Summary of R-Learner and DR-Learner Analysis using Orthogonal Statistical Learning

Yonghan Jung

University of Illinois Urbana-Champaign
yonghan@illinois.edu
https://www.yonghanjung.me/

September 4, 2025

We give a step-by-step review of empirical excess risk for Neyman-orthogonal losses. Each technical statement is followed by two working examples: a DR-learner for the ATE (AIPW squared loss) and a constant-effect R-learner.

Contents

1 Setup & Preliminaries

We study a population risk $L(\tau, \eta)$, where the *target* $\tau \in \mathcal{T}$ and the *nuisance* $\eta \in \mathcal{H}$ live in normed spaces $(\mathcal{T}, \|\cdot\|_{\mathcal{T}})$ and $(\mathcal{H}, \|\cdot\|_{\mathcal{H}})$, respectively. Throughout, η_0 denotes the true nuisance. We define the (possibly non-unique) *oracle minimizer*

$$\tau_0 \in \arg\min_{\tau \in \mathcal{T}} L(\tau, \eta_0), \tag{1}$$

which we assume is nonempty.

Directional derivatives. For a functional F and direction h, the (Gâteaux) derivative with respect to a variable x at x_0 is

$$\nabla_x F(x_0)[h] \triangleq \lim_{t \to 0} \frac{F(x_0 + th) - F(x_0)}{t}, \tag{2}$$

and second derivatives $\nabla_x^2 F(x_0)[h_1, h_2]$ are defined analogously; mixed derivatives such as $\nabla_\eta \nabla_\tau L$ will be used for orthogonality.

Sample splitting and plug-in. We assume a two-way split into independent folds of approximately equal size: one to learn $\hat{\eta}$ (using data $\mathcal{D}_\eta$), and one to learn $\hat{\tau}$ by minimizing $L(\tau, \hat{\eta})$ over τ, i.e.,

$$\tau^*_{\hat{\eta}} \triangleq \arg\min_{\tau} L(\tau, \hat{\eta}), \qquad \text{so that} \qquad \tau_0 = \tau^*_{\eta_0}.$$

This separation prevents overfitting-induced bias when we later linearize around (τ_0, η_0).

Target-class statistical term. Let $R_{\mathcal{T}}(\tau; \eta, \epsilon) \geq 0$ be a data-dependent rate function such that, with probability at least $1 - \epsilon$,

$$L(\tau, \eta) - L(\tau^*_\eta, \eta) \leq R_{\mathcal{T}}(\tau; \eta, \epsilon). \tag{3}$$

You may instantiate $R_{\mathcal{T}}$ via localized complexity (e.g., critical radius) or algorithm-specific bounds; we keep it abstract to highlight how nuisance error propagates into target error.

Goal and norms. Our goal is to upper bound the *target error* $\|\tau - \tau_0\|^2_{\mathcal{T}}$. When we write $\|\cdot\|_p$ we mean the $L_p(P)$ norm with respect to the underlying distribution.

1.1 Examples (R- and DR-learners)

We use standard notation: $T \in \{0, 1\}$ (treatment), X (covariates), Y (outcome). The estimand is the CATE

$$\tau_0(X) \triangleq \mathbb{E}[Y(1) - Y(0) \mid X], \tag{4}$$

under the usual *positivity* ($c \leq \pi_0(X) \leq 1 - c$ a.s.) and i.i.d. sampling. We assume

$$Y(t) \perp\!\!\!\perp T \mid X \implies \mathbb{E}[Y(t) \mid X] = \mathbb{E}[Y \mid t, X], \ \forall t \in \{0, 1\}. \tag{5}$$

