DEEP MINIMAX ESTIMATION FOR NONPARAMETRIC INSTRUMENTAL VARIABLES REGRESSION

Anonymous authors
Paper under double-blind review

ABSTRACT

We study nonparametric instrumental-variable (IV) mean, expectile, and quantile regression using a common neural minimax estimator. ReLU networks represent the structural function and a bounded instrument-indexed critic. Estimation minimizes the largest absolute empirical residual-score moment without separately estimating a first-stage conditional law. Mean and expectile scores enter directly, whereas the quantile indicator is approximated by an integrated kernel. For exact empirical minimizers, we establish nonasymptotic bounds separating statistical error, approximation errors for both networks, and quantile-smoothing error. Under the stated regularity conditions and Hölder smoothness ζ, the expected adversarial-moment risk admits an upper bound of order $n^{-\zeta/(d_{\max}+2\zeta)}$, up to logarithmic factors, where $d_{\max}$ is the larger of the regressor and instrument dimensions. Additional localization, critic-link, and spectral conditions yield structural $L^2(P_X)$ upper bounds. Simulation studies assess finite-sample structural recovery, while an application to randomized-encouragement diary data illustrates differences among fitted response curves.

1 INTRODUCTION

Valid instrumental variables (IVs), together with appropriate identification conditions, allow structural relationships to be studied when regressors are associated with unobserved determinants of the outcome. An average structural response may not capture the asymmetry relevant to a particular application. Expectile and quantile restrictions provide alternative targets: the former balance weighted residual magnitudes, while the latter balance residual probabilities. Let Y denote the outcome, X the possibly endogenous regressor, and Z the instrument. We study functions f_0^m, f_0^e, and f_0^q satisfying, respectively,

$$\begin{aligned}\mathbb{E}\{Y - f_0^m(X) \mid Z\} &= 0,\\ \mathbb{E}[\psi_\tau\{Y - f_0^e(X)\} \mid Z] &= 0,\\ \mathbb{P}\{Y - f_0^q(X) \le 0 \mid Z\} &= \tau,\end{aligned}$$

almost surely, where $\tau \in (0,1)$ and ψ_τ is the asymmetric least-squares score defined in (2). The restrictions for mean and expectile regression coincide at $\tau = 1/2$; otherwise, the three targets need not agree.

Nonparametric IV estimation must address both identification and ill-posedness. In mean regression, a compact conditional-expectation operator can map large structural errors to small observable moment errors; structural recovery rates then depend on smoothness and the degree of ill-posedness (Hall & Horowitz, 2005; Darolles et al., 2011; Chen & Reiss, 2011). Expectile and quantile moment restrictions require score-specific identification conditions. Their scores also differ in regularity: the expectile score is continuous and globally Lipschitz with a linear envelope, while the quantile score is a bounded, discontinuous indicator score (Horowitz & Lee, 2007). We use sub-Gaussian outcome control for the mean and expectile analyses, and integrated-kernel smoothing for quantile estimation.

Classical nonparametric IV (NPIV) estimators regularize inverse problems by sieves, spectral methods, kernels, or Tikhonov regularization (Newey & Powell, 2003; Carrasco et al., 2007; Darolles et al., 2011). Machine-learning approaches also use learned conditional distributions or features in

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

two-stage procedures (Hartford et al., 2017; Xu et al., 2021; Kim et al., 2025), or estimate conditional moments through an adversarial game (Dikkala et al., 2020). IV expectile estimation already has semiparametric and variational precedents (Sobotka et al., 2013; Bennett & Kallus, 2023). The question studied here is how approximation of both neural players and quantile smoothing affect finite-sample risk under a common bounded-critic criterion, and what further conditions are needed for structural recovery.

We represent the structural function and the instrument-indexed critic by ReLU deep neural networks (DNNs), and minimize the largest absolute empirical residual-score moment over the critic class. This construction requires no preliminary estimation of a conditional law or conditional-expectation operator. Our analysis concerns exact empirical minimizers; the experiments use a penalized alternating-training implementation. The main contributions are:

- We prove nonasymptotic oracle inequalities separating statistical error, structural-network and critic-network approximation errors, and quantile-smoothing error. Under the stated regularity conditions and Hölder smoothness ζ, we obtain the expected adversarial-moment risk bound $\widetilde{O}(n^{-\zeta/(d_{\max}+2\zeta)})$, where $d_{\max}$ is the larger input dimension.

- We convert the original-score moment-risk bounds to structural L^2 upper bounds under additional localization, critic-link, and spectral conditions. Mean IV uses the conditional-expectation operator; expectile and quantile IV use weighted local linearizations. These results distinguish moment estimation from ill-posed inversion.

- Simulations evaluate structural recovery within three designs. A randomized-encouragement diary application compares fitted response curves across scores and levels; its structural interpretation requires assumptions on instrument validity and is limited by modest instrument relevance, clustered sampling, and a discrete outcome.

1.1 Paper organization

Sections 2–4 present notation, estimators, and theory; Sections 5–7 present experiments and conclusions. Detailed definitions, assumptions, proofs, and experimental specifications appear in Appendices B–G.

2 Preliminaries

Let $(Y_i, X_i, Z_i)_{i=1}^n$ be independent and identically distributed (i.i.d.) copies of (Y, X, Z) with joint law P, where $Y \in \mathbb{R}$, $X \in \mathcal{X} \subseteq \mathbb{R}^d$, and $Z \in \mathcal{Z} \subseteq \mathbb{R}^p$. Write P_X, P_Z, and P_{XZ} for the marginal and joint laws of the covariates, and set $d_{\max} = \max(d, p)$. We use standard L^ν and supremum norms throughout. The relation $a \lesssim b$ means $a \le Cb$ uniformly over varying quantities, and O_P denotes boundedness in probability. For a function f on a set $\mathcal{A}$, write $\|f\|_{\infty,\mathcal{A}} = \sup_{x\in\mathcal{A}} |f(x)|$, suppressing $\mathcal{A}$ only when its domain is unambiguous. For a probability measure P and $1 \le \nu < \infty$, define $\|f\|_{L^\nu(P)} = (\int |f|^\nu \, dP)^{1/\nu}$, and set $\|f\|_\infty = \|f\|_{\infty,\mathrm{dom}(f)}$. We write P_n for the empirical measure. The relation $a \lesssim b$ means that $a \le Cb$ for some constant $C > 0$, uniformly over the varying quantities, and $a \asymp b$ means that both $a \lesssim b$ and $b \lesssim a$. For deterministic sequences, $a_n = \mathcal{O}(b_n)$ means that $|a_n| \le Cb_n$ for all sufficiently large n and some constant $C > 0$. For random sequences, $A_n = O_P(b_n)$ means that A_n/b_n is bounded in probability. The notation $\widetilde{\mathcal{O}}(b_n)$ suppresses logarithmic factors multiplying b_n. Table 1 summarizes the structural and critic classes, population risks, smoothing bandwidth, and network architecture for the mean, expectile, and quantile formulations.

Table 1: Principal notation.

Symbol	Meaning
$f_0, \widehat{f}$	structural function and its estimator
$\mathcal{F}^*, \mathcal{G}^*$	population structural and critic classes
$\mathcal{F}_n, \mathcal{G}_n$	sample-size-dependent neural classes
$\mathcal{L}^m_{\mathcal{V}}$	population mean adversarial risk over $\mathcal{V}$
$\mathcal{L}^e_{\mathcal{V}}$	population expectile adversarial risk over $\mathcal{V}$
$\mathcal{L}^q_{\mathcal{V},\phi}$	unsmoothed population quantile risk
$\mathcal{L}^q_{\mathcal{V},h}$	smoothed population quantile risk
h	smoothing bandwidth
$\mathcal{W}, \mathcal{D}, \mathcal{S}$	network width, depth, and number of nonzero parameters
$\mathcal{B}_{\mathrm{par}}, \mathcal{B}$	parameter and network-output bounds

Expectile loss and score. Fix $\tau \in (0,1)$ and define

$$w_\tau(u) = \tau \mathbf{1}\{u > 0\} + (1-\tau)\mathbf{1}\{u \le 0\}, \qquad \rho^e_\tau(u) = w_\tau(u)u^2. \tag{1}$$

The derivative of the asymmetric least-squares loss away from zero is the expectile score

$$\psi_\tau(u) = 2w_\tau(u)u. \tag{2}$$

The definition at zero is unambiguous because both one-sided values of the score are zero. The map ψ_τ is continuous, piecewise linear, and globally Lipschitz with constant $2\max\{\tau, 1-\tau\} \le 2$. It also satisfies $|\psi_\tau(u)| \le 2|u|$. Our normalization makes $\psi_{1/2}(u) = u$, so the expectile score at $\tau = 1/2$ is exactly the mean-regression score. The expectile and quantile levels nevertheless index different distributional functionals away from this special case (Newey & Powell, 1987).

Definitions of kernel functions, integrated-kernel quantile smoothing, and covering numbers are collected in Appendix B.

3 Method

We first define the common moment criterion and its identification requirements, then give the three empirical estimators.

3.1 A common minimax formulation

For a candidate structural function f, define

$$\begin{aligned}
\vartheta^m_f(X,Y) &= Y - f(X),\\
\vartheta^e_f(X,Y) &= \psi_\tau\{Y - f(X)\},\\
\vartheta^q_f(X,Y) &= \phi_\tau\{Y - f(X)\} = \tau - \mathbf{1}\{Y - f(X) \le 0\}.
\end{aligned}$$

For a fixed score label $s \in \{m, e, q\}$, the corresponding structural target f^s_0 satisfies

$$\mathbb{E}\{\vartheta^s_{f^s_0}(X,Y) \mid Z\} = 0, \tag{3}$$

Multiplying by an admissible critic $g(Z)$ and taking expectations yields $\mathbb{E}\{\vartheta^s_{f^s_0}(X,Y)g(Z)\} = 0$, provided the product is integrable. A sufficiently rich critic class also gives the converse, motivating a loss based on the largest unconditional moment.

Let $\mathcal{F}^*$ and $\mathcal{G}^*$ denote the population structural and critic classes. For any candidate f, define its population adversarial-moment risk by

$$\mathcal{L}^s_{\mathcal{G}^*}(f) = \sup_{g \in \mathcal{G}^*} \left|\mathbb{E}\{\vartheta^s_f(X,Y)g(Z)\}\right|.$$

The absolute value treats moment violations of either sign symmetrically. Write $M_s(f)(Z) = \mathbb{E}\{\vartheta^s_f(X,Y) \mid Z\}$ for the conditional-moment map.

Assumption 3.1 (Score-specific identification). *Fix the score $s \in \{m, e, q\}$ for the model under consideration. The following conditions hold for that score.*

1. Conditional-moment identification. *There exists $f_0^s \in \mathcal{F}^*$ such that $M_s(f_0^s) = 0$ almost surely, and the moment map is injective at zero on $\mathcal{F}^*$:*
$$M_s(f) = 0 \text{ almost surely} \quad \Longrightarrow \quad f = f_0^s \text{ in } L^2(P_X).$$
2. Critic separation. *For every $f \in \mathcal{F}^*$, $M_s(f) \in L^2(P_Z)$, and every nonzero moment in the range of M_s is detected by $\mathcal{G}^*$:*
$$\|M_s(f)\|_{L^2(P_Z)} > 0 \quad \Longrightarrow \quad \sup_{g \in \mathcal{G}^*} |\mathbb{E}\{M_s(f)(Z)g(Z)\}| > 0.$$

For $s = q$, both conditions concern the original indicator score $\vartheta_f^q = \phi_\tau\{Y - f(X)\}$, not its smoothed version at a fixed bandwidth $h > 0$.

The first clause identifies the structural target; the second makes zero adversarial risk equivalent to the conditional restriction (Lemma 4.1). Both conditions are score specific. Appendix C.4 gives sufficient completeness conditions; Section 4.1 distinguishes the original quantile target from the smoothed population target f_h^* at fixed $h > 0$.

For estimation, we replace the population classes by ReLU network classes $\mathcal{F}$ and $\mathcal{G}$. The generic empirical risk is

$$\widehat{\mathcal{L}}_{\mathcal{G}}^s(f) = \sup_{g \in \mathcal{G}} \left| \frac{1}{n} \sum_{i=1}^n \vartheta_f^s(X_i, Y_i) g(Z_i) \right|, \tag{4}$$

The estimator minimizes this risk over $f \in \mathcal{F}$. For quantiles, we first replace the indicator score by a kernel approximation, as described below.

3.2 IV MEAN REGRESSION

Suppose that

$$Y = f_0^m(X) + \epsilon, \qquad \mathbb{E}(\epsilon \mid Z) = 0. \tag{5}$$

Because X may be associated with ϵ, $\mathbb{E}(\epsilon \mid X) = 0$ is not imposed; consequently, $\mathbb{E}(Y \mid X)$ need not equal $f_0^m(X)$. In this model, (4) gives

$$\widehat{\mathcal{L}}_{\mathcal{G}}^m(f) = \sup_{g \in \mathcal{G}} \left| \frac{1}{n} \sum_{i=1}^n \{Y_i - f(X_i)\} g(Z_i) \right|, \tag{6}$$

$$\widehat{f}_m \in \operatorname*{argmin}_{f \in \mathcal{F}} \widehat{\mathcal{L}}_{\mathcal{G}}^m(f). \tag{7}$$

3.3 IV EXPECTILE REGRESSION

Fix an expectile level $\tau \in (0, 1)$ and suppose that

$$Y = f_0^e(X) + \epsilon, \qquad \mathbb{E}\{\psi_\tau(\epsilon) \mid Z\} = 0, \tag{8}$$

where ψ_τ is defined in (2). Equivalently,

$$\mathbb{E}\{\psi_\tau(Y - f_0^e(X)) \mid Z\} = 0. \tag{9}$$

The corresponding empirical criterion and estimator are

$$\widehat{\mathcal{L}}_{\mathcal{G}}^e(f) = \sup_{g \in \mathcal{G}} \left| \frac{1}{n} \sum_{i=1}^n \psi_\tau\{Y_i - f(X_i)\} g(Z_i) \right|, \tag{10}$$

$$\widehat{f}_e \in \operatorname*{argmin}_{f \in \mathcal{F}} \widehat{\mathcal{L}}_{\mathcal{G}}^e(f). \tag{11}$$

Since ψ_τ is continuous and globally Lipschitz, this criterion requires no smoothing parameter and is directly amenable to gradient-based optimization, using a Clarke generalized derivative when a residual is exactly zero. Its score has a linear, rather than bounded, envelope, so the empirical-process analysis uses the same sub-Gaussian outcome condition as mean regression. At $\tau = 1/2$, $\psi_{1/2}(u) = u$, and the criterion reduces exactly to the IV mean criterion.

3.4 IV QUANTILE REGRESSION AND SMOOTHING

Fix a quantile level $\tau \in (0,1)$ and suppose that

$$Y = f_0^q(X) + \epsilon, \qquad \mathbb{P}(\epsilon \le 0 \mid Z) = \tau. \tag{12}$$

Define the quantile score $\phi_\tau(u) = \tau - \mathbf{1}\{u \le 0\}$. Model (12) is then equivalent to

$$\mathbb{E}\{\phi_\tau(Y - f_0^q(X)) \mid Z\} = 0. \tag{13}$$

We use this equality to define the IV quantile moment. If the conditional distribution of ϵ given Z is continuous at zero, it is equivalent to zero being a conditional τ-quantile; this continuity is implied by the density condition used for the quantile theory. Using $\vartheta_f^q(X,Y) = \phi_\tau\{Y - f(X)\}$ in (4) gives the natural IV quantile criterion. Its indicator is discontinuous in the network output $f(X_i)$, which precludes ordinary gradient updates.

To remove the indicator discontinuity, let $G(t) = \int_{-\infty}^{t} K(u)\,du$ and $\mathbf{I}_h(v) = G(-v/h)$ for bandwidth $h > 0$, with kernel conditions in Appendix B. Define

$$\phi_{\tau,h}(v) = \tau - \mathbf{I}_h(v). \tag{14}$$

For bounded K, the score $\phi_{\tau,h}$ is globally Lipschitz; the logistic choice is smooth in the residual. Smaller bandwidths reduce smoothing error but increase the score's Lipschitz constant. This replacement facilitates gradient-based updates without asserting differentiability of the full minimax objective, which also involves ReLU networks, an absolute value, and a supremum. Related smoothing devices are studied by Whang (2006); Kaplan & Sun (2017); Kaplan (2022). Replacing the indicator score in (13) by (14) leads to the empirical risk and estimator

$$\widehat{\mathcal{L}}_{\mathcal{G},h}^q(f) = \sup_{g\in\mathcal{G}} \left| \frac{1}{n}\sum_{i=1}^{n} \phi_{\tau,h}\{Y_i - f(X_i)\} g(Z_i) \right|, \tag{15}$$

$$\widehat{f}_h \in \operatorname*{argmin}_{f\in\mathcal{F}} \widehat{\mathcal{L}}_{\mathcal{G},h}^q(f). \tag{16}$$

Unlike the mean and expectile criteria, (15) targets a smoothed population moment. Section 4 bounds this risk and transfers the bound to the original quantile criterion.

4 THEORETICAL ANALYSIS

This section derives finite-sample bounds for the population adversarial-moment risks and then converts them to strong structural norms under additional link and source conditions. A common risk bound covers all three scores, with an additional smoothing term for quantiles.

Write P for integration over a fresh observation (Y, X, Z). When a fitted function appears in a population moment or risk, this integration holds the training sample, and hence the fitted function, fixed. Thus $\mathcal{L}(\widehat{f})$ remains random through the training sample; the outer expectation in $\mathbb{E}\{\mathcal{L}(\widehat{f})\}$ averages over that sample. Expectations of empirical-process quantities also average over any auxiliary randomness, such as Rademacher variables.

