Supplement to " Optimal Regularized Online Allocation by Adaptive Re-Solving "

12. Proofs of Main Results

12.1. Proof of Lemma 1

We prove the bound of the deterministic optimal solution. Consider $\Omega'_\mu = \{-\partial r(a)|a \in \mathcal{Z}\}$. The bounded subgradient in Assumption 2 suggests that the dual variable region Ω'_μ we defined is bounded by G. We explain this definition by the optimal conditions of stochastic programming. Note that for problem (2.8), μ is unconstrained. The optimal condition suggests that we can take

$$\vartheta^*(-\mu^*) = \mathbb{E} b_t g_t(b_t^\top(\lambda^* + \mu^*)),$$

if we assume the Fubini theorem holds. Then, by the Fenchel conjugate function, we have $\mu^* = -\vartheta(\mathbb{E} b_t \tilde{x}_t)$. This shows that by defining Ω'_μ, we indeed define the possible region that contains optimal solution μ^*, i.e., $\mu^* \in \Omega'_\mu \subseteq \Omega_\mu$. Thus we have $\|\mu^*\|_\infty \le G$.

For the second bound of $\|\lambda^*\|_\infty$, we only need to check that $d^\top \lambda^* \le 2(\bar{f} + \bar{r})$ always holds. Otherwise if $d^\top \lambda^* > 2(\bar{f} + \bar{r})$, we have

$$\begin{aligned} D(\boldsymbol{\lambda}^*, d) &= \mathbb{E}\sup_x \{f_t(x) - (\lambda^* + \mu^*)^\top b_t x_t\} + \sup_a \{r(a) + a^\top \mu^*\} + d^\top \lambda^* \ge \mathbb{E} f_t(0) + r(0) + d^\top \lambda^* \\ &> (\bar{f} + \bar{r}) \ge D(\mathbf{0}, d), \end{aligned}$$

which suggests that $\boldsymbol{\lambda}^*$ is not optimal. Thus we have $d^\top \lambda^* \le 2(\bar{f} + \bar{r})$, i.e., $\|\lambda^*\|_\infty \le \frac{2(\bar{f}+\bar{r})}{\underline{d}}$. The bound of empirical optimal solution $\boldsymbol{\lambda}_T^*$ follows exactly the same argument.

The following proposition on the growth of second order term $s(\nu, d)$ will be useful in the development of our theory.

Proposition 4. Under Assumptions 1-3, the second order objective function $s(\nu, d)$ in stochastic program satisfies the following growth condition:

$$\underline{\mathcal{L}}_s \|\nu - \nu^*\|_2^2 \le s(\nu) \le \bar{\mathcal{L}}_s \|\nu - \nu^*\|_2^2, \tag{12.1}$$

where the constant $\underline{\mathcal{L}}_s := \sigma_{\min} \underline{\mathcal{L}}_f / 2$, $\bar{\mathcal{L}}_s := \bar{b}^2 \bar{\mathcal{L}}_f / 2$.

Proof: By definition, we have

$$\begin{aligned} s(\nu) &= D(\nu, \mu^*, d) - D(\nu^*, \mu^*, d) - \nabla_\nu D(\nu^*, \mu^*, d)^\top (\nu - \nu^*) \\ &= \int_0^1 \left[\nabla D_\nu(z(\nu - \nu^*) + \nu^*, \mu^*, d) - \nabla D_\nu(\nu^*, \mu^*, d)\right]^\top (\nu - \nu^*) dz, \end{aligned}$$

where $\nabla_\nu D(\nu, d) = \mathbb{E} b_t f_t^*(b_t^\top v)$. Then for any z, we have

$$
\begin{aligned}
&\left[\nabla D_\nu(z(\nu-\nu^*)+\nu^*,\mu^*,d)-\nabla D_\nu(\nu^*,\mu^*,d)\right]^\top(\nu-\nu^*)\\
\le&\left\|\mathbb{E}b_t g_t(b_t^\top(z(\mu+\lambda-\mu^*-\lambda^*)+\mu^*+\lambda^*)-\mathbb{E}b_t g_t(b_t^\top(\mu^*+\lambda^*))\right\|_2\|\nu-\nu^*\|_2\\
\le&\left\|z\bar{\mathcal{L}}_f\bar{b}\mathbb{E}\left[b_t^\top(\mu+\lambda-\mu^*-\lambda^*)\right]\right\|_2\|\nu-\nu^*\|\\
\le& z\bar{\mathcal{L}}_f\bar{b}^2\|\nu-\nu^*\|^2,
\end{aligned}
$$

where the second inequality is by Assumption 3.1 when conditioned on b_t. By the integral of z we have $s(\nu)\le\bar{\mathcal{L}}_s\|\nu-\nu^*\|_2^2$. For the other direction, it is also clear that

$$
\begin{aligned}
&\left[\nabla D_\nu(z(\nu-\nu^*)+\nu^*,\mu^*,d)-\nabla D_\nu(\nu^*,\mu^*,d)\right]^\top(\nu-\nu^*)\\
=&\left(\mathbb{E}b_t g_t(b_t^\top(z(\mu+\lambda-\mu^*-\lambda^*)+\mu^*+\lambda^*)-\mathbb{E}b_t g_t(b_t^\top(\mu^*+\lambda^*))\right)^\top(\nu-\nu^*)\\
=&\mathbb{E}\left[\mathbb{E}\left[\langle\nabla f_t^*(b_t^\top(z(\mu+\lambda-\mu^*-\lambda^*)+\mu^*+\lambda^*))-\nabla f_t^*(b_t^\top(\mu^*+\lambda^*)),b_t^\top(\mu+\lambda-\mu^*-\lambda^*)\rangle\right]\middle|b_t\right]\\
\ge& z\underline{\mathcal{L}}_f\mathbb{E}\left\|b_t^\top(\mu+\lambda-\mu^*-\lambda^*)\right\|_2^2\ge z\underline{\mathcal{L}}_f\sigma_{\min}\|\mu+\lambda-\mu^*-\lambda^*\|_2^2.
\end{aligned}
$$

12.2. Proof of Proposition 1: Empirical Risk Minimization

In this proof, we generalize our discussion to a broader setting: we consider the convergence of the classical ERM method. In many areas of statistics and machine learning research, we aim to solve the following problem, named Empirical Risk Minimization (ERM): given T empirical convex risk functions $\ell(\boldsymbol{\lambda},\xi_t):\mathbb{R}^m\to\mathbb{R},t\in[T]$ where ξ_t are i.i.d realizations from an unknown distribution, with its population version $\mathfrak{L}(\boldsymbol{\lambda})=\mathbb{E}_\xi\ell(\boldsymbol{\lambda},\xi_t)$, we seek to find a good parameter $\widehat{\boldsymbol{\lambda}}$ by minimizing the empirical risk

$$
\widehat{\boldsymbol{\lambda}}=\arg\min_{\lambda\in\mathbb{R}^m}\bar{\ell}(\boldsymbol{\lambda},\{\xi_t\}_{t=1}^T)=\arg\min_{\lambda\in\mathbb{R}^m}\frac{1}{T}\sum_{t=1}^T\ell(\boldsymbol{\lambda},\xi_t)
$$

as a proxy of the parameter that minimizes the population risk $\boldsymbol{\lambda}^*=\arg\min_{\boldsymbol{\lambda}}\mathfrak{L}(\boldsymbol{\lambda})$. In statistics, such $\widehat{\boldsymbol{\lambda}}$ is also called M-estimator.

The following approach will show that, under second-order-growth condition of $\mathfrak{L}(\boldsymbol{\lambda})$, the ERM method provides a estimate $\widehat{\boldsymbol{\lambda}}$ with optimal convergence rate $\mathbb{E}\left\|\widehat{\boldsymbol{\lambda}}-\boldsymbol{\lambda}^*\right\|^2=O(\frac{1}{T})$

Assumption 6 **(Lipschitz continuity)**. *Suppose the subgradient of each $\ell(\boldsymbol{\lambda},\xi_t)$ satisfies*

$$
\|\nabla\ell(\boldsymbol{\lambda},\xi_t)\|\le L.
$$

Assumption 7 **(Second order growth)**. *The population risk satisfies the following second order growth*

$$
\langle\nabla\mathfrak{L}(\boldsymbol{\lambda})-\nabla\mathfrak{L}(\boldsymbol{\lambda}^*),\boldsymbol{\lambda}-\boldsymbol{\lambda}^*\rangle\ge\underline{\mathcal{L}}_\ell\|\boldsymbol{\lambda}-\boldsymbol{\lambda}^*\|_2^2.
$$

Assumption 8 **(Smoothness of the first moment)**. *For any $\boldsymbol{\lambda}\in\mathbb{R}^m$, with $\delta=\|\boldsymbol{\lambda}-\boldsymbol{\lambda}^*\|$, we have*

$$
\mathbb{E}\sup_{\boldsymbol{\lambda}'\in\mathbb{B}(\boldsymbol{\lambda}^*,\delta)}\left\|\nabla\ell(\boldsymbol{\lambda}',\xi_t)-\nabla\ell(\boldsymbol{\lambda}^*,\xi_t)\right\|\le M\delta,
$$

where $\mathbb{B}(\boldsymbol{\lambda}^*,r)=\left\{\boldsymbol{\lambda}':\left\|\boldsymbol{\lambda}'-\boldsymbol{\lambda}^*\right\|\le r\right\}$.

To investigate the convergence of $\widehat{\boldsymbol{\lambda}}$, one classical approach is to compute the statistical complexity of function group $\{\ell(\boldsymbol{\lambda},\xi_t),\boldsymbol{\lambda}\in\Theta\}$ to get a uniform bound of $\sup_{\boldsymbol{\lambda}\in\Theta}\left[\bar{\ell}(\boldsymbol{\lambda},\{\xi_t\}_{t=1}^T)-\mathfrak{L}(\boldsymbol{\lambda})\right]$. However, this approach may fail to reach the optimal convergence rate because it neglects the second-order information. Instead, we consider the second-order part of losses and improve our analyses by a localized argument near the optimal solution $\boldsymbol{\lambda}^*$. Equipped with localized Rademacher complexity, which shares a similar idea as Bartlett et al. (2005), we are able to derive a sharp local probabilistic bound of $\|\boldsymbol{\lambda}-\boldsymbol{\lambda}^*\|$. To fix this idea, we define the second-order part of our loss function:

$$s(\boldsymbol{\lambda},\xi_t)=\ell(\boldsymbol{\lambda},\xi_t)-\langle\nabla\ell(\boldsymbol{\lambda}^*,\xi_t),\boldsymbol{\lambda}-\boldsymbol{\lambda}^*\rangle-\ell(\boldsymbol{\lambda}^*,\xi_t),\ \bar{s}(\boldsymbol{\lambda})=\frac{1}{T}\sum_{t=1}^T s(\boldsymbol{\lambda},\xi_t),$$

with its population version $S(\boldsymbol{\lambda})=\mathbb{E}s(\boldsymbol{\lambda},\xi_t)=\mathfrak{L}(\boldsymbol{\lambda})-\langle\nabla\mathfrak{L}(\boldsymbol{\lambda}),\boldsymbol{\lambda}-\boldsymbol{\lambda}^*\rangle-\mathfrak{L}(\boldsymbol{\lambda}^*)$. Define localized Rademacher complexity of s within a small neighbourhood of $\boldsymbol{\lambda}^*$ as

$$\mathcal{R}_\varepsilon=\mathbb{E}_\xi\mathbb{E}_\sigma\left[\sup_{\boldsymbol{\lambda}\in\mathbb{B}(\boldsymbol{\lambda}^*,\varepsilon)}\frac{1}{T}\sum_{t=1}^T\sigma_t s(\boldsymbol{\lambda},\xi_t)\right],$$

where σ_t are independent Rademacher random variables. We have the following result:

Proposition 5. Under Assumption 6-8, the following inequality holds

$$\mathcal{R}_\varepsilon\le\sqrt{2m\log(3K)}\frac{2L\varepsilon}{\sqrt{T}}+\frac{M\varepsilon^2}{K},$$

for any constant $K>0$. Consequently, if $\varepsilon\ge\frac{64\sqrt{2}L}{\sqrt{T}\underline{\mathcal{L}}_\ell}\sqrt{\log\frac{100M}{\underline{\mathcal{L}}_\ell}}$, we have the following probabilistic bound:

$$\mathbb{P}\left(\left\|\widehat{\boldsymbol{\lambda}}-\boldsymbol{\lambda}^*\right\|\ge\varepsilon\right)\le m\exp\left(-\frac{T\underline{\mathcal{L}}_\ell^2\varepsilon^2}{8L^2m}\right)+\exp\left(-\frac{T\underline{\mathcal{L}}_\ell^2\varepsilon^2}{5000L^2}\right).$$

Proof Define $\mathcal{K}$ as a $\frac{\varepsilon}{K}$-cover of the set $\boldsymbol{\lambda}\in\mathbb{B}(\boldsymbol{\lambda}^*,\varepsilon)$, then by the covering number of a ball, it is valid that $\log|\mathcal{K}|\le m\log(3K)$. Define a projection $\mathcal{K}(\boldsymbol{\lambda})$ that project each $\boldsymbol{\lambda}\in\mathbb{B}(\boldsymbol{\lambda}^*,\varepsilon)$ onto the closest element in the cover $\mathcal{K}$. By Assumption 6, we have a uniform bound $|s(\boldsymbol{\lambda},\xi_t)|\le 2L\varepsilon$. Then it follows that:

$$\begin{aligned}\mathcal{R}_\varepsilon=&\mathbb{E}_\xi\mathbb{E}_\sigma\left[\sup_{\boldsymbol{\lambda}\in\mathbb{B}(\boldsymbol{\lambda}^*,\varepsilon)}\frac{1}{T}\sum_{t=1}^T\sigma_t s(\boldsymbol{\lambda},\xi_t)\right]\\ \le&\mathbb{E}_\xi\left[\mathbb{E}_\sigma\sup_{\boldsymbol{\lambda}\in\mathcal{K}}\frac{1}{T}\sum_{t=1}^T\sigma_t s(\boldsymbol{\lambda},\xi_t)+\mathbb{E}_\sigma\sup_{\boldsymbol{\lambda}\in\mathbb{B}(\boldsymbol{\lambda}^*,\varepsilon),\boldsymbol{\lambda}'=\mathcal{K}(\boldsymbol{\lambda})}\frac{1}{T}\sum_{t=1}^T\sigma_t(s(\boldsymbol{\lambda},\xi_t)-s(\boldsymbol{\lambda}',\xi_t))\right]\\ \le&\sqrt{2m\log(3K)}\frac{2L\varepsilon}{\sqrt{T}}+\mathbb{E}_\xi\left[\mathbb{E}_\sigma\sup_{\boldsymbol{\lambda}\in\mathbb{B}(\boldsymbol{\lambda}^*,\varepsilon),\boldsymbol{\lambda}'=\mathcal{K}(\boldsymbol{\lambda})}\frac{1}{T}\sum_{t=1}^T\sigma_t(s(\boldsymbol{\lambda},\xi_t)-s(\boldsymbol{\lambda}',\xi_t))\right],\end{aligned}$$

where the second inequality is by Massart's finite class lemma. We focus on controlling the second term by computing the first order moment of $\left|s(\boldsymbol{\lambda},\xi_t)-s(\boldsymbol{\lambda}',\xi_t)\right|$:

$$
\begin{aligned}
&\mathbb{E} \sup_{\boldsymbol{\lambda}\in\mathbb{B}(\boldsymbol{\lambda}^*,\varepsilon),\boldsymbol{\lambda}'=\mathcal{K}(\boldsymbol{\lambda})} \left|s(\boldsymbol{\lambda},\xi_t)-s(\boldsymbol{\lambda}',\xi_t)\right| = \mathbb{E} \sup_{\boldsymbol{\lambda}\in\mathbb{B}(\boldsymbol{\lambda}^*,\varepsilon),\boldsymbol{\lambda}'=\mathcal{K}(\boldsymbol{\lambda})} \left|\ell(\boldsymbol{\lambda},\xi_t)-\ell(\boldsymbol{\lambda}',\xi_t)-\langle\nabla\ell(\boldsymbol{\lambda}^*,\xi_t),\boldsymbol{\lambda}-\boldsymbol{\lambda}'\rangle\right| \\
&= \mathbb{E} \sup_{\boldsymbol{\lambda}\in\mathbb{B}(\boldsymbol{\lambda}^*,\varepsilon),\boldsymbol{\lambda}'=\mathcal{K}(\boldsymbol{\lambda})} \left(\int_0^1 v^\top(\nabla\ell(\boldsymbol{\lambda}'+vz,\xi_t)-\nabla\ell(\boldsymbol{\lambda}^*,\xi_t))dz\right), \text{ where } v=\boldsymbol{\lambda}-\boldsymbol{\lambda}' \\
&\le \sup\|v\|\cdot\mathbb{E} \sup_{\boldsymbol{\lambda}\in\mathbb{B}(\boldsymbol{\lambda}^*,\varepsilon),\boldsymbol{\lambda}'=\mathcal{K}(\boldsymbol{\lambda})} \int_0^1 \left\|\nabla\ell(\boldsymbol{\lambda}'+vz,\xi_t)-\nabla\ell(\boldsymbol{\lambda}^*,\xi_t)\right\| dz \\
&\le \frac{\varepsilon}{K}\int_0^1 \mathbb{E} \sup_{\boldsymbol{\lambda}\in\mathbb{B}(\boldsymbol{\lambda}^*,\varepsilon),\boldsymbol{\lambda}'=\mathcal{K}(\boldsymbol{\lambda})} \left\|\nabla\ell(\boldsymbol{\lambda}'+vz,\xi_t)-\nabla\ell(\boldsymbol{\lambda}^*,\xi_t)\right\| dz \le \frac{M\varepsilon^2}{K},
\end{aligned}
$$

where the last inequality we use the Assumption 8. Then it follows that

$$
\begin{aligned}
&\mathbb{E}_\xi\left[E_\sigma \sup_{\boldsymbol{\lambda}\in\mathbb{B},\boldsymbol{\lambda}'=\mathcal{K}(\boldsymbol{\lambda})} \frac{1}{T}\sum_{t=1}^T \sigma_t(s(\boldsymbol{\lambda},\xi_t)-s(\boldsymbol{\lambda}',\xi_t))\right] \le \frac{\sum_{t=1}^T \mathbb{E}_\xi \sup_{\boldsymbol{\lambda}\in\mathbb{B},\boldsymbol{\lambda}'=\mathcal{K}(\boldsymbol{\lambda})}\left|s(\boldsymbol{\lambda},\xi_t)-s(\boldsymbol{\lambda}',\xi_t)\right|}{T} \\
&\le \frac{M\varepsilon^2}{K},
\end{aligned}
$$

which proves the first statement. We then prove the second statement by choosing a suitable K.

If $\widehat{\boldsymbol{\lambda}}$ which minimizes the empirical risk satisfies $\left\|\widehat{\boldsymbol{\lambda}}-\boldsymbol{\lambda}^*\right\|\ge\varepsilon$, then by the convexity of $\bar{\ell}$, there exists a $\boldsymbol{\lambda}\in\mathbb{B}(\boldsymbol{\lambda}^*,\varepsilon)$ such that $\bar{\ell}(\boldsymbol{\lambda})-\bar{\ell}(\boldsymbol{\lambda}^*)\le 0$. Together with the second order growth Assumption 7, we have

$$
\begin{aligned}
\bar{s}(\boldsymbol{\lambda})-S(\boldsymbol{\lambda}) &= \bar{\ell}(\boldsymbol{\lambda})-\bar{\ell}(\boldsymbol{\lambda}^*)-(\mathfrak{L}(\boldsymbol{\lambda})-\mathfrak{L}(\boldsymbol{\lambda}^*))-\langle\nabla\bar{\ell}(\boldsymbol{\lambda}^*)-\nabla\mathfrak{L}(\boldsymbol{\lambda}^*),\boldsymbol{\lambda}-\boldsymbol{\lambda}^*\rangle \\
&\le -\frac{\mathcal{L}_\ell}{2}\varepsilon^2+\left\|\nabla\bar{\ell}(\boldsymbol{\lambda}^*)-\nabla\mathfrak{L}(\boldsymbol{\lambda}^*)\right\|\varepsilon.
\end{aligned}
$$

By Hoeffding's concentration inequality, we have

$$
\mathbb{P}\left(\left\|\nabla\bar{\ell}(\boldsymbol{\lambda}^*)-\nabla\mathfrak{L}(\boldsymbol{\lambda}^*)\right\|\ge\frac{\mathcal{L}_\ell}{4}\varepsilon\right)\le m\exp\left(-\frac{T\mathcal{L}_\ell^2\varepsilon^2}{8L^2m}\right).
$$

Define the event that inequality $\left\|\nabla\bar{\ell}(\boldsymbol{\lambda}^*)-\nabla\mathfrak{L}(\boldsymbol{\lambda}^*)\right\|\ge\frac{\mathcal{L}_\ell}{4}\varepsilon$ holds as $\mathcal{E}_1$. Then under $\{\left\|\widehat{\boldsymbol{\lambda}}-\boldsymbol{\lambda}^*\right\|\ge\varepsilon\}\cap\mathcal{E}_1^c$, we have

$$
\sup_{\boldsymbol{\lambda}\in\mathbb{B}(\boldsymbol{\lambda}^*,\varepsilon)}|\bar{s}(\boldsymbol{\lambda})-S(\boldsymbol{\lambda})|\ge\frac{\mathcal{L}_\ell}{4}\varepsilon^2.
$$

Choosing $K=\frac{32M}{\mathcal{L}_\ell}$. When $\varepsilon\ge\frac{64\sqrt{2}L}{\sqrt{T}\mathcal{L}_\ell}\sqrt{\log\frac{100M}{\mathcal{L}_\ell}}$, we have

$$
2\mathcal{R}_\varepsilon\le\frac{\mathcal{L}_\ell}{8}\varepsilon^2,
$$

thus we have the following inequality:

$$
\sup_{\boldsymbol{\lambda}\in\mathbb{B}(\boldsymbol{\lambda}^*,\varepsilon)}|\bar{s}(\boldsymbol{\lambda})-S(\boldsymbol{\lambda})|\ge 2\mathcal{R}_\varepsilon+\frac{\mathcal{L}_\ell}{8}\varepsilon^2.
$$

By the convergence theory of empirical process (Koltchinskii 2011, Boucheron et al. 2005),

$$
\mathbb{P}\left(\sup_{\boldsymbol{\lambda}\in\mathbb{B}(\boldsymbol{\lambda}^*,\varepsilon)}|\bar{s}(\boldsymbol{\lambda})-S(\boldsymbol{\lambda})|\ge 2\mathcal{R}_\varepsilon+\frac{6L\varepsilon z}{\sqrt{T}}\right)\le\exp\left(-\frac{z^2}{2}\right),
$$

thus we conclude that $\mathbb{P}(\{\left\|\widehat{\boldsymbol{\lambda}}-\boldsymbol{\lambda}^*\right\|\geq\varepsilon\}\cap\mathcal{E}_1^c)\leq\exp\left(-\frac{T\underline{\mathcal{L}}_\ell^2\varepsilon^2}{5000L^2}\right)$, i.e.,

$$\mathbb{P}\left(\left\|\widehat{\boldsymbol{\lambda}}-\boldsymbol{\lambda}^*\right\|\geq\varepsilon\right)\leq m\exp\left(-\frac{T\underline{\mathcal{L}}_\ell^2\varepsilon^2}{8L^2m}\right)+\exp\left(-\frac{T\underline{\mathcal{L}}_\ell^2\varepsilon^2}{5000L^2}\right).$$

THEOREM 8. *The following bound holds for the convergence rate of $\widehat{\boldsymbol{\lambda}}$:*

$$\mathbb{E}\left\|\widehat{\boldsymbol{\lambda}}-\boldsymbol{\lambda}^*\right\|^2\leq\left(\frac{512L^2}{\underline{\mathcal{L}}_\ell^2}\log\frac{100M}{\underline{\mathcal{L}}_\ell}+\frac{8m^2L^2+5000L^2}{\underline{\mathcal{L}}_\ell^2}\right)\frac{1}{T}.$$

Proof This is a direct consequence of Proposition 5. By the integral formula of expectation, we have

$$\begin{aligned}
&\mathbb{E}\left\|\widehat{\boldsymbol{\lambda}}-\boldsymbol{\lambda}^*\right\|^2=\int_0^\infty\mathbb{P}(\left\|\widehat{\boldsymbol{\lambda}}-\boldsymbol{\lambda}^*\right\|\geq\sqrt{z})dz\\
&\leq\left(\frac{512L^2}{\underline{\mathcal{L}}_\ell^2}\log\frac{100M}{\underline{\mathcal{L}}_\ell}\right)\frac{1}{T}+\int_c^\infty\mathbb{P}(\left\|\widehat{\boldsymbol{\lambda}}-\boldsymbol{\lambda}^*\right\|^2\geq z)dz,\ \text{where } c=\left(\frac{512L^2}{\underline{\mathcal{L}}_\ell^2}\log\frac{100M}{\underline{\mathcal{L}}_\ell}\right)\frac{1}{T}\\
&\leq\left(\frac{512L^2}{\underline{\mathcal{L}}_\ell^2}\log\frac{100M}{\underline{\mathcal{L}}_\ell}\right)\frac{1}{T}+\int_0^\infty m\exp\left(-\frac{T\underline{\mathcal{L}}_\ell^2z}{8L^2m}\right)+\exp\left(-\frac{T\underline{\mathcal{L}}_\ell^2z}{5000L^2}\right)dz\\
&\leq\left(\frac{512L^2}{\underline{\mathcal{L}}_\ell^2}\log\frac{100M}{\underline{\mathcal{L}}_\ell}+\frac{8m^2L^2+5000L^2}{\underline{\mathcal{L}}_\ell^2}\right)\frac{1}{T},
\end{aligned}$$

which finishes the proof.

