CONFORMAL INFERENCE FOR MISSING DATA UNDER MULTIPLE ROBUST LEARNING

Wenlu Tang[1], Hongni Wang[2], Xingcai Zhou[3], Bei Jiang[1] and Linglong Kong[1]

[1]University of Alberta, [2]Shandong University of Finance and Economics, and [3]Nanjing Audit University

Supplementary Material

This appendix presents the technical proofs of the theoretical results in the paper.

S1 Main Content Extensions

S1.1 Algorithm for CM-MRL

We provide pseudocode for the proposed CM-MRL in Algorithm 1.

Algorithm 1 CM-MRL: Conformal Prediction with Multiple Robust Learning

Require: Data $\{(X_i, R_i, Y_i)\}_{i=1}^n$; nominal coverage rate $1-\alpha \in (0,1)$; candidate models $\Pi = \{\hat{\pi}^j\}_{j=1}^J$, $\mathcal{F} = \{\hat{f}^k\}_{k=1}^K$; imputations T.

Ensure: Prediction interval $\hat{C}_{\mathrm{MR}}(x)$ with nominal coverage $1-\alpha$.

1: Randomly split the data set into $\mathcal{I}_{\mathrm{tra}}$ and $\mathcal{I}_{\mathrm{cal}}$ with sizes n_{tra} and n_{cal}.

2: On $\mathcal{I}_{\mathrm{tra}}$, fit $\{\hat{\pi}^j\}_{j=1}^J$ and $\{\hat{f}^k\}_{k=1}^K$, compute EL weights $\{\hat{w}_i\}_{i\in\mathcal{I}_{\mathrm{tra}},\, R_i=1}$ that balance the model moments, and obtain the MR predictor $\hat{\mu}_{\mathrm{MR}}(x)$.

3: For each $i \in \mathcal{I}_{\mathrm{cal}}$ with $R_i = 1$, compute the conformity score $\hat{\varepsilon}_i$.

4: For each $i \in \mathcal{I}_{\mathrm{cal}}$ with $R_i = 0$ and for each $k = 1, \ldots, K$, obtain:

1. Draw $Y_i^t(\hat{b}^k) \sim \hat{f}^k(\cdot \mid X_i)$ for $t = 1, \ldots, T$.
2. Imputed score $\hat{\varepsilon}_i^{(k)} = \dfrac{1}{T}\sum_{t=1}^{T} \left| Y_i^t(\hat{b}^k) - \hat{\mu}_{\mathrm{MR}}(X_i) \right|$ and $\hat{q}_{1-\alpha}^{(k)}$.

5: For each $k = 1, \ldots, K$, obtain $\hat{q}_{1-\alpha}^{(k)}$ by (2.11)

6: Compute $\hat{\theta}^j$, $\hat{\xi}^k$, and thus EL weights $\{\hat{d}_i\}_{i\in\mathcal{I}_{\mathrm{cal}},\, R_i=1}$ by (2.14).

7: Solve for $\hat{q}_{\mathrm{MR}}$ as the root of the weighted score equation (2.15).

8: **Output prediction interval.** For any x, return

$$\hat{C}_{\mathrm{MR}}(x) = \Big[\hat{\mu}_{\mathrm{MR}}(x) - \hat{q}_{\mathrm{MR}}, \;\; \hat{\mu}_{\mathrm{MR}}(x) + \hat{q}_{\mathrm{MR}}\Big].$$

S1.2 Additional experiment settings and results

Settings in Experiment I

Missingness follows the MAR propensity $\mathrm{logit}\{\pi_0(X)\} = 3.5 - 5.0\, X_2$, yielding approximately 58% observed outcomes on average and about 42%

missing. We observe $R = \mathbb{1}\{Y \text{ observed}\}$ and (X, R, YR). We vary the error law through three scenarios:

(A) homoskedastic Gaussian noise, $\varepsilon_Y \sim \mathcal{N}(0, 1)$ with $\sigma = 1$;

(B) heavy tails, $\varepsilon_Y \sim t_3$ scaled to unit variance with $\sigma = 1$;

(C) heteroskedasticity, $\varepsilon_Y \sim \mathcal{N}\big(0,\ (0.6 + 0.2|X_1|)^2\big)$ with $\sigma = 1$ (variance rises with $|X_1|$).