4.1 POPULATION CRITERIA, IDENTIFICATION, AND ESTIMATES

We define the score-specific population criteria and show that, under Assumption 3.1, the original criteria uniquely identify the structural targets. This distinction separates identification from the approximation induced by quantile smoothing.

Quantile criteria. For fixed $\tau \in (0,1)$ and critic class $\mathcal{V}$, $\mathcal{L}_{\mathcal{V},\phi}^q$ denotes the original-score risk from Section 3.1. Write $Q_h(f,g) = \mathbb{E}[\phi_{\tau,h}\{Y - f(X)\}g(Z)]$ and $\mathcal{L}_{\mathcal{V},h}^q(f) = \sup_{g\in\mathcal{V}} |Q_h(f,g)|$ for the smoothed moment and risk. Hats denote their empirical counterparts, as in (15).

The mean and expectile criteria are the corresponding instances of the common moment construction in Section 3.1, without a smoothing parameter. Their population and empirical functionals are recorded in Appendix C.1. Lemma 4.1 establishes identification for all three original population risks; Appendix E.1 provides the proof.

Lemma 4.1 (Sufficiency for point identification)**.** *Under the score-s instance of Assumption 3.1, for $s \in \{m, e, q\}$ and $f \in \mathcal{F}^*$,*

$$\mathcal{L}^s_{\mathcal{G}^*}(f) = 0 \iff M_s(f) = 0 \text{ almost surely} \iff f = f^s_0 \text{ in } L^2(P_X),$$

where $\mathcal{L}^q_{\mathcal{G}^*} = \mathcal{L}^q_{\mathcal{G}^*,\phi}$. It follows that f^s_0 is point identified and, up to P_X-null sets, is the unique minimizer of the corresponding original population adversarial risk, whose minimum is zero.

Primitive sufficient conditions are score specific: completeness on the difference class suffices for the mean model, and corresponding pathwise weighted completeness conditions suffice for the nonlinear expectile and quantile maps. Ordinary mean completeness alone does not imply these weighted conditions. The complete statements and the critic-separation argument are given in Appendix C.4.

Score-specific estimands. Lemma 4.1 gives a unique zero-risk target for each original population criterion. Quantile smoothing introduces the additional population target

$$f^*_h \in \operatorname*{argmin}_{f \in \mathcal{F}^*} \mathcal{L}^q_{\mathcal{G}^*,h}(f). \tag{17}$$

At fixed $h > 0$, f^*_h need not equal the structural target f^q_0. Throughout the theory, $\widehat{f}_h$ is the empirical estimator in (15) with $(\mathcal{F}, \mathcal{G}) = (\mathcal{F}_n, \mathcal{G}_n)$. The mean and expectile estimators likewise minimize their respective empirical adversarial risks over the same neural classes. Their exact score-specific displays are in Appendix C.1. Neither requires an intermediate smoothed population target. Below, f_0 denotes f^q_0, f^m_0, or f^e_0 when the model is clear from context. The risk bounds below concern exact global minimizers of the unpenalized empirical criteria over the stated DNNs. The finite-step, penalized alternating training used in the experiments is not assumed to return these minimizers; its optimization error and the effect of its penalties are not covered by these bounds.

4.2 Network Scaling

The complete smoothness, support, boundedness, and score-specific outcome regularity conditions are stated in Appendix C.2. In brief, the population classes are Hölder balls of smoothness ζ on compact support cubes, the neural classes are uniformly bounded, Y is conditionally sub-Gaussian for the mean and expectile results, and the quantile results use a bounded conditional density and kernel. Recall that $d_{\max} = \max(d, p)$, which includes the common-dimension case $p = d$, and set $\varepsilon_n = n^{-\zeta/(d_{\max}+2\zeta)}$. For both players, take the closed, output-clipped sparse ReLU classes

$$\mathcal{F}_n = \mathcal{N}_d(\mathcal{W}_n, \mathcal{D}_n, \mathcal{S}_n, \mathcal{B}_{\mathrm{par},n}, B_0), \qquad \mathcal{G}_n = \mathcal{N}_p(\mathcal{W}_n, \mathcal{D}_n, \mathcal{S}_n, \mathcal{B}_{\mathrm{par},n}, B_0), \tag{18}$$

where $\mathcal{W}_n, \mathcal{D}_n, \mathcal{S}_n$ bound width, depth, and the number of nonzero parameters; $\mathcal{B}_{\mathrm{par},n}$ and B_0 bound parameter magnitudes and outputs. With sufficiently large fixed multiplicative constants, take

$$\begin{aligned} \mathcal{D}_n &\asymp \log n, & \mathcal{W}_n &\asymp n^{d_{\max}/(d_{\max}+2\zeta)}, \\ \mathcal{S}_n &\asymp n^{d_{\max}/(d_{\max}+2\zeta)} \log n, & \mathcal{B}_{\mathrm{par},n} &\asymp n^{d_{\max}/(d_{\max}+2\zeta)}. \end{aligned} \tag{19}$$

With sufficiently large fixed constants, these classes contain uniformly accurate clipped ReLU approximants to both Hölder balls at order ε_n. The approximation display, attainment of the empirical minimax optima, existence of f^*_h, and measurable-selection details are given in Appendix C.5.

4.3 Population Risk Bounds

Define the one-sided critic approximation error

$$d_{\to,1}(\mathcal{G}^*, \mathcal{G}_n) = \sup_{g^* \in \mathcal{G}^*} \inf_{g \in \mathcal{G}_n} \|g - g^*\|_{L^1(P_Z)}. \tag{20}$$

The complete score-specific oracle inequalities are stated in Appendix C.3. The quantile inequality has the representative decomposition

$$\begin{aligned} \mathbb{E}\Big[\mathcal{L}^q_{\mathcal{G}^*,h}(\widehat{f}_h)\Big] \lesssim \left\{ \frac{\mathcal{S}_n \mathcal{D}_n \log(e + n\mathcal{B}_{\mathrm{par},n}\mathcal{W}_n\mathcal{D}_n/h)}{n} \right\}^{1/2} + \inf_{f \in \mathcal{F}_n} \|f - f^q_0\|_\infty \\ + d_{\to,1}(\mathcal{G}^*, \mathcal{G}_n) + h. \end{aligned}$$

Thus the four terms isolate statistical error, structural-network approximation, critic-network approximation, and quantile smoothing. The score's $1/h$ Lipschitz factor enters the statistical bound only logarithmically through the entropy calculation. Bounded conditional density makes the population approximation bound independent of $1/h$, allowing h to match the neural-approximation scale.

Theorem 4.1 (Common population-risk rate). *Suppose all three score-specific instances of Assumption 3.1 and Assumption C.1 hold, together with Assumption C.2 for mean and expectile regression and the boundedness clause of Assumption C.2 and Assumption C.3 for quantile regression. Use (18)–(19). For quantile regression, take*

$$h \asymp \varepsilon_n = n^{-\zeta/(d_{\max}+2\zeta)}.$$

Then, with $r_n = n^{-\zeta/(d_{\max}+2\zeta)}(\log n)^{3/2}$,

$$\mathbb{E}\{\mathcal{L}^m_{\mathcal{G}^*}(\widehat{f}_m)\} + \mathbb{E}\{\mathcal{L}^e_{\mathcal{G}^*}(\widehat{f}_e)\} + \mathbb{E}\{\mathcal{L}^q_{\mathcal{G}^*,h}(\widehat{f}_h)\} + \mathbb{E}\{\mathcal{L}^q_{\mathcal{G}^*,\phi}(\widehat{f}_h)\} = O(r_n).$$

Only the quantile estimator requires the bandwidth choice to keep smoothing error at the neural-approximation scale. These are conditional-moment bounds, not yet bounds on $\|\widehat{f} - f_0\|_{L^2(P_X)}$.

4.4 Conversion to strong structural norms

The preceding risks measure conditional-moment violations. We now state sufficient conditions under which they control $\|\widehat{f} - f_0\|_{L^2(P_X)}$, thereby separating ill-posed inversion from estimation of the moment criterion. For linear NPIV, this distinction is classically expressed through compact conditional-expectation operators, singular-value decay, sieve measures of ill-posedness, and source or approximation restrictions (Hall & Horowitz, 2005; Blundell et al., 2007; Darolles et al., 2011; Chen & Reiss, 2011). Nonlinear and possibly nonsmooth conditional-moment models, including IV expectile and quantile regression, additionally require a local comparison between the nonlinear moment map and its derivative (Horowitz & Lee, 2007; Chen & Pouzo, 2012; Chen et al., 2014).

For each score, write the linearized moment change as $-T(f - f_0)$, where T is the forward operator at f_0: the conditional-expectation operator for the mean score and the appropriate residual-weighted operator for the expectile and quantile scores. The full operator definitions appear in Appendix C.6. Assumption C.4 there requires a compact injective linearization with singular values σ_j; a projected critic-norming bound with exponent b; a projected tangential condition for the nonlinear scores; and a spectral source restriction with index β. Here β controls coefficients in the singular basis, whereas the Hölder index ζ controls neural approximation; the two restrictions are distinct. The following theorem additionally assumes that the estimator belongs to $\mathcal{U}_\beta$ with probability tending to one. This localization condition is separate from the preceding risk bounds; weak moment convergence alone does not establish it.

Theorem 4.2 (Strong-norm rates under ill-posedness). *Let* $r_n = n^{-\zeta/(d_{\max}+2\zeta)}(\log n)^{3/2}$. *Suppose Assumption C.4 holds and* $\mathbb{P}(\widehat{f} \in \mathcal{U}_\beta) \to 1$. *In the quantile model, let* $\widehat{f} = \widehat{f}_h$ *and impose the conditions and canonical bandwidth of Theorem C.2. In the mean model, let* $\widehat{f} = \widehat{f}_m$ *and impose the conditions of Theorem C.3. In the expectile model, let* $\widehat{f} = \widehat{f}_e$ *and impose the conditions of Theorem C.4. In all three cases, use the network scaling (19).*

1. *If the inverse singular values grow at most polynomially,* $\sigma_j^{-1} \lesssim j^a$ *for some* $a > 0$, *then*

$$\begin{aligned}\|\widehat{f} - f_0\|_{L^2(P_X)} &= O_P\left(r_n^{\beta/(a+b+\beta)}\right) \\ &= O_P\left(n^{-\zeta\beta/\{(d_{\max}+2\zeta)(a+b+\beta)\}}(\log n)^{3\beta/\{2(a+b+\beta)\}}\right).\end{aligned}$$

2. *If the inverse singular values grow at most exponentially,* $\sigma_j^{-1} \lesssim \exp(c_2 j^a)$ *for some* $a > 0$ *and* $c_2 > 0$, *then*

$$\|\widehat{f} - f_0\|_{L^2(P_X)} = O_P\{(\log n)^{-\beta/a}\}.$$

Polynomial and exponential singular-value decay describe the familiar mildly and severely ill-posed regimes (Hall & Horowitz, 2005; Chen & Reiss, 2011). The theorem needs only one-sided bounds

for strong-norm upper rates; classifying the exact degree of ill-posedness or proving matching lower bounds requires more. Its dependence on a and b separates structural ill-posedness from the cost of representing projected moments by the critic class. The proof sketch in Appendix D explains the spectral truncation argument, with details in Appendix E. Throughout, "minimax" refers to the min–max criterion, not a claim of statistical minimax optimality.

5 Simulation studies

The simulations examine how test-set errors change with sample size in mean, expectile, and quantile IV regression. All three designs use ten-dimensional regressors and instruments; the expectile and quantile experiments consider $\tau \in \{0.25, 0.50, 0.75\}$. Data-generating designs, the Monte Carlo protocol, complete result tables, and training details appear in Appendix G. MSE and MAE compare fits with the score-specific structural truth, not observed outcomes, on a fixed independent test sample of 10,000 points. Results average 100 independent Monte Carlo replications; Monte Carlo standard errors (MCSEs) are conditional on this test design.

5.1 Simulation results

5.1.1 Mean IV results

Appendix Table 2 and Appendix Figure 2 show that mean-IV MSE and MAE decline throughout the sample-size grid, with the largest gains at smaller sample sizes. Both errors continue to fall beyond $n = 10{,}000$, and the accompanying MCSEs decrease as well.

5.1.2 Expectile IV results

Appendix Tables 3 and 4, together with Appendix Figure 3, show decreasing test-set errors at all three expectile levels. MSE falls by approximately 89–91% between $n = 2{,}000$ and $n = 50{,}000$, and both MSE and MAE decrease at every adjacent sample size. Errors at the two tail levels are similar, and the central expectile has the smallest error throughout the grid. At $\tau = 0.50$, the implemented score is the mean residual, so this path describes a mean-score fit under the bounded design. Its data-generating process (DGP) differs from the Gaussian mean-IV design used above. The quantile tails below show intermediate fluctuations. Since the disturbance laws also differ across designs, this contrast does not isolate the contribution of score continuity or smoothing.

5.1.3 Quantile IV results

Appendix Tables 5 and 6, together with Appendix Figure 4, show lower quantile-IV errors at $n = 50{,}000$ than at $n = 2{,}000$ for every reported quantile in the $d = p = 10$ design. Median MSE and MAE decline monotonically, whereas the two tail quantiles fluctuate at intermediate sample sizes. At $n = 50{,}000$, MSE is 0.0781, 0.0464, and 0.0688 for $\tau = 0.25, 0.50, 0.75$, respectively, and the corresponding MAE values are 0.1817, 0.1494, and 0.1707. The median therefore has the lowest error at the largest sample size on this fixed test design.

6 Real-data analysis

We compare the three score families using a diary study of 63 participants (1,023 usable person-days). Randomized daily encouragement is the instrument, pro-environmental behavior (PEB) the endogenous regressor, and eudaimonic well-being (EWB) the outcome. We use within-participant centered variables; the critic also includes lagged PEB, lagged EWB, and the day index. These augmented working moments require additional validity assumptions beyond randomization. Details appear in Appendix G.

6.1 Results

Encouragement was assigned on 52.7% of retained days; mean PEB and EWB were 3.154 and 3.585, respectively. Realized PEB averaged 0.082 scale points higher on encouraged days, consistent with the positive but modest first stage reported by Caldaroni et al. (2026).

Figure 1 reports contrasts $\widehat{f}(x) - \widehat{f}(0)$ on the original EWB scale, where $x = 0$ denotes a participant's average PEB level. The mean curve is increasing: relative to that level, the estimated EWB contrast is -0.40 at a one-point lower PEB deviation and 0.12 at a one-point higher deviation. The separately fitted central expectile and mean curves differ by at most 0.014 on the evaluation grid. Their close agreement is consistent with the common score at $\tau = 0.50$.

The noncentral expectile and quantile curves vary across levels and need not agree at a common τ because they represent distinct IV moment targets. Appendix G.4 reports the detailed pointwise contrasts.

The boundary behavior is sensitive to participant resampling. Structural interpretation is further qualified by the exclusion restriction, the modest first stage, 63 clusters, clustered sampling, and the discrete outcome. A full discussion appears in Appendix G.4.

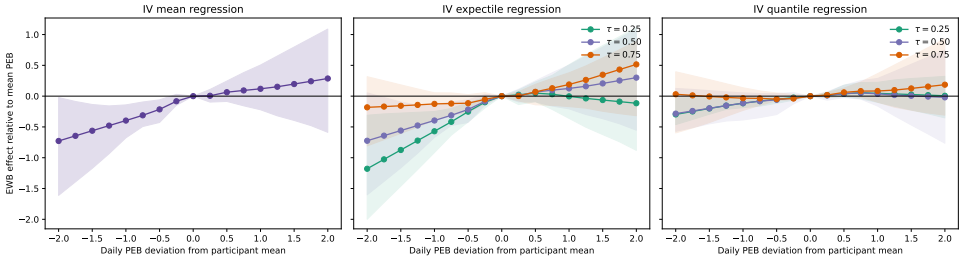

Figure 1: Estimated IV mean (left), expectile (center), and quantile (right) EWB contrasts over within-person PEB deviations, with pointwise percentile ranges from 20 participant-level bootstrap refits. Each curve is centered at zero PEB deviation. The ranges summarize sensitivity to participant resampling and are not confidence bands.

7 Conclusion

The proposed framework estimates nonparametric IV mean, expectile, and quantile targets through score-specific moments and a neural critic. The oracle inequalities separate statistical and approximation errors, with an additional smoothing term for quantiles. Under the stated conditions and common Hölder smoothness, the expected adversarial-moment risks of the theoretical sieve estimators have the same upper-bound order, up to logarithmic factors. Additional localization, link, and spectral conditions yield structural $L^2(P_X)$ upper bounds. Test-set errors generally decrease with sample size in the three simulation designs. In the diary application, the central expectile closely matches the separately fitted mean curve, while the noncentral curves vary across levels. Their structural interpretation depends on the working moment restrictions, and the modest first stage, 63 clusters, and discrete outcome limit the analysis. Theory for finite-step penalized training, clustered sampling, and inference remains to be developed.

AI use statement

Generative AI tools were used to assist with transcribing this manuscript into the ICLR 2027 LaTeX template, adapting figure and table formatting, and diagnosing compilation issues. They were not used to generate the paper's scientific claims, theoretical results, experimental data, or conclusions. All AI-assisted changes have been reviewed by the authors before submission. The authors take responsibility for the final content of this work, including text, claims, and artifacts produced with the aid of generative AI.

References

Javier Alejo, Antonio F. Galvao, and Gabriel Montes-Rojas. A first-stage representation for instrumental variables quantile regression. *The Econometrics Journal*, 26(3):350–377, 2023. doi: 10.1093/ectj/utad010.

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. doi: 10.1093/jrsssb/qkad025.

Andrew Bennett, Nathan Kallus, and Tobias Schnabel. Deep generalized method of moments for instrumental variable analysis. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

Andrew Bennett, Nathan Kallus, Xiaojie Mao, Whitney Newey, Vasilis Syrgkanis, and Masatoshi Uehara. Minimax instrumental variable regression and L_2 convergence guarantees without identification or closedness. In Gergely Neu and Lorenzo Rosasco (eds.), *Proceedings of the Thirty Sixth Conference on Learning Theory*, volume 195 of *Proceedings of Machine Learning Research*, pp. 2291–2318. PMLR, 2023.