By simply equating $\ell(\boldsymbol{\lambda},\xi_t)$ with Fenchel conjugate $f_t^*(\boldsymbol{\lambda})$ in online convex allocation, we are able to prove the optimal dual convergence rate $O(\frac{1}{T})$. Notice that, here, our Assumption 4 is equivalent to the Assumption 8 we used in the proof.

12.3. Proof of Proposition 2

For any given $\varepsilon>0$, we define the neighbourhood of ν^* for given ε as

$$\Omega_\nu(\varepsilon):=\{\nu:\|\nu-\nu^*\|_\infty\leq 4H\varepsilon\}.$$

We then construct a good event $\mathcal{E}(\varepsilon)$ with pro only depends on ε that under this good event, the convex function $\bar{s}_T(\nu,d)$ is larger than a quadratic function in $\Omega_\nu(\varepsilon)$, which serves as a lower bound of dual function. The construction of this good event $\mathcal{E}(\varepsilon)$ is based on the following splitting scheme and concentration of objective function:

1. We first split $\Omega_\nu(\varepsilon)$ into multiple cubes layer by layer, and in every single cube, we control the difference of second-order terms between all the ν in the cube and the central point of the cube.
2. Then, we uniformly control the deviation of second-order terms for all central points.

We now discuss the second order term $\bar{s}_T(\nu,d)$ defined in (3.1). To derive an uniform lower bound of $\bar{s}_T(\nu,d)$, we do the following split on $\Omega_\nu(\varepsilon)$ according to Huber (1967). Define set $\Omega_\nu^k(\varepsilon)=\left\{\nu\middle|\,\|\nu-\nu^*\|_\infty\leq q^k4H\varepsilon\right\},0\leq k\leq N$, where $q\in(0,1)$ and $N\in\mathbb{N}_+$ will be identified later. This split divides $\Omega_\nu(\varepsilon)$ into N layers $\{\Omega_\nu^{k-1}(\varepsilon)\backslash\Omega_\nu^k(\varepsilon)\}_{k=1}^N$ and a center cube $\Omega_\nu^N(\varepsilon)$. We then split each layer into disjoint cubes $\{\bar{\Omega}^{kl}(\varepsilon)\}_{l=1}^{l_k}$ with edges of length $(1-q)q^{k-1}4H\varepsilon$, and denote the center cube by $\bar{\Omega}^{N1}(\varepsilon)$. Huber (1967) shows that there are at most $(2N)^m$ cubes. This split is not unique to get the desired convergence order, but it makes our result tight enough. The center of each cube $\bar{\Omega}^{kl}(\varepsilon)$ is defined as ν_{kl}. Define $\bar{\nu}_{kl} = \arg\max_{\nu\in\bar{\Omega}^{kl}(\varepsilon)} \|\nu-\nu^*\|_2$, and

$$\Gamma_t^{kl} = \max_{\nu\in\bar{\Omega}^{kl}(\varepsilon)} \left[s_t(\nu_{kl},d) - s_t(\nu,d)\right]. \tag{12.2}$$

Then for $k\in\{0,\dots,N-1\}$, and $\forall\nu\in\bar{\Omega}^{kl}(\varepsilon)$, $\bar{s}_T$ can be decomposed as

$$\begin{aligned}\bar{s}_T(\nu,d) =& \frac{1}{T}\sum_{t=1}^T s_t(\nu,d) - \frac{1}{T}\sum_{t=1}^T s_t(\nu_{kl},d) + \frac{1}{T}\sum_{t=1}^T s_t(\nu_{kl},d)\\ \geq& \underbrace{\mathbb{E}s_t(\nu_{kl},d) - \mathbb{E}\Gamma_t^{kl}}_{12.3.1} + \underbrace{-\frac{1}{T}\sum_{t=1}^T \Gamma_t^{kl} + \mathbb{E}\Gamma_t^{kl}}_{12.3.2} + \underbrace{\frac{1}{T}\sum_{t=1}^T s_t(\nu_{kl},d) - \mathbb{E}s_t(\nu_{kl},d)}_{12.3.3}.\end{aligned} \tag{12.3}$$

We study the lower bounds of these 3 terms in (12.3) respectively.

Lower bound of 12.3.1:

$$\begin{aligned}\mathbb{E}\Gamma_t^{kl} =& \mathbb{E}\max_{\nu\in\bar{\Omega}^{kl}(\varepsilon)} \left[f_t^*(b_t^\top(\lambda_{kl}+\mu_{kl})) - f_t^*(b_t^\top(\lambda+\mu)) - g_t(b_t^\top(\lambda^*+\mu^*))^\top b_t^\top(\lambda_{kl}+\mu_{kl}-\lambda-\mu)\right]\\ =& \mathbb{E}\max_{\nu\in\bar{\Omega}^{kl}(\varepsilon)} \left[\int_0^1 v_1^\top(\nu)\left[g_t(b_t^\top(\lambda+\mu)+v_1\cdot z) - g_t(b_t^\top(\lambda^*+\mu^*))\right]dz\right]\\ \leq& \bar{b}\max_{\nu\in\bar{\Omega}^{kl}(\varepsilon)} \|\nu-\nu_{kl}\|_2 \cdot \left[\int_0^1 \mathbb{E}\max_{\nu\in\bar{\Omega}^{kl}(\varepsilon)} \left\|g_t(b_t^\top(\nu)+v_1\cdot z) - g_t(b_t^\top(\nu^*))\right\| dz\right]\\ \leq& L_1\bar{b}^2(\max_{\nu\in\bar{\Omega}^{kl}(\varepsilon)} \|\nu_{kl}-\nu\|_2)\cdot\|\bar{\nu}_{kl}-\nu^*\|_2,\end{aligned} \tag{12.4}$$

where $v_1(\nu) = b_t^\top(\lambda_{kl}+\mu_{kl}-\lambda-\mu)$ is the direction vector, and the second inequality is obtained by Assumption 4.

According to Proposition 4, we have

$$\mathbb{E}s_t(\nu_{kl},d) \geq \underline{\mathcal{L}}_s \|\nu_{kl}-\nu^*\|_2^2.$$

So for the first term, it is clear that

$$-\mathbb{E}\Gamma_t^{kl} + \mathbb{E}s_t(\nu_{kl},d) \geq \underline{\mathcal{L}}_s(\|\nu_{kl}-\nu^*\|_2^2 - L_1\bar{b}^2(\max_{\nu\in\bar{\Omega}^{kl}(\varepsilon)} \|\nu_{kl}-\nu\|_2)\cdot\|\bar{\nu}_{kl}-\nu^*\|_2. \tag{12.5}$$

Lower bound of 12.3.2: Since the gradients $\|g_t\|_\infty$ is bounded by D, by the integral form of Γ_{kl} in the second equality of 12.4, we also have:

$$\left\|\Gamma_t^{kl}\right\|_2 \leq 2\sqrt{n}\bar{b}D\max_{\nu\in\bar{\Omega}^{kl}(\varepsilon)} \|\nu-\nu_{kl}\|_2,$$

for any $t\in[T]$. Define event

$$\mathcal{E}_{kl,1}(\varepsilon_1) = \left\{-\frac{1}{T}\sum_{t=1}^T \Gamma_t^{kl} + \mathbb{E}\Gamma_t^{kl} < -2\varepsilon_1\sqrt{n}\bar{b}D\max_{\nu\in\bar{\Omega}^{kl}(\varepsilon)} \|\nu-\nu_{kl}\|_2\right\}. \tag{12.6}$$

Then according to Hoeffding's inequality, $\mathbb{P}(\mathcal{E}_{kl,1}(\varepsilon_1)) \leq \exp(-\frac{T\varepsilon_1^2}{2})$

Lower bound of 12.3.3: We calculate the norm of each $s_t(\nu_{kl}, d)$:

$$\begin{aligned}\|s_t(\nu_{kl}, d)\|_2 =& \left\| \left[\int_0^1 v_2^\top \left[g_t(b_t(\lambda^* + \mu^*) + v_2 \cdot z) - g_t(b_t(\lambda^* + \mu^*)) \right] dz \right] dz \right\|_2 \\ \leq& 2\sqrt{n}\bar{b}D \left\| \bar{\nu}_{kl} - \nu^* \right\|_2,\end{aligned} \tag{12.7}$$

for any $t \in [T]$, where $v_2 = b_t^\top(\nu_{kl} - \nu^*)$ is the direction vector. Define event

$$\mathcal{E}_{kl,2}(\varepsilon_2) = \left\{ \frac{1}{T} \sum_{t=1}^T s_t(\nu_{kl}, d) - \mathbb{E}s_t(\nu_{kl}, d) < -2\varepsilon_2\sqrt{n}\bar{b}D \left\| \bar{\nu}_{kl} - \nu^* \right\|_2 \right\}. \tag{12.8}$$

Then we have $\mathbb{P}(\mathcal{E}_{kl,2}) \leq \exp(-\frac{T\varepsilon_2^2}{2})$ by Hoeffding's inequality. Now we would like to make all the quantities in the lower bound uniform by leveraging the splitting scheme. From the split, we have

$$\begin{aligned}\max_{\nu \in \bar{\Omega}^{kl}(\varepsilon)} \|\nu - \nu_{kl}\|_2 &= \sqrt{m}(1-q)q^{k-1}4H\varepsilon, \\ \|\nu^* - \nu_{kl}\|_2 &\geq q^k 4H\varepsilon.\end{aligned}$$

And also

$$\begin{aligned}\|\nu^* - \bar{\nu}_{kl}\|_2 &\leq \|\nu^* - \nu_{kl}\|_2 + \max_{\bar{\Omega}^{kl}(\varepsilon)} \|\nu - \nu_{kl}\|_2 \\ \max_{\nu \in \bar{\Omega}^{kl}(\varepsilon)} \|\nu - \nu_{kl}\|_2 &\leq \frac{\sqrt{m}(1-q)}{q} \|\nu^* - \nu_{kl}\|_2 \leq \frac{\sqrt{m}(1-q)}{q} \|\nu^* - \bar{\nu}_{kl}\|_2.\end{aligned}$$

Thus we have the following result for the 12.3.1 term in (12.5).

$$\begin{aligned}-\mathbb{E}\Gamma_t^{kl} + \mathbb{E}s_t(\nu_{kl}, d) \geq& \underline{\mathcal{L}}_s(\|\nu_{kl} - \nu^*\|_2^2 - L_1\bar{b}^2(\max_{\nu \in \bar{\Omega}^{kl}(\varepsilon)} \|\nu_{kl} - \nu\|_2) \cdot \|\bar{\nu}_{kl} - \nu^*\|_2 \\ \geq& \frac{\underline{\mathcal{L}}_s}{\left(1 + \frac{\sqrt{m}(1-q)}{q}\right)^2} \|\bar{\nu}_{kl} - \nu^*\|_2^2 - \frac{\sqrt{m}(1-q)}{q} \cdot L_1\bar{b}^2 \|\bar{\nu}_{kl} - \nu^*\|_2^2.\end{aligned}$$

So there exists $\underline{q} = \frac{\sqrt{m}}{\sqrt{m} + 1 \wedge \frac{\underline{\mathcal{L}}_s}{4L_1\bar{b}^2}}$ such that when $q \geq \underline{q}$, $\frac{\sqrt{m}(1-q)}{q} \leq 1 \wedge \frac{\underline{\mathcal{L}}_s}{4L_1\bar{b}^2}$, and

$$\frac{\underline{\mathcal{L}}_s}{(1 + \frac{\sqrt{m}(1-q)}{q})^2} - \frac{\sqrt{m}(1-q)}{q} \cdot L_1\bar{b}^2 \geq \underline{\mathcal{L}}_s/2.$$

Choose $q = \underline{q} \vee \frac{1}{2}$. Then for the 12.3.1 term in (12.5) we have

$$-\mathbb{E}\Gamma_t^{kl} + \mathbb{E}s_t(\nu_{kl}, d) \geq \frac{\underline{\mathcal{L}}_s}{2} \|\bar{\nu}_{kl} - \nu^*\|_2^2. \tag{12.9}$$

Let $\varepsilon_1 = \varepsilon_2 = \sqrt{m\log m}\varepsilon$. For 12.3.2, under event $\mathcal{E}_{kl,1}^c(\varepsilon_1)$ in (12.6) we have

$$\begin{aligned}-\frac{1}{T}\sum_{t=1}^T \Gamma_t^{kl} + \mathbb{E}\Gamma_t^{kl} \geq& -2\varepsilon\sqrt{nm\log m}\bar{b}D(\max_{\nu \in \bar{\Omega}^{kl}(\varepsilon)} \|\nu - \nu^*\|_2) \\ \geq& -2\varepsilon\sqrt{nm\log m}\bar{b}D\frac{\sqrt{m}(1-q)}{q} \|\bar{\nu}_{kl} - \nu^*\|_2.\end{aligned} \tag{12.10}$$

For 12.3.3, under event $\mathcal{E}^c_{kl,2}(\varepsilon)$ in (12.8) we have

$$\frac{1}{T}\sum_{t=1}^{T} s_t(\nu_{kl},d) - \mathbb{E}s_t(\nu_{kl},d) \geq -2\varepsilon\sqrt{nm\log m}\bar{b}D\left\|\bar{\nu}_{kl}-\nu^*\right\|_2. \tag{12.11}$$

Now we combine second order lower bounds in (12.9), (12.10), (12.11) together under the desired good event

$$\mathcal{E}(\varepsilon) = \cap_{k=1}^{N}\cap_l\left(\mathcal{E}^c_{kl,1}(\varepsilon)\cap\mathcal{E}^c_{kl,2}(\varepsilon)\right),$$

where we choose N by setting the radius of $\bar{\Omega}^{N1}(\varepsilon)$: $\sqrt{m}q^N 4H\varepsilon \leq 2H\varepsilon$, i.e.,

$$N = \lceil\log_q(\frac{1}{2\sqrt{m}})\rceil \leq \frac{4L_1\bar{b}^2}{\underline{\mathcal{L}}_s}\sqrt{m}\log\sqrt{m}.$$

Under $\mathcal{E}(\varepsilon)$, for any $\nu\in\Omega_\nu(\varepsilon)$ satisfying $\left\|\nu-\nu^*\right\|_2 > 2H\varepsilon$, there exists $k=\{0,\dots,N-1\}$ and l such that $\nu\in\bar{\Omega}^{kl}(\varepsilon)$, and

$$\begin{aligned}\bar{s}_T(\nu,d) \geq &\frac{\underline{\mathcal{L}}_s}{2}\left\|\bar{\nu}_{kl}-\nu^*\right\|_2^2 - 2\varepsilon\sqrt{n}\bar{b}D\left(1+\frac{\sqrt{m}(1-q)}{q}\right)\left\|\bar{\nu}_{kl}-\nu^*\right\|_2\\ \geq &\frac{\underline{\mathcal{L}}_s}{2}\left\|\bar{\nu}_{kl}-\nu^*\right\|_2^2 - 4\varepsilon\sqrt{nm\log m}\bar{b}D\left\|\bar{\nu}_{kl}-\nu^*\right\|_2.\end{aligned}$$

Compute the probability of $\mathcal{E}(\varepsilon)$ we can show that

$$\begin{aligned}\mathbb{P}(\mathcal{E}(\varepsilon)) \geq &1-\sum_{0\leq k\leq N-1,\ l}\left(\mathbb{P}(\mathcal{E}_{kl,1}(\varepsilon))+\mathbb{P}(\mathcal{E}_{kl,2}(\varepsilon))\right)\\ \geq &1-2(2\lceil\log_q(\frac{1}{2\sqrt{m}})\rceil)^m\exp(-\frac{m\log mT\varepsilon^2}{2}) \geq 1-2\exp\left(-\frac{m\log m(T\varepsilon^2-1)}{4}\right).\end{aligned}$$

The following Lemma can show the concentration of the first-order term:

Lemma 6. *Under Assumptions 1-3, the concentration of the gradient in the first order term $\bar{\phi}_{T,\nu}(\nu^*,d)$ satisfies*

$$\mathbb{P}(\left\|\bar{\phi}_T(\nu^*,d)-\nabla D_\nu(\nu^*,d)\right\|_2 > \epsilon) \leq 2m\exp\left(-\frac{T\epsilon^2}{2m\sqrt{n}\bar{b}D}\right), \tag{12.12}$$

for any $\epsilon>0$.

Proof: According to Hoeffding's inequality, we have

$$\mathbb{P}(\left|\left(\bar{\phi}_T(\nu^*,d)\right)_i-\left(\nabla D_\nu(\nu^*,d)\right)_i\right| > \epsilon/\sqrt{m}) \leq 2\exp\left(-\frac{T\epsilon^2}{2m\sqrt{n}\bar{b}D}\right)$$

for $\forall i\in[m]$. Combining all m dimensions together, we conclude that

$$\mathbb{P}\left(\left\|\frac{1}{T}\sum_{t=1}^{T}\phi_t(\boldsymbol{\lambda}^*,d)-\nabla D(\boldsymbol{\lambda}^*,d)\right\| > \varepsilon\right) \leq 2m\exp\left(-\frac{T\varepsilon^2}{2m\sqrt{n}\bar{b}D}\right).$$

For the first order term, denote event $\mathcal{E}_0(\varepsilon_0) = \{\left\|\bar{\phi}_T(\boldsymbol{\lambda}, d) - \nabla D(\boldsymbol{\lambda}^*, d)\right\|_2 > \varepsilon_0\}$. Take $\varepsilon_0 = \varepsilon\sqrt{nm\log m}\bar{b}D$. Then by Lemma 6, we have

$$\mathbb{P}(\mathcal{E}_0(\varepsilon_0)) \leq 2m\exp\left(-\frac{T\sqrt{n}\log m\bar{b}D\varepsilon^2}{2}\right).$$

Under event $\mathcal{E}_0^c(\varepsilon)\cap\mathcal{E}(\varepsilon)$, we have

$$\begin{aligned}\bar{D}_T(\boldsymbol{\lambda}, d) - \bar{D}_T(\boldsymbol{\lambda}^*, d) &\geq \bar{s}_T(\nu, d) + \left\langle \bar{\phi}_{T,\nu}(\nu^*, d) - \nabla_\nu D(\boldsymbol{\lambda}^*, d), \nu - \nu^*\right\rangle \\ &\geq \frac{\sigma_{\min}\underline{\mathcal{L}}_f}{4}\left\|\nu' - \nu^*\right\|_2^2 - 5\varepsilon\sqrt{nm\log m}\bar{b}D\left\|\bar{\nu}_{kl} - \nu^*\right\|_2 \\ &= \frac{\sigma_{\min}\underline{\mathcal{L}}_f}{4}\left\|\nu' - \nu^*\right\|_2^2 - \frac{\sigma_{\min}\underline{\mathcal{L}}_f}{4}\cdot 2H\varepsilon\left\|\nu' - \nu^*\right\|_2,\end{aligned} \tag{12.13}$$

where we define $H = 10\sqrt{nm\log m}\bar{b}D/(\sigma_{\min}\underline{\mathcal{L}}_f)$.

We now show how the first inequality leads to the probabilistic bound of $\left\|\nu - \nu^*\right\|_2$. By the definition of $\bar{s}_T(\nu, d)$, if the first inequality holds, an argument similar to (3.2) will lead to

$$\begin{aligned}\bar{D}_T(\boldsymbol{\lambda}, d) - \bar{D}_T(\boldsymbol{\lambda}^*, d) &\geq \bar{s}_T(\nu, d) + \left\langle \bar{\phi}_T(\boldsymbol{\lambda}^*, d), \boldsymbol{\lambda} - \boldsymbol{\lambda}^*\right\rangle \\ &\geq \bar{s}_T(\nu, d) + \left\langle \bar{\phi}_{T,\nu}(\nu^*, d) - \nabla_\nu D(\boldsymbol{\lambda}^*, d), \nu - \nu^*\right\rangle \\ &\geq \frac{\sigma_{\min}\underline{\mathcal{L}}_f}{4}\left(\left\|\nu' - \nu^*\right\|_2^2 - 2H\varepsilon\left\|\nu' - \nu^*\right\|_2\right),\end{aligned} \tag{12.14}$$

with probability at least $1 - 2m\exp(-\frac{T\sqrt{n}\log m\bar{b}D\varepsilon^2}{2})$, where we use the optimality of $\boldsymbol{\lambda}^*$ and the concentration of gradient. If ν_T^* is part of the optimal solution, then we claim that, the dual optimal solution ν_T^* must have $\left\|\nu_T^* - \nu^*\right\|_2 \leq 2H\varepsilon$. Otherwise:

1. If $2H\varepsilon < \left\|\nu_T^* - \nu^*\right\|_2 \leq 4H\varepsilon$, then there will be a $\nu_T^{*\prime}$ such that $\left\|\nu_T^{*\prime} - \nu^*\right\|_2 \geq \left\|\nu_T^* - \nu^*\right\|_2 > 2H\varepsilon$, and

$$\bar{D}_T(\boldsymbol{\lambda}_T^*, d) - \bar{D}_T(\boldsymbol{\lambda}^*, d) \geq \frac{\sigma_{\min}\underline{\mathcal{L}}_f}{4}\left(\left\|\nu_T^{*\prime} - \nu^*\right\|_2^2 - 2H\varepsilon\left\|\nu_T^{*\prime} - \nu^*\right\|_2\right) > 0,$$

which contradicts the optimality of $\boldsymbol{\lambda}_T^*$.