This variation of (σ, ε_Y) by scenario is shown in Table 1.

To study model misspecification and robustness, we provide two candidate propensity models and two candidate outcome models. Both propensity models are correctly specified: the first is the true logistic model, and the second is a larger logistic model that includes the true model as a special case.

$$
\begin{aligned}
\text{Propensity:}\ & \operatorname{logit}\{\pi^1(\alpha^1)\} = \alpha_1^1 + \alpha_2^1 X_2 \quad \text{(correct)},\\
& \operatorname{logit}\{\pi^2(\alpha^2)\} = \alpha_1^2 + \alpha_2^2 X_1 + \alpha_3^2 X_2 + \alpha_4^2 X_3 + \alpha_5^2 X_4 \quad \text{(correct, larger)};\\
\text{Outcome:}\ & a^1(\gamma^1) = \gamma_1^1 + \gamma_2^1 X_1 + \gamma_3^1 X_2 + \gamma_4^1 X_3 + \gamma_5^1 X_4 \quad \text{(correct)},\\
& a^2(\gamma^2) = \gamma_1^2 + \gamma_2^2 X_1 + \gamma_3^2 X_2 + \gamma_4^2 X_3 \quad \text{(misspecified)}.
\end{aligned}
$$

Thus, in Experiment I, the model misspecification considered here arises from the candidate outcome model in some settings, not from the two candidate propensity score models. We examine the six settings in Table 2, which vary the candidates available to the learner; at least one candidate model is correct in each setting.

We study various methods under these settings over 100 Monte Carlo replications with simulated sample size $n = 3000$. For each replicate, we split one third of the sample into the training set and one third into the calibration set, fit candidate models on the training set, and construct prediction intervals on the calibration set at target marginal coverage $1-\alpha = 0.9$. The performance is then validated on testing data of size $n_{tes} = 1000$. We compare our **CM–MRL** (double-calibrated, multiple-robust conformal) against:

- **Impute–SC**: impute then split conformal using the available $a^{(\cdot)}$;
- **SC**: complete-case split conformal on observed pairs only;
- **WCCQR**: weighted conformalized quantile regression.
- **WCCQR-CV**: weighted conformalized quantile regression with cross-validation.

Settings in Experiment II

The covariates include a mixture of continuous and binary components:

$$X_0, X_1, X_3 \sim N(0,1), \quad X_2 \sim \text{Bernoulli}(0.5), \quad X_4 \sim \text{Unif}(-1,1), \quad X_5 \sim \text{Bernoulli}(0.3).$$

The conditional mean is nonlinear,

$$f(X) = 3 + 1.5X_0 + 2\sin(X_1) + X_2 + 0.8X_3X_4 + 0.5\,\mathbf{1}(X_5 = 1),$$

and the outcome is generated by

$$Y = f(X) + \sigma(X)\varepsilon, \qquad \sigma(X) = 0.8 + 0.6\,\mathbb{1}(X_0 > 0), \qquad \varepsilon \sim N(0, 1),$$

so that the data generation is nonlinear and heteroskedastic.

Missingness is MAR: $R = 1$ indicates Y is observed and $R = 0$ indicates Y is missing. We generate $R \sim \text{Bernoulli}(\pi(X))$ under three scenarios:

- Scenario A (mild MAR): $\text{logit}\{\pi(X)\} = -0.3 + 0.6X_0 - 0.8X_2$.
- Scenario B (moderate MAR, nonlinear): $\text{logit}\{\pi(X)\} = -0.5 + 0.9X_0 - 1.1X_2 + 0.4(X_1^2 - 1)$.
- Scenario C (hard MAR, interactions / weaker overlap): $\text{logit}\{\pi(X)\} = -1.0 + 1.2X_0 - 1.2X_2 + 0.8X_3X_4 - 0.5\,\mathbb{1}(X_5 = 1)$.

