

# Adaptive Estimation and Inference in Conditional Moment Models via the Discrepancy Principle

Jiyuan Tan\*      Vasilis Syrgkanis†

## Abstract

We study adaptive estimation and inference in linear inverse problems defined by conditional moment restrictions. Regularization stabilizes such problems, but the regularization parameter depends on a source parameter that is almost never known. We develop a discrepancy principle that chooses this parameter from empirical losses and prescribed bounds on statistical error, without knowing the source parameter. The framework covers Regularized DeepIV (RDIV) and the Tikhonov Regularized Adversarial Estimator (TRAE). Under the stated realizability conditions and comparable fitting and calibration error bounds, both one-step procedures attain strong- and weak-error bounds that match known-source rates for source parameters up to one. We also extend the framework to iterative procedures under realizability of the population iterates, widening the range of source parameters covered by the rate comparison. Finally, we use these adaptive estimators to construct an asymptotically normal doubly robust estimator that adapts to the well-posedness of the primal and dual problems.

## 1 Introduction

Ill-posed inverse problems arise throughout causal inference and econometrics [Horowitz, 2014b], for example in nonparametric instrumental variable (IV) regression [Newey and Powell, 2003, Darolles et al., 2011, Newey, 2013, Ai and Chen, 2003], proximal causal inference [Miao et al., 2024b, Tchetgen Tchetgen et al., 2024], and missing-not-at-random

---

\*Jiyuan Tan was supported by NSF Award IIS-2337916.

†Vasilis Syrgkanis was supported by NSF Award IIS-2337916.

problems [d’Haultfoeuille, 2010, Miao et al., 2024a, Breunig and Haan, 2021, Li et al., 2023]. We study estimating parameters defined by

$$\theta_0 = \mathbb{E} [\tilde{m}(W; h_0)],$$

where  $h \mapsto \tilde{m}(W; h)$  is a known linear map. The nuisance function  $h_0$  solves the conditional moment problem

$$\mathbb{E}[h_0(X) \mid Z = z] = r_0(z), \quad (1.1)$$

where  $X$  and  $Z$  are  $W$ -measurable random variables. The function  $r_0$  represents the expectation of a known linear map:  $\mathbb{E}[m(W; q)] = \mathbb{E}[r_0(Z)q(Z)]$  for the test functions  $q$  that Section 2 introduces.

*Remark (notation).*  $r_0 \in \tilde{\mathcal{Q}}$  is the Riesz representer of the continuous linear functional  $q \mapsto \mathbb{E}[m(W; q)]$ , satisfying  $\mathbb{E}[m(W; q)] = \mathbb{E}[r_0(Z)q(Z)]$  for every  $q \in \tilde{\mathcal{Q}}$ . In Proposition 3.5, Theorem 3.6, Proposition 3.8, Theorem 3.9, Lemma B.1, Lemma B.2, Lemma B.3, Lemma B.4, Corollary B.5, Lemma B.6, Lemma B.8, Corollary B.9, Lemma B.10, a result cited in the proof of Lemma B.2:  $\tilde{\mathcal{Q}} = L_2(Z)$  and  $r_0(z) = \mathbb{E}[Y \mid Z = z]$ . The letter  $q$  is locally reused for function-valued arguments in  $\tilde{\mathcal{Q}}$  and for a deterministic scalar bias exponent in  $(1, \infty)$ ; these uses denote distinct objects. In Theorem 3.1: In  $B_n(\lambda) = C_b \lambda^q$ ,  $q > 1$  is a fixed scalar exponent. In Theorem 3.6, Theorem 3.14: In the application of the generic discrepancy guarantee, the scalar exponent is  $q = p = \min\{\beta + 1, 2\}$ . In Theorem 3.9, Theorem 3.18: In the proof’s scalar abbreviation and application of the generic discrepancy guarantee,  $q = q_t = \min\{\beta + 1, 2t\}$ . In Corollary 3.19: In the dual minimization,  $q \in \mathcal{Q}$  is the candidate function of  $Z$ ; in  $\hat{\theta}(h, q)$  it is the dual nuisance argument.

Two established approaches estimate nuisance functions from conditional moment restrictions. Linear inverse problem (LIP) methods formulate the restriction as an operator equation and regularize its solution [Carrasco et al., 2007, Darolles et al., 2011, Bennett et al., 2023]. Sieve minimum distance methods approximate the unknown functions in finite-dimensional spaces and minimize a sample criterion [Ai and Chen, 2003, Chen and Pouzo, 2012]. We follow the LIP approach. Let  $\Pi_{\tilde{\mathcal{Q}}}$  denote the mean-square projection onto a closed test space. The operator  $\mathcal{T}h = \Pi_{\tilde{\mathcal{Q}}}\mathbb{E}[h(X) \mid Z = \cdot]$  then turns (1.1) into the equation  $\mathcal{T}h_0 = r_0$ . We estimate the minimum-norm solution of this equation with regularized estimators over general hypothesis classes.

While many works use Tikhonov regularization for this inverse problem [Newey, 2013, Li et al., 2024, Bennett et al., 2026, 2023], a central challenge in applying LIP methods is selecting the regularization parameter. Known-smoothness rate analyses for the Tikhonov regularized adversarial estimator (TRAЕ) [Bennett et al., 2026, 2023] and Regularized DeepIV

(RDIV) [Li et al., 2024] express the regularization level in terms of the source parameter—often encoded by the  $\beta$ -source condition (i.e.,  $h_0 = (\mathcal{T}^*\mathcal{T})^{\beta/2} w_0$  for some  $w_0 \in \tilde{\mathcal{H}}$ ). Li et al. [2024] also study model selection for RDIV. Our procedure selects the regularization level within a fixed hypothesis class. In practice,  $\beta$  is rarely known. A misspecified regularization parameter can lead to suboptimal convergence rates or even divergence. The L-curve offers a geometric criterion for choosing the regularization parameter from the solution and residual norms [Hansen, 1992]. Cross-validation (CV) is also widely used, but it requires repeatedly minimizing the loss and is computationally expensive. Moreover, CV typically selects the parameter that minimizes the weak metric; however, this method fails to directly imply a bound in the strong metric. This paper addresses this gap by introducing a discrepancy-principle framework for adaptively selecting a hyperparameter that does not require knowing the exact  $\beta$ .

*Remark (notation).*  $\beta > 0$  is a source exponent satisfying  $h_0 = (\mathcal{T}^*\mathcal{T})^{\beta/2} w_0$  for some  $w_0 \in \tilde{\mathcal{H}}$ . In Corollary 3.19: The two applications use potentially different source exponents:  $h_0 = (\mathcal{T}^*\mathcal{T})^{\beta_h/2} w_h$  and  $q_0 = (\mathcal{T}\mathcal{T}^*)^{\beta_q/2} w_q$ , with source elements in their respective Hilbert spaces;  $\beta_m = \max\{\beta_h, \beta_q\}$ .  $w_0 \in \tilde{\mathcal{H}}$  is a chosen source witness satisfying  $h_0 = (\mathcal{T}^*\mathcal{T})^{\beta/2} w_0$ , with no uniqueness or unit-norm requirement. In Proposition 3.5, Proposition 3.8: When  $h_0 = 0$ , take  $w_0 = 0$  in the initial-acceptance argument.

We primarily use the discrepancy principle (DP) [Morozov, 1966] from classical inverse problems. DP adaptively selects the regularization parameter  $\lambda$  so that the empirical loss is on the same order as the estimated noise level. This data-driven process balances empirical fit and statistical noise. In particular, we extend the classical DP to address the following challenges in our setting. First, classical LIP assumes known operators; in our setting the conditional expectation operator  $\mathcal{T}$  is unknown and must be estimated, either explicitly or implicitly, before solving the inverse problem. Second, classical DP assumes fixed, known noise in  $r_0$ , whereas our noise comes from empirical-process fluctuations and depends on the estimator; see, e.g., the localized concentration bounds in Foster and Syrgkanis [2023, Lemma 14]. Third, classical DP is typically analyzed over full Hilbert spaces with closed-form regularized solutions, while we work with general hypothesis classes (e.g., neural networks) that satisfy inductive-bias assumptions.

*Remark (notation).*  $\lambda$  is the scalar regularization parameter. In the classical Hilbert-space construction (2.3) and Definition 2.2, finite penalties satisfy  $0 < \lambda < \infty$ , with  $\lambda = \infty$  denoting the zero-estimator convention. The value  $\lambda = 0$  is admitted only for the separate finite-dimensional computational convention of Section D.1, which selects a minimum-Euclidean-norm solution on the retained spectral subspace; it does not define an unregularized solution

of (2.3) for arbitrary  $r_\delta$ . In Theorem 3.1, Proposition 3.5, Theorem 3.6, Proposition 3.8, Theorem 3.9, Proposition 3.13, Theorem 3.14, Proposition 3.17, Theorem 3.18, Corollary 3.19: The tuning rule searches the deterministic finite grid  $\Lambda_n = \{2\rho^k : k = 0, \dots, K_n\}$ , where  $1/2 \leq \rho < 1$  and  $K_n = \lceil \log(2n^2)/\log(1/\rho) \rceil$ ; its selected value is denoted  $\lambda_{\text{dp}}$ . In Proposition D.1:  $\lambda > 0$  is a candidate penalty for the unrestricted quadratic program, before the numerical spectral restriction. In Lemma C.2: In each concentration application in the proof,  $\lambda \in \Lambda_n \cap [\lambda_{\text{IT}}^\circ, 2]$  is a fixed candidate penalty.

In this paper, we propose a discrepancy-principle-based tuning rule without prior information about smoothness. Our main contributions are as follows.

- We give a uniform discrepancy rule together with an oracle guarantee that apply to a large class of estimators, and we apply them to RDIV and TRAE. For source parameters up to one, and when fitting and calibration errors are comparable, we show that the discrepancy principle selects a regularization parameter with the same strong- and weak-error rates as in the known-source analysis.
- We additionally extend the adaptive estimator to iterated regularized RDIV and TRAE, and we show that, for a prescribed iteration count  $t$ , the adaptive iterated estimator matches the known-source rates over the range  $\beta \leq 2t - 1$ .
- We combine adaptive TRAE fits of the primal and dual problems into a doubly robust estimator of  $\theta_0$ . We show that this estimator is asymptotically normal and can adapt to the well-posedness of either problem.
- We evaluate the rule with kernel classes in two simulation designs, where the best fixed penalty changes by three orders of magnitude across settings. With a data-driven noise level, the selected penalty follows the best fixed penalty of each setting, and adaptive TRAE has smaller strong, weak, and target errors than any single fixed penalty.

*Remark (notation).*  $t \geq 1$  is a prescribed integer iteration count held fixed as  $n$  varies and chosen independently of the data and the source parameter  $\beta$ .

The rest of the paper proceeds as follows. Section 2 sets up the problem, fixes notation, and reviews linear inverse problems and the discrepancy principle. Section 3 develops the adaptive regularization framework and proves the guarantees for RDIV, TRAE, and doubly robust estimation. Section 4 evaluates the rule in simulations with kernel classes, and the conclusion discusses implications and open directions.

## 1.1 Related Work

**Nonparametric instrumental variables.** Nonparametric IV estimation provides a central example of a statistical inverse problem. Sieve and kernel estimators address the ill-posedness of the conditional expectation operator through finite-dimensional approximation or regularization [Newey and Powell, 2003, Hall and Horowitz, 2005, Darolles et al., 2011, Chen and Pouzo, 2012]. The resulting convergence rates depend on both the regularity of the structural function and the operator [Chen and Reiss, 2011, Florens et al., 2011]. Recent estimators accommodate flexible function classes through kernel methods [Singh et al., 2019, Muandet et al., 2020], minimax formulations of the moment restrictions [Lewis and Syrgkanis, 2018, Dikkala et al., 2020, Liao et al., 2020, Bennett and Kallus, 2023], or two-stage learning of conditional expectations [Hartford et al., 2017, Xu et al., 2021, Li et al., 2024]. Our analysis builds on the regularized least-squares estimator of Li et al. [2024] and the regularized adversarial estimator of Bennett et al. [2023].

**Adaptive estimation in IV models.** Data-driven regularization has an established literature in nonparametric IV models. Johannes and Schwarz [2011] develop a model-selection procedure that adapts over smoothness classes under restrictions on the conditional expectation operator. Horowitz [2014a] select the dimension of a series estimator through an empirical criterion and establish adaptation to unknown smoothness. Breunig and Johannes [2016] combine model selection and Lepski’s method to estimate linear functionals at adaptive rates. More recently, Chen et al. [2025] develop data-driven sieve dimension choices that attain optimal sup-norm rates for structural functions and their derivatives, together with adaptive uniform confidence bands. These procedures adapt the dimension of a sieve or series expansion. We instead fix a general hypothesis class and adapt the Tikhonov penalty itself, combining localized loss bounds for RDIV and TRAE [Li et al., 2024, Bennett et al., 2023] with a discrepancy search over a geometric grid.

**Discrepancy principles and iterative regularization.** The discrepancy principle chooses the amount of regularization by comparing an observed residual with a noise threshold [Morozov, 1966, 1984]. Classical inverse-problem theory relates this rule to source conditions and explains how iterated Tikhonov regularization extends the range of smoothness that the estimator can use [Engl et al., 1996, Nair et al., 2003]. Statistical analyses account for the fluctuations of the residual itself. Blanchard et al. [2018] establish oracle adaptation in the strong norm for residual stopping rules in linear inverse problems. Celisse and Wahl [2021] study discrepancy rules for kernel spectral estimators, including Tikhonov regularization, and

derive oracle inequalities and adaptive rates for ordinary and smoothed residuals. [Harrach et al. \[2020\]](#) analyze regularization with noise estimated from repeated measurements. Our conditional moment setting adds estimation of the operator and optimization over hypothesis classes. The projected RDIV discrepancy and the adversarial discrepancy account for these statistical errors within a common finite-grid selection argument. The iterative analyses apply this argument to a prescribed number of Tikhonov steps.

**Inference on linear functionals.** Inference on functionals of conditional moment solutions connects this regularization problem to semiparametric estimation [[Ai and Chen, 2003](#)]. [Bennett et al. \[2026\]](#) characterize strongly identified linear functionals and construct debiased estimators from regularized primal and dual fits. [Bennett et al. \[2023\]](#) develop source-condition double robustness, under which sufficiently favorable regularity of either inverse problem can support asymptotic normality, and analyze iterated adversarial Tikhonov estimators. We use their primal–dual representation and remainder analysis with discrepancy-selected nuisance estimators. This yields a sufficient condition for asymptotic normality expressed in the primal and dual source parameters while choosing the regularization levels from the data.

## 2 Preliminaries

**Problem formulation.** We observe independent copies of a random variable  $W$  and estimate a linear functional whose nuisance function is defined by a LIP. The target is

$$\theta_0 = \mathbb{E}[\tilde{m}(W; h_0)],$$

with a known linear functional  $h \mapsto \tilde{m}(W; h)$ . Let  $X$  and  $Z$  be two  $W$ -measurable random variables, and let  $\tilde{\mathcal{H}}$  and  $\tilde{\mathcal{Q}}$  be closed linear subspaces of  $L_2(X)$  and  $L_2(Z)$ , respectively, where  $L_2(X)$  denotes square-integrable functions of  $X$ . Let  $\mathcal{H} \subset \tilde{\mathcal{H}}$  and  $\mathcal{Q} \subset \tilde{\mathcal{Q}}$  be subsets of these function spaces. The classes  $\mathcal{H}$  and  $\mathcal{Q}$  may consist of bounded subsets of reproducing kernel Hilbert spaces, neural networks, or parametric families. We define the linear operator  $\mathcal{T}h = \Pi_{\tilde{\mathcal{Q}}} \mathbb{E}[h(X) \mid Z = \cdot]$ , where  $\Pi_{\mathcal{V}}$  denotes mean-square projection onto a closed subspace  $\mathcal{V}$ . The nuisance function  $h_0$  is defined as the minimum-norm solution in  $\tilde{\mathcal{H}}$  to the LIP

$$\mathcal{T}h_0 = r_0, \tag{2.1}$$

where  $r_0$  is the Riesz representer of the continuous linear functional  $q \mapsto \mathbb{E}[m(W; q)]$ , so that

$$\mathbb{E}[m(W; q)] = \mathbb{E}[r_0(Z)q(Z)], \quad \forall q \in \tilde{\mathcal{Q}}.$$

We make the following assumption to ensure that solutions exist.

*Remark (notation).*  $\mathcal{H}$  denotes the primal hypothesis class of functions of  $X$ . In Proposition [D.1](#):  $\mathcal{H}$  is the full RKHS generated by  $k_{\mathcal{H}}$ , before restriction to the retained numerical spectral span; its RKHS norm specifies the minimum-norm representative among functions with the same fitting-sample evaluations.  $L$  denotes the method-specific population loss functional, defined using population expectations. In Lemma [B.1](#), Lemma [B.2](#), Lemma [B.3](#), a result cited in the proof of Lemma [B.2](#):  $L(h) = \mathbb{E}[(Y - \mathcal{T}h(Z))^2] = \sigma_{Y|Z}^2 + \|\mathcal{T}(h - h_0)\|^2$ , where  $\sigma_{Y|Z}^2 = \mathbb{E}[(Y - r_0(Z))^2]$ . In Proposition [3.13](#), Theorem [3.14](#), Proposition [3.17](#), Theorem [3.18](#), Corollary [3.19](#), Lemma [C.1](#), Lemma [C.2](#), Corollary [C.3](#), Lemma [C.4](#), Lemma [C.2](#), Corollary [C.6](#), Lemma [C.7](#), a result cited in the proof of Lemma [C.1](#):  $L(h) = \sup_{f \in \mathcal{F}} \mathbb{E}[2m(W; f) - 2h(X)f(Z) - f(Z)^2]$ ; under the stated closeness assumption,  $L(h) = \|\mathcal{T}(h - h_0)\|^2$  for  $h \in \mathcal{H}$ . In a result cited in the proof of Lemma [C.2](#):  $L(h) = \sup_{f \in \mathcal{F}} \mathbb{E}[2m(W; f) - 2h(X)f(Z) - f(Z)^2]$ ; under the stated closeness assumption,  $L(h) = \|\mathcal{T}(h - h_0)\|^2$  for  $h \in \mathcal{H}$ .

**Assumption 2.1.** We have  $r_0 \in \mathcal{R}(\mathcal{T}) := \{\mathcal{T}h : h \in \tilde{\mathcal{H}}\}$ .

Given an estimated solution  $\hat{h}$  of [\(2.1\)](#), the strong and weak metrics are defined as

$$\|\hat{h} - h_0\|_{\text{strong}} = \|\hat{h} - h_0\|_{L_2}, \quad \|\hat{h} - h_0\|_{\text{weak}} = \|\mathcal{T}(\hat{h} - h_0)\|_{L_2}.$$

The weak metric measures error only after applying  $\mathcal{T}$ , so ill-posedness lets a small weak error coexist with a large strong one. The adjoint of  $\mathcal{T}$  is

$$\mathcal{T}^*q = \Pi_{\tilde{\mathcal{H}}} \mathbb{E}[q(Z) \mid X = \cdot].$$

Let  $a_0$  be the Riesz representer of  $h \mapsto \mathbb{E}[\tilde{m}(W; h)]$ , and let  $q_0$  denote the minimum-norm solution, when it exists, of the dual inverse problem

$$\mathcal{T}^*q_0 = a_0.$$

*Remark (notation).* The notation  $a_0$  is local to either the dual inverse problem or the scalar recursion for  $(a_j)_{j \geq 0}$ . In Corollary [3.19](#):  $a_0 \in \tilde{\mathcal{H}}$  is the Riesz representer satisfying  $\mathbb{E}[\tilde{m}(W; h)] = \mathbb{E}[a_0(X)h(X)]$  for every  $h \in \tilde{\mathcal{H}}$ , and is the right-hand side of  $\mathcal{T}^*q_0 = a_0$ . In Lemma [B.8](#): In the proof's recursion  $a_j = C(1 + a_{j-1})$ ,  $a_0 = 0$  is a scalar initial value, distinct from the Riesz representer in the dual inverse problem.

When the dual problem has a solution, estimates of  $h_0$  and  $q_0$  yield the DR estimator [\[Bennett et al., 2023\]](#) of  $\theta_0$ ,

$$\hat{\theta}(\hat{h}, \hat{q}) = \mathbb{E}_n \left[ \tilde{m}(W; \hat{h}) + m(W; \hat{q}) - \hat{q}(Z)\hat{h}(X) \right].$$

For Corollary 3.19, this empirical average is taken over an evaluation sample independent of the data used to fit and tune the nuisance estimators.

**Linear inverse problems.** Classical inverse-problem theory studies a noisy observation  $r_\delta$  of  $r_0$ , often with a known operator  $\mathcal{T}$  and error bound  $\|r_\delta - r_0\| \leq \delta$ . Inverting an ill-posed operator can amplify this perturbation. To illustrate this point, suppose that the operator  $\mathcal{T}$  admits a countable singular value decomposition

$$\mathcal{T}h = \sum_{i=1}^{\infty} \sigma_i \langle h, v_i \rangle u_i, \quad (2.2)$$

where  $\sigma_1 \geq \sigma_2 \geq \dots > 0$  are the positive singular values, and  $\{v_i\}_{i=1}^{\infty}$  and  $\{u_i\}_{i=1}^{\infty}$  form orthonormal bases of  $(\ker \mathcal{T})^\perp$  and  $(\ker \mathcal{T}^*)^\perp$ , respectively. The minimum-norm solution of (2.1) is

$$h_0 = \sum_{i=1}^{\infty} \frac{\langle r_0, u_i \rangle}{\sigma_i} v_i.$$

The inversion amplifies a perturbation along  $u_k$  by  $1/\sigma_k$ . Tikhonov regularization controls this amplification by solving

$$h_\lambda^\delta = \arg \min_h \{ \|\mathcal{T}h - r_\delta\|^2 + \lambda \|h\|^2 \}. \quad (2.3)$$

Larger values of  $\lambda$  shrink the solution toward zero; smaller values allow greater noise amplification. The discrepancy principle balances these effects through the residual.

**Definition 2.2** (Discrepancy Principle). Given constants  $\ell \in (1, \infty)$  and  $\kappa \in (1, \infty)$ , choose the regularization parameter  $\lambda$  as follows:

1. If  $\|r_\delta\| \leq \kappa\delta$ , choose  $h = 0$ , which corresponds to  $\lambda = \infty$ .
2. If  $\|r_\delta\| > \kappa\delta$ , choose  $\lambda$  such that

$$\|\mathcal{T}h_\lambda^\delta - r_\delta\| \leq \kappa\delta \leq \|\mathcal{T}h_{\lambda'}^\delta - r_\delta\|,$$

for some  $\lambda' \in [\lambda, \ell\lambda]$ .

*Remark (notation).*  $r_\delta \in \tilde{\mathcal{Q}}$  is a perturbed right-hand side satisfying  $\|r_\delta - r_0\| \leq \delta$ .  $\delta \geq 0$  denotes the prescribed bound in  $\|r_\delta - r_0\| \leq \delta$  in the classical inverse problem, and the radius of the population-norm restriction  $\|f\| \leq \delta$  when used as the argument of  $\mathfrak{R}_n(\delta; \mathcal{F})$ .  $h_\lambda^\delta$  is the minimizer over  $h \in \tilde{\mathcal{H}}$  of  $\|\mathcal{T}h - r_\delta\|^2 + \lambda \|h\|^2$  for  $0 < \lambda < \infty$ , with  $h_\infty^\delta = 0$  under the discrepancy rule.

*Remark (notation).*  $\ell$  takes its meaning from its local binding: a prescribed residual-bracketing factor, an integer iteration index, or a real-valued loss function on  $\mathbb{R}$ . In Lemma C.2: In  $M_{\leq j} = \max_{1 \leq \ell \leq j} \{\lambda, \|f_\ell^*\|_\infty^2\}$  and the subsequent references to  $f_\ell^*$ ,  $\ell$  ranges over the integers  $1, \dots, j$ . In a result cited in the proof of Lemma C.2: Put  $M = \max\{1, B\}$ . Here  $\ell : \mathbb{R} \rightarrow \mathbb{R}$  is the globally Lipschitz loss  $\ell(v) = M^2(\max\{-1, \min\{v, 1\}\})^2$ , with Lipschitz constant  $2M^2$ ; boundedness gives  $\ell(h(X)/M) = h(X)^2$  for  $h \in \mathcal{H}$ . In Lemma B.1: In the likelihood concentration argument,  $\ell : \mathbb{R} \rightarrow \mathbb{R}$  is the fixed loss  $\ell(v) = -\log(\max\{1 + v, 2^{-1/2}\})$ ; it is  $\sqrt{2}$ -Lipschitz, satisfies  $\ell(0) = 0$ , and obeys  $\ell(f_g - 1) = -\log f_g$ . In Lemma C.2: Put  $M = \max\{1, B\}$ . In the normalized squared-loss argument,  $\ell : \mathbb{R} \rightarrow \mathbb{R}$  is the loss  $\ell(u) = M^2(\max\{-1, \min\{u, 1\}\})^2$ , which is  $2M^2$ -Lipschitz and satisfies  $\ell(h(X)/M) = h(X)^2$  for  $h \in \mathcal{H}$ .

The first rule uses the default estimator  $h = 0$  when the observed right-hand side is no larger than the prescribed noise threshold. The second rule makes the residual  $\|\mathcal{T}h_\lambda^\delta - r_\delta\|$  comparable to that threshold. This residual is measured after applying  $\mathcal{T}$ ; controlling the reconstruction error in the strong metric additionally requires regularity, such as a source condition. The safety factor  $\kappa > 1$  makes the second rule feasible: the true solution satisfies  $\|\mathcal{T}h_0 - r_\delta\| = \|r_0 - r_\delta\| \leq \delta < \kappa\delta$ , so a residual at the threshold is attainable.

In our statistical setting, both  $\mathcal{T}$  and  $r_0$  are unknown, and this ignorance adds perturbations that the classical rule does not face. With an explicit operator estimate  $\hat{\mathcal{T}}_n$ , the identity  $\hat{\mathcal{T}}_n h_0 = r_0 + e_n$ , where  $e_n = (\hat{\mathcal{T}}_n - \mathcal{T})h_0$ , exhibits one such perturbation. Calibration must also absorb empirical fluctuations and, for the raw RDIV squared loss, an irreducible outcome-noise floor. An adversarial estimator never forms  $\hat{\mathcal{T}}_n$  at all, reaching the operator instead through a minimax objective, so its discrepancy carries a different set of errors. One search rule covers both, with a statistical error threshold that may depend on  $\lambda$ .

To analyze Tikhonov regularization, we introduce the source condition, a standard assumption in the inverse-problem literature.

**Assumption 2.3** (Source Condition). Problem (2.1) satisfies the  $\beta$ -source condition for some  $\beta > 0$ , i.e., there exists  $w_0 \in \tilde{\mathcal{H}}$  such that the minimum-norm solution of (2.1) satisfies  $h_0 = (\mathcal{T}^* \mathcal{T})^{\beta/2} w_0$ .

If  $\mathcal{T}$  is compact, then by (2.2), Assumption 2.3 implies that

$$\sum_{i=1}^{\infty} \frac{\langle h_0, v_i \rangle^2}{\sigma_i^{2\beta}} < \infty,$$

which quantifies the decay of the solution coefficients relative to the singular values. This assumption is standard [Morozov, 1966, Darolles et al., 2011, Li et al., 2024, Bennett et al., 2026, 2023].

**Notation.** Throughout the paper, we use  $\|\cdot\|$  for the  $L_2$  norm with respect to the input distribution, i.e.,

$$\|h\| := \|h\|_{L_2} = \sqrt{\mathbb{E}[h(X)^2]}, \quad \forall h \in L_2(X)$$

and the  $L_1$  norm  $\|h\|_1 = \mathbb{E}[|h(X)|]$ . The same notation applies to functions of  $Z$  under their input distribution. Write  $\|h\|_{2,n} = \{\mathbb{E}_n[h(X)^2]\}^{1/2}$  for the empirical norm. For real numbers,  $a \vee b = \max\{a, b\}$ , and  $\log$  denotes the natural logarithm. We use  $\mathbb{E}_n[\cdot]$  to denote an empirical average, with subscripts indicating the sample when several independent samples are used. We use the asymptotic-order notation  $o(\cdot)$ ,  $O(\cdot)$ ,  $\Omega(\cdot)$ , and  $\Theta(\cdot)$ ;  $o_p(1)$  denotes convergence to zero in probability. For a function class  $\mathcal{F}$ , let  $\text{star}(\mathcal{F}) = \{\gamma f : \gamma \in [0, 1], f \in \mathcal{F}\}$  and  $\mathcal{F} - \mathcal{F} = \{f - f' : f, f' \in \mathcal{F}\}$ . For independent observations  $X_1, \dots, X_n$  from the input distribution, define the expected Rademacher complexity and its localized version by

$$\begin{aligned} \mathfrak{R}_n(\mathcal{F}) &= \mathbb{E}_{X_{1:n}, \epsilon} \left[ \sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i f(X_i) \right| \right], \\ \mathfrak{R}_n(\delta; \mathcal{F}) &= \mathbb{E}_{X_{1:n}, \epsilon} \left[ \sup_{f \in \mathcal{F}, \|f\| \leq \delta} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i f(X_i) \right| \right], \end{aligned}$$

where the  $\epsilon_i$  are independent Rademacher variables, independent of the observations. The localized class uses the population norm. A deterministic critical radius  $\delta_{n,\mathcal{F}} > 0$  satisfies

$$\delta_{n,\mathcal{F}}^2 \geq \mathfrak{R}_n(\delta_{n,\mathcal{F}}; \text{star}(\mathcal{F} - \mathcal{F})).$$

For classes fitted on an independent preliminary sample, we apply this definition conditionally on that sample and use a common deterministic upper bound.

*Remark (notation).*  $\mathcal{F}$  is a function class on the input domain used in the surrounding complexity or loss definition. In Corollary 3.19: For the primal application,  $\mathcal{F} \subset \tilde{\mathcal{Q}}$  is the adversarial test class; for the dual application the corresponding test class is  $\mathcal{S} \subseteq \tilde{\mathcal{H}}$ , with the primal and dual spaces exchanged. In Proposition D.1:  $\mathcal{F}$  is the full RKHS generated by  $k_{\mathcal{F}}$ , and its empirical supremum includes the penalty  $-\eta \|f\|_{\mathcal{F}}^2$  with  $\eta > 0$ . In Proposition 3.13, Theorem 3.14, Proposition 3.17, Theorem 3.18, Lemma C.1, Lemma C.2, Corollary C.3, Lemma C.4, Lemma C.2, Corollary C.6, Lemma C.7, a result cited in the proof of Lemma C.1, a result cited in the proof of Lemma C.2:  $\mathcal{F} \subset \tilde{\mathcal{Q}}$  is the test class in  $L_n(h) = \sup_{f \in \mathcal{F}} \mathbb{E}_n[2m(W; f) - 2h(X)f(Z) - f(Z)^2]$  and its population counterpart; it contains  $\mathcal{T}(h_0 - h)$  for every  $h \in \mathcal{H}$  under Assumption 3.10.  $\gamma$  is locally bound to its construction, with  $\gamma \in [0, 1]$  denoting the scalar multiplier in  $\text{star}(\mathcal{F}) = \{\gamma f : \gamma \in [0, 1], f \in \mathcal{F}\}$ . In Theorem 3.1:  $\gamma \in (0, 1]$  is a fixed exponent in the strong-error comparison and the resulting oracle bound. In Theorem 3.6, Theorem 3.14: When applying the oracle guarantee,

the exponent is  $\gamma = \min\{\beta, 1\}/(1 + \min\{\beta, 1\})$ . In Theorem 3.9, Theorem 3.18: When applying the oracle guarantee, the exponent is  $\gamma = \beta/(\beta + 1)$ . In Lemma B.7: In the spectral bound,  $\gamma > 0$  is an arbitrary real exponent in  $\sup_{0 \leq u \leq 1} u^\gamma \{\lambda/(u + \lambda)\}^{2j} \leq C_{\gamma,j} \lambda^{\min\{\gamma, 2j\}}$ ; it is then specialized to  $\beta$  and  $\beta + 1$ . In Proposition D.1:  $\gamma \in \mathbb{R}^N$  is the critic coefficient vector in  $f_\gamma(\cdot) = \sum_{j=1}^N \gamma_j k_{\mathcal{F}}(Z_j, \cdot)$ , over which the finite-dimensional critic objective is maximized.  $b = \min\{\beta, 1\} \in (0, 1]$  is the truncated source exponent.

### 3 Adaptive Regularization via the Discrepancy Principle

We now state a finite-grid oracle guarantee for the discrepancy rule applied to estimators of (2.1), and then check its conditions for two estimators: RDIV, which estimates the operator explicitly, and TRAE, which estimates the weak loss through an adversarial problem.

#### 3.1 Discrepancy Principle for Conditional Moment Problems

Consider a family of candidate estimators  $\{\hat{h}_\lambda : \lambda \in \Lambda_n\}$  and a nonnegative calibration statistic  $D_n(\hat{h}_\lambda)$ . Fix  $\lambda_0 = 2$  and  $\rho \in [1/2, 1)$ , independent of  $n$ , and set

$$K_n = \left\lceil \frac{\log(2n^2)}{\log(1/\rho)} \right\rceil, \quad \Lambda_n = \{\lambda_k = 2\rho^k : k = 0, \dots, K_n\}.$$

Then  $|\Lambda_n| = O(\log n)$  and  $\lambda_{K_n} \leq n^{-2}$ . For a prescribed noise level  $A_n(\lambda)$ , we use the threshold  $\frac{3}{2}A_n(\lambda)$ . The rule computes the candidates from large to small regularization and returns the first point whose calibration statistic lies below this threshold.

*Remark (notation).*  $\hat{h}$  denotes an estimated primal solution, with  $\hat{h}_\lambda$  denoting a candidate at penalty  $\lambda$ ,  $\hat{h}_{j,\lambda}$  denoting an empirical iterate, and  $\hat{h}_{\lambda_{\text{dp}}}$  or  $\hat{h}_{t,\lambda_{\text{dp}}}$  denoting the corresponding discrepancy-selected estimate. In Proposition 3.8, Theorem 3.9, Lemma B.8, Corollary B.9, Lemma B.10:  $\hat{h}_{0,\lambda} = 0$  and  $\hat{h}_{j,\lambda} = \arg \min_{h \in \mathcal{H}} \{L_n(h) + \lambda \|h - \hat{h}_{j-1,\lambda}\|_{2,n}^2\}$  for  $j = 1, \dots, t$ , with the RDIV squared loss; the candidate used for discrepancy selection is  $\hat{h}_{t,\lambda}$ . In Corollary 3.19:  $\hat{h} = \hat{h}_{\lambda_{\text{dp}}}$  is the one-step adaptive adversarial Tikhonov estimate, with the primal tuning rule run at confidence level  $\xi_n/2$  and with the common deterministic sequence  $\delta_n$  specified for both nuisance problems. Its evaluations use jointly measurable versions relative to the full nuisance-training and tuning sigma-field  $\mathcal{G}_n$  and the observation, representing the fitted function for almost every training realization and agreeing almost surely with its evaluations at every observation in the finite evaluation sample. In Proposition 3.5:  $\hat{h}_\lambda = \arg \min_{h \in \mathcal{H}} \{L_n(h) + \lambda \mathbb{E}_n[h(X)^2]\}$  with  $L_n(h) = \mathbb{E}_{n,\text{fit}}[(Y - (\hat{\mathcal{T}}h)(Z))^2]$ . In Proposition 3.13:  $\hat{h}_\lambda = \arg \min_{h \in \mathcal{H}} \{L_n(h) + \lambda \mathbb{E}_n[h(X)^2]\}$ , where  $L_n(h) = \sup_{f \in \mathcal{F}} \mathbb{E}_n[2m(W; f) - 2h(X)f(Z) -$

$f(Z)^2]$ . In Proposition 3.17, Theorem 3.18, Lemma C.2, Corollary C.6, Lemma C.7:  $\hat{h}_{0,\lambda} = 0$  and  $\hat{h}_{j,\lambda} = \arg \min_{h \in \mathcal{H}} \{L_n(h) + \lambda \|h - \hat{h}_{j-1,\lambda}\|_{2,n}^2\}$  for  $j = 1, \dots, t$ , with the adversarial loss; the candidate used for discrepancy selection is  $\hat{h}_{t,\lambda}$ .  $\hat{h}_\lambda$  denotes the candidate solution at penalty  $\lambda$ , with the discrepancy search using the family indexed by  $\lambda \in \Lambda_n$  and the candidate construction specified by the estimation method. In Proposition 3.13, Theorem 3.14, Lemma C.1, Lemma C.2, Corollary C.3, Lemma C.4, a result cited in the proof of Lemma C.1:  $\hat{h}_\lambda$  is a measurable selected minimizer of  $L_n(h) + \lambda \mathbb{E}_n[h(X)^2]$  over  $h \in \mathcal{H}$ , where  $L_n(h) = \sup_{f \in \mathcal{F}} \mathbb{E}_n[2m(W; f) - 2h(X)f(Z) - f(Z)^2]$ . In Proposition 3.5, Theorem 3.6, Lemma B.2, Lemma B.3, Lemma B.4, Corollary B.5, Lemma B.6, a result cited in the proof of Lemma B.2:  $\hat{h}_\lambda$  is a measurable selected minimizer over  $h \in \mathcal{H}$  of  $\mathbb{E}_{n,\text{fit}}[(Y - \hat{\mathcal{T}}h(Z))^2] + \lambda \mathbb{E}_{n,\text{fit}}[h(X)^2]$ . In Proposition 3.8, Theorem 3.9, Lemma B.8: The candidate supplied to the search is  $\hat{h}_{t,\lambda}$ , with  $\hat{h}_{0,\lambda} = 0$  and measurable selected minimizers  $\hat{h}_{j,\lambda} \in \arg \min_{h \in \mathcal{H}} \{L_n(h) + \lambda \|h - \hat{h}_{j-1,\lambda}\|_{2,n}^2\}$  for  $j = 1, \dots, t$ , using the RDIV fitting loss. In Corollary 3.19: The primal candidate is the TRAE measurable selected minimizer of  $L_n(h) + \lambda \mathbb{E}_n[h(X)^2]$  over  $\mathcal{H}$ ; the corresponding dual construction exchanges the primal and dual spaces and operators. For the selected primal fit  $\hat{h}$ , take jointly measurable fitted evaluations representing the fitted function for almost every nuisance-training and tuning realization and agreeing almost surely with  $\hat{h}(X(W_i))$  simultaneously for  $i = 1, \dots, n$  in the evaluation sample. In Proposition D.1:  $\hat{h}_\lambda$  minimizes  $L_{N,\eta}(h) + \lambda \mathbb{E}_N[h(X)^2]$  over the full RKHS  $\mathcal{H}$ , taking the minimum-RKHS-norm representative among functions with the same fitting-sample evaluations; it has the form  $\hat{h}_\lambda(\cdot) = \sum_{j=1}^N \hat{\alpha}_{\lambda,j} k_{\mathcal{H}}(X_j, \cdot)$  with coefficients solving (D.1). In Proposition 3.17, Theorem 3.18, Lemma C.2, Corollary C.6, Lemma C.7: The candidate supplied to the search is  $\hat{h}_{t,\lambda}$ , with  $\hat{h}_{0,\lambda} = 0$  and measurable selected minimizers  $\hat{h}_{j,\lambda} \in \arg \min_{h \in \mathcal{H}} \{L_n(h) + \lambda \|h - \hat{h}_{j-1,\lambda}\|_{2,n}^2\}$  for  $j = 1, \dots, t$ , using the TRAE adversarial loss.  $D_n$  is the nonnegative discrepancy functional supplied to the search, evaluated at the candidate for each  $\lambda \in \Lambda_n$ , with its functional definition specified by the estimation method. In Proposition 3.5, Theorem 3.6, Proposition 3.8, Theorem 3.9, Lemma B.2, Corollary B.5, Lemma B.6, Lemma B.8, Corollary B.9, Lemma B.10:  $D_n(h) = \mathbb{E}_{n,\text{cal}}[(\hat{r}(Z) - \hat{\mathcal{T}}h(Z))^2]$ , the projected calibration residual; the same functional is evaluated at  $\hat{h}_{t,\lambda}$  for iterative candidates. In Corollary 3.19: For each nuisance problem, the discrepancy is its empirical adversarial loss:  $D_n = L_n$  for the primal problem, with the corresponding loss after exchanging the primal and dual spaces and operators for the dual problem. In Proposition 3.13, Theorem 3.14, Proposition 3.17, Theorem 3.18, Lemma C.4, Corollary C.6, Lemma C.7, a result cited in the proof of Lemma C.1:  $D_n(h) = L_n(h) = \sup_{f \in \mathcal{F}} \mathbb{E}_n[2m(W; f) - 2h(X)f(Z) - f(Z)^2]$ , with the same loss evaluated at the final iterate for iterative candidates.  $A_n : \Lambda_n \rightarrow [0, \infty)$  supplies the noise level in the acceptance threshold

$\frac{3}{2}A_n(\lambda)$ , with prescribed statistical-error bounds in the theoretical constructions and data-dependent noise-floor estimates in the experimental construction. In Theorem 3.1:  $A_n$  is prescribed, and is required to be nonincreasing only over  $\Lambda_n \cap [\lambda^\circ, 2]$ . In Proposition 3.5, Theorem 3.6:  $A_n = A_n^R$ , where  $A_n^R(\lambda) = \delta_n^2/\lambda + \nu_n^2$  for  $\lambda \in \Lambda_n$ . In Proposition 3.8, Theorem 3.9:  $A_n = A_{n,t}^R$ , where  $A_{n,t}^R(\lambda) = s_{n,t}^R/\lambda + \nu_n^2$  and  $s_{n,t}^R = a_t\delta_n^2$ . In Proposition 3.13, Theorem 3.14, Corollary 3.19:  $A_n = A_n^T$ , where  $A_n^T(\lambda) = \delta_n^2$  for  $\lambda \in \Lambda_n$ , using the TRAE accuracy scale for the relevant problem. In Proposition 3.17, Theorem 3.18:  $A_n = A_{n,t}^T$ , where  $A_{n,t}^T(\lambda) = \delta_n^2 + \min\{\lambda, s_{n,t}^T/\lambda\}$  and  $s_{n,t}^T = b_t\delta_n^2$ ; on the oracle segment  $[\lambda_{IT}^\circ, 2]$  this equals  $\delta_n^2 + s_{n,t}^T/\lambda$ .