2. If $\left\|\nu_T^* - \nu^*\right\|_2 > 4H\varepsilon$, since $\bar{D}_T(\boldsymbol{\lambda}^*, d) - \bar{D}_T(\boldsymbol{\lambda}^*, d) = 0$ and $\bar{D}_T(\boldsymbol{\lambda}_T^*, d) - \bar{D}_T(\boldsymbol{\lambda}^*, d) \leq 0$, by the convexity of $\bar{D}_T$ we can always find a $\tilde{\boldsymbol{\lambda}}$ such that $2H\varepsilon < \left\|\tilde{\nu} - \nu^*\right\|_2 \leq 4H\varepsilon$ and $\bar{D}_T(\tilde{\boldsymbol{\lambda}}, d) - \bar{D}_T(\boldsymbol{\lambda}^*, d) \leq 0$. However, according to (12.14), we have

$$\bar{D}_T(\tilde{\boldsymbol{\lambda}}, d) - \bar{D}_T(\boldsymbol{\lambda}^*, d) \geq \frac{\sigma_{\min}\underline{\mathcal{L}}_f}{4}\left(\left\|\tilde{\nu}' - \nu^*\right\|_2^2 - 2H\varepsilon\left\|\tilde{\nu}' - \nu^*\right\|_2\right) > 0,$$

which also ends up with a contradiction.

To better present our idea to readers, we draw a figure here. This clearly shows that when our empirical function is lower bounded by a quadratic function with $H\varepsilon$ as the axis of symmetry, we essentially have $\left\|\nu_T^* - \nu^*\right\|_2 \leq 2H\varepsilon$.

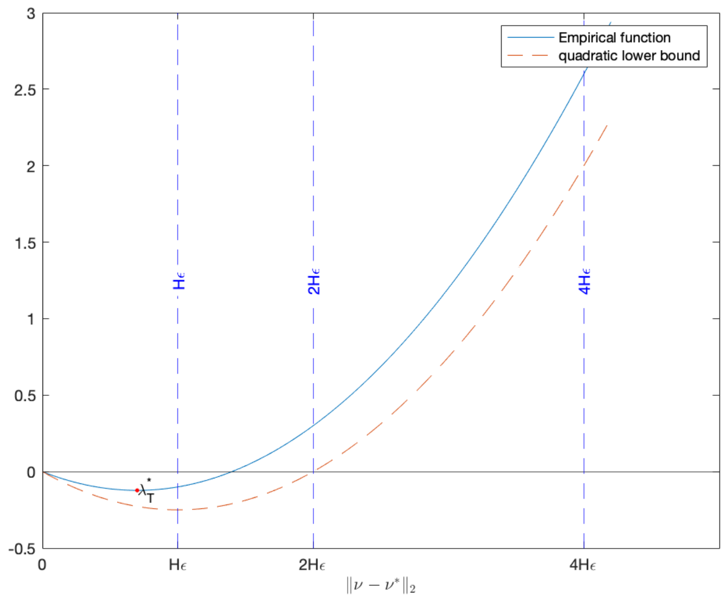

Figure 1 The value of $\bar{D}_T$ with respect to $\|\nu-\nu^*\|_2$. Since the quadratic lower bound has $H\varepsilon$ as the axis of symmetry, we have $\bar{D}_T(\boldsymbol{\lambda}, d) - \bar{D}_T(\boldsymbol{\lambda}^*, d) > 0$ for $\|\nu-\nu^*\|_2 > 2H\varepsilon$.

12.4. Proof of Theorem 1

We first show the dual convergence bound for a fixed $d \in \Omega_d$. By the tail expectation formula, for constant $H > 0$, we have

$$\mathbb{E}\,\|\nu_T^* - \nu^*\|_2^2 = 4H^2 \int_0^\infty \mathbb{P}(\|\nu_T^* - \nu^*\|_2^2 > 4H^2 z)dz.$$

According to the probabilistic bound in Proposition 2, for any $z > 0$,

$$\mathbb{P}(\|\nu_T^* - \nu^*\|_2^2 > 4H^2 z) \le 2m \exp\left(-\frac{T\sqrt{n}\log m\bar{b}D\varepsilon^2}{2}\right) \vee 1 + 2\exp\left(-\frac{m\log m(T\varepsilon^2 - 1)}{4}\right) \vee 1.$$

Then, calculating the integral, we get

$$\begin{aligned}
&\mathbb{E}(\|\nu_T^* - \nu^*\|_2^2) = H^2 \int_0^\infty \mathbb{P}(\|\nu_T^* - \nu^*\|_2^2 \ge 4H^2 z)dz \\
\le& \frac{400nm\log m\bar{b}^2D^2}{\sigma_{\min}^2 \underline{\mathcal{L}}_f^2} \int_{2/(T\sqrt{n}\bar{b}D)}^\infty \left[2\exp\left(-\frac{T\sqrt{n}\log m\bar{b}Dz}{2} + \log m\right) + \frac{2}{T\sqrt{n}\bar{b}D}\right]dz \\
&+ \frac{400nm\log m\bar{b}^2D^2}{\sigma_{\min}^2 \underline{\mathcal{L}}_f^2} \int_{\frac{1}{T}}^\infty \left[2\exp\left(-\frac{m\log m(Tz-1)}{4}\right) + \frac{1}{T}\right]dz \\
\le& C_1 \frac{\bar{b}^2D^2}{\sigma_{\min}^2 \underline{\mathcal{L}}_f^2} \frac{nm\log m}{T}.
\end{aligned}$$

For the optimality of the $O(T^{-1})$ order, let us consider a non-regularized case when $x \in [0,1]$ and $f_t(x) := f(x, \xi_t) := -(x - 2\xi_t)^2/4 + \xi_t^2$, with the single constraint $d = 1/2$ and cost $b_t = 1$. The dual problem is

$$D_t(\lambda) = \begin{cases} \frac{1}{2}\lambda & \text{if } \lambda > \xi_t \\ -\frac{1}{4} + \xi_t - \frac{1}{2}\lambda & \text{if } \lambda < \xi_t - \frac{1}{2} \\ \lambda^2 - 2(\xi_t - \frac{1}{4})\lambda + \xi_t^2 & \text{if } \xi_t - \frac{1}{2} \le \lambda \le \xi_t. \end{cases}$$

Let ξ_t be any distribution varies within $[1/2,\ 3/4]$ with variance $\sigma_\xi^2 > 0$. Then, for any t, we have $\xi_t - 1/4 \in [1/4,\ 1/2] \subseteq [\xi_t - 1/2, \xi_t]$. Thus, for the sample average $\bar{D}_T(\lambda) := T^{-1}\sum_{t=1}^T D_t(\lambda)$, when $\lambda \in [1/4,\ 1/2]$, $\bar{D}_T(\lambda) := \lambda^2 - 2(\bar{\xi_T} - 1/4)\lambda + \bar{\xi_T^2}$ with the optimal solution being $\lambda_T^* := \bar{\xi_T} - 1/4$. We have $\mathbb{E}(\lambda_T^* - \lambda^*)^2 \ge \mathrm{Var}(\bar{\xi_T}) = \sigma_\xi^2/T$. This shows that our $O(T^{-1})$ dual convergence rate is indeed optimal.

However, both the aforementioned analyses and previous works (Li and Ye 2021) only care about a fixed $d \in \Omega_d$, which neglects the intricate dependence of d when d changes in the process according to the accumulated data $\mathcal{H}_t$. We are going to show that, by applying the chaining argument on Ω_d, we can prove the dual convergence uniformly for all $d \in \Omega_d$. We would like to thank an anonymous reviewer for pointing out a possible approach to achieving a uniform bond, which greatly inspired our subsequent proof. It is worth noting that our argument can be used for both (i) localized Rademacher complexity; and (ii) proving a second-order lower bound of empirical function, to show the uniform bound of dual convergence. Here, we show the approach (ii) for demonstration, but the same argument can be applied to (i).

Proof of the uniform bound

We start from the step (12.14):

$$\begin{aligned} \bar{D}_T(\boldsymbol{\lambda}, d) - \bar{D}_T(\boldsymbol{\lambda}^*(d), d) &\ge \bar{s}_T(\nu, d) + \left\langle \bar{\phi}_T(\boldsymbol{\lambda}^*(d), d), \boldsymbol{\lambda} - \boldsymbol{\lambda}^*(d) \right\rangle \\ &\ge \bar{s}_T(\nu, d) + \left\langle \bar{\phi}_{T,\nu}(\nu^*(d), d) - \nabla_\nu D(\boldsymbol{\lambda}^*(d), d), \nu - \nu^* \right\rangle. \end{aligned}$$

We first show the uniform concentration of the gradient, i.e., the first-order term. Notice that $\bar{\phi}_{T,\nu}(\nu^*(d), d) - \nabla_\nu D(\boldsymbol{\lambda}^*(d), d) = -\sum_{t=1}^T \frac{1}{T} b_t g_t(b_t^\top \nu^*(d)) + \mathbb{E} b_t g_t(b_t^\top \nu^*(d))$. To ease the notation, we shall write $\bar{\Phi}_{i,T}(d) = \sum_{t=1}^T \frac{1}{T}\left(b_t g_t(b_t^\top \nu^*(d))\right)_i$, and write $\Phi_i(d) = \mathbb{E}\left(b_t g_t(b_t^\top \nu^*(d))\right)_i$, with each element $\Phi_{it}(d) = \left(b_t g_t(b_t^\top \nu^*(d))\right)_i$. At each dimension i, we have

$$\begin{aligned} &\sup_{d\in\Omega_d} \left| \sum_{t=1}^T \frac{1}{T}\left(b_t g_t(b_t^\top \nu^*(d))\right)_i - \mathbb{E}\left(b_t g_t(b_t^\top \nu^*(d))\right)_i \right| = \sup_{d\in\Omega_d} \left| \bar{\Phi}_{i,T}(d) - \Phi_i(d) \right| \\ &= \sup_{d\in\Omega_d} \left| \bar{\Phi}_{i,T}(d) - \Phi_i(d) \right| - \mathbb{E}\sup_{d\in\Omega_d} \left| \bar{\Phi}_{i,T}(d) - \Phi_i(d) \right| + \mathbb{E}\sup_{d\in\Omega_d} \left| \bar{\Phi}_{i,T}(d) - \Phi_i(d) \right|. \end{aligned}$$

Then, by the bouned difference condition (Koltchinskii 2011), we have the concentration

$$\mathbb{P}\left(\sup_{d\in\Omega_d} \left| \bar{\Phi}_{i,T}(d) - \Phi_i(d) \right| - \mathbb{E}\sup_{d\in\Omega_d} \left| \bar{\Phi}_{i,T}(d) - \Phi_i(d) \right| \ge \epsilon_1/\sqrt{m} \right) \le \exp\left(-\frac{T\epsilon_1^2}{2m\sqrt{n}\bar{b}D} \right).$$

We now compute the Rademacher complexity of $\bar{\Phi}_{i,T}(d)$ to control the third term:

$$\mathbb{E}\sup_{d\in\Omega_d} \left| \bar{\Phi}_{i,T}(d) - \Phi_i(d) \right| \le 2\mathcal{R}(\Phi_i, \Omega_d) := 2\mathbb{E}\left[\mathbb{E}_\sigma \sup_{d\in\Omega_d} \frac{1}{T}\sum_{t=1}^T \sigma_t \Phi_{i,t}(d) \right].$$

We now consider two cases: in the first case, we assume f_t is linear such that $f_t(x_t) = v_t^\top x_t$. Then, we have

$$f_t^*(b_t^\top \nu) := \sum_{i=1}^n \left(v_{it} - b_{it}^\top \nu\right)^+,$$

with the gradient $\Phi_{i,t}(d) = -\sum_{j=1}^n b_{ji,t}\mathbb{I}(v_{it} - b_{it}^\top \nu^*(d) > 0)$, which is composed of indicator functions with finite VC dimension (Vapnik and Chervonenkis 1971). In the space Ω_d, if we define the semi-norm $L_2(\Phi_i, \mathbb{P}_T)(d_1, d_2) = \sqrt{\left(\sum_{t=1}^T \left(\Phi_{i,t}(d_1) - \Phi_{i,t}(d_2)\right)^2\right)/T}$, then the semi-norm can be bounded by $L_2(\Phi_i, \mathbb{P}_T)(d_1, d_2) \le \sqrt{n}\bar{b}\sqrt{\left(\sum_{t=1}^T \left(\mathbb{I}\left(v_{it} - b_{it}^\top \nu^*(d_1) > 0\right) - \mathbb{I}\left(v_{it} - b_{it}^\top \nu^*(d_2) > 0\right)\right)^2\right)/T}$, which is also a semi-norm defined on the function class $\mathbf{G}_1 := \left\{\mathbb{I}\left(v_{it} - b_{it}^\top \nu^*(d) > 0\right) : d \in \Omega_d\right\}$. This shows that, the covering number $N(\epsilon, \Omega_d, L_2(\Phi_i, \mathbb{P}_T))$ can be bounded by

$$N(\epsilon, \Omega_d, L_2(\Phi_i, \mathbb{P}_T)) \le N(\epsilon, \Omega_d, \sqrt{n}\bar{b}L_2(\mathbf{G}_1, \mathbb{P}_T)),$$

which can be uniformly controlled by the VC dimension on $\mathbf{G}_1$, see, e.g., van der Vaart and Wellner (1996), van der Vaart and Wellner (2011), and Giné and Nickl (2021). Then, by Dudley's chaining argument on d, we have

$$\begin{aligned}
\mathcal{R}(\Phi_i, \Omega_d) &\le \frac{C}{\sqrt{T}}\int_0^{2\sqrt{n}\bar{b}} \sqrt{\log N\left(\epsilon, \Omega_d, L_2(\Phi_i, \mathbb{P}_T)\right)}d\epsilon \le \frac{C}{\sqrt{T}}\int_0^{2\sqrt{n}\bar{b}} \sqrt{\log N\left(\epsilon/(\sqrt{n}\bar{b}), \Omega_d, L_2(\mathbf{G}_1, \mathbb{P}_T)\right)}d\epsilon \\
&\le \frac{C\sqrt{n}\bar{b}}{\sqrt{T}}\int_0^1 \sqrt{\log(1 + c(1/\epsilon)^{2m})}d\epsilon \le \frac{C\bar{b}\sqrt{mn}}{\sqrt{T}},
\end{aligned}$$

where we used the VC dimension bound VC-dim$(\mathbf{G}_1) \le m + 1$. In the following discussion, we will treat the diameter of Ω_d as $O(\bar{d})$ for simplicity. This gives the bound on the Rademacher complexity in the linear case. We now consider the second case, i.e., the function is general. In this case, since we assume that g_t is Lipschitz continuous near $\nu^*(d)$, we have $L_2(\mathbb{P}_T)(d_1, d_2) \le L_1\bar{b}^2\left\|\nu^*(d_1) - \nu^*(d_1)\right\|_2 \le \frac{L_1\bar{b}^2}{\sigma_{\min}\underline{\mathcal{L}}_f}\left\|d_1 - d_2\right\|_2$. Here, we use the continuity of $\nu^*(d_1)$, which will be proved in Lemma 7. This leads to the bound of the covering number $N(\epsilon, \Omega_d, L_2(\Phi_i, \mathbb{P}_T)) \le N(\sigma_{\min}\underline{\mathcal{L}}_f\epsilon/(L_1\bar{b}^2), \Omega_d, L_2)$. We then have

$$\begin{aligned}
\mathcal{R}(\Phi_i, \Omega_d) &\le \frac{C}{\sqrt{T}}\int_0^D \sqrt{\log N\left(\sigma_{\min}\underline{\mathcal{L}}_f\epsilon/(L_1\bar{b}^2), \Omega_d, L_2\right)}d\epsilon \\
&\le \frac{CL_1\bar{b}^2}{\sigma_{\min}\underline{\mathcal{L}}_f\sqrt{T}}\int_0^{\sigma_{\min}\underline{\mathcal{L}}_f\bar{d}/(L_1\bar{b}^2)} \sqrt{\log(1 + c(\bar{d}/\epsilon)^{2m})}d\epsilon \le \frac{CL_1\bar{b}^2\bar{d}\sqrt{m}}{\sigma_{\min}\underline{\mathcal{L}}_f\sqrt{T}},
\end{aligned}$$

where in the integral, we only consider the part that ϵ is smaller than $\sigma_{\min}\underline{\mathcal{L}}_f\delta_1$, which gives $\left\|\nu^*(d_1) - \nu^*(d_1)\right\|_2 \le \delta_1$ for a given small constant δ_1 because the large part in the integral can always be bounded by the constant $\bar{d}/(L_1\bar{b}^2) \cdot \sqrt{\log\left(1 + c(\bar{d}/(\sigma_{\min}\underline{\mathcal{L}}_f\delta_1))^{2m}\right)} \lesssim \sqrt{m} \cdot \bar{d}\sqrt{\log(\bar{d}/\delta_1)}/(L_1\bar{b}^2)$. In the following computation, we will use the bound in the second case (for the general functions) because this can show more general results reflecting the role of smoothness L_1. For the linear problem, L_1 will be substituted by $\sqrt{n}$. This leads to the bound that our gradient can be controlled by

$$\begin{aligned}\sup_{d\in\Omega_d}\left\|\bar{\phi}_{T,\nu}(\nu^*(d),d)-\nabla_\nu D(\boldsymbol{\lambda}^*(d),d)\right\|_2 &\le \epsilon_1+\sqrt{\sum_{i=1}^m\left(\mathbb{E}\sup_{d\in\Omega_d}\left|\bar{\Phi}_{i,T}(d)-\Phi_i(d)\right|\right)^2}\\ &\le \epsilon_1+\frac{CL_1\bar{b}^2\bar{d}m}{\sigma_{\min}\underline{\mathcal{L}}_f\sqrt{T}},\end{aligned} \tag{12.15}$$

will probability at least $1-m\exp\left(-\frac{T\epsilon_1^2}{2m\sqrt{n}\bar{b}D}\right)$.

We now consider the uniform lower bound of $\bar{s}_T(\nu,d)$, for ν in a local region close to the optimal solution $\nu^*(d)$, say, $\mathbb{B}(\nu^*(d),\varepsilon)=\{\nu:\|\nu-\nu^*(d)\|_2\le\varepsilon\}$. Notice that, the expectation of $\bar{s}_T(\nu,d)$, i.e., $s(\nu,d)$, as is shown in Proposition 4, shares the following lower bound for every $d\in\Omega_d$: $s(\nu,d)\ge \sigma_{\min}\underline{\mathcal{L}}_f/2\,\|\nu-\nu^*(d)\|_2^2$. This gives

$$\begin{aligned}\bar{s}_T(\nu,d)&=\bar{s}_T(\nu,d)-s(\nu,d)+s(\nu,d)\ge\bar{s}_T(\nu,d)-s(\nu,d)+\frac{\sigma_{\min}\underline{\mathcal{L}}_f}{2}\|\nu-\nu^*(d)\|_2^2\\ &\ge\frac{\sigma_{\min}\underline{\mathcal{L}}_f}{2}\|\nu-\nu^*(d)\|_2^2+\inf_{d,\nu\in\mathbb{B}(\nu^*(d),\varepsilon)}[\bar{s}_T(\nu,d)-s(\nu,d)]-\mathbb{E}\inf_{d,\nu\in\mathbb{B}(\nu^*(d),\varepsilon)}[\bar{s}_T(\nu,d)-s(\nu,d)]\\ &\quad+\mathbb{E}\inf_{d,\nu\in\mathbb{B}(\nu^*(d),\varepsilon)}[\bar{s}_T(\nu,d)-s(\nu,d)].\end{aligned}$$

For ease of exposition, we now consider ν as $\nu=\nu^*(d)+\Delta\nu$, where we confine $\Delta\nu\in\mathbb{B}(0,\varepsilon)$, which will be denoted by $\mathbb{B}(\varepsilon)$ for simplicity. We then have, for every $d\in\Omega_d$, $\Delta\nu\in\mathbb{B}(\varepsilon)$,

$$\mathbb{P}\left(\inf_{d,\Delta\nu\in\mathbb{B}(\varepsilon)}[\bar{s}_T(\nu,d)-s(\nu,d)]-\mathbb{E}\inf_{d,\Delta\nu\in\mathbb{B}(\varepsilon)}[\bar{s}_T(\nu,d)-s(\nu,d)]\le-\epsilon_2\right)\le\exp\left(-\frac{T\epsilon_2^2}{8n\bar{b}^2D^2\varepsilon^2}\right). \tag{12.16}$$

where we use the bounded difference condition (Koltchinskii 2011) and the norm bound of $s_t(\nu^*(d)+\Delta\nu,d)$ as in (12.7)

$$\begin{aligned}|s_t(\nu^*(d)+\Delta\nu,d)|&=\left|\left[\int_0^1 v_2^\top\left[g_t(b_t^\top(\lambda^*+\mu^*)+v_2\cdot z)-g_t(b_t^\top(\lambda^*+\mu^*))\right]dz\right]dz\right|\\ &\le 2\sqrt{n}\bar{b}D\|\Delta\nu\|_2\le2\sqrt{n}\bar{b}D\varepsilon,\end{aligned}$$

where the direction vector $v_2=b_t^\top\Delta\nu$, and we use the fact that $\|v_2\|_2\le\bar{b}\|\Delta\nu\|_2$, and $\|g_t\|_2\le\sqrt{n}D$. We now consider the expectation of the minima, $\mathbb{E}\inf_{d,\Delta\nu\in\mathbb{B}(\varepsilon)}[\bar{s}_T(\nu,d)-s(\nu,d)]$, where we can again use the Rademacher complexity to control it. It is clear that

$$\left|\mathbb{E}\inf_{d,\Delta\nu\in\mathbb{B}(\varepsilon)}[\bar{s}_T(\nu,d)-s(\nu,d)]\right|\le2\mathcal{R}\left(s,\Omega_d\times\mathbb{B}(\varepsilon)\right):=2\mathbb{E}\left[\mathbb{E}_\sigma\sup_{d,\Delta\nu\in\mathbb{B}(\varepsilon)}\frac{1}{T}\sum_{t=1}^T\sigma_t s_t(\nu^*(d)+\Delta\nu,d)\right].$$

We now use Dudley's chaining argument on both $\Delta\nu$ and d to establish a sharp bound of $\mathcal{R}\left(s,\Omega_d\times\mathbb{B}(\varepsilon)\right)$. We first construct a chain on $\Delta\nu$, say $\{\pi_k^\nu(\Delta\nu)\}_{k=1}^\infty$, where each two elements in the chain is bounded by $\left\|\pi_{k-1}^\nu(\Delta\nu)-\pi_k^\nu(\Delta\nu)\right\|_2\le2^{-k}\varepsilon$ and $\pi_0^\nu(\Delta\nu)=0$, $\pi_\infty^\nu(\Delta\nu)=\Delta\nu$.