Scenario C induces stronger selection and weaker overlap in parts of the covariate space, which is particularly challenging for weighting-based calibration.

In each repetition, the data are split into training (50%), calibration (30%), and testing (20%) sets. We compare CM–MRL with the following approaches that use different strategies to quantify uncertainty under missing outcomes:

- **Han-Bootstrap**: the bootstrap predictive interval method of Han (2016).
- **CP-Logit**: weighted conformal prediction under MAR with propensity estimated by logistic regression.
- **CP-NN**: weighted conformal prediction with propensity estimated by a neural-network classifier.
- **CP-RF**: weighted conformal prediction with propensity estimated by random forest.

S1.3 Settings in real data analysis

In the real data ACTG 175, the average age of the subjects is 35 years with a standard deviation of 8.7 years. The cohort includes 1522 white subjects and 617 nonwhite subjects. The subjects include 1171 males and 368 females. Among the patients, 1253 have a history of antiretroviral treatment, and 776 are off-treatment before 96 weeks. Due to some subjects dropping out of the study, CD4 counts at 96 ± 5 weeks are missing for 797 subjects (missingness rate of 37%).

We compare our method with weighted conformal prediction methods whose propensity scores are estimated using machine-learning methods via

Table 1: Summary statistics for ACTG175 in real analysis overall and by missingness indicator R. Continuous variables are reported as mean (SD); categorical variables as n (%).

Variable	**Level**	**Overall**	**Observed ($R = 1$)**	**Missing ($R = 0$)**	**Type**
treat	0	532 (24.9%)	321 (23.9%)	211 (26.5%)	Categorical
	1	1607 (75.1%)	1021 (76.1%)	586 (73.5%)	Categorical
race	0	1522 (71.2%)	991 (73.8%)	531 (66.6%)	Categorical
	1	617 (28.8%)	351 (26.2%)	266 (33.4%)	Categorical
age		35.25 (8.71)	35.34 (8.66)	35.09 (8.80)	Continuous (mean(sd))
gender	0	368 (17.2%)	218 (16.2%)	150 (18.8%)	Categorical
	1	1771 (82.8%)	1124 (83.8%)	647 (81.2%)	Categorical
wtkg		75.13 (13.26)	74.78 (12.57)	75.70 (14.34)	Continuous (mean(sd))
offtrt	0	1363 (63.7%)	1075 (80.1%)	288 (36.1%)	Categorical
	1	776 (36.3%)	267 (19.9%)	509 (63.9%)	Categorical
symptom	0	1769 (82.7%)	1106 (82.4%)	663 (83.2%)	Categorical
	1	370 (17.3%)	236 (17.6%)	134 (16.8%)	Categorical
cd40		350.50 (118.57)	354.99 (116.91)	342.95 (121.03)	Continuous (mean(sd))
Outcome summary					
N		2139	1342	797	
cd496		328.57 (174.66)	328.57 (174.66)		Continuous (mean(sd))

the SuperLearner package in R.

- **Split conformal**: The mean regression parameters are estimated by the multiple robust method. We then impute the missing values and use simple split conformal prediction to construct the prediction interval.

https://cran.r-project.org/web/packages/SuperLearner/index.html

- **CP-NN**: The propensity $\hat{\pi}(x)$ is estimated using a neural network classifier. The resulting IPW weights are used to compute the weighted calibration quantile.

- **CP-RF**: The IPW weights based on the random forest propensity are applied to observed calibration residuals to obtain the weighted quantile.