---

**Algorithm 1** Adaptive Regularization via Discrepancy Principle

---

- 1: **Input:** Candidate estimators  $\hat{h}_\lambda$ , discrepancy  $D_n$ , noise level  $A_n$ .
  - 2: **Input:** Initial parameter  $\lambda_0 = 2$ , search factor  $1/2 \leq \rho < 1$ .
  - 3: Set  $K_n = \lceil \log(2n^2)/\log(1/\rho) \rceil$ .
  - 4: **for**  $k = 0, \dots, K_n$  **do**
  - 5:   Set  $\lambda = \lambda_0\rho^k$  and compute  $\hat{h}_\lambda$ .
  - 6:   Evaluate  $D_n(\hat{h}_\lambda)$ .
  - 7:   **if**  $D_n(\hat{h}_\lambda) \leq \frac{3}{2}A_n(\lambda)$  **then**
  - 8:     **Output:**  $\lambda$  and  $\hat{h}_\lambda$ ; terminate.
  - 9:   **end if**
  - 10: **end for**
  - 11: **Output:**  $\lambda_{K_n}$  and  $\hat{h}_{\lambda_{K_n}}$  (fallback, used only if no candidate clears its threshold).
- 

*Remark (notation).* In the discrepancy grid,  $\rho \in [1/2, 1)$  is the fixed search factor, independent of  $n$ ; the notation  $\rho$  in the experimental designs of Section 4 instead denotes their instrument-strength parameter, equal to the correlation of  $(X, Z)$  in the linear Gaussian design.

If no candidate clears its threshold, the rule returns the terminal point  $\lambda_{K_n}$ , so the selected estimator is defined for every outcome. Theorem 3.1 and Propositions 3.5 and 3.13 show that the rule accepts a candidate, and so never reaches the fallback, on the events they describe.

When the rule accepts a noninitial parameter, the selected point  $\lambda_{dp}$  and its rejected predecessor  $\lambda' = \lambda_{dp}/\rho \in [\lambda_{dp}, 2\lambda_{dp}]$  satisfy

$$D_n(\hat{h}_{\lambda_{dp}}) \leq \frac{3}{2}A_n(\lambda_{dp}), \quad D_n(\hat{h}_{\lambda'}) > \frac{3}{2}A_n(\lambda'). \quad (3.1)$$

Acceptance at  $\lambda_{dp}$  bounds the weak error; rejection at  $\lambda'$ , or acceptance at an oracle point, keeps  $\lambda_{dp}$  from falling too far.

*Remark (notation).*  $\lambda_{\text{dp}} = 2\rho^{k_*}$ , where  $k_* = \min\{k \in \{0, \dots, K_n\} : D_n(\widehat{h}_{2\rho^k}) \leq \frac{3}{2}A_n(2\rho^k)\}$  if this set is nonempty, and  $k_* = K_n$  otherwise, using the candidates, discrepancy, and noise level supplied to the rule. In Proposition 3.5, Theorem 3.6: The rule uses the one-step RDIV candidates  $\widehat{h}_\lambda$ , the projected calibration discrepancy, and  $A_n^R(\lambda) = \delta_n^2/\lambda + \nu_n^2$ . In Proposition 3.13, Theorem 3.14: The rule uses the one-step TRAE candidates,  $D_n = L_n$ , and  $A_n^T(\lambda) = \delta_n^2$ . In Corollary 3.19: The primal and dual TRAE problems are tuned separately at confidence levels  $\xi_n/2$  using a common deterministic accuracy bound  $\delta_n$ ; the dual output is  $\lambda_{\text{dp}}^{\text{dual}}$ . In Proposition 3.8: The rule uses the terminal iterative RDIV candidates  $\widehat{h}_{t,\lambda}$ , the projected calibration discrepancy, and  $A_{n,t}^R(\lambda) = s_{n,t}^R/\lambda + \nu_n^2$ ; its selected parameter is denoted  $\lambda_{\text{dp},t}^R$ . In Proposition 3.17: The rule uses the terminal iterative TRAE candidates and  $A_{n,t}^T(\lambda) = \delta_n^2 + \min\{\lambda, s_{n,t}^T/\lambda\}$ .

The statistic  $D_n(\widehat{h}_\lambda)$  plays the role of the squared residual in Definition 2.2, and  $\frac{3}{2}A_n(\lambda)$  plays the role of  $\kappa^2\delta^2$ . For RDIV,  $D_n$  is a squared residual computed from estimates of the operator and right-hand side. For TRAE, it is an empirical adversarial loss. The statistical bounds in (3.3) turn a small discrepancy into a small squared weak error. When the first accepted parameter has a predecessor, the two inequalities in (3.1) bracket the threshold at adjacent grid points with ratio  $1/\rho \in (1, 2]$ .

The following theorem states conditions on the discrepancy bound and threshold under which the rule has an oracle guarantee. Throughout,  $\xi \in (0, 1)$  is a generic confidence level; we absorb fixed numerical rescalings of  $\xi$  into the rate bounds and relabel it after intersecting finitely many event families.

*Remark (notation).*  $\xi \in (0, 1)$  is a prescribed failure-probability level, with fixed numerical rescalings absorbed into the rate bounds and the level relabelled after intersecting finitely many event families. In Corollary 3.19: The confidence level used in each primal and dual tuning rule is  $\xi_n/2$ , where  $\xi_n \downarrow 0$  is deterministic; the simultaneous nuisance bounds consequently have failure probability at most  $\xi_n$ .

**Theorem 3.1** (Oracle guarantee for the discrepancy rule). *Let  $A_n : \Lambda_n \rightarrow [0, \infty)$ , and let  $\mathcal{S}_n \subseteq (0, 2]$  be a stability region. Fix constants  $C_A, C_V, C_b > 0$  and exponents  $\gamma \in (0, 1]$  and  $q > 1$ , and write  $B_n(\lambda) = C_b\lambda^q$ . Suppose that, on an event  $\mathcal{E}_n$  with probability at least  $1 - \xi$ , the discrepancy bound*

$$D_n(\widehat{h}_\lambda) \leq A_n(\lambda) + B_n(\lambda), \quad (3.2)$$

*holds simultaneously over  $\lambda \in \Lambda_n \cap \mathcal{S}_n$ , and every  $\lambda \in \Lambda_n \cap \mathcal{S}_n$  satisfies*

$$\begin{aligned} \|\mathcal{T}(\widehat{h}_\lambda - h_0)\|^2 &\leq C_A\{D_n(\widehat{h}_\lambda) + A_n(\lambda)\}, \\ \|\widehat{h}_\lambda - h_0\|^2 &\leq C_V \left[ \frac{D_n(\widehat{h}_\lambda) + A_n(\lambda)}{\lambda} + \{D_n(\widehat{h}_\lambda) + A_n(\lambda)\}^\gamma \right], \end{aligned} \quad (3.3)$$

Suppose that some grid point  $\lambda^\circ$  satisfies  $\Lambda_n \cap [\lambda^\circ, 2] \subseteq \mathcal{S}_n$ , that  $A_n$  is nonincreasing over the grid points in  $[\lambda^\circ, 2]$ , and that

$$B_n(\lambda^\circ) \leq \frac{1}{2}A_n(\lambda^\circ). \quad (3.4)$$

Then Algorithm 1 accepts a candidate before exhausting the grid, and returns a parameter satisfying  $\lambda_{\text{dp}} \geq \lambda^\circ$  with

$$\|\mathcal{T}(\widehat{h}_{\lambda_{\text{dp}}} - h_0)\|^2 \leq O\{A_n(\lambda^\circ)\}, \quad (3.5)$$

$$\|\widehat{h}_{\lambda_{\text{dp}}} - h_0\|^2 \leq O\{A_n(\lambda^\circ)/\lambda^\circ + A_n(\lambda^\circ)^\gamma\}. \quad (3.6)$$

*Remark.* On  $\mathcal{E}_n$ , the algorithm terminates through its acceptance branch at an index  $k \leq k^\circ$ , where  $\lambda^\circ = 2\rho^{k^\circ}$  and  $k^\circ \leq K_n$ ; acceptance at  $k = K_n$  is allowed, and both error bounds hold on this event.

*Remark (notation).*  $\mathcal{S}$  is a fixed population function class whose local definition specifies its role and whether its functions have argument  $X$  or  $(X, Z)$ . In Corollary 3.19:  $\mathcal{S}$  is the dual test class appearing in  $L_n^{\text{dual}}(q) = \sup_{s \in \mathcal{S}} \mathbb{E}_n[2\tilde{m}(W; s) - 2q(Z)s(X) - s(X)^2]$ . In Lemma C.2, a result cited in the proof of Lemma C.2:  $\mathcal{S} = \text{star}(\mathcal{H} - \mathcal{H})$  is the star-shaped difference class used to verify the anchored critical-radius condition. In Lemma B.1: In the likelihood concentration argument,  $\mathcal{S} = \text{star}(\mathcal{G}_0 - \mathcal{G}_0)$  is the fixed star-shaped difference class of likelihood-transform functions on  $(X, Z)$ , with norm  $\|\cdot\|_{L_2(X, Z)}$ .  $\mathcal{S}_n \subseteq (0, 2]$  is the prescribed stability domain for the generic comparisons, which are required at grid points in  $\Lambda_n \cap \mathcal{S}_n$ . In Theorem 3.6, Theorem 3.14:  $\mathcal{S}_n = (0, 2]$ . In Theorem 3.9:  $\mathcal{S}_n = [\lambda_{\text{IR}}^\circ, 2]$  in the application of the generic comparison argument. In Theorem 3.18:  $\mathcal{S}_n = [\lambda_{\text{IT}}^\circ, 2]$  in the application of the generic comparison argument.  $\mathcal{E}_n$  is a common event of probability at least  $1 - \xi$  on which the relevant comparisons hold simultaneously over their stated function and grid domains. In Proposition 3.13, Theorem 3.14, Lemma C.1, Lemma C.2, Corollary C.3, Lemma C.4:  $\mathcal{E}_n$  is the common TRAE concentration event formed using  $\eta_n = \xi/(K_n + 1)$  and a union bound over  $\Lambda_n$ ; it includes the uniform loss comparison, the localized squared-norm comparison, and  $L_n(h_0) \leq C\delta_n^2$ .  $\lambda^\circ \in \Lambda_n$  is any oracle comparison point satisfying  $\Lambda_n \cap [\lambda^\circ, 2] \subseteq \mathcal{S}_n$ , nonincreasing  $A_n$  on that grid segment, and  $B_n(\lambda^\circ) \leq \frac{1}{2}A_n(\lambda^\circ)$ ; no uniqueness or computable selection rule is required.  $B$  is a fixed finite nonnegative uniform envelope bound for the quantities specified in the applicable boundedness assumptions. In Proposition 3.13, Theorem 3.14, Proposition 3.17, Theorem 3.18, Corollary 3.19, Lemma C.1, Lemma C.2, Corollary C.3, Lemma C.4, Lemma C.2, Corollary C.6, Lemma C.7, a result cited in the proof of Lemma C.1, a result cited in the proof of Lemma C.2:  $|m(W; f)|$ ,  $|h(X)|$ , and  $|f(Z)|$  are almost surely bounded by  $B$ , uniformly over  $h \in \mathcal{H}$  and  $f \in \mathcal{F}$ ; for a dual fit the corresponding bound uses the exchanged primal and dual objects. In

Proposition 3.5, Theorem 3.6, Proposition 3.8, Theorem 3.9, Lemma B.1, Lemma B.2, Corollary B.5, Lemma B.6, Lemma B.8, Corollary B.9, Lemma B.10, a result cited in the proof of Lemma B.2:  $|Y| \leq B$  almost surely,  $|h(x)| \leq B$  for every  $h \in \mathcal{H}$  on the common support of the conditional densities, and  $|\widehat{r}(z)| \leq B$ .

The rule uses the statistical-error bound  $A_n$  without evaluating the unknown bias bound  $B_n$ . It compares  $D_n$  with  $\frac{3}{2}A_n$ . Condition (3.4) guarantees that such a point exists, so the rule accepts  $\lambda^\circ$  whenever it reaches it; because the search runs from large to small regularization, it stops at some  $\lambda_{\text{dp}} \geq \lambda^\circ$ . Acceptance at  $\lambda_{\text{dp}}$  caps the discrepancy there by  $\frac{3}{2}A_n(\lambda_{\text{dp}})$ , and monotonicity carries  $A_n(\lambda)$  and  $A_n(\lambda)/\lambda$  back to  $\lambda^\circ$ , which gives (3.5).

We will specify  $D_n$ ,  $A_n$ , and  $\mathcal{S}_n$  for RDIV and TRAE in the following subsections, writing  $V_n$  there for the method-specific component of  $A_n(\lambda)/\lambda$ . The proof is in Section A.

**Noise calibration.** The tuning rule takes valid statistical-error bounds as inputs. In the following analyses,  $\delta_n$  and  $\nu_n$  denote deterministic upper bounds multiplied by fixed constants large enough to absorb the concentration and stability constants. These multipliers may depend on the boundedness, density-ratio, and mean-square-continuity constants, but are independent of  $n$ ,  $\lambda$ , and the unknown source parameter. They are fixed before fitting the candidate estimators. The constants  $a_t$  and  $b_t$  used for the iterative procedures also absorb the propagation constants for the prescribed iteration count. The discrepancy bounds in the Appendix show that such choices exist. This convention establishes adaptation to the source parameter conditional on valid bounds for the function classes and preliminary nuisances. Estimating all of these bounds from the data requires a separate calibration analysis.

*Remark (notation).*  $\delta_n$  is a prescribed deterministic, method-specific statistical-error bound varying with  $n$ , calibrated before fitting by fixed multipliers independent of  $n$ ,  $\lambda$ , and the unknown source parameter. In Proposition 3.8, Theorem 3.9, Lemma B.8, Corollary B.9, Lemma B.10: The iteration-grid confidence allocation is  $\eta_{n,t} = \xi/\{t(K_n + 1)\}$ ; the resulting  $\log t$  contribution is absorbed into a constant adjustment of  $\delta_n$  for fixed  $t$ . In Proposition 3.13, Theorem 3.14, Lemma C.1, Lemma C.2, Corollary C.3, Lemma C.4, a result cited in the proof of Lemma C.1, a result cited in the proof of Lemma C.2:  $\delta_n$  is the TRAE scale, unrelated to the RDIV scales, satisfying  $\delta_n \geq \max\{\widetilde{\delta}_n(\eta_n), n^{-1/2}\}$  and  $\delta_n = o(1)$ , with  $\eta_n = \xi/(K_n + 1)$ . In Corollary 3.19:  $\delta_n$  is a common deterministic TRAE bound for both primal and dual problems, each tuned at confidence level  $\xi_n/2$  for a deterministic  $\xi_n \downarrow 0$ , and additionally satisfies  $\delta_n = o(n^{-\alpha_{\text{dr}}})$ , where  $\alpha_{\text{dr}} = (1 + \min\{\beta_m, 1\})/(2 + 4 \min\{\beta_m, 1\})$  and  $\beta_m = \max\{\beta_h, \beta_q\}$ . In Proposition 3.17, Theorem 3.18, Lemma C.2, Corollary C.6, Lemma C.7:  $\delta_n$  is the TRAE scale satisfying  $\delta_n \geq \max\{\widetilde{\delta}_n(\eta_{n,t}), n^{-1/2}\}$  and  $\delta_n = o(1)$ , with  $\eta_{n,t} = \xi/\{t(K_n + 1)\}$ . In Proposition 3.5, Theorem 3.6, Proposition 3.8, Theorem 3.9,

Lemma B.1, Lemma B.2, Lemma B.3, Lemma B.4, Corollary B.5, Lemma B.6, Lemma B.8, Corollary B.9, Lemma B.10, a result cited in the proof of Lemma B.2:  $\delta_n$  is the RDIV fitting scale satisfying  $\delta_n \geq \delta_{n,\mathcal{G}} + \delta_{n,\mathcal{H}} + \delta_{n,\xi}$  and  $\delta_n + \nu_n = o(1)$ .  $\nu_n$  is a prescribed deterministic RDIV projected-discrepancy accuracy bound, fixed before fitting the candidates, with  $\nu_n \geq \delta_{n,\mathcal{G}} + \delta_{n,r} + \delta_{n,\xi}$  and  $\delta_n + \nu_n = o(1)$ ; its calibration multipliers are independent of  $n$ ,  $\lambda$ , and the unknown source parameter. In Proposition 3.8, Theorem 3.9, Lemma B.8, Corollary B.9, Lemma B.10: The bounds use the iteration-grid confidence allocation  $\eta_{n,t} = \xi/\{t(K_n + 1)\}$ ; the resulting  $\log t$  contribution is absorbed into a constant adjustment of  $\nu_n$  for the prescribed fixed  $t$ .  $a_t \geq 1$  is a deterministic calibrated RDIV propagation constant depending only on the prescribed  $t$  and fixed population quantities, independent of  $n$  and  $\lambda$ .  $b_t \geq 2$  is a calibrated deterministic propagation constant depending only on the prescribed iteration count  $t$  and the fixed boundedness and mean-square-continuity constants, independently of  $n$ ,  $\beta$ , and  $\|w_0\|$ . In Lemma C.2:  $b_t$  is the terminal member of a nondecreasing sequence with  $b_0 = 0$  and  $b_j \geq 2$  for  $j = 1, \dots, t$ ; in the proof one may choose  $b_j = C(1 + b_{j-1})$  for a sufficiently large fixed  $C \geq 2$ .

## 3.2 Adaptive Regularized DeepIV Estimator

In the IV setting,  $r_0(z) = \mathbb{E}[Y \mid Z = z]$  and  $\mathcal{T}h = \mathbb{E}[h(X) \mid Z = \cdot]$ ; thus, in this subsection,  $\tilde{\mathcal{Q}} = L_2(Z)$  and  $\mathcal{T}$  is the conditional expectation operator. The RDIV estimator first estimates the conditional density of  $X$  given  $Z$  with respect to a common dominating measure  $\mu$ , then uses the resulting operator in a penalized least-squares problem [Li et al., 2024]. For the analysis that follows, split the observations into three independent samples  $I_{\text{pre}}, I_{\text{fit}}, I_{\text{cal}}$ , each of size  $n$ . Fixed positive sample fractions give the same rates in the total sample size. On  $I_{\text{pre}}, \text{fit}$

$$\hat{g} = \arg \max_{g \in \mathcal{G}} \mathbb{E}_{n,\text{pre}}[\log g(X \mid Z)], \quad (\hat{\mathcal{T}}h)(z) = \int h(x) \hat{g}(x \mid z) \mu(dx),$$

and a regression estimate  $\hat{r}$  of  $r_0$  from  $Y$  on  $Z$ . Write  $\mathbb{E}_n = \mathbb{E}_{n,\text{fit}}$  within this subsection. The RDIV fit is

$$L_n(h) = \mathbb{E}_n[(Y - \hat{\mathcal{T}}h(Z))^2], \quad \hat{h}_\lambda = \arg \min_{h \in \mathcal{H}} \{L_n(h) + \lambda \mathbb{E}_n[h(X)^2]\}.$$

We assume these empirical minima are attained and measurable.

*Remark (notation).*  $Y$  is the scalar outcome, with  $r_0(z) = \mathbb{E}[Y \mid Z = z]$  in the IV setting. In Proposition 3.5, Theorem 3.6, Proposition 3.8, Theorem 3.9, Lemma B.1, Lemma B.2, Lemma B.3, Lemma B.4, Corollary B.5, Lemma B.6, Lemma B.8, Corollary B.9, Lemma B.10,

a result cited in the proof of Lemma B.2:  $|Y| \leq B$  almost surely for a fixed  $B < \infty$ .  $\hat{r}$  is the pilot regression estimate of  $r_0(z) = \mathbb{E}[Y \mid Z = z]$  fitted from  $Y$  on  $Z$  using  $I_{\text{pre}}$ . In conditioning expressions,  $I_{\text{pre}}$  denotes the observed preliminary sample, so conditioning on  $I_{\text{pre}}$  means conditioning on  $\sigma(W_i : i \in I_{\text{pre}})$ .  $I_{\text{fit}}$  is the fitting block of the three-way RDIV split, with  $|I_{\text{fit}}| = n$ , and  $\mathbb{E}_n = \mathbb{E}_{n,\text{fit}}$  in the RDIV analysis.  $I_{\text{cal}}$  is the calibration block of the three-way RDIV split, with  $|I_{\text{cal}}| = n$ .  $L_n$  denotes the empirical fitting loss on  $\mathcal{H}$ , with  $\mathbb{E}_n$  taken over the fitting observations of the applicable procedure. In Proposition 3.5, Theorem 3.6, Proposition 3.8, Theorem 3.9, Lemma B.2, Lemma B.3, Lemma B.4, Lemma B.6, Lemma B.7, Lemma B.8, Corollary B.9, Lemma B.10, a result cited in the proof of Lemma B.2:  $L_n(h) = \mathbb{E}_{n,\text{fit}}[(Y - \hat{\mathcal{T}}h(Z))^2]$ ; the iterative procedure retains this loss and changes the penalty center. In Proposition 3.13, Theorem 3.14, Proposition 3.17, Theorem 3.18, Corollary 3.19, Lemma C.1, Lemma C.2, Corollary C.3, Lemma C.4, Lemma C.2, Corollary C.6, Lemma C.7, a result cited in the proof of Lemma C.1:  $L_n(h) = \sup_{f \in \mathcal{F}} \mathbb{E}_n[2m(W; f) - 2h(X)f(Z) - f(Z)^2]$ , with  $\mathbb{E}_n$  the TRAE fitting-sample average; the iterative procedure retains this loss and changes the penalty center.

The population squared loss decomposes into an outcome-noise level and the squared weak error:

$$\sigma_{Y|Z}^2 := \mathbb{E}[(Y - r_0(Z))^2], \quad L(h) := \mathbb{E}[(Y - \mathcal{T}h(Z))^2] = \sigma_{Y|Z}^2 + \|\mathcal{T}(h - h_0)\|^2.$$

The raw loss exceeds the squared weak error by the constant  $\sigma_{Y|Z}^2$ , so it cannot be compared directly with a threshold for statistical error. We instead calibrate on  $I_{\text{cal}}$  with the projected residual

$$D_n(h) = \mathbb{E}_{n,\text{cal}}[(\hat{r}(Z) - \hat{\mathcal{T}}h(Z))^2]. \quad (3.7)$$

The preliminary sample makes  $\hat{r}$  and  $\hat{\mathcal{T}}$  independent of  $I_{\text{cal}}$ . At their population values, with ground-truth  $r_0$  and  $\mathcal{T}$ , for a fixed  $h$ ,

$$\mathbb{E}[D_n(h)] = \|r_0 - \mathcal{T}h\|^2 = \|\mathcal{T}(h - h_0)\|^2.$$

Calibration therefore requires one regression fit for  $\hat{r}$  plus one loss evaluation per candidate, on top of the  $O(\log n)$  regularized fits.

**Assumption 3.2** (Solution realizability). The minimum-norm solution  $h_0$  lies in  $\mathcal{H}$ .

**Assumption 3.3** (Conditional-density realizability and nuisance rate). The conditional density  $g_0$  belongs to  $\mathcal{G}$ , whose elements are conditional densities with respect to a common dominating measure. For fixed constants  $c_0, C_G > 0$ ,

$$g_0(x \mid z) \geq c_0, \quad \sup_{g \in \mathcal{G}} \|g\|_\infty \leq C_G.$$

For a prescribed  $\beta$ -independent accuracy bound  $\delta_{n,r} = \delta_{n,r}(\xi) \geq 0$ , the nuisance satisfies

$$\|\hat{r} - r_0\|_{L_2(Z)} \leq \delta_{n,r}$$

with probability at least  $1 - \xi$ .

*Remark (notation).*  $g_0(\cdot | z) \in \mathcal{G}$  is the true conditional density of  $X$  given  $Z = z$  with respect to  $\mu$ , normalized by  $\int g_0(x | z) \mu(dx) = 1$  for  $P_Z$ -almost every  $z$ , with  $g_0(x | z) \geq c_0$  on the common conditional-density support.  $c_0 > 0$  is a fixed finite constant, independent of sample size, such that  $g_0(x | z) \geq c_0$  on the common conditional-density support.  $C_G > 0$  is a fixed finite constant, independent of sample size, such that  $\sup_{g \in \mathcal{G}} \|g\|_\infty \leq C_G$  on the common conditional-density support.  $\delta_{n,r} = \delta_{n,r}(\xi) \geq 0$  is a prescribed deterministic,  $\beta$ -independent accuracy bound such that  $\|\hat{r} - r_0\|_{L_2(Z)} \leq \delta_{n,r}$  with probability at least  $1 - \xi$  over the preliminary sample. In Lemma B.1: The pilot-regression accuracy bound is evaluated at failure probability  $\xi/2$  and relabelled  $\delta_{n,r}$ , so here  $\delta_{n,r}$  denotes  $\delta_{n,r}(\xi/2)$  and satisfies  $\|\hat{r} - r_0\|_{L_2(Z)} \leq \delta_{n,r}$  with probability at least  $1 - \xi/2$  over the preliminary sample.

**Assumption 3.4** (Boundedness and critical radius). For a fixed  $B < \infty$ ,  $|Y| \leq B$  almost surely, every  $h \in \mathcal{H}$  satisfies  $|h(x)| \leq B$  on the common support of the conditional densities, and  $|\hat{r}(z)| \leq B$ . Clipping the nuisance to  $[-B, B]$  weakly decreases its  $L_2$  error. For  $g \in \mathcal{G}$ , write

$$(\mathcal{T}_g h)(z) = \int h(x) g(x | z) \mu(dx).$$

Define the likelihood-transform class

$$\mathcal{G}_0 = \left\{ (x, z) \mapsto \sqrt{\frac{g(x | z) + g_0(x | z)}{2g_0(x | z)}} : g \in \mathcal{G} \right\}.$$

Let  $\eta_n = \xi/(K_n + 1)$ . The classes below are measurable. The  $\beta$ -independent sequence  $\delta_{n,\mathcal{G}} = \delta_{n,\mathcal{G}}(\xi)$  upper-bounds the critical radius of  $\mathcal{G}_0$ , while  $\delta_{n,\mathcal{H}} = \delta_{n,\mathcal{H}}(\eta_n)$  is a common upper bound on the critical radii of

$$\mathcal{H} \quad \text{and} \quad \mathcal{T}_g \mathcal{H} := \{\mathcal{T}_g h : h \in \mathcal{H}\}, \quad g \in \mathcal{G},$$

under their respective input distributions. Moreover,

$$\delta_{n,\mathcal{G}} \geq C \sqrt{\frac{\log \log(en) + \log(1/\xi)}{n}}, \quad \delta_{n,\mathcal{H}} \geq C \sqrt{\frac{\log \log(en) + \log(1/\eta_n)}{n}}.$$

*Remark.* For  $\mathcal{H}$  and, for each fixed  $g \in \mathcal{G}$ ,  $\mathcal{T}_g \mathcal{H}$ , the empirical-process suprema used in the critical-radius analysis are measurable under their respective input distributions.

*Remark (notation).* For every  $g \in \mathcal{G}$ ,  $\mathcal{T}_g$  maps  $h \in \mathcal{H}$  to the  $P_Z$ -square-integrable function  $(\mathcal{T}_g h)(z) = \int h(x)g(x | z) \mu(dx)$ , with  $z$  understood  $P_Z$ -almost everywhere.  $\mathcal{G}_0$  is the measurable class  $\{(x, z) \mapsto \sqrt{\{g(x | z) + g_0(x | z)\}/\{2g_0(x | z)\}} : g \in \mathcal{G}\}$  on the common conditional-density support, with input distribution the true joint law of  $(X, Z)$  and positive denominator there.  $\eta_n = \xi/(K_n + 1)$  is the deterministic grid confidence allocation at the prescribed failure level  $\xi$ . In Corollary 3.19: For each of the primal and dual tuning rules, the prescribed failure level is  $\xi_n/2$ , so  $\eta_n = \xi_n/\{2(K_n + 1)\}$ , where  $\xi_n \downarrow 0$  is deterministic.  $\delta_{n,\mathcal{G}}$  is a prescribed deterministic,  $\beta$ -independent upper bound on the critical radius of  $\mathcal{G}_0$  under the joint input distribution of  $(X, Z)$ , evaluated at the failure probability allocated to the conditional-density fit and satisfying  $\delta_{n,\mathcal{G}} \geq C\sqrt{\{\log \log(en) + \log(1/\xi)\}/n}$ . In Lemma B.1: The conditional-density bound is evaluated at failure probability  $\xi/2$  and relabelled  $\delta_{n,\mathcal{G}}$ , so here  $\delta_{n,\mathcal{G}}$  denotes  $\delta_{n,\mathcal{G}}(\xi/2)$ .  $\delta_{n,\mathcal{H}} = \delta_{n,\mathcal{H}}(\eta_n)$  is a prescribed deterministic common upper bound on the critical radii of  $\mathcal{H}$  under the distribution of  $X$  and of  $\mathcal{T}_g \mathcal{H}$  under the distribution of  $Z$ , uniformly for  $g \in \mathcal{G}$ , with  $\eta_n = \xi/(K_n + 1)$  and  $\delta_{n,\mathcal{H}} \geq C\sqrt{\{\log \log(en) + \log(1/\eta_n)\}/n}$ . In Proposition 3.8, Theorem 3.9, Lemma B.8, Corollary B.9, Lemma B.10: The iterative concentration comparisons allocate failure probability  $\eta_{n,t} = \xi/\{t(K_n + 1)\}$  over iteration-grid pairs; the resulting confidence adjustment is absorbed into  $\delta_n$  for fixed  $t$ .

The density, boundedness, and complexity conditions follow the RDIV analysis of Li et al. [2024]; the nuisance rate controls the additional regression used for calibration. Define calibrated scales satisfying

$$\begin{aligned} \delta_{n,\xi} &:= \sqrt{\frac{\log((K_n + 1)/\xi)}{n}}, & \delta_n &\geq \delta_{n,\mathcal{G}} + \delta_{n,\mathcal{H}} + \delta_{n,\xi}, \\ \nu_n &\geq \delta_{n,\mathcal{G}} + \delta_{n,r} + \delta_{n,\xi}, & \delta_n + \nu_n &= o(1). \end{aligned} \quad (3.8)$$

The scale  $\delta_n$  controls the RDIV fitting error, whereas  $\nu_n$  controls the projected discrepancy.

Under the calibration convention above, Corollary B.5 gives a fixed constant  $C_R > 0$  such that, on the simultaneous RDIV event,

$$D_n(\hat{h}_\lambda) \leq \frac{\delta_n^2}{\lambda} + \nu_n^2 + C_R \|w_0\|^2 \lambda^p, \quad p = \min\{\beta + 1, 2\}, \quad (3.9)$$

for every  $\lambda \in \Lambda_n$ . The first two terms are the noise level used by the rule, while the last term is the regularization bias used only in the analysis. These bounds fit Theorem 3.1 with

$$A_n^R(\lambda) = \frac{\delta_n^2}{\lambda} + \nu_n^2, \quad V_n^R(\lambda) = \frac{\delta_n^2}{\lambda^2}, \quad \lambda \in \Lambda_n. \quad (3.10)$$

The rule therefore compares  $D_n$  with  $\frac{3}{2}A_n^R(\lambda)$ . As  $\lambda$  decreases, the fitting bound  $\delta_n^2/\lambda$  grows and the threshold loosens accordingly. The Appendix proves the error bounds in (3.3) with

$\gamma = \alpha := \min\{\beta, 1\}/(1 + \min\{\beta, 1\})$ . For a sufficiently small fixed  $c_R > 0$ , let  $\lambda_R^\circ$  be the first grid point no larger than

$$r_{n,\beta}^R = c_R \max\{\delta_n^{2/(p+1)}, \nu_n^{2/p}\}.$$

Then (3.4) holds and

$$A_n^R(\lambda_R^\circ) = O\{\delta_n^{2p/(p+1)} + \nu_n^2\}, \quad V_n^R(\lambda_R^\circ) = O\{\delta_n^{2(p-1)/(p+1)}\}. \quad (3.11)$$

*Remark (notation).*  $\alpha$  is local notation whose scalar-exponent and RKHS-coefficient meanings are distinct. In Proposition D.1:  $\alpha \in \mathbb{R}^N$  is the optimization variable representing  $h_\alpha(\cdot) = \sum_{j=1}^N \alpha_j k_{\mathcal{H}}(X_j, \cdot)$ ; its fitting-sample evaluation vector is  $K_{\mathcal{H}}\alpha$ . In Theorem 3.6, Lemma A.1, Lemma B.6, Lemma C.2, Lemma C.4:  $\alpha = \min\{\beta, 1\}/(1 + \min\{\beta, 1\}) = (p-1)/p$  is the scalar interpolation exponent. Whenever the defining set is nonempty,  $\lambda_R^\circ$  is the deterministic oracle point  $2\rho^{k^\circ}$ , where  $k^\circ = \min\{k \in \{0, \dots, K_n\} : 2\rho^k \leq c_R \max\{\delta_n^{2/(p+1)}, \nu_n^{2/p}\}\}$  and  $c_R > 0$  is a sufficiently small fixed constant; this point exists for all sufficiently large  $n$ .

Termination requires the last grid point to clear its threshold, which happens once the sample size is large enough. For RDIV the required sample size is

$$N_R(B) := \max\left\{2, \left\lceil \frac{8B^2}{3 \log 2} \right\rceil\right\},$$

which equals 4 under the normalization  $B = 1$ .

**Proposition 3.5.** *Suppose Assumptions 2.1 to 3.4 hold. For every  $n \geq N_R(B)$ , Algorithm 1, with the projected statistic (3.7) and noise level (3.10), accepts a candidate after at most  $K_n + 1 = O(\log n)$  regularized fits, with probability at least  $1 - \xi$ . If  $h_0 = 0$ , the initial parameter is accepted. In general, the returned parameter either is the initial parameter or is a noninitial grid point satisfying (3.1).*

*Additional conditions.* For the conditional fitting argument in Lemma B.2, use the following measurable-version convention. Let  $Q$  denote the law of one observation  $W$ , let  $a$  denote the preliminary observations, and put  $Q_n^{\mathbf{w}} = n^{-1} \sum_{i=1}^n \delta_{w_i}$  for a fresh fitting tuple  $\mathbf{w} = (w_1, \dots, w_n)$ . Define

$$\ell_{\lambda,a}(h; w) = \{Y(w) - (\widehat{\mathcal{T}}_a h)(Z(w))\}^2 + \lambda h(X(w))^2, \quad S_{n,\lambda,q}(a, \mathbf{w}) = \sup_{\substack{h \in \mathcal{H} \\ \|h - h_0\| \leq q}} |(Q_n^{\mathbf{w}} - Q)\{\ell_{\lambda,a}(h) - \ell_{\lambda,a}(h_0)\}|$$

For each  $n \geq 1$ ,  $\lambda \in \Lambda_n$  and positive rational  $q$ ,  $S_{n,\lambda,q}$  admits an extended nonnegative version jointly measurable in the preliminary observations and the fresh fitting tuple. This version agrees with  $S_{n,\lambda,q}(a, \cdot)$  for  $Q^n$ -almost every tuple for almost every preliminary realization, and agrees almost surely with  $S_{n,\lambda,q}(I_{\text{pre}}, (W_i)_{i \in I_{\text{fit}}})$  at the actual independent fitting block. All losses

use the same function representatives as the empirical objective. This is only a measurability and evaluation-coherence convention; no deviation bound, rate, countability or separability is assumed. For the fitted conditional density in the RDIV preliminary construction, use a coherent measurable version at independent calibration observations. For each  $n \geq 1$ , writing  $a = I_{\text{pre}}$ , the selected density  $\hat{g}_{n,a}(x \mid z)$  admits a nonnegative version  $\bar{g}_n(a, x, z)$  jointly measurable in the preliminary observations,  $x$  and  $z$ . For almost every preliminary realization  $a$ , for  $P_Z$ -almost every  $z$ , the functions  $\bar{g}_n(a, \cdot, z)$  and  $\hat{g}_{n,a}(\cdot \mid z)$  agree  $\mu$ -almost everywhere. Almost surely, this same agreement holds at every actual calibration observation  $Z_i$ ,  $i \in I_{\text{cal}}$ . The empirical objectives and selected maximizer remain those defined using the original supplied representatives: this convention does not replace the optimizer. It supplies only measurability and coherence of the fitted density evaluated on the independent calibration sample; it assumes no normalization event, probability inequality, rate, countability or separability. For the preliminary regression estimate used in the RDIV calibration loss, take a coherent measurable version on the independent calibration sample. For each  $n \geq 1$ , writing  $a = I_{\text{pre}}$ , the selected  $L_2(Z)$  pilot has a version  $\bar{r}_n(a, z)$  jointly measurable in the preliminary observations and  $z$ . For almost every preliminary realization  $a$ ,  $\bar{r}_n(a, z) = \hat{r}_{n,a}(z)$  for  $P_Z$ -almost every  $z$ . Almost surely, the same equality holds simultaneously at every actual calibration observation  $Z_i$ ,  $i \in I_{\text{cal}}$ , for the supplied representatives used in  $D_n$ . This is only a measurability and evaluation-coherence convention. It imposes no exact joint measurability on every supplied representative, does not replace an empirical objective or selected optimizer, and supplies no boundedness event, probability inequality, error rate, countability or separability assumption. The existing pilot accuracy and envelope assumptions remain unchanged. Use the following well-definedness convention for the likelihood-transform critical-radius and concentration arguments in Assumption 3.4 and Lemma B.1, separately from the existing hypothesis-class, density-image and fitted-process conventions. Let  $Q$  be the true joint law of  $(X, Z)$ , let  $\mathcal{G}_0$  be the likelihood-transform class already defined in Assumption 3.4, and put

$$\mathcal{A}_r = \{f \in \text{star}(\mathcal{G}_0 - \mathcal{G}_0) : \|f\|_{L_2(Q)} \leq r\}.$$

For every integer  $n \geq 1$  and every  $r > 0$ , the three processes

$$\sup_{f \in \mathcal{G}_0} |(Q_n - Q)f|, \quad \sup_{f \in \mathcal{A}_r} |(Q_n - Q)f|, \quad \sup_{f \in \mathcal{A}_r} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i f(X_i, Z_i) \right|$$

regarded as extended nonnegative quantities, are measurable under  $Q^n$  for the first two and under the product of  $Q^n$  with the independent Rademacher law for the third. The functions use the same fixed density representatives as the likelihood objective. This convention only makes these empirical processes and their expectations well defined. It supplies no concentration

event, deviation inequality, probability bound, rate, countability, separability, or additional statistical restriction.

*Remark.* For every preliminary sample size  $n \geq 1$ , the empirical log likelihood  $g \mapsto \mathbb{E}_{n,\text{pre}}[\log g(X | Z)]$  attains its maximum over  $\mathcal{G}$  almost surely, and  $\hat{g} \in \mathcal{G}$  is a measurable maximizer selected using only  $I_{\text{pre}}$ . For each fixed  $n \geq N_R(B)$ , all probabilistic conclusions, including initial acceptance when  $h_0 = 0$ , hold on the same event  $\mathcal{E}_n^R$  of probability at least  $1 - \xi$ . For every  $\lambda \in \Lambda_n$ , the empirical RDIV minimum is attained and  $\hat{h}_\lambda$  is a measurable selected minimizer of  $L_n(h) + \lambda \mathbb{E}_n[h(X)^2]$  over  $h \in \mathcal{H}$ .

