\|\delta_i(K^{(i+1)}\widetilde{f}) \leq \|\rho\|(V^{(i+1)}\delta(\widetilde{f}))_i = (\Gamma\delta(\widetilde{f}))_i \tag{34}$$

where $\widetilde{f}\colon \mathcal{B}^c \to \mathbb{R}$ is the restriction of f to $\mathcal{B}^c$. The assertion then follows from the Azuma-Hoeffding theorem [34, Theorem 4.1] which we recite as Theorem 10 to make this paper self-contained.

Proof of Proposition 6 For arbitrary $z, z' \in \mathcal{Z}^d$ it holds

$$|f(z) - f(z')| \leq \delta_j(f)\rho(z_j, z_j'), \qquad \forall\, i \in [d] \tag{35}$$

and thus, by summing over all indices we get

$$|f(z) - f(z')| \leq \frac{1}{d}\sum_{j\in[d]} \delta_j(f)\rho(z_j, z_j') \tag{36}$$

Let $x, z \in \mathcal{Z}^d$ with $x^{[d]\setminus\{i\}} = z^{[d]\setminus\{i\}}$ be given for some $i \in [d]$. Recall the action (8) of Markov kernels $K^{(i+1)}$ is an expected value with respect to conditional distributions $\mu^{(i,d]}(dy^{(i,d]}|x^{[i]})$.

Because ν^d has no atoms, $\nu^d|\mathcal{A}^c$ also has no atoms. Therefore, there is a unique KR-rearrangement $\widehat{T}$ with $\widehat{T}_\sharp\nu^d = \nu^d|\mathcal{A}^c$. Then $\widetilde{T} = T \circ \widehat{T}$ is a KR-rearrangement with

$$\widetilde{T}_\sharp\nu^d = \mu|\mathcal{B}^c \tag{37}$$

by Lemma 9 and we have $\widetilde{T}(\widehat{x}) = x$. Lemma 3 implies

$$\mu^{(i,d]}(dy^{(i,d]}|\mathcal{B}^c, x^{[i]}) = \widetilde{T}(\widehat{x}^{[i]}, \cdot)_\sharp\nu^{d-i} \tag{38}$$

An analogous expression holds for the distribution conditioned on z. We have therefore found two transport functions pushing the reference measure to the respective conditional distributions. By Lemma 5, a coupling of the conditional distributions is then given by

$$P^{[i]}_{x,z} = (\widetilde{T}^{(i,d]}(\widehat{x}^{[i]}, \cdot), \widetilde{T}^{(i,d]}(\widehat{z}^{[i]}, \cdot))_\sharp\nu^{d-i} \tag{39}$$

Using a change of measure we find

$$\begin{aligned}
&K^{(i+1)}f(x) - K^{(i+1)}f(z)\\
&\quad= \int P^{[i]}_{x,z}(du^{(i,d]}, dv^{(i,d]})\big(f(x^{[i]}u^{(i,d]}) - f(z^{[i]}v^{(i,d]})\big) &(40)\\
&\quad= \int \big(f(x^{[i]}\widetilde{T}^{(i,d]}(\widehat{x}^{[i]}, \tau)) - f(z^{[i]}\widetilde{T}^{(i,d]}(\widehat{z}^{[i]}, \tau))\big)\nu^{d-i}(\tau) &(41)\\
&\quad\leq \frac{\delta_i(f)}{d}\rho(x_i, z_i) + \sum_{j\in(i,d]}\frac{\delta_j(f)}{d}\int \rho\big(\widetilde{T}^{(i,d]}(\widehat{x}^{[i]}, \tau)_j, \widetilde{T}^{(i,d]}(\widehat{z}^{[i]}, \tau)_j\big)\nu^{d-i}(\tau) &(42)\\
&\quad\leq \frac{\delta_i(f)}{d}\rho(x_i, z_i) + \sum_{j\in(i,d]}\frac{\delta_j(f)}{d}L_{ij}\rho(x_i, z_i) &(43)
\end{aligned}$$