Richard Blundell, Xiaohong Chen, and Dennis Kristensen. Semi-nonparametric IV estimation of shape-invariant engel curves. *Econometrica*, 75(6):1613–1669, 2007. doi: 10.1111/j.1468-0262.2007.00808.x.

Silvia Caldaroni, Maria Gerbino, Florian Schmiedek, Andreas B. Neubauer, Lucia Manfredi, Fulvio Gregori, Concetta Pastorelli, Giuseppe Corbelli, and Antonio Zuffianò. The positive effect of pro-environmental behavior on eudaimonic well-being in young adults: A daily diary study using the within-person encouragement design. *Journal of Personality*, 94(1):124–136, 2026. doi: 10.1111/jopy.13021.

Marine Carrasco, Jean-Pierre Florens, and Éric Renault. Linear inverse problems in structural econometrics: Estimation based on spectral decomposition and regularization. In James J. Heckman and Edward E. Leamer (eds.), *Handbook of Econometrics*, volume 6B, pp. 5633–5751. Elsevier, 2007.

Xiaohong Chen and Demian Pouzo. Estimation of nonparametric conditional moment models with possibly nonsmooth generalized residuals. *Econometrica*, 80(1):277–321, 2012.

Xiaohong Chen and Markus Reiss. On rate optimality for ill-posed inverse problems in econometrics. *Econometric Theory*, 27(3):497–521, 2011.

Xiaohong Chen, Victor Chernozhukov, Sokbae Lee, and Whitney K. Newey. Local identification of nonparametric and semiparametric models. *Econometrica*, 82(2):785–809, 2014. doi: 10.3982/ECTA9988.

Victor Chernozhukov and Christian Hansen. An IV model of quantile treatment effects. *Econometrica*, 73(1):245–261, 2005.

Victor Chernozhukov and Christian Hansen. Instrumental quantile regression inference for structural and treatment effect models. *Journal of Econometrics*, 132(2):491–525, 2006.

Victor Chernozhukov and Christian Hansen. Instrumental variable quantile regression: A robust inference approach. *Journal of Econometrics*, 142(1):379–398, 2008.

Victor Chernozhukov, Guido W. Imbens, and Whitney K. Newey. Instrumental variable estimation of nonseparable models. *Journal of Econometrics*, 139(1):4–14, 2007.

Serge Darolles, Yanqin Fan, Jean-Pierre Florens, and Eric Renault. Nonparametric instrumental regression. *Econometrica*, 79(5):1541–1565, 2011.

Nishanth Dikkala, Greg Lewis, Lester W. Mackey, and Vasilis Syrgkanis. Minimax estimation of conditional moment models. In *Advances in Neural Information Processing Systems*, volume 33, pp. 12248–12262, 2020.

Xingdong Feng, Xin He, Yuling Jiao, Lican Kang, and Caixing Wang. Deep nonparametric quantile regression under covariate shift. *Journal of Machine Learning Research*, 25(385):1–50, 2024.

Marcelo Fernandes, Víctor Henriques, and Eduardo Fonseca Mendes. Estimation risk in conditional expectiles. arXiv preprint arXiv:2609.02673, 2026.

Yuri R. Fonseca, Caio F. L. Peixoto, and Yuri F. Saporito. Nonparametric instrumental variable regression through stochastic approximate gradients. In *Advances in Neural Information Processing Systems*, volume 37, 2024. doi: 10.52202/079017-4189.

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.

Jason Hartford, Greg Lewis, Kevin Leyton-Brown, and Matt Taddy. Deep IV: A flexible approach for counterfactual prediction. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1414–1423, 2017.

Joel L. Horowitz and Sokbae Lee. Nonparametric instrumental variables estimation of a quantile regression model. *Econometrica*, 75(4):1191–1208, 2007.

David M Kaplan. Smoothed instrumental variables quantile regression. *The Stata Journal*, 22(2): 379–403, 2022.

David M Kaplan and Yixiao Sun. Smoothed estimating equations for instrumental variables quantile regression. *Econometric Theory*, 33(1):105–157, 2017.

Masahiro Kato, Masaaki Imaizumi, Kenichiro McAlinn, Haruo Kakehi, and Shota Yasui. Learning causal models from conditional moment restrictions by importance weighting. In *International Conference on Learning Representations*, 2022.

Juno Kim, Dimitri Meunier, Arthur Gretton, Taiji Suzuki, and Zhu Li. Optimality and adaptivity of deep neural features for instrumental variable regression. In *International Conference on Learning Representations*, 2025.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.

Greg Lewis and Vasilis Syrgkanis. Adversarial generalized method of moments. *CoRR*, abs/1803.07164, 2018.

Oscar Hernan Madrid Padilla, Wesley Tansey, and Yanzhen Chen. Quantile regression with ReLU networks: Estimators and minimax rates. *Journal of Machine Learning Research*, 23(247):1–42, 2022.

Krikamol Muandet, Arash Mehrjou, Si Kai Lee, and Anant Raj. Dual instrumental variable regression. In *Advances in Neural Information Processing Systems*, volume 33, pp. 2710–2721, 2020.

Whitney K. Newey and James L. Powell. Asymmetric least squares estimation and testing. *Econometrica*, 55(4):819–847, 1987. doi: 10.2307/1911031.

Whitney K Newey and James L Powell. Instrumental variable estimation of nonparametric models. *Econometrica*, 71(5):1565–1578, 2003.

Rahul Singh, Maneesh Sahani, and Arthur Gretton. Kernel instrumental variable regression. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

Fabian Sobotka, Rosalba Radice, Giampiero Marra, and Thomas Kneib. Estimating the relationship between women's education and fertility in Botswana by using an instrumental variable approach to semiparametric expectile regression. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 62(1):25–45, 2013. doi: 10.1111/j.1467-9876.2012.01050.x.

Yoon-Jae Whang. Smoothed empirical likelihood methods for quantile regression models. *Econometric Theory*, 22(2):173–205, 2006.

Liyuan Xu, Yutian Chen, Siddharth Srinivasan, Nando de Freitas, Arnaud Doucet, and Arthur Gretton. Learning deep features in instrumental variable regression. In *International Conference on Learning Representations*, 2021.

APPENDIX

Appendix A presents the related work. Appendix B collects notation and technical definitions, and Appendix C gives the detailed identification, network-approximation, and strong-norm conditions moved from the main text. Appendix D outlines the proofs, Appendix E contains the complete proofs, and Appendix F records auxiliary results. Appendix G presents the configurations for the numerical experiments.

A RELATED WORK

Identification and regularization in NPIV are studied through the conditional-expectation operator in Newey & Powell (2003); Hall & Horowitz (2005); Darolles et al. (2011); Carrasco et al. (2007); Chen & Reiss (2011). These works relate structural recovery rates to smoothness and ill-posedness. Our strong-norm analysis uses this operator perspective to interpret bounds on a learned moment criterion.

Several machine-learning IV estimators replace fixed sieves with learned conditional laws, learned features, kernels, or dual moment objectives. Deep IV estimates the conditional treatment distribution (Hartford et al., 2017), whereas deep-feature IV learns instrument and treatment representations in two stages (Xu et al., 2021); recent theory studies optimality and adaptivity of the latter construction (Kim et al., 2025). Kernel IV uses kernel embeddings and two-stage regularization (Singh et al., 2019). Dual IV uses a single-stage dual objective (Muandet et al., 2020), while importance weighting with an estimated conditional density ratio converts conditional restrictions into unconditional moments (Kato et al., 2022). Adversarial GMM and DeepGMM use learned critics to detect or combine moment violations (Lewis & Syrgkanis, 2018; Bennett et al., 2019), and Dikkala et al. (2020) derive finite-sample guarantees for general conditional-moment games. Strong L^2 guarantees for penalized minimax IV estimators are available under source and realizability conditions (Bennett et al., 2023), while functional stochastic-gradient methods give excess-risk bounds for neural or kernel NPIV estimators (Fonseca et al., 2024).

Expectile regression originates from the asymmetric least-squares characterization of Newey & Powell (1987). Sobotka et al. (2013) develop a two-stage semiparametric IV approach with flexible additive covariate effects and apply it to women's education and fertility in Botswana. Section 6.3 of Bennett & Kallus (2023) treats parametric IV mean and expectile regression using the same asymmetric residual moment as here, up to normalization, and establishes consistency, asymptotic normality, and semiparametric efficiency under its stated conditions. Their variational method of moments (VMM) uses a residual-dependent quadratic critic penalty motivated by optimal moment weighting. In contrast, our theoretical criterion takes the largest absolute moment over a bounded neural critic class. We analyze nonparametric approximation of both players, rather than propose a new IV expectile conditional-moment formulation. In a different setting, Fernandes et al. (2026) study estimation risk for conditional expectiles in time-series scale models. They estimate scale parameters by quasi-maximum likelihood and innovation expectiles from standardized residuals, accounting for first-step estimation error. Our IV target is defined by a residual-score restriction and does not require a preliminary scale model.

For IV quantiles, the foundational literature establishes identification, estimation, and inference under rank restrictions (Chernozhukov & Hansen, 2005; 2006; 2008). Nonparametric and sieve approaches formulate IV quantile regression as a nonlinear conditional-moment model with a nonsmooth generalized residual (Horowitz & Lee, 2007; Chernozhukov et al., 2007; Chen & Pouzo, 2012). Alejo et al. (2023) derive a density-weighted first-stage representation for studying quantile-specific instrument relevance and Jacobian identification conditions. The variational method of moments also treats instrumental quantile regression as an example of its general conditional-moment framework (Bennett & Kallus, 2023). Smoothing has been used for quantile and IV quantile estimating equations (Whang, 2006; Kaplan & Sun, 2017; Kaplan, 2022), while ReLU quantile regression has been analyzed without endogenous regressors (Madrid Padilla et al., 2022). Our quantile estimator places an integrated-kernel score inside the common neural minimax criterion used for all three targets, and its risk bound retains the bandwidth error.

B DEFINITIONS

ReLU network classes. Let $\mathrm{ReLU}(t) = \max\{t, 0\}$ denote the ReLU activation, applied coordinatewise to vectors. In this generic definition, let $d \in \mathbb{N}$ denote the input dimension. The structural and critic instances below use input dimensions d and p, respectively. Let $\mathcal{D} \geq 2$ denote the number of affine layers, so the network has $\mathcal{D} - 1$ hidden layers. Set $w_0 = d$ and $w_{\mathcal{D}} = 1$, and let $w_1, \ldots, w_{\mathcal{D}-1} \in \mathbb{N}$ be the hidden-layer widths. Write $\mathcal{W} = \max_{1 \leq j \leq \mathcal{D}-1} w_j$ for the maximal hidden-layer width. For $A_j \in \mathbb{R}^{w_j \times w_{j-1}}$ and $b_j \in \mathbb{R}^{w_j}$, $j = 1, \ldots, \mathcal{D}$, define recursively

$$\begin{aligned} z_0 &= x, \\ z_j &= \mathrm{ReLU}(A_j z_{j-1} + b_j), \quad j = 1, \ldots, \mathcal{D} - 1, \\ f_\Theta(x) &= A_{\mathcal{D}} z_{\mathcal{D}-1} + b_{\mathcal{D}}, \end{aligned} \tag{21}$$

where $\Theta = \{(A_\ell, b_\ell) : 1 \leq \ell \leq \mathcal{D}\}$. Define the network size and parameter magnitude by

$$\|\Theta\|_0 = \sum_{\ell=1}^{\mathcal{D}} \left\{ \|\operatorname{vec}(A_\ell)\|_0 + \|b_\ell\|_0 \right\}, \tag{22}$$

$$\|\Theta\|_\infty = \max_{1 \leq \ell \leq \mathcal{D}} \left\{ \|\operatorname{vec}(A_\ell)\|_\infty, \|b_\ell\|_\infty \right\}. \tag{23}$$

For positive constants $\mathcal{S}, \mathcal{B}_{\mathrm{par}}$, and $\mathcal{B}$, define

$$\mathcal{N}_d(\mathcal{W}, \mathcal{D}, \mathcal{S}, \mathcal{B}_{\mathrm{par}}, \mathcal{B}) = \left\{ f_\Theta : \mathbb{R}^d \to \mathbb{R} \;\middle|\; \begin{array}{ll} \max\limits_{1 \leq \ell \leq \mathcal{D}-1} w_\ell \leq \mathcal{W}, & \|\Theta\|_0 \leq \mathcal{S}, \\ \|\Theta\|_\infty \leq \mathcal{B}_{\mathrm{par}}, & \|f_\Theta\|_\infty \leq \mathcal{B} \end{array} \right\}. \tag{24}$$

The structural and adversarial classes are restrictions of $\mathcal{N}_d$ and $\mathcal{N}_p$ to $\mathcal{X}$ and $\mathcal{Z}$, respectively.

Hölder classes. Let $\alpha = (\alpha_1, \ldots, \alpha_d) \in \mathbb{N}_0^d$ be a multi-index, write $|\alpha|_1 = \sum_{j=1}^d \alpha_j$ and $D^\alpha v = \partial^{|\alpha|_1} v / (\partial x_1^{\alpha_1} \cdots \partial x_d^{\alpha_d})$, and let $\zeta = k + \gamma$, where $k \in \mathbb{N}_0$ and $\gamma \in (0, 1]$. This convention covers every $\zeta > 0$; when ζ is an integer, take $\gamma = 1$ and $k = \zeta - 1$. For $v : [0,1]^d \to \mathbb{R}$, define

$$\|v\|_{\mathcal{H}^\zeta} = \max_{|\alpha|_1 \leq k} \|D^\alpha v\|_\infty + \max_{|\alpha|_1 = k} \sup_{\substack{x, y \in [0,1]^d \\ x \neq y}} \frac{|D^\alpha v(x) - D^\alpha v(y)|}{\|x - y\|_2^\gamma}. \tag{25}$$

The Hölder ball of smoothness ζ and radius B is $\mathcal{H}^\zeta([0,1]^d, B) = \{v : [0,1]^d \to \mathbb{R} : \|v\|_{\mathcal{H}^\zeta} \leq B\}$.

Kernel functions. A function $K : \mathbb{R} \to [0, \infty)$ is called a kernel if it is measurable and satisfies

$$K(u) = K(-u), \qquad \int_{\mathbb{R}} K(u)\, du = 1, \qquad \mu_1(K) := \int_{\mathbb{R}} |u| K(u)\, du < \infty. \tag{26}$$

We also assume $\|K\|_\infty < \infty$; compactly supported kernels are allowed. For a bandwidth $h > 0$, the rescaled and integrated kernels are $K_h(u) = h^{-1} K(u/h)$ and $G(t) = \int_{-\infty}^t K(u)\, du$, respectively. In the numerical experiments (Sections 5–6), we use the standard logistic density and its integrated kernel,

$$K(u) = \frac{e^{-u}}{(1 + e^{-u})^2}, \qquad G(t) = \frac{1}{1 + e^{-t}}. \tag{27}$$

Kernel smoothing of the quantile score. The discontinuous map $v \mapsto \mathbf{1}\{v \leq 0\}$ is approximated by

$$\mathbf{I}_h(v) = G(-v/h) = \int_{\mathbb{R}} \mathbf{1}\{v + hu \leq 0\} K(u)\, du. \tag{28}$$

For every $v \neq 0$, $\mathbf{I}_h(v) \to \mathbf{1}\{v \leq 0\}$ as $h \downarrow 0$. The resulting smoothed quantile score $\phi_{\tau,h}(v) = \tau - \mathbf{I}_h(v)$ approximates $\phi_\tau(v) = \tau - \mathbf{1}\{v \leq 0\}$.

Covering numbers. Let $(\mathcal{M}, D)$ be a metric space and $\mathcal{V} \subseteq \mathcal{M}$. For $v \in \mathcal{M}$ and $\delta > 0$, write $B_D(v, \delta) = \{u \in \mathcal{M} : D(u, v) \leq \delta\}$. The δ-covering number of $\mathcal{V}$ under D is

$$\mathcal{N}(\delta, \mathcal{V}, D) = \inf \left\{ m \in \mathbb{N} : \begin{array}{c} \text{there exist } v_1, \ldots, v_m \in \mathcal{M} \text{ such that} \\ \mathcal{V} \subseteq \bigcup_{j=1}^{m} B_D(v_j, \delta) \end{array} \right\}. \tag{29}$$

C Additional theoretical details

C.1 Score-specific criterion and estimator details

Mean criterion. The mean model has a parallel construction without a smoothing parameter. Define the population and empirical moment functionals by

$$Q^m(f, g) = \mathbb{E}[\{Y - f(X)\}g(Z)], \tag{30}$$

$$\widehat{Q}^m(f, g) = \frac{1}{n} \sum_{i=1}^{n} \{Y_i - f(X_i)\}g(Z_i). \tag{31}$$

For any critic class $\mathcal{V}$, their adversarial counterparts are $\mathcal{L}^m_{\mathcal{V}}(f) = \sup_{g \in \mathcal{V}} |Q^m(f, g)|$ and $\widehat{\mathcal{L}}^m_{\mathcal{V}}(f) = \sup_{g \in \mathcal{V}} |\widehat{Q}^m(f, g)|$.

Expectile criterion. The expectile model also has no smoothing parameter. Define

$$Q^e(f, g) = \mathbb{E}[\psi_\tau\{Y - f(X)\}g(Z)], \tag{32}$$

$$\widehat{Q}^e(f, g) = \frac{1}{n} \sum_{i=1}^{n} \psi_\tau\{Y_i - f(X_i)\}g(Z_i). \tag{33}$$

For any critic class $\mathcal{V}$, let $\mathcal{L}^e_{\mathcal{V}}(f) = \sup_{g \in \mathcal{V}} |Q^e(f, g)|$ and $\widehat{\mathcal{L}}^e_{\mathcal{V}}(f) = \sup_{g \in \mathcal{V}} |\widehat{Q}^e(f, g)|$.