S1.4 Interval Length under Misspecification

Let $\tilde{\mu} : \mathcal{X} \to \mathbb{R}$ be any predictor fitted on the training split with a possibly misspecified model. On the calibration split, define $\tilde{\varepsilon}_i = |Y_i - \tilde{\mu}(X_i)|$ and let $q_{1-\alpha}(\tilde{\varepsilon})$ be the $1-\alpha$-quantile of $\tilde{\varepsilon}$. Construct the CM–MRL weight vector $\{\hat{d}_i\}_{i\in\mathcal{I}_{\mathrm{cal}},\, R_i=1}$ by solving the empirical-likelihood program (2.14), which enforces (2.4)–(2.5). Let $\hat{q}_{\mathrm{MR}}$ solve the weighted score equation (2.15), and define the CM–MRL interval length $\mathrm{Len}_{\mathrm{MR}} = 2\,\hat{q}_{\mathrm{MR}}$.

For any *single-model* calibration scheme S that enforces only one of the moment blocks in (2.4)–(2.5), i.e., either a single propensity constraint for some j, or a single $\psi_{1-\alpha}$-moment constraint for some k, let $\hat{q}_S$ denote the corresponding weighted quantile estimator and $\mathrm{Len}_{\mathrm{S}} = 2\,\hat{q}_S$ its interval length.

Theorem 1 (Interval-length under misspecification)**.** *Assume (C1)–(C5)* *hold. For every $\epsilon > 0$ and $\delta \in (0,1)$, there exists N such that for all $n \geq N$,*

$$\mathbb{P}\big(\,\mathrm{Len}_{\mathrm{MR}} \leq \mathrm{Len}_{\mathrm{S}} + \epsilon\big) \;\geq\; 1 - \delta.$$

Thus, even when $\tilde{\mu}$ is misspecified, the CM–MRL interval is asymptotically no longer than any single-model interval that satisfies only one of (2.4)–(2.5).

Remark 1 (Condition for strict dominance)**.** If the extra balancing moments in (2.4)–(2.5) (beyond the single constraint used by S) have nonzero linear correlation with Z, i.e., they explain residual variation in $\mathbb{1}\{\tilde{\varepsilon} \leq q_{1-\alpha}(\tilde{\varepsilon})\}$, then $\kappa_{\mathrm{MR}} < \kappa_S$ and the dominance is strict. Equality occurs only if the added moments are L^2-orthogonal to Z.

Remark 2 (From fixed center to consistent center)**.** If $\tilde{\mu} = \hat{\mu}_{\mathrm{MR}}$ and $\hat{\mu}_{\mathrm{MR}} \to \mu_0$ in probability, then $q_{1-\alpha}(\tilde{\varepsilon}) \to q^*_{1-\alpha}$ and Theorem 1 reduces to Theorem 2 with the same variance ordering. The result thus implies shorter or equal asymptotic lengths at the nominal coverage when the MR center is consistent, and the no-longer-than guarantee remains valid even under misspecification of the center.

S1.5 Extension to quantile conformity scores

On the first calibration split, for each observed response $Y_i\,(R_i = 1)$, we compute the quantile conformity score,

$$\hat{S}_i = \max\left\{Y_i - \hat{f}_U\left(X_i\right), \hat{f}_L\left(X_i\right) - Y_i\right\},$$

where $\hat{f}_U$ and $\hat{f}_L$ are the lower and upper quantile estimates. Usually, $\hat{f}_U$ is the $1-\alpha/2$ conditional quantile estimator and $\hat{f}_L$ is the $\alpha/2$ conditional quantile estimator.

For the second calibration, for each missing $Y_i\,(R_i = 0)$, we impute Y_i from each outcome model k and compute imputed scores

$$\hat{S}_i^{(k)} = \frac{1}{T}\sum_{t=1}^{T} \max\left\{Y_i^t(\hat{b}^k) - \hat{f}_U\left(X_i\right), \hat{f}_L\left(X_i\right) - Y_i^t(\hat{b}^k)\right\},$$

$$Y_i^t(\hat{b}^k) \sim f_k\left(\cdot \mid X_i; \hat{b}^k\right).$$

Following the paper's second calibration in (2.11), with $\hat{\varepsilon}$ replaced by $\hat{S}$, for each k we obtain a model-wise score quantile $\hat{q}_{1-\alpha}^{(k)}$ by minimizing the $1-\alpha$-check loss over the whole calibration set,

$$\hat{q}_{1-\alpha}^{(k)} \in \arg\min_{q} \frac{1}{n_{\mathrm{cal}}} \sum_{i\in\mathcal{I}_{\mathrm{cal}}} \left\{R_i \rho_{1-\alpha}\left(\hat{S}_i^{(k)} - q\right) + \left(1-R_i\right)\rho_{1-\alpha}\left(\hat{S}_i^{(k)} - q\right)\right\}.$$