Then, by Dudley's chaining argument on $\Delta\nu$, we have

$$\mathcal{R}\left(s, \Omega_d \times \mathbb{B}(\varepsilon)\right) \leq \sum_{k=1}^{\infty} \mathbb{E}\left[\mathbb{E}_\sigma \sup_{\pi^\nu_{k-1}(\Delta\nu), \pi^\nu_k(\Delta\nu)} \sup_d \frac{1}{T}\sum_{t=1}^{T} \sigma_t\left(s_t(\nu^*(d)+\pi^\nu_{k-1}, d) - s_t(\nu^*(d)+\pi^\nu_k, d)\right)\right].$$

For each k, the elements in $\sup_{\pi^\nu_{k-1}(\Delta\nu), \pi^\nu_k(\Delta\nu)}$ are finite and the number can be controlled by $N^2(2^{-k}\varepsilon, \mathbb{B}(\varepsilon), L_2)$. Thus, we can use chaining again for d given each k. For each k, we have

$$\begin{aligned}&\left|s_t(\nu^*(d)+\pi^\nu_{k-1}, d) - s_t(\nu^*(d)+\pi^\nu_k, d)\right| = \left|\left[\int_0^1 v_3^\top\left[g_t(b_t^\top(\nu^*(d)+\pi^\nu_{k-1}) + v_3\cdot z) - g_t(b_t^\top \nu^*(d))\right]dz\right]\right| \\ &\leq 2\sqrt{n}\bar{b}D\left\|\pi^\nu_{k-1} - \pi^\nu_k\right\| \leq 2\sqrt{n}\bar{b}D\cdot 2^{-k}\varepsilon,\end{aligned}$$

which serves as an envelope of each $s_t(\nu^*(d)+\pi^\nu_{k-1}, d) - s_t(\nu^*(d)+\pi^\nu_k, d)$. Here the direction vector $v_3 = b_t^\top(\pi^\nu_k - \pi^\nu_{k-1})$. Still, we first consider the linear case. If f_t is linear, we have $f_t(x) = v_t^\top x$ and $f_t^*(b_t^\top\nu) = \sum_{i=1}^n (v_{it} - b_{it}^\top\nu)^+$, and $g_{jt}(b_t^\top\nu) = -\mathbb{I}(v_{j,t} - b_{jt}^\top\nu > 0)$.

$$s_t(\nu^*(d)+\Delta\nu, d) = f_t^*(b_t^\top(\nu^*(d)+\Delta\nu)) - f_t^*(b_t^\top\nu^*(d)) + \sum_{i=1}^{m}\sum_{j=1}^{n} b_{ji,t}\mathbb{I}(v_{it} - b_{it}^\top\nu^*(d) > 0)\Delta\nu_i.$$

We now control the semi-norm defined using $s_t(\nu^*(d)+\pi^\nu_{k-1}, d) - s_t(\nu^*(d)+\pi^\nu_k, d)$ for each k. It is clear that for any fixed chain element π^ν_k, π^ν_{k-1} and any d_1, d_2, we have

$$\begin{aligned}&\left|\left(s_t(\nu^*(d_1)+\pi^\nu_{k-1}, d_1) - s_t(\nu^*(d_1)+\pi^\nu_k, d_1)\right) - \left(s_t(\nu^*(d_2)+\pi^\nu_{k-1}, d_2) - s_t(\nu^*(d_2)+\pi^\nu_k, d_2)\right)\right| \\ &\leq n\max_i\left|\left(v_{it} - b_{it}^\top(\nu^*(d_1)+\pi^\nu_{k-1})\right)^+ - \left(v_{it} - b_{it}^\top(\nu^*(d_1)+\pi^\nu_k)\right)^+\right. \\ &\left.- \left(\left(v_{it} - b_{it}^\top(\nu^*(d_2)+\pi^\nu_{k-1})\right)^+ - \left(v_{it} - b_{it}^\top(\nu^*(d_2)+\pi^\nu_k)\right)^+\right)\right| \\ &+ \bar{b}2^{-k}\varepsilon n\max_i\left|\mathbb{I}\left(v_{it} - b_{it}^\top\nu^*(d_1) > 0\right) - \mathbb{I}\left(v_{it} - b_{it}^\top\nu^*(d_2) > 0\right)\right|.\end{aligned}$$

Then, the semi-norm defined using $s_t(\nu^*(d)+\pi^\nu_{k-1}, d) - s_t(\nu^*(d)+\pi^\nu_k, d)$ can be controlled by the semi-norm defined using the function class

$$\mathbf{G}_2 := \left\{\left(v_{it} - b_{it}^\top(\nu^*(d)+\pi^\nu_{k-1})\right)^+ - \left(v_{it} - b_{it}^\top(\nu^*(d)+\pi^\nu_k)\right)^+ : d\in\Omega_d\right\},$$

and the semi-norm defined using $\mathbf{G}_1 := \left\{\mathbb{I}\left(v_{it} - b_{it}^\top\nu^*(d) > 0\right) : d\right\}$, i.e.,

$$L_2(s_t, \mathbb{P}_T) \leq n\max_{\mathbb{P}_T} L_2(\mathbf{G}_2, \mathbb{P}_T) + \bar{b}2^{-n}\varepsilon n\max_{\mathbb{P}_T} L_2(\mathbf{G}_1, \mathbb{P}_T),$$

and thus

$$N(\epsilon, \Omega_d, L_2(s_t, \mathbb{P}_T)) \leq \max_{\mathbb{P}_T} N(\epsilon/2, \Omega_d, \bar{b}2^{-n}\varepsilon n L_2(\mathbf{G}_1, \mathbb{P}_T))\cdot\max_{\mathbb{P}_T} N(\epsilon/2, \Omega_d, nL_2(\mathbf{G}_2, \mathbb{P}_T)),$$

where the inequality holds because we can construct the first layer of covering with $\epsilon/2$ and the second layer of covering on each ball of the first covering with $\epsilon/2$. For $L_2(\mathbf{G}_1, \mathbb{P}_T)$, we have VC-dim$(\mathbf{G}_1) \leq m+1$; for $L_2(\mathbf{G}_2, \mathbb{P}_T)$, the subgraph of elements in $\mathbf{G}_2$ can be covered by

$$\mathrm{sub}(\mathbf{G}_2) \subseteq \mathrm{sub}\left(v_{it} - b_{it}^\top(\nu^*(d)+\pi^\nu_{k-1})\right)\cup\mathrm{sub}\left(-v_{it} + b_{it}^\top(\nu^*(d)+\pi^\nu_k)\right)\cup\mathrm{sub}\left(b_{it}^\top(\pi^\nu_k - \pi^\nu_{k-1})\right)\cup\mathrm{sub}\left(0\right),$$

which shows that the function class $\mathbf{G}_2$ also has the VC dimension bounded by $\text{VC-dim}(\mathbf{G}_2) \lesssim m$ according to van der Vaart and Wellner (1996). Moreover, s_t and $\mathbf{G}_2$ are uniformly bounded by $2\sqrt{n}\bar{b}D \cdot 2^{-k}\varepsilon$. Together we have

$$
\begin{aligned}
&\mathbb{E}\left[\mathbb{E}_\sigma \sup_{\pi^\nu_{k-1}(\Delta\nu),\pi^\nu_k(\Delta\nu)} \sup_d \frac{1}{T}\sum_{t=1}^T \sigma_t\left(s_t(\nu^*(d)+\pi^\nu_{k-1},d) - s_t(\nu^*(d)+\pi^\nu_k,d)\right)\right] \\
&\quad\le \frac{C\sqrt{n}\bar{b}D\varepsilon}{2^k\sqrt{T}}\int_0^1 \sqrt{\log\left(N^2(2^{-k}\varepsilon,\mathbb{B}(\varepsilon),L_2)\cdot N\left(\sqrt{n}\bar{b}D2^{-k}\varepsilon\epsilon,\Omega_d,L_2(s_t,\mathbb{P}_T)\right)\right)}d\epsilon \\
&\quad\le \frac{C\sqrt{n}\bar{b}D\varepsilon\sqrt{\log N(2^{-k}\varepsilon,\mathbb{B}(\varepsilon),L_2)}}{2^k\sqrt{T}} \\
&\quad+ \frac{C\sqrt{n}\bar{b}D\varepsilon}{2^k\sqrt{T}}\left(\int_0^1 \sqrt{\log N(\epsilon/2,\Omega_d,\sqrt{n}L_2(\mathbf{G}_1,\mathbb{P}_T))} + \sqrt{\log N(\sqrt{n}\bar{b}D2^{-k}\varepsilon\epsilon/2,\Omega_d,\sqrt{n}L_2(\mathbf{G}_2,\mathbb{P}_T))}d\epsilon\right) \\
&\quad\le \frac{C\sqrt{n}\bar{b}D\varepsilon\left(\sqrt{\log N(2^{-k}\varepsilon,\mathbb{B}(\varepsilon),L_2)}+\sqrt{mn}\right)}{2^k\sqrt{T}} \\
&\quad\le \frac{C\sqrt{n}\bar{b}D\varepsilon\left(\sqrt{mn}+\sqrt{\log N(2^{-k}\varepsilon,\mathbb{B}(\varepsilon),L_2)}\right)}{2^k\sqrt{T}},
\end{aligned}
\tag{12.17}
$$

where for each n, the covering number is at most $N^2(2^{-k}\varepsilon,\mathbb{B}(\varepsilon),L_2)$ times the covering number on d. Here, we omit the maximum over $\mathbb{P}_T$ since the VC dimension bound is uniform. Summing up k from $k=1$ to ∞ will result in the bound of the whole Rademacher complexity $\mathcal{R}\left(s,\Omega_d\times\mathbb{B}(\varepsilon)\right)$:

$$
\mathcal{R}\left(s,\Omega_d\times\mathbb{B}(\varepsilon)\right) \le \frac{C\sqrt{n}\bar{b}D\varepsilon\left(\sqrt{mn}+\int_0^1\sqrt{\log N\left(\epsilon,\mathbb{B}(1),L_2\right)}d\epsilon\right)}{\sqrt{T}} \le \frac{C\sqrt{n}\bar{b}D\varepsilon\sqrt{mn}}{\sqrt{T}}. \tag{12.18}
$$

Similarly, if f_t is general, we have stronger smoothness condition on f_t^*, which means that for different d_1, d_2, the semi-norm can be bounded given that

$$
\begin{aligned}
&\left|\left(s_t(\nu^*(d_1)+\pi^\nu_{k-1},d_1)-s_t(\nu^*(d_1)+\pi^\nu_k,d_1)\right)-\left(s_t(\nu^*(d_2)+\pi^\nu_{k-1},d_2)-s_t(\nu^*(d_2)+\pi^\nu_k,d_2)\right)\right| \\
&=\left|\left[\int_0^1 v_3^\top\left[g_t(b_t(\nu^*(d_1)+\pi^\nu_{k-1})+v_3\cdot z)-g_t(b_t(\nu^*(d_2)+\pi^\nu_{k-1})+v_3\cdot z)-g_t(b_t^\top\nu^*(d_1))+g_t(b_t^\top\nu^*(d_2))\right]dz\right]\right| \\
&\le 2L_1\sqrt{n}\bar{b}^2D\left\|\pi^\nu_{k-1}-\pi^\nu_k\right\|_2\left\|\nu^*(d_1)-\nu^*(d_2)\right\|_2 \le 2L_1\sqrt{n}\bar{b}^2D\cdot 2^{-k}\varepsilon\left\|\nu^*(d_1)-\nu^*(d_2)\right\|_2,
\end{aligned}
$$

which also gives a controlled covering number for each k with the smoothness of $\nu^*(d)$:

$$
N(\epsilon,\Omega_d,L_2(s_t,\mathbb{P}_T)) \le N(\sigma_{\min}\underline{\mathcal{L}}_f\epsilon/(2L_1\sqrt{n}\bar{b}^2D2^{-k}\varepsilon),\Omega_d,L_2).
$$

We then have the Rademacher complexity bound using the same argument as (12.17) and (12.18) in the first case:

$$
\mathcal{R}\left(s,\Omega_d\times\mathbb{B}(\varepsilon)\right) \le \frac{C\sqrt{n}\bar{b}D\varepsilon\left(L_1\sqrt{m}\bar{b}\bar{d}/(\sigma_{\min}\underline{\mathcal{L}}_f)+\int_0^1\sqrt{\log N\left(\epsilon,\mathbb{B}(1),L_2\right)}d\epsilon\right)}{\sqrt{T}} \le \frac{CL_1\sqrt{mn}\bar{b}^2\bar{d}D\varepsilon}{\sigma_{\min}\underline{\mathcal{L}}_f\sqrt{T}}.
$$

Here, we use the fact that

$$\int_0^1 \sqrt{\log\left(N\left(\sqrt{n}\bar{b}D2^{-k}\varepsilon\epsilon, \Omega_d, L_2(s_t, \mathbb{P}_T)\right)\right)}d\epsilon \leq \int_0^1 \sqrt{\log N(\sigma_{\min}\underline{\mathcal{L}}_f\epsilon/(2L_1\bar{b}), \Omega_d, L_2)}d\epsilon \leq C\frac{L_1\bar{b}\bar{d}\sqrt{m}}{\sigma_{\min}\underline{\mathcal{L}}_f}.$$

Associate this result with (12.16), we have the uniform lower bound of

$$\begin{aligned}\bar{s}_T(\nu,d) &\geq \frac{\sigma_{\min}\underline{\mathcal{L}}_f}{2}\left\|\nu-\nu^*(d)\right\|_2^2 + \inf_{d,\nu\in\mathbb{B}(\nu^*(d),\varepsilon)}\left[\bar{s}_T(\nu,d)-s(\nu,d)\right] - \mathbb{E}\inf_{d,\nu\in\mathbb{B}(\nu^*(d),\varepsilon)}\left[\bar{s}_T(\nu,d)-s(\nu,d)\right] \\ &\quad + \mathbb{E}\inf_{d,\nu\in\mathbb{B}(\nu^*(d),\varepsilon)}\left[\bar{s}_T(\nu,d)-s(\nu,d)\right] \\ &\geq \frac{\sigma_{\min}\underline{\mathcal{L}}_f}{2}\left\|\nu-\nu^*(d)\right\|_2^2 - \epsilon_2 - \frac{CL_1\sqrt{mn}\bar{b}^2\bar{d}D\varepsilon}{\sigma_{\min}\underline{\mathcal{L}}_f\sqrt{T}},\end{aligned}$$

which gives the uniform lower bound of the whole dual function by applying (12.15):

$$\begin{aligned}&\bar{D}_T(\boldsymbol{\lambda},d) - \bar{D}_T(\boldsymbol{\lambda}^*(d),d) \geq \bar{s}_T(\nu,d) + \left\langle \bar{\phi}_{T,\nu}(\nu^*(d),d) - \nabla_\nu D(\boldsymbol{\lambda}^*(d),d), \nu-\nu^*\right\rangle. \\ &\geq \frac{\sigma_{\min}\underline{\mathcal{L}}_f}{2}\left\|\nu-\nu^*(d)\right\|_2^2 - \left(\epsilon_1 + \frac{CL_1\bar{b}^2\bar{d}m}{\sigma_{\min}\underline{\mathcal{L}}_f\sqrt{T}}\right)\left\|\nu-\nu^*(d)\right\|_2 - \epsilon_2 - \frac{CL_1\sqrt{mn}\bar{b}^2\bar{d}D\varepsilon}{\sigma_{\min}\underline{\mathcal{L}}_f\sqrt{T}},\end{aligned}$$

which holds uniformly for all $d\in\Omega_d$ and $\nu\in\mathbb{B}(\nu^*(d),\varepsilon)$ with probability at least $1-m\exp\left(-\frac{T\epsilon_1^2}{2m\sqrt{n}\bar{b}D}\right)-\exp\left(-\frac{T\epsilon_2^2}{8n\bar{b}^2D^2\varepsilon^2}\right)$. Choosing $\epsilon_1=\sigma_{\min}\underline{\mathcal{L}}_f/8\varepsilon$ and $\epsilon_2=\sigma_{\min}\underline{\mathcal{L}}_f\varepsilon^2/64$, it is clear that when $\varepsilon\geq\frac{CL_1\sqrt{mn}\bar{b}^2\bar{d}D\varepsilon}{\sigma_{\min}\underline{\mathcal{L}}_f\sqrt{T}}$, we have

$$\bar{D}_T(\boldsymbol{\lambda},d) - \bar{D}_T(\boldsymbol{\lambda}^*(d),d) \geq \frac{\sigma_{\min}\underline{\mathcal{L}}_f}{2}\left(\left\|\nu-\nu^*(d)\right\|_2 - \frac{\varepsilon}{4}\right)^2 - \frac{\varepsilon^2}{4}.$$

If the $\boldsymbol{\lambda}_T^*$ is the optimal solution, we must have $\sup_{d\in\Omega_d}\left\|\nu_T^*(d)-\nu^*(d)\right\|_2\leq\frac{3\varepsilon}{4}$ following the same reason as Proposition 2. We then use the tail expectation formula again on the probability bound to get the result

$$\mathbb{E}\sup_{d\in\Omega_d}\left\|\nu_T^*(d)-\nu^*(d)\right\|_2^2 = \int_0^\infty \mathbb{P}\left(\sup\left\|\nu_T^*-\nu^*\right\|_2^2 > z\right)dz \lesssim C_1\frac{L_1^2\bar{b}^4\bar{d}^2D^2}{\sigma_{\min}^2\underline{\mathcal{L}}_f^2}\frac{nm\log m}{T},$$

where an additional $L_1^2\bar{b}^2\bar{d}^2$ term appears due to the chaining argument on d.

12.5. Proof of Theorem 2

We first prove this result for a fixed d. The uniform result can be similarly proved by simply using the argument in the proof of Theorem 1. Recall that, by the proof of Proposition 2, the convex function $\bar{D}_T$ is larger than a quadratic function in a neighborhood of ν^* with a high probability claimed there. Then, for any ϵ satisfying $\epsilon<2H^2\varepsilon^2\sigma_{\min}\underline{\mathcal{L}}_f$, with the same high probability, the ϵ-optimal solution must belong to $\Omega_\nu(\varepsilon)$, because, for all the points in the border $\left\|\nu-\nu^*\right\|_2=4H\varepsilon$, we already have $\bar{D}_T(\boldsymbol{\lambda},d)-\bar{D}_T(\boldsymbol{\lambda}^*,d)\geq 2H^2\varepsilon^2\sigma_{\min}\underline{\mathcal{L}}_f$. Then, with the same high probability, it follows that

$$\epsilon \geq \bar{D}_T(\boldsymbol{\lambda}_T^\epsilon,d) - \bar{D}_T(\boldsymbol{\lambda}^*,d) \geq \frac{\sigma_{\min}\underline{\mathcal{L}}_f}{4}\left\|\nu_T^{\epsilon\prime}-\nu^*\right\|_2^2 - \frac{\sigma_{\min}\underline{\mathcal{L}}_f}{4}\cdot 2H\varepsilon\left\|\nu_T^{\epsilon\prime}-\nu^*\right\|_2,$$

which suggests that $\|\nu_T^\epsilon - \nu^*\|_2 \le \|\nu_T^{\epsilon\prime} - \nu^*\|_2 \le H\varepsilon + \left(H^2\varepsilon^2 + 4\epsilon/(\sigma_{\min}\underline{\mathcal{L}}_f)\right)^{1/2}$. Still, applying the tail expectation formula, we get

$$\begin{aligned}
\mathbb{E}(\|\nu_T^\epsilon - \nu^*\|_2^2) = 4H^2 &\int_0^{\frac{2\epsilon}{H^2\sigma_{\min}\underline{\mathcal{L}}_f}} \mathbb{P}(\|\nu_T^\epsilon - \nu^*\|_2 \ge 2H\sqrt{z})dz \\
&+ 4H^2 \int_{\frac{2\epsilon}{H^2\sigma_{\min}\underline{\mathcal{L}}_f}}^{\infty} \mathbb{P}(\|\nu_T^\epsilon - \nu^*\|_2 \ge 2H\sqrt{z})dz \\
&\le \frac{8\epsilon}{\sigma_{\min}\underline{\mathcal{L}}_f} + 4H^2 \int_{\frac{\epsilon}{H^2\underline{\mathcal{L}}_D}}^{\infty} \mathbb{P}(\|\nu_T^\epsilon - \nu^*\|_2 \ge 2H\sqrt{z})dz.
\end{aligned}$$

Let $2H\sqrt{z} = H\varepsilon + \sqrt{H^2\varepsilon^2 + \frac{4\epsilon}{\sigma_{\min}\underline{\mathcal{L}}_f}}$. When $z > \frac{\epsilon}{H^2\sigma_{\min}\underline{\mathcal{L}}_f}$, we have $\epsilon < 2H^2\varepsilon^2\sigma_{\min}\underline{\mathcal{L}}_f$, thus $\mathbb{P}(\|\nu_T^\epsilon - \nu^*\|_2 \ge 2H\sqrt{z})$ can be bounded by Proposition 2. Also when $2H\sqrt{z} = H\varepsilon + \sqrt{H^2\varepsilon^2 + \frac{2\epsilon}{\sigma_{\min}\underline{\mathcal{L}}_f}}$, we have $\varepsilon^2 \ge z - \frac{2\epsilon}{H^2\sigma_{\min}\underline{\mathcal{L}}_f}$. By the integral of z, we get the second part of the bound. We note that this analysis can also be applied to the uniform bound in the proof of Theorem 1. For the uniform bound, there will be an additional $L_1^2\bar{b}^2\bar{d}^2$ term due to the chaining argument.

12.6. Proof of Corollary 1

We first prove this result for a fixed d. The uniform result can be similarly proved by simply using the argument in the proof of Theorem 1. Recall the proof of Theorem 2 that when ε satisfying $\epsilon < 2H^2\varepsilon^2\sigma_{\min}\underline{\mathcal{L}}_f$, with high probability the deterministic ϵ-optimal solution must be in $\Omega_\nu(\varepsilon)$. Similarly, for the stochastic ϵ-optimal solution, we try to confine it in a larger region so that with high probability $\mathbb{E}\left[\|\nu_T^\epsilon - \nu_T^*\|_2^2 \middle| \bar{D}_T\right]$ can still be bounded by ε. Notice that, although our Proposition 2 only focus on $\Omega_\nu(\varepsilon)$, it also bring us information outside $\Omega_\nu(\varepsilon)$. For any ε and ϵ, under the event when Proposition 2 holds, and any $\bar{D}_T$ we have:

1. If $\bar{D}_T(\boldsymbol{\lambda}_T^\epsilon, d) - \bar{D}_T(\boldsymbol{\lambda}^*, d) \le 2H^2\varepsilon^2\sigma_{\min}\underline{\mathcal{L}}_f$, then $\|\nu_T^\epsilon - \nu_T^*\|_2 \le 4H\varepsilon$.
2. If $\bar{D}_T(\boldsymbol{\lambda}_T^\epsilon, d) - \bar{D}_T(\boldsymbol{\lambda}^*, d) > 2H^2\varepsilon^2\sigma_{\min}\underline{\mathcal{L}}_f$, then we have $\|\nu_T^\epsilon - \nu_T^*\|_2 \le \frac{2}{H\varepsilon\sigma_{\min}\underline{\mathcal{L}}_f}(\bar{D}_T(\boldsymbol{\lambda}_T^\epsilon, d) - \bar{D}_T(\boldsymbol{\lambda}^*, d))$. Because the convex function $\bar{D}_T(\boldsymbol{\lambda}, d) - \bar{D}_T(\boldsymbol{\lambda}^*, d) = 0$ when $\boldsymbol{\lambda} = \boldsymbol{\lambda}^*$, and when $\|\nu - \nu^*\|_2 = 4H\varepsilon$, $\bar{D}_T(\boldsymbol{\lambda}, d) - \bar{D}_T(\boldsymbol{\lambda}^*, d) \ge 2H^2\varepsilon^2\sigma_{\min}\underline{\mathcal{L}}_f$.