The mean and expectile estimators use the same neural classes:

$$\widehat{f}_m \in \operatorname*{argmin}_{f \in \mathcal{F}_n} \widehat{\mathcal{L}}^m_{\mathcal{G}_n}(f). \tag{34}$$

$$\widehat{f}_e \in \operatorname*{argmin}_{f \in \mathcal{F}_n} \widehat{\mathcal{L}}^e_{\mathcal{G}_n}(f). \tag{35}$$

C.2 Regularity assumptions for the risk bounds

Recall that P_X and P_Z denote the marginal laws of X and Z. The assumptions below specify the smoothness, boundedness, and local regularity conditions used to derive the rates.

Assumption C.1 (Smoothness and support)**.** *Without loss of generality, normalize the compact supports so that $\mathcal{X} \subseteq [0, 1]^d$ and $\mathcal{Z} \subseteq [0, 1]^p$. For some $\zeta > 0$ and $0 < B_0 < \infty$,*

$$\mathcal{F}^* = \mathcal{H}^\zeta([0, 1]^d, B_0), \qquad \mathcal{G}^* = \mathcal{H}^\zeta([0, 1]^p, B_0).$$

Assumption C.2 (Uniform boundedness and mean/expectile-outcome tails)**.** *Because the Hölder norm dominates the sup norm, Assumption C.1 already bounds the population classes by B_0. With the same constant, require that*

$$\sup_{f \in \mathcal{F}_n} \|f\|_{\infty,[0,1]^d} \leq B_0, \qquad \sup_{g \in \mathcal{G}_n} \|g\|_{\infty,[0,1]^p} \leq B_0.$$

For the mean- and expectile-regression results, impose in addition that Y is uniformly conditionally sub-Gaussian given Z: there is $K_Y < \infty$ such that

$$\operatorname*{ess\,sup}_{z \in \mathcal{Z}} \mathbb{E}\big[\exp\{Y^2/K_Y^2\} \mid Z = z\big] \leq 2. \tag{36}$$

Assumption C.3 (Conditional density regularity). *The conditional law of Y given $(X, Z) = (x, z)$ admits a density $p_{Y|X,Z}(\cdot \mid x, z)$. There exists $M_0 < \infty$ such that*

$$\sup_{\substack{x\in\mathcal{X},\ z\in\mathcal{Z}\\ u\in\mathbb{R}}} p_{Y|X,Z}(u \mid x, z) \le M_0.$$

The kernel satisfies (26) and $\|K\|_\infty < \infty$.

Assumption C.1 ensures uniform ReLU approximation of both players. Assumption C.2 bounds the network classes and imposes a uniform conditional sub-Gaussian tail on the outcome in the mean and expectile models. In particular, $\mathbb{P}(|Y| > t \mid Z) \le 2\exp(-t^2/K_Y^2)$ almost surely and all conditional moments of Y are uniformly finite. The quantile score is bounded, so the tail condition is not needed for the quantile results. The expectile score grows linearly in the residual, which is why it retains the sub-Gaussian condition rather than introducing a separate heavy-tail analysis. Assumption C.3, which is needed only for quantile regression, makes the conditional smoothed moment Lipschitz in f and bounds the population effect of indicator smoothing by order h.

C.3 Complete oracle inequalities

Theorem C.1 (Smoothed quantile risk bound). *Suppose that the score-q instance of Assumption 3.1, Assumptions C.1 and C.3, and the boundedness clause of Assumption C.2 hold. Let $\mathcal{F}_n$ and $\mathcal{G}_n$ be ReLU DNN classes with parameters $(\mathcal{W}_n, \mathcal{D}_n, \mathcal{S}_n, \mathcal{B}_{\mathrm{par},n})$. For every $0 < h \le 1$,*

$$\mathbb{E}\Big[\mathcal{L}^q_{\mathcal{G}^*,h}(\widehat{f}_h)\Big] \lesssim \left\{\frac{\mathcal{S}_n\mathcal{D}_n\log(e + n\mathcal{B}_{\mathrm{par},n}\mathcal{W}_n\mathcal{D}_n/h)}{n}\right\}^{1/2} + \inf_{f\in\mathcal{F}_n}\|f - f_0^q\|_\infty + d_{\to,1}(\mathcal{G}^*, \mathcal{G}_n) + h. \tag{37}$$

Under (18)–(19) and

$$h \asymp \varepsilon_n = n^{-\zeta/(d_{\max}+2\zeta)}, \tag{38}$$

the expected smoothed population risk satisfies

$$\mathbb{E}\Big[\mathcal{L}^q_{\mathcal{G}^*,h}(\widehat{f}_h)\Big] = O\Big(n^{-\zeta/(d_{\max}+2\zeta)}(\log n)^{3/2}\Big). \tag{39}$$

The next theorem transfers this bound to the original quantile criterion.

Theorem C.2 (Unsmoothed quantile excess risk). *Under the conditions of Theorem C.1, the network scaling (18)–(19), and the bandwidth (38),*

$$\mathbb{E}\Big[\mathcal{L}^q_{\mathcal{G}^*,\phi}(\widehat{f}_h)\Big] = O\Big(n^{-\zeta/(d_{\max}+2\zeta)}(\log n)^{3/2}\Big). \tag{40}$$

Because the mean score is Lipschitz without smoothing, it admits the parallel oracle bound below.

Theorem C.3 (Mean-regression oracle and excess-risk bound). *Suppose that the score-m instance of Assumption 3.1 and Assumptions C.1 and C.2 hold. Let $\mathcal{F}_n$ and $\mathcal{G}_n$ be ReLU DNN classes with the parameters used above. Then*

$$\mathbb{E}\Big[\mathcal{L}^m_{\mathcal{G}^*}(\widehat{f}_m)\Big] \lesssim \left\{\frac{\mathcal{S}_n\mathcal{D}_n\log(e + n\mathcal{B}_{\mathrm{par},n}\mathcal{W}_n\mathcal{D}_n)}{n}\right\}^{1/2} + \inf_{f\in\mathcal{F}_n}\|f - f_0^m\|_\infty + d_{\to,1}(\mathcal{G}^*, \mathcal{G}_n). \tag{41}$$

Under (18)–(19), it satisfies

$$\mathbb{E}\Big[\mathcal{L}^m_{\mathcal{G}^*}(\widehat{f}_m)\Big] = O\Big\{n^{-\zeta/(d_{\max}+2\zeta)}(\log n)^{3/2}\Big\}. \tag{42}$$

Because the expectile score is globally Lipschitz and has the same linear envelope as the mean score, it admits the following parallel result under the same sub-Gaussian outcome condition.

Theorem C.4 (Expectile-regression oracle and excess-risk bound)**.** *Suppose that the score-e instance of Assumption 3.1 and Assumptions C.1 and C.2 hold. Let $\mathcal{F}_n$ and $\mathcal{G}_n$ be ReLU DNN classes with the parameters used above. Then*

$$\mathbb{E}\Big[\mathcal{L}^e_{\mathcal{G}^*}(\widehat{f}_e)\Big] \lesssim \left\{\frac{\mathcal{S}_n\mathcal{D}_n\log(e+n\mathcal{B}_{\mathrm{par},n}\mathcal{W}_n\mathcal{D}_n)}{n}\right\}^{1/2} + \inf_{f\in\mathcal{F}_n}\|f-f^e_0\|_\infty + d_{\to,1}(\mathcal{G}^*,\mathcal{G}_n). \tag{43}$$

Under (18)–(19),

$$\mathbb{E}\Big[\mathcal{L}^e_{\mathcal{G}^*}(\widehat{f}_e)\Big] = O\Big\{n^{-\zeta/(d_{\max}+2\zeta)}(\log n)^{3/2}\Big\}. \tag{44}$$

Since $\mathcal{L}^e_{\mathcal{G}^*}(f^e_0)=0$, both displays are also excess-risk bounds.

C.4 Sufficient identification conditions

For the mean score, $M_m(f) = -\mathbb{E}\{f(X) - f^m_0(X) \mid Z\}$. Hence completeness of the conditional law of X given Z on the difference class $\mathcal{F}^* - f^m_0$ is sufficient for clause 1 of Assumption 3.1 (Newey & Powell, 2003). On $\mathcal{F}^* - f^m_0$, completeness is equivalent to injectivity of the conditional-expectation operator. Requiring completeness on all of $L^2(P_X)$ would be stronger than necessary.

For expectile regression, put $R^e_0 = Y - f^e_0(X)$ and $\Delta_f = f - f^e_0$. The continuity and piecewise linearity of ψ_τ give the exact path identity

$$M_e(f)(Z) = -\mathbb{E}[\overline{w}_{\tau,f}(X,Y)\Delta_f(X)\,|\,Z], \qquad \overline{w}_{\tau,f}(X,Y) = \int_0^1 2w_\tau\{R^e_0 - t\Delta_f(X)\}\,dt, \tag{45}$$

where validity gives $M_e(f^e_0)=0$. Moreover,

$$2\min\{\tau,1-\tau\} \le \overline{w}_{\tau,f}(X,Y) \le 2\max\{\tau,1-\tau\}.$$

Thus a primitive sufficient condition for global expectile identification is pathwise weighted completeness: for every $f\in\mathcal{F}^*$, the conditional expectation on the right-hand side of (45) vanishes almost surely only if $f = f^e_0$ in $L^2(P_X)$. At $\tau = 1/2$, the weight equals one and this condition reduces to the mean completeness condition. Away from $\tau = 1/2$, ordinary mean completeness does not by itself imply the required weighted completeness because the path weight depends on (X,Y) as well as on the candidate f.

For quantile regression, the moment map is nonlinear (Horowitz & Lee, 2007). If the conditional density exists along every segment joining f^q_0 and f, let

$$\overline{p}_f(X,Z) = \int_0^1 p_{Y|X,Z}\{f^q_0(X) + t(f(X) - f^q_0(X)) \mid X, Z\}\,dt.$$

The conditional mean-value identity gives

$$M_q(f) = -\mathbb{E}\big[\overline{p}_f(X,Z)\{f(X) - f^q_0(X)\}\,\big|\,Z\big],$$

where validity implies $M_q(f^q_0)=0$. A primitive sufficient condition for global quantile identification is therefore pathwise weighted completeness: for every $f\in\mathcal{F}^*$, the right-hand side vanishes almost surely only if $f = f^q_0$ in $L^2(P_X)$. The condition is global because $\overline{p}_f$ averages the density along the entire segment from f^q_0 to f. When the conditional density is continuous, its local limit is $p_{Y|X,Z}\{f^q_0(X) \mid X, Z\}$, which yields the derivative-weighted operator underlying the modified-completeness condition for local identification in Chernozhukov et al. (2007). This distinction is substantive: ordinary mean completeness alone does not establish global quantile identification. Moreover, injectivity of the derivative is only a first-order rank condition. Together with suitable differentiability and control of the nonlinear remainder, it can yield local identification; global identification instead requires restrictions over all relevant candidates or paths (Chernozhukov et al., 2007; Chen et al., 2014).

Clause 2 of Assumption 3.1 holds automatically for the ideal critic class in Assumption C.1: the linear span of a positive-radius Hölder ball on the compact instrument domain is dense in $L^2(P_Z)$, so a nonzero $M_s(f)$ cannot be orthogonal to every $g\in\mathcal{G}^*$. For the theoretical classes used here, the substantive identification requirement is therefore the score-specific injectivity in clause 1; critic separation ensures that the population minimax loss inherits it.

C.5 Network approximation and existence details

The constants in (19) are large enough to contain the networks in Lemma F.2 for input dimensions d and p at accuracy ε_n. That construction has width of order $\varepsilon_n^{-d_{\max}/\zeta}$ and an additional $\log(1/\varepsilon_n)$ factor in the number of nonzero parameters. A fixed ReLU subnetwork implements clipping to $[-B_0, B_0]$ without changing the orders in (19). Hence,

$$\sup_{v\in\mathcal{F}^*}\inf_{f\in\mathcal{F}_n}\|f-v\|_\infty+\sup_{g^*\in\mathcal{G}^*}\inf_{g\in\mathcal{G}_n}\|g-g^*\|_\infty\lesssim\varepsilon_n. \tag{46}$$

For each n, each fixed-depth, fixed-width parameterization satisfying the sparsity and parameter bounds is a finite union of compact sets. The output constraint is closed, and the empirical criteria are continuous in the parameters. Because only finitely many depths and widths occur, the inner maxima and outer minima are attained. The Hölder balls $\mathcal{F}^*$ and $\mathcal{G}^*$ are compact in the uniform norm, and the smoothed population moment is jointly continuous on their product. The resulting supremum risk is therefore continuous on $\mathcal{F}^*$, so a smoothed population minimizer f_h^* exists. The measurable maximum theorem supplies measurable selections of the empirical minimizers; such selections are used throughout.

C.6 Strong structural norms: operators and regularity

For mean regression, define the conditional-expectation operator $(T_m\delta)(z)=\mathbb{E}\{\delta(X)\mid Z=z\}$, where $T_m: L^2(P_X)\to L^2(P_Z)$. For expectile regression, define

$$\omega_{\tau,0}(x,z)=\mathbb{E}[2w_\tau\{Y-f_0^e(X)\}\,|\,X=x,Z=z]\,, \tag{47}$$

$$(T_e\delta)(z)=\mathbb{E}[\omega_{\tau,0}(X,Z)\delta(X)\,|\,Z=z]\,. \tag{48}$$

The bounds on w_τ make $T_e: L^2(P_X)\to L^2(P_Z)$ bounded. If the conditional law assigns no mass to $Y=f_0^e(X)$, dominated convergence and the piecewise linear expectile score imply, for every $\delta\in L^2(P_X)$,

$$\left\|\frac{M_e(f_0^e+t\delta)-M_e(f_0^e)}{t}+T_e\delta\right\|_{L^2(P_Z)}\longrightarrow 0\qquad(t\to 0). \tag{49}$$

For quantile regression, fix a jointly measurable version of the conditional density and let

$$(T_q\delta)(z)=\mathbb{E}\big[p_{Y|X,Z}\{f_0^q(X)\mid X,Z\}\delta(X)\,\big|\,Z=z\big]\,.$$

Boundedness in Assumption C.3 and the pointwise continuity imposed in clause 1 of Assumption C.4 imply, by dominated convergence, that, for each $\delta\in L^2(P_X)$,

$$\left\|\frac{M_q(f_0^q+t\delta)-M_q(f_0^q)}{t}+T_q\delta\right\|_{L^2(P_Z)}\longrightarrow 0\qquad(t\to 0).$$

For either nonlinear score, these pointwise derivative conditions alone need not yield a Fréchet remainder bound uniformly over local directions. The projected tangential condition in Assumption C.4 supplies the uniform local comparison used in the proof. Such derivative-weighted operators govern local identification and local ill-posedness in nonlinear IV expectile and quantile models (Chernozhukov et al., 2007; Chen et al., 2014). Write M for M_m, M_e, or M_q, T for the corresponding operator, f_0 for f_0^m, f_0^e, or f_0^q, and $\mathcal{R}(f)$ for $\mathcal{L}^m_{\mathcal{G}^*}(f)$, $\mathcal{L}^e_{\mathcal{G}^*}(f)$, or $\mathcal{L}^q_{\mathcal{G}^*,\phi}(f)$. For the mean score, $M_m(f)-M_m(f_0^m)=-T_m(f-f_0^m)$. Hence T_m, T_e, and T_q are defined as positive forward operators, while the corresponding moment map has first-order term $-T(f-f_0^s)$; the sign disappears upon taking the norms below.

Regularization is intrinsic to this problem. If T is compact with singular system

$$(\sigma_j,u_j,v_j)_{j\ge1},\qquad \sigma_j\downarrow 0,$$

take $\delta^{(j)}=u_j$. Then $\|\delta^{(j)}\|_{L^2(P_X)}=1$, while $\|T\delta^{(j)}\|_{L^2(P_Z)}=\sigma_j\to 0$. Hence convergence of the entire forward moment—the linearized moment in the expectile or quantile model—in $L^2(P_Z)$ does not imply uniform structural convergence over unrestricted directions. This is the standard ill-posedness mechanism in NPIV models (Hall & Horowitz, 2005; Darolles et al., 2011).

Assumption C.4 (Local link and spectral regularity)**.** *For some $\beta>0$, there is a local parameter set $\mathcal{U}_\beta$ containing f_0 such that, for every $f\in\mathcal{U}_\beta$, $f-f_0\in L^2(P_X)$ and $M(f)-M(f_0)\in L^2(P_Z)$. The conditions below hold for the score under consideration.*

1. Linearization and singular system. *In the quantile model, the jointly measurable density version defining T_q is continuous in its outcome argument at $f_0^q(x)$ for P_{XZ}-almost every (x, z). In the expectile model, $\mathbb{P}\{Y = f_0^e(X) \mid X, Z\} = 0$ almost surely, and $T = T_e$ is defined by (48). The bounded linear operator $T : L^2(P_X) \to L^2(P_Z)$ is compact and injective. It has singular triplets $(\sigma_j, u_j, v_j)_{j\geq 1}$ satisfying $Tu_j = \sigma_j v_j$, $T^* v_j = \sigma_j u_j$, and $\sigma_1 \geq \sigma_2 \geq \cdots > 0$, where $(u_j)_{j\geq 1}$ is a complete orthonormal basis of $L^2(P_X)$, and $(v_j)_{j\geq 1}$ is orthonormal in $L^2(P_Z)$. Let Π_J be the orthogonal projection onto $V_J = \operatorname{span}(v_1, \ldots, v_J)$.*

2. Projected critic norming. *For some $b \geq 0$, the critic risk controls each projected original conditional moment:*

$$\|\Pi_J\{M(f) - M(f_0)\}\|_{L^2(P_Z)} \lesssim J^b \mathcal{R}(f), \qquad J \geq 1, \quad f \in \mathcal{U}_\beta.$$

3. Projected tangential condition. *For the expectile and quantile models, the nonlinear moment map also obeys the following projected tangential condition, where $s = e$ in the expectile model and $s = q$ in the quantile model:*

$$\|\Pi_J\{M_s(f) - M_s(f_0^s) + T_s(f - f_0^s)\}\|_{L^2(P_Z)} \leq c_T \|\Pi_J T_s(f - f_0^s)\|_{L^2(P_Z)},$$

for every $J \geq 1$ and $f \in \mathcal{U}_\beta$, where $0 \leq c_T < 1$. This clause is vacuous for the linear mean model.