**Theorem 3.6.** *Under the assumptions of Proposition 3.5, let  $\lambda_{\text{dp}}$  be its output and write*

$$p = \min\{\beta + 1, 2\}, \quad \alpha = \frac{\min\{\beta, 1\}}{1 + \min\{\beta, 1\}}, \quad \varepsilon_{n,\beta} = \delta_n^{2p/(p+1)} + \nu_n^2.$$

*For every  $n \geq N_R(B)$ , with probability at least  $1 - \xi$ ,*

$$\|\mathcal{T}(\hat{h}_{\lambda_{\text{dp}}} - h_0)\|^2 \leq O(\varepsilon_{n,\beta}), \quad \|\hat{h}_{\lambda_{\text{dp}}} - h_0\|^2 \leq O(\delta_n^{2(p-1)/(p+1)} + \varepsilon_{n,\beta}^\alpha) \leq O(\varepsilon_{n,\beta}^\alpha).$$

*If  $h_0 = 0$ , the squared strong and weak errors are both  $O(\delta_n^2)$ .*

*Additional conditions. For the conditional fitting argument in Lemma B.2, use the following measurable-version convention. Let  $Q$  denote the law of one observation  $W$ , let  $a$  denote the preliminary observations, and put  $Q_n^{\mathbf{w}} = n^{-1} \sum_{i=1}^n \delta_{w_i}$  for a fresh fitting tuple  $\mathbf{w} = (w_1, \dots, w_n)$ . Define*

$$\ell_{\lambda,a}(h; w) = \{Y(w) - (\hat{\mathcal{T}}_a h)(Z(w))\}^2 + \lambda h(X(w))^2, \quad S_{n,\lambda,q}(a, \mathbf{w}) = \sup_{\substack{h \in \mathcal{H} \\ \|h - h_0\| \leq q}} |(Q_n^{\mathbf{w}} - Q)\{\ell_{\lambda,a}(h) - \ell_{\lambda,a}(h_0)\}|$$

*For each  $n \geq 1$ ,  $\lambda \in \Lambda_n$  and positive rational  $q$ ,  $S_{n,\lambda,q}$  admits an extended nonnegative version jointly measurable in the preliminary observations and the fresh fitting tuple. This version agrees with  $S_{n,\lambda,q}(a, \cdot)$  for  $Q^n$ -almost every tuple for almost every preliminary realization, and agrees almost surely with  $S_{n,\lambda,q}(I_{\text{pre}}, (W_i)_{i \in I_{\text{fit}}})$  at the actual independent fitting block. All losses use the same function representatives as the empirical objective. This is only a measurability and evaluation-coherence convention; no deviation bound, rate, countability or separability is assumed. For the fitted conditional density in the RDIV preliminary construction, use a coherent measurable version at independent calibration observations. For each  $n \geq 1$ , writing  $a = I_{\text{pre}}$ , the selected density  $\hat{g}_{n,a}(x | z)$  admits a nonnegative version  $\bar{g}_n(a, x, z)$  jointly measurable in the preliminary observations,  $x$  and  $z$ . For almost every preliminary realization  $a$ , for  $P_Z$ -almost every  $z$ , the functions  $\bar{g}_n(a, \cdot, z)$  and  $\hat{g}_{n,a}(\cdot | z)$  agree  $\mu$ -almost everywhere. Almost surely, this same agreement holds at every actual calibration observation*

$Z_i, i \in I_{\text{cal}}$ . The empirical objectives and selected maximizer remain those defined using the original supplied representatives: this convention does not replace the optimizer. It supplies only measurability and coherence of the fitted density evaluated on the independent calibration sample; it assumes no normalization event, probability inequality, rate, countability or separability. For the preliminary regression estimate used in the RDIV calibration loss, take a coherent measurable version on the independent calibration sample. For each  $n \geq 1$ , writing  $a = I_{\text{pre}}$ , the selected  $L_2(Z)$  pilot has a version  $\bar{r}_n(a, z)$  jointly measurable in the preliminary observations and  $z$ . For almost every preliminary realization  $a$ ,  $\bar{r}_n(a, z) = \hat{r}_{n,a}(z)$  for  $P_Z$ -almost every  $z$ . Almost surely, the same equality holds simultaneously at every actual calibration observation  $Z_i, i \in I_{\text{cal}}$ , for the supplied representatives used in  $D_n$ . This is only a measurability and evaluation-coherence convention. It imposes no exact joint measurability on every supplied representative, does not replace an empirical objective or selected optimizer, and supplies no boundedness event, probability inequality, error rate, countability or separability assumption. The existing pilot accuracy and envelope assumptions remain unchanged. Use the following well-definedness convention for the likelihood-transform critical-radius and concentration arguments in Assumption 3.4 and Lemma B.1, separately from the existing hypothesis-class, density-image and fitted-process conventions. Let  $Q$  be the true joint law of  $(X, Z)$ , let  $\mathcal{G}_0$  be the likelihood-transform class already defined in Assumption 3.4, and put

$$\mathcal{A}_r = \{f \in \text{star}(\mathcal{G}_0 - \mathcal{G}_0) : \|f\|_{L_2(Q)} \leq r\}.$$

For every integer  $n \geq 1$  and every  $r > 0$ , the three processes

$$\sup_{f \in \mathcal{G}_0} |(Q_n - Q)f|, \quad \sup_{f \in \mathcal{A}_r} |(Q_n - Q)f|, \quad \sup_{f \in \mathcal{A}_r} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i f(X_i, Z_i) \right|$$

regarded as extended nonnegative quantities, are measurable under  $Q^n$  for the first two and under the product of  $Q^n$  with the independent Rademacher law for the third. The functions use the same fixed density representatives as the likelihood objective. This convention only makes these empirical processes and their expectations well defined. It supplies no concentration event, deviation inequality, probability bound, rate, countability, separability, or additional statistical restriction.

*Remark.* For each fixed  $n \geq N_R(B)$ , the displayed error bounds and, when  $h_0 = 0$ , both sharper  $O(\delta_n^2)$  bounds hold jointly with probability at least  $1 - \xi$  over the preliminary, fitting, and calibration samples.

If  $\nu_n = O(\delta_n)$ , fitting and projected calibration operate at the same statistical rate, and the theorem gives

$$\|\hat{h}_{\lambda_{\text{dp}}} - h_0\|^2 = O(\delta_n^{2(p-1)/(p+1)}), \quad \|\mathcal{T}(\hat{h}_{\lambda_{\text{dp}}} - h_0)\|^2 = O(\delta_n^{2p/(p+1)}).$$

In particular, for  $0 < \beta \leq 1$ , the squared strong-error rate is  $O(\delta_n^{2\beta/(\beta+2)})$ , which matches the noniterative RDIV rate obtained when  $\beta$  is known [Li et al., 2024, Theorem 7]. For  $\beta > 1$ , the projected-residual bias bound uses smoothness up to  $\beta = 1$ , giving the squared strong-error rate  $O(\delta_n^{2/3})$ . A nuisance with  $\delta_{n,r} = O(\delta_n)$  suffices for these rates. More generally, the nuisance rate enters the calibration term through  $\nu_n^2$ . The proof applies Theorem 3.1 with (3.10)–(3.11) by establishing the discrepancy bound and the weak- and strong-error bounds at every grid point for RDIV.

### 3.3 Adaptive Iterative Regularized DeepIV Estimator

Iterated Tikhonov regularization extends the range of source smoothness used by RDIV and improves the convergence rates [Li et al., 2024]. Fix an integer  $t \geq 1$ , chosen independently of the data and the source parameter  $\beta$ , and initialize

$$h_{0,\lambda}^* = \hat{h}_{0,\lambda} = 0.$$

For  $j = 1, \dots, t$ , define the population and empirical iterates by

$$h_{j,\lambda}^* = \arg \min_{h \in \mathcal{H}} \left\{ \|\mathcal{T}h - r_0\|^2 + \lambda \|h - h_{j-1,\lambda}^*\|^2 \right\}, \quad (3.12)$$

$$\hat{h}_{j,\lambda} = \arg \min_{h \in \mathcal{H}} \left\{ L_n(h) + \lambda \|h - \hat{h}_{j-1,\lambda}\|_{2,n}^2 \right\}. \quad (3.13)$$

We assume that each empirical minimum is attained and that the selected minimizers are measurable. The first iterate equals the RDIV estimator in Section 3.2. The iterative candidate at  $\lambda$  is  $\hat{h}_{t,\lambda}$ , and its discrepancy remains the projected calibration loss  $D_n(\hat{h}_{t,\lambda})$  in (3.7).

*Remark (notation).*  $h_{j,\lambda}^*$  is the population minimizer over  $\mathcal{H}$  for positive finite  $\lambda$ , initialized by  $h_{0,\lambda}^* = 0$  and recursively defined for positive integer  $j$ ; the iterative procedures use  $j = 1, \dots, t$ . In Theorem 3.18, Lemma C.2, Corollary C.6, Lemma C.7:  $h_{j,\lambda}^* \in \mathcal{H}$  for  $j = 1, \dots, t$  and  $\lambda \in \Lambda_n \cap [\lambda_{\text{IT}}^\circ, 2]$ , for every  $n$  for which the oracle point is defined. In Theorem 3.9, Lemma B.8, Corollary B.9, Lemma B.10:  $h_{j,\lambda}^* \in \mathcal{H}$  for  $j = 1, \dots, t$  and  $\lambda \in \Lambda_n \cap [\lambda_{\text{IR}}^\circ, 2]$ , for every  $n$  for which the oracle point is defined.

For the  $t(K_n + 1)$  fitted iterates, set

$$\eta_{n,t} = \frac{\xi}{t(K_n + 1)}.$$

We continue to write  $\delta_n$  and  $\nu_n$  as in (3.8). Because  $t$  is fixed, the resulting  $\log t$  term is a constant adjustment to  $\delta_n$  and  $\nu_n$ .

Let  $a_t \geq 1$  be a deterministic constant that depends only on the prescribed iteration count  $t$  and fixed population quantities. Define

$$q_t = \min\{\beta + 1, 2t\}, \quad s_{n,t}^R = a_t \delta_n^2.$$

We then use

$$A_{n,t}^R(\lambda) = \frac{s_{n,t}^R}{\lambda} + \nu_n^2, \quad V_{n,t}^R(\lambda) = \frac{s_{n,t}^R}{\lambda^2}, \quad (3.14)$$

for the noise level and stability bound and run Algorithm 1 with the threshold  $\frac{3}{2}A_{n,t}^R(\lambda)$ . For a sufficiently small fixed  $c_{\text{IR}} > 0$ , let  $\lambda_{\text{IR}}^\circ$  be the first grid point no larger than

$$r_{n,t,\beta}^{\text{IR}} = c_{\text{IR}} \max\{(s_{n,t}^R)^{1/(q_t+1)}, \nu_n^{2/q_t}\}.$$

For all sufficiently large  $n$ , this point satisfies (3.4), and

$$A_{n,t}^R(\lambda_{\text{IR}}^\circ) = O_t\{(s_{n,t}^R)^{q_t/(q_t+1)} + \nu_n^2\}, \quad V_{n,t}^R(\lambda_{\text{IR}}^\circ) = O_t\{(s_{n,t}^R)^{(q_t-1)/(q_t+1)}\}. \quad (3.15)$$

*Remark (notation).* Whenever the defining set is nonempty,  $\lambda_{\text{IR}}^\circ$  is the deterministic oracle point  $2\rho^{k^\circ}$ , where  $k^\circ = \min\{k \in \{0, \dots, K_n\} : 2\rho^k \leq c_{\text{IR}} \max\{(s_{n,t}^R)^{1/(q_t+1)}, \nu_n^{2/q_t}\}\}$  and  $c_{\text{IR}} > 0$  is a sufficiently small fixed constant; this point exists for all sufficiently large  $n$ .

**Assumption 3.7** (Oracle-region iterative RDIV realizability). For every  $n$  for which  $\lambda_{\text{IR}}^\circ$  is defined,  $h_{j,\lambda}^* \in \mathcal{H}$  for  $j = 1, \dots, t$  and  $\lambda \in \Lambda_n \cap [\lambda_{\text{IR}}^\circ, 2]$ .

On this segment, Corollary B.9 gives

$$D_n(\widehat{h}_{t,\lambda}) \leq A_{n,t}^R(\lambda) + C_{\text{IR}}\|w_0\|^2 \lambda^{q_t}, \quad \lambda \in \Lambda_n \cap [\lambda_{\text{IR}}^\circ, 2]. \quad (3.16)$$

**Proposition 3.8.** Suppose the assumptions of Proposition 3.5 hold. For every  $n \geq N_{\text{R}}(B)$ , Algorithm 1, applied to  $\widehat{h}_{t,\lambda}$  with the noise level (3.14), accepts a candidate with probability at least  $1 - \xi$ . It solves at most  $t(K_n + 1) = O(t \log n)$  regularized least-squares problems. If  $h_0 = 0$ , the initial parameter is accepted. Otherwise, the output is either the initial parameter or a noninitial point satisfying (3.1) with  $D_n(\widehat{h}_{t,\lambda})$  and  $A_{n,t}^R(\lambda)$ .

Additional conditions. For the conditional fitting argument in Lemma B.2, use the following measurable-version convention. Let  $Q$  denote the law of one observation  $W$ , let  $a$  denote the preliminary observations, and put  $Q_n^{\mathbf{w}} = n^{-1} \sum_{i=1}^n \delta_{w_i}$  for a fresh fitting tuple  $\mathbf{w} = (w_1, \dots, w_n)$ . Define

$$\ell_{\lambda,a}(h; w) = \{Y(w) - (\widehat{\mathcal{T}}_a h)(Z(w))\}^2 + \lambda h(X(w))^2, \quad S_{n,\lambda,q}(a, \mathbf{w}) = \sup_{\substack{h \in \mathcal{H} \\ \|h - h_0\| \leq q}} |(Q_n^{\mathbf{w}} - Q)\{\ell_{\lambda,a}(h) - \ell_{\lambda,a}(h_0)\}|$$

For each  $n \geq 1$ ,  $\lambda \in \Lambda_n$  and positive rational  $q$ ,  $S_{n,\lambda,q}$  admits an extended nonnegative version jointly measurable in the preliminary observations and the fresh fitting tuple. This version agrees with  $S_{n,\lambda,q}(a, \cdot)$  for  $Q^n$ -almost every tuple for almost every preliminary realization, and agrees almost surely with  $S_{n,\lambda,q}(I_{\text{pre}}, (W_i)_{i \in I_{\text{fit}}})$  at the actual independent fitting block. All losses use the same function representatives as the empirical objective. This is only a measurability and evaluation-coherence convention; no deviation bound, rate, countability or separability is assumed. For the fitted conditional density in the RDIV preliminary construction, use a coherent measurable version at independent calibration observations. For each  $n \geq 1$ , writing  $a = I_{\text{pre}}$ , the selected density  $\hat{g}_{n,a}(x \mid z)$  admits a nonnegative version  $\bar{g}_n(a, x, z)$  jointly measurable in the preliminary observations,  $x$  and  $z$ . For almost every preliminary realization  $a$ , for  $P_Z$ -almost every  $z$ , the functions  $\bar{g}_n(a, \cdot, z)$  and  $\hat{g}_{n,a}(\cdot \mid z)$  agree  $\mu$ -almost everywhere. Almost surely, this same agreement holds at every actual calibration observation  $Z_i$ ,  $i \in I_{\text{cal}}$ . The empirical objectives and selected maximizer remain those defined using the original supplied representatives: this convention does not replace the optimizer. It supplies only measurability and coherence of the fitted density evaluated on the independent calibration sample; it assumes no normalization event, probability inequality, rate, countability or separability. For the preliminary regression estimate used in the RDIV calibration loss, take a coherent measurable version on the independent calibration sample. For each  $n \geq 1$ , writing  $a = I_{\text{pre}}$ , the selected  $L_2(Z)$  pilot has a version  $\bar{r}_n(a, z)$  jointly measurable in the preliminary observations and  $z$ . For almost every preliminary realization  $a$ ,  $\bar{r}_n(a, z) = \hat{r}_{n,a}(z)$  for  $P_Z$ -almost every  $z$ . Almost surely, the same equality holds simultaneously at every actual calibration observation  $Z_i$ ,  $i \in I_{\text{cal}}$ , for the supplied representatives used in  $D_n$ . This is only a measurability and evaluation-coherence convention. It imposes no exact joint measurability on every supplied representative, does not replace an empirical objective or selected optimizer, and supplies no boundedness event, probability inequality, error rate, countability or separability assumption. The existing pilot accuracy and envelope assumptions remain unchanged. Use the following well-definedness convention for the likelihood-transform critical-radius and concentration arguments in Assumption 3.4 and Lemma B.1, separately from the existing hypothesis-class, density-image and fitted-process conventions. Let  $Q$  be the true joint law of  $(X, Z)$ , let  $\mathcal{G}_0$  be the likelihood-transform class already defined in Assumption 3.4, and put

$$\mathcal{A}_r = \{f \in \text{star}(\mathcal{G}_0 - \mathcal{G}_0) : \|f\|_{L_2(Q)} \leq r\}.$$

For every integer  $n \geq 1$  and every  $r > 0$ , the three processes

$$\sup_{f \in \mathcal{G}_0} |(Q_n - Q)f|, \quad \sup_{f \in \mathcal{A}_r} |(Q_n - Q)f|, \quad \sup_{f \in \mathcal{A}_r} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i f(X_i, Z_i) \right|$$

regarded as extended nonnegative quantities, are measurable under  $Q^n$  for the first two and under the product of  $Q^n$  with the independent Rademacher law for the third. The functions use the same fixed density representatives as the likelihood objective. This convention only makes these empirical processes and their expectations well defined. It supplies no concentration event, deviation inequality, probability bound, rate, countability, separability, or additional statistical restriction.

*Remark.* For each fixed  $n \geq N_R(B)$  and prescribed  $t$ , there is an event of probability at least  $1 - \xi$  on which the rule accepts a candidate and satisfies the stated output alternatives, including acceptance of  $\lambda_0 = 2$  when  $h_0 = 0$ .

**Theorem 3.9** (Adaptive iterative RDIV). *Under the assumptions of Proposition 3.8 and Assumption 3.7, let  $\lambda_{\text{dp},t}^R$  denote the selected parameter and set*

$$\varepsilon_{n,t,\beta}^R = (s_{n,t}^R)^{q_t/(q_t+1)} + \nu_n^2.$$

For every fixed  $t$  and every  $n \geq N_R(B)$ , with probability at least  $1 - \xi$ ,

$$\|\mathcal{T}(\hat{h}_{t,\lambda_{\text{dp},t}^R} - h_0)\|^2 \leq O_t(\varepsilon_{n,t,\beta}^R), \quad \|\hat{h}_{t,\lambda_{\text{dp},t}^R} - h_0\|^2 \leq O_t((s_{n,t}^R)^{(q_t-1)/(q_t+1)} + (\varepsilon_{n,t,\beta}^R)^{\beta/(\beta+1)}).$$

Here  $O_t(\cdot)$  allows constants that depend on the prescribed iteration count and fixed population quantities. In particular,  $a_t$  is independent of  $n$ . If  $\nu_n = O(\delta_n)$  and  $2t \geq \beta + 1$ , the bounds reduce to

$$\|\hat{h}_{t,\lambda_{\text{dp},t}^R} - h_0\|^2 = O_t(\delta_n^{2\beta/(\beta+2)}), \quad \|\mathcal{T}(\hat{h}_{t,\lambda_{\text{dp},t}^R} - h_0)\|^2 = O_t(\delta_n^{2(\beta+1)/(\beta+2)}).$$

*Additional conditions.* For the conditional fitting argument in Lemma B.2, use the following measurable-version convention. Let  $Q$  denote the law of one observation  $W$ , let  $a$  denote the preliminary observations, and put  $Q_n^{\mathbf{w}} = n^{-1} \sum_{i=1}^n \delta_{w_i}$  for a fresh fitting tuple  $\mathbf{w} = (w_1, \dots, w_n)$ . Define

$$\ell_{\lambda,a}(h; w) = \{Y(w) - (\hat{\mathcal{T}}_a h)(Z(w))\}^2 + \lambda h(X(w))^2, \quad S_{n,\lambda,q}(a, \mathbf{w}) = \sup_{\substack{h \in \mathcal{H} \\ \|h - h_0\| \leq q}} |(Q_n^{\mathbf{w}} - Q)\{\ell_{\lambda,a}(h) - \ell_{\lambda,a}(h_0)\}|$$

For each  $n \geq 1$ ,  $\lambda \in \Lambda_n$  and positive rational  $q$ ,  $S_{n,\lambda,q}$  admits an extended nonnegative version jointly measurable in the preliminary observations and the fresh fitting tuple. This version agrees with  $S_{n,\lambda,q}(a, \cdot)$  for  $Q^n$ -almost every tuple for almost every preliminary realization, and agrees almost surely with  $S_{n,\lambda,q}(I_{\text{pre}}, (W_i)_{i \in I_{\text{fit}}})$  at the actual independent fitting block. All losses use the same function representatives as the empirical objective. This is only a measurability and evaluation-coherence convention; no deviation bound, rate, countability or separability

is assumed. For the fitted conditional density in the RDIV preliminary construction, use a coherent measurable version at independent calibration observations. For each  $n \geq 1$ , writing  $a = I_{\text{pre}}$ , the selected density  $\hat{g}_{n,a}(x \mid z)$  admits a nonnegative version  $\bar{g}_n(a, x, z)$  jointly measurable in the preliminary observations,  $x$  and  $z$ . For almost every preliminary realization  $a$ , for  $P_Z$ -almost every  $z$ , the functions  $\bar{g}_n(a, \cdot, z)$  and  $\hat{g}_{n,a}(\cdot \mid z)$  agree  $\mu$ -almost everywhere. Almost surely, this same agreement holds at every actual calibration observation  $Z_i$ ,  $i \in I_{\text{cal}}$ . The empirical objectives and selected maximizer remain those defined using the original supplied representatives: this convention does not replace the optimizer. It supplies only measurability and coherence of the fitted density evaluated on the independent calibration sample; it assumes no normalization event, probability inequality, rate, countability or separability. For the preliminary regression estimate used in the RDIV calibration loss, take a coherent measurable version on the independent calibration sample. For each  $n \geq 1$ , writing  $a = I_{\text{pre}}$ , the selected  $L_2(Z)$  pilot has a version  $\bar{r}_n(a, z)$  jointly measurable in the preliminary observations and  $z$ . For almost every preliminary realization  $a$ ,  $\bar{r}_n(a, z) = \hat{r}_{n,a}(z)$  for  $P_Z$ -almost every  $z$ . Almost surely, the same equality holds simultaneously at every actual calibration observation  $Z_i$ ,  $i \in I_{\text{cal}}$ , for the supplied representatives used in  $D_n$ . This is only a measurability and evaluation-coherence convention. It imposes no exact joint measurability on every supplied representative, does not replace an empirical objective or selected optimizer, and supplies no boundedness event, probability inequality, error rate, countability or separability assumption. The existing pilot accuracy and envelope assumptions remain unchanged. Use the following well-definedness convention for the likelihood-transform critical-radius and concentration arguments in Assumption 3.4 and Lemma B.1, separately from the existing hypothesis-class, density-image and fitted-process conventions. Let  $Q$  be the true joint law of  $(X, Z)$ , let  $\mathcal{G}_0$  be the likelihood-transform class already defined in Assumption 3.4, and put

$$\mathcal{A}_r = \{f \in \text{star}(\mathcal{G}_0 - \mathcal{G}_0) : \|f\|_{L_2(Q)} \leq r\}.$$

For every integer  $n \geq 1$  and every  $r > 0$ , the three processes

$$\sup_{f \in \mathcal{G}_0} |(Q_n - Q)f|, \quad \sup_{f \in \mathcal{A}_r} |(Q_n - Q)f|, \quad \sup_{f \in \mathcal{A}_r} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i f(X_i, Z_i) \right|$$

regarded as extended nonnegative quantities, are measurable under  $Q^n$  for the first two and under the product of  $Q^n$  with the independent Rademacher law for the third. The functions use the same fixed density representatives as the likelihood objective. This convention only makes these empirical processes and their expectations well defined. It supplies no concentration event, deviation inequality, probability bound, rate, countability, separability, or additional statistical restriction.

*Remark.* The displayed discrepancy inequality holds on  $\mathcal{E}_{n,t}^{\text{IR}}$ , with  $\mathbb{P}(\mathcal{E}_{n,t}^{\text{IR}}) \geq 1 - \xi$ , simultaneously for every  $\lambda \in \Lambda_n \cap [\lambda_{\text{IR}}^\circ, 2]$ .

The condition  $2t \geq \beta + 1$  marks the range in which a fixed  $t$ -step filter exploits the full source smoothness: iterating widens that range from  $\beta \leq 1$  to  $\beta \leq 2t - 1$ , and the discrepancy rule still selects  $\lambda$  without knowing  $\beta$ .

### 3.4 Adaptive Tikhonov Regularized Adversarial Estimator

The RDIV bounds include the error from estimating the operator explicitly. TRAE uses a minimax formulation to control the weak loss without fitting a conditional density. We reuse  $L_n$  and  $L$  for the TRAE losses and write  $\mathbb{E}_n$  for the average over its fitting sample of size  $n$ . For a general linear moment functional  $m(W; f)$ , its empirical loss is

$$L_n(h) = \sup_{f \in \mathcal{F}} \mathbb{E}_n [2m(W; f) - 2h(X)f(Z) - f(Z)^2],$$

where  $\mathcal{F} \subset \tilde{\mathcal{Q}}$  is the test class. The population loss is

$$L(h) = \sup_{f \in \mathcal{F}} \mathbb{E} [2m(W; f) - 2h(X)f(Z) - f(Z)^2].$$

The TRAE estimator then minimizes the penalized empirical risk:

$$\hat{h}_\lambda = \arg \min_{h \in \mathcal{H}} \{L_n(h) + \lambda \mathbb{E}_n[h(X)^2]\}.$$

We assume that the empirical loss is measurable, its minimum is attained, and the selected minimizer is measurable. The following assumption makes  $L(h)$  equal to the squared weak error.

**Assumption 3.10** (Closeness). For all  $h \in \mathcal{H}$ ,  $\mathcal{T}(h_0 - h) \in \mathcal{F}$ .

Under Assumption 3.10, the population loss simplifies to the desired squared error:

$$\begin{aligned} L(h) &= \sup_{f \in \mathcal{F}} \mathbb{E} [2(r_0(Z) - (\mathcal{T}h)(Z))f(Z) - f(Z)^2] \\ &= \sup_{f \in \mathcal{F}} \mathbb{E} [-\{f(Z) - (r_0(Z) - (\mathcal{T}h)(Z))\}^2 + \{r_0(Z) - (\mathcal{T}h)(Z)\}^2] \\ &= \|\mathcal{T}h - r_0\|^2 = \|\mathcal{T}(h - h_0)\|^2, \end{aligned}$$

where the first equality uses  $f \in \tilde{\mathcal{Q}}$  and hence  $\mathbb{E}[h(X)f(Z)] = \langle \mathcal{T}h, f \rangle$ . The next two assumptions control fluctuations of the empirical loss.

**Assumption 3.11** (Mean-square continuity). For all  $f \in \mathcal{F}$ ,  $\mathbb{E}[m(W; f)^2] \leq O(\|f\|^2)$ .

*Remark.* There exists a fixed  $C < \infty$ , independent of  $f$ ,  $n$ , and  $\lambda$ , such that  $\mathbb{E}[m(W; f)^2] \leq C\|f\|^2$  for every  $f \in \mathcal{F}$ .

**Assumption 3.12** (Boundedness and critical radius). Let

$$\mathcal{H} \cdot \mathcal{F} = \{(x, z) \mapsto h(x)f(z) : h \in \mathcal{H}, f \in \mathcal{F}\}, \quad m \circ \mathcal{F} = \{w \mapsto m(w; f) : f \in \mathcal{F}\}.$$

For a fixed  $B < \infty$ , suppose that  $|m(W; f)|$ ,  $|h(X)|$ , and  $|f(Z)|$  are a.s. bounded by  $B$ , uniformly over  $h \in \mathcal{H}$ ,  $f \in \mathcal{F}$ . Suppose in addition that, for every integer  $n \geq 1$  and every  $r > 0$ , the localized squared-loss suprema

$$\sup_{\substack{h \in \mathcal{H} \\ \|h - h_0\| \leq r}} |(\mathbb{E}_n - \mathbb{E})[h(X)^2 - h_0(X)^2]| \quad \text{and} \quad \sup_{\substack{h \in \mathcal{H} \\ \|h - h_0\| \leq r}} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i \{h(X_i)^2 - h_0(X_i)^2\} \right|$$

are measurable random variables under the fitting-sample law and its product with the independent Rademacher law, respectively. Regard the functions in  $\mathcal{H} \cdot \mathcal{F}$  and  $\mathcal{F}$  as functions of  $W$  through  $(X, Z)$  and  $Z$ , respectively, and put

$$\mathcal{U} = \text{star}(m \circ \mathcal{F}) \times \text{star}(\mathcal{H} \cdot \mathcal{F}) \times \text{star}(\mathcal{F}), \quad \|u\|_{L_2(W; \mathbb{R}^3)}^2 = \mathbb{E} \left[ \sum_{j=1}^3 u_j(W)^2 \right].$$

Define the loss with zero comparator by

$$\text{clip}_B(a) = \max\{-B, \min\{a, B\}\}, \quad \ell_B(u) = 2u_1 - 2u_2 - \text{clip}_B(u_3)^2, \quad \ell_B(0) = 0.$$

Suppose also that, for every integer  $n \geq 1$  and every  $r > 0$ , the localized loss suprema

$$\sup_{\substack{u \in \mathcal{U} \\ \|u\|_{L_2(W; \mathbb{R}^3)} \leq r}} |(\mathbb{E}_n - \mathbb{E})[\ell_B(u(W)) - \ell_B(0)]| \quad \text{and} \quad \sup_{\substack{u \in \mathcal{U} \\ \|u\|_{L_2(W; \mathbb{R}^3)} \leq r}} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i \{\ell_B(u(W_i)) - \ell_B(0)\} \right|$$

are measurable random variables under the fitting-sample law and its product with the independent Rademacher law, respectively. Whenever the primal and dual roles are exchanged, suppose that these same two localized loss suprema for the exchanged problem are measurable under its fitting-sample law and its product with the independent Rademacher law, respectively. For iterative TRAE, suppose in addition that the following fixed-population-center squared-loss processes are measurable. For every integer  $n \geq 1$ , prescribed integer  $t \geq 1$ , grid point  $\lambda \in \Lambda_n \cap [\lambda_{\text{IT}}^\circ, 2]$  for which  $\lambda_{\text{IT}}^\circ$  is defined, radius  $r > 0$ , and  $j = 1, \dots, t$ , put

$$a_{j,\lambda}(h; X) = \{h(X) - h_{j-1,\lambda}^*(X)\}^2 - \{h_{j,\lambda}^*(X) - h_{j-1,\lambda}^*(X)\}^2, \quad d_{j,\lambda}(h; X) = \{h(X) - h_{j,\lambda}^*(X)\}^2,$$

where the population iterates are defined by (3.12). For each  $v \in \{a_{j,\lambda}, d_{j,\lambda}\}$ , suppose that both

$$\sup_{\substack{h \in \mathcal{H} \\ \|h - h_{j,\lambda}^*\| \leq r}} |(\mathbb{E}_n - \mathbb{E})v(h; X)| \quad \text{and} \quad \sup_{\substack{h \in \mathcal{H} \\ \|h - h_{j,\lambda}^*\| \leq r}} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i v(h; X_i) \right|$$

are measurable under the fitting-sample law and its product with the independent Rademacher law, respectively, using the same function representatives as the empirical penalties. The zero-iterate comparison uses  $\hat{h}_{0,\lambda} = h_{0,\lambda}^* = 0$ .  $\tilde{\delta}_n(\xi) = \Omega \left( \sqrt{(\log \log(en) + \log(1/\xi))/n} \right)$  is a common upper bound on the critical radii of  $\text{star}(\mathcal{H} \cdot \mathcal{F})$ ,  $\text{star}(m \circ \mathcal{F})$ ,  $\text{star}(\mathcal{F})$ ,  $\text{star}(\mathcal{H})$ , and  $\text{star}(\mathcal{H} - \mathcal{H})$ .

*Remark.* The lower bound denoted by  $\Omega$  means that there exists a constant  $c > 0$ , independent of  $n$  and  $\xi$ , such that for every integer  $n \geq 1$  and every  $\xi \in (0, 1)$ ,  $\tilde{\delta}_n(\xi) \geq c \sqrt{(\log \log(en) + \log(1/\xi))/n}$ .

*Remark (notation).*  $\tilde{\delta}_n(\xi)$  is a prescribed common critical-radius upper bound indexed by sample size  $n$  and failure probability  $\xi \in (0, 1)$  for the five star-shaped classes specified in Assumption 3.12.

Assumption 3.12 yields the faster convergence rates [Foster and Syrgkanis, 2023]. Because its function classes and nuisance requirements differ from those of RDIV, its critical-radius scale is unrelated to the scales in (3.8); within this subsection,  $\delta_n$  always means the TRAE scale. For the finite grid  $\Lambda_n$ , set  $\eta_n = \xi/(K_n + 1)$  and choose  $\delta_n$  so that

$$\delta_n \geq \max\{\tilde{\delta}_n(\eta_n), n^{-1/2}\} = \Omega \left( \sqrt{\frac{\log \log(en) + \log((K_n + 1)/\xi)}{n}} \right), \quad \delta_n = o(1). \quad (3.17)$$

Under the confidence convention above, the concentration bounds for each fixed  $\lambda$  hold jointly with failure probability at most  $\eta_n$ . A union bound over the  $K_n + 1$  candidates then gives a common event  $\mathcal{E}_n$  with probability at least  $1 - \xi$ . Because  $K_n = O(\log n)$ , the grid contributes  $O(\log \log(en))$  inside the logarithmic confidence term.

For TRAE, the uniform loss comparison has an additive remainder  $O(\delta_n^2)$  over the regularization grid. We therefore take  $D_n = L_n$  and use the common noise level and stability bound

$$A_n^T(\lambda) = \delta_n^2, \quad V_n^T(\lambda) = \frac{\delta_n^2}{\lambda}. \quad (3.18)$$

The threshold is  $\frac{3}{2}A_n^T(\lambda)$ . The constant noise level has the order of the squared weak-error bound available when  $\beta$  is known [Bennett et al., 2023], while the stability bound measures the regularization needed to turn that weak control into a strong-error bound.

Under the calibration convention above, Corollary C.3 gives a fixed constant  $C_T > 0$  such that, on  $\mathcal{E}_n$ ,

$$L_n(h_0) \leq \delta_n^2, \quad L_n(\hat{h}_\lambda) \leq \delta_n^2 + C_T \|w_0\|^2 \lambda^{\min\{\beta+1, 2\}}, \quad \lambda \in \Lambda_n. \quad (3.19)$$

Write  $p = \min\{\beta + 1, 2\}$ . For all sufficiently large  $n$ , let  $\lambda_T^\circ$  be the first grid point no larger than

$$r_{n,\beta}^T = c_T \delta_n^{2/p},$$

where  $c_T > 0$  is a sufficiently small fixed constant. Then the regularization bias at  $\lambda_T^\circ$  is at most  $A_n^T(\lambda_T^\circ)/2$ , and

$$A_n^T(\lambda_T^\circ) = \delta_n^2, \quad V_n^T(\lambda_T^\circ) = O(\delta_n^{2(p-1)/p}). \quad (3.20)$$

*Remark (notation).* For all sufficiently large  $n$ ,  $\lambda_T^\circ$  is the deterministic oracle point  $2\rho^{k^\circ}$ , where  $k^\circ = \min\{k \in \{0, \dots, K_n\} : 2\rho^k \leq c_T \delta_n^{2/p}\}$ ,  $c_T > 0$  is a sufficiently small fixed constant, and  $\delta_n$  is the TRAE scale in (3.17).

As for RDIV, termination requires a minimum sample size, here

$$N_T(B) := \max\{2, \lceil 2B^2 \rceil\},$$

which equals 2 under the normalization  $B = 1$ .

**Proposition 3.13.** *Suppose that Assumptions 2.1, 2.3, 3.2, 3.10, 3.11, and 3.12 hold. Run Algorithm 1 with inputs  $\lambda_0 = 2$ , a fixed  $1/2 \leq \rho < 1$ , and noise level  $A_n^T(\lambda) = \delta_n^2$ , where  $\delta_n$  satisfies (3.17). For every  $n \geq N_T(B)$ , the algorithm accepts a candidate after at most  $K_n + 1 = O(\log n)$  optimization problems, with probability at least  $1 - \xi$ . The returned parameter either is  $\lambda_0$  through the initial-acceptance branch or is a noninitial grid point satisfying (3.1).*

*Additional condition.* Interpret the uniform almost-sure boundedness in Assumption 3.12 using one common null set for the whole function classes: there exists a set  $N$  with  $\mathbb{P}(W \in N) = 0$  such that, for every  $w \notin N$ , every  $h \in \mathcal{H}$  and every  $f \in \mathcal{F}$ , all three inequalities  $|h(X(w))| \leq B$ ,  $|f(Z(w))| \leq B$  and  $|m(w; f)| \leq B$  hold simultaneously. The exceptional set does not depend on  $h$  or  $f$ . Apply this convention separately to each primal or exchanged problem whenever that problem is assumed to satisfy Assumption 3.12.

*Remark.* For each fixed  $n \geq N_T(B)$ , acceptance within the stated number of optimization problems and the stated alternative for the returned parameter hold jointly on  $\mathcal{E}_n$ , with probability at least  $1 - \xi$  over the fitting sample. The empirical loss  $L_n$  is measurable, and the asserted attainment and measurable selection concern the minimum of  $h \mapsto L_n(h) + \lambda \mathbb{E}_n[h(X)^2]$  over  $h \in \mathcal{H}$ . For a fixed constant  $C_T > 0$ , both displayed inequalities hold on  $\mathcal{E}_n$ , simultaneously for all  $\lambda \in \Lambda_n$ , under the stated calibration convention.

**Theorem 3.14.** *Under the assumptions of Proposition 3.13, let  $\lambda_{\text{dp}}$  be chosen by Algorithm 1. For every  $n \geq N_T(B)$ , with probability at least  $1 - \xi$ ,*

$$\|\widehat{h}_{\lambda_{\text{dp}}} - h_0\|^2 \leq O\left(\delta_n^{2\min\{\beta, 1\}/(1+\min\{\beta, 1\})}\right), \quad \|\mathcal{T}(\widehat{h}_{\lambda_{\text{dp}}} - h_0)\|^2 \leq O\left(\delta_n^2\right),$$

where  $\delta_n$  is defined in (3.17).

*Additional condition.* Interpret the uniform almost-sure boundedness in Assumption 3.12 using one common null set for the whole function classes: there exists a set  $N$  with  $\mathbb{P}(W \in N) = 0$  such that, for every  $w \notin N$ , every  $h \in \mathcal{H}$  and every  $f \in \mathcal{F}$ , all three inequalities  $|h(X(w))| \leq B$ ,  $|f(Z(w))| \leq B$  and  $|m(w; f)| \leq B$  hold simultaneously. The exceptional set does not depend on  $h$  or  $f$ . Apply this convention separately to each primal or exchanged problem whenever that problem is assumed to satisfy Assumption 3.12.

*Remark.* With  $\beta$  and  $\xi$  fixed as  $n \rightarrow \infty$ , for each  $n \geq N_T(B)$  the displayed strong- and weak-error bounds hold jointly on a fitting-sample event of probability at least  $1 - \xi$ .

When the two methods have comparable  $\delta_n$  and  $\nu_n = O(\delta_n)$  for RDIV, the exponent in Theorem 3.14 improves on the one in Theorem 3.6: the squared strong-error rate  $O(\delta_n^{2\min\{\beta, 1\}/(1+\min\{\beta, 1\})})$  recovers the known-smoothness rate of Bennett et al. [2023] under Assumption 3.10. The proof controls the inner maximization with localized concentration inequalities and compares the empirical optimizer with  $h_0$  through the variational source inequality; this yields the discrepancy bound and the weak- and strong-error bounds that Theorem 3.1 requires, with (3.18)–(3.20).

### 3.5 Adaptive Iterative Tikhonov Regularized Adversarial Estimator

The iterative TRAE uses the same population iterates (3.12) and initializes  $\widehat{h}_{0,\lambda} = 0$ . For  $j = 1, \dots, t$ , its empirical recursion is

$$\widehat{h}_{j,\lambda} = \arg \min_{h \in \mathcal{H}} \left\{ L_n(h) + \lambda \|h - \widehat{h}_{j-1,\lambda}\|_{2,n}^2 \right\},$$

where  $L_n$  is the adversarial loss in Section 3.4. The first iterate equals the noniterative TRAE. The higher-qualification analysis uses the oracle-region condition in Assumption 3.16; we evaluate the critical-radius bound in (3.17) at  $\eta_{n,t} = \xi/\{t(K_n + 1)\}$ .