Now the EL weights $\left\{\hat{d}_i\right\}_{i\in\mathcal{I}_{\mathrm{cal}}\,,R_i=1}$ on complete calibration cases can be derived using the same EL program (2.14), but with the moment vector built from propensity balancing moments and score-quantile balancing moments. Thus, we can define $\hat{q}_{MR}$ as the root of the weighted score equation

$$\sum_{i\in\mathcal{I}_{\mathrm{cal}},R_i=1} \hat{d}_i \psi_{1-\alpha}\left(\hat{S}_i - \hat{q}_{MR}\right) = 0.$$

Finally, we can construct the missing-data MR-CQR interval using both quantile estimators and $\hat{q}_{MR}$:

$$\hat{C}_{MR}(x) = \left[\hat{f}_L(x) - \hat{q}_{MR}, \hat{f}_U(x) + \hat{q}_{MR}\right].$$

For the above conformalized quantile-based interval, we can derive theoretical results similar to those in Section 3 under the following conditions.

(C1*) (X_i, R_i, Y_i) are i.i.d., split independent, MAR $R \perp Y \mid X$, and overlap

$$0 < \inf_x \pi_0(x) \leq \sup_x \pi_0(x) < 1.$$

(C2*) Assume the cdf F_{S^*} is continuous in a neighborhood of $q^*_{1-\alpha}$. The density $f_{S^*}\left(q^*_{1-\alpha}\right) > 0$ and f_{S^*} is continuously differentiable at $q^*_{1-\alpha}$.

(C3*) At least one propensity model in Π is correctly specified and at least one outcome/imputation model in $\mathcal{F}$ is correctly specified for $Y \mid X$. In other words, the MR weighted quantile estimators $\hat{f}_L, \hat{f}_U$ are consistent for $f_{0,L}, f_{0,U}$ with convergence rate $\sqrt{n}$, i.e., $\left\|\hat{f}_L - f_{0,L}\right\|_\infty + \left\|\hat{f}_U - f_{0,U}\right\|_\infty = o_p\left(n^{-1/2}\right)$.

(C4*) Assume that $\{\pi_j(\cdot;\alpha)\}$ and the outcome model classes are Donsker classes. Thus, the induced class $\{(x,y) \mapsto \psi_{1-\alpha}\left(\max\{y - f_U(x), f_L(x) - y\} - q\right)\}$ also comes from a Donsker class.

(C5*) Quantile check losses for $\gamma \in \{\alpha/2, 1-\alpha/2\}$ and subgradients $\psi_\gamma(u) =$

$\gamma - \mathbf{1}(u < 0)$ are bounded. The conditional density $f_{Y|X}(\cdot \mid x)$ exists and is bounded and bounded away from 0.

Then, we can derive the same asymptotic results as shown in Section 3.

Theorem 2. *Under (C1*)-(C5*), the EL-calibrated estimator $\hat{q}_{MR}$ built from the CQR score satisfies $\hat{q}_{MR} \xrightarrow{p} q^*_{1-\alpha}$. Moreover, if f_{S^*} is continuously differentiable at $q^*_{1-\alpha}$, then $\sqrt{m_{\mathrm{cal}}}\left(\hat{q}_{MR} - q^*_{1-\alpha}\right) \Rightarrow \mathcal{N}\left(0, \frac{\alpha(1-\alpha)}{\left\{f_{S^*}\left(q^*_{1-\alpha}\right)\right\}^2} \cdot \kappa_{MR}\right)$. where $\kappa_{MR} \leq 1$ is the same type of EL efficiency factor as in the paper.*