4. Spectral source restriction. *For this spectral smoothness index β,*

$$\sum_{j\geq 1} j^{2\beta} \langle f - f_0, u_j \rangle^2_{L^2(P_X)} \lesssim 1, \qquad f \in \mathcal{U}_\beta.$$

Clause 2 is a projected norming condition for the adversarial loss. It is implied by

$$\sup_{\substack{v \in V_J \\ \|v\|_{L^2(P_Z)} \leq 1}} \inf\{a > 0 : v \in a\mathcal{G}^*\} \lesssim J^b,$$

because Hilbert-space duality then yields the projected-moment bound. The factor J^b measures the cost of representing $L^2(P_Z)$-normalized sieve directions by the bounded critic class; increasingly oscillatory directions may require larger Hölder norms. This factor is analogous to a sieve norm-equivalence factor, but it differs from σ_J^{-1}, which measures structural ill-posedness (Blundell et al., 2007; Chen & Reiss, 2011; Dikkala et al., 2020).

Clause 3 is a projected tangential-cone condition preventing the nonlinear expectile or quantile remainder from overwhelming its first-order term; related local conditions are central to nonlinear nonparametric identification and estimation (Chen & Pouzo, 2012; Chen et al., 2014). Clauses 1 and 3 also verify local identification for the expectile and quantile maps on $\mathcal{U}_\beta$: if $M_s(f) = M_s(f_0^s)$ for the corresponding score s, then $(1 - c_T)\|\Pi_J T_s(f - f_0^s)\|_{L^2(P_Z)} = 0$ for every J, and completeness of $(u_j)_{j\geq 1}$ gives $f = f_0^s$ in $L^2(P_X)$. The corresponding implication for the mean map follows directly from linearity and injectivity of T_m.

Clause 4 is a spectral-ellipsoid restriction controlling the high-frequency tail of $f - f_0$. It is a source-type approximation condition in the singular basis; the local set $\mathcal{U}_\beta$ need not be an open $L^2(P_X)$ ball. The spectral index β controls coefficients in the singular basis of T, whereas the Hölder index ζ controls neural approximation. They describe different restrictions and are not identified with one another.

D Proof sketch

This appendix outlines the arguments underlying Theorems 4.1 and 4.2; Appendix E contains the complete proofs. We first control the smoothed criterion and transfer the resulting bound to the original quantile criterion. The mean, expectile, and strong-norm results require the modifications described below.

D.1 Proof strategy for Theorem C.1

For brevity, write $L^*(f) = \mathcal{L}^q_{\mathcal{G}^*,h}(f)$, $L_n(f) = \mathcal{L}^q_{\mathcal{G}_n,h}(f)$, and $\widehat{L}_n(f) = \widehat{\mathcal{L}}^q_{\mathcal{G}_n,h}(f)$. Define

$$\mathcal{E}_{\mathrm{sta},n} = \sup_{f\in\mathcal{F}_n} |\widehat{L}_n(f) - L_n(f)|, \qquad \mathcal{E}_{\mathcal{F},n} = \inf_{f\in\mathcal{F}_n} L_n(f), \qquad \mathcal{E}_{\mathcal{G},n} = \sup_{f\in\mathcal{F}_n} \{L^*(f) - L_n(f)\}.$$

Inserting L_n and $\widehat{L}_n$ and using the empirical optimality of $\widehat{f}_h$ yields

$$L^*(\widehat{f}_h) \le 2\mathcal{E}_{\mathrm{sta},n} + \mathcal{E}_{\mathcal{F},n} + \mathcal{E}_{\mathcal{G},n}. \tag{50}$$

Statistical error. For the composite class $\mathscr{H}_h = \{(x,y,z) \mapsto \phi_{\tau,h}\{y - f(x)\}g(z) : f \in \mathcal{F}_n,\ g \in \mathcal{G}_n\}$, symmetrization yields a Rademacher process indexed by both neural players. The sample-level Lipschitz constant of the smoothed score is $O(h^{-1})$, but product covers place this factor inside the entropy logarithm. Lemma D.2 gives the resulting bound; there is no multiplicative h^{-1} loss in the statistical rate.

Approximation terms. We use a network approximation to the unsmoothed target f_0^q as the comparator. Unlike the sample score, the smoothed population moment is Lipschitz in f with a constant independent of h, by the conditional density bound. Together with the $O(h)$ moment at f_0^q, this gives $\mathcal{E}_{\mathcal{F},n} \lesssim \inf_{f\in\mathcal{F}_n} \|f - f_0^q\|_\infty + h$. Lemma D.3 bounds the critic term by the directed distance (20). Only the comparison $L^*(f) \le L_n(f) + O(\varepsilon_n)$ is needed: the neural critics need not lie in $\mathcal{G}^*$, and no reverse approximation is required. The uniform approximation in Lemma F.2 supplies this one-sided $L^1(P_Z)$ control. It follows that $\mathcal{E}_{\mathcal{F},n} + \mathcal{E}_{\mathcal{G},n} \lesssim \varepsilon_n$ at $h \asymp \varepsilon_n$. Substitution into (50) proves Theorem C.1.

D.2 Proof strategy for Theorem C.2

The transfer to the original score rests on the uniform smoothing bound

$$\sup_{f\in\mathcal{F}_n\cup\mathcal{F}^*} \left|\mathcal{L}^q_{\mathcal{G}^*,\phi}(f) - \mathcal{L}^q_{\mathcal{G}^*,h}(f)\right| \lesssim h. \tag{51}$$

Evaluating this bound at $\widehat{f}_h$, taking expectations over the training sample, and applying Theorem C.1 add only an $O(h)$ term. The choice $h \asymp \varepsilon_n$ therefore gives Theorem C.2. Since the original risk at f_0^q is zero, the result is also an excess-risk bound.

Remark D.1 (Mean and expectile regression)**.** *The proofs of Theorems C.3 and C.4 use the same deterministic decomposition as* (50)*. The identity residual and the expectile score are globally Lipschitz without smoothing, so the entropy bound has no h^{-1} factor and the transfer step* (51) *is absent. Both scores have a linear envelope in Y, so the conditional sub-Gaussian assumption controls their empirical envelopes; the expectile Lipschitz and envelope constants depend only on the fixed level τ.*

Strong-norm conversion. The three original-risk bounds and Markov's inequality give $\mathcal{R}(\widehat{f}) = O_P(r_n)$. On the localization event $\widehat{f} \in \mathcal{U}_\beta$, Assumption C.4 controls the first J singular components by projected moments and the remaining components by the source condition:

$$\|\widehat{f} - f_0\|^2_{L^2(P_X)} \lesssim \sigma_J^{-2} J^{2b} \mathcal{R}(\widehat{f})^2 + J^{-2\beta}.$$

Balancing the two terms gives $J \asymp r_n^{-1/(a+b+\beta)}$ under polynomial inverse growth. Under exponential inverse growth, take $J = \lfloor c\{\log(1/r_n)\}^{1/a} \rfloor$ with $0 < c < c_2^{-1/a}$, so the stochastic term is smaller than the spectral tail. The localization event has probability tending to one, which yields the unconditional rates in Theorem 4.2.

D.3 Formal component bounds

The following lemmas formalize the decomposition, statistical bound, and approximation comparison used above.

Lemma D.1 (Deterministic oracle decomposition)**.** *Let $\mathcal{F}$ and $\mathcal{G}$ be nonempty approximation classes, $\widehat{f} \in \operatorname{argmin}_{f\in\mathcal{F}} \widehat{\mathcal{L}}_{\mathcal{G}}(f)$. Let $\mathcal{L}_{\mathcal{G}^*}$, $\mathcal{L}_{\mathcal{G}}$, and $\widehat{\mathcal{L}}_{\mathcal{G}}$ denote the three corresponding risks for either the smoothed quantile criterion or the mean or expectile criterion. Then*

$$\mathcal{L}_{\mathcal{G}^*}(\widehat{f}) \le \underbrace{2\sup_{f\in\mathcal{F}}\left|\widehat{\mathcal{L}}_{\mathcal{G}}(f) - \mathcal{L}_{\mathcal{G}}(f)\right|}_{\mathcal{E}_{\mathrm{sta}}} + \underbrace{\inf_{f\in\mathcal{F}} \mathcal{L}_{\mathcal{G}}(f)}_{\mathcal{E}_{\mathcal{F}}} + \underbrace{\sup_{f\in\mathcal{F}}\{\mathcal{L}_{\mathcal{G}^*}(f) - \mathcal{L}_{\mathcal{G}}(f)\}}_{\mathcal{E}_{\mathcal{G}}}.$$

Statistical error. The next lemma controls statistical error uniformly over both minimax players. The logarithmic dependence on h^{-1} arises from the Lipschitz constant of the smoothed quantile score.

Lemma D.2 (Uniform statistical error)**.** *Suppose the observations are i.i.d., K satisfies (26), $\|K\|_\infty < \infty$, and $\mathcal{F}$ and $\mathcal{G}$ are ReLU DNN classes with width at most $\mathcal{W}$, depth at most $\mathcal{D}$, size at most $\mathcal{S}$, parameter bound at most $\mathcal{B}_{\mathrm{par}} \ge 1$, and uniform output bound at most B_0. If $0 < h \le 1$ and $\widehat{\mathcal{L}}_{\mathcal{G}}$ and $\mathcal{L}_{\mathcal{G}}$ denote the smoothed quantile risks at bandwidth h, then*

$$\mathbb{E}\sup_{f\in\mathcal{F}}\left|\widehat{\mathcal{L}}_{\mathcal{G}}(f) - \mathcal{L}_{\mathcal{G}}(f)\right| \lesssim \left\{\frac{\mathcal{S}\mathcal{D}\log(e + n\mathcal{B}_{\mathrm{par}}\mathcal{W}\mathcal{D}/h)}{n}\right\}^{1/2}.$$

Approximation error. Lemma D.3 converts approximation error for the structural and critic classes into criterion error.

Lemma D.3 (Smoothed-quantile criterion approximation)**.** *Suppose the score-q instance of Assumption 3.1, Assumption C.3, and the boundedness clause of Assumption C.2 hold. Let $\mathcal{F}$ and $\mathcal{G}$ be nonempty classes bounded by B_0 on the respective support cubes. For the smoothed quantile criterion and $0 < h \le 1$,*

$$\mathcal{E}_{\mathcal{F}} \lesssim \inf_{f\in\mathcal{F}} \|f - f_0^q\|_\infty + h, \quad \mathcal{E}_{\mathcal{G}} \le d_{\to,1}(\mathcal{G}^*, \mathcal{G}).$$

E Proofs

This appendix proves the identification, risk, and strong-norm results in Section 4. The proofs combine the oracle decomposition with statistical and approximation bounds developed in Appendix D.

E.1 Proof of score-specific point identification

Proof of Lemma 4.1. Fix $s \in \{m, e, q\}$ and $f \in \mathcal{F}^*$. By the tower property, for every $g \in \mathcal{G}^*$,

$$\mathbb{E}\{\vartheta_f^s(X, Y)g(Z)\} = \mathbb{E}\{M_s(f)(Z)g(Z)\}.$$

If $M_s(f) = 0$ almost surely, every unconditional moment tested by the critic class vanishes, so $\mathcal{L}^s_{\mathcal{G}^*}(f) = 0$, with $\mathcal{L}^q_{\mathcal{G}^*} = \mathcal{L}^q_{\mathcal{G}^*,\phi}$. Conversely, suppose that this risk is zero. If $M_s(f)$ were nonzero in $L^2(P_Z)$, the critic-separation clause of Assumption 3.1 would yield $g \in \mathcal{G}^*$ for which $|\mathbb{E}\{M_s(f)(Z)g(Z)\}| > 0$, contradicting the zero risk. Hence $M_s(f) = 0$ almost surely. Conditional-moment identification then gives $f = f_0^s$ in $L^2(P_X)$. Finally, validity at f_0^s and the tower property imply $\mathcal{L}^s_{\mathcal{G}^*}(f_0^s) = 0$, proving the reverse implication and uniqueness up to P_X-null sets. □

E.2 Proof of the oracle decomposition

Proof of Lemma D.1. Fix $f \in \mathcal{F}$. Add and subtract the empirical and population criteria over $\mathcal{G}$:

$$\begin{aligned}\mathcal{L}_{\mathcal{G}^*}(\widehat{f}) = {} & \{\mathcal{L}_{\mathcal{G}^*}(\widehat{f}) - \mathcal{L}_{\mathcal{G}}(\widehat{f})\} + \{\mathcal{L}_{\mathcal{G}}(\widehat{f}) - \widehat{\mathcal{L}}_{\mathcal{G}}(\widehat{f})\} \\ & + \{\widehat{\mathcal{L}}_{\mathcal{G}}(\widehat{f}) - \widehat{\mathcal{L}}_{\mathcal{G}}(f)\} + \{\widehat{\mathcal{L}}_{\mathcal{G}}(f) - \mathcal{L}_{\mathcal{G}}(f)\} \\ & + \mathcal{L}_{\mathcal{G}}(f).\end{aligned}$$

Empirical optimality makes the third term nonpositive. The second and fourth terms are bounded by $\mathcal{E}_{\mathrm{sta}}$, and the first is bounded by $\mathcal{E}_{\mathcal{G}}$. Taking the infimum over $f \in \mathcal{F}$ gives $\mathcal{L}_{\mathcal{G}^*}(\widehat{f}) \le 2\mathcal{E}_{\mathrm{sta}} + \mathcal{E}_{\mathcal{F}} + \mathcal{E}_{\mathcal{G}}$. □

E.3 Proof of the statistical-error bound

Proof of Lemma D.2. Let P_n and P denote the empirical and population measures, and write $\ell_{f,g}(x,y,z) = \phi_{\tau,h}\{y - f(x)\}g(z)$ and $\mathscr{H}_h = \{\ell_{f,g} : f \in \mathcal{F}, g \in \mathcal{G}\}$. The elementary inequality $|\sup_g a_g - \sup_g b_g| \le \sup_g |a_g - b_g|$ yields $\mathcal{E}_{\mathrm{sta}} \le \sup_{f\in\mathcal{F},g\in\mathcal{G}} |(P_n - P)\ell_{f,g}|$. By symmetrization, for independent Rademacher variables $\xi_1, \dots, \xi_n$,

$$\mathbb{E}\mathcal{E}_{\mathrm{sta}} \le 2\mathbb{E} \sup_{\ell\in\mathscr{H}_h} \left| \frac{1}{n} \sum_{i=1}^{n} \xi_i \ell(X_i, Y_i, Z_i) \right|.$$

Let $L_K = \|K\|_\infty$. Since $\phi_{\tau,h}$ is L_K/h-Lipschitz and all critic functions are uniformly bounded,

$$\|\ell_{f_1,g_1} - \ell_{f_2,g_2}\|_\infty \lesssim h^{-1}\|f_1 - f_2\|_\infty + \|g_1 - g_2\|_\infty.$$

It follows that the product of an $O(h\delta)$-cover of $\mathcal{F}$ and an $O(\delta)$-cover of $\mathcal{G}$ is a δ-cover of $\mathscr{H}_h$. Specifically, for suitable constants $c_1, c_2 > 0$,

$$\mathcal{N}(\delta, \mathscr{H}_h, \|\cdot\|_\infty) \le \mathcal{N}(c_1 h\delta, \mathcal{F}, \|\cdot\|_\infty) \times \mathcal{N}(c_2\delta, \mathcal{G}, \|\cdot\|_\infty).$$

Applying Lemma F.3 to both network classes gives, for $0 < \delta < 1$,

$$\log \mathcal{N}(\delta, \mathscr{H}_h, \|\cdot\|_\infty) \lesssim \mathcal{SD} \log\left(\frac{\mathcal{B}_{\mathrm{par}}\mathcal{WD}}{h\delta}\right).$$

The entropy of the $\mathcal{G}$-factor is dominated by the displayed quantity because $h \le 1$. For $\delta \ge 1$, monotonicity of covering numbers yields the combined bound

$$\log \mathcal{N}(\delta, \mathscr{H}_h, \|\cdot\|_\infty) \lesssim \mathcal{SD} \log\left(e + \frac{\mathcal{B}_{\mathrm{par}}\mathcal{WD}}{h(\delta \wedge 1)}\right).$$

Conditionally on the observations, the empirical $L^2(P_n)$ metric is bounded by the uniform metric. Hence the truncated Dudley entropy bound implies, for $0 < \delta_{\mathrm{tr}} < \min\{1, 2B_0\}$,

$$\begin{aligned}\mathbb{E}\mathcal{E}_{\mathrm{sta}} &\lesssim \delta_{\mathrm{tr}} + \frac{1}{\sqrt{n}} \int_{\delta_{\mathrm{tr}}}^{2B_0} \sqrt{\mathcal{SD} \log\left(e + \frac{\mathcal{B}_{\mathrm{par}}\mathcal{WD}}{h(u \wedge 1)}\right)}\, du \\ &\lesssim \delta_{\mathrm{tr}} + \left\{ \frac{\mathcal{SD} \log(e + \mathcal{B}_{\mathrm{par}}\mathcal{WD}/(h\delta_{\mathrm{tr}}))}{n} \right\}^{1/2}.\end{aligned}$$

The last step bounds the logarithm on $[\delta_{\mathrm{tr}}, 2B_0]$ by its value at δ_{tr}. Taking $\delta_{\mathrm{tr}} = n^{-1}$ and absorbing the resulting n^{-1} term yields

$$\mathbb{E}\mathcal{E}_{\mathrm{sta}} \lesssim \left\{ \frac{\mathcal{SD} \log(e + n\mathcal{B}_{\mathrm{par}}\mathcal{WD}/h)}{n} \right\}^{1/2}.$$

□

E.4 Proofs for the approximation bounds

Proof of Lemma D.3. For the structural approximation term, fix $f \in \mathcal{F}$ and define $\Delta(X) = f(X) - f_0^q(X)$. Lemma F.5 and the conditional quantile restriction imply, uniformly over $g \in \mathcal{G}$,

$$|Q_h(f, g)| \lesssim \|\Delta\|_\infty + h.$$

Taking the infimum over $f \in \mathcal{F}$ proves the first claim.