*Remark (notation).*  $\hat{h}_{j,\lambda}$  denotes a selected measurable empirical minimizer in the recursion initialized by  $\hat{h}_{0,\lambda} = 0$ , with  $j = 1, \dots, t$  and the empirical penalty norm evaluated on the fitting sample. In Proposition 3.8, Theorem 3.9, Lemma B.8, Corollary B.9, Lemma B.10: The loss in the recursion is  $L_n(h) = \mathbb{E}_{n,\text{fit}}[(Y - \hat{\mathcal{T}}h(Z))^2]$ . In Proposition 3.17, Theorem 3.18, Lemma C.2, Corollary C.6, Lemma C.7: The loss in the recursion is the adversarial loss  $L_n(h) = \sup_{f \in \mathcal{F}} \mathbb{E}_n[2m(W; f) - 2h(X)f(Z) - f(Z)^2]$ . In Lemma B.8, Lemma C.7: The empirical penalty center remains  $\hat{h}_{j-1,\lambda}$ ; the fixed-center concentration comparisons use population iterates as centers.

**Assumption 3.15** (Mean-square continuity on differences). For all  $f, f' \in \mathcal{F}$ ,

$$\mathbb{E} [m(W; f - f')^2] \leq C \|f - f'\|^2$$

for a fixed constant  $C < \infty$ .

Assumption 3.15 extends Assumption 3.11 from single test functions to their differences, which the recursive loss comparison needs; the iterative analysis uses both. It holds for  $m(W; f) = Yf(Z)$  when  $Y$  is bounded. As for iterative RDIV, each empirical minimum is attained and the selected minimizers are measurable.

For the same weak-error component, the iterative TRAE analysis offers two statistical controls [Bennett et al., 2023, Theorem 6]. Let  $b_t \geq 2$  be a deterministic constant that depends on the prescribed iteration count  $t$  and the fixed class bounds, independently of  $\beta$  and  $\|w_0\|$ , and write  $s_{n,t}^T = b_t \delta_n^2$  and  $q_t = \min\{\beta + 1, 2t\}$ . Define

$$A_{n,t}^T(\lambda) = \delta_n^2 + \min \left\{ \lambda, \frac{s_{n,t}^T}{\lambda} \right\}, \quad V_{n,t}^T(\lambda) = \frac{s_{n,t}^T}{\lambda^2}. \quad (3.21)$$

Keeping the minimum retains the sharper of the two weak-metric bounds but makes the threshold nonmonotone on the full grid. The oracle point defined below satisfies  $\lambda_{\text{IT}}^\circ \geq \sqrt{s_{n,t}^T}$ . Hence, on the oracle region  $[\lambda_{\text{IT}}^\circ, 2]$ , the noise level equals  $\delta_n^2 + s_{n,t}^T/\lambda$  and is nonincreasing, as required by Theorem 3.1. Because the oracle point is accepted, the first-crossing rule stops in this region and never uses the lower branch. For all sufficiently large  $n$ , let  $\lambda_{\text{IT}}^\circ$  be the first grid point no larger than

$$r_{n,t}^{\text{IT}} = c_{\text{IT}}(s_{n,t}^T)^{1/(q_t+1)},$$

where  $c_{\text{IT}} > 0$  is a sufficiently small fixed constant. Then  $\lambda_{\text{IT}}^\circ \geq \sqrt{s_{n,t}^T}$ , the regularization bias at  $\lambda_{\text{IT}}^\circ$  is at most  $A_{n,t}^T(\lambda_{\text{IT}}^\circ)/2$ , and

$$A_{n,t}^T(\lambda_{\text{IT}}^\circ) = O_t((s_{n,t}^T)^{q_t/(q_t+1)}), \quad V_{n,t}^T(\lambda_{\text{IT}}^\circ) = O_t((s_{n,t}^T)^{(q_t-1)/(q_t+1)}). \quad (3.22)$$

*Remark (notation).* For all sufficiently large  $n$ ,  $\lambda_{\text{IT}}^\circ$  is the deterministic oracle point  $2\rho^{k^\circ}$ , where  $k^\circ = \min\{k \in \{0, \dots, K_n\} : 2\rho^k \leq c_{\text{IT}}(s_{n,t}^\text{T})^{1/(q_t+1)}\}$  and  $c_{\text{IT}} > 0$  is a sufficiently small fixed constant.

**Assumption 3.16** (Oracle-region iterative TRAE realizability). For every  $n$  for which  $\lambda_{\text{IT}}^\circ$  is defined,  $h_{j,\lambda}^* \in \mathcal{H}$  for  $j = 1, \dots, t$  and  $\lambda \in \Lambda_n \cap [\lambda_{\text{IT}}^\circ, 2]$ .

On this grid segment, the discrepancy bound in Corollary C.6 gives

$$L_n(\hat{h}_{t,\lambda}) \leq A_{n,t}^\text{T}(\lambda) + C_{\text{IT}}\|w_0\|^2\lambda^{q_t}, \quad \lambda \in \Lambda_n \cap [\lambda_{\text{IT}}^\circ, 2]. \quad (3.23)$$

Every grid point in this segment satisfies  $\lambda \geq \sqrt{s_{n,t}^\text{T}}$ . Theorem 3.1 therefore applies once the TRAE loss comparison gives the weak- and strong-error bounds at every grid point on that segment.

The corresponding minimum sample size for the iterative rule is

$$N_{\text{IT}}(B) := \max\{2, \lceil 8B^2 \rceil\},$$

which equals 8 under the normalization  $B = 1$ .

**Proposition 3.17.** *Suppose the assumptions of Proposition 3.13 hold. For every  $n \geq N_{\text{IT}}(B)$ , Algorithm 1, applied to  $\hat{h}_{t,\lambda}$  with the noise level (3.21), accepts a candidate after at most  $t(K_n + 1) = O(t \log n)$  minimax problems, with probability at least  $1 - \xi$ . The returned parameter is either the initial point or a noninitial point satisfying (3.1) with the iterative loss and threshold.*

*Additional condition. Interpret the uniform almost-sure boundedness in Assumption 3.12 using one common null set for the whole function classes: there exists a set  $N$  with  $\mathbb{P}(W \in N) = 0$  such that, for every  $w \notin N$ , every  $h \in \mathcal{H}$  and every  $f \in \mathcal{F}$ , all three inequalities  $|h(X(w))| \leq B$ ,  $|f(Z(w))| \leq B$  and  $|m(w; f)| \leq B$  hold simultaneously. The exceptional set does not depend on  $h$  or  $f$ . Apply this convention separately to each primal or exchanged problem whenever that problem is assumed to satisfy Assumption 3.12.*

*Remark.* For each fixed  $n \geq N_{\text{IT}}(B)$  and prescribed  $t$ , acceptance within the stated number of minimax problems and the stated alternative for the returned parameter hold jointly on  $\mathcal{E}_{n,t}^{\text{IT}}$ , with probability at least  $1 - \xi$  over the fitting sample.

**Theorem 3.18** (Adaptive iterative TRAE). *Under the assumptions of Proposition 3.17 and Assumptions 3.15 and 3.16, let  $\lambda_{\text{dp},t}^\text{T}$  denote the selected parameter, with  $s_{n,t}^\text{T}$  and  $q_t$  defined above. For every fixed  $t$  and every  $n \geq N_{\text{IT}}(B)$ , with probability at least  $1 - \xi$ ,*

$$\|\mathcal{T}(\hat{h}_{t,\lambda_{\text{dp},t}^\text{T}} - h_0)\|^2 \leq O_t((s_{n,t}^\text{T})^{q_t/(q_t+1)}), \quad \|\hat{h}_{t,\lambda_{\text{dp},t}^\text{T}} - h_0\|^2 \leq O_t((s_{n,t}^\text{T})^{(q_t-1)/(q_t+1)}).$$

Because  $t$  is fixed,  $s_{n,t}^T = b_t \delta_n^2 = O_t(\delta_n^2)$ , so  $b_t$  affects only the constants hidden by  $O_t(\cdot)$ . If  $2t \geq \beta + 1$ , these bounds become

$$\|\widehat{h}_{t,\lambda_{\text{dp},t}^T} - h_0\|^2 = O_t(\delta_n^{2\beta/(\beta+2)}), \quad \|\mathcal{T}(\widehat{h}_{t,\lambda_{\text{dp},t}^T} - h_0)\|^2 = O_t(\delta_n^{2(\beta+1)/(\beta+2)}).$$

*Additional condition. Interpret the uniform almost-sure boundedness in Assumption 3.12 using one common null set for the whole function classes: there exists a set  $N$  with  $\mathbb{P}(W \in N) = 0$  such that, for every  $w \notin N$ , every  $h \in \mathcal{H}$  and every  $f \in \mathcal{F}$ , all three inequalities  $|h(X(w))| \leq B$ ,  $|f(Z(w))| \leq B$  and  $|m(w; f)| \leq B$  hold simultaneously. The exceptional set does not depend on  $h$  or  $f$ . Apply this convention separately to each primal or exchanged problem whenever that problem is assumed to satisfy Assumption 3.12.*

*Remark.* The constants hidden by  $O_t(\cdot)$  may depend on the prescribed  $t$ , fixed population quantities including  $\beta$  and  $\|w_0\|$ , and the fixed grid ratio  $\rho$ , but are independent of  $n$  and the candidate  $\lambda$ . Under the assumptions of Theorem 3.18, the displayed inequality holds simultaneously for all  $\lambda \in \Lambda_n \cap [\lambda_{\text{IT}}^\circ, 2]$  on  $\mathcal{E}_{n,t}^{\text{IT}}$ , where  $\mathbb{P}(\mathcal{E}_{n,t}^{\text{IT}}) \geq 1 - \xi$ .

As for RDIV, a fixed  $t$  widens the range of source smoothness that the TRAE rate exploits to  $\beta \leq 2t - 1$ . In the low-smoothness regime, the one-step result in Theorem 3.14 remains the sharper of the two.

### 3.6 Application to Doubly Robust Estimation

Estimating  $\theta_0$  itself calls for the dual nuisance function as well, and the method of Section 3.4 applies to the dual problem unchanged:

$$\theta_0 = \mathbb{E}[m(W; q_0)], \quad \mathcal{T}^* q_0 = a_0,$$

where  $a_0$  is the Riesz representer of the linear functional  $h \mapsto \mathbb{E}[\tilde{m}(W; h)]$ . For a test class  $\mathcal{S} \subseteq \tilde{\mathcal{H}}$ , define the empirical loss

$$L_n^{\text{dual}}(q) = \sup_{s \in \mathcal{S}} \mathbb{E}_n [2\tilde{m}(W; s) - 2q(Z)s(X) - s(X)^2]$$

and the adversarial Tikhonov regularization estimator,

$$\widehat{q} = \arg \min_{q \in \mathcal{Q}} \{L_n^{\text{dual}}(q) + \lambda_{\text{dp}}^{\text{dual}} \mathbb{E}_n [q(Z)^2]\},$$

where we select  $\lambda_{\text{dp}}^{\text{dual}}$  with the discrepancy principle. Given estimates of  $h_0$  and  $q_0$  from the adaptive Tikhonov method, let  $\mathbb{E}_{n,\text{eval}}$  denote an empirical average over an evaluation sample

of size  $n$  that is independent of the data used to train and tune both nuisance estimators. We define the DR estimator of  $\theta_0$  as

$$\widehat{\theta}(h, q) = \mathbb{E}_{n, \text{eval}} [\widetilde{m}(W; h) + m(W; q) - h(X)q(Z)].$$

The next corollary establishes its asymptotic normality.

*Remark (notation).*  $\widehat{q}$  denotes the dual empirical adversarial Tikhonov fit at the discrepancy-selected dual penalty. In Corollary 3.19: The dual discrepancy rule is run at confidence level  $\xi_n/2$ , using the common deterministic sequence  $\delta_n$  that satisfies (3.17) for both nuisance problems. Take jointly measurable fitted evaluations of  $\widehat{q}$  representing the fitted function for almost every nuisance-training and tuning realization and agreeing almost surely with  $\widehat{q}(Z(W_i))$  simultaneously for  $i = 1, \dots, n$  in the evaluation sample.

**Corollary 3.19.** *Suppose that the assumptions of Theorem 3.14 hold for both  $h_0$  and  $q_0$  with (potentially different) source conditions  $\beta_h, \beta_q > 0$ . Choose a deterministic sequence  $\xi_n \downarrow 0$ , and run the primal and dual tuning rules at confidence levels  $\xi_n/2$ . Let a common deterministic sequence  $\delta_n$  satisfy (3.17) for both problems at these confidence levels. Assume that, for some  $\gamma_h, \gamma_q > 0$ , the score maps satisfy*

$$\begin{aligned} \mathbb{E} [(\widetilde{m}(W; h_1) - \widetilde{m}(W; h_2))^2] &\leq C \|h_1 - h_2\|^{\gamma_h}, \\ \mathbb{E} [(m(W; q_1) - m(W; q_2))^2] &\leq C \|q_1 - q_2\|^{\gamma_q} \end{aligned}$$

for all  $h_1, h_2 \in \mathcal{H}$  and  $q_1, q_2 \in \mathcal{Q}$ . Let  $\beta_m = \max\{\beta_h, \beta_q\}$  and assume that  $\delta_n$  satisfies

$$\delta_n = o(n^{-\alpha_{\text{dr}}}), \quad \alpha_{\text{dr}} := \frac{1 + \min\{\beta_m, 1\}}{2 + 4 \min\{\beta_m, 1\}}.$$

Let  $\widehat{h}, \widehat{q}$  be the adaptive adversarial Tikhonov estimators, trained and tuned independently of the evaluation sample used in  $\widehat{\theta}(\widehat{h}, \widehat{q})$ . For each  $n$ , write  $\mathcal{G}_n$  for the full nuisance-training and tuning sigma-field. Take jointly measurable versions of the fitted nuisance evaluations, representing the fitted functions for almost every training realization and agreeing almost surely with their evaluations at every observation in the finite evaluation sample. Write

$$S(w; h, q) = \widetilde{m}(w; h) + m(w; q) - h(X(w))q(Z(w)), \quad d_n(w) = S(w; \widehat{h}, \widehat{q}) - S(w; h_0, q_0).$$

Suppose that, for each  $n$ ,  $d_n$  admits a version jointly measurable in  $\mathcal{G}_n$  and a fresh observation, agreeing with  $d_n$  for almost every training realization and almost every fresh

observation, and agreeing almost surely with  $d_n(W_i)$  simultaneously for  $i = 1, \dots, n$  in the evaluation sample. Finally, assume that  $0 < \sigma_0^2 := \mathbb{E}[\psi_0(W)^2] < \infty$ . Then,

$$\sqrt{n}(\hat{\theta}(\hat{h}, \hat{q}) - \theta_0) = \frac{1}{\sqrt{n}} \sum_{i=1}^n \psi_0(W_i) + o_p(1),$$

where  $\psi_0(W) = \tilde{m}(W; h_0) + m(W; q_0) - q_0(Z)h_0(X) - \theta_0$ . Consequently,  $\sqrt{n}(\hat{\theta}(\hat{h}, \hat{q}) - \theta_0) \rightsquigarrow N(0, \sigma_0^2)$ .

*Additional condition. Interpret the uniform almost-sure boundedness in Assumption 3.12 using one common null set for the whole function classes: there exists a set  $N$  with  $\mathbb{P}(W \in N) = 0$  such that, for every  $w \notin N$ , every  $h \in \mathcal{H}$  and every  $f \in \mathcal{F}$ , all three inequalities  $|h(X(w))| \leq B$ ,  $|f(Z(w))| \leq B$  and  $|m(w; f)| \leq B$  hold simultaneously. The exceptional set does not depend on  $h$  or  $f$ . Apply this convention separately to each primal or exchanged problem whenever that problem is assumed to satisfy Assumption 3.12.*

*Remark (notation).*  $\hat{\theta}$  denotes an empirical estimate of  $\theta_0$ , with its nuisance arguments displayed when the doubly robust score functional is intended. In Corollary 3.19:  $\hat{\theta}(h, q) = \mathbb{E}_{n, \text{eval}}[\tilde{m}(W; h) + m(W; q) - h(X)q(Z)]$ , and the estimator in the conclusion is its evaluation at  $(h, q) = (\hat{h}, \hat{q})$ ; here  $q \in \mathcal{Q}$  is a nuisance-function argument.  $\beta_h > 0$  is a source exponent satisfying  $h_0 = (\mathcal{T}^* \mathcal{T})^{\beta_h/2} w_h$  for some  $w_h \in \tilde{\mathcal{H}}$ ; it need not be the largest exponent for which such a representation exists.  $\beta_q > 0$  is a source exponent satisfying  $q_0 = (\mathcal{T} \mathcal{T}^*)^{\beta_q/2} w_q$  for some  $w_q \in \tilde{\mathcal{Q}}$ ; it need not be the largest exponent for which such a representation exists.

Taking  $\xi_n = 1/\log(en)$  drives the simultaneous probability of the nuisance bounds to one at a cost of only  $O(\log \log(en))$  in the logarithmic confidence term. The sufficient condition for asymptotic normality depends on  $\beta_m$  alone, which automatically selects the better source-condition parameter from the primal and dual problems. For  $\beta_m \geq 1$  the requirement becomes  $\delta_n = o(n^{-1/3})$ , which some nonparametric classes meet; for every fixed  $\beta_m > 0$  the required exponent stays below  $1/2$  and approaches  $1/2$  as  $\beta_m \downarrow 0$ .

## 4 Experiments

The guarantees in Section 3 compare the selected estimator with a regularization level that depends on the unknown source parameter. A simulation can test this comparison only when the best regularization level changes from one problem to another. We therefore vary the instrument strength, the noise level, the sample size, and the shape of the structural function, and we compare the discrepancy rule with the best fixed penalty in each setting. We use kernel hypothesis and critic classes. For these classes each candidate  $\hat{h}_\lambda$  is the exact penalized

solution and has a closed form, so the comparison separates the choice of  $\lambda$  from optimization error. With neural networks, each candidate comes from a separate training run, and in our runs the empirical TRAE loss increased in about 30% of the steps from one grid point to the next smaller one. An exact penalized solution cannot behave this way. We leave the joint treatment of optimization error and tuning to future work.

## 4.1 Design

**Linear Gaussian design.** Let  $Z$ ,  $V$ , and  $\varepsilon$  be independent standard normal variables, and set

$$X = \rho Z + \sqrt{1 - \rho^2} V, \quad Y = h_s(X) + aV + \sigma\varepsilon,$$

with  $a = 0.5$ . The structural function is a Gaussian bump of width  $s$ , centered and scaled to unit variance under  $X \sim N(0, 1)$ :

$$h_s(x) = \frac{\exp\{-x^2/(2s^2)\} - m_s}{d_s}, \quad m_s = \frac{s}{\sqrt{s^2 + 1}}, \quad d_s^2 = \frac{s}{\sqrt{s^2 + 2}} - \frac{s^2}{s^2 + 1}.$$

The disturbance  $aV + \sigma\varepsilon$  is correlated with  $X$  and has mean zero given  $Z$ , so  $h_s$  solves  $\mathbb{E}[Y - h(X) \mid Z] = 0$ . In the basis of normalized Hermite polynomials, the conditional expectation operator is diagonal with singular values  $\rho^k$ ,  $k \geq 0$ . The operator is therefore injective, and a smaller  $\rho$  gives a more ill-posed problem. Section D shows that  $h_s$  satisfies Assumption 2.3 if and only if

$$\beta < \beta_{\max}(s, \rho) := \frac{\log(1 + s^2)}{-2 \log \rho}.$$

Over the regimes below,  $\beta_{\max}$  ranges from 0.46 at  $(s, \rho) = (0.7, 0.65)$  to 17.2 at  $(s, \rho) = (2.2, 0.95)$ , so the source parameter lies on both sides of one.

**Nonlinear design.** The second design adds a binary treatment, a nonlinear first stage, and heteroskedastic noise. With  $Z$ ,  $V$ , and  $\varepsilon$  as above, let  $A \mid Z \sim \text{Bernoulli}\{\text{expit}(Z)\}$  and

$$X = \rho\{Z + 0.6 \sin(1.5Z)\} + 0.5(A - \tfrac{1}{2}) + \sqrt{1 - \rho^2} V, \quad Y = g(X) + A\tau(X) + aV + \sigma\kappa(Z)\varepsilon,$$

where  $g(x) = \tanh(x) + 0.8 \exp\{-(x - 0.5)^2/0.5\}$ ,  $\tau(x) = 1 + 0.5 \sin x$ , and  $\kappa(z)^2 = 0.5 + \text{expit}(z)$ . The hypothesis takes  $(A, X)$  and the instrument is  $(A, Z)$ , so that  $h_0(a, x) = g(x) + a\tau(x)$ . Given  $A$ , the first stage is a Gaussian location family whose location is increasing in  $z$  with range  $\mathbb{R}$ , so the operator is again injective.

**Regimes.** We cross  $\rho \in \{0.65, 0.8, 0.95\}$ ,  $\sigma \in \{0.1, 1, 3\}$ , and  $n \in \{500, 1000, 2000, 4000, 8000\}$ . In the linear design, the share of  $\text{Var}(Y)$  due to noise grows from about 20% at  $\sigma = 0.1$  to about 90% at  $\sigma = 3$ . For TRAE, the linear design also crosses  $s \in \{0.7, 2.2\}$ , which gives 90 regimes; the comparison across the two designs uses  $s = 0.7$ , so each design has 45 regimes. Each regime uses 50 data sets for TRAE and the first 20 of them for RDIV. The data sets are nested in  $n$  and share random numbers across the other factors, so all comparisons between methods and between regimes are paired.

## 4.2 Estimators and Noise Levels

Both estimators use the kernel  $k(u, u') = 1 + \exp(-\|u - u'\|^2/2)$  for the hypothesis and, in TRAE, for the critic. The constant term allows functions that do not vanish at infinity, such as  $h_s$  and its conditional mean. Outcomes are standardized on the fitting block. Each hypothesis is restricted to the eigenvectors of its Gram matrix whose eigenvalues exceed  $\sqrt{N\epsilon_{\text{mach}}}$  times the largest one, where  $N$  is the size of the fitting block.

TRAE fits on the first half of the sample. Its critic is the RKHS of the same kernel with penalty  $\eta\|f\|^2$ ,  $\eta = 10^{-3}$ , which we profile out in closed form (Section D.1). RDIV splits the sample into four blocks of size  $n/4$ , used for the preliminary fits, the fit, calibration, and evaluation. On the preliminary block, a mixture density network estimates the conditional density of the continuous regressor given the instrument, and kernel ridge regression gives  $\hat{r}$ . Fifty draws per observation from the fitted density define  $\hat{\mathcal{T}}$ . The second stage is then a kernel ridge problem, and  $D_n$  is the projected residual (3.7) on the calibration block.

The search uses the grid of Algorithm 1 with  $\lambda_0 = 2$  and search factor 1/2, so that  $\lambda_k = 2^{1-k}$  and the grid has between 16 and 26 points (Table 4). On every data set we fit all grid points and  $\lambda = 0$  once. The discrepancy rule and every fixed penalty are evaluated on these same fits, so two methods that select the same  $\lambda$  return the same function.

**Noise level.** The theory takes  $A_n$  from prescribed bounds whose constants are unknown. We compare two choices. The *fixed-scale rule* uses (3.17) with constant one,

$$\delta_N^2 = \frac{\log \log(eN) + \log\{(K_N + 1)/\xi\}}{N}, \quad \xi = 0.1,$$

with  $A_n = \delta_N^2$  for TRAE and  $A_n(\lambda) = \delta_N^2/\lambda + \nu_N^2$  for RDIV, where  $\nu_N$  has the same form on the calibration block. The *noise-floor rule* sets  $A_n(\lambda)$  to an estimate of the value the statistic would take if the residual of  $\hat{h}_\lambda$  carried no information about  $Z$ . For TRAE, let  $\mathbf{e}$  be the residual of  $\hat{h}_\lambda$  on the fitting block. Multiplying the part the critic does not explain,

$(I - P_\eta)\mathbf{e}$ , by independent random signs removes any dependence on  $Z$  and keeps the scale of the residual. The expected discrepancy of the result is

$$A_n(\lambda) = \frac{1}{N} \sum_{i=1}^N [P_\eta]_{ii} [(I - P_\eta)\mathbf{e}]_i^2.$$

For RDIV,  $A_n(\lambda)$  adds the variance of  $\widehat{\mathcal{T}}\widehat{h}_\lambda$  on the calibration block, the variance of  $\widehat{r}$ , and an estimate of  $\|(\mathcal{T} - \widehat{\mathcal{T}})\widehat{h}_\lambda\|^2$ ; Section D gives the formulas. Both noise-floor levels are bounded below by  $1/N$ . We fixed the noise-floor rule as the primary rule on pilot data sets, before drawing the main ones. It is a data-driven choice of  $A_n$ , and the theorems, which take  $A_n$  as prescribed, do not apply to it directly.

**Comparison.** For a method and a regime, let  $R$  be the mean squared error of the method over the data sets of the regime, and let  $R^*$  be the smallest such value among the fixed grid penalties in this regime. We report the *risk ratio*  $R/R^*$ . A ratio of one means that the method does as well as the best penalty for this regime, chosen with knowledge of the truth. In practice a fixed penalty is set once and used for every problem. We therefore also report the *best single fixed penalty*, the grid value with the smallest geometric mean ratio over the regimes. It is chosen after all results are known, which favors the fixed approach.

We compute the strong error  $\|\widehat{h} - h_0\|^2$  and the weak error  $\|\mathcal{T}(\widehat{h} - h_0)\|^2$  by quadrature, using the analytic structural function and operator. The target error is  $(\widehat{\theta} - \theta_0)^2$ , where  $\theta_0 = \mathbb{E}[h_0(1, X) - h_0(0, X)]$  in the nonlinear design and  $\theta_0 = \mathbb{E}[h_0(X + 1) - h_0(X)]$  in the linear design. The estimate  $\widehat{\theta}$  is the plug-in average over held-out observations.

### 4.3 Results

**TRAЕ in the linear design.** Table 1 summarizes the 90 regimes. Across them, the best fixed penalty for the strong error ranges from 2 to  $2^{-9}$ , and Figure 1 shows that the selected penalty moves with it; the correlation between their logarithms is 0.90. No single penalty works in every regime. Each fixed grid value is more than 8 times worse than the best penalty in some regime, while the discrepancy rule is never more than 1.83 times worse. On the same data sets, the rule has lower strong error than the best single fixed penalty  $\lambda = 2^{-4}$  in 58 of the 90 regimes. The paired difference exceeds two standard errors in 43 regimes in favor of the rule and in 12 against it. The rule also has lower risk than  $\lambda = 0$  in all 90 regimes. Its largest ratios occur at  $\rho = 0.65$  and  $n = 500$ , where the best penalty sits at the top of the grid and the rule stops at a smaller one.

Table 1: TRAE with the kernel class in the linear Gaussian design, 90 regimes. Each entry summarizes the risk ratio over regimes: its geometric mean, median, and largest value, and the share of regimes with ratio at most 1.5. The best single fixed penalty is chosen after seeing the results, separately for each error.

| Error  | Method                                | Geometric mean | Median | Worst | Within 1.5 |
|--------|---------------------------------------|----------------|--------|-------|------------|
| Strong | Noise-floor rule                      | 1.15           | 1.13   | 1.83  | 96%        |
|        | Fixed-scale rule                      | 1.24           | 1.21   | 2.52  | 92%        |
|        | Best single fixed, $\lambda = 2^{-4}$ | 1.56           | 1.24   | 8.17  | 61%        |
|        | $\lambda = 0$                         | 172            | 173    | 7908  | 0%         |
| Weak   | Noise-floor rule                      | 1.25           | 1.23   | 1.98  | 90%        |
|        | Fixed-scale rule                      | 1.59           | 1.51   | 3.56  | 50%        |
|        | Best single fixed, $\lambda = 2^{-5}$ | 1.56           | 1.27   | 7.75  | 64%        |
|        | $\lambda = 0$                         | 105            | 90     | 5179  | 0%         |




Figure 1: Selected penalty against the best fixed penalty for the strong error, one point per regime. The selected penalty is the geometric mean over the data sets of a regime. TRAE uses 90 linear and 45 nonlinear regimes; RDIV uses 45 regimes of each design. The lowest RDIV points are regimes in which some searches reached the end of the grid. The dashed line is the diagonal.

The fixed-scale rule is less accurate, with geometric mean ratios 1.24 for the strong error and 1.59 for the weak error. Its constant also matters: with constant 0.5 or 2 in place of one, the worst strong-error ratio becomes 14.8 or 8.2. The noise-floor rule has no such constant. At the threshold where the noise-floor rule stopped, the analytic  $h_0$  would also have passed in 86% of the data sets. The search never reached the end of the grid and accepted the initial value  $\lambda_0 = 2$  in 3% of the data sets. Figure 2 shows the strong error against  $n$  for  $s = 0.7$ . The adaptive curve stays close to the best fixed penalty of each regime, while each fixed penalty is accurate only for some  $(\rho, \sigma, n)$ :  $\lambda = 2^{-1}$  at small  $n$  and high noise, and  $\lambda = 2^{-9}$  at large  $n$  and low noise.

**Two designs, TRAE and RDIV.** Table 2 adds the nonlinear design, RDIV, and the target error. TRAE improves on the best single fixed penalty in both designs and for all three errors. Its worst ratio is at most 2.19, compared with 4.72 to 13.43 for the best single fixed penalty. Its smallest margin is for the weak error in the nonlinear design, 1.31 against 1.41.

RDIV improves on the best single fixed penalty in the linear design for all three errors. In the nonlinear design it improves the target error, 1.22 against 1.40, is close on the strong error, 1.30 against 1.26, and is worse on the weak error, 1.64 against 1.23. Its search is also less stable at both ends of the sample-size range. In the linear design it accepted  $\lambda_0$  in 53% of the data sets at  $n = 500$ , and it reached the end of the grid in 7% of the data sets overall and in 17% to 20% at  $n = 8000$ , mostly at the lowest noise level. Two features of RDIV may explain the difference. It fits and calibrates on different blocks with an estimated operator, so its statistic need not decrease along the grid. It increased by more than 5% in 18% of the grid steps, while the TRAE statistic never did. Its noise floor also depends on the estimated operator error.

**Scope.** We chose the regimes so that the best penalty changes across them, since this is the setting in which a tuning rule is needed. Within a single regime, the best fixed penalty for that regime is about as accurate as the rule, and usually more accurate. The Gaussian designs violate the boundedness conditions of our theorems, and the noise-floor rule uses a data-driven  $A_n$ . The experiments therefore show how the rule behaves with exact penalized solutions under conditions that the theory does not cover.




Figure 2: Strong error of TRAE against the sample size in the linear design with  $s = 0.7$ , for three fixed penalties and the noise-floor rule. The dashed curve is the best fixed penalty of each regime. Bars show  $\pm 1$  Monte Carlo standard error over 50 data sets.

Table 2: Geometric mean risk ratio over 45 regimes per design, with the largest ratio in parentheses. The linear design uses  $s = 0.7$ . “Fixed” is the best single fixed penalty, chosen after seeing the results, separately for each estimator, design, and error. The smaller geometric mean in each pair is in bold.

| Estimator | Design    | Strong error       |                    | Weak error         |                    | Target error       |              |
|-----------|-----------|--------------------|--------------------|--------------------|--------------------|--------------------|--------------|
|           |           | Adaptive           | Fixed              | Adaptive           | Fixed              | Adaptive           | Fixed        |
| TRAЕ      | Linear    | <b>1.16</b> (1.73) | 1.60 (7.78)        | <b>1.24</b> (1.90) | 1.59 (7.75)        | <b>1.23</b> (2.19) | 1.66 (13.43) |
|           | Nonlinear | <b>1.07</b> (1.30) | 1.48 (8.36)        | <b>1.31</b> (1.87) | 1.41 (4.72)        | <b>1.12</b> (1.80) | 1.58 (9.58)  |
| RDIV      | Linear    | <b>1.34</b> (3.69) | 1.43 (8.18)        | <b>1.38</b> (3.53) | 1.58 (12.93)       | <b>1.39</b> (7.81) | 1.56 (8.00)  |
|           | Nonlinear | 1.30 (2.61)        | <b>1.26</b> (3.57) | 1.64 (2.62)        | <b>1.23</b> (3.36) | <b>1.22</b> (2.60) | 1.40 (7.90)  |

## 5 Conclusion

We develop a discrepancy rule for choosing the regularization level in conditional moment models when the source smoothness is unknown. The rule sets its threshold from prescribed bounds on statistical error and applies to both RDIV and TRAЕ. Its oracle guarantee reduces the analysis of any such rule to two ingredients: a discrepancy bound and weak- and strong-error bounds, both holding at every grid point. This general principle applies to the noniterative and iterative versions of RDIV and TRAЕ, and it yields adaptive rates that match known-smoothness comparisons within the smoothness ranges stated in the theorems when the underlying statistical errors are comparable. The threshold still requires valid statistical-error bounds. In simulations with kernel classes, where every candidate is an exact penalized solution, the rule followed the best fixed penalty across settings in which that penalty changed by three orders of magnitude, and it stayed close to the best penalty of each setting. These results use a data-driven noise level. A calibration of the noise level that the theory covers, and the role of optimization error for neural network classes, remain open.

## References

- Chunrong Ai and Xiaohong Chen. Efficient estimation of models with conditional moment restrictions containing unknown functions. *Econometrica*, 71(6):1795–1843, 2003. doi: 10.1111/1468-0262.00470.
- Andrew Bennett and Nathan Kallus. The variational method of moments. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 85(3):810–841, 2023.

- Andrew Bennett, Nathan Kallus, Xiaojie Mao, Whitney Newey, Vasilis Syrgkanis, and Masatoshi Uehara. Source condition double robust inference on functionals of inverse problems. *arXiv preprint arXiv:2307.13793*, 2023.
- Andrew Bennett, Nathan Kallus, Xiaojie Mao, Whitney K. Newey, Vasilis Syrgkanis, and Masatoshi Uehara. Inference on strongly identified functionals of weakly identified functions. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 88(3):998–1028, 2026. doi: 10.1093/jrssb/qkaf075.
- Gilles Blanchard, Marc Hoffmann, and Markus Reiß. Optimal adaptation for early stopping in statistical inverse problems. *SIAM/ASA Journal on Uncertainty Quantification*, 6(3):1043–1075, 2018. doi: 10.1137/17M1154096.
- Christoph Breunig and Peter Haan. Nonparametric regression with selectively missing covariates. *Journal of Econometrics*, 223(1):28–52, 2021.
- Christoph Breunig and Jan Johannes. Adaptive estimation of functionals in nonparametric instrumental regression. *Econometric Theory*, 32(3):612–654, 2016. doi: 10.1017/S0266466614000966.
- Marine Carrasco, Jean-Pierre Florens, and Eric Renault. *Linear inverse problems in structural econometrics: Estimation based on spectral decomposition and regularization*, pages 5633–5751. Elsevier, 2007. ISBN 978-0-444-53200-8. doi: 10.1016/s1573-4412(07)06077-1.
- Alain Celisse and Martin Wahl. Analyzing the discrepancy principle for kernelized spectral filter learning algorithms. *Journal of Machine Learning Research*, 22(76):1–59, 2021. URL <https://www.jmlr.org/papers/v22/20-358.html>.
- Xiaohong Chen and Demian Pouzo. Estimation of Nonparametric Conditional Moment Models With Possibly Nonsmooth Generalized Residuals. *Econometrica*, 80(1):277–321, 2012. ISSN 1468-0262. doi: 10.3982/ECTA7888.
- Xiaohong Chen and Markus Reiss. On rate optimality for ill-posed inverse problems in econometrics. *Econometric Theory*, 27(3):497–521, June 2011. ISSN 1469-4360, 0266-4666. doi: 10.1017/S0266466610000381.
- Xiaohong Chen, Timothy Christensen, and Sid Kankanala. Adaptive estimation and uniform confidence bands for nonparametric structural functions and elasticities. *The Review of Economic Studies*, 92(1):162–196, 2025. doi: 10.1093/restud/rdae025.

- Serge Darolles, Yanqin Fan, Jean-Pierre Florens, and Eric Renault. Nonparametric instrumental regression. *Econometrica*, 79(5):1541–1565, 2011. doi: 10.3982/ECTA6539.
- Nishanth Dikkala, Greg Lewis, Lester Mackey, and Vasilis Syrgkanis. Minimax estimation of conditional moment models. *Advances in Neural Information Processing Systems*, 33: 12248–12262, 2020.
- Xavier d’Haultfoeulle. A new instrumental method for dealing with endogenous selection. *Journal of Econometrics*, 154(1):1–15, 2010.
- Heinz Werner Engl, Martin Hanke, and Andreas Neubauer. *Regularization of inverse problems*, volume 375. Springer Science & Business Media, 1996.
- Jean-Pierre Florens, Jan Johannes, and Sébastien van Bellegem. Identification and estimation by penalization in nonparametric instrumental regression. *Econometric Theory*, 27(3): 472–496, June 2011. ISSN 1469-4360, 0266-4666. doi: 10.1017/S026646661000037X.
- Dylan J. Foster and Vasilis Syrgkanis. Orthogonal statistical learning. *The Annals of Statistics*, 51(3):879–908, 2023. doi: 10.1214/23-AOS2258.
- Peter Hall and Joel L. Horowitz. Nonparametric methods for inference in the presence of instrumental variables. *The Annals of Statistics*, 33(6):2904–2929, 2005. doi: 10.1214/009053605000000714.
- Per Christian Hansen. Analysis of discrete ill-posed problems by means of the L-curve. *SIAM Review*, 34(4):561–580, 1992. doi: 10.1137/1034115.
- Bastian Harrach, Tim Jahn, and Roland Potthast. Beyond the Bakushinkii veto: regularising linear inverse problems without knowing the noise distribution. *Numerische Mathematik*, 145:581–603, 2020. doi: 10.1007/s00211-020-01122-2.
- Jason Hartford, Greg Lewis, Kevin Leyton-Brown, and Matt Taddy. Deep iv: A flexible approach for counterfactual prediction. In *International Conference on Machine Learning*, pages 1414–1423. PMLR, 2017.
- Joel L. Horowitz. Adaptive nonparametric instrumental variables estimation: Empirical choice of the regularization parameter. *Journal of Econometrics*, 180(2):158–173, 2014a. doi: 10.1016/j.jeconom.2014.03.006.
- Joel L Horowitz. Ill-posed inverse problems in economics. *Annu. Rev. Econ.*, 6(1):21–51, 2014b.

- Jan Johannes and Maik Schwarz. Partially adaptive nonparametric instrumental regression by model selection. *Journal of the Indian Statistical Association*, 49(2):149–175, 2011. URL <https://arxiv.org/abs/1003.3128>.
- Greg Lewis and Vasilis Syrgkanis. Adversarial generalized method of moments. *arXiv preprint arXiv:1803.07164*, 2018.
- Wei Li, Wang Miao, and Eric J. Tchetgen Tchetgen. Non-parametric inference about mean functionals of non-ignorable non-response data without identifying the joint distribution. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 85(3):913–935, 2023. doi: 10.1093/jrsssb/qkad047.
- Zihao Li, Hui Lan, Vasilis Syrgkanis, Mengdi Wang, and Masatoshi Uehara. Regularized deepiv with model selection. *arXiv preprint arXiv:2403.04236*, 2024. URL <https://arxiv.org/abs/2403.04236>.
- Luofeng Liao, You-Lin Chen, Zhuoran Yang, Bo Dai, Mladen Kolar, and Zhaoran Wang. Provably efficient neural estimation of structural equation models: An adversarial approach. *Advances in Neural Information Processing Systems*, 33:8947–8958, 2020.
- Wang Miao, Lan Liu, Yilin Li, Eric J. Tchetgen Tchetgen, and Zhi Geng. Identification and semiparametric efficiency theory of nonignorable missing data with a shadow variable. *ACM/IMS Journal of Data Science*, 1(2):1–23, 2024a. doi: 10.1145/3592389. Article 5.
- Wang Miao, Xu Shi, Yilin Li, and Eric J. Tchetgen Tchetgen. A confounding bridge approach for double negative control inference on causal effects. *Statistical Theory and Related Fields*, 8(4):262–273, 2024b. doi: 10.1080/24754269.2024.2390748.
- V. A. Morozov. On the solution of functional equations by the method of regularization. *Doklady Akademii Nauk SSSR*, 167(3):510–512, 1966. URL <https://www.mathnet.ru/eng/dan32161>.
- Vladimir A. Morozov. *Methods for Solving Incorrectly Posed Problems*. Springer, New York, 1984. doi: 10.1007/978-1-4612-5280-1.
- Krikamol Muandet, Arash Mehrjou, Si Kai Lee, and Anant Raj. Dual instrumental variable regression. *Advances in Neural Information Processing Systems*, 33:2710–2721, 2020.
- M. Thamban Nair, E. Scheck, and Ulrich Tautenhahn. Morozov’s discrepancy principle under general source conditions. *Zeitschrift für Analysis und ihre Anwendungen*, 22(1):199–214, 2003. doi: 10.4171/ZAA/1140.