S1.6 Extension of Condition (C4) to DML

Let $D_i = (X_i, R_i, R_iY_i)$ and let $\varepsilon^* = |Y - \mu_0(X)|$ be the oracle conformity score with $1-\alpha$-quantile $q^*_{1-\alpha}$. Recall $\psi_{1-\alpha}(u) = 1 - \alpha - \mathbf{1}(u < 0)$. The quantile $q^*_{1-\alpha}$ is equivalently characterized as the unique solution to

$$\Psi(q) := \mathbb{E}\big[\psi_{1-\alpha}(\varepsilon^* - q)\big] = 0, \qquad \Psi'(q^*_{1-\alpha}) = -f_{\varepsilon^*}(q^*_{1-\alpha}) < 0. \tag{S1.1}$$

Under MAR, ε^* is observed only when $R = 1$. Define two nuisance objects:

$$\pi_0(x) = \mathbb{P}(R = 1 \mid X = x), \qquad m_0(x, q) = \mathbb{E}\big[\psi_{1-\alpha}(\varepsilon^* - q) \mid X = x\big].$$

Consider the orthogonal (doubly robust) score

$$\varphi(W; q, \eta) := m(X, q) + \frac{R}{\pi(X)}\Big\{\psi_{1-\alpha}(\varepsilon - q) - m(X, q)\Big\}, \qquad \eta = (\pi, m), \tag{S1.2}$$

where $\varepsilon = |Y - \mu_0(X)|$ when $R = 1$. Then

$$\mathbb{E}\big[\varphi(W; q, \eta_0)\big] = \mathbb{E}\big[\psi_{1-\alpha}(\varepsilon^* - q)\big] = \Psi(q), \tag{S1.3}$$

so $q^*_{1-\alpha}$ is also the unique root of $\mathbb{E}[\varphi(W; q, \eta_0)] = 0$.

Lemma 1 (Neyman orthogonality). *Let $\Psi_\varphi(q, \eta) := \mathbb{E}[\varphi(W; q, \eta)]$ with φ defined in* (S1.2)*. Then the Gateaux derivative of $\Psi_\varphi(q, \eta)$ with respect to η vanishes at $(q, \eta) = (q^*_{1-\alpha}, \eta_0)$:*

$$\partial_\eta \Psi_\varphi(q^*_{1-\alpha}, \eta)\big|_{\eta=\eta_0} = 0.$$

Proof. Write $\psi = \psi_{1-\alpha}(\varepsilon^* - q^*_{1-\alpha})$ and $m_0(X) = m_0(X, q^*_{1-\alpha})$ for brevity. For a perturbation $m_t = m_0 + t\Delta_m$,

$$\begin{aligned}\frac{d}{dt}\mathbb{E}\left[m_t(X) + \frac{R}{\pi_0(X)}\{\psi - m_t(X)\}\right]_{t=0} &= \mathbb{E}\left[\Delta_m(X)\Big(1 - \frac{R}{\pi_0(X)}\Big)\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[1 - \frac{R}{\pi_0(X)} \mid X\right] \Delta_m(X)\right] = 0,\end{aligned}$$

since $\mathbb{E}[R \mid X] = \pi_0(X)$. For a perturbation $\pi_t = \pi_0 + t\Delta_\pi$,

$$\frac{d}{dt}\mathbb{E}\left[m_0(X) + \frac{R}{\pi_t(X)}\{\psi - m_0(X)\}\right]_{t=0} = \mathbb{E}\left[-\frac{R}{\pi_0(X)^2}\{\psi - m_0(X)\}\Delta_\pi(X)\right] = 0,$$

because $\mathbb{E}[\psi - m_0(X) \mid X] = 0$ by definition of $m_0(X, q^*_{1-\alpha})$. Combining the two derivatives yields the claim. $\square$