We conclude that under the event when Proposition 2 holds, for any $\epsilon < 2H^2\varepsilon^2\sigma_{\min}\underline{\mathcal{L}}_f$,

$$\mathbb{E}_{\mathcal{B}}\left[\|\nu_T^\epsilon - \nu^*\|_2^2 \middle| \bar{D}_T\right] \le 16H^2\varepsilon^2 + \frac{4\sqrt{m\left(2\frac{\bar{f}+\bar{r}}{\underline{d}} + G\right)}}{H\varepsilon\sigma_{\min}\underline{\mathcal{L}}_f} \cdot \epsilon$$

because $\|\nu_T^\epsilon - \nu^*\|_2 \le 4\sqrt{m\left(2\frac{\bar{f}+\bar{r}}{\underline{d}} + G\right)}$. The RHS term has a minimum value

$$z_0 = 3 \cdot 8\epsilon^{\frac{2}{3}}\left(m\left(2\frac{\bar{f}+\bar{r}}{\underline{d}} + G\right)\right)^{\frac{1}{3}} / (\sigma_{\min}\underline{\mathcal{L}}_f)^{\frac{2}{3}}$$

when $\varepsilon_0 = \epsilon^{\frac{1}{3}}\left(m\left(2\frac{\bar{f}+\bar{r}}{\underline{d}}+G\right)\right)^{\frac{1}{6}}/(2H(\sigma_{\min}\underline{\mathcal{L}}_f)^{\frac{1}{3}})$. When the RHS term is larger than is minimum value, we can always take the corresponding ε at the right side where $\varepsilon > \varepsilon_0$ and it follows that

$$z = 16H^2\varepsilon^2 + \frac{4\sqrt{m\left(2\frac{\bar{f}+\bar{r}}{\underline{d}}+G\right)}}{H\varepsilon\sigma_{\min}\underline{\mathcal{L}}_f}\cdot\epsilon \le 48H^2\varepsilon^2.$$

Then by the tail expectation formula we have

$$\begin{aligned}
\mathbb{E}_{\mathcal{B},\mathcal{P}}\|\nu_T^\epsilon-\nu^*\|_2^2 &= \int_0^{z_0}\mathbb{P}(\mathbb{E}_{\mathcal{B}}\left[\|\nu_T^\epsilon-\nu^*\|_2^2\middle|\bar{D}_T\right]\ge z)dz + \int_{z_0}^{\infty}\mathbb{P}(\mathbb{E}_{\mathcal{B}}\left[\|\nu_T^\epsilon-\nu^*\|_2^2\middle|\bar{D}_T\right]\ge z)dz\\
&\le z_0 + \int_{z_0}^{\infty}\left[2m\exp\left(-\frac{T\sqrt{n}\log m\bar{b}Dz/(48H^2)}{2}\right)\vee 1\right]dz\\
&+ \int_{z_0}^{\infty}\left[2\exp\left(-\frac{m\log m(Tz/(48H^2)-1)}{4}\right)\vee 1\right]dz.\\
&\le z_0 + C_2\frac{\bar{b}^2D^2}{\sigma_{\min}^2\underline{\mathcal{L}}_f^2}\frac{nm\log m}{T}.
\end{aligned}$$

For the uniform bound, there will be an additional $L_1^2\bar{b}^2\bar{d}^2$ term due to the chaining argument.

12.7. Proof of Lemma 2

Recall the Lagrangian of program (2.5). By duality, we have

$$\begin{aligned}
R^*(\mathcal{P}) &:= \mathbb{E}_{\mathcal{P}}\left[\max_{x_t\in\mathcal{X}}\sum_{t=1}^T f_t(x_t) + T\cdot r\left(\frac{\sum_{t=1}^T b_tx_t}{T}\right),\ s.t.\ \sum_{t=1}^T b_tx_t \preceq dT\right]\\
&\le \mathbb{E}\sum_{t=1}^T\left[f_t(\tilde{x}_t(\nu^*)) + r\left(\tilde{a}(\mu^*)\right) + (\tilde{a}(\mu^*) - b_t\tilde{x}_t(\nu^*))^\top\mu^* + (d - b_t\tilde{x}_t(\nu^*))^\top\lambda^*\right]\\
&= T\cdot h(\nu^*)
\end{aligned}$$

12.8. Proof of Proposition 3

Since r is proper, by Fenchel conjugate function, the definition of $\hat{\mu}_T$ implies

$$\begin{aligned}
r\left(\frac{\sum_{t=1}^T b_tx_t}{T}\right) + \hat{\mu}_T^\top\frac{\sum_{t=1}^T b_tx_t}{T} =& r^*(-\hat{\mu}_T) - \hat{\mu}_T^\top\mathbb{E}b_t\tilde{x}_t(\nu^*) + \hat{\mu}_T^\top\mathbb{E}b_t\tilde{x}_t(\nu^*)\\
\ge& r\left(\tilde{a}(\mu^*)\right) + \hat{\mu}_T^\top\tilde{a}(\mu^*).
\end{aligned}$$

Combined with $R(A|\mathcal{P}) = \mathbb{E}_{A,\mathcal{P}}\left[\sum_{t=1}^T f_t(x_t) + T\cdot r\left(\frac{\sum_{t=1}^T b_tx_t}{T}\right)\right]$, we have

$$R(A|\mathcal{P}) \ge \mathbb{E}\left[\sum_{t=1}^T f_t(x_t) + Tr\left(\tilde{a}(\mu^*)\right) + T\hat{\mu}_T^\top\tilde{a}(\mu^*) - T\hat{\mu}_T^\top\frac{\sum_{t=1}^T b_tx_t}{T}\right].$$

The Assumption 2 suggests that

$$\|\hat{\mu}_T\|_2 \le \sqrt{m}G,\ \text{and}\ \|a\|_2 = \|\vartheta^*(-\mu)\|_2 \le \sqrt{n}D\bar{b}.$$

Thus

$$
\begin{aligned}
R\left(A|\mathcal{P}\right) \geq & \mathbb{E}\left[\sum_{t=1}^{T}\left[f_t(\tilde{x}_t(\nu_{t-1})) + r\left(\tilde{a}(\mu^*)\right)\right] + \left\langle \hat{\mu}_T, \sum_{t=1}^{T}(\tilde{a}(\mu^*) - b_t x_t)\right\rangle\right] \\
= & \mathbb{E}\left[\sum_{t=1}^{T} h(\nu_{t-1}) + \left\langle \hat{\mu}_T - \mu^*, \sum_{t=1}^{T}(\tilde{a}(\mu^*) - b_t x_t)\right\rangle - \left\langle \lambda^*, \sum_{t=1}^{T}(d - b_t x_t)\right\rangle\right]
\end{aligned}
$$

Combined with (5.2), we can show that

$$
R^*(\mathcal{P}) - R\left(A|\mathcal{P}\right) \leq \mathbb{E}\left[\sum_{t=1}^{T} h(\nu^*) - h(\nu_{t-1}) + \left\langle \lambda^*, \sum_{t=1}^{T}(d - b_t x_t)\right\rangle\right],
$$

or, equivalent, in a two-phase form:

$$
R^*(\mathcal{P}) - R\left(A|\mathcal{P}\right) \leq \mathbb{E}\left[\sum_{t=1}^{\tau} h(\nu^*) - h(\nu_{t-1}) + \left\langle \lambda^*, \sum_{t=1}^{\tau}(d - b_t x_t)\right\rangle\right] + 2(\bar{f} + \bar{r} + \sqrt{mn}GD\bar{b})(T - \tau).
$$

We conclude the proof.

12.9. Proof of Lemma 3

For technical convenience, we assume that, for each non-binding dimensions $i \in I_{\mathrm{NB}}$, the updated constraint d_{it} never exceeds the threshold $\bar{d}$ (the uniform bound defined in Assumption 1) at all iterations. This is a mid assumption both for theory and in practice. Indeed, if d_{it} is larger than the $\bar{d}$, this means that the constraint d_{it} is very loose so that its impact to the optimization problem is negligible. In this case, such a constraint can essentially be discarded. We start with a lemma on the continuity of dual optimal solution to prove Lemma 3.

Lemma 7 **(Continuity of dual optimal solution)**. *Under Assumption 1, 2, 3, for the stochasitc program* $\min_{\mu, \lambda \succeq 0} D(\boldsymbol{\lambda}, d') = \mathbb{E} f_t^*(b_t^\top(\mu + \lambda)) + r^*(-\mu) + d'^\top \lambda$, *let* d' *be* $d'_1, d'_2 \in \Omega_d$ *separately, then the corresponding optimal solution* $\nu^*(d'_1), \nu^*(d'_2)$ *satisfies*

$$
\left\|\nu^*(d'_1) - \nu^*(d'_2)\right\|_2^2 \leq \frac{1}{\sigma_{\min}^2 \underline{\mathcal{L}}_f^2} \left\|d'_1 - d'_2\right\|_2^2.
$$

If further d'_1, d'_2 *identify the same binding/non-binding dimensions, then*

$$
\left\|\nu^*(d'_1) - \nu^*(d'_2)\right\|_2^2 \leq \frac{1}{\sigma_{\min}^2 \underline{\mathcal{L}}_f^2} \sum_{i \in I_B}(d'_{1i} - d'_{2i})^2,
$$

where the binding dimension I_B *is with respect to* d'_1 *and* d'_2.

Proof of Lemma 7 By Proposition 4 and the uniform assumption on d, we have

$$
\begin{aligned}
D(\nu^*(d'_2), \mu^*(d'_1), d'_1) - D(\boldsymbol{\lambda}^*(d'_1), d'_1) &\geq \frac{1}{2}\sigma_{\min}\underline{\mathcal{L}}_f \left\|\nu^*(d'_2) - \nu^*(d'_1)\right\|_2^2 \\
D(\nu^*(d'_1), \mu^*(d'_2), d'_2) - D(\boldsymbol{\lambda}^*(d'_2), d'_2) &\geq \frac{1}{2}\sigma_{\min}\underline{\mathcal{L}}_f \left\|\nu^*(d'_1) - \nu^*(d'_2)\right\|_2^2.
\end{aligned}
$$

Summing up two inequality we have

$$(d_1' - d_2')^\top (\nu^*(d_2') - \nu^*(d_1')) \geq \sigma_{\min}\underline{\mathcal{L}}_f \left\|\nu^*(d_2') - \nu^*(d_1')\right\|_2^2, \tag{12.19}$$

or equivalently, $\sum_{i\in I_{\mathrm{B}}}(d_{1i}' - d_{2i}')(\nu_i^*(d_2') - \nu_i^*(d_1')) \geq \sigma_{\min}\underline{\mathcal{L}}_f \left\|\nu^*(d_2') - \nu^*(d_1')\right\|_2^2$ if further d_1', d_2' share the same binding/non-binding dimensions. From (12.19) we can show that

$$\begin{aligned}\sigma_{\min}\underline{\mathcal{L}}_f \left\|\nu^*(d_2') - \nu^*(d_1')\right\|_2^2 &\leq (d_1' - d_2')^\top (\nu^*(d_2') - \nu^*(d_1')) \leq \|d_1' - d_2'\|_2 \left\|\nu^*(d_2') - \nu^*(d_1')\right\|_2 \\ \left\|\nu^*(d_2') - \nu^*(d_1')\right\|_2 &\leq \frac{1}{\sigma_{\min}\underline{\mathcal{L}}_f} \|d_1' - d_2'\|_2 .\end{aligned}$$

Thus we get the first statement. For the second statement, we focus on the binding dimensions:

$$\begin{aligned}\sigma_{\min}\underline{\mathcal{L}}_f \left\|\nu^*(d_2') - \nu^*(d_1')\right\|_2^2 &\leq \sum_{i\in I_{\mathrm{B}}}(d_{1i}' - d_{2i}')(\nu_i^*(d_2') - \nu_i^*(d_1')) \\ &\leq \sqrt{\sum_{i\in I_{\mathrm{B}}}(d_{1i}' - d_{2i}')^2}\sqrt{\sum_{i\in I_{\mathrm{B}}}(\nu_i^*(d_2') - \nu_i^*(d_1'))^2} \\ &\leq \sqrt{\sum_{i\in I_{\mathrm{B}}}(d_{1i}' - d_{2i}')^2}\left\|\nu^*(d_2') - \nu^*(d_1')\right\|_2 ,\end{aligned}$$

which completes the proof of Lemma 7.

Then, we return to Lemma 3 and consider the original constraints d and the its binding/non-binding dimensions: $I_{\mathrm{B}} = \{i | d_i - \mathbb{E}\left(b_t\tilde{x}_t(\nu^*)\right)_i = 0\}$, and $I_{\mathrm{NB}} = \{i | d_i - \mathbb{E}\left(b_t\tilde{x}_t(\nu^*)\right)_i > 0\}$. Here we write the corresponding optimal solution to $\min_{\mu,\lambda\succeq 0} D(\boldsymbol{\lambda}, d)$ as $\boldsymbol{\lambda}^*$, and write $\boldsymbol{\lambda}^*(d')$ if we change d to d'. Then if $i\in I_{\mathrm{B}}$ and i changes to non-binding dimensions for d', by Lemma 7 and Assumption 5.2, for any $\|d' - d\|_2 \leq \delta_0 \wedge \frac{\underline{\lambda}}{2\mathcal{L}_r}$ we have

$$\|d - d'\|_2 \geq \sigma_{\min}\underline{\mathcal{L}}_f \left\|\nu^*(d') - \nu^*\right\|_2 \geq \sigma_{\min}\underline{\mathcal{L}}_f \left(|\lambda_i^* - 0| - |\mu_i - \mu_i(d')|\right) \geq \frac{1}{2}\sigma_{\min}\underline{\mathcal{L}}_f\underline{\lambda}, \tag{12.20}$$

where $\underline{\lambda} = \min\{\lambda_i^* | i\in I_{\mathrm{B}}\}$. If on the other hand, $i\in I_{\mathrm{NB}}$ and i changes to binding dimensions for d', by Assumption 4, we have

$$\begin{aligned}\mathbb{E}\left\|\nu^*(d') - \nu^*\right\|_2 &\geq \frac{1}{2\bar{b}^2 L_1}\left|\mathbb{E}\left(b_t\tilde{x}_t(\nu^*(d'))\right)_i - \mathbb{E}\left(b_t\tilde{x}_t(\nu^*)\right)_i\right| = \frac{1}{2\bar{b}^2 L_1}\left|d_i' - \mathbb{E}\left(b_t\tilde{x}_t(\nu^*)\right)_i\right| \\ &\geq \frac{1}{2\bar{b}^2 L_1}\left(\left|d_i - \mathbb{E}\left(b_t\tilde{x}_t(\nu^*)\right)_i\right| - |d_i' - d_i|\right).\end{aligned}$$

Denote the minimum of remaining resources in non-binding dimensions by

$$\gamma = \min_{i\in I_{\mathrm{NB}}}\left\{d_i - \mathbb{E}\left(b_t\tilde{x}_t(\nu^*)\right)_i\right\}.$$

By Lemma 7 we have

$$\|d - d'\|_2 \geq \sigma_{\min}\underline{\mathcal{L}}_f\mathbb{E}\left\|\nu^*(d') - \nu^*\right\|_2 \geq \frac{\sigma_{\min}\underline{\mathcal{L}}_f}{2\bar{b}^2 L_1}\left(\gamma - |d_i' - d_i|\right) \geq \frac{\sigma_{\min}\underline{\mathcal{L}}_f}{2\bar{b}^2 L_1}\left(\gamma - \|d - d'\|_2\right),$$

i.e., $\|d-d'\|_2 \geq \frac{\gamma\sigma_{\min}\underline{\mathcal{L}}_f}{\sigma_{\min}\underline{\mathcal{L}}_f+2\bar{b}^2L_1}$. Combined with (12.20), taking

$$\delta_d = \frac{1}{\sqrt{m}} \cdot \left(\frac{\gamma\sigma_{\min}\underline{\mathcal{L}}_f}{\sigma_{\min}\underline{\mathcal{L}}_f + 2\bar{b}^2 L_1}\right) \wedge \left(\frac{1}{2}\sigma_{\min}\underline{\mathcal{L}}_f\underline{\lambda}\right) \wedge \delta_0 \wedge \frac{\underline{\lambda}}{2\bar{\mathcal{L}}_r},$$

we can conclude that when $|d_i - d_i'| \leq \delta_d$, the binding/non-binding dimensions will never change. Moreover, enlarging the constraint in a non-binding dimension will never change this constraint to the binding dimension. So, for the non-binding dimensions, $d_i' - d_i$ can be any large. This finishes the proof.

12.10. Proof of lemma 4

Proof of this lemma under frequent resolving is similar in spirit as that in Li and Ye (2021), but with different induction and dual convergence rate. Here we focus on the infrequent re-solving scheme in Algorithm 4. All the proof for the infrequent re-solving scheme is valid for the frequent resolving case. we define $d_t = d_{T_j}$ if $T_j \leq t < T_{j+1}$. Then d_t will only update when $t = T_j$ for some j. Without loss of generality, we assume $\frac{1}{\rho}$ is a integer and $T = \left(\frac{1}{\rho}\right)^K$ for some integer K. Denote the ratio $C_\rho = \rho/(1-\rho)$. Also assume $C_4 = O(mn\log m)$ in Condition 1. Decomposing the perturbation of dual variables will lead to

$$\mathbb{E}\left[\sum_{t=1}^{\tau}\|\nu_{t-1}-\nu^*\|_2^2\right] \leq 2\mathbb{E}\left[\sum_{t=1}^{\tau}\|\nu_{t-1}-\nu^*(d_{t-1})\|_2^2 + \|\nu^*(d_{t-1})-\nu^*\|_2^2\right].$$

By the definition of stopping time τ and Condition 1, the first term in the RHS has

$$\mathbb{E}\left[\sum_{t=1}^{\tau}\|\nu_{t-1}-\nu^*(d_{t-1})\|_2^2\right] \leq \sum_{j=1}^{J} 2C_4\frac{1}{T_j}\cdot(T_{j+1}-T_j) \text{ or } C_4\frac{1}{T-T_j}\cdot(T_{j+1}-T_j) \leq 2C_4(\log T + 1).$$

Notice that, we need a good initialization: $\mathbb{E}\|\nu_0 - \nu^*\|^2 = O(1/T)$ to reach this condition for the first $(1-\rho)T$ terms in infrequent resolving case. For the second term, we apply lemma 7 to it.

$$2\mathbb{E}\left[\sum_{t=1}^{\tau}\|\nu^*(d_{t-1})-\nu^*\|_2^2\right] \leq \frac{2}{\sigma_{\min}^2\underline{\mathcal{L}}_f^2}\mathbb{E}\left[\sum_{t=1}^{\tau}\sum_{i\in I_{\mathrm{B}}}(d_{it}-d_i)^2\right].$$

Thus we transform the perturbation of $\nu^*(d_t)$ into the deviation of d_t in the binding dimensions.

To ease our analysis, we define a new sequence d_t'

$$d_t' = \begin{cases} d_t, & \text{if } t \leq \tau \\ d_{t-1}, & \text{if } t > \tau \end{cases},$$

which shares the same stopping time with d_t and define $\tau_i = \min_t\{T - \left\lceil \frac{m\bar{b}}{\underline{d}} \right\rceil\} \cup \{t | d_{it}' \notin \mathcal{D}_i\}$ for $i \in [m]$ as the stopping time on each dimension with $\tau = \min\{\tau_1, ..., \tau_m\}$. We follow a similar approach as Li and Ye (2021) to bound the stopping time. The distinction is that our analysis is more refined and we use a martingale argument that is different from Li and Ye (2021) to study the infrequent resolving scheme.

We first consider the binding dimensions. For any $i \in I_{\mathrm{B}}$, we derive:

$$
\begin{aligned}
d'_{i,T_{j+1}} &= d'_{i,T_j} + \frac{\sum_{k=T_j+1}^{T_{j+1}} \left[d'_{i,T_j} - \left(b_k \tilde{x}_k(\nu_{T_j})\right)_i \right]}{T - T_{j+1}} \mathbb{I}(\tau > T_j) \\
\mathbb{E}\left(d'_{i,T_{j+1}} - d_i\right)^2 &= \mathbb{E}\left(d'_{i,T_j} - d_i\right)^2 + \underbrace{\mathbb{E}\frac{\left(\sum_{k=T_j+1}^{T_{j+1}} \left[d'_{i,T_j} - \left(b_k \tilde{x}_k(\nu_{T_j})\right)_i \right]\right)^2}{(T - T_{j+1})^2} \mathbb{I}(\tau > T_j)}_{A'} \\
&\quad + \underbrace{2\mathbb{E}\frac{\left(d'_{i,T_j} - d_i\right)\left(\sum_{k=T_j+1}^{T_{j+1}} \left[d'_{i,T_j} - \left(b_k \tilde{x}_k(\nu^*(d_{T_j}))\right)_i \right]\right)}{T - T_{j+1}} \mathbb{I}(\tau > T_j)}_{B'} \\
&\quad + \underbrace{2\mathbb{E}\frac{\left(d'_{i,T_j} - d_i\right)\left(\sum_{k=T_j+1}^{T_{j+1}} \left(b_k \tilde{x}_k(\nu_{T_j})\right)_i - \left(b_k \tilde{x}_k(\nu^*(d_{T_j}))\right)_i\right)}{T - T_{j+1}} \mathbb{I}(\tau > T_j)}_{C'} .
\end{aligned} \tag{12.21}
$$

For the term A' we have

$$
\begin{aligned}
A' &= \mathbb{E}\frac{\left(\sum_{k=T_j+1}^{T_{j+1}} \left[d'_{i,T_j} - \mathbb{E}\left[\left(b_k \tilde{x}_k(\nu_{T_j})\right)_i | \mathcal{H}_{T_j}\right] + \mathbb{E}\left[\left(b_k \tilde{x}_k(\nu_{T_j})\right)_i | \mathcal{H}_{T_j}\right] - \left(b_k \tilde{x}_k(\nu_{T_j})\right)_i \right]\right)^2}{(T - T_{j+1})^2} \mathbb{I}(\tau > T_j) \\
&\leq 2\mathbb{E}\frac{\left(\sum_{k=T_j+1}^{T_{j+1}} d'_{i,T_j} - \mathbb{E}\left[\left(b_k \tilde{x}_k(\nu_{T_j})\right)_i | \mathcal{H}_{T_j}\right]\right)^2}{(T - T_{j+1})^2} \\
&\quad + 2\mathbb{E}\frac{\left(\sum_{k=T_j+1}^{T_{j+1}} \mathbb{E}\left[\left(b_k \tilde{x}_k(\nu_{T_j})\right)_i | \mathcal{H}_{T_j}\right] - \left(b_k \tilde{x}_k(\nu_{T_j})\right)_i\right)^2}{(T - T_{j+1})^2} \\
&\leq 2\mathbb{E}\frac{\sum_{k=T_j+1}^{T_{j+1}} \left(\mathbb{E}\left[\left(b_k \tilde{x}_k(\nu^*(d_{T_j}))\right)_i - \left(b_k \tilde{x}_k(\nu_{T_j})\right)_i \middle| \mathcal{H}_{T_j}\right]\right)^2}{T - T_{j+1}} \\
&\leq 2\bar{b}^4 L_1^2 \mathbb{E}\left\|\nu^*(d_{T_j}) - \nu_{T_j}\right\|_2^2 \\
&\leq \frac{C_4 \bar{b}^4 L_1^2 + 4nD^2\bar{b}^2}{T - T_j} + \frac{C_4 \bar{b}^4 L_1^2}{T_j}.
\end{aligned}
$$

For the second inequality, since $i \in I_{\mathrm{B}}$ and $d_t \in \sigma(\mathcal{H}_t)$, conditioned on past history $\mathcal{H}_{T_j}$, we always have

$$
d'_{i,T_j} - \mathbb{E}\left[\left(b_k \tilde{x}_k(\nu^*(d_{T_j}))\right)_i \middle| \mathcal{H}_{T_j}\right] = 0, \text{ for any } k \geq T_j + 1.
$$

This also indicates that $B' = 0$. For the third inequality, we use Assumption 4. For the term C', we apply Assumption 4 and Condition 1:

$$
\begin{aligned}
C' &= 2\mathbb{E}\left[\mathbb{E}\frac{\left(d'_{i,T_j} - d_i\right)\left(\sum_{k=T_j+1}^{T_{j+1}} \left(b_k \tilde{x}_k(\nu_{T_j})\right)_i - \left(b_k \tilde{x}_k(\nu^*(d_{T_j}))\right)_i\right)}{T - T_{j+1}} \mathbb{I}(\tau > T_j) \middle| \mathcal{H}_{T_j}\right] \\
&\leq 2\bar{b}^2 L_1 \mathbb{E}\frac{\left|d'_{i,T_j} - d_i\right| \sum_{k=T_j+1}^{T_{j+1}} \left\|\nu_{T_j} - \nu^*(d_{T_j})\right\|}{T - T_{j+1}} \\
&\leq 2\sqrt{C_4}\bar{b}^2 L_1 \sqrt{\mathbb{E}\left(d'_{i,T_j} - d_i\right)^2} \sqrt{\frac{1}{T - T_j} + \frac{1}{T_j}}.
\end{aligned}
$$