which shows

$$\delta_i(K^{(i+1)}f) \leq \frac{1}{d}\Big(\delta_i(f) + \sum_{j\in(i,d]} L_{ij}\delta_j(f)\Big) \tag{44}$$

for good inputs. We have thus found a Wasserstein matrix $V^{(i+1)}$ for $K^{(i+1)}$ with entries

$$V^{(i+1)}_{ij} = \begin{cases} 0 & \text{if } i > j\\ d^{-1} & \text{if } i = j\\ d^{-1}L_{ij} & \text{if } i < j \end{cases} \tag{45}$$

in row i which shows the assertion.

Proof of Theorem 7 For any hypothesis $h \in \mathcal{H}$, we have

$$\begin{aligned}
\mathcal{R}(h) - \mathcal{R}_m(h, \mathcal{D}_m) &= \mathbb{E}_{Z\sim\mu}[L(h,Z) - \mathcal{R}_m(h,\mathcal{D}_m)] && \text{(46a)}\\
&= \mathbb{E}_{Z\sim\mu}\Big[\Big(L(h,Z) - \mathcal{R}_m(h,\mathcal{D}_m)\Big)\mathbf{1}\{Z \notin \mathcal{B}\}\Big] \\
&\qquad + \mathbb{E}_{Z\sim\mu}\Big[\Big(L(h,Z) - \mathcal{R}_m(h,\mathcal{D}_m)\Big)\mathbf{1}\{Z \in \mathcal{B}\}\Big] && \text{(46b)}\\
&\le \mathbb{E}_{Z\sim\mu}\Big[\Big(L(h,Z) - \mathcal{R}_m(h,\mathcal{D}_m)\Big)\mathbf{1}\{Z \notin \mathcal{B}\}\Big] + \xi && \text{(46c)}\\
&\le \mathbb{E}_{Z\sim\mu|\mathcal{B}^c}[L(h,Z)] - \mathcal{R}_m(h,\mathcal{D}_m) + \xi && \text{(46d)}
\end{aligned}$$

where in (46c) we have used that pointwise loss is in $[0,1]$. Note that the underlying distribution of the risk $\mathcal{R}(h)$ is μ, while $\mathcal{D}_m$ are drawn from $\mu|\mathcal{B}^c$. The above inequality reconciles this such that a concentration argument for the conditional distribution becomes applicable. For any PAC-Bayes posterior distribution ζ and any $\beta > 0$, this implies

$$\begin{aligned}
\mathcal{R}(\zeta) - \mathcal{R}_m(\zeta, \mathcal{D}_m) &= \mathbb{E}_{h\sim\zeta}\mathbb{E}_{Z\sim\mu}[L(h,Z) - \mathcal{R}_m(h,\mathcal{D}_m)] && \text{(47a)}\\
&\le \mathbb{E}_{h\sim\zeta}\big[\mathbb{E}_{Z\sim\mu|\mathcal{B}^c}[L(h,Z)] - \mathcal{R}_m(h,\mathcal{D}_m)\big] + \xi && \text{(47b)}\\
&= \frac{1}{\beta}\mathbb{E}_{h\sim\zeta}\big[\beta(\mathbb{E}_{Z\sim\mu|\mathcal{B}^c}[L(h,Z)] - \mathcal{R}_m(h,\mathcal{D}_m))\big] + \xi && \text{(47c)}\\
&\le \frac{1}{\beta}\log \mathbb{E}_{h\sim\pi}\Big[\exp\big(\beta(\mathbb{E}_{Z\sim\mu|\mathcal{B}^c}[L(h,Z)] - \mathcal{R}_m(h,\mathcal{D}_m))\big)\Big] \\
&\qquad + \frac{1}{\beta}\,\mathrm{KL}[\zeta : \pi] + \xi && \text{(47d)}
\end{aligned}$$