Let the calibration index set be partitioned into K folds $\{\mathcal{I}_1, \ldots, \mathcal{I}_K\}$ of equal size, and let $k(i)$ denote the fold containing i. For each fold k, estimate nuisances $\widehat{\eta}^{(-k)} = (\widehat{\pi}^{(-k)}, \widehat{m}^{(-k)})$ using data excluding fold k with any ML method. Define the cross-fitted sample moment

$$\widehat{\Psi}_{\text{cf}}(q) := \frac{1}{n_{\text{cal}}} \sum_{i \in \mathcal{I}_{\text{cal}}} \varphi\big(W_i; q, \widehat{\eta}^{(-k(i))}\big), \qquad \widehat{q}_{\text{cf}} \text{ solves } \ \widehat{\Psi}_{\text{cf}}(\widehat{q}_{\text{cf}}) = 0. \quad \text{(S1.4)}$$

When the main algorithm uses a single train/calibration split, it already implements a special case of sample splitting; cross-fitting is an optional refinement to reuse data more efficiently.

Replacement for (C4): a DML-type condition. We replace (C4) by the following high-level requirement.

1. $\pi_0(x)$ is bounded away from 0 and 1, and $\mathbb{E}|\varphi(W; q^*_{1-\alpha}, \eta_0)|^{2+\delta} < \infty$ for some $\delta > 0$.

2. The nuisance estimators $\widehat{\eta}^{(-k)}$ are trained excluding fold k, so that conditional on the training sample, $\{\varphi(W_i; \cdot, \widehat{\eta}^{(-k(i))}) : i \in \mathcal{I}_{\text{cal}}\}$ are independent across i.

3. There exists a deterministic neighborhood $\mathcal{N}$ of $q^*_{1-\alpha}$ such that

$$\sup_{q \in \mathcal{N}} \|\widehat{m}^{(-k)}(\cdot, q) - m_0(\cdot, q)\|_{L_2(P)} = o_p(1), \qquad \|\widehat{\pi}^{(-k)} - \pi_0\|_{L_2(P)} = o_p(1),$$

uniformly over folds k, and the second-order remainder condition holds:

$$\sqrt{m_{\text{cal}}} \left(\sup_{q \in \mathcal{N}} \|\widehat{m}^{(-k)}(\cdot, q) - m_0(\cdot, q)\|_{L_2(P)} \right) \cdot \|\widehat{\pi}^{(-k)} - \pi_0\|_{L_2(P)} = o_p(1), \quad \text{(S1.5)}$$

uniformly over k. A sufficient condition is $\sup_{q\in\mathcal{N}} \|\widehat{m}^{(-k)}(\cdot,q) - m_0(\cdot,q)\|_{L_2(P)} = o_p(m_{\mathrm{cal}}^{-1/4})$ and $\|\widehat{\pi}^{(-k)} - \pi_0\|_{L_2(P)} = o_p(m_{\mathrm{cal}}^{-1/4})$.

The next theorem shows that the DML/cross-fitted estimator $\widehat{q}_{\mathrm{cf}}$ achieves the same root-m_{cal} limit as the complete-case empirical quantile, without requiring Donsker conditions.

Theorem 3. *Assume (C1)–(C3) and (C5) from the main text, replace (C4) by (C4†), and keep (C2). Then $\widehat{q}_{\mathrm{cf}} \xrightarrow{p} q^*_{1-\alpha}$. If f_{ε^*} is continuously differentiable at $q^*_{1-\alpha}$, then*

$$\sqrt{m_{\mathrm{cal}}}\big(\widehat{q}_{\mathrm{cf}} - q^*_{1-\alpha}\big) \;\Rightarrow\; \mathcal{N}\left(0,\ \frac{\mathrm{Var}\big(\varphi(W; q^*_{1-\alpha}, \eta_0)\big)}{\{f_{\varepsilon^*}(q^*_{1-\alpha})\}^2}\right).$$

Proof. We give a step-by-step derivation in the style of Chernozhukov et al. (2018).