Here the first inequality is because of Assumption 4, and the second inequality is from Condition 1 and Cauchy inequality. Here in the derivation, we can treat $\{\nu_t\}$ as a new sequence generated by $\{d'_t\}$, which has the same value with the original one when $t \le \tau$, and takes $\nu_t = \mathcal{B}_t(\mathcal{H}_t, d'_t)$ when $t > \tau$. We then get the recurrence relation of $d'_{i,T_j} - d_i$:

$$\begin{aligned}\mathbb{E}\left(d'_{i,T_{j+1}} - d_i\right)^2 \le & \mathbb{E}\left(d'_{i,T_j} - d_i\right)^2 + \frac{C_4 \bar{b}^4 L_1^2 + 4nD^2\bar{b}^2}{T - T_j} + \frac{C_4\bar{b}^4 L_1^2}{T_j} \\ & + 2\sqrt{C_4}\bar{b}^2 L_1 \sqrt{\mathbb{E}\left(d'_{it} - d_i\right)^2}\sqrt{\frac{1}{T-T_j} + \frac{1}{T_j}}.\end{aligned}$$

Since $d_0 = d$, assigning $C_4 = O(mn\log m)$ and by induction we have

$$\mathbb{E}\left(d'_{i,T_j} - d_i\right)^2 \le C_5 C_\rho mn \log m D^2 \bar{b}^4 L_1^2 \frac{\rho^{-j} - 1}{T}$$

So, we have

$$\begin{aligned}2\mathbb{E}\left[\sum_{t=1}^{\tau} \|\nu^*(d_{t-1}) - \nu^*\|_2^2\right] \le & 2\mathbb{E}\left[\sum_{t=1}^{\tau}\sum_{i\in I_{\mathrm{B}}} (d_{i,t-1} - d_i)^2\right] \\ \le & 2m\mathbb{E}\sum_{j=1}^{J}(T_j - T_{j-1})\left[\left(d'_{i,T_j} - d_i\right)^2\right] \le C_5 C_\rho m^2 n \log m D^2 \bar{b}^4 L_1^2 \log T + C,\end{aligned}$$

$$\text{and } \mathbb{E}\left[\sum_{t=1}^{\tau}\|\nu_{t-1} - \nu^*\|_2^2\right] \le \left(2C_2 + \frac{mC_3}{\underline{\mathcal{L}}_D^2}\right)\log T + 2C_2,$$

Notice that, by Condition 1, C_4 can take as small as $O(L_1^2 mn\log m/(\sigma_{\min}^2 \underline{\mathcal{L}}_f^2))$, which completes the proof that

$$\mathbb{E}\left[\sum_{t=1}^{\tau}\|\nu_t - \nu^*\|_2^2\right] \le C \cdot C_\rho \frac{m^2 n \log m D^2 \bar{b}^4 L_1^4 \log T}{\sigma_{\min}^2 \underline{\mathcal{L}}_f^2} + C.$$

In the following discussion, we will treat C_ρ as a constant and thus can be omitted. In the case of frequent resolving, we only need to substitute the recurrence relation of $d'_{it} - d_i$ with:

$$\mathbb{E}\left(d'_{i,t+1} - d_i\right)^2 \le \mathbb{E}\left(d'_{it} - d_i\right)^2 + \frac{\left(\bar{d} + \sqrt{n}D\bar{b}\right)^2}{(T-t-1)^2} + \frac{2\sqrt{2C_2 L_1}\bar{b}^2\sqrt{\frac{1}{t+1} + \frac{1}{T-t}}\sqrt{\mathbb{E}\left(d'_{it} - d_i\right)^2}}{T-t-1},$$

which can be derived by the same analysis as the infrequent resolving case. Since $d_0 = d$, by induction we have $\mathbb{E}\left(d'_{it} - d_i\right)^2 \le C_3 \frac{t+1}{(T+1)(T-t)}$, where $C_3 = \left(2 \cdot \left(\bar{d} + \sqrt{n}D\bar{b}\right)^2 \vee \left(2\sqrt{2C_2 L_1}\bar{b}^2\right) + 1\right)^2$. The lemma still holds.

12.11. Proof of lemma 5

We prove this Lemma under an infrequent resolving setting. The frequent resolving can be handled analogously. Since $\tau = \min\{\tau_1, ..., \tau_m\}$, we only need to show $\mathbb{E}(T - \tau_i) \le C\log T$ for any i in binding dimensions and non-binding dimensions. By the definition of ρ, and τ, the algorithm will only hit the stopping time after the first update of d_t, i.e, $\tau > T_1$. For the binding dimensions, applying Chebyshev's inequality, we have

$$\begin{aligned}\mathbb{E}(T-\tau_i) \leq \sum_{i=1}^{T} \mathbb{P}\left(\tau_i \leq t\right) \leq 1+\frac{\sqrt{n} D \bar{b}}{\underline{d}}+\sum_{j=2}^{J} \mathbb{P}\left(\tau_i \leq T_j\right)\left(T_j-T_{j-1}\right) \\ \leq 1+\frac{\sqrt{n} D \bar{b}}{\underline{d}}+\sum_{j=2}^{J} \mathbb{P}\left(\left|d_{i, T_j}-d_i\right| \geq \delta_d\right)\left(T_j-T_{j-1}\right) \\ \leq 1+\frac{\sqrt{n} D \bar{b}}{\underline{d}}+\sum_{j=2}^{J}\left(\frac{\mathbb{E}\left(d_{i, T_{j+1}}^{\prime}-d_i\right)^2}{\delta_d^2}\right)\left(\rho^{j-1}(1-\rho) T\right) \\ \leq C+\frac{\sqrt{n} D \bar{b}}{\underline{d}}+\frac{C_5 m^2 n D^2 \bar{b}^4 L_1^4}{\delta_d^2} \log T .\end{aligned} \tag{12.22}$$

For the non-binding dimensions, $\mathcal{D}$ ensures that binding/non-binding dimensions remain unchanged when $d' \in \mathcal{D}$. Then for $d' \in \mathcal{D}$, we define

$$\tilde{d}_i^{\prime}= \begin{cases}d_i^{\prime}, & \text { if } i \in I_{\mathrm{B}} \\ d_i-\delta_d, & \text { if } i \in I_{\mathrm{NB}}\end{cases} .$$

We know that $\nu^*(d')=\nu^*(\tilde{d}')$ because the non-binding constraints are loose, then

$$\mathbb{E}\left(b_t \tilde{x}_t\left(\nu^*\left(d^{\prime}\right)\right)\right)_i=\mathbb{E}\left(b_t \tilde{x}_t\left(\nu^*\left(\tilde{d}^{\prime}\right)\right)\right)_i<d_i-\delta_d .$$

Recall that $\tilde{x}_k(\cdot) \perp\!\!\!\perp \mathcal{H}_t$, thus $\mathbb{E}\left[\left(b_k \tilde{x}_k\left(\nu^*\left(d_t\right)\right)_i \mid \mathcal{H}_t\right]<d_i-\delta_d\right.$ for any $k \geq t+1$, $i \in I_{\mathrm{NB}}$ and $d_t \in \mathcal{D}$. This implies that

$$\begin{aligned}\mathbb{P}\left(\tau_i \leq T_j\right) \leq & \mathbb{P}\left(\sum_{t=1}^{t^{\prime}}\left(b_t \tilde{x}_t\left(\nu_{t-1}\right)\right)_i \geq t^{\prime}\left(d_i-\delta_d\right)+T \delta_d \text { for some } 1 \leq t^{\prime} \leq T_j\right) \\ \leq & \mathbb{P}\left(\sum_{t=1}^{t^{\prime}}\left[\left(b_t \tilde{x}_t\left(\nu_{t-1}\right)\right)_i-\mathbb{E}\left[\left(b_t \tilde{x}_t\left(\nu^*\left(d_{t-1}\right)\right)_i \mid \mathcal{H}_{t-1}\right]\right] \geq T \delta_d \text { for some } 1 \leq t^{\prime} \leq T_j\right)\right. \\ \leq & \mathbb{P}\left(\sum_{t=1}^{t^{\prime}}\left[\left(b_t \tilde{x}_t\left(\nu_{t-1}\right)\right)_i-\mathbb{E}\left[\left(b_t \tilde{x}_t\left(\nu_{t-1}\right)\right)_i \mid \mathcal{H}_{t-1}\right]\right]+\right. \\ & \left.\sum_{t=1}^{t^{\prime}}\left|\mathbb{E}\left[\left(b_t \tilde{x}_t\left(\nu^*\left(d_{t-1}\right)\right)_i \mid \mathcal{H}_{t-1}\right]-\mathbb{E}\left[\left(b_t \tilde{x}_t\left(\nu_{t-1}\right)\right)_i \mid \mathcal{H}_{t-1}\right]\right| \geq T \delta_d \text { for some } 1 \leq t^{\prime} \leq T_j\right)\right. \\ \leq & \mathbb{P}\left(\sum_{t=1}^{t^{\prime}}\left[\left(b_t \tilde{x}_t\left(\nu_{t-1}\right)\right)_i-\mathbb{E}\left[\left(b_t \tilde{x}_t\left(\nu_{t-1}\right)\right)_i \mid \mathcal{H}_{t-1}\right]\right] \geq \frac{T \delta_d}{2} \text { for some } 1 \leq t^{\prime} \leq T_j\right) \\ & +\mathbb{P}\left(\sum_{t=1}^{t^{\prime}}\left|\mathbb{E}\left[\left(b_t \tilde{x}_t\left(\nu^*\left(d_{t-1}\right)\right)_i \mid \mathcal{H}_{t-1}\right]-\mathbb{E}\left[\left(b_t \tilde{x}_t\left(\nu_{t-1}\right)\right)_i \mid \mathcal{H}_{t-1}\right]\right| \geq \frac{T \delta_d}{2} \text { for some } 1 \leq t^{\prime} \leq T_j\right) .\right.\end{aligned}$$

Since sequences in the last two lines are martingales/sub-martingales, we use Doob's martingale inequality and get the following derivation:

$$\begin{aligned}\mathbb{P}\left(\tau_i \leq T_j\right) \leq & \frac{4}{T^2\delta_d^2}\sum_{t=1}^{T_j}\mathbb{E}\left[(b_t\tilde{x}_t\left(\nu_{t-1}\right))_i - \mathbb{E}\left[(b_t\tilde{x}_t\left(\nu_{t-1}\right))_i|\mathcal{H}_{t-1}\right]\right]^2 \\ & + \frac{4}{T^2\delta_d^2}\mathbb{E}\sum_{t=1}^{T_j}\left[|\mathbb{E}\left[(b_t\tilde{x}_t(\nu^*(d_{t-1}))_i|\mathcal{H}_{t-1}\right] - \mathbb{E}\left[(b_t\tilde{x}_t\left(\nu_{t-1}\right))_i|\mathcal{H}_{t-1}\right]|\right]^2.\end{aligned} \tag{12.23}$$

By Assumption 4, we have

$$\begin{aligned}\mathbb{P}\left(\tau_i \leq T_j\right) \leq & \frac{16n\bar{b}^2D^2T_j}{T^2\delta_d^2} + \frac{8L_1^2\bar{b}^4}{T^2\delta_d^2}\sum_{t=1}^{T_j}\mathbb{E}\left\|\nu_{t-1} - \nu^*(d_{t-1})\right\|_2^2 \\ \leq & \frac{16n\bar{b}^2D^2T_j}{T^2\delta_d^2} + \frac{C_5m^2n\log mD^2\bar{b}^4L_1^4\log T}{T^2\delta_d^2},\end{aligned}$$

where for $j=1$, we use the assumption on initialization. We now go back to calculate the $\mathbb{E}(T-\tau_i)$:

$$\begin{aligned}\mathbb{E}(T-\tau_i) \leq & 2 + \frac{\sqrt{n}D\bar{b}}{\underline{d}} + \sum_{j=2}^{J}\mathbb{P}\left(\tau_i \leq T_j\right)\left(T_j - T_{j-1}\right) \\ \leq & C + \frac{\sqrt{n}D\bar{b}}{\underline{d}} + \sum_{j=2}^{J}\frac{16n\bar{b}^2D^2T_j(T_j-T_{j-1})}{T^2\delta_d^2} + \frac{C_5(T_j-T_{j-1})}{T^2\delta_d^2}(m^2n\log mD^2\bar{b}^4L_1^2\log T) \\ \leq & C + \frac{\sqrt{n}D\bar{b}}{\underline{d}} + \frac{C_5n\bar{b}^2D^2\log T}{\delta_d^2} + \frac{C_5m^2n\log m\bar{b}^4D^2L_1^4\log T}{\delta_d^2}.\end{aligned} \tag{12.24}$$

Notice that, by Condition 1, C_4 can take as small as $O(L_1^2mn\log m/(\sigma_{\min}^2\underline{\mathcal{L}}_f^2))$. Putting together (12.22) and (12.24) we conclude the proof of lemma 5:

$$\mathbb{E}(T-\tau) \leq \frac{Cm^2n\log m\bar{b}^4D^2L_1^4\log T}{\delta_d^2\sigma_{\min}^2\underline{\mathcal{L}}_f^2}.$$

12.12. Proof of Theorem 3

Equipped with Lemma 4 and 5, we now continue sketching the proof of Theorem 3. By Proposition 3, it suffices to bound the two terms there.

Proof of Theorem 3 The proof continues from Proposition 3.

Step 1: bounding R.1. By Fenchel conjugate function, we re-write the bridging function $h(\nu)$ by

$$\begin{aligned}h(\nu) = & \mathbb{E}\left[f_t(\tilde{x}_t(\nu)) + r\left(\tilde{a}(\mu)\right) + \left(\tilde{a}(\mu) - b_t\tilde{x}_t(\nu)\right)^\top\mu^* + \left(d - b_t\tilde{x}_t(\nu)\right)^\top\lambda^*\right] \\ = & \mathbb{E}\left[f_t^*(b_t^\top(\lambda+\mu)) + r^*(-\mu)\right] + \mathbb{E}(\mu^*-\mu)^\top(\tilde{a}(\mu) - b_t\tilde{x}_t(\nu)) + \mathbb{E}(\lambda^*-\lambda)^\top(d - b_t\tilde{x}_t(\nu)) \\ = & \mathbb{E}\left[f_t^*(b_t^\top(\lambda+\mu)) + r^*(-\mu)\right] - \mathbb{E}\left[g_t(b_t^\top(\lambda+\mu))^\top b_t^\top(\lambda+\mu-\lambda^*-\mu^*)\right. \\ & \left. -\vartheta^*(-\mu)^\top(\mu-\mu^*) + d^\top(\lambda-\lambda^*)\right] \\ = & s(\nu,d) - \left\langle\nabla_\nu D(\nu,\mu^*,d) - \nabla_\nu D(\nu^*,\mu^*,d), (\nu-\nu^*)\right\rangle.\end{aligned}$$

By Assumption 2 and 3, we get

$$
\begin{aligned}
h(\nu^*)-h(\nu) =& s(\nu^*,d)-s(\nu,d)+\langle\nabla_\nu D(\nu,\mu^*,d)-\nabla_\nu D(\nu^*,\mu^*,d),(\nu-\nu^*)\rangle\\
&+\langle\nabla s(\nu,d)-\nabla s(\nu^*,d),\nu-\nu^*\rangle\le\left(2\bar{\mathcal{L}}_s-\underline{\mathcal{L}}_s\right)\|\nu-\nu^*\|_2^2\\
&=(\bar{b}^2\bar{\mathcal{L}}_f-\frac{1}{2}\sigma_{\min}\underline{\mathcal{L}}_f)\|\nu-\nu^*\|_2^2.
\end{aligned}
$$

Then Lemma 4 gives rise to the following bound.

$$
\mathbb{E}\left[\sum_{t=1}^{\tau}h(\nu^*)-h(\nu_{t-1})\right]\le O(\log T).
$$

Step 2: bounding R.2. This term can be controlled by the definition of stopping time and Lemma 5.

$$
\begin{aligned}
&\mathbb{E}\left[2(\bar{f}+\bar{r}+C_3)(T-\tau)+\left\langle\lambda^*,\sum_{t=1}^{\tau}(d-b_tx_t)\right\rangle\right]\\
&=\mathbb{E}\left[2(\bar{f}+\bar{r}+C_3)(T-\tau)+\langle\lambda^*,d_\tau(T-\tau)-d(T-\tau)\rangle\right]\\
&\le\mathbb{E}\left[2(\bar{f}+2\bar{r}+C_3)(T-\tau)+\sum_{i\in I_{\mathrm{B}}}\lambda_i^*(d_i+\delta_d)(T-\tau)\right]\\
&\le(2\bar{f}+2\bar{r}+2C_3+(\|d\|+\sqrt{m}\delta_d)\frac{2(\bar{f}+\bar{r})}{\underline{d}})\mathbb{E}(T-\tau)=O(\log T).
\end{aligned}
$$

Thus we finish the proof.

Step 3: bounding R.3. This term requires the most effort. It concerns the combined effects of variable splitting and complementary slackness. The following lemma is important for bounding this term.

LEMMA 8. *Suppose $i\in I_{NB}$. Under Assumptions 1-5, Algorithm 1 with selected dual optimizer $\{\mathcal{B}_t\}_{t\ge1}$ satisfying Condition 1 and stopping time (5.5) ensures*

$$
\mathbb{E}\left\|\hat{\mu}_{I_{NB},T}-\mu^*_{I_{NB}}\right\|_2^2\le O(\frac{\log T}{T}),\ \textit{and}\ \mathbb{E}\left\|\sum_{t=1}^{\tau}(\tilde{a}(\mu^*)-b_tx_t)\right\|_2^2\le O(T\log T).
$$

The proof Lemma 8 exploits the local smoothness of r and $\tilde{x}_t$ with the help of the optimality of μ^*, i.e., $\tilde{a}(\mu^*)=\mathbb{E}b_t\tilde{x}_t(\nu^*)$. We first check the binding dimensions: for all the binding dimensions $i\in I_{\mathrm{B}}$, since $\mathbb{E}(b_t\tilde{x}_t(\nu^*))_i=d_i$, by the definition of stopping time, it can be shown that

$$
\begin{aligned}
\mathbb{E}\left\langle\hat{\mu}_{I_{\mathrm{B}},T}-\mu^*_{I_{\mathrm{B}}},\sum_{t=1}^{\tau}\tilde{a}(\mu^*)_{I_{\mathrm{B}}}-(b_tx_t)_{I_{\mathrm{B}}}\right\rangle=&\mathbb{E}\left\langle\hat{\mu}_{I_{\mathrm{B}},T}-\mu^*_{I_{\mathrm{B}}},\sum_{t=1}^{\tau}d_{I_{\mathrm{B}}}-(b_tx_t)_{I_{\mathrm{B}}}\right\rangle\\
&\le\mathbb{E}\sqrt{mn}GD\bar{b}(d+\delta_d)(T-\tau)\\
&=O(\log T).
\end{aligned}\tag{12.25}
$$

We then go back to the non-binding dimension $i\in I_{\mathrm{NB}}$ with the help of Lemma 8. By Cauchy–Schwarz inequality, we get

$$
\mathbb{E}\left[\left\langle\hat{\mu}_{I_{\mathrm{NB}},T}-\mu^*_{I_{\mathrm{NB}}},\sum_{t=1}^{\tau}\tilde{a}(\mu^*)_{I_{\mathrm{NB}}}-(b_tx_t)_{I_{\mathrm{NB}}}\right\rangle\right]\le\left(\mathbb{E}\left\|\hat{\mu}_{I_{\mathrm{NB}},T}-\mu^*_{I_{\mathrm{NB}}}\right\|_2^2\mathbb{E}\left\|\sum_{t=1}^{\tau}(\mathbb{E}b_t\tilde{x}_t(\nu^*)-b_tx_t)_{I_{\mathrm{NB}}}\right\|_2^2\right)^{1/2}.\tag{12.26}
$$

Thus R.3 can be controlled by $\log T$. The proof is concluded. The constant factor is as large as

$$\mathring{C} \lesssim \frac{m^2 n \log m \bar{b}^6 D^4 \bar{d}^2 L_1^6}{\delta_d^2 \sigma_{\min}^2 \underline{\mathcal{L}}_f^2} \cdot \left((\bar{f} + \bar{r} + \sqrt{mn} G D \bar{b}) \vee \frac{n \bar{b}^2 G^2 + m \bar{\mathcal{L}}_r^2 \delta_0^2}{\delta_0^2} \right)$$

from the proof of Theorem 1 and Lemma 8. Generally, taking $m \lesssim n$ we have the order of $\mathring{C} = O(m^2 n^2 \log m)$

12.13. Proof of Lemma 8

For the $\mathbb{E} \left\| \hat{\mu}_{I_{\text{NB}},T} - \mu^*_{I_{\text{NB}}} \right\|_2^2$, $i \in I_{\text{NB}}$, the optimality of μ^* implies $\tilde{a}(\mu^*) = \mathbb{E} b_t \tilde{x}_t(\nu^*)$, thus by conjugate we have

$$\begin{aligned}
\mathbb{E} \left\| \hat{\mu}_{I_{\text{NB}},T} - \mu^*_{I_{\text{NB}}} \right\|_2^2 &= \mathbb{E} \left\| \nabla_{I_{\text{NB}}} r\left(\tilde{a}(\mu^*)\right) - \nabla_{I_{\text{NB}}} r\left(\frac{\sum_{t=1}^T b_t x_t}{T} \right) \right\|_2^2 \\
&\le n \bar{b}^2 G^2 \mathbb{P} \left(\left\| \mathbb{E} b_t \tilde{x}_t(\nu^*) - \frac{\sum_{t=1}^T b_t x_t}{T} \right\| \ge \delta_0 \right) + \mathbb{E} m \bar{\mathcal{L}}_r^2 \left\| \tilde{a}(\mu^*) - \frac{\sum_{t=1}^T b_t x_t}{T} \right\|_2^2 \\
&\le \frac{n \bar{b}^2 G^2 + m \bar{\mathcal{L}}_r^2 \delta_0^2}{\delta_0^2} \mathbb{E} \left\| \frac{\sum_{t=1}^T b_t x_t - \mathbb{E} b_t \tilde{x}_t(\nu^*)}{T} - \frac{\sum_{t=1}^\tau \mathbb{E}\left[b_t \tilde{x}_t(\nu_{t-1}) | \mathcal{H}_{t-1} \right]}{T} + \frac{\sum_{t=1}^\tau \mathbb{E}\left[b_t \tilde{x}_t(\nu_{t-1}) | \mathcal{H}_{t-1} \right]}{T} \right\|_2^2 \\
&\le 3 C_0 \, \mathbb{E} \left\| \frac{\sum_{t=1}^\tau b_t \tilde{x}_t(\nu_{t-1}) - \mathbb{E}\left[b_t \tilde{x}_t(\nu_{t-1}) | \mathcal{H}_{t-1} \right]}{T} \right\|_2^2 \text{ (part 12.13.1)} \\
&+ 3 C_0 \, \mathbb{E} \left\| \frac{\sum_{t=1}^\tau \mathbb{E}\left[b_t \tilde{x}_t(\nu_{t-1}) | \mathcal{H}_{t-1} \right] - \mathbb{E} b_t \tilde{x}_t(\nu^*)}{T} \right\|_2^2 \text{ (part 12.13.2)} \\
&+ 3 C_0 \, \mathbb{E} \left\| \frac{\sum_{t=\tau+1}^T b_t x_t - \mathbb{E} b_t \tilde{x}_t(\nu^*)}{T} \right\|_2^2 \text{ (part 12.13.3).}
\end{aligned}$$

For the part 12.13.1, notice that this is a martingale series. Thus we have

$$\begin{aligned}
\mathbb{E} \left\| \frac{\sum_{t=1}^\tau b_t \tilde{x}_t(\nu_{t-1}) - \mathbb{E}\left[b_t \tilde{x}_t(\nu_{t-1}) | \mathcal{H}_{t-1} \right]}{T} \right\|_2^2 &\le \frac{\sum_{t=1}^T \operatorname{Var}(b_t \tilde{x}_t(\nu_{t-1}) - \mathbb{E}\left[b_t \tilde{x}_t(\nu_{t-1}) | \mathcal{H}_{t-1} \right])}{T^2} \\
&\le \frac{n D^2 \bar{b}^2}{T}.
\end{aligned}$$

For the part 12.13.2, applying Assumption 4 we can yield

$$\begin{aligned}
\mathbb{E} \left\| \frac{\sum_{t=1}^\tau \mathbb{E}\left[b_t \tilde{x}_t(\nu_{t-1}) | \mathcal{H}_{t-1} \right] - b_t \tilde{x}_t(\nu^*)}{T} \right\|_2^2 &\le \mathbb{E} \frac{\tau \sum_{t=1}^\tau \left\| \mathbb{E}\left[b_t \tilde{x}_t(\nu_{t-1}) | \mathcal{H}_{t-1} \right] - b_t \tilde{x}_t(\nu^*) \right\|_2^2}{T^2} \\
&\le \frac{\mathbb{E} \sum_{t=1}^T \left\| \mathbb{E}\left[b_t \tilde{x}_t(\nu_{t-1}) | \mathcal{H}_{t-1} \right] - b_t \tilde{x}_t(\nu^*) \right\|_2^2 \mathbb{I}(t \le \tau)}{T} \\
&= \frac{\sum_{t=1}^T \mathbb{E} \left\| \left[\mathbb{E}\left[b_t \tilde{x}_t(\nu_{t-1}) - b_t \tilde{x}_t(\nu^*) | \mathcal{H}_{t-1}, b_t \right] \mathbb{I}(t \le \tau) \right] \right\|_2^2}{T} \\
&\overset{\text{(a)}}{\le} L_1^2 \bar{b}^2 \frac{\sum_{t=1}^T \mathbb{E}\left[\left\| \nu_{t-1} - \nu^* \right\|_2^2 \mathbb{I}(t \le \tau) \right]}{T} = L_1^2 \bar{b}^2 \frac{\mathbb{E}\left[\sum_{t=1}^\tau \left\| \nu_{t-1} - \nu^* \right\|_2^2 \right]}{T} \\
&\overset{\text{(b)}}{\le} \frac{L_1^6 \bar{b}^2 m^2 n \log m D^2 \bar{b}^4 \log T}{\sigma_{\min}^2 \underline{\mathcal{L}}_f^2 T}.
\end{aligned}$$

(a) is by Assumption 4 and the fact $\{t \le \tau\} \in \sigma(\mathcal{H}_{t-1})$, and $\tilde{x}_t(\cdot) \perp\!\!\!\perp \nu_{t-1}$. (b) is by Lemma 4.