1.1.1 R-Learner

The Robinson decomposition posits

$$Y = f_0(X) + T\,\tau_0(X) + \epsilon_Y, \quad \mathbb{E}[\epsilon_Y \mid T, X] = 0, \tag{6}$$

$$T = \pi_0(X) + \epsilon_X, \quad \mathbb{E}[\epsilon_X \mid X] = 0, \tag{7}$$

and with $m_0(X) \triangleq \mathbb{E}[Y \mid X]$ we have $m_0(X) = f_0(X) + \pi_0(X)\tau_0(X)$. Hence

$$Y - m_0(X) \;=\; (T - \pi_0(X))\,\tau_0(X) + \epsilon_Y. \tag{8}$$

Thus, viewing τ_0 as an OLS-type coefficient in a residualized regression, we define

$$L_{\mathrm{R}}(\tau, \eta_0 \triangleq \{m_0, \pi_0\}) \;\triangleq\; \mathbb{E}\Big[\{Y - m_0(X) - (T - \pi_0(X))\tau(X)\}^2\Big], \tag{9}$$

so that $\tau_0 \in \arg\min_\tau L_{\mathrm{R}}(\tau, \eta_0)$.

1.1.2 DR-Learner

We define following nuisances:

$$\mu_0(T, X) \triangleq \mathbb{E}[Y \mid T, X], \qquad \omega_0(T, X) \triangleq \frac{2T - 1}{P(T \mid X)}. \tag{10}$$

Define the pseudo-outcome

$$\varphi(V; \eta_0 \triangleq \{\mu_0, \pi_0\}) \triangleq \omega_0(T, X)\{Y - \mu_0(T, X)\} + \mu_0(1, X) - \mu_0(0, X), \tag{11}$$

and the squared-loss objective

$$L_{\mathrm{DR}}(\tau, \eta) \;\triangleq\; \mathbb{E}\Big[\{\varphi(V; \eta) - \tau(X)\}^2\Big]. \tag{12}$$

This loss is centered at the CATE in virtue of $\mathbb{E}[\varphi(V; \eta_0) \mid X] = \tau_0(X)$.

2 Assumptions

We now state structural conditions that yield fast rates. The exposition follows the orthogonal-statistical-learning (OSL) template: first-order optimality at truth, curvature in τ, and orthogonality to damp the impact of nuisance error.

Assumption 1 (**First-order optimality in** τ)**.** Moving away from τ_0 cannot reduce the population risk at the true nuisance:

$$\nabla_\tau L(\tau_0, \eta_0)[h_\tau] \;\geq\; 0 \quad \text{for all feasible directions } h_\tau \text{ from } \tau_0. \tag{13}$$

Assumption 2 (**Strong convexity (quadratic growth) in** τ)**.** There exist constants $\lambda > 0$, $\kappa \geq 0$, and $r \in [0, 1)$ such that, for any $\bar{\tau}$ on the line segment between τ and τ_0,

$$\nabla^2_\tau L(\bar{\tau}, \eta)[\tau - \tau_0, \tau - \tau_0] \;\geq\; \lambda\|\tau - \tau_0\|^2_{\mathcal{T}} \;-\; \kappa\|\eta - \eta_0\|^{\frac{4}{1+r}}_{\mathcal{H}}. \tag{14}$$

Rationale: The risk function $L(\tau,\eta)$ is in a bowl-shape over τ. The κ term allows mild curvature deterioration when $\eta \neq \eta_0$; the exponent $4/(1+r)$ is chosen to balance mixed terms via Young's inequality later.

2.1 Assumption checks for the examples

We verify that the R- and DR-losses satisfy the above, clarifying how positivity yields curvature and how residualization/DR construction yields orthogonality.

2.1.1 R-Learner: assumptions hold

First-order optimality. With $\tilde{Y} \triangleq Y - m_0(X)$, $\tilde{T} \triangleq T - \pi_0(X)$,

$$\nabla_\tau L_{\mathrm{R}}(\tau_0,\eta_0)[h_\tau] = -2\,\mathbb{E}\Big[(\tilde{Y} - \tilde{T}\tau_0)\,\tilde{T}\,h_\tau(X)\Big] = -2\,\mathbb{E}\Big[\mathbb{E}[\epsilon_Y \mid T,X]\tilde{T}\,h_\tau(X)\Big] = 0. \tag{15}$$

Hence Assumption 1 holds.