by Donsker and Varadhan's variational formula [2, Lemma 2.2]. Focusing on the first term, we find

$$\begin{aligned}
\exp\big(\beta(\mathbb{E}_{Z\sim\mu|\mathcal{B}^c}[L(h,Z)] - \mathcal{R}_m(h,\mathcal{D}_m))\big) &= \exp\Big(\frac{\beta}{m}\sum_{k\in[m]}\big(\mathbb{E}_{Z\sim\mu|\mathcal{B}^c}[L(h,Z)] - L(h,Z^{(k)})\big)\Big) && \text{(48a)}\\
&= \prod_{k\in[m]}\exp\Big(\frac{\beta}{m}\big(\mathbb{E}_{Z\sim\mu|\mathcal{B}^c}[L(h,Z)] - L(h,Z^{(k)})\big)\Big) && \text{(48b)}
\end{aligned}$$

Each structured datum $Z^{(k)}$ is drawn independently from $\mu|\mathcal{B}^c$. By Proposition 6 there exists a Wasserstein dependency matrix $\Gamma = \frac{\|\rho\|}{d}D$ for $\mu|\mathcal{B}^c$ where D has entries (17). Then

$$\begin{aligned}
&\mathbb{E}_{\mathcal{D}_m\sim(\mu|\mathcal{B}^c)^m}\prod_{k\in[m]}\exp\Big(\frac{\beta}{m}\big(\mathbb{E}_{Z\sim\mu|\mathcal{B}^c}[L(h,Z)] - L(h,Z^{(k)})\big)\Big)\\
&\qquad= \prod_{k\in[m]}\mathbb{E}_{Z^{(k)}\sim(\mu|\mathcal{B}^c)}\exp\Big(\frac{\beta}{m}\big(\mathbb{E}_{Z\sim\mu|\mathcal{B}^c}[L(h,Z)] - L(h,Z^{(k)})\big)\Big) && \text{(49a)}\\
&\qquad= \prod_{k\in[m]}\mathbb{E}_{Z^{(k)}\sim\mu|\mathcal{B}^c}\Big[\exp\Big(\frac{\beta}{m}\big(\mathbb{E}_{Z\sim\mu|\mathcal{B}^c}[L(h,Z)] - L(h,Z^{(k)})\big)\Big)\Big] && \text{(49b)}\\
&\qquad\le \prod_{k\in[m]}\exp\left(\frac{\beta^2}{8m^2}\|\Gamma\delta(\widetilde{L}(h,\cdot))\|_2^2\right) \text{ by Theorem 4} && \text{(49c)}\\
&\qquad= \exp\left(\frac{\beta^2}{8m}\|\Gamma\delta(\widetilde{L}(h,\cdot))\|_2^2\right) && \text{(49d)}\\
&\qquad\le \exp\left(\frac{\beta^2}{8m}\|\Gamma\widetilde{\delta}\|_2^2\right) && \text{(49e)}
\end{aligned}$$

Denote the shorthand

$$U = \mathbb{E}_{\mathcal{D}_m\sim(\mu|\mathcal{B}^c)^m}\left[\exp\big(\beta(\mathbb{E}_{Z\sim\mu|\mathcal{B}^c}[L(h,Z)] - \mathcal{R}_m(h,\mathcal{D}_m))\big)\right] \tag{50}$$

By Markov's inequality it holds

$$\mathbb{P}_{\mathcal{D}_m \sim (\mu|\mathcal{B}^c)^m}\left[\exp\left(\beta(\mathbb{E}_{Z\sim\mu|\mathcal{B}^c}[L(h,Z)] - \mathcal{R}_m(h,\mathcal{D}_m))\right) \geq \frac{1}{\delta}U\right] \leq \delta \tag{51}$$

and combining this with (49) we have

$$\exp\left(\beta(\mathbb{E}_{Z\sim\mu|\mathcal{B}^c}[L(h,Z)] - \mathcal{R}_m(h,\mathcal{D}_m))\right) \leq \frac{1}{\delta}\exp\left(\frac{\beta^2}{8m}\|\Gamma\widetilde{\delta}\|_2^2\right) \tag{52}$$