For the part 12.13.3, since $\|b_t x_t - \mathbb{E} b_t \tilde{x}_t(\nu^*)\|_2^2 \le n D^2 \bar{b}^2$, by Lemma 5, we have

$$\begin{aligned}\mathbb{E}\left\|\frac{\sum_{t=\tau+1}^T b_t x_t - \mathbb{E} b_t \tilde{x}_t(\nu^*)}{T}\right\|_2^2 \le & \frac{\mathbb{E}\left[(T-\tau)\sum_{t=\tau+1}^T \|b_t x_t - \mathbb{E} b_t \tilde{x}_t(\nu^*)\|^2\right]}{T^2} \\ \le & nD^2\bar{b}^2 \frac{\mathbb{E}(T-\tau)}{T} \le \frac{Cm^2n^2 \log m \bar{b}^6 D^4 L_1^6}{\delta_d^2 \sigma_{\min}^2 \mathcal{L}_f^2} \frac{\log T}{T}.\end{aligned}$$

We then go back to control the next term $\mathbb{E}\left\|\sum_{t=1}^{\tau}(\tilde{a}(\mu^*) - b_t x_t)\right\|_2^2$. This can be bounded exactly by part 12.13.1 and 12.13.2. From the argument above, we show that $\mathbb{E}\left\|\sum_{t=1}^{\tau}(\tilde{a}(\mu^*) - b_t x_t)\right\|_2^2$ is controlled by $O(T\log T)$. Thus we finish the proof.

12.14. Proof of Theorem 4 and 5

Suppose $m \le n$. We specify a non-regularized case where $f_t(x) = -\sum_{i=1}^n \left[\frac{1}{4}(x_i - 2\xi_{it})^2 + \xi_{it}^2\right]$, with fixed cost $b_t = [I_m, \mathbf{0}_{m\times(n-m)}]$, average resource capacity $d = \frac{1}{2}D\mathbf{1}_m$, and ξ_{it} are i.i.d for every t, but may be different for each dimension i, following a distribution within $[\frac{1}{2}D, \frac{3}{4}D]$ and we take $\mathbb{E}\xi_{it} = \frac{3}{8}D$ for example. Then the dual problem is

$$D_t(\lambda) = \sum_{i=1}^m D_{it}(\lambda_i), \text{where } D_{it}(\lambda_i) = \begin{cases} \frac{1}{2}D\lambda_i & \text{if } \lambda_i > \xi_{it} \\ -\frac{1}{4}D + \xi_{it} - \frac{1}{2}D\lambda_i & \text{if } \lambda_i < \xi_{it} - \frac{1}{2}D \\ \lambda_i^2 - 2(\xi_t - \frac{1}{4}D)\lambda_i + \xi_{it}^2 & \text{if } \xi_{it} - \frac{1}{2}D \le \lambda_i \le \xi_{it} \end{cases}.$$

Suppose λ^* is the optimal solution to the deterministic problem $\min_{\lambda\ge 0} D(\lambda) = \mathbb{E}D_t(\lambda)$. Without loss of generality, we assume that our dual variable λ_i is taken within $\left[\frac{1}{4}D, \frac{1}{2}D\right]$ since we know that $\lambda_i^* = \mathbb{E}\xi_t - \frac{1}{4}D = \frac{3}{8}D \in \left[\frac{1}{4}D, \frac{1}{2}D\right]$. Then, we have

$$D_t(\lambda) = f_t^*(\lambda) + d^\top \lambda = \sum_{i=1}^m \left[\lambda_i^2 - 2(\xi_{it} - \frac{1}{4}D)\lambda_i + \xi_{it}^2\right].$$

For the dual-based police $\{\lambda_t\}_{t=0}^{T-1}$, the corresponding primal variable for $1 \le i \le m$ is $x_{it} = \tilde{x}_{it}(\lambda_{i,t-1}) = 2\xi_{it} - 2\lambda_{i,t-1}$ or void if the resource is depleted. We have the following regret:

$$\begin{aligned}\text{Regret}(A) =& R^*(\mathcal{P}) - R(A|\mathcal{P}) \\ =& \mathbb{E}\left[\max_{x_t\in[0,D]}\left\{\sum_{t=1}^T f_t(x_t) \text{ s.t. } \sum_{t=1}^T b_t x_t \le \frac{1}{2}D\mathbf{1}_m T\right\}\right] - \mathbb{E}\left[\sum_{t=1}^T f_t(x_t)\right] \\ =& \mathbb{E}\left[\min_{\lambda\ge 0}\left\{\sum_{t=1}^T D_t(\lambda)\right\}\right] - \mathbb{E}\left[\sum_{t=1}^T f_t(x_t)\right] \\ =& \mathbb{E}\left[\sum_{t=1}^T D_t(\lambda_T^*)\right] - \mathbb{E}\left[\sum_{t=1}^T f_t(x_t)\right].\end{aligned}$$

Define the corresponding

$$h(\lambda) = \mathbb{E}\left[f_t(\tilde{x}_t(\lambda)) + \langle d - b_t\tilde{x}_t(\lambda), \lambda^*\rangle\right] = D(\lambda) - \langle \nabla D(\lambda), \lambda - \lambda^*\rangle.$$

We have $h(\lambda^*) = D(\lambda^*)$ and $h(\lambda^*) - h(\lambda) = (\lambda^* - \lambda)^2$. For the quadratic function D_t, we always have $D_t(\lambda_1) - D_t(\lambda_2) = \nabla D_t(\lambda_2)(\lambda_1 - \lambda_2) + (\lambda_1 - \lambda_2)^2$. Thus it follows that

$$\begin{aligned}
\text{Regret}(A) &= \mathbb{E}\left[\sum_{t=1}^{T} D_t(\lambda_T^*)\right] - TD(\lambda^*) + TD(\lambda^*) - \mathbb{E}\left[\sum_{t=1}^{T} f_t(x_t)\right] \\
&= \mathbb{E}\left[\sum_{t=1}^{T} D_t(\lambda_T^*) - D_t(\lambda^*)\right] + Th(\lambda^*) - \mathbb{E}\left[\sum_{t=1}^{T} f_t(x_t)\right] \\
&= -\mathbb{E}\left[\sum_{t=1}^{T} [\nabla D_t(\lambda_T^*)(\lambda^* - \lambda_T^*)] + T(\lambda^* - \lambda_T^*)^2\right] + Th(\lambda^*) - \mathbb{E}\left[\sum_{t=1}^{T} f_t(x_t)\right] \\
&= -T\mathbb{E}(\lambda^* - \lambda_T^*)^2 + Th(\lambda^*) - \mathbb{E}\left[\sum_{t=1}^{T} f_t(x_t)\right].
\end{aligned}$$

By the dual convergence in Theorem 1, we know that the first term $T\mathbb{E}(\lambda^* - \lambda_T^*)^2$ can be bounded by a constant. Now, we handle the second term by controlling the stopping time. Define the stopping time $\tau_0 = \min_i \left\{t \in [T] \middle| \sum_{j=1}^{t} x_{i,j} \geq \frac{1}{2}DT - D\right\} \cup \{T\}$. Then when $t \leq \tau_0$, we always have $x_{it} = \tilde{x}_{it}(\lambda_{i,t-1}) = 2\xi_{it} - 2\lambda_{i,t-1}$, and $0 \leq \sum_{t=\tau_0+1}^{T} x_{it} \leq D$ for $t > \tau$. Then we have

$$\begin{aligned}
\mathbb{E}\left[\sum_{t=1}^{T} f_t(x_t)\right] &\leq \mathbb{E}\left[\sum_{t=1}^{\tau_0} f_t(\tilde{x}_t(\lambda_{t-1})) + \langle \frac{1}{2}D\mathbf{1}_m - b_t\tilde{x}_t(\lambda_{t-1}), \lambda^*\rangle\right] + \mathbb{E}\left[\sum_{t=\tau_0+1}^{T} f_t(x_t) + \langle \frac{1}{2}D\mathbf{1}_m - x_t(\lambda), \lambda^*\rangle\right] \\
&\leq \mathbb{E}\sum_{t=1}^{\tau} h(\lambda_{t-1}) + \mathbb{E}\left[\sum_{t=\tau_0+1}^{T} \frac{3}{4}D\mathbf{1}_m^\top x_t + \frac{1}{2}D\mathbf{1}_m^\top \lambda^*\right] \\
&\leq \mathbb{E}\sum_{t=1}^{\tau} h(\lambda_{t-1}) + \frac{3}{16}mD^2\mathbb{E}[T - \tau_0] + \frac{3}{4}mD^2.
\end{aligned}$$

The first inequality is because of the resource constraint, and the second one is because $f_t(x) \leq \nabla f_t(0)(x - 0) \leq \frac{3}{4}D\mathbf{1}_m^\top x$. If we specify $\lambda_{i,t-1} = \xi_{i,t}$ when the resource constraints are violated, we also have $\mathbb{E}\left[\sum_{t=1}^{T} f_t(x_t)\right] \leq \mathbb{E}\sum_{t=1}^{T} h(\lambda_{t-1})$. Then

$$\begin{aligned}
Th(\lambda^*) - \mathbb{E}\left[\sum_{t=1}^{T} f_t(x_t)\right] &\geq \mathbb{E}\left[\sum_{t=1}^{\tau_0} h(\lambda^*) - h(\lambda_{t-1})\right] + \mathbb{E}(h(\lambda^*) - \frac{3}{16}mD^2)\mathbb{E}[T - \tau_0] - \frac{3}{4}mD^2 \\
&= \mathbb{E}\left[\sum_{t=1}^{\tau_0} (\lambda^* - \lambda_{t-1})^2\right] + \frac{5}{64}mD^2\mathbb{E}[T - \tau_0] - \frac{3}{4}mD^2,
\end{aligned} \tag{12.27}$$

or $Th(\lambda^*) - \mathbb{E}\left[\sum_{t=1}^{T} f_t(x_t)\right] \geq \mathbb{E}\left[\sum_{t=1}^{T} (\lambda^* - \lambda_{t-1})^2\right]$. Applying van Trees inequality to the estimation of λ^* (Li and Ye 2021) by changing the distribution of ξ_{it}, we can prove the Theorem 4. To prove the

Theorem 5, we only need to show the stopping time $\mathbb{E}[T-\tau_0] \geq \Omega(\sqrt{T})$ given the convergence condition. To this end, we consider the two-point distribution $\mathbb{P}(\xi_{it} = \frac{1}{2}D) = \mathbb{P}(\xi_{it} = \frac{3}{4}D) = \frac{1}{2}$. This proof is inspired by Arlotto and Gurvich (2019). Denote $t' = \lfloor T - \sqrt{T} \rfloor$. We show that $\mathbb{P}(\tau_0 \leq t')$ is larger that a constant c so that $\mathbb{E}\tau_0 \leq (1-c)T + c(T-\sqrt{T}) \leq T - c\sqrt{T}$. Now, we focus on one dimension and assume $m = n = 1$. The stopping time of one dimension can be extended to multiple dimensions. This gives

$$\begin{aligned}
\mathbb{P}(\tau_0 \leq t') &= \mathbb{P}\left(\sum_{t=1}^{t'} 2(\xi_t - \lambda_{t-1}) \geq \frac{DT}{2} - D\right) \\
&\geq \mathbb{P}\left(\left\{\sum_{t=1}^{t'} 2(\xi_t - \lambda^*) \geq \frac{DT}{2} - D + \varepsilon D\sqrt{t'}\right\} \cap \left\{\sum_{t=1}^{t'} |\lambda_{t-1} - \lambda^*| < \varepsilon D\sqrt{t'}\right\}\right) \\
&\geq \mathbb{P}\left(\left\{\sum_{t=1}^{t'} 2(\xi_t - \lambda^*) \geq \frac{DT}{2} - D + \varepsilon D\sqrt{t'}\right\}\right) - \mathbb{P}\left(\sum_{t=1}^{t'} |\lambda_{t-1} - \lambda^*| \geq \varepsilon D\sqrt{t'}\right).
\end{aligned}$$

With the condition $\mathbb{E}|\lambda_t - \lambda^*| \leq c_2 D/\sqrt{t+1}$, we have $\mathbb{P}\left(\sum_{t=1}^{t'} |\lambda_{t-1} - \lambda^*| \geq \varepsilon D\sqrt{t'}\right) \leq \frac{2c_2}{\varepsilon}$ by Chebyshev's inequality. Then it holds that

$$\begin{aligned}
\mathbb{P}(\tau_0 \leq t') &\geq \mathbb{P}\left(\left\{\sum_{t=1}^{t'} 2(\xi_t - \lambda^*) \geq \frac{DT}{2} - D + \varepsilon D\sqrt{t'}\right\}\right) - \frac{2c_2}{\varepsilon} \\
&= \mathbb{P}\left(\left\{\sum_{t=1}^{t'} \frac{4}{D}(\xi_t - \frac{D}{2}) \geq \frac{t'}{2} + (T-t') - 2 + 2\varepsilon\sqrt{t'}\right\}\right) - \frac{2c_2}{\varepsilon} \\
&\geq \mathbb{P}\left(\left\{\sum_{t=1}^{t'} \frac{4}{D}(\xi_t - \frac{D}{2}) \geq \frac{t'}{2} + (1+2\varepsilon)\sqrt{t'}\right\}\right) - \frac{2c_2}{\varepsilon},
\end{aligned}$$

where $\sum_{t=1}^{t'} \frac{4}{D}(\xi_t - \frac{D}{2})$ follows the binomial distribution $B(t', \frac{1}{2})$, with mean $\mu = \frac{t'}{2}$ and standard deviation $\sigma = \frac{\sqrt{t'}}{2}$. The second inequality is because $T - t' \leq \sqrt{T} + 1$ and $\sqrt{T} - \sqrt{t'} \leq \sqrt{T} - \sqrt{T - \sqrt{T}} = \frac{\sqrt{T}}{\sqrt{T-\sqrt{T}}+\sqrt{T}} \leq 1$. For the binomial distribution, $\mathbb{P}(X \geq \mu + x\sigma)$ converge to $\Phi(-x)$ for any x with known $O(\frac{1}{\sqrt{n}})$ speed by Berry-Esseen CLT where $\Phi(x)$ is the distribution function of standard normal distribution. We let $c_2 = \sup_{\varepsilon>0} \varepsilon\Phi(-2-4\varepsilon)/4$. Then there exists $\varepsilon_0 > 0$ such that when T is large enough, $\mathbb{P}\left(\left\{\sum_{t=1}^{t'} \frac{4}{D}(\xi_t - \frac{D}{2}) \geq \frac{t'}{2} + (1+2\varepsilon_0)\sqrt{t'}\right\}\right) \geq \frac{3c_2}{\varepsilon_0}$, which indicates that $\mathbb{P}(\tau_0 \leq t') \geq \frac{c_2}{\varepsilon_0}$. This makes our proof complete.

12.15. Proof of Theorem 6

Theorem 6 can be proved following the same path as in the proof of Theorem 3. The key is that, with a good initialization, Lemma 4 and 5 still hold. This has actually been mentioned in the proof of Lemma 4 and 5 and thus omitted here.

12.16. Proof of Theorem 7

To prove the $O(\log^2 T)$ bound, we revise the previous proof of Lemma 4 and 5 and re-compute some important rates. For online gradient descent, without loss of generality, suppose $d_t \in \Omega_d$ before the stopping time τ. Then By Assumption 3 and the stochastic gradient descent approach in Rakhlin et al. (2012), for $t \geq T_j$,we have

$$\begin{aligned}\mathbb{E}\left(\left\|\nu_{t+1}-\nu^*(d_{T_j})\right\|_2^2+\left\|\mu_{t+1}-\mu^*(d_{T_j})\right\|_2^2\right) &\leq \mathbb{E}\left\|\nu_t-\eta_{t+1}\nabla D_{t+1,\nu}(\nu_t,\mu_t,d_{T_j})-\nu^*(d_{T_j})\right\|_2^2\\ &+\mathbb{E}\left\|\mu_{t+1}-\mu^*(d_{T_j})-\eta_{t+1}\nabla D_{t+1,\mu}(\nu_t,\mu_t,d_{T_j})\right\|_2^2,\end{aligned}$$

and we also have

$$\begin{aligned}&\mathbb{E}\left\|\nu_{t+1}-\nu^*(d_{T_j})\right\|_2^2\leq \mathbb{E}\left\|\nu_t-\eta_{t+1}\nabla D_{\nu}(\nu_t,\mu_t,d_{T_j})-\nu^*(d_{T_j})\right\|_2^2, \text{ by projection}\\ &\mathbb{E}\left\|\nu_{t+1}-\nu^*(d_{T_j})\right\|_2^2\leq (1-2\eta_{t+1}\sigma_{\min}\underline{\mathcal{L}}_f)\mathbb{E}\left\|\nu_t-\nu^*(d_{T_j})\right\|_2^2+\eta_{t+1}^2 n\bar{b}^2D^2.\end{aligned}$$

And hence, by choosing $\eta_t = 1/(\sigma_{\min}\underline{\mathcal{L}}_f(t-T_j+1))$, we have

$$\mathbb{E}\left[\left\|\nu_t-\nu^*(d_{T_j})\right\|_2^2\middle|\mathcal{H}_{T_j}\right]\leq \frac{4n\bar{b}^2D^2}{\sigma_{\min}^2\underline{\mathcal{L}}_f^2(t-T_j+1)}, \tag{12.28}$$

where $\nu^*(d_{T_j})$ is part of the optimal solution to $D(\boldsymbol{\lambda}, d_{T_j})$. See Rakhlin et al. (2012) for a detailed approach. By this convergence rate, we have the following rates

$$\mathbb{E}\left[\sum_{t=1}^{\tau}\|\nu_{t-1}-\nu^*(d_{t-1})\|_2^2\right]\leq \sum_{j=1}^{J}\frac{4n\bar{b}^2D^2(\log(T_{j+1}-T_j)+1)}{\sigma_{\min}^2\underline{\mathcal{L}}_f^2}\leq \frac{Cn\bar{b}^2D^2(\log^2 T)}{\sigma_{\min}^2\underline{\mathcal{L}}_f^2}.$$

We now revisit the previous induction of d_t in (12.21). For the three parts A', B', C', we have the new rates:

$$\begin{aligned}A' &= \mathbb{E}\frac{\left(\sum_{k=T_j+1}^{T_{j+1}}\left[d'_{i,T_j}-\mathbb{E}\left[\left(b_k\tilde{x}_k(\nu^*(d_{T_j}))\right)_i|\mathcal{H}_{T_j}\right]+\mathbb{E}\left[\left(b_k\tilde{x}_k(\nu^*(d_{T_j}))\right)_i|\mathcal{H}_{T_j}\right]-(b_k\tilde{x}_k(\nu_{k-1}))_i\right]\right)^2}{(T-T_{j+1})^2}\mathbb{I}(\tau>T_j)\\ &\leq 2\mathbb{E}\frac{\sum_{k=T_j+1}^{T_{j+1}}\left(\mathbb{E}\left[\left(b_k\tilde{x}_k(\nu^*(d_{T_j}))\right)_i-(b_k\tilde{x}_k(\nu_{k-1}))_i\middle|\mathcal{H}_{T_j}\right]\right)^2}{T-T_{j+1}}\\ &\leq \frac{2n\bar{b}^4L_1^2D^2}{\sigma_{\min}^2\underline{\mathcal{L}}_f^2}\frac{\log(T_{j+1}-T_j)}{T-T_{j+1}}\\ &\leq \frac{C_4n\bar{b}^4L_1^2D^2\log T}{\sigma_{\min}^2\underline{\mathcal{L}}_f^2\rho^jT}.\end{aligned}$$

And for B', we also have:

$$d'_{i,T_j}-\mathbb{E}\left[\left(b_k\tilde{x}_k(\nu^*(d_{T_j}))\right)_i\middle|\mathcal{H}_{T_j}\right]=0, \text{ for any } k\geq T_j+1.$$

For part C', it follows that

$$\begin{aligned} C' &= 2\mathbb{E}\left[\mathbb{E}\frac{\left(d'_{i,T_j}-d_i\right)\left(\sum_{k=T_j+1}^{T_{j+1}}(b_k\tilde{x}_k(\nu_{k-1}))_i-\left(b_k\tilde{x}_k(\nu^*(d_{T_j}))\right)_i\right)}{T-T_{j+1}}\mathbb{I}(\tau>T_j)\middle|\mathcal{H}_{T_j}\right] \\ &\le 2\bar{b}^2L_1\mathbb{E}\frac{\left|d'_{i,T_j}-d_i\right|\sum_{k=T_j+1}^{T_{j+1}}\left\|\nu_{k-1}-\nu^*(d_{T_j})\right\|}{T-T_{j+1}} \\ &\le \frac{C\sqrt{n}\bar{b}^3DL_1}{\sigma_{\min}\underline{\mathcal{L}}_f}\sqrt{\mathbb{E}\left(d'_{i,T_j}-d_i\right)^2}\sqrt{\frac{1}{T-T_j}} \\ &\le \frac{C\sqrt{n}\bar{b}^3DL_1}{\sigma_{\min}\underline{\mathcal{L}}_f}\sqrt{\mathbb{E}\left(d'_{i,T_j}-d_i\right)^2}\sqrt{\frac{1}{\rho^jT}}. \end{aligned}$$