Strong convexity. We have

$$\nabla^2_\tau L(\bar{\tau},\eta)[\tau-\tau_0,\tau-\tau_0] = 2\mathbb{E}\Big[\{\tau(X)-\tau_0(X)\}^2\,\{T-\pi(X)\}^2\Big], \tag{16}$$

where

$$\mathbb{E}[(T-\pi(X))^2 \mid X] = \underbrace{\mathrm{Var}(T\mid X)}_{=\pi_0(X)(1-\pi_0(X))} + \big(\pi_0(X)-\pi(X)\big)^2 \;\geq\; \pi_0(X)(1-\pi_0(X)). \tag{17}$$

Therefore,

$$\begin{aligned}
\nabla^2_\tau L(\bar{\tau},\eta)[\tau-\tau_0,\tau-\tau_0] &= 2\mathbb{E}\Big[\{\tau(X)-\tau_0(X)\}^2\,\{T-\pi(X)\}^2\Big] && (18)\\
&\geq 2\mathbb{E}\Big[\{\tau(X)-\tau_0(X)\}^2\mathrm{Var}(T\mid X)\Big] && (19)\\
&= 2\mathbb{E}\Big[\{\tau(X)-\tau_0(X)\}^2\pi_0(X)\{1-\pi_0(X)\}\Big] && (20)\\
&\geq 2E\Big[\{\tau(X)-\tau_0(X)\}^2 c\{1-c\}\Big] && (21)\\
&= 2c(1-c)\|\tau-\tau_0\|_2^2. && (22)
\end{aligned}$$

Hence, Assumption 2 holds with $\lambda = 2c(1-c)$ and $\kappa = 0$ (taking $\|\cdot\|_{\mathcal{T}} = \|\cdot\|_2$).

2.1.2 DR-Learner: assumptions hold

First-order optimality.

$$\nabla_\tau L_{\mathrm{DR}}(\tau_0,\eta_0)[h_\tau] = -2\,\mathbb{E}[\{\varphi(V;\eta_0)-\tau_0(X)\}\,h_\tau(X)] = 0, \tag{23}$$

since $\mathbb{E}[\varphi(V;\eta_0)\mid X] = \tau_0(X)$.

Strong convexity. We first note that

$$\nabla_\tau L_{\mathrm{DR}}(\tau_0,\eta_0)[\tau-\tau_0] = -2\mathbb{E}[\{\varphi(V;\eta)-\tau(X)\}\{\tau(X)-\tau_0(X)\}]. \tag{24}$$

This gives

$$\nabla_\tau^2 L_{\mathrm{DR}}(\tau,\eta)[\tau-\tau_0,\tau-\tau_0] = 2\|\tau-\tau_0\|_2^2, \tag{25}$$

which shows that $\kappa = 0$ and $\lambda = 2$ (taking $\|\cdot\|_{\mathcal{T}} = \|\cdot\|_2$).

3 Main Result

Theorem 1 (**Fast Rate Convergence**)**.** Suppose Assumption 1 and 2 hold. Then,

$$\|\tau-\tau_0\|_{\mathcal{T}}^2 \leq \tfrac{2}{\lambda} R_{\mathcal{T}}(\hat\tau;\hat\eta,\epsilon) + \tfrac{2}{\lambda}\{\nabla_\tau L(\tau_0,\eta_0)[\hat\tau-\tau_0] - \nabla_\tau L(\tau_0,\hat\tau)[\hat\tau-\tau_0]\} + \tfrac{\kappa}{\lambda}\|\eta-\eta_0\|_{\mathcal{H}}^{\frac{4}{1+r}}. \tag{26}$$

Proof of Thm. 1. By applying the Taylor's expansion and rearranging, we have

$$\tfrac{1}{2}\nabla_\tau^2 L(\bar\tau,\hat\eta)[(\hat\tau-\tau_0)^2] = L(\hat\tau,\hat\eta) - L(\tau_0,\hat\eta) - \nabla_\tau L(\tau_0,\hat\tau)[\hat\tau-\tau_0],$$

where $\bar\tau$ is on the line segment between $\hat\tau$ and τ_0.
Using Assumption 2, we have

$$\frac{\lambda}{2}\|\tau-\tau_0\|_{\mathcal{T}}^2 \leq \underbrace{L(\hat\tau,\hat\eta) - L(\tau_0,\hat\eta)}_{R_{\mathcal{T}}(\hat\tau;\hat\eta,\epsilon)} - \nabla_\tau L(\tau_0,\hat\tau)[\hat\tau-\tau_0] + \frac{\kappa}{2}\|\eta-\eta_0\|_{\mathcal{H}}^{\frac{4}{1+r}}.$$