with probability at least $1-\delta$ over the sample. Using (47) we thus have

$$\begin{aligned}\mathcal{R}(\zeta) - \mathcal{R}_m(\zeta,\mathcal{D}_m) &\leq \frac{1}{\beta}\Big(\log \mathbb{E}_{h\sim\pi}\Big[\frac{1}{\delta}\exp\left(\frac{\beta^2}{8m}\|\Gamma\widetilde{\delta}\|_2^2\right)\Big] + \mathrm{KL}[\zeta:\pi]\Big) + \xi &\text{(53a)}\\ &= \frac{\beta}{8m}\|\Gamma\widetilde{\delta}\|_2^2 + \frac{1}{\beta}\Big(\log\frac{1}{\delta} + \mathrm{KL}[\zeta:\pi]\Big) + \xi &\text{(53b)}\end{aligned}$$

Ideally, we would minimize the right hand side with respect to β. However, this would mean to have β depend on ζ and we thus would not have a uniform bound for all posterior distributions.

Instead, [37] approaches the problem by defining a sequence of constant $(\delta_j,\beta_j)_{j\in\mathbb{N}_0}$ and bounding the probability that the bound does not hold for any sequence element. Since in the opposite (high-probability) case, the bound holds for all sequence elements, an optimal one can subsequently be chosen dependent on the posterior.

For all $j\in\mathbb{N}_0$, define

$$\delta_j = \delta 2^{-(j+1)}, \qquad \beta_j = 2^j\sqrt{\frac{8m\log\frac{1}{\delta}}{\|\Gamma\widetilde{\delta}\|_2^2}} \tag{54}$$

which are independent of ζ. Now consider the event E_j that

$$\exp\left(\beta_j(\mathbb{E}_{X\sim\mu|\mathcal{B}^c}[\ell(h,X)] - \mathcal{R}_m(h,\mathcal{D}_m))\right) \geq \frac{1}{\delta_j}\exp\left(\frac{\beta_j^2}{8m}\|\Gamma\widetilde{\delta}\|_2^2\right) \tag{55}$$

By the above argument leading up to (52), the probability for E_j under a random sample of the conditioned data distribution $\mu|\mathcal{B}^c$ is at most δ_j. Therefore, the probability that any E_j occurs is bounded by

$$\mathbb{P}\big(\bigcup_{j\in\mathbb{N}_0} E_j\big) \leq \sum_{j\in\mathbb{N}_0}\mathbb{P}(E_j) \leq \sum_{j\in\mathbb{N}_0}\delta_j = \delta \tag{56}$$

Thus, for all posteriors ζ with probability at least $1-\delta$ none of the events (55) occurs. We may therefore select an index j dependent on ζ to obtain a sharper risk certificate which still holds with probability at least $1-\delta$ over the sample conditioned on the good set. For a fixed posterior ζ, the optimizer of (53b) would be

$$\beta^* = \frac{1}{\|\Gamma\widetilde{\delta}\|_2}\sqrt{8m(\log\frac{1}{\delta} + \mathrm{KL}[\zeta:\pi])} \tag{57}$$

Equating this to (55) and rounding down to the nearest integer gives

$$j^* = \left\lfloor \frac{1}{2}\log_2\left(1 + \frac{\mathrm{KL}[\zeta:\pi]}{\log\frac{1}{\delta}}\right)\right\rfloor \tag{58}$$