Thus, we then get the recurrence relation of $d'_{i,T_j}-d_i$:

$$\begin{aligned} \mathbb{E}\left(d'_{i,T_{j+1}}-d_i\right)^2 \le& \mathbb{E}\left(d'_{i,T_j}-d_i\right)^2+\frac{C_4n\bar{b}^4L_1^2D^2\log T}{\sigma_{\min}^2\underline{\mathcal{L}}_f^2\rho^jT} \\ &+\frac{C\sqrt{n}\bar{b}^3DL_1}{\sigma_{\min}\underline{\mathcal{L}}_f}\sqrt{\mathbb{E}\left(d'_{i,T_j}-d_i\right)^2}\sqrt{\frac{1}{\rho^jT}}. \end{aligned}$$

Since $d_0=d$, by induction we have

$$\mathbb{E}\left(d'_{i,T_j}-d_i\right)^2 \le C_5C_\rho nD^2\bar{b}^6L_1^2\frac{\log T}{\sigma_{\min}^2\underline{\mathcal{L}}_f^2\rho^jT}.$$

So, we have

$$\begin{aligned} 2\mathbb{E}\left[\sum_{t=1}^{\tau}\|\nu^*(d_{t-1})-\nu^*\|^2\right] \le& 2\mathbb{E}\left[\sum_{t=1}^{\tau}\sum_{i\in I_{\mathrm{B}}}(d_{i,t-1}-d_i)^2\right] \\ \le& 2m\mathbb{E}\sum_{j=1}^{J}(T_j-T_{j-1})\left[\left(d'_{i,T_j}-d_i\right)^2\right]\le C_5\frac{mnD^2\bar{b}^6L_1^2\log^2T}{\sigma_{\min}^2\underline{\mathcal{L}}_f^2}+C, \end{aligned}$$

and $\mathbb{E}\left[\sum_{t=1}^{\tau}\|\nu_{t-1}-\nu^*\|_2^2\right]\le\frac{Cmn\bar{b}^6D^2L_1^2(\log^2T)}{\sigma_{\min}^2\underline{\mathcal{L}}_f^2}+2C_2,$

Extending this computation to the control of stopping time, we also have

$$\mathbb{E}(T-\tau)\le\frac{Cmn\bar{b}^6D^2L_1^2\log^2T}{\delta_d^2\sigma_{\min}^2\underline{\mathcal{L}}_f^2},$$

where the proof essentially follows Lemma 5. Similarly, following the proof of Lemma 8, we also have

$$\mathbb{E}\left\|\hat{\mu}_{I_{\mathrm{NB}},T}-\mu^*_{I_{\mathrm{NB}}}\right\|_2^2\le O\left(\frac{\log^2T}{T}\right),\ \text{and}\ \mathbb{E}\left\|\sum_{t=1}^{\tau}(\tilde{a}(\mu^*)-b_tx_t)\right\|_2^2\le O(T\log^2T).$$

Combining these together, we have the regret upper bound

$$\mathrm{Regret}(A)\le\widetilde{C}\log^2T,$$

for some constant $\widetilde{C}=O(mn^2)$ depending on the values in Assumptions 1-5. Here, the additional n factor comes from the complementary slackness bound by our problem's scale, as shown in Lemma 8.

13. Numerical Experiments

The implementation details on multiple input models are as follows: the dual updates are calculated by closed-form solutions to Equation (4.3) under input I-III and by *cvxpy* (Diamond and Boyd (2016)) under input IV. See Table 1 for parameter settings of different inputs. For each T, we randomly generate T observations from distribution, run each algorithm in an online fashion, and keep a record of their output. The regret is calculated as the difference between the offline optimal (solved by *cvxpy*) and the online output. We report the average regret over 10 repetitions for all the following graphs.

Input	$f_t(x)$	$r(x)$	b_{it}	d_i
I	$a_t^\top x$	$-\kappa \left\|x - d/2\right\|_2^2$	$\mathrm{U}(0,1)$	0.1
II	$a_t^\top x$	$-\kappa \left\|x - d/2\right\|_2^2$	Bernoulli(p_i)	$\mathrm{U}(0.25, 0.75)$
III	$-\frac{1}{4}x^2 + \xi_t x$	0	1	0.5
IV	$-\frac{a_t}{4}x^2 + \frac{a_t x}{2}$	$\kappa \min_i x_i/d_i$	$\mathrm{U}(0, 0.5)$	0.3

Table 1 Parameter Settings of Inputs

Input model I: Online welfare maximization with costs, independent reward, and resource consumption. The reward functions are linear as $f_t(x) = a_t^\top x$. The regularization function is the ℓ_2 loss $r(x) = -\kappa \left\|x - d/2\right\|_2^2$, which corresponds to the application of online welfare maximization with square costs. The reward coefficients a_t's and the constraint coefficients b_t's are i.i.d. random vectors with dimension $m = 6$. Sepecifically, a_{it} is generated from the uniform distribution $U(0, 10)$, and b_{it} is generated from the uniform distribution $U(0, 1)$. κ is set to 0.001.

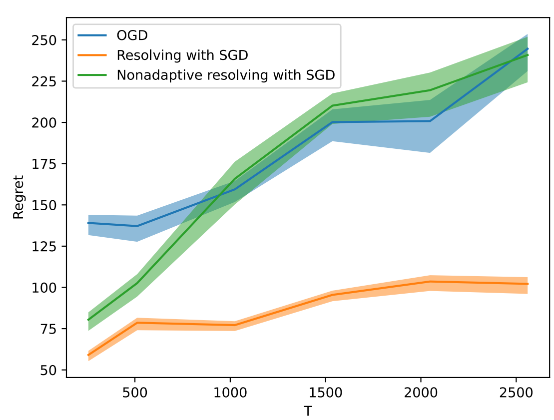

Figure 2 Regret versus horizon (T) under Input I. OGD stands for online gradient descent in Balseiro et al. (2020); resolving with SGD is our Algorithm 2; nonadaptive resolving with SGD is the nonadaptive version (i.e., without updating the constraints) of Algorithm 2.

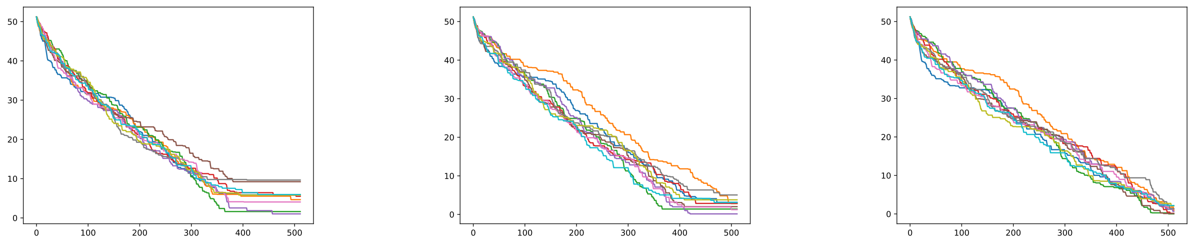

(a) Online gradient descent (b) Nonadaptive version of Algorithm 2 (c) Algorithm 2

Figure 3 Remaining resource of one binding dimension versus time under input model I with $T = 512$. Ten curves are displayed, each of which corresponds to one simulation.

To illustrate how the regret scales with the time horizon T, we evaluate the algorithms with different T chosen from $\{256, 512, 1024, 1536, 2048, 2560\}$. We find that Resolving with SGD (Algorithm 2) shows logarithmic regret, while its counterpart without constraint update ($d_t \equiv d$ in Equation 4.2) shows a much worse regret. We name the latter algorithm as the "Nonadaptive resolving with SGD". The online gradient descent (OGD) method in Balseiro et al. (2020) exhibits a $O(\sqrt{T})$ regret as indicated in their theoretical findings. The regret comparison between the algorithms can be found in Figure 2. In Figure 3, we plot the dynamics of resource consumption for one binding dimension of the aforementioned algorithms. Ten curves are displayed, each of which corresponds to one simulation. Being adaptive to the level of remaining resources, Algorithm 2 controls carefully the constraint consumption to ensure that the resources are consumed at a steady rate till they are used up. In comparison, both the OGD and the nonadaptive version of Algorithm 2 stop allocating resources too early, demonstrating the benefits of the constraint updates, which exploit the history of past actions.

Input model II: Online welfare maximization with costs, dependent reward and resource consumption. The parameter setting below is based on Balseiro et al. (2022). The reward functions and the regularization function are the same as in input I, whereas input II considers the case when the reward coefficients a_t are random variables conditional to the constraint coefficients b_ts. We set $a_t = \text{Proj}_{[0,10]}\{\theta_t^\top b_t + \delta_t \mathbf{1}\}$, where θ_t is generated from a multi-variate Gaussian distribution $N(0, \text{diag}(1))$, and δ_t is generated from the standard Gaussian distribution $N(0,1)$. The constraint coefficients b_{it}'s are generated from Bernoulli distribution with probability parameter p_i with $p_i = (1+\alpha)/2$, and α is generated from the beta distribution $\text{Beta}(1,3)$. The average resource constraints d_i's are generated from the uniform distribution $\text{U}(0.25, 0.75)$. κ is set to 0.001.

Similar to the setting of input I, we evaluate the algorithms under input II with different T's and fix $m = 6$. The regret performances and resource consumption are displayed in Figure 4 and Figure 5, respectively. Among the three algorithms (Algorithm 2, the nonadaptive Algorithm 2 and the OGD method in Balseiro et al. (2020)), Algorithm 2 achieves a logarithmic regret, the nonadaptive Algorithm 2 suffers from a higher regret while the regret of OGD grows in a much faster speed.

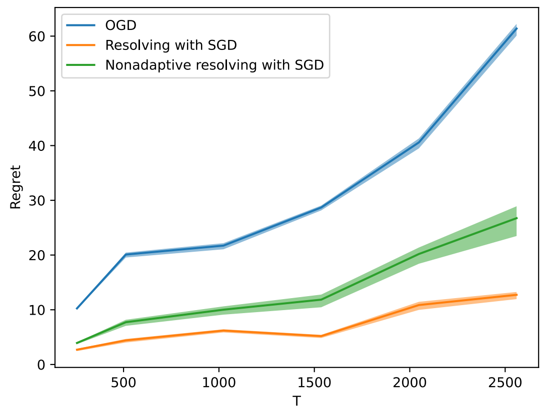

Figure 4 Regret versus horizon (T) under Input II. OGD stands for online gradient descent in Balseiro et al. (2020); resolving with SGD is Algorithm 2; nonadaptive resolving with SGD is the nonadaptive version of Algorithm 2.

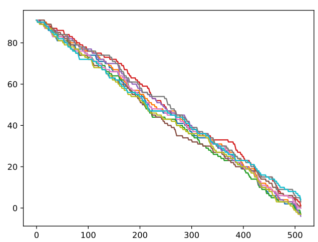

(a) Online gradient descent

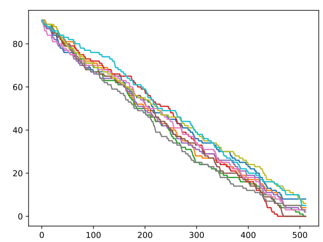

(b) Nonadaptive version of Algorithm 2

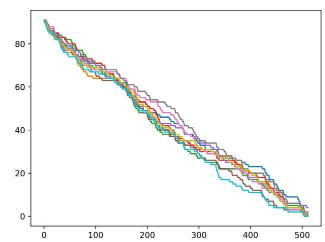

(c) Algorithm 2

Figure 5 Remaining resource of one binding dimension versus time under input model II with $T = 512$. Ten curves are displayed, each of which corresponds to one simulation.

Input model III: Non-regularized online convex resource allocation with one resource. In this model, we assess the algorithms' performance under a non-regularized special case, where there is only one resource, the reward function $f_t(x) = f_t(x, \xi_t) = -\frac{1}{4}x^2 + \xi_t x$, the constraint $d = \frac{1}{2}$ and cost $b_t = 1$. The random variable ξ_t follows a two-point distribution that takes value in $\{\frac{1}{2}, \frac{3}{4}\}$ with equal probability, i.e., $\mathbb{P}[\xi_t = \frac{1}{2}] = \mathbb{P}[\xi_t = \frac{3}{4}] = 0.5$. This special case is used in the proof of Theorem 5.

For input model III, the optimal solution to Problem (2.8) admits a closed-form due to the simple distribution, which also leads to a relatively small regret compared to other input models. We compare further with the "No learning" algorithm, which is the convex version of Algorithm 1 in Li and Ye (2021). It requires the computation of optimal dual solutions, while neither Adaptive (Algorithm 2) nor OGD needs this information. The regret comparison is shown in Figure 6a. The empirical performance corroborates the theoretical results in that the infrequent resolving saves computation significantly with a relatively small performance loss. We further explain the reason for the performance advantage by plotting the remaining time before stopping in Figure 6b. Benchmark algorithms without constraint update (No learn-

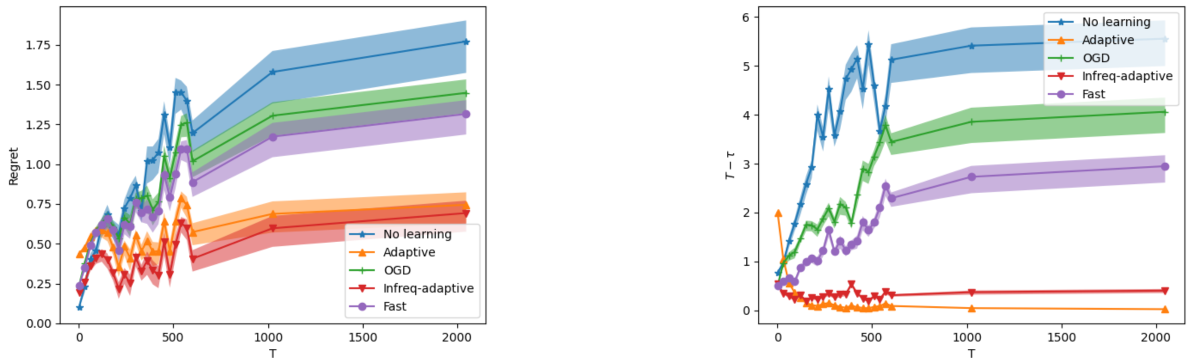

(a) Regret versus horizon (T) under input III. (b) Remaining time ($T - \tau$) versus horizon (T) under input III

Figure 6 Performance evaluation under input III. No learning is the convex version of Algorithm 1 in Li and Ye (2021); OGD is online gradient descent in Balseiro et al. (2020); Adaptive is Algorithm 2; infreq-adaptive is Algorithm 4; Fast is Algorithm 5.

ing, OGD) stop allocating resource $O(\sqrt{T})$ steps earlier than Adaptive (Algorithm 2) and Infreq-adaptive (Algorithm 4), which leads to the terrible regret performance. In comparison, Fast (Algorithm 5) uses both the updated constraints and the original constraints, leading to a less accurate dual solution compared to Algorithm 4 and worse performance.

Input model IV: Online convex allocation with max-min regularizer. To demonstrate that our algorithm also works well for nonsmooth regularizers, we set up input model IV. The reward functions are $f_t(x) = -\frac{a_t}{4}x^2 + \frac{a_t}{2}x$, where $a_t \in \mathbb{R}$ is generated from the uniform distribution $U(0,10)$. Each dimension of the tth constraint coefficient, i.e., $b_{it}, i \in [m]$, is generated from the uniform distribution $U(0,0.5)$. We set $m = 12$ in this input model. The regularization function is $r(x) = \kappa \min_i x_i/d_i$ so that the minimum allocation is not too small. κ is set to $[0, 0.01, 0.1]$ to compare the performance of algorithms under different regularization levels.

In Figure 7, regret curves under different κs of Nonadaptive (Algorithm 3 with OGD), Adaptive (Algorithm 2), Infreq-adaptive (Algorithm 4) and Fast (Algorithm 5) are plotted. The regret of Algorithm 2 and Algorithm 4 grow slower than the other two, which shows the advantage of using updated constraint information. It is observed in Figure 7d that the per-step reward decreases as we regularize the solution by setting $\kappa > 0$, which is consistent with the intuition that regularization has a negative impact on the revenue. Also, it is observed that the reward decreases when the regularization level further increases, but we also observe a slight increase for Fast (Algorithm 5) when κ increases from 0.01 to 0.1, which indicates that a proper regularization level may not decrease the reward too much. It would be interesting to study the problem of choosing the appropriate regularization level in the future.

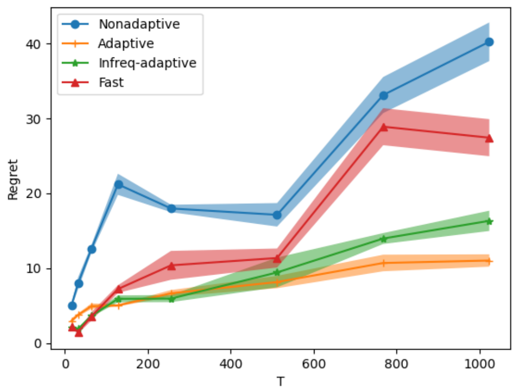

(a) $\kappa = 0$

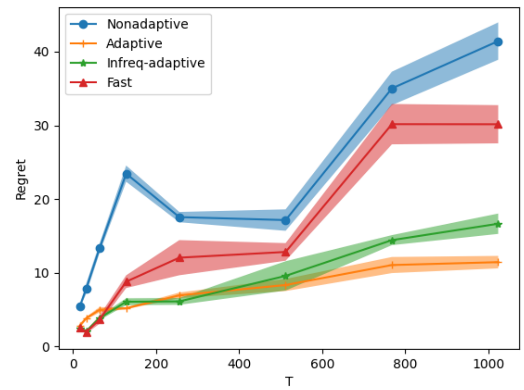

(b) $\kappa = 0.01$

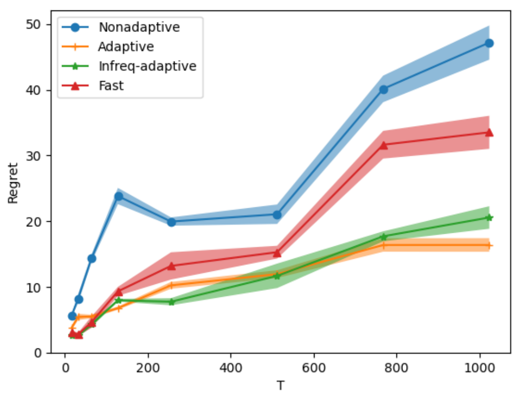

(c) $\kappa = 0.1$

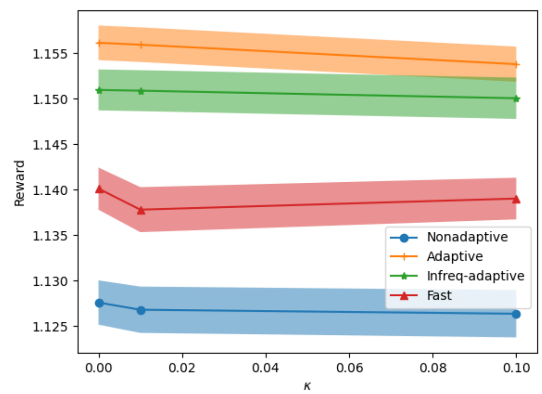

(d) Per-step reward ($\sum_t f_t(x_t)/T$) versus κ when $T = 1024$.

Figure 7 Regret versus horizon (T) under input IV with different κs in Figure 7a. Nonadaptive is Algorithm 3 with OGD; Adaptive is Algorithm 2; Infreq-adaptive is Algorithm 4; Fast is Algorithm 5. Figure 7d shows the impact of different regularization levels on the per-step reward (with 95% confidence interval).

14. More Discussions

14.1. Computational cost

Specifically, for strongly convex dual objectives, Our algorithm of frequent resolving requires computing gradients for $O(T^2)$ times in total; for more general dual objectives, it requires $O(T^4)$ times of gradient computation in total. With a good initialization, we can unevenly divide time horizon T into $\log T$ epochs and only solve empirical dual optimization at the beginning of each epoch. This infrequent algorithm only takes gradient computations for $O(T \log T)$ times for strongly convex problems (which is nearly linear) and $O(T^3 \log T)$ for general problems, while delivering an optimal regret bound. The fast algorithm we established has only linear computational cost $O(T)$, which is comparable with OGD but reaches sub-optimal regret $O(\log^2 T)$.

14.2. Fenchel conjugate of regularizers

Here, we provide Fenchel conjugates for three commonly used regularizers.

1. ℓ_1**-loss**: $r(a) := -\kappa \|a\|_1$. Define $\mathcal{Z} := \{a | \|a\|_\infty \leq L\}$, $h(a) = r(a) - \mu^\top a$ then $r^*(\mu) = \max_{a\in\mathcal{Z}}\{r(a) - \mu^\top a\} := \max_{a\in\mathcal{Z}}\{h(a)\}$. We have the subgradient: $\nabla h(a) = -\kappa\,\mathrm{sign}(a) - \mu$. Since $\|\mu\|_\infty \leq G = \kappa$, we know that $h(a)$ takes its maximum only when $a = 0$. the conjugate r^* in Ω_μ is thus of the form $r^*(\mu) = 0$ when $\kappa \geq G = \kappa$.
2. **Max-min loss**: $r(a) := \kappa \min_i(a_i/d_i)$. Define $\mathcal{Z} := \{a | |a_i| \leq d_i L,\ i \in [m]\}$, and $z_i = a_i/d_i$. Then we define the function to be maximized as $h(z) := r(a) - \mu^\top a = \kappa \min(z_i) - \sum \mu_i d_i z_i$, for the region $\|z\|_\infty \leq L$. By computing the subgradient of each dimension, we know that the optimal z that maximizes the $h(z)$ must have:

$$z_i = \begin{cases} L & \text{if } \mu_i d_i < 0 \\ -L & \text{if } \mu_i d_i > \kappa \\ \min_i(z_i) & \text{if } \mu_i d_i \in [0, \kappa] \end{cases}.$$

   Therefor, we have $z_i = \min_i(z_i)$ if $\mu_i \geq 0$. Moreover, whether $\min_i(z_i)$ takes L or $-L$ depends on the value $\kappa - \sum d_i (\mu_i)_+$. If $\kappa - \sum d_i (\mu_i)_+ > 0$, then we will have $\min_i(z_i) = L$, otherwise $\min_i(z_i) = -L$. Thus, the conjugate r^* in Ω_μ is of the form $r^*(\mu) = L \cdot \left(\left|\kappa - \sum d_i (\mu_i)_+\right| - \sum d_i (\mu_i)_-\right)$.
3. **Negative max loss**: $r(a) := -\kappa \max_i(a_i/d_i)$. Define $\mathcal{Z} := \{a | |a_i| \leq d_i L,\ i \in [m]\}$, and $z_i = a_i/d_i$. Then we have $h(z) := r(a) - \mu^\top a = -\kappa \max(z_i) - \sum \mu_i d_i z_i$. Following an analogous argument as in the max-min loss, we have $r^*(\mu) = L \cdot \left(\left|\kappa + \sum d_i (\mu_i)_-\right| + \sum d_i (\mu_i)_+\right)$.