Since $\nabla_\tau L(\tau_0,\eta_0)[\hat\tau-\tau_0] \geq 0$ by Assumption 1, we have

$$\frac{\lambda}{2}\|\tau-\tau_0\|_{\mathcal{T}}^2 \leq R_{\mathcal{T}}(\hat\tau;\hat\eta,\epsilon) + \{\nabla_\tau L(\tau_0,\eta_0)[\hat\tau-\tau_0] - \nabla_\tau L(\tau_0,\hat\tau)[\hat\tau-\tau_0]\} + \frac{\kappa}{2}\|\eta-\eta_0\|_{\mathcal{H}}^{\frac{4}{1+r}}. \tag{27}$$

□

The middle difference

$$\{\nabla_\tau L(\tau_0,\eta_0) - \nabla_\tau L(\tau_0,\hat\eta)\}[\hat\tau-\tau_0] \tag{28}$$

is the *nuisance leakage of the first-order optimality condition.* It is the main channel through which nuisance error affects the target. Under Neyman orthogonality, the leakage is *higher than first order* in $\|\hat\eta-\eta_0\|$ (typically quadratic or a product of nuisance errors), so $\hat\tau$ inherits only a higher-order remainder rather than linear bias. In particular, for the DR-learner it factors into a product of nuisance errors (yielding *double robustness*), whereas for the R-learner it enables *fast rates* once the nuisances are sufficiently accurate. We quantify these forms below for each loss.

3.1 Nuisance Leakage: R-learner

Theorem 2 (Error Analysis: R-learner). Suppose Assumption 1 and 2 hold with $\|\cdot\|_{\mathcal{T}} = \|\cdot\|_2$. Let $a \triangleq \|\tau_0\|_\infty^2$ and $\lambda \triangleq 2c(1-c)$, where c is a constant satisfying $c \le \pi_0(X) \le 1-c$. Then, with probability $1-\epsilon$,

$$\|\hat{\tau} - \tau_0\|_2^2 \le \frac{4}{\lambda}\mathcal{R}_{\mathcal{T}}(\hat{\tau}; \hat{\eta}, \epsilon) + \frac{32a}{\lambda^2}\|\hat{\pi} - \pi_0\|_4^4 + \frac{32}{\lambda^2}\|\hat{m} - m_0\|_4^2\|\hat{\pi} - \pi_0\|_4^2. \tag{29}$$

Proof of Thm. 2. Let $h_\tau(X) \triangleq \hat{\tau}(X) - \tau_0(X)$. Let $\delta_m(X) \triangleq (m(X) - m_0(X))$ and $\delta_\pi(X) \triangleq (\pi(X) - \pi_0(X))$.
Note, the first-order risk function is the following:

$$\nabla_\tau L_{\mathrm{R}}[\tau, \eta](h_\tau) = -2\mathbb{E}[\{Y - m(X) - \tau(X)(T - \pi(X))\} \cdot \{T - \pi(X)\} \cdot h_\tau(X)], \tag{30}$$