For the critic approximation term, use $|\phi_{\tau,h}| \le 1$ and define $a_f(g) = |Q_h(f, g)|$. The elementary comparison gives

$$\sup_{g^*\in\mathcal{G}^*} a_f(g^*) - \sup_{g\in\mathcal{G}} a_f(g) \le \sup_{g^*\in\mathcal{G}^*} \inf_{g\in\mathcal{G}} |a_f(g^*) - a_f(g)|.$$

Because $|\phi_{\tau,h}| \le 1$,

$$|a_f(g_1) - a_f(g_2)| \le |Q_h(f, g_1) - Q_h(f, g_2)| \le \|g_1 - g_2\|_{L^1(P_Z)}.$$

Taking the supremum over f yields $\mathcal{E}_{\mathcal{G}} \le d_{\to,1}(\mathcal{G}^*, \mathcal{G})$. □

E.5 PROOFS OF THE MAIN RISK BOUNDS

Proof of Theorem C.1. In Lemma D.1, take $\mathcal{L}_{\mathcal{V}} = \mathcal{L}^q_{\mathcal{V},h}$ and $\widehat{\mathcal{L}}_{\mathcal{V}} = \widehat{\mathcal{L}}^q_{\mathcal{V},h}$. Taking expectations in Lemma D.1 and applying Lemmas D.2 and D.3 yields

$$\mathbb{E}\{\mathcal{L}^q_{\mathcal{G}^*,h}(\widehat{f}_h)\} \lesssim \left\{\frac{\mathcal{S}_n\mathcal{D}_n\log(e+n\mathcal{B}_{\mathrm{par},n}\mathcal{W}_n\mathcal{D}_n/h)}{n}\right\}^{1/2} + \inf_{f\in\mathcal{F}_n}\|f-f^q_0\|_\infty + h + d_{\to,1}(\mathcal{G}^*,\mathcal{G}_n).$$

This establishes (37). By (46) and $h \asymp \varepsilon_n$, the approximation terms have the required order. Under (38) and (19), $\log(e+n\mathcal{B}_{\mathrm{par},n}\mathcal{W}_n\mathcal{D}_n/h) \lesssim \log n$. Substituting $\mathcal{D}_n \lesssim \log n$ and $\mathcal{S}_n \lesssim n^{d_{\max}/(d_{\max}+2\zeta)}\log n$ therefore yields

$$\left\{\frac{\mathcal{S}_n\mathcal{D}_n\log(e+n\mathcal{B}_{\mathrm{par},n}\mathcal{W}_n\mathcal{D}_n/h)}{n}\right\}^{1/2} \lesssim n^{-\zeta/(d_{\max}+2\zeta)}(\log n)^{3/2}.$$

The approximation terms have a smaller logarithmic order, which proves (39). □

Proof of Theorem C.2. By boundedness of $\mathcal{G}^*$ and Lemma F.5,

$$\sup_{f\in\mathcal{F}_n\cup\mathcal{F}^*}\left|\mathcal{L}^q_{\mathcal{G}^*,\phi}(f)-\mathcal{L}^q_{\mathcal{G}^*,h}(f)\right| \le \sup_{\substack{f\in\mathcal{F}_n\cup\mathcal{F}^*\\ g\in\mathcal{G}^*}}\mathbb{E}[|g(Z)|\,|\mathbf{1}\{Y-f(X)\le 0\}-\mathbf{I}_h\{Y-f(X)\}|] \lesssim h.$$

Because $\mathcal{L}^q_{\mathcal{G}^*,\phi}(f^q_0)=0$, evaluating the preceding uniform bound at $\widehat{f}_h$ yields

$$\mathcal{L}^q_{\mathcal{G}^*,\phi}(\widehat{f}_h) \lesssim \mathcal{L}^q_{\mathcal{G}^*,h}(\widehat{f}_h) + h.$$

Taking expectations, applying Theorem C.1, and using (38) proves (40). □

Proof of Theorem C.3. We apply Lemma D.1 with $\mathcal{L}_{\mathcal{V}} = \mathcal{L}^m_{\mathcal{V}}$ and $\widehat{\mathcal{L}}_{\mathcal{V}} = \widehat{\mathcal{L}}^m_{\mathcal{V}}$. For $\ell^m_{f,g}(x,y,z) = \{y-f(x)\}g(z)$, write $V_n^2 = n^{-1}\sum_{i=1}^n Y_i^2$. For any two pairs of networks,

$$\|\ell^m_{f_1,g_1}-\ell^m_{f_2,g_2}\|_{L^2(P_n)} \le B_0\|f_1-f_2\|_\infty + (V_n+B_0)\|g_1-g_2\|_\infty.$$

Under (36), the loss class has an integrable envelope, so symmetrization applies. Define

$$\mathcal{E}^m_{\mathrm{sta},n} = \sup_{f\in\mathcal{F}_n}\left|\widehat{\mathcal{L}}^m_{\mathcal{G}_n}(f)-\mathcal{L}^m_{\mathcal{G}_n}(f)\right|.$$

Conditionally on the observations, the same product-covering argument as in Lemma D.2, now in the empirical $L^2(P_n)$ metric, gives the entropy bound with the additional scale factor $1+V_n$. Specifically, if $\mathscr{H}_m = \{\ell^m_{f,g} : f\in\mathcal{F}_n, g\in\mathcal{G}_n\}$, then, for $0<u\lesssim 1+V_n$,

$$\log\mathcal{N}\{u,\mathscr{H}_m,L^2(P_n)\} \lesssim \mathcal{S}_n\mathcal{D}_n\log\left[e+\frac{\mathcal{B}_{\mathrm{par},n}\mathcal{W}_n\mathcal{D}_n(1+V_n)}{u\wedge(1+V_n)}\right].$$

Apply the truncated conditional Dudley bound with truncation level $(1+V_n)/n$, and rescale the entropy integral by $u=(1+V_n)t$. The resulting conditional Rademacher bound is at most a constant times

$$(1+V_n)\left\{\frac{\mathcal{S}_n\mathcal{D}_n\log(e+n\mathcal{B}_{\mathrm{par},n}\mathcal{W}_n\mathcal{D}_n)}{n}\right\}^{1/2}.$$

Condition (36) implies $\mathbb{E}(Y^2)\le 2K_Y^2$, and Jensen's inequality gives $\mathbb{E}V_n \le \sqrt{2}K_Y$. Averaging the conditional bound therefore gives

$$\mathbb{E}\mathcal{E}^m_{\mathrm{sta},n} \lesssim \left\{\frac{\mathcal{S}_n\mathcal{D}_n\log(e+n\mathcal{B}_{\mathrm{par},n}\mathcal{W}_n\mathcal{D}_n)}{n}\right\}^{1/2}.$$

Moreover, $|Q^m(f,g)-Q^m(f^m_0,g)| \le B_0\|f-f^m_0\|_\infty$. Because $\mathbb{E}(Y-f^m_0(X)\mid Z)=0$, this controls the comparator risk over every bounded network critic, irrespective of whether that critic belongs to $\mathcal{G}^*$. The critic approximation term is at most $(\sqrt{2}K_Y+B_0)d_{\to,1}(\mathcal{G}^*,\mathcal{G}_n)$. Indeed, (36) and conditional Jensen imply $\mathbb{E}(|Y|\mid Z)\le\sqrt{2}K_Y$ almost surely, so for every pair g_1,g_2, $\mathbb{E}[|Y-f(X)|\,|g_1(Z)-g_2(Z)|]\le(\sqrt{2}K_Y+B_0)\|g_1-g_2\|_{L^1(P_Z)}$. Combining these three bounds with the oracle decomposition establishes (41). Finally, (46), (19), and $\mathcal{L}^m_{\mathcal{G}^*}(f^m_0)=0$ give (42). □

Proof of Theorem C.4. Apply Lemma D.1 with $\mathcal{L}_{\mathcal{V}} = \mathcal{L}^e_{\mathcal{V}}$ and $\widehat{\mathcal{L}}_{\mathcal{V}} = \widehat{\mathcal{L}}^e_{\mathcal{V}}$. For

$$\ell^e_{f,g}(x,y,z) = \psi_\tau\{y - f(x)\}g(z), \qquad V_n^2 = \frac{1}{n}\sum_{i=1}^n Y_i^2,$$

the score bounds in Lemma F.4 imply

$$\|\ell^e_{f_1,g_1} - \ell^e_{f_2,g_2}\|_{L^2(P_n)} \le 2B_0\|f_1 - f_2\|_\infty + 2(V_n + B_0)\|g_1 - g_2\|_\infty. \tag{52}$$

The loss class has an integrable envelope under (36), so symmetrization applies. Define

$$\mathcal{E}^e_{\mathrm{sta},n} = \sup_{f\in\mathcal{F}_n}\left|\widehat{\mathcal{L}}^e_{\mathcal{G}_n}(f) - \mathcal{L}^e_{\mathcal{G}_n}(f)\right|.$$

Combining (52) with the same product-covering and truncated conditional Dudley argument used in the proof of Theorem C.3 gives

$$\mathbb{E}\mathcal{E}^e_{\mathrm{sta},n} \lesssim \left\{\frac{\mathcal{S}_n\mathcal{D}_n\log(e + n\mathcal{B}_{\mathrm{par},n}\mathcal{W}_n\mathcal{D}_n)}{n}\right\}^{1/2}, \tag{53}$$

because $\mathbb{E}V_n \le \sqrt{2}K_Y$.

For the structural approximation term, validity at f_0^e, boundedness of the network critics, and the global Lipschitz property of ψ_τ yield

$$|Q^e(f,g) - Q^e(f_0^e,g)| \le 2B_0\|f - f_0^e\|_\infty.$$

For the critic approximation term, $|\psi_\tau\{Y - f(X)\}| \le 2(|Y| + B_0)$ and conditional Jensen give

$$|Q^e(f,g_1) - Q^e(f,g_2)| \le 2(\sqrt{2}K_Y + B_0)\|g_1 - g_2\|_{L^1(P_Z)}.$$

Substitution of these three bounds into the deterministic oracle decomposition proves (43); fixed constants depending on τ, B_0, and K_Y are absorbed by $\lesssim$. Finally, (46) and (19) give (44) by the same calculation as in the mean case. □

Proof of Theorem 4.2. Equations (40) and (42), together with (44), imply $\mathbb{E}\{\mathcal{R}(\widehat{f})\} = O(r_n)$ in the quantile, mean, and expectile models, respectively. Since $\mathcal{R} \ge 0$, Markov's inequality implies $\mathcal{R}(\widehat{f}) = O_P(r_n)$. Let $A_n = \{\widehat{f} \in \mathcal{U}_\beta\}$. On A_n, put $\delta = \widehat{f} - f_0 = \sum_{j\ge1}\delta_j u_j$. For every integer $J \ge 1$, clause 2 of Assumption C.4 gives $\|\Pi_J T_m\delta\|_{L^2(P_Z)} \lesssim J^b\mathcal{R}(\widehat{f})$ in the mean model, because $M_m(\widehat{f}) - M_m(f_0^m) = -T_m\delta$. In the expectile and quantile models, the triangle inequality and clause 3 give

$$\begin{aligned}(1 - c_T)\|\Pi_J T_s\delta\|_{L^2(P_Z)} &\le \|\Pi_J\{M_s(\widehat{f}) - M_s(f_0^s)\}\|_{L^2(P_Z)}\\ &\lesssim J^b\mathcal{R}(\widehat{f}).\end{aligned}$$

Hence, in each model,

$$\|\Pi_J T\delta\|_{L^2(P_Z)} \lesssim J^b\mathcal{R}(\widehat{f}). \tag{54}$$

In the expectile and quantile models, the implicit constant additionally depends on $(1 - c_T)^{-1}$.

Parseval's identity, monotonicity of (σ_j), (54), and clause 4 of Assumption C.4 now yield, on A_n,

$$\begin{aligned}\|\delta\|^2_{L^2(P_X)} &= \sum_{j\le J}\delta_j^2 + \sum_{j>J}\delta_j^2\\ &\le \sigma_J^{-2}\sum_{j\le J}\sigma_j^2\delta_j^2 + J^{-2\beta}\sum_{j>J}j^{2\beta}\delta_j^2\\ &= \sigma_J^{-2}\|\Pi_J T\delta\|^2_{L^2(P_Z)} + J^{-2\beta}\sum_{j>J}j^{2\beta}\delta_j^2\\ &\lesssim \sigma_J^{-2}J^{2b}\mathcal{R}(\widehat{f})^2 + J^{-2\beta}.\end{aligned}$$

Under mild ill-posedness, take the deterministic sequence, for all sufficiently large n, $J_n = \max\{1, \lfloor r_n^{-1/(a+b+\beta)} \rfloor\}$. Since $\sigma_{J_n}^{-1} \lesssim J_n^a$ and $\mathcal{R}(\widehat{f}) = O_P(r_n)$, the two square-root terms in the preceding display have orders

$$J_n^{a+b}\mathcal{R}(\widehat{f}) = O_P\{r_n^{\beta/(a+b+\beta)}\}, \qquad J_n^{-\beta} \asymp r_n^{\beta/(a+b+\beta)}.$$

Therefore, $\mathbf{1}_{A_n}\|\widehat{f} - f_0\|_{L^2(P_X)} = O_P\{r_n^{\beta/(a+b+\beta)}\}$. Under severe ill-posedness, for all sufficiently large n, take $J_n = \max\{1, \lfloor c\{\log(1/r_n)\}^{1/a} \rfloor\}$, where $0 < c < c_2^{-1/a}$. Then $\kappa = 1 - c_2 c^a > 0$ and

$$J_n^b \exp(c_2 J_n^a) r_n \lesssim \{\log(1/r_n)\}^{b/a} r_n^{\kappa} = o\Big(\{\log(1/r_n)\}^{-\beta/a}\Big).$$

It follows from $\mathcal{R}(\widehat{f}) = O_P(r_n)$ that the first square-root term is $o_P(J_n^{-\beta})$, so the same display gives $\mathbf{1}_{A_n}\|\widehat{f} - f_0\|_{L^2(P_X)} = O_P\{(\log(1/r_n))^{-\beta/a}\} = O_P\{(\log n)^{-\beta/a}\}$, because $\log(1/r_n) \asymp \log n$. Finally, define $s_n = r_n^{\beta/(a+b+\beta)}$ under mild ill-posedness and $s_n = (\log n)^{-\beta/a}$ under severe ill-posedness. For either asserted rate and every $C > 0$,

$$\mathbb{P}\{\|\widehat{f} - f_0\|_{L^2(P_X)} > C s_n\} \le \mathbb{P}(A_n^c) + \mathbb{P}\{A_n,\ \|\widehat{f} - f_0\|_{L^2(P_X)} > C s_n\}.$$

The first term tends to zero and the second is uniformly small for large C, which proves the unconditional O_P statements. □

F Auxiliary Results

This appendix records the external approximation result, its implication for the network sieves, and the covering and smoothing bounds used in the proofs.

Lemma F.1 (Hölder approximation by ReLU DNNs)**.** *Suppose $v \in \mathcal{H}^\zeta([0,1]^d, B)$, where $\zeta = k+\gamma$, $k \in \mathbb{N}_0$, and $\gamma \in (0,1]$. For every $\varepsilon \in (0,1)$, there exists a ReLU network ψ_ε whose width $\mathcal{W}_\varepsilon$, depth $\mathcal{D}_\varepsilon$, size $\mathcal{S}_\varepsilon$, and parameter bound $\mathcal{B}_{\mathrm{par},\varepsilon}$ satisfy*

$$\mathcal{W}_\varepsilon \lesssim \varepsilon^{-d/\zeta}, \qquad \mathcal{D}_\varepsilon \lesssim \log(1/\varepsilon), \qquad \mathcal{S}_\varepsilon \lesssim \varepsilon^{-d/\zeta}\log(1/\varepsilon), \qquad \mathcal{B}_{\mathrm{par},\varepsilon} \lesssim \varepsilon^{-d/\zeta},$$

and $\|v - \psi_\varepsilon\|_{\infty,[0,1]^d} \lesssim \varepsilon$. This result is Theorem 5 of Feng et al. (2024).