Denote this number before rounding by r, i.e. $j^* = \lfloor r \rfloor$. For any real number r it holds $r-1 \leq \lfloor r \rfloor \leq r$. Therefore

$$\frac{1}{2}\sqrt{1 + \frac{\mathrm{KL}[\zeta:\pi]}{\log\frac{1}{\delta}}} = 2^{r-1} \leq 2^{j^*} \leq 2^r = \sqrt{1 + \frac{\mathrm{KL}[\zeta:\pi]}{\log\frac{1}{\delta}}} \tag{59}$$

which gives the following bounds on u_{j^*}

$$\frac{1}{2}\sqrt{\frac{8m(\log\frac{1}{\delta} + \mathrm{KL}[\zeta:\pi])}{\|\Gamma\widetilde{\delta}\|_2^2}} \leq u_{j^*} \leq \sqrt{\frac{8m(\log\frac{1}{\delta} + \mathrm{KL}[\zeta:\pi])}{\|\Gamma\widetilde{\delta}\|_2^2}} \tag{60}$$

Likewise, we bound

$$\mathrm{KL}[\zeta : \pi] + \log \frac{1}{\delta_{j^*}} = \mathrm{KL}[\zeta : \pi] + \log \frac{2}{\delta} + j^* \log 2 \tag{61a}$$

$$\leq \mathrm{KL}[\zeta : \pi] + \log \frac{2}{\delta} + \frac{\log 2}{2} \log_2 \Big(1 + \frac{\mathrm{KL}[\zeta : \pi]}{\log \frac{1}{\delta}} \Big) - \log 2 \tag{61b}$$

$$= \mathrm{KL}[\zeta : \pi] + \log \frac{1}{\delta} + \frac{1}{2} \log \Big(1 + \frac{\mathrm{KL}[\zeta : \pi]}{\log \frac{1}{\delta}} \Big) \tag{61c}$$

$$= \mathrm{KL}[\zeta : \pi] + \log \frac{1}{\delta} + \frac{1}{2} \log \Big(\log \frac{1}{\delta} + \mathrm{KL}[\zeta : \pi] \Big) - \frac{1}{2} \log \log \frac{1}{\delta} \tag{61d}$$

The assumption $\delta \leq \exp(-e^{-1})$ yields $-\log \log \frac{1}{\delta} \leq 1$ and because $x + 1 \leq \exp(x)$ for all $x \in \mathbb{R}$, we find

$$\mathrm{KL}[\zeta : \pi] + \log \frac{1}{\delta_{j^*}} \leq \mathrm{KL}[\zeta : \pi] + \log \frac{1}{\delta} + \frac{1}{2} \Big(\log \big(\log \frac{1}{\delta} + \mathrm{KL}[\zeta : \pi] \big) + 1 \Big) \tag{62a}$$

$$\leq \mathrm{KL}[\zeta : \pi] + \log \frac{1}{\delta} + \frac{1}{2} \big(\log \frac{1}{\delta} + \mathrm{KL}[\zeta : \pi] \big) \tag{62b}$$

$$= \frac{3}{2} \big(\log \frac{1}{\delta} + \mathrm{KL}[\zeta : \pi] \big) \tag{62c}$$

We can now use the bounds (62c) and (60) in (53b) to bound the expected generalization error

$$\mathcal{R}(\zeta) - \mathcal{R}_m(\zeta, \mathcal{D}_m) \leq \frac{u_{j^*}}{8m} \|\Gamma \widetilde{\delta}\|_2^2 + \frac{1}{u_{j^*}} \Big(\log \frac{1}{\delta_{j^*}} + \mathrm{KL}[\zeta : \pi] \Big) + \xi \tag{63a}$$

$$\leq \frac{u_{j^*}}{8m} \|\Gamma \widetilde{\delta}\|_2^2 + \frac{3}{2u_{j^*}} \Big(\log \frac{1}{\delta} + \mathrm{KL}[\zeta : \pi] \Big) + \xi \tag{63b}$$

$$\leq \frac{1}{2} \|\Gamma \widetilde{\delta}\|_2 \sqrt{\frac{\log \frac{1}{\delta} + \mathrm{KL}[\zeta : \pi]}{2m}} + \frac{3}{2} \|\Gamma \widetilde{\delta}\|_2 \sqrt{\frac{\log \frac{1}{\delta} + \mathrm{KL}[\zeta : \pi]}{2m}} + \xi \tag{63c}$$

$$= 2 \|\Gamma \widetilde{\delta}\|_2 \sqrt{\frac{\log \frac{1}{\delta} + \mathrm{KL}[\zeta : \pi]}{2m}} + \xi \tag{63d}$$