We note that $\nabla_\tau L_{\mathrm{R}}[\tau_0, \eta_0](h_\tau) = 0$, as shown in the first-order optimality condition analysis. To analyze the leakage, we rewrite a few terms here:

$$Y - m - \tau_0(T - \pi) = \underbrace{Y - m_0 - \tau_0(T - \pi_0)}_{\epsilon_Y} - \delta_m + \tau_0\delta_\pi \tag{31}$$

$$T - \pi = T - \pi_0 - \delta_\pi. \tag{32}$$

Then, we can rewrite the first-order risk as follows:

$$\begin{aligned}
\nabla_\tau L_{\mathrm{R}}[\tau_0, \eta](h_\tau) &= -2\mathbb{E}[\{\epsilon_Y - \delta_m + \tau_0\delta_\pi\} \cdot (T - \pi_0 - \delta_\pi) \cdot h_\tau] && (33)\\
&= 2\mathbb{E}[\{\tau_0\delta_\pi^2 - \delta_m\delta_\pi\}h_\tau] && (34)\\
&\le 2|\mathbb{E}[\tau_0\delta_\pi^2 h_\tau]| + 2|\mathbb{E}[\delta_m\delta_\pi h_\tau]| && (35)\\
&\le 2\|\delta_\pi\|_4^2 \cdot \|\tau_0\|_\infty \cdot \|h_\tau\|_2 + 2\|\delta_m\|_4 \cdot \|\delta_\pi\|_4 \cdot \|h_\tau\|_2 && (36)\\
&= 2\|h_\tau\|_2 \cdot \left(\|\tau_0\|_\infty\|\delta_\pi\|_4^2 + \|\delta_m\|_4 \cdot \|\delta_\pi\|_4\right). && (37)
\end{aligned}$$

Then, for any $\alpha > 0$, Young's inequality (with $p = q = 2$) gives

$$\begin{aligned}
&\nabla_\tau L(\tau_0, \eta_0)[\hat{\tau} - \tau_0] - \nabla_\tau L(\tau_0, \hat{\eta})[\hat{\tau} - \tau_0] && (38)\\
&\le 2\|h_\tau\|_2 \cdot \left(\|\tau_0\|_\infty\|\delta_\pi\|_4^2 + \delta_m\|_4 \cdot \|\delta_\pi\|_4\right) && (39)\\
&\le \alpha\|h_\tau\|_2^2 + \frac{1}{\alpha}\left(\|\tau_0\|_\infty\|\delta_\pi\|_4^2 + \|\delta_m\|_4 \cdot \|\delta_\pi\|_4\right)^2 && (40)\\
&= \alpha\|h_\tau\|_2^2 + \frac{2}{\alpha}\|\tau_0\|_\infty^2\|\delta_\pi\|_4^4 + \frac{2}{\alpha}\|\delta_m\|_4^2\|\delta_\pi\|_4^2. && (41)
\end{aligned}$$

Choose $\alpha = \lambda/4$. Let $\mathcal{R}_{\mathcal{T}} \triangleq \mathcal{R}_{\mathcal{T}}(\hat{\tau}; \hat{\eta}, \epsilon)$. Then, by Thm. 1, we have

$$\begin{aligned}
\|h_\tau\|_2^2 &\le \frac{2}{\lambda}\mathcal{R}_{\mathcal{T}} + \frac{2}{\lambda}\frac{\lambda}{4}\|h_\tau\|_2^2 + \frac{16a}{\lambda^2}\|\delta_\pi\|_2^2 + \frac{16}{\lambda^2}\|\delta_m\|_4^2\|\delta_\pi\|_4^2, && (42)\\
\implies \frac{1}{2}\|h_\tau\|_2^2 &\le \frac{2}{\lambda}\mathcal{R}_{\mathcal{T}} + \frac{16a}{\lambda^2}\|\delta_\pi\|_2^2 + \frac{16}{\lambda^2}\|\delta_m\|_4^2\|\delta_\pi\|_4^2, && (43)
\end{aligned}$$

which completes the proof. □

3.2 Nuisance Leakage: DR-learner

Theorem 3 (**Error Analysis: DR-learner**)**.** Suppose Assumption 1 and 2 hold with $\|\cdot\|_{\mathcal{T}} = \|\cdot\|_2$. Then, with probability $1-\epsilon$,

$$\|\hat{\tau} - \tau_0\|_2^2 \leq 2\mathcal{R}_{\mathcal{T}}(\hat{\tau}; \hat{\eta}, \epsilon) + 8\|\omega - \omega_0\|_4^2 \|\mu_0 - \mu\|_4^2. \tag{44}$$