- Whitney K. Newey. Nonparametric Instrumental Variables Estimation. *The American Economic Review*, 103(3):550–556, 2013. ISSN 0002-8282.
- Whitney K. Newey and James L. Powell. Instrumental variable estimation of nonparametric models. *Econometrica*, 71(5):1565–1578, 2003. doi: 10.1111/1468-0262.00459.
- Rahul Singh, Maneesh Sahani, and Arthur Gretton. Kernel instrumental variable regression. *Advances in Neural Information Processing Systems*, 32, 2019.
- Eric J. Tchetgen Tchetgen, Andrew Ying, Yifan Cui, Xu Shi, and Wang Miao. An introduction to proximal causal learning. *Statistical Science*, 39(3):375–390, 2024. doi: 10.1214/23-STS911.
- Liyuan Xu, Yutian Chen, Siddarth Srinivasan, Nando de Freitas, Arnaud Doucet, and Arthur Gretton. Learning deep features in instrumental variable regression. In *International Conference on Learning Representations*, 2021. URL [https://openreview.net/forum?id=sy4Kg\\_ZQmS7](https://openreview.net/forum?id=sy4Kg_ZQmS7).

## A Proof of the General Oracle Guarantee

*Proof of Theorem 3.1.* The proof has three steps. First, we prove that the oracle point is always accepted. The rule therefore stops no later than that point, and acceptance together with monotonicity turns the error bounds into the stated rates.

First, by (3.2) and (3.4),

$$D_n(\widehat{h}_{\lambda^\circ}) \leq A_n(\lambda^\circ) + B_n(\lambda^\circ) \leq \frac{3}{2}A_n(\lambda^\circ).$$

Thus  $\lambda^\circ$  is accepted, so some candidate clears its threshold and the fallback is not reached. Second, since Algorithm 1 searches from larger to smaller regularization and stops at the first acceptance, it returns a point  $\lambda_{\text{dp}} \geq \lambda^\circ$ . Third, every grid point between  $\lambda^\circ$  and 2 lies in  $\mathcal{S}_n$ , so (3.3) applies at the selected point. The rule accepted  $\lambda_{\text{dp}}$ , so  $D_n(\widehat{h}_{\lambda_{\text{dp}}}) \leq \frac{3}{2}A_n(\lambda_{\text{dp}})$  and

$$D_n(\widehat{h}_{\lambda_{\text{dp}}}) + A_n(\lambda_{\text{dp}}) \leq \frac{5}{2}A_n(\lambda_{\text{dp}}).$$

Because  $A_n$  is nonincreasing over the grid points in  $[\lambda^\circ, 2]$  and  $\lambda \mapsto 1/\lambda$  decreases, the product  $A_n(\lambda)/\lambda$  is nonincreasing there as well, so

$$A_n(\lambda_{\text{dp}}) \leq A_n(\lambda^\circ), \quad \frac{A_n(\lambda_{\text{dp}})}{\lambda_{\text{dp}}} \leq \frac{A_n(\lambda^\circ)}{\lambda^\circ}.$$

Substituting the last two displays into (3.3) proves (3.5).  $\square$

The following consequence of the source condition allows us to compare both one-step estimators directly with  $h_0$ .

**Lemma A.1** (Variational source inequality). *Let  $\mathcal{A} = \mathcal{T}^*\mathcal{T}$ ,  $b = \min\{\beta, 1\}$ ,  $p = 1 + b$ , and  $\alpha = b/(1 + b)$ . Under Assumption 2.3, every  $u \in \widetilde{\mathcal{H}}$  satisfies*

$$|\langle h_0, u \rangle| \leq \|w_0\| \|\mathcal{T}u\|^b \|u\|^{1-b}.$$

Consequently, for every fixed  $\epsilon \in (0, 1)$  and  $\lambda \in (0, 2]$ ,

$$\begin{aligned} 2\lambda |\langle h_0, u \rangle| &\leq \epsilon \{\|\mathcal{T}u\|^2 + \lambda \|u\|^2\} + C_\epsilon \|w_0\|^2 \lambda^p, \\ 2|\langle h_0, u \rangle| &\leq \epsilon \|u\|^2 + C_\epsilon \|w_0\|^{2/(1+b)} \|\mathcal{T}u\|^{2b/(1+b)}. \end{aligned}$$

*Proof.* For  $0 < \beta \leq 1$ , spectral interpolation and  $\|\mathcal{A}\| \leq 1$  give

$$|\langle h_0, u \rangle| = |\langle w_0, \mathcal{A}^{\beta/2} u \rangle| \leq \|w_0\| \|\mathcal{A}^{1/2} u\|^\beta \|u\|^{1-\beta}.$$

For  $\beta > 1$ , the positive self-adjoint operator  $\mathcal{A}$  has spectrum in  $[0, 1]$  because  $\|\mathcal{A}\| \leq 1$ . Let  $\mu_u$  denote its spectral measure associated with  $u$ . Since  $s^\beta \leq s$  for  $s \in [0, 1]$ ,

$$\|\mathcal{A}^{\beta/2}u\|^2 = \int_{[0,1]} s^\beta d\mu_u(s) \leq \int_{[0,1]} s d\mu_u(s) = \|\mathcal{A}^{1/2}u\|^2.$$

Since  $\|\mathcal{A}^{1/2}u\| = \|\mathcal{T}u\|$ , these two cases prove the first display. To obtain the first consequence, set  $E_\lambda = \|\mathcal{T}u\|^2 + \lambda\|u\|^2$  and observe that

$$2\lambda|\langle h_0, u \rangle| \leq 2\|w_0\|\lambda^{(1+b)/2}E_\lambda^{1/2}.$$

The inequality  $2uv \leq \epsilon u^2 + \epsilon^{-1}v^2$  gives the first consequence with  $C_\epsilon = \epsilon^{-1}$ , uniformly over  $b \in (0, 1]$ . Its coefficient is independent of  $\beta$  and  $\|w_0\|$ . Applying the weighted form of the same inequality to  $2\|w_0\|\|\mathcal{T}u\|^b\|u\|^{1-b}$  gives the second consequence.  $\square$

## B Proofs of Section 3.2

Throughout the appendix,  $R^\lambda(h) = L(h) + \lambda\|h\|^2$  and  $R_n^\lambda(h) = L_n(h) + \lambda\|h\|_{2,n}^2$ , where

$$\|h\|_{2,n} := \{\mathbb{E}_n[h(X)^2]\}^{1/2}.$$

In this section,  $\mathbb{E}_n$  is the average on  $I_{\text{fit}}$ ,  $L(h) = \sigma_{Y|Z}^2 + \|\mathcal{T}(h - h_0)\|^2$ , and  $D_n$  is the projected discrepancy in (3.7). We use the scales in (3.8). Constants may depend on the fixed bounds and population quantities; they are uniform over  $n$  and the regularization grid.

**Lemma B.1** (Preliminary RDIV nuisance and operator bounds). *Under Assumptions 3.3 and 3.4, with probability at least  $1 - \xi$ ,*

$$\mathbb{E}_Z[H^2\{\widehat{g}(\cdot | Z), g_0(\cdot | Z)\}] \leq C\delta_{n,\mathcal{G}}^2, \tag{B.1}$$

$$\begin{aligned} \|\widehat{r} - r_0\|_{L_2(Z)} &\leq \delta_{n,r}, \\ \sup_{h \in \mathcal{H}} \|(\widehat{\mathcal{T}} - \mathcal{T})h\|_{L_2(Z)} &\leq C\delta_{n,\mathcal{G}}, \end{aligned} \tag{B.2}$$

$$\|(\widehat{\mathcal{T}} - \mathcal{T})(h - h')\|_{L_1(Z)} \leq C\delta_{n,\mathcal{G}}\|h - h'\|, \quad h, h' \in \mathcal{H}. \tag{B.3}$$

Moreover, if

$$\widehat{L}(h) := \mathbb{E}[(Y - \widehat{\mathcal{T}}h(Z))^2 | I_{\text{pre}}],$$

then, uniformly over  $h, h' \in \mathcal{H}$ ,

$$|\{\widehat{L}(h) - \widehat{L}(h')\} - \{L(h) - L(h')\}| \leq C\delta_{n,\mathcal{G}}\|h - h'\|. \tag{B.4}$$

Additional conditions. For the conditional fitting argument in Lemma B.2, use the following measurable-version convention. Let  $Q$  denote the law of one observation  $W$ , let  $a$  denote the preliminary observations, and put  $Q_n^{\mathbf{w}} = n^{-1} \sum_{i=1}^n \delta_{w_i}$  for a fresh fitting tuple  $\mathbf{w} = (w_1, \dots, w_n)$ . Define

$$\ell_{\lambda,a}(h; w) = \{Y(w) - (\hat{\mathcal{T}}_a h)(Z(w))\}^2 + \lambda h(X(w))^2, \quad S_{n,\lambda,q}(a, \mathbf{w}) = \sup_{\substack{h \in \mathcal{H} \\ \|h - h_0\| \leq q}} |(Q_n^{\mathbf{w}} - Q)\{\ell_{\lambda,a}(h) - \ell_{\lambda,a}(h_0)\}|$$

For each  $n \geq 1$ ,  $\lambda \in \Lambda_n$  and positive rational  $q$ ,  $S_{n,\lambda,q}$  admits an extended nonnegative version jointly measurable in the preliminary observations and the fresh fitting tuple. This version agrees with  $S_{n,\lambda,q}(a, \cdot)$  for  $Q^n$ -almost every tuple for almost every preliminary realization, and agrees almost surely with  $S_{n,\lambda,q}(I_{\text{pre}}, (W_i)_{i \in I_{\text{fit}}})$  at the actual independent fitting block. All losses use the same function representatives as the empirical objective. This is only a measurability and evaluation-coherence convention; no deviation bound, rate, countability or separability is assumed. For the fitted conditional density in the RDIV preliminary construction, use a coherent measurable version at independent calibration observations. For each  $n \geq 1$ , writing  $a = I_{\text{pre}}$ , the selected density  $\hat{g}_{n,a}(x | z)$  admits a nonnegative version  $\bar{g}_n(a, x, z)$  jointly measurable in the preliminary observations,  $x$  and  $z$ . For almost every preliminary realization  $a$ , for  $P_Z$ -almost every  $z$ , the functions  $\bar{g}_n(a, \cdot, z)$  and  $\hat{g}_{n,a}(\cdot | z)$  agree  $\mu$ -almost everywhere. Almost surely, this same agreement holds at every actual calibration observation  $Z_i$ ,  $i \in I_{\text{cal}}$ . The empirical objectives and selected maximizer remain those defined using the original supplied representatives: this convention does not replace the optimizer. It supplies only measurability and coherence of the fitted density evaluated on the independent calibration sample; it assumes no normalization event, probability inequality, rate, countability or separability. For the preliminary regression estimate used in the RDIV calibration loss, take a coherent measurable version on the independent calibration sample. For each  $n \geq 1$ , writing  $a = I_{\text{pre}}$ , the selected  $L_2(Z)$  pilot has a version  $\bar{r}_n(a, z)$  jointly measurable in the preliminary observations and  $z$ . For almost every preliminary realization  $a$ ,  $\bar{r}_n(a, z) = \hat{r}_{n,a}(z)$  for  $P_Z$ -almost every  $z$ . Almost surely, the same equality holds simultaneously at every actual calibration observation  $Z_i$ ,  $i \in I_{\text{cal}}$ , for the supplied representatives used in  $D_n$ . This is only a measurability and evaluation-coherence convention. It imposes no exact joint measurability on every supplied representative, does not replace an empirical objective or selected optimizer, and supplies no boundedness event, probability inequality, error rate, countability or separability assumption. The existing pilot accuracy and envelope assumptions remain unchanged. Use the following well-definedness convention for the likelihood-transform critical-radius and concentration arguments in Assumption 3.4 and Lemma B.1, separately from the existing hypothesis-class, density-image and fitted-process conventions. Let  $Q$  be the true joint law of

$(X, Z)$ , let  $\mathcal{G}_0$  be the likelihood-transform class already defined in Assumption 3.4, and put

$$\mathcal{A}_r = \{f \in \text{star}(\mathcal{G}_0 - \mathcal{G}_0) : \|f\|_{L_2(Q)} \leq r\}.$$

For every integer  $n \geq 1$  and every  $r > 0$ , the three processes

$$\sup_{f \in \mathcal{G}_0} |(Q_n - Q)f|, \quad \sup_{f \in \mathcal{A}_r} |(Q_n - Q)f|, \quad \sup_{f \in \mathcal{A}_r} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i f(X_i, Z_i) \right|$$

regarded as extended nonnegative quantities, are measurable under  $Q^n$  for the first two and under the product of  $Q^n$  with the independent Rademacher law for the third. The functions use the same fixed density representatives as the likelihood objective. This convention only makes these empirical processes and their expectations well defined. It supplies no concentration event, deviation inequality, probability bound, rate, countability, separability, or additional statistical restriction.

*Additional condition.* For another result of the paper: Use the following well-definedness convention for the likelihood-transform critical-radius and concentration arguments in Assumption 3.4 and Lemma B.1, separately from the existing hypothesis-class, density-image and fitted-process conventions. Let  $Q$  be the true joint law of  $(X, Z)$ , let  $\mathcal{G}_0$  be the likelihood-transform class already defined in Assumption 3.4, and put

$$\mathcal{A}_r = \{f \in \text{star}(\mathcal{G}_0 - \mathcal{G}_0) : \|f\|_{L_2(Q)} \leq r\}.$$

For every integer  $n \geq 1$  and every  $r > 0$ , the three processes

$$\sup_{f \in \mathcal{G}_0} |(Q_n - Q)f|, \quad \sup_{f \in \mathcal{A}_r} |(Q_n - Q)f|, \quad \sup_{f \in \mathcal{A}_r} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i f(X_i, Z_i) \right|$$

regarded as extended nonnegative quantities, are measurable under  $Q^n$  for the first two and under the product of  $Q^n$  with the independent Rademacher law for the third. The functions use the same fixed density representatives as the likelihood objective. This convention only makes these empirical processes and their expectations well defined. It supplies no concentration event, deviation inequality, probability bound, rate, countability, separability, or additional statistical restriction.

*Remark (notation).*  $H^2(g_1, g_2) = \int (\sqrt{g_1(x)} - \sqrt{g_2(x)})^2 \mu(dx)$  is defined for probability densities  $g_1, g_2$  with respect to the common measure  $\mu$  on  $\mathcal{X}$ ; for conditional densities it is evaluated at each fixed  $z$ , with any averaging over  $Z$  written separately.

*Proof.* Allocate failure probability  $\xi/2$  to the conditional-density estimator and  $\xi/2$  to the pilot regression. As elsewhere, evaluate the corresponding deterministic accuracy bounds at these confidence levels and relabel them  $\delta_{n,\mathcal{G}}$  and  $\delta_{n,r}$ . We establish the density bound directly using the Localized Concentration lemma of Li et al. [2024]. Write  $P = \mathbb{E}_{X,Z}$  and  $P_n = \mathbb{E}_{n,\text{pre}}$ , and define

$$f_g(x, z) = \sqrt{\frac{g(x | z) + g_0(x | z)}{2g_0(x | z)}}, \quad \mathcal{U} = \{f_g - 1 : g \in \mathcal{G}\}.$$

Realizability gives  $f_{g_0} = 1$ , so  $0 \in \mathcal{U}$ . The density bounds give

$$2^{-1/2} \leq f_g \leq M, \quad M = \sqrt{\frac{C_G/c_0 + 1}{2}}.$$

Thus  $\mathcal{U}$  is uniformly bounded by a fixed constant. Define the globally Lipschitz loss

$$\ell(v) = -\log(\max\{1 + v, 2^{-1/2}\}).$$

The maximum map is 1-Lipschitz, and the derivative of  $-\log t$  has absolute value at most  $\sqrt{2}$  for  $t \geq 2^{-1/2}$ . Hence  $\ell$  is  $\sqrt{2}$ -Lipschitz,  $\ell(0) = 0$ , and  $\ell(f_g - 1) = -\log f_g$ .

We verify the localized complexity condition for this application. Put  $\mathcal{S} = \text{star}(\mathcal{G}_0 - \mathcal{G}_0)$  and choose a deterministic critical radius  $d \leq \delta_{n,\mathcal{G}}$  as supplied by Assumption 3.4. By the critical-radius definition,  $\mathfrak{R}_n(d; \mathcal{S}) \leq d^2$ . For  $s \geq d$ , star-shapedness implies that  $(d/s)v \in \mathcal{S}$  and  $\|(d/s)v\|_{L_2(X,Z)} \leq d$  whenever  $v \in \mathcal{S}$  and  $\|v\|_{L_2(X,Z)} \leq s$ . Scaling inside the Rademacher supremum therefore gives

$$\mathfrak{R}_n(s; \mathcal{S}) \leq \frac{s}{d} \mathfrak{R}_n(d; \mathcal{S}) \leq sd.$$

Since  $1 \in \mathcal{G}_0$ , we have  $\text{star}(\mathcal{U}) \subseteq \mathcal{S}$ . Taking  $s = \delta_{n,\mathcal{G}}$  yields

$$\mathfrak{R}_n(\delta_{n,\mathcal{G}}; \text{star}(\mathcal{U})) \leq \delta_{n,\mathcal{G}} d \leq \delta_{n,\mathcal{G}}^2.$$

Apply the cited localized concentration lemma with input  $(X, Z)$ , one coordinate, class  $\mathcal{U}$ , fixed comparator 0, loss  $\ell$ , and failure probability  $\xi/2$ . The measurability and logarithmic radius floor are supplied by Assumption 3.4, with the confidence evaluation specified in the proof. Consequently, for a fixed constant  $A > 0$ , with probability at least  $1 - \xi/2$ ,

$$|(P_n - P)(-\log f_g)| \leq A \{ \delta_{n,\mathcal{G}} \|f_g - 1\|_{L_2(X,Z)} + \delta_{n,\mathcal{G}}^2 \}, \quad g \in \mathcal{G}.$$

To use empirical optimality, note that  $P_n \log g_0(X | Z)$  is finite by the density bounds. Since  $\hat{g}$  maximizes the empirical log likelihood and all candidate densities are bounded

above,  $\widehat{g}(X_i | Z_i) > 0$  at every preliminary observation, almost surely. For  $a > 0$ , the arithmetic-geometric mean inequality gives  $(a + 1)/2 \geq \sqrt{a}$ , and hence

$$-\log \sqrt{\frac{a+1}{2}} \leq -\frac{1}{4} \log a.$$

Taking  $a = \widehat{g}(X_i | Z_i)/g_0(X_i | Z_i)$  and averaging gives

$$P_n(-\log f_{\widehat{g}}) \leq -\frac{1}{4} P_n \log \frac{\widehat{g}(X | Z)}{g_0(X | Z)} \leq 0.$$

For every  $g \in \mathcal{G}$ , conditional density normalization gives

$$Pf_g^2 = \mathbb{E}_Z \int \frac{g(x | Z) + g_0(x | Z)}{2} \mu(dx) = 1.$$

Since  $\log t \leq t - 1$  for  $t > 0$ , it follows that

$$P(-\log f_g) \geq 1 - Pf_g = \frac{1}{2} P(f_g - 1)^2.$$

On the concentration event, put  $b = \|f_{\widehat{g}} - 1\|_{L_2(X, Z)}$ . Combining these inequalities and using Young's inequality gives

$$\frac{b^2}{2} \leq P(-\log f_{\widehat{g}}) \leq A\delta_{n, \mathcal{G}}b + A\delta_{n, \mathcal{G}}^2 \leq \frac{b^2}{4} + (A^2 + A)\delta_{n, \mathcal{G}}^2.$$

Thus  $b^2 \leq 4(A^2 + A)\delta_{n, \mathcal{G}}^2$ .

Finally, for  $t \geq 0$  and  $f = \sqrt{(t^2 + 1)/2}$ , we have  $f \leq t + 1$  and

$$|t - 1| = \frac{2(f + 1)}{t + 1} |f - 1| \leq 4|f - 1|.$$

Apply this pointwise with  $t = \sqrt{\widehat{g}/g_0}$  and integrate under the true joint distribution to obtain

$$\begin{aligned} \mathbb{E}_Z [H^2\{\widehat{g}(\cdot | Z), g_0(\cdot | Z)\}] &= P \left( \sqrt{\frac{\widehat{g}(X | Z)}{g_0(X | Z)}} - 1 \right)^2 \\ &\leq 16b^2 \leq 64(A^2 + A)\delta_{n, \mathcal{G}}^2. \end{aligned}$$

This proves (B.1) with probability at least  $1 - \xi/2$ . The nuisance-rate assumption gives  $\|\widehat{r} - r_0\|_{L_2(Z)} \leq \delta_{n, r}$  with probability at least  $1 - \xi/2$ . Therefore, without requiring independence between  $\widehat{g}$  and  $\widehat{r}$ , the union bound shows that both inequalities hold simultaneously with probability at least  $1 - \xi$ . We work on this intersection for the rest of the proof. Here

$$H^2(g_1, g_2) = \int \{\sqrt{g_1(x)} - \sqrt{g_2(x)}\}^2 \mu(dx).$$

Let  $P_X$  be the true marginal distribution of  $X$  and let  $\hat{P}_X$  be the marginal induced by  $Z \sim P_Z$  and  $X | Z \sim \hat{g}(\cdot | Z)$ . The density bounds in Assumption 3.3 give  $\hat{g}(x | z)/g_0(x | z) \leq C_G/c_0$  and hence

$$\frac{d\hat{P}_X}{dP_X} \leq \frac{C_G}{c_0}.$$

The total-variation–Hellinger inequality and  $\|h\|_\infty \leq B$  give

$$\|(\hat{\mathcal{T}} - \mathcal{T})h\|^2 \leq B^2 \mathbb{E}_Z \left[ \left\{ \int |\hat{g}(x | Z) - g_0(x | Z)| \mu(dx) \right\}^2 \right] \leq CB^2 \delta_{n,\mathcal{G}}^2,$$

which proves (B.2). For  $u = h - h'$ , factor  $\hat{g}/g_0 - 1 = (\sqrt{\hat{g}/g_0} - 1)(\sqrt{\hat{g}/g_0} + 1)$ . Cauchy–Schwarz and the conditional density-ratio bound yield

$$\begin{aligned} \|(\hat{\mathcal{T}} - \mathcal{T})u\|_1 &\leq \mathbb{E} \left[ |u(X)| \left| \sqrt{\frac{\hat{g}(X | Z)}{g_0(X | Z)}} - 1 \right| \left\{ \sqrt{\frac{\hat{g}(X | Z)}{g_0(X | Z)}} + 1 \right\} \right] \\ &\leq C\|u\| \mathbb{E}_Z [H^2\{\hat{g}(\cdot | Z), g_0(\cdot | Z)\}]^{1/2} \leq C\delta_{n,\mathcal{G}}\|u\|, \end{aligned}$$

which proves (B.3). Li et al. [2024] use the same operator argument in their MLE analysis.

Finally, write  $e_h = (\hat{\mathcal{T}} - \mathcal{T})h$ . Pointwise,

$$\begin{aligned} (Y - \hat{\mathcal{T}}h)^2 - (Y - \hat{\mathcal{T}}h')^2 - \{(Y - \mathcal{T}h)^2 - (Y - \mathcal{T}h')^2\} \\ = -(\hat{\mathcal{T}} - \mathcal{T})u(2Y - \hat{\mathcal{T}}h - \hat{\mathcal{T}}h') + \mathcal{T}u(e_h + e_{h'}). \end{aligned}$$

The expectation of the first term is at most  $4B\|(\hat{\mathcal{T}} - \mathcal{T})u\|_1$ ; the second is at most  $\|\mathcal{T}u\|(\|e_h\| + \|e_{h'}\|)$ . Apply (B.2), (B.3), and conditional Jensen to obtain (B.4).  $\square$

**Lemma B.2** (Uniform RDIV empirical comparisons). *Under Assumptions 2.1 to 3.4, there is an event  $\mathcal{E}_n^R$  with probability at least  $1 - \xi$  on which, simultaneously for every  $\lambda \in \Lambda_n$  and  $h \in \mathcal{H}$ ,*

$$\left| \{\hat{R}^\lambda(h) - \hat{R}^\lambda(h_0)\} - \{R_n^\lambda(h) - R_n^\lambda(h_0)\} \right| \leq C\{\delta_{n,\mathcal{H}}\|h - h_0\| + \delta_{n,\mathcal{H}}^2\}, \quad (\text{B.5})$$

$$\left| D_n(\hat{h}_\lambda) - \|\mathcal{T}(\hat{h}_\lambda - h_0)\|^2 \right| \leq C\nu_n\|\mathcal{T}(\hat{h}_\lambda - h_0)\| + C\nu_n^2, \quad (\text{B.6})$$

where  $\hat{R}^\lambda(h) = \hat{L}(h) + \lambda\|h\|^2$ .

*Additional conditions. For the conditional fitting argument in Lemma B.2, use the following measurable-version convention. Let  $Q$  denote the law of one observation  $W$ , let  $a$  denote the preliminary observations, and put  $Q_n^{\mathbf{w}} = n^{-1} \sum_{i=1}^n \delta_{w_i}$  for a fresh fitting tuple  $\mathbf{w} = (w_1, \dots, w_n)$ . Define*

$$\ell_{\lambda,a}(h; w) = \{Y(w) - (\hat{\mathcal{T}}_a h)(Z(w))\}^2 + \lambda h(X(w))^2, \quad S_{n,\lambda,q}(a, \mathbf{w}) = \sup_{\substack{h \in \mathcal{H} \\ \|h - h_0\| \leq q}} |(Q_n^{\mathbf{w}} - Q)\{\ell_{\lambda,a}(h) - \ell_{\lambda,a}(h_0)\}|$$

For each  $n \geq 1$ ,  $\lambda \in \Lambda_n$  and positive rational  $q$ ,  $S_{n,\lambda,q}$  admits an extended nonnegative version jointly measurable in the preliminary observations and the fresh fitting tuple. This version agrees with  $S_{n,\lambda,q}(a, \cdot)$  for  $Q^n$ -almost every tuple for almost every preliminary realization, and agrees almost surely with  $S_{n,\lambda,q}(I_{\text{pre}}, (W_i)_{i \in I_{\text{fit}}})$  at the actual independent fitting block. All losses use the same function representatives as the empirical objective. This is only a measurability and evaluation-coherence convention; no deviation bound, rate, countability or separability is assumed. For the fitted conditional density in the RDIV preliminary construction, use a coherent measurable version at independent calibration observations. For each  $n \geq 1$ , writing  $a = I_{\text{pre}}$ , the selected density  $\hat{g}_{n,a}(x \mid z)$  admits a nonnegative version  $\bar{g}_n(a, x, z)$  jointly measurable in the preliminary observations,  $x$  and  $z$ . For almost every preliminary realization  $a$ , for  $P_Z$ -almost every  $z$ , the functions  $\bar{g}_n(a, \cdot, z)$  and  $\hat{g}_{n,a}(\cdot \mid z)$  agree  $\mu$ -almost everywhere. Almost surely, this same agreement holds at every actual calibration observation  $Z_i$ ,  $i \in I_{\text{cal}}$ . The empirical objectives and selected maximizer remain those defined using the original supplied representatives: this convention does not replace the optimizer. It supplies only measurability and coherence of the fitted density evaluated on the independent calibration sample; it assumes no normalization event, probability inequality, rate, countability or separability. For the preliminary regression estimate used in the RDIV calibration loss, take a coherent measurable version on the independent calibration sample. For each  $n \geq 1$ , writing  $a = I_{\text{pre}}$ , the selected  $L_2(Z)$  pilot has a version  $\bar{r}_n(a, z)$  jointly measurable in the preliminary observations and  $z$ . For almost every preliminary realization  $a$ ,  $\bar{r}_n(a, z) = \hat{r}_{n,a}(z)$  for  $P_Z$ -almost every  $z$ . Almost surely, the same equality holds simultaneously at every actual calibration observation  $Z_i$ ,  $i \in I_{\text{cal}}$ , for the supplied representatives used in  $D_n$ . This is only a measurability and evaluation-coherence convention. It imposes no exact joint measurability on every supplied representative, does not replace an empirical objective or selected optimizer, and supplies no boundedness event, probability inequality, error rate, countability or separability assumption. The existing pilot accuracy and envelope assumptions remain unchanged. Use the following well-definedness convention for the likelihood-transform critical-radius and concentration arguments in Assumption 3.4 and Lemma B.1, separately from the existing hypothesis-class, density-image and fitted-process conventions. Let  $Q$  be the true joint law of  $(X, Z)$ , let  $\mathcal{G}_0$  be the likelihood-transform class already defined in Assumption 3.4, and put

$$\mathcal{A}_r = \{f \in \text{star}(\mathcal{G}_0 - \mathcal{G}_0) : \|f\|_{L_2(Q)} \leq r\}.$$

For every integer  $n \geq 1$  and every  $r > 0$ , the three processes

$$\sup_{f \in \mathcal{G}_0} |(Q_n - Q)f|, \quad \sup_{f \in \mathcal{A}_r} |(Q_n - Q)f|, \quad \sup_{f \in \mathcal{A}_r} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i f(X_i, Z_i) \right|$$

regarded as extended nonnegative quantities, are measurable under  $Q^n$  for the first two and under the product of  $Q^n$  with the independent Rademacher law for the third. The functions use the same fixed density representatives as the likelihood objective. This convention only makes these empirical processes and their expectations well defined. It supplies no concentration event, deviation inequality, probability bound, rate, countability, separability, or additional statistical restriction.

*Additional conditions.* For a result cited in the proof of Lemma B.2: For the conditional fitting argument in Lemma B.2, use the following measurable-version convention. Let  $Q$  denote the law of one observation  $W$ , let  $a$  denote the preliminary observations, and put  $Q_n^{\mathbf{w}} = n^{-1} \sum_{i=1}^n \delta_{w_i}$  for a fresh fitting tuple  $\mathbf{w} = (w_1, \dots, w_n)$ . Define

$$\ell_{\lambda,a}(h; w) = \{Y(w) - (\hat{\mathcal{T}}_a h)(Z(w))\}^2 + \lambda h(X(w))^2, \quad S_{n,\lambda,q}(a, \mathbf{w}) = \sup_{\substack{h \in \mathcal{H} \\ \|h - h_0\| \leq q}} |(Q_n^{\mathbf{w}} - Q)\{\ell_{\lambda,a}(h) - \ell_{\lambda,a}(h_0)\}|$$

For each  $n \geq 1$ ,  $\lambda \in \Lambda_n$  and positive rational  $q$ ,  $S_{n,\lambda,q}$  admits an extended nonnegative version jointly measurable in the preliminary observations and the fresh fitting tuple. This version agrees with  $S_{n,\lambda,q}(a, \cdot)$  for  $Q^n$ -almost every tuple for almost every preliminary realization, and agrees almost surely with  $S_{n,\lambda,q}(I_{\text{pre}}, (W_i)_{i \in I_{\text{fit}}})$  at the actual independent fitting block. All losses use the same function representatives as the empirical objective. This is only a measurability and evaluation-coherence convention; no deviation bound, rate, countability or separability is assumed. For a result cited in the proof of Lemma B.2: For the fitted conditional density in the RDIV preliminary construction, use a coherent measurable version at independent calibration observations. For each  $n \geq 1$ , writing  $a = I_{\text{pre}}$ , the selected density  $\hat{g}_{n,a}(x | z)$  admits a nonnegative version  $\bar{g}_n(a, x, z)$  jointly measurable in the preliminary observations,  $x$  and  $z$ . For almost every preliminary realization  $a$ , for  $P_Z$ -almost every  $z$ , the functions  $\bar{g}_n(a, \cdot, z)$  and  $\hat{g}_{n,a}(\cdot | z)$  agree  $\mu$ -almost everywhere. Almost surely, this same agreement holds at every actual calibration observation  $Z_i$ ,  $i \in I_{\text{cal}}$ . The empirical objectives and selected maximizer remain those defined using the original supplied representatives: this convention does not replace the optimizer. It supplies only measurability and coherence of the fitted density evaluated on the independent calibration sample; it assumes no normalization event, probability inequality, rate, countability or separability. For a result cited in the proof of Lemma B.2: For the preliminary regression estimate used in the RDIV calibration loss, take a coherent measurable version on the independent calibration sample. For each  $n \geq 1$ , writing  $a = I_{\text{pre}}$ , the selected  $L_2(Z)$  pilot has a version  $\bar{r}_n(a, z)$  jointly measurable in the preliminary observations and  $z$ . For almost every preliminary realization  $a$ ,  $\bar{r}_n(a, z) = \hat{r}_{n,a}(z)$  for  $P_Z$ -almost every  $z$ . Almost surely, the same equality holds simultaneously at every actual

calibration observation  $Z_i$ ,  $i \in I_{\text{cal}}$ , for the supplied representatives used in  $D_n$ . This is only a measurability and evaluation-coherence convention. It imposes no exact joint measurability on every supplied representative, does not replace an empirical objective or selected optimizer, and supplies no boundedness event, probability inequality, error rate, countability or separability assumption. The existing pilot accuracy and envelope assumptions remain unchanged. For a result cited in the proof of Lemma B.2: Use the following well-definedness convention for the likelihood-transform critical-radius and concentration arguments in Assumption 3.4 and Lemma B.1, separately from the existing hypothesis-class, density-image and fitted-process conventions. Let  $Q$  be the true joint law of  $(X, Z)$ , let  $\mathcal{G}_0$  be the likelihood-transform class already defined in Assumption 3.4, and put

$$\mathcal{A}_r = \{f \in \text{star}(\mathcal{G}_0 - \mathcal{G}_0) : \|f\|_{L_2(Q)} \leq r\}.$$

For every integer  $n \geq 1$  and every  $r > 0$ , the three processes

$$\sup_{f \in \mathcal{G}_0} |(Q_n - Q)f|, \quad \sup_{f \in \mathcal{A}_r} |(Q_n - Q)f|, \quad \sup_{f \in \mathcal{A}_r} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i f(X_i, Z_i) \right|$$

regarded as extended nonnegative quantities, are measurable under  $Q^n$  for the first two and under the product of  $Q^n$  with the independent Rademacher law for the third. The functions use the same fixed density representatives as the likelihood objective. This convention only makes these empirical processes and their expectations well defined. It supplies no concentration event, deviation inequality, probability bound, rate, countability, separability, or additional statistical restriction.

*Remark (notation).*  $\mathcal{E}_n^{\text{R}}$  is the intersection of the preliminary, fitting, and calibration events under the paper's confidence convention, with probability at least  $1 - \xi$ , on which the fitting comparisons hold simultaneously for every  $h \in \mathcal{H}$  and  $\lambda \in \Lambda_n$ , and the projected calibration comparisons hold simultaneously for the fitted candidates  $\hat{h}_\lambda$ ,  $\lambda \in \Lambda_n$ .

*Proof.* First condition on  $I_{\text{pre}}$ , so that  $\hat{\mathcal{T}}$  is fixed and the expected critical-radius bounds apply to the independent fitting sample. For each fixed  $\lambda \in \Lambda_n$ , apply the localized concentration inequality in Foster and Syrgkanis [2023, Lemma 14] to the two-coordinate class

$$h \mapsto \{\hat{\mathcal{T}}h(Z), h(X)\}$$

at the fixed comparator  $h_0$  and to the loss  $\phi_\lambda(y, u, v) = (y - u)^2 + \lambda v^2$ . This loss is uniformly Lipschitz on the bounded range because  $\lambda \leq 2$ . The critical-radius bounds in Assumption 3.4 give

$$\left| \{\hat{R}^\lambda(h) - \hat{R}^\lambda(h_0)\} - \{R_n^\lambda(h) - R_n^\lambda(h_0)\} \right| \leq C \left[ \delta_{n, \mathcal{H}} \{\|\hat{\mathcal{T}}(h - h_0)\| + \|h - h_0\|\} + \delta_{n, \mathcal{H}}^2 \right].$$

Conditional Jensen and the density-ratio bound in the proof of Lemma B.1 give  $\|\widehat{\mathcal{T}}u\| \leq (C_G/c_0)^{1/2}\|u\|$ . Thus, a union bound over the  $K_n + 1$  deterministic grid points gives (B.5). The definition of  $\eta_n$  makes its total failure probability at most  $\xi$ .

For calibration, condition on  $I_{\text{pre}}$  and  $I_{\text{fit}}$ , and set

$$q_\lambda = \widehat{r} - \widehat{\mathcal{T}}\widehat{h}_\lambda, \quad u_\lambda = r_0 - \mathcal{T}\widehat{h}_\lambda.$$

The functions  $q_\lambda$  are fixed under this conditioning and bounded by  $2B$ . Bernstein's inequality and a union bound over  $\Lambda_n$  give, with conditional failure probability at most  $\xi$ ,

$$|D_n(\widehat{h}_\lambda) - \|q_\lambda\|^2| \leq C\{\delta_{n,\xi}\|q_\lambda\| + \delta_{n,\xi}^2\}, \quad \lambda \in \Lambda_n.$$

The simultaneous preliminary-sample bounds in Lemma B.1 give

$$\|q_\lambda - u_\lambda\| \leq \|\widehat{r} - r_0\| + \|(\widehat{\mathcal{T}} - \mathcal{T})\widehat{h}_\lambda\| \leq \delta_{n,r} + C\delta_{n,\mathcal{G}}.$$

The identity  $|\|q_\lambda\|^2 - \|u_\lambda\|^2| \leq 2\|u_\lambda\|\|q_\lambda - u_\lambda\| + \|q_\lambda - u_\lambda\|^2$ , together with  $\|q_\lambda\| \leq \|u_\lambda\| + \delta_{n,r} + C\delta_{n,\mathcal{G}}$ , proves (B.6). Combining the preliminary, fitting, and calibration events under the confidence convention gives the stated probability. All remaining RDIV arguments take place on  $\mathcal{E}_n^{\text{R}}$ .  $\square$

**Lemma B.3** (RDIV population–empirical risk comparison). *On  $\mathcal{E}_n^{\text{R}}$ , uniformly over  $h \in \mathcal{H}$  and  $\lambda \in \Lambda_n$ ,*

$$|R^\lambda(h) - R^\lambda(h_0) - \{R_n^\lambda(h) - R_n^\lambda(h_0)\}| \leq C\{\delta_n\|h - h_0\| + \delta_n^2\}.$$

*Proof.* Apply (B.4) and (B.5). Since  $\delta_n \geq \delta_{n,\mathcal{G}} + \delta_{n,\mathcal{H}}$ , their sum is bounded by the displayed right-hand side. The fitting event already holds simultaneously over the finite grid.  $\square$

**Lemma B.4** (Direct RDIV control). *On  $\mathcal{E}_n^{\text{R}}$ , for every  $\lambda \in \Lambda_n$ ,*

$$\|\mathcal{T}(\widehat{h}_\lambda - h_0)\|^2 + \lambda\|\widehat{h}_\lambda - h_0\|^2 \leq C\left\{\frac{\delta_n^2}{\lambda} + \|w_0\|^2\lambda^p\right\}, \quad p = \min\{\beta + 1, 2\}.$$

*Proof.* Set  $e_\lambda = \widehat{h}_\lambda - h_0$ . Since  $h_0 \in \mathcal{H}$ , empirical optimality and Lemma B.3 give

$$\|\mathcal{T}e_\lambda\|^2 + \lambda\|e_\lambda\|^2 \leq C\{\delta_n\|e_\lambda\| + \delta_n^2\} + 2\lambda|\langle h_0, e_\lambda \rangle|.$$

Apply the first consequence of Lemma A.1 and  $C\delta_n\|e_\lambda\| \leq \lambda\|e_\lambda\|^2/4 + C\delta_n^2/\lambda$ . Since  $\lambda \leq 2$ , the remaining  $\delta_n^2$  term is covered by  $C\delta_n^2/\lambda$ . Absorbing the quadratic terms proves the claim.  $\square$

**Corollary B.5** (RDIV discrepancy bound). *Under the assumptions of Theorem 3.6, the calibrated RDIV scales satisfy (3.9) simultaneously for every  $\lambda \in \Lambda_n$  on  $\mathcal{E}_n^R$ , for a fixed constant  $C_R > 0$ .*

*Additional conditions. For the conditional fitting argument in Lemma B.2, use the following measurable-version convention. Let  $Q$  denote the law of one observation  $W$ , let  $a$  denote the preliminary observations, and put  $Q_n^{\mathbf{w}} = n^{-1} \sum_{i=1}^n \delta_{w_i}$  for a fresh fitting tuple  $\mathbf{w} = (w_1, \dots, w_n)$ . Define*

$$\ell_{\lambda,a}(h; w) = \{Y(w) - (\widehat{T}_a h)(Z(w))\}^2 + \lambda h(X(w))^2, \quad S_{n,\lambda,q}(a, \mathbf{w}) = \sup_{\substack{h \in \mathcal{H} \\ \|h - h_0\| \leq q}} |(Q_n^{\mathbf{w}} - Q)\{\ell_{\lambda,a}(h) - \ell_{\lambda,a}(h_0)\}|$$

*For each  $n \geq 1$ ,  $\lambda \in \Lambda_n$  and positive rational  $q$ ,  $S_{n,\lambda,q}$  admits an extended nonnegative version jointly measurable in the preliminary observations and the fresh fitting tuple. This version agrees with  $S_{n,\lambda,q}(a, \cdot)$  for  $Q^n$ -almost every tuple for almost every preliminary realization, and agrees almost surely with  $S_{n,\lambda,q}(I_{\text{pre}}, (W_i)_{i \in I_{\text{fit}}})$  at the actual independent fitting block. All losses use the same function representatives as the empirical objective. This is only a measurability and evaluation-coherence convention; no deviation bound, rate, countability or separability is assumed. For the fitted conditional density in the RDIV preliminary construction, use a coherent measurable version at independent calibration observations. For each  $n \geq 1$ , writing  $a = I_{\text{pre}}$ , the selected density  $\widehat{g}_{n,a}(x | z)$  admits a nonnegative version  $\bar{g}_n(a, x, z)$  jointly measurable in the preliminary observations,  $x$  and  $z$ . For almost every preliminary realization  $a$ , for  $P_Z$ -almost every  $z$ , the functions  $\bar{g}_n(a, \cdot, z)$  and  $\widehat{g}_{n,a}(\cdot | z)$  agree  $\mu$ -almost everywhere. Almost surely, this same agreement holds at every actual calibration observation  $Z_i$ ,  $i \in I_{\text{cal}}$ . The empirical objectives and selected maximizer remain those defined using the original supplied representatives: this convention does not replace the optimizer. It supplies only measurability and coherence of the fitted density evaluated on the independent calibration sample; it assumes no normalization event, probability inequality, rate, countability or separability. For the preliminary regression estimate used in the RDIV calibration loss, take a coherent measurable version on the independent calibration sample. For each  $n \geq 1$ , writing  $a = I_{\text{pre}}$ , the selected  $L_2(Z)$  pilot has a version  $\bar{r}_n(a, z)$  jointly measurable in the preliminary observations and  $z$ . For almost every preliminary realization  $a$ ,  $\bar{r}_n(a, z) = \widehat{r}_{n,a}(z)$  for  $P_Z$ -almost every  $z$ . Almost surely, the same equality holds simultaneously at every actual calibration observation  $Z_i$ ,  $i \in I_{\text{cal}}$ , for the supplied representatives used in  $D_n$ . This is only a measurability and evaluation-coherence convention. It imposes no exact joint measurability on every supplied representative, does not replace an empirical objective or selected optimizer, and supplies no boundedness event, probability inequality, error rate, countability or separability assumption. The existing pilot accuracy and envelope assumptions remain unchanged. Use*

the following well-definedness convention for the likelihood-transform critical-radius and concentration arguments in Assumption 3.4 and Lemma B.1, separately from the existing hypothesis-class, density-image and fitted-process conventions. Let  $Q$  be the true joint law of  $(X, Z)$ , let  $\mathcal{G}_0$  be the likelihood-transform class already defined in Assumption 3.4, and put

$$\mathcal{A}_r = \{f \in \text{star}(\mathcal{G}_0 - \mathcal{G}_0) : \|f\|_{L_2(Q)} \leq r\}.$$

For every integer  $n \geq 1$  and every  $r > 0$ , the three processes

$$\sup_{f \in \mathcal{G}_0} |(Q_n - Q)f|, \quad \sup_{f \in \mathcal{A}_r} |(Q_n - Q)f|, \quad \sup_{f \in \mathcal{A}_r} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i f(X_i, Z_i) \right|$$

regarded as extended nonnegative quantities, are measurable under  $Q^n$  for the first two and under the product of  $Q^n$  with the independent Rademacher law for the third. The functions use the same fixed density representatives as the likelihood objective. This convention only makes these empirical processes and their expectations well defined. It supplies no concentration event, deviation inequality, probability bound, rate, countability, separability, or additional statistical restriction.