Note that β^* would attain the optimal value

$$\mathcal{R}(\zeta) - \mathcal{R}_m(\zeta, \mathcal{D}_m) \leq \|\Gamma \widetilde{\delta}\|_2 \sqrt{\frac{\mathrm{KL}[\zeta : \pi] + \log \frac{1}{\delta}}{2m}} + \xi \tag{64}$$

which only differs from the above uniform bound by a factor of two. Finally, recall $\Gamma = \frac{\|\rho\|}{d} D$ where D has entries (17).

B Additional Lemmata

Lemma 9. *Let $T \colon \Omega \to \Omega$ be a measurable function on a measurable space (Ω, Σ) and let ν, μ be measures on Ω with $T_\sharp \nu = \mu$. Let $B \in \Sigma$ be a fixed set with $\mu(B) > 0$ and $A = T^{-1}(B)$ its preimage under T. Then*

$$T_\sharp(\nu | A) = \mu | B\,. \tag{65}$$

Proof. Let $S \in \Sigma$ be arbitrary and let $\widetilde{\mu} = T_\sharp(\nu | A)$. Then

$$\widetilde{\mu}(S) = (\nu | A)(T^{-1}(S)) = \frac{\nu(T^{-1}(S) \cap A)}{\nu(A)} \tag{66}$$

as well as

$$(\mu | B)(S) = \frac{\mu(S \cap B)}{\mu(B)} = \frac{\nu(T^{-1}(S \cap B))}{\nu(A)} \tag{67}$$

Note that

$$x \in T^{-1}(S) \cap T^{-1}(B) \;\Leftrightarrow\; T(x) \in S \wedge T(x) \in B \;\Leftrightarrow\; T(x) \in S \cap B \;\Leftrightarrow\; x \in T^{-1}(S \cap B) \tag{68}$$

thus $T^{-1}(S) \cap T^{-1}(B) = T^{-1}(S \cap B)$ and consequently $\widetilde{\mu}(S) = (\mu | B)(S)$. Since S was arbitrary, this shows the assertion. $\square$

The following theorem exists in various forms in the literature. To make this paper self-contained, we recite the version in [34] which is used to bound moment-generating functions in Proposition 4. Note that we only use the MGF bound (69) in our analysis. However, the concentration inequality (70) also holds analogously under the assumptions of Proposition 4 which may be of independent interest.

Theorem 10 (**Azuma-Hoeffding [34, Theorem 4.1]**). *Let $(M^{(i)})_{i\in[m]}$ be a martingale difference sequence with respect to a filtration $(\Sigma_i)_{i\in[m]}$ of sigma algebras. Suppose that for each $i \in [m]$ there exist Σ_{i-1}-measurable random variables $A^{(i)}$, $B^{(i)}$ such that $A^{(i)} \leq M^{(i)} \leq B^{(i)}$ almost surely. Then for all $\lambda \in \mathbb{R}$ it holds that*

$$\mathbb{E}\Big[\exp\big(\lambda \sum_{i\in[m]} M^{(i)}\big)\Big] \leq \exp\Big(\frac{\lambda^2}{8}\sum_{i\in[m]} \|B^{(i)} - A^{(i)}\|_\infty^2\Big) \tag{69}$$

and consequently, for any $t \geq 0$

$$\mathbb{P}\Big(\Big|\sum_{i\in[m]} M^{(i)}\Big| \geq t\Big) \leq 2\exp\Big(-\frac{2t^2}{\sum_{i\in[m]}\|B^{(i)} - A^{(i)}\|_\infty^2}\Big)\,. \tag{70}$$