Proof of Thm. 3. Let $h_\tau(X) \triangleq \hat{\tau}(X) - \tau_0(X)$. Let $\delta_\mu(XZ) \triangleq (\mu(XZ) - \mu_0(XZ))$ and $\delta_\omega(X) \triangleq (\omega(X) - \omega_0(X))$.
Note

$$\nabla_\tau L_{\mathrm{DR}}(\tau, \eta)[h_\tau] = -2\mathbb{E}[\{\varphi(V;\eta) - \tau\}h_\tau] \leq 2|\mathbb{E}[\{\varphi(V;\eta) - \tau\}h_\tau]|. \tag{45}$$

We note $\nabla_\tau L_{\mathrm{DR}}(\tau_0, \eta_0)[h_\tau] = 0$, as shown in the first-order optimality condition analysis. Also,

$$\begin{aligned}
&|\mathbb{E}[\{\varphi(V;\eta) - \tau_0\}h_\tau]]| && (46)\\
&= |\mathbb{E}[\{\omega(TX)\{Y - \mu(TX)\} + \omega_0(TX)\mu(TX) - \omega_0(TX)\mu_0(TX)\}h_\tau(X)]| && (47)\\
&= |\mathbb{E}[\{\omega(TX)\{\mu_0(TX) - \mu(TX)\} + \omega_0(TX)\{\mu(TX) - \mu_0(TX)\}\}h_\tau(X)]| && (48)\\
&= |\mathbb{E}[\{\omega(TX) - \omega_0(TX)\}\{\mu_0(TX) - \mu(TX)\}h_\tau(X)]| && (49)\\
&\leq \|h_\tau\|_2 \|(\omega - \omega_0)(\mu_0 - \mu)\|_2 && (50)\\
&\leq \|h_\tau\|_2 \|\omega - \omega_0\|_4 \|\mu_0 - \mu\|_4. && (51)
\end{aligned}$$

Then, for any $\alpha > 0$, Young's inequality (with $p = q = 2$) gives

$$\begin{aligned}
&2|\mathbb{E}[\{\varphi(V;\eta) - \tau_0\}h_\tau]]| && (52)\\
&\leq \alpha\|h_\tau\|_2^2 + \frac{1}{\alpha}\|\omega - \omega_0\|_4^2 \|\mu_0 - \mu\|_4^2. && (53)
\end{aligned}$$

Choose $\alpha = \lambda/4$. Let $\mathcal{R}_{\mathcal{T}} \triangleq \mathcal{R}_{\mathcal{T}}(\hat{\tau}; \hat{\eta}, \epsilon)$. Then, by Thm. 1, we have

$$\begin{aligned}
&\|h_\tau\|_2^2 \leq \frac{2}{\lambda}\mathcal{R}_{\mathcal{T}} + \frac{2}{\lambda}\frac{\lambda}{4}\|h_\tau\|_2^2 + \frac{16}{\lambda^2}\|\omega - \omega_0\|_4^2 \|\mu_0 - \mu\|_4^2 && (54)\\
&\implies \frac{1}{2}\|h_\tau\|_2^2 \leq \frac{2}{\lambda}\mathcal{R}_{\mathcal{T}} + \frac{16}{\lambda^2}\|\omega - \omega_0\|_4^2 \|\mu_0 - \mu\|_4^2, && (55)
\end{aligned}$$

which completes the proof. □

4 Universal Analysis with Orthogonality

In this section, we present a general result: fast rate convergence can be achieved whenever the risk function satisfies certain smoothness conditions and, at the same time, its first-order condition remains stable under small errors in nuisance estimation. To formalize this idea, we introduce the following assumptions together with their rationales.