Lemma F.2. *Under Assumption C.1, the boundedness clause of Assumption C.2, and the scaling in (18)–(19), (46) holds.*

Proof. Apply Lemma F.1 with $\varepsilon = A\varepsilon_n$, first in dimension d and then in dimension p, for a sufficiently large fixed constant A. This rescaling absorbs the constant in the width bound and accommodates the prescribed width order in (19); the approximation error remains $O(\varepsilon_n)$. Because $d, p \le d_{\max}$, both Hölder balls admit approximating networks with width and parameter bound $O(n^{d_{\max}/(d_{\max}+2\zeta)})$, size $O\{n^{d_{\max}/(d_{\max}+2\zeta)}\log n\}$, and depth $O(\log n)$. If necessary, compose the approximating network with $t \mapsto -B_0 + \mathrm{ReLU}(t + B_0) - \mathrm{ReLU}(t - B_0)$. This map clips to $[-B_0, B_0]$, adds only a fixed number of parameters and layers, and cannot increase the uniform approximation error because the target lies in that interval. The resulting networks belong to $\mathcal{F}_n$ and $\mathcal{G}_n$ for sufficiently large constants in (19). Finally, $\|u\|_{L^1(P_Z)} \le \|u\|_\infty$, which proves (46). □

Lemma F.3 (Covering number of ReLU DNNs)**.** *Let $\mathcal{N}_{\mathrm{ReLU}}(\mathcal{W}, \mathcal{D}, \mathcal{S}, \mathcal{B}_{\mathrm{par}})$ be a class of ReLU networks on $[0,1]^d$ with width at most $\mathcal{W}$, depth at most $\mathcal{D}$, at most $\mathcal{S}$ nonzero parameters, and parameter magnitudes bounded by $\mathcal{B}_{\mathrm{par}} \ge 1$. For $0 < \delta \le 1$,*

$$\log \mathcal{N}(\delta, \mathcal{N}_{\mathrm{ReLU}}(\mathcal{W}, \mathcal{D}, \mathcal{S}, \mathcal{B}_{\mathrm{par}}), \|\cdot\|_\infty) \lesssim \mathcal{S}\mathcal{D}\log\left(\frac{\mathcal{B}_{\mathrm{par}}\mathcal{W}\mathcal{D}}{\delta}\right).$$

This result is Lemma 20 of Feng et al. (2024).

Lemma F.4 (Expectile score bounds and linearization)**.** *For every $\tau \in (0,1)$, the score ψ_τ defined in (2) is continuous, globally Lipschitz with constant $2\max\{\tau, 1-\tau\} \le 2$, and satisfies*

$$|\psi_\tau(u)| \le 2|u|.$$

If the score-e conditional restriction holds, then (45) holds for every candidate f. If, in addition, $\mathbb{P}\{Y = f_0^e(X) \mid X, Z\} = 0$ almost surely, then (49) holds for every $\delta \in L^2(P_X)$.

Proof. On the two half-lines, $\psi_\tau(u)$ has slopes 2τ and $2(1-\tau)$, and both pieces meet at zero. This proves the Lipschitz and envelope bounds. Since ψ_τ is absolutely continuous, the fundamental theorem of calculus gives, with $R_0^e = Y - f_0^e(X)$ and $\Delta_f = f - f_0^e$,

$$\psi_\tau(R_0^e - \Delta_f) - \psi_\tau(R_0^e) = -\Delta_f \int_0^1 2w_\tau(R_0^e - t\Delta_f)\, dt.$$

The value assigned to the derivative at zero is immaterial to the integral. Taking conditional expectations given Z and using validity proves the path identity.

For the derivative statement, the no-atom condition ensures that $R_0^e \neq 0$ almost surely. Hence, for almost every observation,

$$\frac{\psi_\tau(R_0^e - t\delta(X)) - \psi_\tau(R_0^e)}{t} \longrightarrow -2w_\tau(R_0^e)\delta(X).$$

The absolute value of the quotient is bounded by $2|\delta(X)|$. Dominated convergence in L^2, followed by contraction of conditional expectation, proves (49); iterated conditioning identifies its limit with $-T_e\delta$. □

Lemma F.5 (Smoothing and translation bounds)**.** *Suppose Assumption C.3 and the boundedness clause of Assumption C.2 hold and K satisfies (26). For $0 < h \leq 1$, uniformly over $f, f_1, f_2 \in \mathcal{F}_n \cup \mathcal{F}^*$,*

$$\mathbb{E}[|\mathbf{I}_h\{Y - f_1(X)\} - \mathbf{I}_h\{Y - f_2(X)\}| \,|\, X, Z] \lesssim |f_1(X) - f_2(X)|,$$
$$\mathbb{E}\,|\mathbf{1}\{Y - f(X) \leq 0\} - \mathbf{I}_h\{Y - f(X)\}| \lesssim h.$$

Proof. For fixed (x, z), let $H_h(a \mid x, z) = \mathbb{E}\{\mathbf{I}_h(Y - a) \mid X = x, Z = z\}$. Let $F_{x,z}$ denote the conditional distribution function of Y. The integral representation of $\mathbf{I}_h$ and Tonelli's theorem gives $H_h(a \mid x, z) = \int F_{x,z}(a - hu)K(u)\, du$. For $a_2 > a_1$, Assumption C.3 therefore implies

$$0 \leq H_h(a_2 \mid x, z) - H_h(a_1 \mid x, z) = \int \{F_{x,z}(a_2 - hu) - F_{x,z}(a_1 - hu)\}K(u)\, du$$
$$\leq M_0(a_2 - a_1).$$

Because $a \mapsto \mathbf{I}_h(Y - a)$ is monotone, the conditional expectation of its absolute increment equals the absolute increment of H_h. Hence

$$|H_h(f_1(x) \mid x, z) - H_h(f_2(x) \mid x, z)| \leq M_0|f_1(x) - f_2(x)|.$$

Integrating over (X, Z) proves the first inequality. For the second, the integral representation of $\mathbf{I}_h$, Jensen's inequality, and Tonelli's theorem give, conditionally on (X, Z),

$$\mathbb{E}[|\mathbf{1}\{Y - f(X) \leq 0\} - \mathbf{I}_h\{Y - f(X)\}| \,|\, X, Z]$$
$$\leq \int_{\mathbb{R}} \mathbb{P}(\mathbf{1}\{Y - f(X) \leq 0\} \neq \mathbf{1}\{Y - f(X) + hu \leq 0\} \,|\, X, Z)\, K(u)\, du$$
$$\leq M_0 h \int_{\mathbb{R}} |u|K(u)\, du = M_0\mu_1(K)h.$$

Taking expectations yields the second inequality. □

G Experimental Details

This appendix specifies preprocessing, network architectures, and the penalized updates used in the numerical experiments.

G.1 Data-generating designs

All three designs use the nonlinear base function

$$f_0(x) = \exp(0.1\boldsymbol{\alpha}_{\text{sim}}^\top x) + 0.2(1 + \boldsymbol{\beta}_{\text{sim}}^\top x),$$

where $\boldsymbol{\alpha}_{\mathrm{sim}}$ and $\boldsymbol{\beta}_{\mathrm{sim}}$ are fixed index vectors. For reproducibility, let $a_{\mathrm{sim},j} = 0.5 + 0.1(j-1)$ and $b_{\mathrm{sim},j} = 1.4 - 0.1(j-1)$, $j = 1, \ldots, 10$, and set $\boldsymbol{\alpha}_{\mathrm{sim}} = \mathbf{a}_{\mathrm{sim}}/\|\mathbf{a}_{\mathrm{sim}}\|_2$ and $\boldsymbol{\beta}_{\mathrm{sim}} = \mathbf{b}_{\mathrm{sim}}/\|\mathbf{b}_{\mathrm{sim}}\|_2$. Thus, before normalization,

$$\mathbf{a}_{\mathrm{sim}} = (0.5, 0.6, \ldots, 1.4)^\top, \qquad \mathbf{b}_{\mathrm{sim}} = (1.4, 1.3, \ldots, 0.5)^\top.$$

These vectors are held fixed across sample sizes and replications. On an independent test sample of size $N_{\mathrm{test}} = 10{,}000$, we evaluate discrepancies from the generating function by mean squared error (MSE),

$$\mathrm{MSE} = \frac{1}{N_{\mathrm{test}}} \sum_{i=1}^{N_{\mathrm{test}}} \{\widehat{f}(X_i) - f_0(X_i)\}^2.$$

We also report mean absolute error (MAE),

$$\mathrm{MAE} = \frac{1}{N_{\mathrm{test}}} \sum_{i=1}^{N_{\mathrm{test}}} \left|\widehat{f}(X_i) - f_0(X_i)\right|.$$

Here the comparison function is the score-specific structural truth. It is the displayed f_0 in the mean and quantile designs, and it is $f^e_{0,\tau} = f_0 + c^e_\tau$ in the expectile design defined below. For a metric M, let $R = 100$ denote the number of independent Monte Carlo replications, and let M_r denote its value in replication r, $r = 1, \ldots, R$. The reported Monte Carlo standard error (MCSE) is

$$\mathrm{MCSE}(\overline{\mathsf{M}}) = \left\{ \frac{1}{R(R-1)} \sum_{r=1}^{R} (\mathsf{M}_r - \overline{\mathsf{M}})^2 \right\}^{1/2}, \qquad \overline{\mathsf{M}} = R^{-1} \sum_{r=1}^{R} \mathsf{M}_r.$$

In the mean-regression design, $d = p = 10$, $Z \sim N(0, I_d)$, and Z is independent of $(\epsilon_m, U^\top)^\top$. Let $\mathbf{1}_d$ denote the d-vector of ones and let I_d denote the $d \times d$ identity matrix. The vector $(\epsilon_m, U^\top)^\top \in \mathbb{R}^{d+1}$ is jointly Gaussian with mean zero and covariance

$$\Sigma_m = \begin{pmatrix} 1 & \rho_m \mathbf{1}_d^\top \\ \rho_m \mathbf{1}_d & I_d \end{pmatrix}, \qquad \rho_m = 0.1.$$

Set $X = Z + U$ and $Y = f_0(X) + \epsilon_m$. Thus X is endogenous while $\mathbb{E}(\epsilon_m \mid Z) = 0$.

In the expectile-regression design, $d = p = 10$ and all primitive draws are bounded in order to retain the sub-Gaussian condition used by the expectile theory. Let $V \sim N(0,1)$ conditional on $|V| \le 3$, put

$$v_3 = 1 - \frac{6\phi(3)}{2\Phi(3) - 1}, \qquad T = V/\sqrt{v_3},$$

and draw the coordinates of $Z \in \mathbb{R}^p$, the coordinates of $U \in \mathbb{R}^d$, and an auxiliary scalar $\widetilde{U}$ independently from the law of T. Let $\boldsymbol{\rho}_e = (0.5/\sqrt{d})\mathbf{1}_d$, and let $A \in \mathbb{R}^{d \times p}$ be the orthogonal factor in the QR decomposition of a 10×10 standard-normal matrix. Define

$$X = AZ + U, \qquad \epsilon_e^{\mathrm{raw}} = \boldsymbol{\rho}_e^\top U + \sqrt{1 - \|\boldsymbol{\rho}_e\|_2^2}\, \widetilde{U}, \qquad Y = f_0(X) + \epsilon_e^{\mathrm{raw}}.$$

For each τ, the shift c^e_τ is the unique solution of

$$\mathbb{E}[\psi_\tau(\epsilon_e^{\mathrm{raw}} - c^e_\tau)] = 0,$$

which we approximate using a fixed antithetic Monte Carlo sample of size $500{,}000$. The numerical shifts used in the experiment are $c^e_{0.25} = -0.4375$, $c^e_{0.50} = 0$, and $c^e_{0.75} = 0.4375$. With the population shifts, the generating expectile function is $f^e_{0,\tau}(x) = f_0(x) + c^e_\tau$ and

$$\mathbb{E}\big[\psi_\tau\{Y - f^e_{0,\tau}(X)\} \mid Z\big] = 0.$$

Because all primitive variables are bounded, X, $\epsilon_e^{\mathrm{raw}}$, and Y are bounded as well; in particular, the conditional sub-Gaussian requirement in Assumption C.2 holds.

In the quantile-regression design, $d = p = 10$, $Z \sim N(0, I_p)$, the coordinates of $U \in \mathbb{R}^d$ are independent t_3 variables, and $\widetilde{U} \sim t_3$; moreover, Z, U, and $\widetilde{U}$ are mutually independent. Set $\boldsymbol{\rho}_q = (0.5/\sqrt{d})\mathbf{1}_d$, so that $\|\boldsymbol{\rho}_q\|_2 = 0.5$, and define

$$\epsilon_\tau^{\mathrm{raw}} = \boldsymbol{\rho}_q^\top U + \sqrt{1 - \|\boldsymbol{\rho}_q\|_2^2}\, \widetilde{U}, \qquad c_\tau = Q_\tau(\epsilon_\tau^{\mathrm{raw}}), \qquad \epsilon_\tau = \epsilon_\tau^{\mathrm{raw}} - c_\tau,$$

where Q_τ denotes the marginal τ-quantile; computationally, c_τ is evaluated from a fixed Monte Carlo sample of size 500,000. Let $A \in \mathbb{R}^{d\times p}$ be the orthogonal factor in the QR decomposition of a 10×10 standard-normal random matrix, held fixed across sample sizes, quantile levels, and replications. Finally, set $X = AZ + U$ and $Y = f_0(X) + \epsilon_\tau$. Then $\mathbb{P}(\epsilon_\tau \le 0 \mid Z) = \tau$. The disturbance laws and support conditions differ across designs, so we compare estimators and sample sizes within each design. The bounded expectile outcome satisfies the sub-Gaussian tail requirement. The Gaussian mean and heavy-tailed quantile designs have unbounded covariate support and therefore differ from the compact-support setting of the theory. For the quantile estimator, we smooth the discontinuous indicator by replacing $\mathbf{1}\{v \le 0\}$, where $v = Y - f(X)$, with $\mathbf{I}_h(v) = G(-v/h) = \{1 + \exp(v/h)\}^{-1}$, where G is the standard logistic distribution function and $K = G'$ is the logistic density kernel in (27). The implemented score is therefore $\phi_{\tau,h}(v) = \tau - \{1+\exp(v/h)\}^{-1}$. Let $Y_{1,r}, \ldots, Y_{n,r}$ denote the training responses in replication r, and define the replication-specific robust response scale $S_r^{(Y)} = \mathrm{IQR}(Y_{1,r}, \ldots, Y_{n,r})/1.349$. The bandwidth is fixed at $h = 0.30$ on the original response scale, so the standardized residual uses $\widetilde{h}_r = 0.30/S_r^{(Y)}$. This construction is integrated-kernel smoothing of the quantile indicator. The logistic density is bounded, symmetric, and has a finite first absolute moment, so it satisfies the kernel conditions.

G.2 Monte Carlo protocol

All three experiments use the common sample-size grid

$$\mathcal{N}_{\mathrm{sim}} = \{2{,}000, 4{,}000, 6{,}000, 8{,}000, 10{,}000, 15{,}000, 20{,}000, 30{,}000, 50{,}000\}.$$

For each design and, where applicable, each expectile or quantile level, we run 100 independent Monte Carlo replications. Each replication draws a training sample of size 50,000; the samples at smaller n are nested prefixes. This common-random-number construction preserves the marginal experiment at each n while reducing extraneous variation in comparisons across sample sizes. The full design therefore contains 900 mean-IV fits, 2,700 expectile-IV fits, and 2,700 quantile-IV fits.

Rather than split each training sample, we use an independent test sample of size 10,000. Its covariates and score-specific structural truth are held fixed across all sample sizes and Monte Carlo replications; the same test design is also used across the three quantile levels and, separately, across the three expectile levels. Each fit standardizes the test covariates using moments computed from its training sample and is then transformed back to the original response scale before comparison with $f_0(\cdot)$. The test set is used only to compute MSE and MAE after training. The tables report replication means with MCSEs in parentheses, and the figures display means with pointwise 1.96-MCSE error bars. The MCSEs describe variation over training replications conditional on the fixed test design; they do not include uncertainty from drawing a new test set.

Mean-IV summary figure.

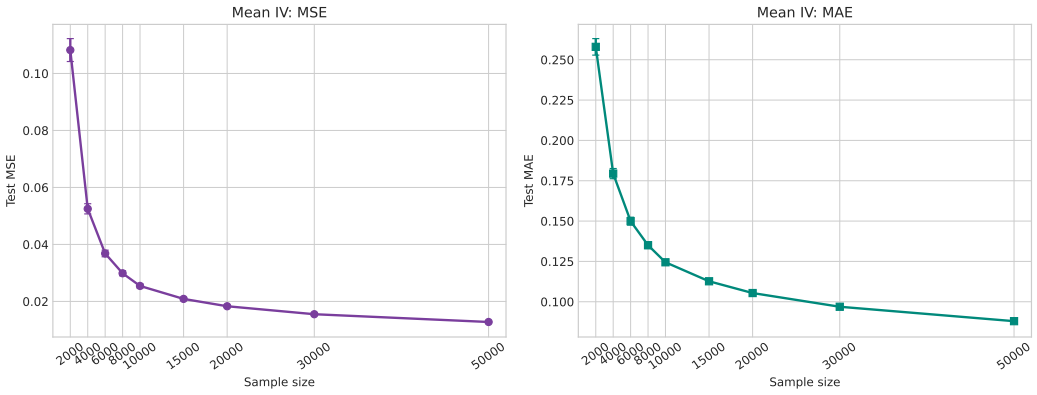

Figure 2: Mean-IV test MSE (left) and MAE (right) across sample sizes, with pointwise 1.96-MCSE error bars over 100 replications.

Expectile-IV summary figure.

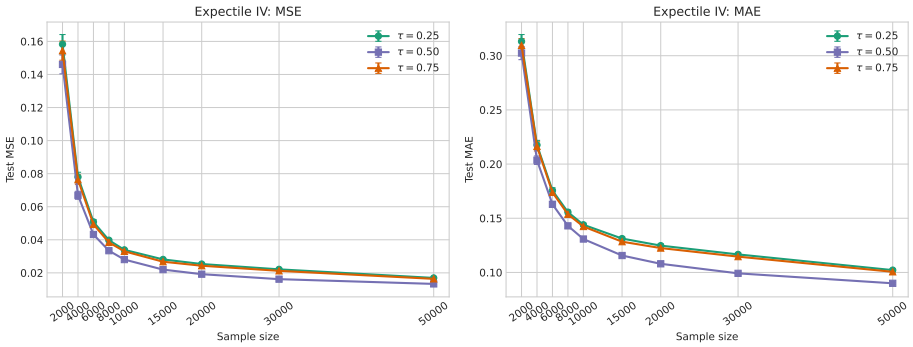

Figure 3: Expectile-IV test MSE (left) and MAE (right) across sample sizes for $\tau = 0.25, 0.50, 0.75$, with pointwise 1.96-MCSE error bars over 100 replications.