*Remark.* The assumptions imported from Theorem 3.6 are the general hypotheses of Proposition 3.5, without imposing  $h_0 = 0$ .

*Proof.* Let  $\zeta_\lambda = \|\mathcal{T}(\hat{h}_\lambda - h_0)\|$ . By Lemmas B.2 and B.4,

$$D_n(\hat{h}_\lambda) \leq \zeta_\lambda^2 + C\nu_n \zeta_\lambda + C\nu_n^2 \leq c_R \left( \frac{\delta_n^2}{\lambda} + \nu_n^2 \right) + C_R \|w_0\|^2 \lambda^p$$

for fixed constants  $c_R, C_R > 0$ , where the second inequality follows from Young's inequality. The constant  $c_R$  depends only on the fixed bounds and population constants. Multiplying both  $\delta_n$  and  $\nu_n$  by  $\sqrt{\max\{c_R, 1\}}$  and relabeling the resulting upper bounds gives (3.9). The noise coefficients are independent of  $\beta$  and  $\|w_0\|$ ; the source dependence remains in the displayed bias term. These fixed multipliers are chosen in advance under the noise-calibration convention in Section 3.  $\square$

*Proof of Proposition 3.5.* Throughout the proof, we work on  $\mathcal{E}_n^R$ . Since the projected residual is bounded by  $2B$ ,  $D_n(\hat{h}_\lambda) \leq 4B^2$ . At the terminal grid point,

$$\frac{\delta_n^2}{\lambda_{K_n}} \geq \frac{\delta_{n,\xi}^2}{\lambda_{K_n}} \geq n \log\{(K_n + 1)/\xi\} \geq n \log 2.$$

For  $n \geq N_R(B)$ , its threshold is therefore at least  $4B^2$ , so the terminal point is accepted and the algorithm stops within  $K_n + 1 = O(\log n)$  fits.

If  $h_0 = 0$ , take the source element  $w_0 = 0$ . Then (3.9) gives

$$D_n(\widehat{h}_2) \leq \delta_n^2/2 + \nu_n^2 \leq \frac{3}{2}(\delta_n^2/2 + \nu_n^2),$$

so the initial point is accepted. In general, let  $k^*$  be the first accepted index. If  $k^* = 0$ , the algorithm returns through the initial-acceptance branch. If  $k^* \geq 1$ , its predecessor is rejected, and  $\rho \geq 1/2$  gives (3.1).  $\square$

**Lemma B.6** (RDIV error bounds). *Under the assumptions of Theorem 3.6, for every  $\lambda \in \Lambda_n$ , on  $\mathcal{E}_n^R$ ,*

$$\begin{aligned} \|\mathcal{T}(\widehat{h}_\lambda - h_0)\|^2 &\leq C\{D_n(\widehat{h}_\lambda) + A_n^R(\lambda)\}, \\ \|\widehat{h}_\lambda - h_0\|^2 &\leq C\left[V_n^R(\lambda) + \{D_n(\widehat{h}_\lambda) + A_n^R(\lambda)\}^\alpha\right], \end{aligned}$$

where the noise level and stability bound are defined in (3.10) and  $\alpha = \min\{\beta, 1\}/\{1 + \min\{\beta, 1\}\}$ .

*Additional conditions. For the conditional fitting argument in Lemma B.2, use the following measurable-version convention. Let  $Q$  denote the law of one observation  $W$ , let  $a$  denote the preliminary observations, and put  $Q_n^{\mathbf{w}} = n^{-1} \sum_{i=1}^n \delta_{w_i}$  for a fresh fitting tuple  $\mathbf{w} = (w_1, \dots, w_n)$ . Define*

$$\ell_{\lambda,a}(h; w) = \{Y(w) - (\widehat{\mathcal{T}}_a h)(Z(w))\}^2 + \lambda h(X(w))^2, \quad S_{n,\lambda,q}(a, \mathbf{w}) = \sup_{\substack{h \in \mathcal{H} \\ \|h - h_0\| \leq q}} |(Q_n^{\mathbf{w}} - Q)\{\ell_{\lambda,a}(h) - \ell_{\lambda,a}(h_0)\}|$$

*For each  $n \geq 1$ ,  $\lambda \in \Lambda_n$  and positive rational  $q$ ,  $S_{n,\lambda,q}$  admits an extended nonnegative version jointly measurable in the preliminary observations and the fresh fitting tuple. This version agrees with  $S_{n,\lambda,q}(a, \cdot)$  for  $Q^n$ -almost every tuple for almost every preliminary realization, and agrees almost surely with  $S_{n,\lambda,q}(I_{\text{pre}}, (W_i)_{i \in I_{\text{fit}}})$  at the actual independent fitting block. All losses use the same function representatives as the empirical objective. This is only a measurability and evaluation-coherence convention; no deviation bound, rate, countability or separability is assumed. For the fitted conditional density in the RDIV preliminary construction, use a coherent measurable version at independent calibration observations. For each  $n \geq 1$ , writing  $a = I_{\text{pre}}$ , the selected density  $\widehat{g}_{n,a}(x \mid z)$  admits a nonnegative version  $\bar{g}_n(a, x, z)$  jointly measurable in the preliminary observations,  $x$  and  $z$ . For almost every preliminary realization  $a$ , for  $P_Z$ -almost every  $z$ , the functions  $\bar{g}_n(a, \cdot, z)$  and  $\widehat{g}_{n,a}(\cdot \mid z)$  agree  $\mu$ -almost everywhere. Almost surely, this same agreement holds at every actual calibration observation  $Z_i$ ,  $i \in I_{\text{cal}}$ . The empirical objectives and selected maximizer remain those defined using the original supplied representatives: this convention does not replace the optimizer. It*

supplies only measurability and coherence of the fitted density evaluated on the independent calibration sample; it assumes no normalization event, probability inequality, rate, countability or separability. For the preliminary regression estimate used in the RDIV calibration loss, take a coherent measurable version on the independent calibration sample. For each  $n \geq 1$ , writing  $a = I_{\text{pre}}$ , the selected  $L_2(Z)$  pilot has a version  $\bar{r}_n(a, z)$  jointly measurable in the preliminary observations and  $z$ . For almost every preliminary realization  $a$ ,  $\bar{r}_n(a, z) = \hat{r}_{n,a}(z)$  for  $P_Z$ -almost every  $z$ . Almost surely, the same equality holds simultaneously at every actual calibration observation  $Z_i$ ,  $i \in I_{\text{cal}}$ , for the supplied representatives used in  $D_n$ . This is only a measurability and evaluation-coherence convention. It imposes no exact joint measurability on every supplied representative, does not replace an empirical objective or selected optimizer, and supplies no boundedness event, probability inequality, error rate, countability or separability assumption. The existing pilot accuracy and envelope assumptions remain unchanged. Use the following well-definedness convention for the likelihood-transform critical-radius and concentration arguments in Assumption 3.4 and Lemma B.1, separately from the existing hypothesis-class, density-image and fitted-process conventions. Let  $Q$  be the true joint law of  $(X, Z)$ , let  $\mathcal{G}_0$  be the likelihood-transform class already defined in Assumption 3.4, and put

$$\mathcal{A}_r = \{f \in \text{star}(\mathcal{G}_0 - \mathcal{G}_0) : \|f\|_{L_2(Q)} \leq r\}.$$

For every integer  $n \geq 1$  and every  $r > 0$ , the three processes

$$\sup_{f \in \mathcal{G}_0} |(Q_n - Q)f|, \quad \sup_{f \in \mathcal{A}_r} |(Q_n - Q)f|, \quad \sup_{f \in \mathcal{A}_r} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i f(X_i, Z_i) \right|$$

regarded as extended nonnegative quantities, are measurable under  $Q^n$  for the first two and under the product of  $Q^n$  with the independent Rademacher law for the third. The functions use the same fixed density representatives as the likelihood objective. This convention only makes these empirical processes and their expectations well defined. It supplies no concentration event, deviation inequality, probability bound, rate, countability, separability, or additional statistical restriction.

*Remark.* The assumptions of Theorem 3.6 mean its standing hypotheses inherited from Proposition 3.5; the bounds here hold for arbitrary  $h_0$ , without imposing  $h_0 = 0$ .

*Proof.* Let  $\zeta_\lambda = \|\mathcal{T}(\hat{h}_\lambda - h_0)\|$ . By (B.6),

$$\zeta_\lambda^2 \leq D_n(\hat{h}_\lambda) + C\nu_n \zeta_\lambda + C\nu_n^2.$$

Since  $\nu_n^2 \leq A_n^R(\lambda)$ , Young's inequality therefore yields

$$\zeta_\lambda^2 \leq C\{D_n(\hat{h}_\lambda) + A_n^R(\lambda)\}.$$

Set  $e_\lambda = \widehat{h}_\lambda - h_0$ . Empirical optimality and Lemma B.3 give

$$\lambda\{\|\widehat{h}_\lambda\|^2 - \|h_0\|^2\} \leq C\{\delta_n\|e_\lambda\| + \delta_n^2\}.$$

Using  $\|\widehat{h}_\lambda\|^2 - \|h_0\|^2 = \|e_\lambda\|^2 + 2\langle h_0, e_\lambda \rangle$  and dividing by  $\lambda$  yield

$$\|e_\lambda\|^2 \leq C\frac{\delta_n}{\lambda}\|e_\lambda\| + C\frac{\delta_n^2}{\lambda} + 2|\langle h_0, e_\lambda \rangle|.$$

Apply the second consequence of Lemma A.1, the weak bound above, and Young's inequality. Since  $\lambda \leq 2$ , this gives

$$\|e_\lambda\|^2 \leq C \left[ \frac{\delta_n^2}{\lambda^2} + \{D_n(\widehat{h}_\lambda) + A_n^R(\lambda)\}^\alpha \right],$$

which is the second claim.  $\square$

*Proof of Theorem 3.6.* If  $h_0 = 0$ , Proposition 3.5 selects  $\lambda = 2$ , and Lemma B.4 gives both squared errors of order  $\delta_n^2$ .

Suppose  $h_0 \neq 0$ . If  $\max\{\delta_n, \nu_n\} > 1$ , boundedness and the contraction property of  $\mathcal{T}$  give the stated rates directly. We may therefore assume  $\max\{\delta_n, \nu_n\} \leq 1$ . The discrepancy bound (3.9) gives (3.2) with  $A_n = A_n^R$  and  $q = p$ . Lemma B.6 establishes (3.3) on  $\mathcal{S}_n = (0, 2]$  with  $\gamma = \alpha$ , since  $D_n \geq 0$  and  $V_n^R(\lambda) \leq A_n^R(\lambda)/\lambda$ . For all sufficiently large  $n$ , the grid point  $\lambda_R^\circ$  defined before (3.11) exists and satisfies (3.4). The noise level in (3.10) is nonincreasing. Moreover,  $A_n^R(\lambda_R^\circ)/\lambda_R^\circ = V_n^R(\lambda_R^\circ) + \nu_n^2/\lambda_R^\circ$ , and  $\lambda_R^\circ \asymp \max\{\delta_n^{2/(p+1)}, \nu_n^{2/p}\}$  with  $\alpha = (p-1)/p$  give  $\nu_n^2/\lambda_R^\circ = O\{V_n^R(\lambda_R^\circ) + \varepsilon_{n,\beta}^\alpha\}$ . Apply Theorem 3.1 and then (3.11) to obtain the refined weak- and strong-error bounds. Boundedness covers the remaining finitely many  $n$ , after we enlarge the fixed constants. Finally,  $(p-1)/p = \alpha$  and  $\varepsilon_{n,\beta} \geq \delta_n^{2p/(p+1)}$  imply its simpler stated form.  $\square$

## B.1 Iterative RDIV

The population recursion (3.12) uses the squared weak loss for both RDIV and TRAE. The following spectral calculation gives its bias bounds.

**Lemma B.7** (Iterative Tikhonov bias). *Let  $\mathcal{A} = \mathcal{T}^*\mathcal{T}$  and  $r_\lambda(u) = \lambda/(u + \lambda)$ . For every  $j \geq 1$ ,*

$$h_0 - h_{j,\lambda}^* = r_\lambda(\mathcal{A})^j h_0 = \mathcal{A}^{\beta/2} r_\lambda(\mathcal{A})^j w_0, \quad (\text{B.7})$$

$$\|h_{j,\lambda}^* - h_0\|^2 \leq C\|w_0\|^2 \lambda^{\min\{\beta, 2j\}}, \quad (\text{B.8})$$

$$\|\mathcal{T}(h_{j,\lambda}^* - h_0)\|^2 \leq C\|w_0\|^2 \lambda^{\min\{\beta+1, 2j\}}. \quad (\text{B.9})$$

Moreover, if  $z_{j,\lambda} = r_\lambda(\mathcal{A})^j w_0$ , then  $\|z_{j,\lambda}\| \leq \|w_0\|$  and

$$\|h_{j,\lambda}^* - h_0\|^2 \leq \|w_0\|^{2/(\beta+1)} \|\mathcal{T}(h_{j,\lambda}^* - h_0)\|^{2\beta/(\beta+1)}. \quad (\text{B.10})$$

*Remark (notation).*  $r_\lambda(u) = \lambda/(u + \lambda)$  for finite  $\lambda > 0$  and  $u \geq 0$ , and  $r_\lambda(\mathcal{A})$  denotes this function applied to  $\mathcal{A}$  by spectral functional calculus.

*Proof.* The normal equation for (3.12) is

$$(\mathcal{A} + \lambda I)h_{j,\lambda}^* = \mathcal{A}h_0 + \lambda h_{j-1,\lambda}^*.$$

Subtracting this identity from  $(\mathcal{A} + \lambda I)h_0$  and iterating from  $h_{0,\lambda}^* = 0$  gives (B.7). Conditional Jensen gives  $\|\mathcal{A}\| \leq 1$ . For every  $\gamma > 0$ ,

$$\sup_{0 \leq u \leq 1} u^\gamma \left( \frac{\lambda}{u + \lambda} \right)^{2j} \leq C_{\gamma,j} \lambda^{\min\{\gamma, 2j\}}.$$

When  $\gamma \leq 2j$ , substitute  $u = \lambda v$  and bound  $v^\gamma/(1+v)^{2j}$ . When  $\gamma > 2j$ , use  $u^{\gamma-2j} \leq 1$  and  $u\lambda/(u+\lambda) \leq \lambda$ . Functional calculus with  $\gamma = \beta$  and  $\gamma = \beta+1$  proves (B.8) and (B.9). Finally, Hölder's inequality for the spectral measure of  $\mathcal{A}$ , applied to  $\mathcal{A}^{\beta/2} z_{j,\lambda}$ , gives (B.10).  $\square$

**Lemma B.8** (Recursive RDIV stability). *Under the assumptions of Theorem 3.9, there is an event  $\mathcal{E}_{n,t}^{\text{IR}}$  with probability at least  $1 - \xi$  on which, simultaneously for  $j = 1, \dots, t$  and  $\lambda \in \Lambda_n \cap [\lambda_{\text{IR}}^\circ, 2]$ ,*

$$\lambda \|\widehat{h}_{j,\lambda} - h_{j,\lambda}^*\|^2 + \|\mathcal{T}(\widehat{h}_{j,\lambda} - h_{j,\lambda}^*)\|^2 \leq a_j \frac{\delta_n^2}{\lambda}, \quad (\text{B.11})$$

$$\|\widehat{h}_{j,\lambda} - h_{j,\lambda}^*\|^2 \leq a_j \frac{\delta_n^2}{\lambda^2}, \quad (\text{B.12})$$

where  $a_j \geq 1$  is a deterministic constant that depends only on  $j$  and fixed population quantities. On the same event, the final fitted candidates also satisfy

$$\left| D_n(\widehat{h}_{t,\lambda}) - \|\mathcal{T}(\widehat{h}_{t,\lambda} - h_0)\|^2 \right| \leq C\nu_n \|\mathcal{T}(\widehat{h}_{t,\lambda} - h_0)\| + C\nu_n^2 \quad (\text{B.13})$$

simultaneously over  $\lambda \in \Lambda_n \cap [\lambda_{\text{IR}}^\circ, 2]$ .

*Additional conditions.* For the conditional fitting argument in Lemma B.2, use the following measurable-version convention. Let  $Q$  denote the law of one observation  $W$ , let  $a$  denote the preliminary observations, and put  $Q_n^\mathbf{w} = n^{-1} \sum_{i=1}^n \delta_{w_i}$  for a fresh fitting tuple  $\mathbf{w} = (w_1, \dots, w_n)$ . Define

$$\ell_{\lambda,a}(h; w) = \{Y(w) - (\widehat{\mathcal{T}}_a h)(Z(w))\}^2 + \lambda h(X(w))^2, \quad S_{n,\lambda,q}(a, \mathbf{w}) = \sup_{\substack{h \in \mathcal{H} \\ \|h - h_0\| \leq q}} |(Q_n^\mathbf{w} - Q)\{\ell_{\lambda,a}(h) - \ell_{\lambda,a}(h_0)\}|$$

For each  $n \geq 1$ ,  $\lambda \in \Lambda_n$  and positive rational  $q$ ,  $S_{n,\lambda,q}$  admits an extended nonnegative version jointly measurable in the preliminary observations and the fresh fitting tuple. This version agrees with  $S_{n,\lambda,q}(a, \cdot)$  for  $Q^n$ -almost every tuple for almost every preliminary realization, and agrees almost surely with  $S_{n,\lambda,q}(I_{\text{pre}}, (W_i)_{i \in I_{\text{fit}}})$  at the actual independent fitting block. All losses use the same function representatives as the empirical objective. This is only a measurability and evaluation-coherence convention; no deviation bound, rate, countability or separability is assumed. For the fitted conditional density in the RDIV preliminary construction, use a coherent measurable version at independent calibration observations. For each  $n \geq 1$ , writing  $a = I_{\text{pre}}$ , the selected density  $\hat{g}_{n,a}(x | z)$  admits a nonnegative version  $\bar{g}_n(a, x, z)$  jointly measurable in the preliminary observations,  $x$  and  $z$ . For almost every preliminary realization  $a$ , for  $P_Z$ -almost every  $z$ , the functions  $\bar{g}_n(a, \cdot, z)$  and  $\hat{g}_{n,a}(\cdot | z)$  agree  $\mu$ -almost everywhere. Almost surely, this same agreement holds at every actual calibration observation  $Z_i$ ,  $i \in I_{\text{cal}}$ . The empirical objectives and selected maximizer remain those defined using the original supplied representatives: this convention does not replace the optimizer. It supplies only measurability and coherence of the fitted density evaluated on the independent calibration sample; it assumes no normalization event, probability inequality, rate, countability or separability. For the preliminary regression estimate used in the RDIV calibration loss, take a coherent measurable version on the independent calibration sample. For each  $n \geq 1$ , writing  $a = I_{\text{pre}}$ , the selected  $L_2(Z)$  pilot has a version  $\bar{r}_n(a, z)$  jointly measurable in the preliminary observations and  $z$ . For almost every preliminary realization  $a$ ,  $\bar{r}_n(a, z) = \hat{r}_{n,a}(z)$  for  $P_Z$ -almost every  $z$ . Almost surely, the same equality holds simultaneously at every actual calibration observation  $Z_i$ ,  $i \in I_{\text{cal}}$ , for the supplied representatives used in  $D_n$ . This is only a measurability and evaluation-coherence convention. It imposes no exact joint measurability on every supplied representative, does not replace an empirical objective or selected optimizer, and supplies no boundedness event, probability inequality, error rate, countability or separability assumption. The existing pilot accuracy and envelope assumptions remain unchanged. For the fixed-population-center fitting argument in Lemma B.8, use the following measurable-version convention in addition to the existing  $h_0$ -centered convention. Let  $Q$  be the law of  $W$ , let  $a$  be the preliminary observations, and let  $Q_n^{\mathbf{w}} = n^{-1} \sum_{i=1}^n \delta_{w_i}$  for a fresh fitting tuple. For every  $n \geq 1$ ,  $t \geq 1$ ,  $\lambda \in \Lambda_n$ ,  $1 \leq j \leq t$ , and positive rational  $q$ , define

$$\ell_{\lambda,a}^{(j)}(h; w) = \{Y(w) - (\hat{T}_a h)(Z(w))\}^2 + \lambda \{h(X(w)) - h_{j-1,\lambda}^*(X(w))\}^2$$

and

$$S_{n,\lambda,q}^{(j)}(a, \mathbf{w}) = \sup_{\substack{h \in \mathcal{H} \\ \|h - h_{j,\lambda}^*\| \leq q}} \left| (Q_n^{\mathbf{w}} - Q) \{ \ell_{\lambda,a}^{(j)}(h) - \ell_{\lambda,a}^{(j)}(h_{j,\lambda}^*) \} \right|.$$

For  $1 \leq k \leq t$ , also define

$$U_{n,\lambda,q}^{(k)}(\mathbf{w}) = \sup_{\substack{h \in \mathcal{H} \\ \|h - h_{k,\lambda}^*\| \leq q}} |(Q_n^{\mathbf{w}} - Q)\{h(X) - h_{k,\lambda}^*(X)\}|^2.$$

Each of these suprema admits an extended nonnegative version jointly measurable in the preliminary observations and fresh fitting tuple (with  $U^{(k)}$  viewed as constant in  $a$ ). The version agrees with the indicated supremum for  $Q^n$ -almost every tuple for almost every preliminary realization and agrees almost surely at the actual independent preliminary and fitting blocks. Use exactly the same function representatives as the empirical objective and empirical squared distance. The zero-iterate comparison is the identity  $\hat{h}_{0,\lambda} = h_{0,\lambda}^* = 0$  and requires no extra process. This is solely measurability and evaluation coherence for the displayed fixed-center processes; it assumes no deviation inequality, rate, event probability, countability, separability, or additional statistical restriction. Use the following well-definedness convention for the likelihood-transform critical-radius and concentration arguments in Assumption 3.4 and Lemma B.1, separately from the existing hypothesis-class, density-image and fitted-process conventions. Let  $Q$  be the true joint law of  $(X, Z)$ , let  $\mathcal{G}_0$  be the likelihood-transform class already defined in Assumption 3.4, and put

$$\mathcal{A}_r = \{f \in \text{star}(\mathcal{G}_0 - \mathcal{G}_0) : \|f\|_{L_2(Q)} \leq r\}.$$

For every integer  $n \geq 1$  and every  $r > 0$ , the three processes

$$\sup_{f \in \mathcal{G}_0} |(Q_n - Q)f|, \quad \sup_{f \in \mathcal{A}_r} |(Q_n - Q)f|, \quad \sup_{f \in \mathcal{A}_r} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i f(X_i, Z_i) \right|$$

regarded as extended nonnegative quantities, are measurable under  $Q^n$  for the first two and under the product of  $Q^n$  with the independent Rademacher law for the third. The functions use the same fixed density representatives as the likelihood objective. This convention only makes these empirical processes and their expectations well defined. It supplies no concentration event, deviation inequality, probability bound, rate, countability, separability, or additional statistical restriction.

*Additional condition.* For another result of the paper: For the fixed-population-center fitting argument in Lemma B.8, use the following measurable-version convention in addition to the existing  $h_0$ -centered convention. Let  $Q$  be the law of  $W$ , let  $a$  be the preliminary observations, and let  $Q_n^{\mathbf{w}} = n^{-1} \sum_{i=1}^n \delta_{w_i}$  for a fresh fitting tuple. For every  $n \geq 1$ ,  $t \geq 1$ ,  $\lambda \in \Lambda_n$ ,  $1 \leq j \leq t$ , and positive rational  $q$ , define

$$\ell_{\lambda,a}^{(j)}(h; w) = \{Y(w) - (\hat{\mathcal{T}}_a h)(Z(w))\}^2 + \lambda \{h(X(w)) - h_{j-1,\lambda}^*(X(w))\}^2$$

and

$$S_{n,\lambda,q}^{(j)}(a, \mathbf{w}) = \sup_{\substack{h \in \mathcal{H} \\ \|h - h_{j,\lambda}^*\| \leq q}} \left| (Q_n^{\mathbf{w}} - Q) \{ \ell_{\lambda,a}^{(j)}(h) - \ell_{\lambda,a}^{(j)}(h_{j,\lambda}^*) \} \right|.$$

For  $1 \leq k \leq t$ , also define

$$U_{n,\lambda,q}^{(k)}(\mathbf{w}) = \sup_{\substack{h \in \mathcal{H} \\ \|h - h_{k,\lambda}^*\| \leq q}} \left| (Q_n^{\mathbf{w}} - Q) \{ h(X) - h_{k,\lambda}^*(X) \}^2 \right|.$$

Each of these suprema admits an extended nonnegative version jointly measurable in the preliminary observations and fresh fitting tuple (with  $U^{(k)}$  viewed as constant in  $a$ ). The version agrees with the indicated supremum for  $Q^n$ -almost every tuple for almost every preliminary realization and agrees almost surely at the actual independent preliminary and fitting blocks. Use exactly the same function representatives as the empirical objective and empirical squared distance. The zero-iterate comparison is the identity  $\widehat{h}_{0,\lambda} = h_{0,\lambda}^* = 0$  and requires no extra process. This is solely measurability and evaluation coherence for the displayed fixed-center processes; it assumes no deviation inequality, rate, event probability, countability, separability, or additional statistical restriction.

*Remark (notation).*  $\{a_j\}_{j=0}^t$  is a deterministic RDIV propagation sequence initialized at  $a_0 = 0$ , with  $a_j \geq 1$  for  $1 \leq j \leq t$ ; each positive-index term depends only on  $j$  and fixed population quantities and is independent of  $n$  and  $\lambda$ .  $\mathcal{E}_{n,t}^{\text{IR}}$  is the common preliminary, recursive fitting, and final calibration event constructed with allocation  $\eta_{n,t} = \xi / \{t(K_n + 1)\}$ , of probability at least  $1 - \xi$ , on which the recursive RDIV stability comparisons hold simultaneously for  $j = 1, \dots, t$  and  $\lambda \in \Lambda_n \cap [\lambda_{\text{IR}}^\circ, 2]$ , and the projected calibration comparison holds for the final candidates on that same grid segment.

*Proof.* Fix  $j$  and  $\lambda$ , and write

$$e_j = \widehat{h}_{j,\lambda} - h_{j,\lambda}^*, \quad e_{j-1} = \widehat{h}_{j-1,\lambda} - h_{j-1,\lambda}^*.$$

Apply the fitting and first-stage comparisons in Lemmas B.1 and B.2 with fixed comparator  $h_{j,\lambda}^*$  and penalty center  $h_{j-1,\lambda}^*$ . The fitting loss is Lipschitz in the two coordinates  $\{\widehat{\mathcal{T}}h(Z), h(X)\}$ , uniformly over these bounded centers. This is the fixed-comparator fitting application of Foster and Syrgkanis [2023, Lemma 14], with comparator  $h_{j,\lambda}^*$  and penalty center  $h_{j-1,\lambda}^*$ . The critical-radius convention for  $\mathcal{H}$  controls its differences from the fixed comparator. A union bound over the deterministic iteration–grid pairs, with probability allocation  $\eta_{n,t}$ , gives

$$\begin{aligned} & \lambda \|e_j\|^2 + \|\mathcal{T}e_j\|^2 \\ & \leq R_{n,j}^{\lambda,*}(\widehat{h}_{j,\lambda}) - R_{n,j}^{\lambda,*}(h_{j,\lambda}^*) + C\{\delta_n \|e_j\| + \delta_n^2\}, \end{aligned} \tag{B.14}$$

where

$$R_{n,j}^{\lambda,*}(h) = L_n(h) + \lambda \|h - h_{j-1,\lambda}^*\|_{2,n}^2.$$

Empirical optimality in (3.13) and the identity

$$\begin{aligned} & \|\widehat{h}_{j,\lambda} - h_{j-1,\lambda}^*\|_{2,n}^2 - \|h_{j,\lambda}^* - h_{j-1,\lambda}^*\|_{2,n}^2 \\ & - \|\widehat{h}_{j,\lambda} - \widehat{h}_{j-1,\lambda}\|_{2,n}^2 + \|h_{j,\lambda}^* - \widehat{h}_{j-1,\lambda}\|_{2,n}^2 = 2\langle e_j, e_{j-1} \rangle_n \end{aligned}$$

show that the empirical difference in (B.14) is at most  $2\lambda\langle e_j, e_{j-1} \rangle_n$ . Apply the localized squared-loss comparison to  $(h - h_{k,\lambda}^*)^2$ , whose comparator and center are fixed. Here the localized squared-loss comparison is the scalar application of Foster and Syrgkanis [2023, Lemma 14] at the fixed comparator  $h_{k,\lambda}^*$ . Young's inequality gives, on the same simultaneous event,

$$\|e_k\|_{2,n} \leq C(\|e_k\| + \delta_n), \quad k = j-1, j.$$

Young's inequality yields, for a fixed constant  $C > 1$  independent of  $n, \lambda, j, \beta$ , and  $\|w_0\|$ ,

$$\lambda\|e_j\|^2 + \|\mathcal{T}e_j\|^2 \leq C\frac{\delta_n^2}{\lambda} + C\lambda\|e_{j-1}\|^2.$$

Starting from  $e_0 = 0$ , choose  $a_0 = 0$  and  $a_j = C(1 + a_{j-1})$ . Induction proves (B.11); division by  $\lambda$  gives (B.12). Hence, for each fixed  $j$ ,  $a_j$  depends only on  $j$  and the fixed constants.

For calibration, condition on the preliminary and fitting samples, so the final candidates are fixed functions. The Bernstein and nuisance-comparison argument in the proof of Lemma B.2, followed by a union bound over the oracle grid segment, gives (B.13). Intersecting this calibration event with the recursive fitting event under the confidence allocation  $\eta_{n,t}$  gives  $\mathcal{E}_{n,t}^{\text{IR}}$  with probability at least  $1 - \xi$ .  $\square$

**Corollary B.9** (Iterative RDIV discrepancy bound). *Under the assumptions of Theorem 3.9, on  $\mathcal{E}_{n,t}^{\text{IR}}$ , simultaneously for every  $\lambda \in \Lambda_n \cap [\lambda_{\text{IR}}^\circ, 2]$ ,*

$$D_n(\widehat{h}_{t,\lambda}) \leq \frac{a_t \delta_n^2}{\lambda} + \nu_n^2 + C_{\text{IR}} \|w_0\|^2 \lambda^{q_t}, \quad q_t = \min\{\beta + 1, 2t\}.$$

*Additional conditions. For the conditional fitting argument in Lemma B.2, use the following measurable-version convention. Let  $Q$  denote the law of one observation  $W$ , let  $a$  denote the preliminary observations, and put  $Q_n^{\mathbf{w}} = n^{-1} \sum_{i=1}^n \delta_{w_i}$  for a fresh fitting tuple  $\mathbf{w} = (w_1, \dots, w_n)$ . Define*

$$\ell_{\lambda,a}(h; w) = \{Y(w) - (\widehat{\mathcal{T}}_a h)(Z(w))\}^2 + \lambda h(X(w))^2, \quad S_{n,\lambda,q}(a, \mathbf{w}) = \sup_{\substack{h \in \mathcal{H} \\ \|h - h_0\| \leq q}} |(Q_n^{\mathbf{w}} - Q)\{\ell_{\lambda,a}(h) - \ell_{\lambda,a}(h_0)\}|$$

For each  $n \geq 1$ ,  $\lambda \in \Lambda_n$  and positive rational  $q$ ,  $S_{n,\lambda,q}$  admits an extended nonnegative version jointly measurable in the preliminary observations and the fresh fitting tuple. This version agrees with  $S_{n,\lambda,q}(a, \cdot)$  for  $Q^n$ -almost every tuple for almost every preliminary realization, and agrees almost surely with  $S_{n,\lambda,q}(I_{\text{pre}}, (W_i)_{i \in I_{\text{fit}}})$  at the actual independent fitting block. All losses use the same function representatives as the empirical objective. This is only a measurability and evaluation-coherence convention; no deviation bound, rate, countability or separability is assumed. For the fitted conditional density in the RDIV preliminary construction, use a coherent measurable version at independent calibration observations. For each  $n \geq 1$ , writing  $a = I_{\text{pre}}$ , the selected density  $\hat{g}_{n,a}(x | z)$  admits a nonnegative version  $\bar{g}_n(a, x, z)$  jointly measurable in the preliminary observations,  $x$  and  $z$ . For almost every preliminary realization  $a$ , for  $P_Z$ -almost every  $z$ , the functions  $\bar{g}_n(a, \cdot, z)$  and  $\hat{g}_{n,a}(\cdot | z)$  agree  $\mu$ -almost everywhere. Almost surely, this same agreement holds at every actual calibration observation  $Z_i$ ,  $i \in I_{\text{cal}}$ . The empirical objectives and selected maximizer remain those defined using the original supplied representatives: this convention does not replace the optimizer. It supplies only measurability and coherence of the fitted density evaluated on the independent calibration sample; it assumes no normalization event, probability inequality, rate, countability or separability. For the preliminary regression estimate used in the RDIV calibration loss, take a coherent measurable version on the independent calibration sample. For each  $n \geq 1$ , writing  $a = I_{\text{pre}}$ , the selected  $L_2(Z)$  pilot has a version  $\bar{r}_n(a, z)$  jointly measurable in the preliminary observations and  $z$ . For almost every preliminary realization  $a$ ,  $\bar{r}_n(a, z) = \hat{r}_{n,a}(z)$  for  $P_Z$ -almost every  $z$ . Almost surely, the same equality holds simultaneously at every actual calibration observation  $Z_i$ ,  $i \in I_{\text{cal}}$ , for the supplied representatives used in  $D_n$ . This is only a measurability and evaluation-coherence convention. It imposes no exact joint measurability on every supplied representative, does not replace an empirical objective or selected optimizer, and supplies no boundedness event, probability inequality, error rate, countability or separability assumption. The existing pilot accuracy and envelope assumptions remain unchanged. Use the following well-definedness convention for the likelihood-transform critical-radius and concentration arguments in Assumption 3.4 and Lemma B.1, separately from the existing hypothesis-class, density-image and fitted-process conventions. Let  $Q$  be the true joint law of  $(X, Z)$ , let  $\mathcal{G}_0$  be the likelihood-transform class already defined in Assumption 3.4, and put

$$\mathcal{A}_r = \{f \in \text{star}(\mathcal{G}_0 - \mathcal{G}_0) : \|f\|_{L_2(Q)} \leq r\}.$$

For every integer  $n \geq 1$  and every  $r > 0$ , the three processes

$$\sup_{f \in \mathcal{G}_0} |(Q_n - Q)f|, \quad \sup_{f \in \mathcal{A}_r} |(Q_n - Q)f|, \quad \sup_{f \in \mathcal{A}_r} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i f(X_i, Z_i) \right|$$

regarded as extended nonnegative quantities, are measurable under  $Q^n$  for the first two and under the product of  $Q^n$  with the independent Rademacher law for the third. The functions use the same fixed density representatives as the likelihood objective. This convention only makes these empirical processes and their expectations well defined. It supplies no concentration event, deviation inequality, probability bound, rate, countability, separability, or additional statistical restriction.

*Proof.* By Lemmas B.7 and B.8 and the squared triangle inequality,

$$\|\mathcal{T}(\hat{h}_{t,\lambda} - h_0)\|^2 \leq C \frac{a_t \delta_n^2}{\lambda} + C \|w_0\|^2 \lambda^{q_t}.$$

The calibration comparison (B.13) applies to the final fitted candidate at each grid point. Young's inequality absorbs its linear term into the displayed weak error and  $\nu_n^2$ . Enlarging the fixed constants if necessary gives (3.16).  $\square$

**Lemma B.10** (Iterative RDIV error bounds). *Under the assumptions of Theorem 3.9, for every  $\lambda \in \Lambda_n \cap [\lambda_{\text{IR}}^\circ, 2]$ , on  $\mathcal{E}_{n,t}^{\text{IR}}$ ,*

$$\begin{aligned} \|\mathcal{T}(\hat{h}_{t,\lambda} - h_0)\|^2 &\leq C_t \{D_n(\hat{h}_{t,\lambda}) + A_{n,t}^{\text{R}}(\lambda)\}, \\ \|\hat{h}_{t,\lambda} - h_0\|^2 &\leq C_t \left[ V_{n,t}^{\text{R}}(\lambda) + \{D_n(\hat{h}_{t,\lambda}) + A_{n,t}^{\text{R}}(\lambda)\}^{\beta/(\beta+1)} \right], \end{aligned}$$

where the noise level and stability bound are defined in (3.14).