Assumption 3 (**Orthogonality at the truth (Neyman orthogonality)**)**.** Small errors in the nuisance η do not change the first-order optimality condition in τ (i.e.,

$\nabla_\tau L(\tau_0, \cdot)$) at the true nuisance η_0.

$$\nabla_\eta \nabla_\tau L(\tau_0, \eta_0)[h_\tau, h_\eta] = 0 \quad \text{for all directions } h_\tau \in \mathcal{T},\ h_\eta \in \mathcal{H}. \tag{56}$$

Assumption 4 (**Curvature Bound**)**.** There exist constants $b_1 > 0$ such that,

$$\nabla_\tau^2 L(\bar{\tau}, \eta)[\tau - \tau_0, \tau - \tau_0] \ \leq\ b_1 \|\tau - \tau_0\|_{\mathcal{T}}^2. \tag{57}$$

Rationale: (with Assumption 2), the curvature of the risk function a quadratic function of $\tau - \tau_0$.

Assumption 5 (**Smoothness on Nuisance**)**.** There exist constants $b_2 > 0$ and $r \in [0, 1]$ defined in Assumption 2 such that,

$$|\nabla_\eta^2 \nabla_\tau L(\tau_0, \bar{\eta})[\tau - \tau_0, \eta - \eta_0, \eta - \eta_0]| \ \leq\ b_2 \|\tau - \tau_0\|_{\mathcal{T}}^{1-r} \|\eta - \eta_0\|_{\mathcal{H}}^2 \tag{58}$$

Rationale: The curvature of the first-order function of the risk is $O(\|\eta - \eta_0\|_{\mathcal{H}}^2)$, allowing mild deterioration when $\tau \neq \tau_0$.

Under these conditions, we obtain the following universal fast rate result.

Theorem 4 (**Fast Rate Convergence - Universal**)**.** Suppose Assumption 1 to 5 hold. Let $\beta_1 = 2/\lambda$ and $\beta_2 = \frac{\lambda}{2}\left(\frac{b_2(1+r)}{4}(\frac{b_2(1-r)}{\lambda})^{\frac{1-r}{1+r}} + \frac{\kappa}{2}\right)$. Then, with probability $1 - \epsilon$,

$$\|\hat{\tau} - \tau_0\|_{\mathcal{T}}^2 \leq \beta_1 R_{\mathcal{T}}(\hat{\tau}; \hat{\eta}, \epsilon) + \beta_2 \|\hat{\eta} - \eta_0\|_{\mathcal{H}}^{\frac{4}{1+r}}. \tag{59}$$

Proof of Thm. 4. By applying the Taylor's expansion and rearranging, we have

$$\tfrac{1}{2}\nabla_\tau^2 L(\bar{\tau}, \hat{\eta})[h_\tau, h_\tau] = L(\hat{\tau}, \hat{\eta}) - L(\tau_0, \hat{\eta}) - \nabla_\tau L(\tau_0, \hat{\tau})[h_\tau],$$

where $\bar{\tau}$ is on the line segment between $\hat{\tau}$ and τ_0.
Using Assumption 2, we have

$$\frac{\lambda}{2}\|h_\tau\|_{\mathcal{T}}^2 \leq \underbrace{L(\hat{\tau}, \hat{\eta}) - L(\tau_0, \hat{\eta})}_{R_{\mathcal{T}}(\hat{\tau}; \hat{\eta}, \epsilon)} - \nabla_\tau L(\tau_0, \hat{\tau})[h_\tau] + \frac{\kappa}{2}\|h_\eta\|_{\mathcal{H}}^{\frac{4}{1+r}}.$$

By applying Taylor's expansion,

$$-\nabla_\tau L(\tau_0, \hat{\tau})[h_\tau] = \underbrace{-\nabla_\tau L(\tau_0, \eta_0)[h_\tau]}_{\leq 0, \text{ by Assumption 1}} - \underbrace{\nabla_\eta \nabla_\tau L(\tau_0, \eta_0)[h_\tau, h_\eta]}_{=0, \text{ by Assumption 3}} - \tfrac{1}{2}\nabla_\eta^2 \nabla_\tau L(\tau_0, \bar{\eta})[h_\tau, h_\eta, h_\eta]. \tag{60}$$