Quantile-IV summary figure.

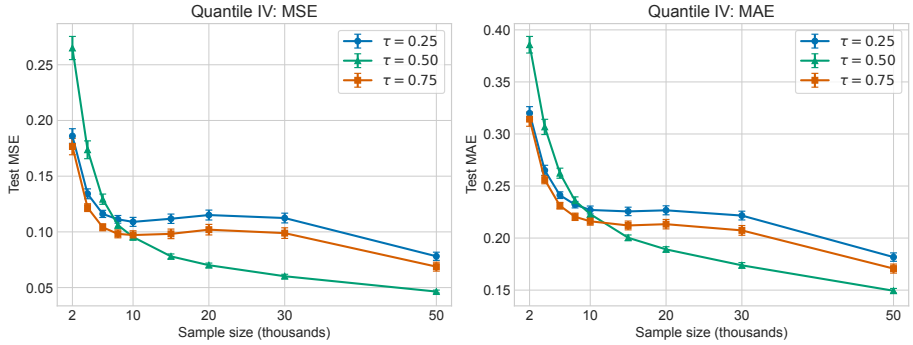

Figure 4: Quantile-IV test MSE (left) and MAE (right) for $d = p = 10$ across sample sizes and $\tau = 0.25, 0.50, 0.75$, with pointwise 1.96-MCSE error bars over 100 replications.

G.3 Detailed simulation result tables

Table 2: Mean-IV test MSE and MAE across the sample-size grid, with MCSE in parentheses.

n	MSE (MCSE)	MAE (MCSE)
2,000	0.1082 (0.0021)	0.2580 (0.0026)
4,000	0.0525 (0.0009)	0.1794 (0.0016)
6,000	0.0368 (0.0006)	0.1499 (0.0013)
8,000	0.0299 (0.0004)	0.1350 (0.0010)
10,000	0.0254 (0.0003)	0.1245 (0.0009)
15,000	0.0209 (0.0003)	0.1127 (0.0007)
20,000	0.0183 (0.0002)	0.1054 (0.0006)
30,000	0.0155 (0.0002)	0.0969 (0.0006)
50,000	0.0128 (0.0001)	0.0880 (0.0005)

Table 3: Expectile-IV test MSE across sample sizes and expectile levels, with MCSE in parentheses.

n	$\tau = 0.25$	$\tau = 0.50$	$\tau = 0.75$
2,000	0.1583 (0.0030)	0.1461 (0.0029)	0.1544 (0.0029)
4,000	0.0780 (0.0015)	0.0669 (0.0012)	0.0764 (0.0014)
6,000	0.0508 (0.0007)	0.0432 (0.0007)	0.0493 (0.0008)
8,000	0.0398 (0.0005)	0.0334 (0.0005)	0.0384 (0.0005)
10,000	0.0339 (0.0003)	0.0280 (0.0004)	0.0330 (0.0004)
15,000	0.0281 (0.0003)	0.0220 (0.0003)	0.0267 (0.0003)
20,000	0.0253 (0.0003)	0.0192 (0.0002)	0.0243 (0.0003)
30,000	0.0221 (0.0003)	0.0162 (0.0002)	0.0212 (0.0003)
50,000	0.0169 (0.0002)	0.0133 (0.0001)	0.0164 (0.0002)

Table 4: Expectile-IV test MAE across sample sizes and expectile levels, with MCSE in parentheses.

n	$\tau = 0.25$	$\tau = 0.50$	$\tau = 0.75$
2,000	0.3132 (0.0032)	0.3025 (0.0031)	0.3095 (0.0032)
4,000	0.2176 (0.0022)	0.2036 (0.0020)	0.2164 (0.0022)
6,000	0.1756 (0.0013)	0.1629 (0.0013)	0.1739 (0.0014)
8,000	0.1557 (0.0009)	0.1430 (0.0010)	0.1536 (0.0010)
10,000	0.1440 (0.0007)	0.1308 (0.0009)	0.1424 (0.0009)
15,000	0.1313 (0.0007)	0.1156 (0.0007)	0.1284 (0.0007)
20,000	0.1248 (0.0008)	0.1080 (0.0006)	0.1225 (0.0007)
30,000	0.1166 (0.0008)	0.0992 (0.0006)	0.1146 (0.0007)
50,000	0.1021 (0.0006)	0.0900 (0.0005)	0.1006 (0.0006)

Table 5: Quantile-IV test MSE for $d = p = 10$ across sample sizes and quantile levels, with MCSE over 100 replications in parentheses.

n	$\tau = 0.25$	$\tau = 0.50$	$\tau = 0.75$
2,000	0.1861 (0.0033)	0.2650 (0.0053)	0.1766 (0.0038)
4,000	0.1342 (0.0023)	0.1737 (0.0041)	0.1218 (0.0018)
6,000	0.1161 (0.0016)	0.1293 (0.0024)	0.1041 (0.0016)
8,000	0.1113 (0.0018)	0.1061 (0.0017)	0.0982 (0.0018)
10,000	0.1090 (0.0021)	0.0954 (0.0014)	0.0972 (0.0020)
15,000	0.1118 (0.0021)	0.0780 (0.0011)	0.0982 (0.0022)
20,000	0.1151 (0.0023)	0.0701 (0.0010)	0.1020 (0.0024)
30,000	0.1124 (0.0022)	0.0602 (0.0009)	0.0989 (0.0024)
50,000	0.0781 (0.0019)	0.0464 (0.0007)	0.0688 (0.0020)

Table 6: Quantile-IV test MAE for $d = p = 10$ across sample sizes and quantile levels, with MCSE over 100 replications in parentheses.

n	$\tau = 0.25$	$\tau = 0.50$	$\tau = 0.75$
2,000	0.3199 (0.0032)	0.3857 (0.0041)	0.3142 (0.0035)
4,000	0.2651 (0.0025)	0.3068 (0.0037)	0.2561 (0.0020)
6,000	0.2412 (0.0016)	0.2623 (0.0025)	0.2311 (0.0016)
8,000	0.2321 (0.0016)	0.2359 (0.0019)	0.2204 (0.0018)
10,000	0.2270 (0.0019)	0.2230 (0.0017)	0.2161 (0.0019)
15,000	0.2256 (0.0021)	0.2003 (0.0014)	0.2121 (0.0022)
20,000	0.2267 (0.0022)	0.1891 (0.0013)	0.2133 (0.0023)
30,000	0.2216 (0.0021)	0.1738 (0.0013)	0.2074 (0.0024)
50,000	0.1817 (0.0021)	0.1494 (0.0011)	0.1707 (0.0022)

G.4 A RANDOMIZED-ENCOURAGEMENT ILLUSTRATION

We analyze the daily-diary data of Caldaroni et al. (2026), collected through a smartphone-based within-person encouragement design. Sixty-three Italian young adults were followed for 21 days. On 11 randomly selected days per participant, the participant received a morning message encouraging more pro-environmental behavior (PEB), such as recycling or reducing plastic, water, and energy consumption. Each evening, participants reported four PEB items and two eudaimonic well-being (EWB) items—meaning in life and connectedness to others—on five-point scales.

We use randomized daily encouragement Z_{it} as an instrument for realized PEB X_{it}, and EWB as the outcome Y_{it}. The original file contains 1,323 person-days. There are 1,086 days with jointly observed PEB and EWB; after constructing one-day lags, our dynamic specification uses 1,023 person-days from all 63 participants. We center PEB, EWB, encouragement, and the lagged variables within participants. The structural regressor is the deviation of daily PEB from the participant-specific mean. The critic uses randomized encouragement, lagged PEB, lagged EWB, and the day index. The working moments therefore use within-person, day-to-day variation rather than between-person differences. Let i index participants and t index study days, and define the four-dimensional critic input $W_{it} = (Z_{it}, X_{i,t-1}, Y_{i,t-1}, t)^\top \in \mathbb{R}^4$, with these variables understood to retain their within-participant centered values. We build the empirical criterion from the working restriction

$$\mathbb{E}\big[\vartheta_f^s(X_{it}, Y_{it})g(W_{it})\big] = 0,$$

using the scores defined in Section 3.1, with smoothing only for quantiles. We estimate one mean curve and expectile and quantile curves at $\tau \in \{0.25, 0.50, 0.75\}$. A structural interpretation requires the score-specific moment restriction to hold after centering and with the additional lagged variables and day index entering the critic. The reported curves average five initializations. Pointwise percentile ranges from 20 participant-level bootstrap refits provide an exploratory assessment of sensitivity to participant resampling. They are not used as confidence bands.

Detailed fitted contrasts. The noncentral expectile curves differ across levels. At a one-point lower PEB deviation, the estimated contrasts are -0.57, -0.40, and -0.13 for $\tau = 0.25, 0.50, 0.75$, respectively; at a one-point higher deviation they are approximately 0.00, 0.13, and 0.19. At a two-point positive deviation, the separation becomes larger, with corresponding estimates -0.12, 0.30, and 0.52. The quantile curves are closer together near a one-point positive deviation, where the estimated contrasts are approximately 0.05, 0.04, and 0.08 at the same three levels, whereas at large negative deviations the lower and median curves fall more than the upper curve. These comparisons concern distinct IV moment targets, so curves at the same τ need not agree across the two score families.

Real-data limitations. The bootstrap ranges widen near the boundaries, indicating greater sensitivity to participant resampling there. The boundary patterns should therefore be interpreted cautiously. The exclusion restriction requires that the morning message affect evening EWB only through realized PEB; a direct motivational effect would violate this condition. The modest first stage and 63 participants limit precision, and the i.i.d. theory does not cover the clustered diary structure. The two-item discrete EWB outcome also does not literally satisfy the continuous-density assumption used in the quantile theory. Expectile estimation avoids that density and smoothing requirement, but it does not remove the identification, clustering, or weak-first-stage qualifications.

G.5 DATASETS AND DATA PREPROCESSING

Simulation Studies. We use the nested training samples and fixed independent test designs specified in Section G.2. All centering and scaling parameters are computed from the training sample and then applied to the test covariates; predictions are returned to the original response scale for evaluation.

- **Mean-IV Design:** We set $d = p = 10$ and standardize X and Z with the training-sample means and standard deviations. We center and scale the continuous response Y in the same way.
- **Expectile-IV Design:** We set $d = p = 10$ and generate Z, U, and the auxiliary disturbance from independent standard-normal draws truncated to $[-3, 3]$ and then rescaled to unit

variance. The endogeneity vector has Euclidean norm 0.5. We standardize X, Z, and Y using training-sample means and standard deviations. A fixed antithetic Monte Carlo sample of size 500,000 determines the expectile shifts for $\tau \in \{0.25, 0.50, 0.75\}$.

- **Quantile-IV Design:** We set $d = p = 10$ and draw the structural disturbance U and auxiliary error from Student $t(3)$ distributions. We standardize the response by its sample median and interquartile range (IQR), $\widetilde{Y} = (Y - \mathrm{median}(Y))/(\mathrm{IQR}(Y)/1.349)$. Before training, 500,000 Monte Carlo draws determine the shifts for $\tau \in \{0.25, 0.50, 0.75\}$.

Real-data pipeline. The empirical application uses daily-diary data from the within-person randomized-encouragement study. We drop values coded as -999, order observations by participant and day, construct one-day lags ($\mathrm{PEB}_{t-1}, \mathrm{EWB}_{t-1}$), discard incomplete person-days, and center all variables within participants. The resulting panel contains 63 participants and 1,023 complete person-days. Centered EWB is the outcome, centered PEB is the endogenous regressor, and the critic inputs are centered encouragement, lagged PEB, lagged EWB, and day.

Table 7: Dimensions and sample sizes of the experimental datasets.

Setting	Dimensions	Sample Size (n)
Mean-IV	$d = 10, p = 10$	$n \in \{2{,}000, \ldots, 50{,}000\}$
Expectile-IV	$d = 10, p = 10$	$n \in \{2{,}000, \ldots, 50{,}000\}$
Quantile-IV	$d = 10, p = 10$	$n \in \{2{,}000, \ldots, 50{,}000\}$
Diary study	$d = 1, p = 4$	1023

G.6 Network Architecture

In all three simulation designs, f_{net} and g_{net} are multilayer perceptrons with two 64-unit ReLU hidden layers. The critic has a linear scalar output. For the structural network, let $f_{\mathrm{raw}}(x)$ denote the final affine output; we apply $f_{\mathrm{net}}(x) = 4\tanh\{f_{\mathrm{raw}}(x)/4\}$, giving $|f_{\mathrm{net}}(x)| < 4$ on the standardized training scale.

For the real-data application, each network has two 32-unit ReLU hidden layers. The structural network takes the scalar endogenous regressor as input, and the critic takes the four-dimensional instrument/covariate vector. The same architecture is used for the mean, expectile, and quantile fits. We use the default PyTorch initialization and do not constrain the network outputs.

Table 8: Detailed specifications of structural (f_{net}) and critic (g_{net}) neural networks.

Setting	Network	Input Dim.	Hidden Layers	Hidden activation
Mean-IV	f_{net}	10	$64 \to 64$	ReLU
	g_{net}	10	$64 \to 64$	ReLU
Expectile-IV	f_{net}	10	$64 \to 64$	ReLU
	g_{net}	10	$64 \to 64$	ReLU
Quantile-IV	f_{net}	10	$64 \to 64$	ReLU
	g_{net}	10	$64 \to 64$	ReLU
Diary study	f_{net}	1	$32 \to 32$	ReLU
	g_{net}	4	$32 \to 32$	ReLU

G.7 Hyperparameter Settings and Optimization

The simulation and empirical fits use the same form of output penalties. Write $(\widetilde{X}_i, \widetilde{Y}_i, \widetilde{Z}_i)$ for the training variables after preprocessing, with $\widetilde{Z}_i$ denoting the critic inputs in either setting. In this paragraph, f and g denote network outputs in these training coordinates, and $\widetilde{\vartheta}^s_f$ denotes the

implemented score, including the logistic approximation for quantiles. For a batch $\mathcal{I}$, define

$$\mathcal{J}_{\mathcal{I}}^{s}(f,g) = \frac{1}{|\mathcal{I}|}\sum_{i\in\mathcal{I}} \widetilde{\vartheta}_f^{s}(\widetilde{X}_i,\widetilde{Y}_i)g(\widetilde{Z}_i) - \frac{\lambda_g}{2|\mathcal{I}|}\sum_{i\in\mathcal{I}} g(\widetilde{Z}_i)^2 + \frac{\lambda_f}{2|\mathcal{I}|}\sum_{i\in\mathcal{I}} f(\widetilde{X}_i)^2.$$

Critic updates ascend this objective with f fixed, and structural updates descend it with g fixed. We use Adam (Kingma & Ba, 2015) with learning rates $\eta_f = \eta_g = 3 \times 10^{-4}$, three critic updates per structural update, and gradient clipping with a norm threshold of 5. The simulations use mini-batches and parameter weight decay 10^{-5}, in addition to the displayed output penalties. The empirical fits use full batches and no parameter weight decay. Finite alternating updates are not asserted to solve the penalized saddle problem exactly.

In the simulations, the critic penalty is $\lambda_g = 5$, and the structural penalty coefficient is $\lambda_f(n) = 0.01\sqrt{2000/n}$, which decreases to zero with n. The mini-batch size is $\min(1024, n)$, and the training budget is

$$T(n) = \max\left(800,\ 30\left\lceil \frac{n}{\min(1024,n)} \right\rceil\right). \tag{55}$$

For quantile IV, the original-response bandwidth is fixed at $h = 0.30$. If $S_r^{(Y)}$ is the robust response scale defined in Section G.1, the standardized training residual is smoothed with $\widetilde{h}_r = 0.30/S_r^{(Y)}$. Thus the implemented logistic term equals $G(-(Y - f(X))/0.30)$ when expressed on the original response scale.

For the empirical application, all three score families use 500 full-batch steps, with $\lambda_g = 5$ and $\lambda_f = 0.01$. The logistic bandwidth $h = 0.20$ is on the original EWB scale and applies only to the quantile fits; the corresponding standardized bandwidth is divided by the training response standard deviation. Reported curves average five restarts after subtracting each fitted curve's value at zero PEB deviation. Twenty participant-level bootstrap refits yield pointwise 2.5th–97.5th percentile ranges as exploratory summaries of sensitivity to participant resampling, without a confidence-interval interpretation.

Table 9: Hyperparameter settings and training controls across settings.

Hyperparameter	Mean IV	Expectile IV	Quantile IV	Diary study
Optimizer (η_f, η_g)	Adam (3×10^{-4})	Adam (3×10^{-4})	Adam (3×10^{-4})	Adam (3×10^{-4})
Weight Decay	10^{-5}	10^{-5}	10^{-5}	–
Critic Update Ratio ($g : f$)	3 : 1	3 : 1	3 : 1	3 : 1
Batch Size	$\min(1024, n)$	$\min(1024, n)$	$\min(1024, n)$	Full batch (1023)
Total Training Steps (T)	$\max(800, 30\lceil n/\text{batch}\rceil)$	$\max(800, 30\lceil n/\text{batch}\rceil)$	$\max(800, 30\lceil n/\text{batch}\rceil)$	500 steps
Gradient Clipping Norm	5.0	5.0	5.0	5.0
Critic Penalty (λ_g)	5.0	5.0	5.0	5.0
Structural Penalty (λ_f)	$0.01\sqrt{2000/n}$	$0.01\sqrt{2000/n}$	$0.01\sqrt{2000/n}$	0.01
Smoothing bandwidth h (original response scale)	–	–	0.30	0.20 (quantile only)