*Additional conditions.* For the conditional fitting argument in Lemma B.2, use the following measurable-version convention. Let  $Q$  denote the law of one observation  $W$ , let  $a$  denote the preliminary observations, and put  $Q_n^{\mathbf{w}} = n^{-1} \sum_{i=1}^n \delta_{w_i}$  for a fresh fitting tuple  $\mathbf{w} = (w_1, \dots, w_n)$ . Define

$$\ell_{\lambda,a}(h; w) = \{Y(w) - (\hat{\mathcal{T}}_a h)(Z(w))\}^2 + \lambda h(X(w))^2, \quad S_{n,\lambda,q}(a, \mathbf{w}) = \sup_{\substack{h \in \mathcal{H} \\ \|h - h_0\| \leq q}} |(Q_n^{\mathbf{w}} - Q)\{\ell_{\lambda,a}(h) - \ell_{\lambda,a}(h_0)\}|$$

For each  $n \geq 1$ ,  $\lambda \in \Lambda_n$  and positive rational  $q$ ,  $S_{n,\lambda,q}$  admits an extended nonnegative version jointly measurable in the preliminary observations and the fresh fitting tuple. This version agrees with  $S_{n,\lambda,q}(a, \cdot)$  for  $Q^n$ -almost every tuple for almost every preliminary realization, and agrees almost surely with  $S_{n,\lambda,q}(I_{\text{pre}}, (W_i)_{i \in I_{\text{fit}}})$  at the actual independent fitting block. All losses use the same function representatives as the empirical objective. This is only a measurability and evaluation-coherence convention; no deviation bound, rate, countability or separability is assumed. For the fitted conditional density in the RDIV preliminary construction, use a coherent measurable version at independent calibration observations. For each  $n \geq 1$ , writing  $a = I_{\text{pre}}$ , the selected density  $\hat{g}_{n,a}(x \mid z)$  admits a nonnegative version  $\bar{g}_n(a, x, z)$

jointly measurable in the preliminary observations,  $x$  and  $z$ . For almost every preliminary realization  $a$ , for  $P_Z$ -almost every  $z$ , the functions  $\bar{g}_n(a, \cdot, z)$  and  $\hat{g}_{n,a}(\cdot | z)$  agree  $\mu$ -almost everywhere. Almost surely, this same agreement holds at every actual calibration observation  $Z_i$ ,  $i \in I_{\text{cal}}$ . The empirical objectives and selected maximizer remain those defined using the original supplied representatives: this convention does not replace the optimizer. It supplies only measurability and coherence of the fitted density evaluated on the independent calibration sample; it assumes no normalization event, probability inequality, rate, countability or separability. For the preliminary regression estimate used in the RDIV calibration loss, take a coherent measurable version on the independent calibration sample. For each  $n \geq 1$ , writing  $a = I_{\text{pre}}$ , the selected  $L_2(Z)$  pilot has a version  $\bar{r}_n(a, z)$  jointly measurable in the preliminary observations and  $z$ . For almost every preliminary realization  $a$ ,  $\bar{r}_n(a, z) = \hat{r}_{n,a}(z)$  for  $P_Z$ -almost every  $z$ . Almost surely, the same equality holds simultaneously at every actual calibration observation  $Z_i$ ,  $i \in I_{\text{cal}}$ , for the supplied representatives used in  $D_n$ . This is only a measurability and evaluation-coherence convention. It imposes no exact joint measurability on every supplied representative, does not replace an empirical objective or selected optimizer, and supplies no boundedness event, probability inequality, error rate, countability or separability assumption. The existing pilot accuracy and envelope assumptions remain unchanged. Use the following well-definedness convention for the likelihood-transform critical-radius and concentration arguments in Assumption 3.4 and Lemma B.1, separately from the existing hypothesis-class, density-image and fitted-process conventions. Let  $Q$  be the true joint law of  $(X, Z)$ , let  $\mathcal{G}_0$  be the likelihood-transform class already defined in Assumption 3.4, and put

$$\mathcal{A}_r = \{f \in \text{star}(\mathcal{G}_0 - \mathcal{G}_0) : \|f\|_{L_2(Q)} \leq r\}.$$

For every integer  $n \geq 1$  and every  $r > 0$ , the three processes

$$\sup_{f \in \mathcal{G}_0} |(Q_n - Q)f|, \quad \sup_{f \in \mathcal{A}_r} |(Q_n - Q)f|, \quad \sup_{f \in \mathcal{A}_r} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i f(X_i, Z_i) \right|$$

regarded as extended nonnegative quantities, are measurable under  $Q^n$  for the first two and under the product of  $Q^n$  with the independent Rademacher law for the third. The functions use the same fixed density representatives as the likelihood objective. This convention only makes these empirical processes and their expectations well defined. It supplies no concentration event, deviation inequality, probability bound, rate, countability, separability, or additional statistical restriction.

*Proof.* Let  $\zeta = \|\mathcal{T}(\hat{h}_{t,\lambda} - h_0)\|$ . By (B.13),

$$\zeta^2 \leq D_n(\hat{h}_{t,\lambda}) + C\nu_n\zeta + C\nu_n^2,$$

and  $\nu_n^2 \leq A_{n,t}^R(\lambda)$ . Young's inequality proves the first bound. The same bound, the weak part of (B.11), and the squared triangle inequality give

$$\|\mathcal{T}(h_{t,\lambda}^* - h_0)\|^2 \leq C_t \{D_n(\hat{h}_{t,\lambda}) + A_{n,t}^R(\lambda)\}.$$

Apply (B.10) to this bound, combine the result with (B.12), and use the squared triangle inequality to obtain the second claim.  $\square$

*Proof of Proposition 3.8.* At the terminal point,  $D_n(\hat{h}_{t,\lambda_{K_n}}) \leq 4B^2$ , while

$$\frac{a_t \delta_n^2}{\lambda_{K_n}} \geq \frac{\delta_{n,\xi}^2}{\lambda_{K_n}} \geq n \log 2.$$

Thus the terminal candidate is accepted for  $n \geq N_R(B)$ . If  $h_0 = 0$ , every population iterate equals  $0 \in \mathcal{H}$ , so the discrepancy-bound argument at  $\lambda = 2$ , with  $w_0 = 0$ , gives initial acceptance on the corresponding concentration event. For  $h_0 \neq 0$ , the first accepted index is either zero or has a rejected predecessor in  $[\lambda, 2\lambda]$ . Each candidate requires  $t$  recursive fits, so the total number is at most  $t(K_n + 1)$ .  $\square$

*Proof of Theorem 3.9.* Write  $s = s_{n,t}^R$  and  $q = q_t$ . Boundedness gives the result directly when  $s > 1$  or  $\nu_n > 1$ , so consider  $s \vee \nu_n \leq 1$ . The discrepancy bound (3.16) gives (3.2) with the bounds in (3.14). Lemma B.10 establishes (3.3) on  $\mathcal{S}_n = [\lambda_{\text{IR}}^\circ, 2]$  with  $\gamma = \beta/(\beta + 1)$ , since  $D_n \geq 0$  and  $V_{n,t}^R(\lambda) \leq A_{n,t}^R(\lambda)/\lambda$ . For all sufficiently large  $n$ , the point  $\lambda_{\text{IR}}^\circ$  defined in the main text exists and satisfies (3.4). The noise level is nonincreasing. Applying Theorem 3.1 gives  $\lambda_{\text{dp},t}^R \geq \lambda_{\text{IR}}^\circ$  and the stated weak bound. For the strong bound, use the sharper method-specific bound in Lemma B.10. Acceptance and the monotonicity of  $A_{n,t}^R$  and  $V_{n,t}^R$  give

$$\begin{aligned} \|\hat{h}_{t,\lambda_{\text{dp},t}^R} - h_0\|^2 &\leq C_t \left[ V_{n,t}^R(\lambda_{\text{dp},t}^R) + \{D_n(\hat{h}_{t,\lambda_{\text{dp},t}^R}) + A_{n,t}^R(\lambda_{\text{dp},t}^R)\}^{\beta/(\beta+1)} \right] \\ &\leq C_t \left[ V_{n,t}^R(\lambda_{\text{IR}}^\circ) + A_{n,t}^R(\lambda_{\text{IR}}^\circ)^{\beta/(\beta+1)} \right]. \end{aligned}$$

Now (3.15) proves the stated strong bound. Boundedness covers the remaining finitely many  $n$ , after we enlarge the fixed constants. When  $2t \geq \beta + 1$ ,  $q = \beta + 1$ . Under  $\nu_n = O(\delta_n)$  and fixed  $t$ , the strong bound has order  $\delta_n^{2\beta/(\beta+2)}$ , and the weak bound has order  $\delta_n^{2(\beta+1)/(\beta+2)}$ .  $\square$

## C Proofs of Section 3.4

Throughout this section,  $L_n$  and  $L$  denote the TRAE losses,  $D_n = L_n$ , and  $\delta_n$  is the TRAE scale in (3.17). The finite grid  $\Lambda_n = \{2\rho^k : k = 0, \dots, K_n\}$  is defined in Algorithm 1, and  $\mathcal{E}_n$  is the common event defined after (3.17). All bounds below hold simultaneously on  $\mathcal{E}_n$ .

**Lemma C.1** (Uniform TRAE loss comparison). *Under Assumptions 3.10 to 3.12, on  $\mathcal{E}_n$ , the following inequalities hold simultaneously for all  $h \in \mathcal{H}$ :*

$$\frac{1}{2}L(h) - O(\delta_n^2) \leq L_n(h) \leq \frac{5}{4}L(h) + O(\delta_n^2).$$

*Additional condition. Interpret the uniform almost-sure boundedness in Assumption 3.12 using one common null set for the whole function classes: there exists a set  $N$  with  $\mathbb{P}(W \in N) = 0$  such that, for every  $w \notin N$ , every  $h \in \mathcal{H}$  and every  $f \in \mathcal{F}$ , all three inequalities  $|h(X(w))| \leq B$ ,  $|f(Z(w))| \leq B$  and  $|m(w; f)| \leq B$  hold simultaneously. The exceptional set does not depend on  $h$  or  $f$ . Apply this convention separately to each primal or exchanged problem whenever that problem is assumed to satisfy Assumption 3.12.*

*Additional condition.* For a result cited in the proof of Lemma C.1: Interpret the uniform almost-sure boundedness in Assumption 3.12 using one common null set for the whole function classes: there exists a set  $N$  with  $\mathbb{P}(W \in N) = 0$  such that, for every  $w \notin N$ , every  $h \in \mathcal{H}$  and every  $f \in \mathcal{F}$ , all three inequalities  $|h(X(w))| \leq B$ ,  $|f(Z(w))| \leq B$  and  $|m(w; f)| \leq B$  hold simultaneously. The exceptional set does not depend on  $h$  or  $f$ . Apply this convention separately to each primal or exchanged problem whenever that problem is assumed to satisfy Assumption 3.12.

*Proof.* Let  $f_h = \mathcal{T}(h_0 - h) \in \mathcal{F}$ . The localized concentration bound used in Bennett et al. [2023, Lemma 9] gives

$$\begin{aligned} L(h) &= \sup_{f \in \mathcal{F}} \mathbb{E} [2(m(W; f) - f(Z)h(X)) - f^2(Z)] \\ &= \mathbb{E} [2(m(W; f_h) - f_h(Z)h(X)) - f_h^2(Z)] \\ &\leq \mathbb{E}_n [2(m(W; f_h) - f_h(Z)h(X)) - f_h^2(Z)] + O(\delta_n^2 + \delta_n \|f_h\|) \\ &\leq \sup_{f \in \mathcal{F}} \mathbb{E}_n [2(m(W; f) - f(Z)h(X)) - f^2(Z)] + O(\delta_n^2 + \delta_n \|f_h\|) \\ &\leq L_n(h) + O(\delta_n^2 + \delta_n \|f_h\|). \end{aligned}$$

Since  $\|f_h\|^2 = L(h)$ , Young's inequality gives

$$L(h) \leq 2L_n(h) + O(\delta_n^2).$$

Similarly,

$$\begin{aligned}
L_n(h) &= \sup_{f \in \mathcal{F}} \mathbb{E}_n [2(m(W; f) - f(Z)h(X)) - f^2(Z)] \\
&\leq \sup_{f \in \mathcal{F}} \{ \mathbb{E} [2(m(W; f) - f(Z)h(X)) - f(Z)^2] + C(\delta_n^2 + \delta_n \|f\|) \} \\
&\leq \sup_{f \in \tilde{\mathcal{Q}}} \mathbb{E} \left[ 2(m(W; f) - f(Z)h(X)) - \frac{4}{5}f^2(Z) \right] + O(\delta_n^2) \\
&= \frac{5}{4} \|\mathcal{T}(h_0 - h)\|^2 + O(\delta_n^2) \\
&= \frac{5}{4} L(h) + O(\delta_n^2).
\end{aligned}$$

Here Young's inequality bounds  $C\delta_n\|f\|$  by  $\|f\|^2/5 + O(\delta_n^2)$ . We then enlarge the supremum from  $\mathcal{F}$  to  $\tilde{\mathcal{Q}}$  and complete the square. The feasible residual  $f_h$  gives the lower bound, and the enlarged supremum gives the upper bound under Assumption 3.10.  $\square$

**Lemma C.2** (Direct TRAE control). *Under the assumptions of Theorem 3.14, on  $\mathcal{E}_n$ , the following bound holds simultaneously for all  $\lambda \in \Lambda_n$ :*

$$L_n(\hat{h}_\lambda) + \frac{\lambda}{2} \|\hat{h}_\lambda - h_0\|^2 \leq C\{\delta_n^2 + \|w_0\|^2 \lambda^p\}, \quad p = \min\{\beta + 1, 2\}.$$

Moreover, for all  $\lambda \in \Lambda_n$ ,

$$\|\hat{h}_\lambda - h_0\|^2 \leq C \left[ \frac{\delta_n^2}{\lambda} + \{L_n(\hat{h}_\lambda) + \delta_n^2\}^\alpha \right], \quad \alpha = \frac{\min\{\beta, 1\}}{1 + \min\{\beta, 1\}}.$$

*Additional condition. Interpret the uniform almost-sure boundedness in Assumption 3.12 using one common null set for the whole function classes: there exists a set  $N$  with  $\mathbb{P}(W \in N) = 0$  such that, for every  $w \notin N$ , every  $h \in \mathcal{H}$  and every  $f \in \mathcal{F}$ , all three inequalities  $|h(X(w))| \leq B$ ,  $|f(Z(w))| \leq B$  and  $|m(w; f)| \leq B$  hold simultaneously. The exceptional set does not depend on  $h$  or  $f$ . Apply this convention separately to each primal or exchanged problem whenever that problem is assumed to satisfy Assumption 3.12.*

*Additional condition.* For a result cited in the proof of Lemma C.2: Interpret the uniform almost-sure boundedness in Assumption 3.12 using one common null set for the whole function classes: there exists a set  $N$  with  $\mathbb{P}(W \in N) = 0$  such that, for every  $w \notin N$ , every  $h \in \mathcal{H}$  and every  $f \in \mathcal{F}$ , all three inequalities  $|h(X(w))| \leq B$ ,  $|f(Z(w))| \leq B$  and  $|m(w; f)| \leq B$  hold simultaneously. The exceptional set does not depend on  $h$  or  $f$ . Apply this convention separately to each primal or exchanged problem whenever that problem is assumed to satisfy Assumption 3.12.

*Proof.* Write  $e_\lambda = \widehat{h}_\lambda - h_0$ . The localized squared-norm inequality in Foster and Syrgkanis [2023, Lemma 14], applied using the squared-loss measurability condition in Assumption 3.12, gives, uniformly over  $h \in \mathcal{H}$ ,

$$|\{\|h\|_{2,n}^2 - \|h_0\|_{2,n}^2\} - \{\|h\|^2 - \|h_0\|^2\}| \leq C\{\delta_n\|h - h_0\| + \delta_n^2\}.$$

<sup>1</sup> The event  $\mathcal{E}_n$  includes this bound and  $L_n(h_0) \leq C\delta_n^2$ . Empirical optimality relative to

---

<sup>1</sup>Put  $M = \max\{1, B\}$ ,  $\mathcal{U} = \{h/M : h \in \mathcal{H}\}$ , and  $u_0 = h_0/M$ . Assumption 3.2 gives  $u_0 \in \mathcal{U}$ , and boundedness gives  $|u(X)| \leq 1$  uniformly over  $u \in \mathcal{U}$  on the common full-probability set. The loss

$$\ell(v) = M^2(\max\{-1, \min\{v, 1\}\})^2$$

is  $L$ -Lipschitz with  $L = 2M^2$ , since clipping is 1-Lipschitz and  $|a^2 - b^2| \leq 2|a - b|$  for  $a, b \in [-1, 1]$ . Moreover,  $\ell(h(X)/M) = h(X)^2$ .

Set  $\mathcal{S} = \text{star}(\mathcal{H} - \mathcal{H})$ . By Assumption 3.12 and (3.17), there is a deterministic critical radius  $d \leq \delta_n$  for  $\mathcal{S}$ . Since  $0 \in \mathcal{S}$ , we have  $\mathcal{S} \subseteq \text{star}(\mathcal{S} - \mathcal{S})$ , so the critical-radius definition implies  $\mathfrak{R}_n(d; \mathcal{S}) \leq d^2$ . For  $s \geq d$ , scaling by  $d/s$  preserves membership in  $\mathcal{S}$  and gives

$$\mathfrak{R}_n(s; \mathcal{S}) \leq \frac{s}{d}\mathfrak{R}_n(d; \mathcal{S}) \leq sd.$$

Also,  $h_0 \in \mathcal{H}$  implies  $\text{star}(\mathcal{H} - h_0) \subseteq \mathcal{S}$ . Thus, for every  $s \geq \delta_n$ ,

$$\begin{aligned} \mathfrak{R}_n(s; \text{star}(\mathcal{U} - u_0)) &= M^{-1}\mathfrak{R}_n(Ms; \text{star}(\mathcal{H} - h_0)) \\ &\leq M^{-1}\mathfrak{R}_n(Ms; \mathcal{S}) \leq sd \leq s\delta_n. \end{aligned}$$

In particular, every  $r \geq \delta_n$  satisfies the anchored fixed-point inequality for the normalized class.

For  $R > 0$ , define

$$Z_n(R) = \sup_{\substack{u \in \mathcal{U} \\ \|u - u_0\| \leq R}} |(\mathbb{E}_n - \mathbb{E})[\ell(u(X)) - \ell(u_0(X))]|.$$

Its empirical and Rademacher suprema are precisely the squared-loss suprema in Assumption 3.12, with localization radius  $MR$ , and hence are measurable. The localized concentration inequality used in the proof of Foster and Syrgkanis [2023, Lemma 14], applied to this measurable loss process and the anchored fixed point, gives universal constants  $c_1, c_2 > 0$  such that, for  $R \geq r \geq \delta_n$  and  $v > 0$ ,

$$\mathbb{P}\{Z_n(R) \geq 16LRr + v\} \leq c_1 \exp\left\{-\frac{c_2nv^2}{L^2R^2 + 2Lv}\right\}.$$

Taking  $v = LRr$  and using  $r \leq R$  gives

$$\mathbb{P}\{Z_n(R) \geq 17LRr\} \leq c_1 \exp\{-c_2nr^2/3\}.$$

Choose  $r = a\delta_n$  with a fixed sufficiently large  $a \geq 1$ . If  $r \geq 2$ , set  $J = 0$ ; otherwise set  $J = \lceil \log(2/r)/\log 2 \rceil$ . Put  $R_j = 2^j r$  for  $j = 0, \dots, J$ . A union bound shows that the measurable event

$h_0 \in \mathcal{H}$  therefore gives

$$L_n(\widehat{h}_\lambda) + \lambda\{\|e_\lambda\|^2 + 2\langle h_0, e_\lambda \rangle\} \leq C\{\delta_n^2 + \lambda\delta_n\|e_\lambda\|\}. \quad (\text{C.1})$$

By Lemma C.1,  $L(\widehat{h}_\lambda) \leq 2L_n(\widehat{h}_\lambda) + C\delta_n^2$ . Apply the first consequence of Lemma A.1 to the inner-product term in (C.1). Young's inequality also gives  $C\lambda\delta_n\|e_\lambda\| \leq \lambda\|e_\lambda\|^2/8 + C\delta_n^2$ . Absorbing the resulting empirical-loss and squared-norm terms proves the first bound.

For the second bound, the lower comparison in Lemma C.1 gives

$$L(\widehat{h}_\lambda) \leq 2L_n(\widehat{h}_\lambda) + C\delta_n^2 \leq C\{L_n(\widehat{h}_\lambda) + \delta_n^2\}.$$

By Assumptions 3.2 and 3.10, taking  $h = h_0$  gives  $0 = \mathcal{T}(h_0 - h_0) \in \mathcal{F}$ . Hence  $L_n(\widehat{h}_\lambda) \geq 0$ . Drop this term from (C.1), divide by  $\lambda$ , and apply Young's inequality to obtain

$$\|e_\lambda\|^2 \leq C\frac{\delta_n^2}{\lambda} + C|\langle h_0, e_\lambda \rangle|.$$

The second consequence of Lemma A.1 and the weak bound above complete the proof.  $\square$

**Corollary C.3** (TRAЕ discrepancy bound). *Under the assumptions of Theorem 3.14, the calibrated TRAЕ scale satisfies (3.19) simultaneously for every  $\lambda \in \Lambda_n$  on  $\mathcal{E}_n$ , for a fixed constant  $C_T > 0$ .*

$$Z_n(R_j) < 17LR_jr, \quad j = 0, \dots, J,$$

has failure probability at most  $(J+1)c_1 \exp\{-c_2nr^2/3\}$ . Since  $r \geq \delta_n \geq n^{-1/2}$ , there is a universal constant  $C_0$  such that  $J+1 \leq C_0 \log(en)$ . Equation (3.17) gives  $n\delta_n^2 \geq 1$  and a fixed positive lower bound for

$$\frac{n\delta_n^2}{\log \log(en) + \log(1/\eta_n)}.$$

Consequently, a fixed sufficiently large  $a$  makes this failure probability at most  $\eta_n/4$ .

On this event, let  $u \in \mathcal{U}$  and write  $b = \|u - u_0\|$ . If  $b \leq r$ , the loss increment is bounded by  $Z_n(r) < 17Lr^2$ . If  $b > r$ , boundedness gives  $b \leq 2$ , so there is an index  $j \in \{1, \dots, J\}$  with  $R_{j-1} < b \leq R_j$ . Then  $R_j < 2b$ , and the loss increment is bounded by  $Z_n(R_j) < 34Lrb$ . Thus, in both cases,

$$|(\mathbb{E}_n - \mathbb{E})[\ell(u(X)) - \ell(u_0(X))]| \leq 34Lr(b+r).$$

Substitution of  $u = h/M$ ,  $L = 2M^2$ , and  $r = a\delta_n$  bounds the squared-norm difference by

$$68M^2a\delta_n\{\|h - h_0\|/M + a\delta_n\} \leq C\{\delta_n\|h - h_0\| + \delta_n^2\},$$

uniformly over  $h \in \mathcal{H}$ . The constant depends only on the fixed boundedness and calibration constants. Allocating  $\eta_n/4$  to this comparison at each grid point uses at most  $(K_n + 1)\eta_n/4 = \xi/4$  of the total failure probability, supplying the squared-norm comparison included in  $\mathcal{E}_n$ .

Additional condition. Interpret the uniform almost-sure boundedness in Assumption 3.12 using one common null set for the whole function classes: there exists a set  $N$  with  $\mathbb{P}(W \in N) = 0$  such that, for every  $w \notin N$ , every  $h \in \mathcal{H}$  and every  $f \in \mathcal{F}$ , all three inequalities  $|h(X(w))| \leq B$ ,  $|f(Z(w))| \leq B$  and  $|m(w; f)| \leq B$  hold simultaneously. The exceptional set does not depend on  $h$  or  $f$ . Apply this convention separately to each primal or exchanged problem whenever that problem is assumed to satisfy Assumption 3.12.

*Proof.* Because  $L(h_0) = 0$ , Lemma C.1 gives  $L_n(h_0) \leq C_0 \delta_n^2$ . The first bound in Lemma C.2 gives

$$L_n(\hat{h}_\lambda) \leq a_T \delta_n^2 + C_T \|w_0\|^2 \lambda^{\min\{\beta+1, 2\}}$$

for fixed constants  $a_T, C_T > 0$ . The constants  $C_0$  and  $a_T$  depend only on the prescribed boundedness and continuity constants, independently of  $\beta$  and  $\|w_0\|$ . Multiplying  $\delta_n$  by  $\sqrt{\max\{C_0, a_T, 1\}}$  and relabeling the resulting upper bound gives (3.19).  $\square$

The next lemma bounds the weak and strong errors at every grid point.

**Lemma C.4** (TRAЕ error bounds). *Suppose that the assumptions of Theorem 3.14 hold. For every  $\lambda \in \Lambda_n$ , on  $\mathcal{E}_n$ ,*

$$\begin{aligned} \|\mathcal{T}(\hat{h}_\lambda - h_0)\|^2 &\leq C \{L_n(\hat{h}_\lambda) + \delta_n^2\}, \\ \|\hat{h}_\lambda - h_0\|^2 &\leq C \left[ V_n^T(\lambda) + \{L_n(\hat{h}_\lambda) + \delta_n^2\}^\alpha \right], \end{aligned}$$

where the noise level and stability bound are defined in (3.18) and  $\alpha = \min\{\beta, 1\}/\{1 + \min\{\beta, 1\}\}$ .

Additional condition. Interpret the uniform almost-sure boundedness in Assumption 3.12 using one common null set for the whole function classes: there exists a set  $N$  with  $\mathbb{P}(W \in N) = 0$  such that, for every  $w \notin N$ , every  $h \in \mathcal{H}$  and every  $f \in \mathcal{F}$ , all three inequalities  $|h(X(w))| \leq B$ ,  $|f(Z(w))| \leq B$  and  $|m(w; f)| \leq B$  hold simultaneously. The exceptional set does not depend on  $h$  or  $f$ . Apply this convention separately to each primal or exchanged problem whenever that problem is assumed to satisfy Assumption 3.12.

*Proof.* Lemma C.1 gives

$$\|\mathcal{T}(\hat{h}_\lambda - h_0)\|^2 = L(\hat{h}_\lambda) \leq 2L_n(\hat{h}_\lambda) + O(\delta_n^2) \leq C \{L_n(\hat{h}_\lambda) + \delta_n^2\}.$$

The second conclusion of Lemma C.2 gives

$$\|\hat{h}_\lambda - h_0\|^2 \leq C \left[ \frac{\delta_n^2}{\lambda} + \{L_n(\hat{h}_\lambda) + \delta_n^2\}^\alpha \right] = C \left[ V_n^T(\lambda) + \{L_n(\hat{h}_\lambda) + \delta_n^2\}^\alpha \right].$$

$\square$

We first establish termination, then combine the discrepancy and error bounds over the grid to prove the rates.

*Proof of Proposition 3.13.* Throughout the proof, we work on  $\mathcal{E}_n$ . Empirical optimality and  $h_0 \in \mathcal{H}$  give

$$L_n(\widehat{h}_{\lambda_{K_n}}) \leq L_n(h_0) + \lambda_{K_n} \|h_0\|_{2,n}^2 \leq \delta_n^2 + B^2 n^{-2}.$$

Here we used (3.19),  $\lambda_{K_n} \leq n^{-2}$ , and the bound  $|h_0| \leq B$ . Since  $\delta_n^2 \geq n^{-1}$  and  $n \geq N_T(B)$ ,

$$B^2 n^{-2} \leq \frac{1}{2} n^{-1} \leq \frac{1}{2} \delta_n^2.$$

Thus the terminal point passes the threshold, and the algorithm returns an estimator after at most  $K_n + 1 = O(\log n)$  optimization problems.

Let the selected index be

$$k^* := \min\{0 \leq k \leq K_n : L_n(\widehat{h}_{2\rho^k}) \leq \tfrac{3}{2}\delta_n^2\},$$

whose existence follows from the preceding argument. If  $k^* = 0$ , the algorithm returns through the initial-acceptance branch. If  $k^* \geq 1$ , its minimality gives

$$L_n(\widehat{h}_{2\rho^{k^*}}) \leq \tfrac{3}{2}\delta_n^2 < L_n(\widehat{h}_{2\rho^{k^*-1}}).$$

Because  $1/2 \leq \rho < 1$ , the companion parameter  $2\rho^{k^*-1} = \lambda_{\text{dp}}/\rho$  belongs to  $[\lambda_{\text{dp}}, 2\lambda_{\text{dp}}]$ . This argument covers both zero and nonzero signals. □

Theorem 3.1 now combines the discrepancy bound with the weak- and strong-error bounds at every grid point. The same argument is uniform over a fixed source ball  $\{(\mathcal{T}^*\mathcal{T})^{\beta/2}w : \|w\| \leq \varrho\}$  when the realizability conditions and the concentration event  $\mathcal{E}_n$  hold uniformly over that ball; the stated theorem is pointwise in  $h_0$ .

*Proof of Theorem 3.14.* Throughout the proof, we work on  $\mathcal{E}_n$  and write  $p = \min\{\beta + 1, 2\}$ . If  $\delta_n > 1$ , boundedness and the contraction property of  $\mathcal{T}$  give both stated rates directly. Suppose  $\delta_n \leq 1$ . The discrepancy bound (3.19) gives (3.2) with  $q = p$  and the bounds in (3.18). Lemma C.4 establishes (3.3) on  $\mathcal{S}_n = (0, 2]$  with  $\gamma = (p - 1)/p$ , since  $D_n = L_n \geq 0$  and  $V_n^T(\lambda) = A_n^T(\lambda)/\lambda$ . For all sufficiently large  $n$ , the point  $\lambda_T^\circ$  defined in the main text exists and satisfies (3.4); the noise level is nonincreasing. Apply Theorem 3.1 and (3.20). The strong bound has order  $\delta_n^{2(p-1)/p}$ , which is the stated rate because  $p - 1 = \min\{\beta, 1\}$ . Boundedness covers the remaining finitely many  $n$ , after we enlarge the fixed constants. □

## C.1 Iterative TRAE

For this subsection,  $\mathcal{E}_{n,t}^{\text{IT}}$  denotes the common concentration event formed with probability allocation  $\eta_{n,t} = \xi / \{t(K_n + 1)\}$ . The additional class star( $\mathcal{H} - \mathcal{H}$ ) in Assumption 3.12 controls the random centers in the iterative penalties.

*Remark (notation).*  $\mathcal{E}_{n,t}^{\text{IT}}$  is the common iterative TRAE concentration event obtained by intersecting the requisite iteration-grid concentration events with failure-probability allocation  $\eta_{n,t} = \xi / \{t(K_n + 1)\}$ , with probability at least  $1 - \xi$  under the paper's confidence convention. In Lemma C.2, Corollary C.6, Lemma C.7: The concentration comparisons are simultaneous over the oracle-relevant grid segment  $\Lambda_n \cap [\lambda_{\text{IT}}^\circ, 2]$ , including the recursive comparisons for  $j = 1, \dots, t$ .

**Lemma C.5** (Iterative TRAE propagation). *Under the assumptions of Theorem 3.18, there is a nondecreasing deterministic sequence with  $b_0 = 0$  and  $b_j \geq 2$ ,  $j = 1, \dots, t$ , where each  $b_j$  depends only on  $j$  and the prescribed boundedness and mean-square-continuity constants, independently of  $\beta$  and  $\|w_0\|$ . This sequence may be calibrated so that, on  $\mathcal{E}_{n,t}^{\text{IT}}$ , simultaneously for  $\lambda \in \Lambda_n \cap [\lambda_{\text{IT}}^\circ, 2]$ ,*

$$\|\widehat{h}_{j,\lambda} - h_{j,\lambda}^*\|^2 \leq C \frac{b_j \delta_n^2}{\lambda^2}, \quad j = 1, \dots, t, \quad (\text{C.2})$$

$$L(\widehat{h}_{t,\lambda}) \leq C \left[ \delta_n^2 + \min \left\{ \lambda, \frac{b_t \delta_n^2}{\lambda} \right\} + \|w_0\|^2 \lambda^{q_t} \right]. \quad (\text{C.3})$$

*Additional conditions. Interpret the uniform almost-sure boundedness in Assumption 3.12 using one common null set for the whole function classes: there exists a set  $N$  with  $\mathbb{P}(W \in N) = 0$  such that, for every  $w \notin N$ , every  $h \in \mathcal{H}$  and every  $f \in \mathcal{F}$ , all three inequalities  $|h(X(w))| \leq B$ ,  $|f(Z(w))| \leq B$  and  $|m(w; f)| \leq B$  hold simultaneously. The exceptional set does not depend on  $h$  or  $f$ . Apply this convention separately to each primal or exchanged problem whenever that problem is assumed to satisfy Assumption 3.12. With  $\mathcal{U}$  and  $\ell_B$  as in Assumption 3.12, for every fixed iteration-grid pair  $(j, \lambda)$  under consideration and every  $r > 0$ , both  $\sup_{u \in \mathcal{U}, \|u - u_{j,\lambda}^*\|_{L_2(W; \mathbb{R}^3)} \leq r} |(\mathbb{E}_n - \mathbb{E})\{\ell_B(u(W)) - \ell_B(u_{j,\lambda}^*(W))\}|$  and  $\sup_{u \in \mathcal{U}, \|u - u_{j,\lambda}^*\|_{L_2(W; \mathbb{R}^3)} \leq r} |n^{-1} \sum_{i=1}^n \epsilon_i \{\ell_B(u(W_i)) - \ell_B(u_{j,\lambda}^*(W_i))\}|$  are measurable under the fitting-sample law and its product with the independent Rademacher law, respectively, where  $f_j^* = \mathcal{T}(h_0 - h_{j,\lambda}^*)$  and  $u_{j,\lambda}^*(W) = (m(W; f_j^*), h_{j,\lambda}^*(X) f_j^*(Z), f_j^*(Z))$ .*

*Additional condition.* For another result of the paper: Interpret the uniform almost-sure boundedness in Assumption 3.12 using one common null set for the whole function classes: there exists a set  $N$  with  $\mathbb{P}(W \in N) = 0$  such that, for every  $w \notin N$ , every  $h \in \mathcal{H}$  and every  $f \in \mathcal{F}$ , all three inequalities  $|h(X(w))| \leq B$ ,  $|f(Z(w))| \leq B$  and  $|m(w; f)| \leq B$  hold

simultaneously. The exceptional set does not depend on  $h$  or  $f$ . Apply this convention separately to each primal or exchanged problem whenever that problem is assumed to satisfy Assumption 3.12.

*Remark (notation).*  $\{b_j\}_{j=0}^t$  is a nondecreasing deterministic propagation sequence with  $b_0 = 0$  and  $b_j \geq 2$  for  $1 \leq j \leq t$ , each positive-index term depending only on  $j$  and the fixed boundedness and mean-square-continuity constants, independently of  $\beta$  and  $\|w_0\|$ .

*Proof.* We adapt the iterative argument of Bennett et al. [2023, Theorem 6] to the finite grid and give the recursive comparison explicitly. Fix  $j$  and  $\lambda$ , and set

$$e_j = \|\widehat{h}_{j,\lambda} - h_{j,\lambda}^*\|, \quad u_j = \|\mathcal{T}(\widehat{h}_{j,\lambda} - h_{j,\lambda}^*)\|, \quad f_j^* = \mathcal{T}(h_0 - h_{j,\lambda}^*).$$

Also define

$$M_{\leq j} = \max_{1 \leq \ell \leq j} \{\lambda, \|f_\ell^*\|_\infty^2\}.$$

By Assumptions 3.10 and 3.16,  $f_\ell^* \in \mathcal{F}$ . The boundedness part of Assumption 3.12 therefore gives  $M_{\leq j} \leq \max\{2, B^2\}$  for every grid point under consideration.

To track the random penalty center, write

$$\begin{aligned} I_{n,j} &= \|h_{j,\lambda}^* - \widehat{h}_{j-1,\lambda}\|_{2,n}^2 - \|\widehat{h}_{j,\lambda} - \widehat{h}_{j-1,\lambda}\|_{2,n}^2, \\ I_j &= \|h_{j,\lambda}^* - h_{j-1,\lambda}^*\|^2 - \|\widehat{h}_{j,\lambda} - h_{j-1,\lambda}^*\|^2. \end{aligned}$$

We derive the excess-loss comparison directly. Write  $g = h_{j,\lambda}^*$ ,  $k = f_j^*$ ,  $\bar{f} = \mathcal{T}(h_0 - \widehat{h}_{j,\lambda})$ , and

$$\ell(h, f) = 2m(W; f) - 2h(X)f(Z) - f(Z)^2.$$

By Assumptions 3.10 and 3.16,  $k, \bar{f} \in \mathcal{F}$ . For every  $h \in \mathcal{H}$  and  $f \in \mathcal{F}$ , the Riesz identity in Section 2 and the definition of the projected operator give

$$\mathbb{E}[\ell(h, f)] = 2\langle \mathcal{T}(h_0 - h), f \rangle - \|f\|^2 = \|\mathcal{T}(h_0 - h)\|^2 - \|f - \mathcal{T}(h_0 - h)\|^2.$$

Indeed,  $f \in \widetilde{\mathcal{Q}}$  implies  $\mathbb{E}[h(X)f(Z)] = \langle \mathcal{T}h, f \rangle$ , and  $\mathbb{E}[m(W; f)] = \langle r_0, f \rangle = \langle \mathcal{T}h_0, f \rangle$ .

Apply the localized concentration inequality of Bennett et al. [2023, Lemma 9] to the product and moment-image classes in Assumption 3.12, and apply its bounded Lipschitz version to the squared test functions. With the deterministic anchors  $g$  and  $k$ , these inequalities give, uniformly over  $h \in \mathcal{H}$  and  $f \in \mathcal{F}$ ,

$$\begin{aligned} &|(\mathbb{E}_n - \mathbb{E})[\ell(h, f) - \ell(g, k)]| \\ &\leq C\delta_n \left\{ \left( \mathbb{E}[m(W; f - k)^2] \right)^{1/2} + \|h(X)f(Z) - g(X)k(Z)\|_{L_2(W)} + \|f - k\| \right\} + C\delta_n^2 \\ &\leq C \left\{ \delta_n \|h - g\| \|k\|_\infty + \delta_n \|f - k\| + \delta_n^2 \right\}. \end{aligned}$$

For the second inequality, Assumption 3.15 bounds the first norm by  $C\|f - k\|$ , while

$$h(X)f(Z) - g(X)k(Z) = h(X)(f - k)(Z) + (h - g)(X)k(Z)$$

implies that the product norm is at most  $B\|f - k\| + \|h - g\|\|k\|_\infty$ . The square map is  $2B$ -Lipschitz on the range of the test functions. Also,  $h_0 \in \mathcal{H}$  and Assumption 3.10 imply  $0 \in \mathcal{F}$ . Applying the same concentration inequality with zero moment, product, and test-function anchors, and using Assumption 3.11, gives

$$|(\mathbb{E}_n - \mathbb{E})\ell(h, f)| \leq C\{\delta_n\|f\| + \delta_n^2\}, \quad h \in \mathcal{H}, f \in \mathcal{F}.$$

Allocate fixed fractions of  $\eta_{n,t}$  to these finitely many concentration inequalities and to the comparison of Lemma C.1 used below, absorbing the resulting fixed constants into the calibrated scale.

Taking  $h = g$  in the increment bound and using the population identity yields

$$\begin{aligned} \mathbb{E}_n[\ell(g, f) - \ell(g, k)] &\leq -\|f - k\|^2 + C\delta_n\|f - k\| + C\delta_n^2 \\ &\leq C\delta_n^2. \end{aligned}$$

Taking the supremum over  $f \in \mathcal{F}$  therefore proves

$$0 \leq L_n(g) - \mathbb{E}_n\ell(g, k) \leq C\delta_n^2,$$

without requiring an empirical maximizing test function. Next,  $\|\bar{f} - k\| = u_j$ , so the increment bound at  $(\hat{h}_{j,\lambda}, \bar{f})$  gives

$$\begin{aligned} L(\hat{h}_{j,\lambda}) - L(g) &\leq \mathbb{E}_n[\ell(\hat{h}_{j,\lambda}, \bar{f}) - \ell(g, k)] + C\{\delta_n e_j\|k\|_\infty + \delta_n u_j + \delta_n^2\} \\ &\leq L_n(\hat{h}_{j,\lambda}) - L_n(g) + C\{\delta_n e_j\|k\|_\infty + \delta_n u_j + \delta_n^2\}. \end{aligned}$$

Empirical optimality relative to  $g \in \mathcal{H}$  gives  $L_n(\hat{h}_{j,\lambda}) - L_n(g) \leq \lambda I_{n,j}$ , including when the penalty center is random. Consequently,

$$L(\hat{h}_{j,\lambda}) - L(h_{j,\lambda}^*) \leq \lambda I_{n,j} + C\{\delta_n e_j\|f_j^*\|_\infty + \delta_n u_j + \delta_n^2\}. \quad (\text{C.4})$$

This comparison uses the localized bounds for  $\text{star}(\mathcal{H} \cdot \mathcal{F})$ ,  $\text{star}(m \circ \mathcal{F})$ , and  $\text{star}(\mathcal{F})$  from Assumption 3.12, together with Assumption 3.15. In particular, that bound controls  $\mathbb{E}[m(W; \hat{f}_{\hat{h}_{j,\lambda}} - f_j^*)^2]$  by  $Cu_j^2$ , where  $f_h = \mathcal{T}(h_0 - h)$ . The population normal equation for  $h_{j,\lambda}^*$  gives the exact quadratic lower bound

$$\lambda e_j^2 + u_j^2 \leq L(\hat{h}_{j,\lambda}) - L(h_{j,\lambda}^*) - \lambda I_j. \quad (\text{C.5})$$

The quadratic identity for (3.12), together with  $L(h) = \|\mathcal{T}(h - h_0)\|^2$  on  $\mathcal{H}$ , proves this bound.

Expanding the squared penalties separates the center error as  $2\langle \widehat{h}_{j,\lambda} - h_{j,\lambda}^*, \widehat{h}_{j-1,\lambda} - h_{j-1,\lambda}^* \rangle_n$ . For the fixed-center squared-loss comparisons, apply the same localized concentration inequality of Bennett et al. [2023, Lemma 9] to the bounded clipped losses corresponding to  $(h - h_{j-1,\lambda}^*)^2$  at comparator  $h_{j,\lambda}^*$  and  $(h - h_{k,\lambda}^*)^2$  at comparator  $h_{k,\lambda}^*$ , for the fixed iteration-grid pairs under consideration. The fixed-population-center squared-loss measurability condition in Assumption 3.12 supplies measurable localized empirical and Rademacher suprema for these applications.