Continuing,

$$\begin{aligned}
-\nabla_\tau L(\tau_0, \hat{\tau})[h_\tau] &= -\nabla_\tau L(\tau_0, \eta_0)[h_\tau] - \tfrac{1}{2}\nabla_\eta^2 \nabla_\tau L(\tau_0, \bar{\eta})[h_\tau, h_\eta, h_\eta] && (61)\\
&\leq -\tfrac{1}{2}\nabla_\eta^2 \nabla_\tau L(\tau_0, \bar{\eta})[h_\tau, h_\eta, h_\eta] && (62)\\
&\leq \tfrac{b_2}{2}\|\tau - \tau_0\|_{\mathcal{T}}^{1-r}\|\eta - \eta_0\|_{\mathcal{H}}^2, \text{ by Assumption 5}. && (63)
\end{aligned}$$

Invoking Young's inequality, for any constant $\alpha > 0$,

$$\tfrac{b_2}{2}\|\tau - \tau_0\|_{\mathcal{T}}^{1-r}\|\eta - \eta_0\|_{\mathcal{H}}^2 \leq \tfrac{b_2\alpha}{2p}\|\tau - \tau_0\|_{\mathcal{T}}^{p(1-r)} + \tfrac{b_2}{2q\alpha^{q/p}}\|\eta - \eta_0\|_{\mathcal{H}}^{2q}. \tag{64}$$

Choose $p = \frac{2}{1-r}$ and $q = \frac{2}{1+r}$. Then, it becomes

$$\tfrac{b_2}{2}\|\tau - \tau_0\|_{\mathcal{T}}^{1-r}\|\eta - \eta_0\|_{\mathcal{H}}^2 \leq \tfrac{b_2\alpha(1-r)}{4}\|\tau - \tau_0\|_{\mathcal{T}}^2 + \tfrac{b_2(1+r)\alpha^{\frac{r-1}{1+r}}}{4}\|\eta - \eta_0\|_{\mathcal{H}}^{\frac{4}{1+r}}. \tag{65}$$

Choose $\alpha = \frac{\lambda}{b_2(1-r)}$. Then, the upper bound will be

$$\tfrac{b_2}{2}\|\tau - \tau_0\|_{\mathcal{T}}^{1-r}\|\eta - \eta_0\|_{\mathcal{H}}^2 \leq \tfrac{\lambda}{4}\|h_\tau\|_{\mathcal{T}}^2 + \tfrac{b_2(1+r)}{4}(\tfrac{b_2(1-r)}{\lambda})^{\frac{1-r}{1+r}}\|h_\eta\|_{\mathcal{H}}^{\frac{4}{1+r}}. \tag{66}$$

Combining, we have

$$-\nabla_\tau L(\tau_0, \hat{\tau})[h_\tau] \leq \tfrac{\lambda}{4}\|h_\tau\|_{\mathcal{T}}^2 + \tfrac{b_2(1+r)}{4}(\tfrac{b_2(1-r)}{\lambda})^{\frac{1-r}{1+r}}\|h_\eta\|_{\mathcal{H}}^{\frac{4}{1+r}}. \tag{67}$$

Then,

$$\frac{\lambda}{2}\|h_\tau\|_{\mathcal{T}}^2 \leq R_{\mathcal{T}}(\hat{\tau}; \hat{\eta}, \epsilon) + \tfrac{\lambda}{4}\|h_\tau\|_{\mathcal{T}}^2 + \tfrac{b_2(1+r)}{4}(\tfrac{b_2(1-r)}{\lambda})^{\frac{1-r}{1+r}}\|h_\eta\|_{\mathcal{H}}^{\frac{4}{1+r}} + \frac{\kappa}{2}\|h_\eta\|_{\mathcal{H}}^{\frac{4}{1+r}}, \tag{68}$$

which implies that

$$\frac{\lambda}{2}\|h_\tau\|_{\mathcal{T}}^2 \leq R_{\mathcal{T}}(\hat{\tau}; \hat{\eta}, \epsilon) + \left(\tfrac{b_2(1+r)}{4}(\tfrac{b_2(1-r)}{\lambda})^{\frac{1-r}{1+r}} + \tfrac{\kappa}{2}\right)\|h_\eta\|_{\mathcal{H}}^{\frac{4}{1+r}}. \tag{69}$$

□

References