Here is the scalar specialization and the penalty expansion. Put  $M = \max\{1, B\}$  and  $\mathcal{U} = \mathcal{H}/M$ . For each deterministic center  $c = h_{j-1,\lambda}^*$  or  $c = h_{k,\lambda}^*$ , use the loss

$$\ell_c(v; X) = \{M \max\{-1, \min\{v, 1\}\} - c(X)\}^2.$$

On the common full-probability set,  $|c(X)| \leq M$ , including  $c = h_{0,\lambda}^* = 0$ , so this loss is  $4M^2$ -Lipschitz in  $v$  and agrees with  $(h(X) - c(X))^2$  at  $v = h(X)/M$ . For a comparator  $g \in \mathcal{H}$ , the inclusion  $\text{star}(\mathcal{H} - g) \subseteq \mathcal{S} := \text{star}(\mathcal{H} - \mathcal{H})$  supplies the anchored complexity. Indeed, if  $d \leq \delta_n$  is a critical radius for  $\mathcal{S}$ , then  $\mathfrak{R}_n(d; \mathcal{S}) \leq d^2$  because  $0 \in \mathcal{S}$  and  $\mathcal{S} \subseteq \text{star}(\mathcal{S} - \mathcal{S})$ . Star scaling gives  $\mathfrak{R}_n(s; \mathcal{S}) \leq sd$  for  $s \geq d$ . Hence, for  $s \geq \delta_n$ ,

$$\mathfrak{R}_n(s; \text{star}(\mathcal{U} - g/M)) \leq M^{-1} \mathfrak{R}_n(Ms; \mathcal{S}) \leq s\delta_n.$$

The measurable squared-loss processes in Assumption 3.12 are exactly the loss-increment processes for the stated centers and comparators, after this normalization. Lemma 9 therefore gives, on the allocated concentration event,

$$|(\mathbb{E}_n - \mathbb{E})[\{h(X) - h_{j-1,\lambda}^*(X)\}^2 - \{h_{j,\lambda}^*(X) - h_{j-1,\lambda}^*(X)\}^2]| \leq C\{\delta_n \|h - h_{j,\lambda}^*\| + \delta_n^2\},$$

and, for  $k = 1, \dots, t$ ,

$$|\|h - h_{k,\lambda}^*\|_{2,n}^2 - \|h - h_{k,\lambda}^*\|^2| \leq C\{\delta_n \|h - h_{k,\lambda}^*\| + \delta_n^2\}.$$

Allocate fixed fractions of  $\eta_{n,t}$  to these scalar comparisons. Their constants depend only on the fixed class bounds.

Write  $\Delta_j = \widehat{h}_{j,\lambda} - h_{j,\lambda}^*$  and  $\Delta_0 = 0$ . Expansion gives

$$I_{n,j} - I_j = (\mathbb{E}_n - \mathbb{E})[\{h_{j,\lambda}^* - h_{j-1,\lambda}^*\}^2 - \{\widehat{h}_{j,\lambda} - h_{j-1,\lambda}^*\}^2] + 2\mathbb{E}_n[\Delta_j(X)\Delta_{j-1}(X)].$$

The first term is at most  $C(\delta_n e_j + \delta_n^2)$ . The second scalar comparison and  $2\delta_n e_k \leq e_k^2 + \delta_n^2$  give  $\|\Delta_k\|_{2,n} \leq C(e_k + \delta_n)$  for  $k \geq 1$ ; this bound also holds for  $k = 0$  because  $\Delta_0 = 0$ . Empirical Cauchy–Schwarz thus bounds the last term by

$$2\|\Delta_j\|_{2,n}\|\Delta_{j-1}\|_{2,n} \leq 2Ce_j e_{j-1} + C\delta_n(e_j + e_{j-1}) + C\delta_n^2.$$

Combining these bounds yields

$$I_{n,j} - I_j \leq C\delta_n(e_j + e_{j-1}) + C\delta_n^2 + 2Ce_j e_{j-1}. \quad (\text{C.6})$$

Combining (C.4)–(C.6) and applying Young’s inequality gives

$$e_j^2 \leq C \frac{\delta_n^2 M_{\leq j}}{\lambda^2} + Ce_{j-1}^2.$$

All terms proportional to  $\delta_n^2$ , including those multiplied by  $\lambda \leq 2$ , are covered by  $\delta_n^2 M_{\leq j}/\lambda$  before division by  $\lambda$ . Since  $e_0 = 0$  and  $M_{\leq j}$  is uniformly bounded, enlarge the fixed constant  $C \geq 2$  if necessary, set  $b_0 = 0$ , and recursively choose  $b_j = C(1 + b_{j-1})$ . Induction gives (C.2). Thus each  $b_j$  depends only on  $j$  and the fixed constants.

For the weak bound, put  $v_j = \|\mathcal{T}(h_{j,\lambda}^* - h_0)\| = \|k\|$ . The absolute concentration inequality established in this proof, evaluated at the feasible residual  $\bar{f}$ , gives

$$L(\hat{h}_{j,\lambda}) \leq L_n(\hat{h}_{j,\lambda}) + C\{\delta_n \|\bar{f}\| + \delta_n^2\}.$$

The assumptions of Lemma C.1 are inherited from the assumptions of Theorem 3.18. Its upper comparison, applied to  $g = h_{j,\lambda}^* \in \mathcal{H}$  with the allocated confidence level, gives

$$L_n(g) \leq \frac{5}{4}L(g) + C\delta_n^2 = \frac{5}{4}v_j^2 + C\delta_n^2.$$

Empirical optimality,  $L(g) = v_j^2$ , and  $\|\bar{f}\| \leq v_j + u_j$  therefore imply

$$\begin{aligned} L(\hat{h}_{j,\lambda}) - L(g) &\leq \lambda I_{n,j} + \frac{1}{4}v_j^2 + C\delta_n(v_j + u_j) + C\delta_n^2 \\ &\leq \lambda I_{n,j} + C\{v_j^2 + \delta_n u_j + \delta_n^2\}, \end{aligned}$$

where the last inequality uses  $\delta_n v_j \leq (v_j^2 + \delta_n^2)/2$ . Combining this with (C.5) gives

$$\lambda e_j^2 + u_j^2 \leq C\{v_j^2 + \delta_n u_j + \delta_n^2\} + \lambda(I_{n,j} - I_j).$$

Substituting (C.6) and absorbing the linear terms gives

$$\frac{\lambda}{2}e_j^2 + \frac{1}{2}u_j^2 \leq C(v_j^2 + \delta_n^2 + \lambda e_{j-1}^2).$$

At the final iteration, boundedness of  $\mathcal{H}$  and (C.2) give the two simultaneous bounds

$$\lambda e_{t-1}^2 \leq C\lambda, \quad \lambda e_{t-1}^2 \leq C \frac{b_t \delta_n^2}{\lambda}.$$

Taking their minimum, applying the squared triangle inequality, and using the weak-bias bound (B.9) prove (C.3). The argument used only  $\lambda \in (0, 2]$ , with the upper endpoint absorbed into fixed constants. Each comparison is applied to the fixed iteration–grid pairs in the oracle-relevant segment. The definition of  $\eta_{n,t}$  and a union bound give the simultaneous event.  $\square$

**Corollary C.6** (Iterative TRAE discrepancy bound). *Under the assumptions of Theorem 3.18, the calibrated iterative TRAE scales satisfy (3.23) simultaneously for all  $\lambda \in \Lambda_n \cap [\lambda_{\text{IT}}^\circ, 2]$  on  $\mathcal{E}_{n,t}^{\text{IT}}$ .*

*Additional condition. Interpret the uniform almost-sure boundedness in Assumption 3.12 using one common null set for the whole function classes: there exists a set  $N$  with  $\mathbb{P}(W \in N) = 0$  such that, for every  $w \notin N$ , every  $h \in \mathcal{H}$  and every  $f \in \mathcal{F}$ , all three inequalities  $|h(X(w))| \leq B$ ,  $|f(Z(w))| \leq B$  and  $|m(w; f)| \leq B$  hold simultaneously. The exceptional set does not depend on  $h$  or  $f$ . Apply this convention separately to each primal or exchanged problem whenever that problem is assumed to satisfy Assumption 3.12.*

*Proof.* On the oracle grid segment,  $\lambda \geq \sqrt{s_{n,t}^{\text{T}}}$ , so  $A_{n,t}^{\text{T}}(\lambda) = \delta_n^2 + b_t \delta_n^2 / \lambda$ . Use the propagated branch  $b_t \delta_n^2 / \lambda$  in (C.3) and apply the upper comparison in Lemma C.1. Calibrate the coefficients of  $\delta_n^2$  and  $b_t \delta_n^2 / \lambda$  into the prescribed scales. The remaining bias term has coefficient  $C_{\text{IT}} \|w_0\|^2$ , which proves (3.23) on this segment.  $\square$

The next lemma establishes error bounds at every grid point in the stable branch of the iterative noise level.

**Lemma C.7** (Iterative TRAE error bounds). *Let  $\lambda \in \Lambda_n \cap [\lambda_{\text{IT}}^\circ, 2]$  satisfy  $\lambda \geq \sqrt{s_{n,t}^{\text{T}}}$ . On  $\mathcal{E}_{n,t}^{\text{IT}}$ ,*

$$\|\mathcal{T}(\widehat{h}_{t,\lambda} - h_0)\|^2 \leq C_t \{D_n(\widehat{h}_{t,\lambda}) + A_{n,t}^{\text{T}}(\lambda)\}, \quad (\text{C.7})$$

$$\|\widehat{h}_{t,\lambda} - h_0\|^2 \leq C_t \left[ \frac{1}{\lambda} \{D_n(\widehat{h}_{t,\lambda}) + A_{n,t}^{\text{T}}(\lambda)\} + \{D_n(\widehat{h}_{t,\lambda}) + A_{n,t}^{\text{T}}(\lambda)\}^{\beta/(\beta+1)} \right], \quad (\text{C.8})$$

$$\lambda \|\widehat{h}_{t,\lambda} - h_{t,\lambda}^*\|^2 + \|\mathcal{T}(\widehat{h}_{t,\lambda} - h_{t,\lambda}^*)\|^2 \leq C_t \{D_n(\widehat{h}_{t,\lambda}) + A_{n,t}^{\text{T}}(\lambda)\}, \quad (\text{C.9})$$

$$\|\mathcal{T}(h_{t,\lambda}^* - h_0)\|^2 \leq C_t \{D_n(\widehat{h}_{t,\lambda}) + A_{n,t}^{\text{T}}(\lambda)\}, \quad (\text{C.10})$$

$$\|\widehat{h}_{t,\lambda} - h_{t,\lambda}^*\|^2 \leq \frac{C_t}{\lambda} \{D_n(\widehat{h}_{t,\lambda}) + A_{n,t}^{\text{T}}(\lambda)\}. \quad (\text{C.11})$$

Additional condition. Interpret the uniform almost-sure boundedness in Assumption 3.12 using one common null set for the whole function classes: there exists a set  $N$  with  $\mathbb{P}(W \in N) = 0$  such that, for every  $w \notin N$ , every  $h \in \mathcal{H}$  and every  $f \in \mathcal{F}$ , all three inequalities  $|h(X(w))| \leq B$ ,  $|f(Z(w))| \leq B$  and  $|m(w; f)| \leq B$  hold simultaneously. The exceptional set does not depend on  $h$  or  $f$ . Apply this convention separately to each primal or exchanged problem whenever that problem is assumed to satisfy Assumption 3.12.

*Remark.* The constant  $C_t$  may depend on the prescribed  $t$  and fixed population quantities, including  $\beta$  and  $\|w_0\|$ , but is independent of  $n$ ,  $\lambda$ , and the realized fitting observations.

*Proof.* Set

$$e = \|\hat{h}_{t,\lambda} - h_{t,\lambda}^*\|, \quad u = \|\mathcal{T}(\hat{h}_{t,\lambda} - h_{t,\lambda}^*)\|, \quad v = \|\mathcal{T}(h_{t,\lambda}^* - h_0)\|.$$

Let

$$\begin{aligned} Q_t^\lambda(h) &= L(h) + \lambda \|h - h_{t-1,\lambda}^*\|^2, \\ \tilde{Q}_{n,t}^\lambda(h) &= L_n(h) + \lambda \|h - h_{t-1,\lambda}^*\|_{2,n}^2, \\ Q_{n,t}^\lambda(h) &= L_n(h) + \lambda \|h - \hat{h}_{t-1,\lambda}\|_{2,n}^2. \end{aligned}$$

The loss comparison in Lemma C.1 and the localized squared-norm comparison, applied at the fixed comparator  $h_{t,\lambda}^* \in \mathcal{H}$ , give

$$\begin{aligned} &Q_t^\lambda(\hat{h}_{t,\lambda}) - Q_t^\lambda(h_{t,\lambda}^*) - \{\tilde{Q}_{n,t}^\lambda(\hat{h}_{t,\lambda}) - \tilde{Q}_{n,t}^\lambda(h_{t,\lambda}^*)\} \\ &\leq L_n(\hat{h}_{t,\lambda}) + \frac{1}{4}v^2 + C\{\delta_n^2 + \lambda\delta_n e\}. \end{aligned}$$

Empirical optimality gives  $Q_{n,t}^\lambda(\hat{h}_{t,\lambda}) \leq Q_{n,t}^\lambda(h_{t,\lambda}^*)$ . Expanding the two empirical penalties yields

$$\left| \{\tilde{Q}_{n,t}^\lambda(\hat{h}_{t,\lambda}) - \tilde{Q}_{n,t}^\lambda(h_{t,\lambda}^*)\} - \{Q_{n,t}^\lambda(\hat{h}_{t,\lambda}) - Q_{n,t}^\lambda(h_{t,\lambda}^*)\} \right| \leq 2\lambda \|e_t\|_{2,n} \|e_{t-1}\|_{2,n},$$

where  $e_j = \hat{h}_{j,\lambda} - h_{j,\lambda}^*$  in this display. The localized norm comparison and (C.2) bound the last expression together with the term  $C\lambda\delta_n e$  in the fixed-center comparison by

$$\epsilon\lambda e^2 + C_{\epsilon,t} \left( \delta_n^2 + \frac{b_t\delta_n^2}{\lambda} \right)$$

for any fixed  $\epsilon \in (0, 1)$ . Thus  $\epsilon$  is the total absorption allocated to both terms. Taking  $\epsilon = 1/4$  and using the population quadratic identity (C.5) yields

$$\frac{3}{4}\lambda e^2 + u^2 \leq L_n(\hat{h}_{t,\lambda}) + \frac{1}{4}v^2 + C_t \left( \delta_n^2 + \frac{b_t\delta_n^2}{\lambda} \right). \quad (\text{C.12})$$

On the branch  $\lambda \geq \sqrt{s_{n,t}^T}$ ,

$$A_{n,t}^T(\lambda) = \delta_n^2 + \frac{s_{n,t}^T}{\lambda}.$$

Since  $D_n = L_n$ , (C.12) therefore gives

$$\frac{3}{4}\lambda e^2 + u^2 \leq \frac{1}{4}v^2 + C_t\{D_n(\hat{h}_{t,\lambda}) + A_{n,t}^T(\lambda)\}. \quad (\text{C.13})$$

The lower comparison in Lemma C.1 implies

$$L(\hat{h}_{t,\lambda}) \leq 2L_n(\hat{h}_{t,\lambda}) + C\delta_n^2 \leq C_t\{D_n(\hat{h}_{t,\lambda}) + A_{n,t}^T(\lambda)\}.$$

This proves (C.7). Since  $v^2 \leq 2L(\hat{h}_{t,\lambda}) + 2u^2$ , substituting (C.13) on the right gives

$$v^2 \leq C_t\{D_n(\hat{h}_{t,\lambda}) + A_{n,t}^T(\lambda)\} + \frac{1}{2}v^2.$$

Absorb the last term to obtain (C.10), and substitute it into (C.13). Multiplying the resulting bound by 4/3 gives (C.9). Dividing (C.9) by  $\lambda$  proves (C.11). Apply (B.10) to (C.10) and combine the result with (C.11) by the squared triangle inequality. This proves (C.8).  $\square$

*Proof of Proposition 3.17.* Throughout the proof, we work on  $\mathcal{E}_{n,t}^{\text{IT}}$ . At the terminal grid point, empirical optimality at the final iteration and  $h_0 \in \mathcal{H}$  give

$$L_n(\hat{h}_{t,\lambda_{K_n}}) \leq L_n(h_0) + \lambda_{K_n}\|h_0 - \hat{h}_{t-1,\lambda_{K_n}}\|_{2,n}^2 \leq \delta_n^2 + 4B^2n^{-2}.$$

Since  $A_{n,t}^T(\lambda) \geq \delta_n^2$ ,  $\delta_n^2 \geq n^{-1}$ , and  $n \geq N_{\text{IT}}(B)$ , the terminal point is accepted. The first accepted index is therefore well defined; it is either zero or has a rejected predecessor in  $[\lambda, 2\lambda]$ . Each grid candidate requires  $t$  minimax problems.  $\square$

*Proof of Theorem 3.18.* Write  $s = s_{n,t}^T$  and  $q = q_t$ . Boundedness gives the result directly when  $s > 1$ , so consider  $s \leq 1$ . The discrepancy bound (3.23) gives (3.2) with the bounds in (3.21). Lemma C.7 establishes (3.3) for every  $\lambda \in \mathcal{S}_n = [\lambda_{\text{IT}}^\circ, 2]$  with  $\gamma = \beta/(\beta + 1)$ . For all sufficiently large  $n$ , the point  $\lambda_{\text{IT}}^\circ$  defined in the main text exists and satisfies (3.4). Since  $\lambda_{\text{IT}}^\circ \geq \sqrt{s}$ , on  $\mathcal{S}_n$ ,

$$A_{n,t}^T(\lambda) = \delta_n^2 + \frac{s}{\lambda}, \quad V_{n,t}^T(\lambda) = \frac{s}{\lambda^2}.$$

The noise level is therefore nonincreasing, and  $b_t \geq 2$  with  $\lambda \leq 2$  give  $A_{n,t}^T(\lambda)/\lambda \leq CV_{n,t}^T(\lambda)$  on that segment. Apply Theorem 3.1 and (3.22). The weak bound follows directly, while the source-bias contribution to the strong bound satisfies

$$(s^{q/(q+1)})^{\beta/(\beta+1)} \leq s^{(q-1)/(q+1)}$$

because  $q \leq \beta + 1$ . This proves the two stated rates. Boundedness covers the remaining finitely many  $n$ , after we enlarge the fixed constants. Finally,  $2t \geq \beta + 1$  gives  $q = \beta + 1$  and the displayed unsaturated rates.  $\square$

*Proof of Corollary 3.19.* Apply Theorem 3.14 to both nuisance estimators at the prescribed confidence levels  $\xi_n/2$ . Their rate bounds then hold on a common event with probability at least  $1 - \xi_n \rightarrow 1$ . Because  $\beta_h, \beta_q > 0$  and  $\delta_n \rightarrow 0$ , these bounds give

$$\|\hat{h} - h_0\| = o_p(1), \quad \|\hat{q} - q_0\| = o_p(1). \quad (\text{C.14})$$

The remaining bias condition is

$$\sqrt{n}\mathbb{E}[(\hat{q}(Z) - q_0(Z))(\hat{h}(X) - h_0(X))] = o_p(1). \quad (\text{C.15})$$

Suppose first that  $\beta_m = \beta_h \geq \beta_q$ . By the adjoint identity and the Cauchy–Schwarz inequality,

$$\begin{aligned} \left| \mathbb{E}[(\hat{q}(Z) - q_0(Z))(\hat{h}(X) - h_0(X))] \right| &\leq \|\hat{h} - h_0\| \|\mathcal{T}^*(\hat{q} - q_0)\| \\ &= O_p\left(\delta_n^{\min\{\beta_m, 1\}/(1+\min\{\beta_m, 1\})} \cdot \delta_n\right) \\ &= O_p\left(\delta_n^{1+\min\{\beta_m, 1\}/(1+\min\{\beta_m, 1\})}\right). \end{aligned}$$

If  $\beta_m = \beta_q \geq \beta_h$ , the symmetric bound

$$\left| \mathbb{E}[(\hat{q}(Z) - q_0(Z))(\hat{h}(X) - h_0(X))] \right| \leq \|\hat{q} - q_0\| \|\mathcal{T}(\hat{h} - h_0)\|$$

gives the same rate. By assumption,  $\delta_n = o\left(n^{-(1+\min\{\beta_m, 1\})/(2+4\min\{\beta_m, 1\})}\right)$ , and hence

$$\left| \mathbb{E}[(\hat{q}(Z) - q_0(Z))(\hat{h}(X) - h_0(X))] \right| = o_p\left(n^{-1/2}\right),$$

which proves (C.15).

Use the jointly measurable version of  $d_n$  stipulated in Corollary 3.19. Its agreement with the fitted score difference for almost every training realization and almost every fresh observation identifies its conditional population integrals with those of the score difference. Its almost sure agreement at every evaluation observation identifies its evaluation average with the difference of the two evaluated score averages. Conditional on the nuisance-training data, the evaluation observations are independent and identically distributed. Boundedness of the nuisance classes and the score continuity assumptions imply

$$\mathbb{E}[d_n(W)^2 \mid \text{training}] \leq C \left\{ \|\hat{h} - h_0\|^{\gamma_h} + \|\hat{q} - q_0\|^{\gamma_q} + \|\hat{h} - h_0\|^2 + \|\hat{q} - q_0\|^2 \right\} = o_p(1).$$

Conditional Chebyshev's inequality therefore gives  $\sqrt{n}(\mathbb{E}_{n,\text{eval}} - \mathbb{E})d_n = o_p(1)$ , where  $\mathbb{E}d_n$  integrates a fresh observation conditional on the training data. The primal and dual moment equations give

$$\mathbb{E}d_n = -\mathbb{E}[(\widehat{h}(X) - h_0(X))(\widehat{q}(Z) - q_0(Z))].$$

Combining this identity with (C.15) yields

$$\sqrt{n}\{\widehat{\theta}(\widehat{h}, \widehat{q}) - \theta_0\} = \frac{1}{\sqrt{n}} \sum_{i=1}^n \psi_0(W_i) + o_p(1).$$

The moment equations also imply  $\mathbb{E}\psi_0(W) = 0$ . The classical central limit theorem applies because  $0 < \mathbb{E}[\psi_0(W)^2] < \infty$ , proving the stated normal limit.  $\square$

## D Additional Experiment Details

**Source condition in the linear design.** Let  $\phi_k = \text{He}_k/\sqrt{k!}$  denote the normalized probabilists' Hermite polynomials, an orthonormal basis of  $L_2(N(0, 1))$ . Since  $(X, Z)$  is a standard bivariate normal pair with correlation  $\rho$ , Mehler's formula gives  $\mathbb{E}[\phi_k(X) | Z] = \rho^k \phi_k(Z)$ , so  $\mathcal{T}^* \mathcal{T} \phi_k = \rho^{2k} \phi_k$ . The function  $h_s$  is even and centered, so  $\langle h_s, \phi_0 \rangle = 0$  and  $\langle h_s, \phi_{2k+1} \rangle = 0$ . For  $k \geq 1$ , a direct calculation gives

$$\langle h_s, \phi_{2k} \rangle = \frac{m_s}{d_s} \frac{(-1)^k \sqrt{(2k)!}}{2^k k! (1 + s^2)^k}.$$

The source condition  $h_s = (\mathcal{T}^* \mathcal{T})^{\beta/2} w_0$  holds if and only if

$$\|w_0\|^2 = \sum_{k \geq 1} \frac{\langle h_s, \phi_{2k} \rangle^2}{\rho^{4k\beta}} = \frac{m_s^2}{d_s^2} \sum_{k \geq 1} \binom{2k}{k} \frac{q^k}{4^k} = \frac{m_s^2}{d_s^2} \{(1 - q)^{-1/2} - 1\}, \quad q = \frac{\rho^{-4\beta}}{(1 + s^2)^2},$$

is finite, that is, if and only if  $q < 1$ . This is the condition  $\beta < \beta_{\max}(s, \rho)$  in Section 4.1. Table 3 gives its values.

Table 3: Upper end  $\beta_{\max}(s, \rho)$  of the source exponents satisfied by  $h_s$  in the linear design. The source condition holds for every  $\beta < \beta_{\max}$  and fails at  $\beta_{\max}$ .

|           | $\rho = 0.65$ | $\rho = 0.8$ | $\rho = 0.95$ |
|-----------|---------------|--------------|---------------|
| $s = 0.7$ | 0.46          | 0.89         | 3.89          |
| $s = 2.2$ | 2.05          | 3.95         | 17.2          |

**Errors.** Strong errors use a trapezoid rule on 2001 points in  $[-8, 8]$  for the standard normal regressor. Weak errors integrate the analytic conditional law of  $X$  given the instrument over a grid in  $(Z, V)$  on  $[-6, 6]^2$ , with  $241 \times 161$  points in the 90-regime study and  $121 \times 81$  points in the two-design study. In the nonlinear design, both errors use a joint grid over  $(A, Z, V)$  weighted by  $\mathbb{P}(A | Z)$ . The plug-in target uses the second half of the sample for TRAE and the last quarter for RDIV. In the nonlinear design, the continuous regressor is standardized with the mean and standard deviation of the first half of the sample.

**RDIV nuisances and noise level.** The mixture density network has 20 components and hidden width 64. It is trained by Adam with learning rate  $10^{-3}$  and weight decay  $10^{-4}$  for at most 300 epochs and stops after 25 epochs without improvement on a quarter of the preliminary block. The pilot  $\hat{r}$  is kernel ridge regression on the preliminary block, with the kernel of Section 4.2 and ridge  $10^{-3}$ . Fifty draws per observation from the fitted density, frozen across candidates, define  $\hat{\mathcal{T}}$ .

Let  $G$  and  $G_{\text{cal}}$  be the matrices that map the coordinates  $\mathbf{b}$  of a hypothesis to  $\hat{\mathcal{T}}h$  on the fitting and calibration blocks, and write  $G^\top G = W \text{diag}(\mu_j) W^\top$ . The candidate is  $\hat{\mathbf{b}}_\lambda = (G^\top G + \lambda I)^{-1} G^\top \mathbf{y}$ . Let  $\hat{\sigma}^2$  be the mean squared difference between  $\mathbf{y}$  and  $\hat{r}(Z)$  on the fitting block. Treating the noise as homoskedastic with variance  $\hat{\sigma}^2$ , the variance of  $\hat{\mathcal{T}}\hat{h}_\lambda$  on the calibration block is

$$\hat{\sigma}^2 \sum_j \frac{\mu_j}{(\mu_j + \lambda)^2} \frac{\|G_{\text{cal}} W \mathbf{e}_j\|^2}{N_{\text{cal}}},$$

where  $\mathbf{e}_j$  is the  $j$ th unit vector. The pilot has the form  $\hat{r}(z) = \mathbf{w}(z)^\top \mathbf{y}_{\text{pre}}$ , and its variance term is  $\hat{\nu}^2 = \hat{\sigma}^2 N_{\text{cal}}^{-1} \sum_{i \in I_{\text{cal}}} \|\mathbf{w}(Z_i)\|^2$ . For the operator error, the difference  $\hat{h}_\lambda(X_i) - \hat{\mathcal{T}}\hat{h}_\lambda(Z_i)$  has conditional mean  $\{(\mathcal{T} - \hat{\mathcal{T}})\hat{h}_\lambda\}(Z_i)$  given  $Z_i$ . On the calibration block, we evaluate the TRAE statistic with the kernel critic on  $Z$  at this difference, subtract its sign-flip level, and take the positive part. The noise level  $A_n(\lambda)$  is the sum of the three terms.

*Remark (notation).*  $\mathbf{y}$  denotes the vector of outcomes on the fitting block, on the outcome scale used by the fit. In Proposition D.1:  $\mathbf{y} = (Y_1, \dots, Y_N)^\top$ , with  $N = n_1$  the fitting-sample size.

**Grids.** Table 4 lists the fitting-block sizes and grid lengths. Every grid starts at  $\lambda_0 = 2$  and halves at each step. The best fixed penalty of a regime is taken over its whole grid, and the best single fixed penalty over the grid values present at every sample size.

**Pilot runs.** Two pilot runs on seeds 1001–1008 fixed the noise-floor rule before the main seeds 2001–2050 were drawn. The first used 12 regimes of the linear design, with  $\rho \in \{0.65, 0.95\}$ ,

Table 4: Fitting-block sizes and grid lengths.  $M = K_N + 1$  is the number of candidates, and the grid ends at  $2^{1-K_N}$ .

| $n$  | $N_{\text{TRAE}}$ | $M$ | Final $\lambda$ | $N_{\text{RDIV}}$ | $M$ | Final $\lambda$ |
|------|-------------------|-----|-----------------|-------------------|-----|-----------------|
| 500  | 250               | 18  | $2^{-16}$       | 125               | 16  | $2^{-14}$       |
| 1000 | 500               | 20  | $2^{-18}$       | 250               | 18  | $2^{-16}$       |
| 2000 | 1000              | 22  | $2^{-20}$       | 500               | 20  | $2^{-18}$       |
| 4000 | 2000              | 24  | $2^{-22}$       | 1000              | 22  | $2^{-20}$       |
| 8000 | 4000              | 26  | $2^{-24}$       | 2000              | 24  | $2^{-22}$       |

$s = 0.7$ ,  $\sigma \in \{0.1, 3\}$ , and  $n \in \{500, 2000, 8000\}$ ; the second used 16 regimes of both designs at  $n \in \{1000, 4000\}$ . For RDIV, the operator-error term was added after the pilot showed that without it the search reached the end of the grid in up to 38% of the data sets. On the pilot, the fixed-scale RDIV level  $\delta_N^2/\lambda + \nu_N^2$  performed about as well on the strong error and worse on the weak and target errors.

**Provenance.** The two studies are **20261001-gaussian-regimes** (90 regimes, TRAE) and **20261001-regimes-v2** (two designs, TRAE and RDIV) in the project’s experiment directory at source revision 5985a2c. Both contain every fitted candidate with its discrepancy, noise level, and errors, so all tables and figures can be rebuilt from them. In both studies the linear-design TRAE fits at  $s = 0.7$  coincide. The kernel computations ran on CPUs; the mixture density networks for RDIV ran on GPUs.

## D.1 RKHS computation

Let  $N = n_1$  be the fitting-sample size,  $\mathbf{y} = (Y_1, \dots, Y_N)^\top$ , and define the Gram matrices

$$[K_{\mathcal{H}}]_{ij} = k_{\mathcal{H}}(X_i, X_j), \quad [K_{\mathcal{F}}]_{ij} = k_{\mathcal{F}}(Z_i, Z_j), \quad P_\eta = K_{\mathcal{F}}(K_{\mathcal{F}} + N\eta I_N)^{-1}.$$

For the minimum-RKHS-norm representative  $h_\alpha(x) = \sum_{j=1}^N \alpha_j k_{\mathcal{H}}(X_j, x)$ , the critic can be profiled out in closed form. Before the spectral restriction below, the RKHS TRAE fit at a positive candidate  $\lambda$  solves the convex quadratic program

$$\hat{\alpha}_\lambda \in \arg \min_{\alpha \in \mathbb{R}^N} \{(\mathbf{y} - K_{\mathcal{H}}\alpha)^\top P_\eta(\mathbf{y} - K_{\mathcal{H}}\alpha) + \lambda \alpha^\top K_{\mathcal{H}}^2 \alpha\}. \quad (\text{D.1})$$

Equivalently, after dropping the constant  $\mathbf{y}^\top P_\eta \mathbf{y}$ , the objective is

$$\alpha^\top (K_{\mathcal{H}} P_\eta K_{\mathcal{H}} + \lambda K_{\mathcal{H}}^2) \alpha - 2\alpha^\top K_{\mathcal{H}} P_\eta \mathbf{y}.$$

Both  $P_\eta$  and  $K_{\mathcal{H}}$  are positive semidefinite, so this objective is convex and its stationarity condition

$$(K_{\mathcal{H}}P_\eta K_{\mathcal{H}} + \lambda K_{\mathcal{H}}^2) \alpha = K_{\mathcal{H}}P_\eta \mathbf{y} \quad (\text{D.2})$$

identifies the minimizer. Proposition D.1 gives its minimum-Euclidean-norm solution,  $\hat{\alpha}_\lambda = (K_{\mathcal{H}}P_\eta K_{\mathcal{H}} + \lambda K_{\mathcal{H}}^2)^\dagger K_{\mathcal{H}}P_\eta \mathbf{y}$ . Each candidate therefore costs a linear solve, with no iterative optimization.

A change of variable gives a well-conditioned computation on a prescribed spectral subspace. Let  $Q_{\mathcal{H}}$  contain the orthonormal eigenvectors of  $K_{\mathcal{H}}$  whose eigenvalues exceed  $\epsilon_{\mathcal{H}}\lambda_{\max}(K_{\mathcal{H}})$ , and let  $D_{\mathcal{H}}$  contain those eigenvalues. The default relative tolerance is  $\epsilon_{\mathcal{H}} = \sqrt{N\epsilon_{\text{mach}}}$ . Restrict the hypothesis to the corresponding kernel span. Its fitting values have the form  $\mathbf{v} = Q_{\mathcal{H}}\mathbf{a}$ , and its coefficients are  $\alpha = Q_{\mathcal{H}}D_{\mathcal{H}}^{-1}\mathbf{a}$ . With

$$M = Q_{\mathcal{H}}^\top P_\eta Q_{\mathcal{H}}, \quad \mathbf{b} = Q_{\mathcal{H}}^\top P_\eta \mathbf{y},$$

the restricted quadratic program becomes

$$\hat{\mathbf{a}}_\lambda = \arg \min_{\mathbf{a}} \{(\mathbf{y} - Q_{\mathcal{H}}\mathbf{a})^\top P_\eta (\mathbf{y} - Q_{\mathcal{H}}\mathbf{a}) + \lambda \|\mathbf{a}\|^2\}, \quad (M + \lambda I)\hat{\mathbf{a}}_\lambda = \mathbf{b}. \quad (\text{D.3})$$

The matrix  $M$  is positive semidefinite with spectrum in  $[0, 1)$ , so the system has condition number at most  $(1 + \lambda)/\lambda$  for  $\lambda > 0$ .

Write  $M = V \text{diag}(m_i) V^\top$  and  $\mathbf{t} = V^\top \mathbf{b}$ . Then the whole regularization path follows from

$$\hat{\mathbf{a}}_\lambda = V \text{diag}\left(\frac{1}{m_i + \lambda}\right) \mathbf{t}, \quad L_{N,\eta}(\hat{h}_\lambda) = \frac{1}{N}(\mathbf{y} - Q_{\mathcal{H}}\hat{\mathbf{a}}_\lambda)^\top P_\eta (\mathbf{y} - Q_{\mathcal{H}}\hat{\mathbf{a}}_\lambda). \quad (\text{D.4})$$

At  $\lambda = 0$ , we use the minimum-Euclidean-norm solution after retaining compressed eigenvalues larger than the absolute numerical tolerance  $10^{-12}$ . The selection statistic and predictions use the same fitted function, since  $K_{\mathcal{H}}\hat{\alpha}_\lambda = Q_{\mathcal{H}}\hat{\mathbf{a}}_\lambda$  on the retained subspace. Kernel and compressed-matrix eigendecompositions are computed once; each subsequent candidate costs at most  $O(N^2)$  operations. The computed fit minimizes the objective on the retained hypothesis span, up to numerical error. This restriction is an additional regularization choice and must be included in the experiment settings.

For RDIV, let  $\mathbf{y}$  be the standardized outcomes on the fitting block, and restrict the hypothesis to the retained span as above, with coordinates  $\mathbf{b}$ . Averaging the hypothesis features over the frozen draws from the fitted density gives a matrix  $G$  with  $G\mathbf{b} = (\hat{\mathcal{T}}h(Z_i))_{i=1}^N$ . The candidate minimizes  $N^{-1}\{\|\mathbf{y} - G\mathbf{b}\|^2 + \lambda \|\mathbf{b}\|^2\}$ , so that

$$(G^\top G + \lambda I)\hat{\mathbf{b}}_\lambda = G^\top \mathbf{y},$$

and one eigendecomposition of  $G^\top G$  gives the whole path. The discrepancy is  $N_{\text{cal}}^{-1} \|\widehat{\mathbf{r}}_{\text{cal}} - G_{\text{cal}} \widehat{\mathbf{b}}_\lambda\|^2$ , where  $G_{\text{cal}}$  is formed in the same way on the calibration block and  $\widehat{\mathbf{r}}_{\text{cal}}$  contains the pilot values there. The retained span and the critic ridge are computational choices; the guarantees for bounded classes do not follow from these reductions alone.

## D.2 RKHS reduction

For the kernel TRAE estimator in Section 4.2, define the ridge-stabilized empirical adversarial loss

$$L_{N,\eta}(h) := \sup_{f \in \mathcal{F}} \left\{ \frac{2}{N} \sum_{i=1}^N \{Y_i - h(X_i)\} f(Z_i) - \frac{1}{N} \sum_{i=1}^N f(Z_i)^2 - \eta \|f\|_{\mathcal{F}}^2 \right\}, \quad (\text{D.5})$$

where  $\mathcal{F}$  is the RKHS generated by  $k_{\mathcal{F}}$  and  $\eta > 0$ . The fitted hypothesis minimizes

$$L_{N,\eta}(h) + \frac{\lambda}{N} \sum_{i=1}^N h(X_i)^2.$$

**Proposition D.1** (RKHS quadratic reduction). *Let  $\mathcal{H}$  and  $\mathcal{F}$  be the RKHSs generated by  $k_{\mathcal{H}}$  and  $k_{\mathcal{F}}$ , and let  $K_{\mathcal{H}}$ ,  $K_{\mathcal{F}}$ , and  $P_\eta$  be defined as in Section D.1. Among functions with the same fitting-sample evaluations, take the representative with minimum RKHS norm. Then the minimizer of*

$$L_{N,\eta}(h) + \lambda \mathbb{E}_N[h(X)^2]$$

*has the form  $\widehat{h}_\lambda(\cdot) = \sum_{j=1}^N \widehat{\alpha}_{\lambda,j} k_{\mathcal{H}}(X_j, \cdot)$ , where  $\widehat{\alpha}_\lambda$  solves (D.1). The minimum-Euclidean-norm coefficient vector is*

$$\widehat{\alpha}_\lambda = (K_{\mathcal{H}} P_\eta K_{\mathcal{H}} + \lambda K_{\mathcal{H}}^2)^\dagger K_{\mathcal{H}} P_\eta \mathbf{y}, \quad (\text{D.6})$$

where  $A^\dagger$  denotes the Moore–Penrose pseudoinverse of  $A$ .

*Proof.* Fix  $h$  and write  $\mathbf{r}_h = (Y_i - h(X_i))_{i=1}^N$ . The representer theorem gives a maximizing critic of the form

$$f_\gamma(\cdot) = \sum_{j=1}^N \gamma_j k_{\mathcal{F}}(Z_j, \cdot).$$

Its fitting-sample evaluations are  $K_{\mathcal{F}} \gamma$  and its squared RKHS norm is  $\gamma^\top K_{\mathcal{F}} \gamma$ . Substitution into (D.5) gives

$$\sup_{\gamma \in \mathbb{R}^N} \left\{ \frac{2}{N} \mathbf{r}_h^\top K_{\mathcal{F}} \gamma - \frac{1}{N} \gamma^\top K_{\mathcal{F}}^2 \gamma - \eta \gamma^\top K_{\mathcal{F}} \gamma \right\}.$$

The first-order condition is

$$K_{\mathcal{F}} \{ (K_{\mathcal{F}} + N\eta I_N) \gamma - \mathbf{r}_h \} = 0.$$

We may therefore take

$$\hat{\gamma}_h = (K_{\mathcal{F}} + N\eta I_N)^{-1} \mathbf{r}_h.$$

Evaluating the objective at this coefficient vector yields

$$L_{N,\eta}(h) = \frac{1}{N} \mathbf{r}_h^\top K_{\mathcal{F}} (K_{\mathcal{F}} + N\eta I_N)^{-1} \mathbf{r}_h = \frac{1}{N} \mathbf{r}_h^\top P_\eta \mathbf{r}_h. \quad (\text{D.7})$$

The objective depends on  $h$  only through its fitting-sample evaluations. Its minimum-RKHS-norm representative consequently has the form  $h_\alpha(\cdot) = \sum_{j=1}^N \alpha_j k_{\mathcal{H}}(X_j, \cdot)$ , with evaluation vector  $K_{\mathcal{H}}\alpha$ . Combining (D.7) with

$$\mathbb{E}_N[h_\alpha(X)^2] = \frac{1}{N} \alpha^\top K_{\mathcal{H}}^2 \alpha$$

and multiplying the objective by  $N$  gives (D.1). Since  $P_\eta$  is positive semidefinite, this is a convex quadratic program. Its first-order condition is

$$(K_{\mathcal{H}} P_\eta K_{\mathcal{H}} + \lambda K_{\mathcal{H}}^2) \alpha = K_{\mathcal{H}} P_\eta \mathbf{y}.$$

The Moore–Penrose solution in (D.6) is the minimum-Euclidean-norm solution of this system.  $\square$