Step 1 (Identification and Jacobian). By (S1.3), $\Psi_\varphi(q,\eta_0) = \Psi(q)$ and $\Psi(q^*_{1-\alpha}) = 0$. Condition (C2) ensures $\partial_q \Psi_\varphi(q^*_{1-\alpha},\eta_0) = \Psi'(q^*_{1-\alpha}) = -f_{\varepsilon^*}(q^*_{1-\alpha}) \neq 0$.

Step 2 (Z-estimator linearization). Since $\widehat{q}_{\mathrm{cf}}$ solves $\widehat{\Psi}_{\mathrm{cf}}(q) = 0$, a mean-value expansion yields

$$0 = \widehat{\Psi}_{\mathrm{cf}}(q^*_{1-\alpha}) + \partial_q \widehat{\Psi}_{\mathrm{cf}}(\widetilde{q})\,(\widehat{q}_{\mathrm{cf}} - q^*_{1-\alpha}), \quad \text{for some } \widetilde{q} \text{ between } \widehat{q}_{\mathrm{cf}} \text{ and } q^*_{1-\alpha}.$$

Thus,

$$\sqrt{m_{\mathrm{cal}}}(\widehat{q}_{\mathrm{cf}} - q^*_{1-\alpha}) = -\Big(\partial_q \widehat{\Psi}_{\mathrm{cf}}(\widetilde{q})\Big)^{-1} \sqrt{m_{\mathrm{cal}}}\,\widehat{\Psi}_{\mathrm{cf}}(q^*_{1-\alpha}). \tag{S1.6}$$

Under (C2) and the uniform consistency in (C4†), $\partial_q \widehat{\Psi}_{\mathrm{cf}}(\widetilde{q}) \to_p -f_{\varepsilon^*}(q^*_{1-\alpha})$.

Step 3 (Decomposition of the empirical term). Write $\mathbb{P}_n$ for the empirical measure on $\mathcal{I}_{\mathrm{cal}}$ and $\mathbb{P}$ for expectation. Decompose

$$\begin{aligned}\widehat{\Psi}_{\mathrm{cf}}(q^*_{1-\alpha}) =& (\mathbb{P}_n - \mathbb{P})\varphi(\cdot; q^*_{1-\alpha}, \eta_0) + \mathbb{P}\big[\varphi(\cdot; q^*_{1-\alpha}, \widehat{\eta}) - \varphi(\cdot; q^*_{1-\alpha}, \eta_0)\big] \\ & + (\mathbb{P}_n - \mathbb{P})\big[\varphi(\cdot; q^*_{1-\alpha}, \widehat{\eta}) - \varphi(\cdot; q^*_{1-\alpha}, \eta_0)\big]\,,\end{aligned}$$

where $\widehat{\eta}$ stands for the fold-specific cross-fitted nuisance in (S1.4).

Step 4 (Second-order bias via orthogonality). A functional Taylor expansion gives

$$\begin{aligned}\mathbb{P}\big[\varphi(\cdot; q^*_{1-\alpha}, \widehat{\eta}) - \varphi(\cdot; q^*_{1-\alpha}, \eta_0)\big] =& \partial_\eta \Psi_\varphi(q^*_{1-\alpha}, \eta_0)\,[\widehat{\eta} - \eta_0] \\ & + O_p\Big(\|\widehat{\pi} - \pi_0\|_{L_2(P)} \cdot \sup_{q\in\mathcal{N}} \|\widehat{m}(\cdot, q) - m_0(\cdot, q)\|_{L_2(P)}\Big).\end{aligned}$$

By Lemma 1, the linear term vanishes, and the remainder is controlled by (S1.5). Hence $\sqrt{m_{\mathrm{cal}}}\,\mathbb{P}[\varphi(\cdot; q^*_{1-\alpha}, \widehat{\eta}) - \varphi(\cdot; q^*_{1-\alpha}, \eta_0)] = o_p(1)$.

Step 5 (Empirical process remainder without Donsker). By cross-fitting, for each fold k the evaluation sample is independent of $\widehat{\eta}^{(-k)}$. Conditional on the training sample,

$$\begin{aligned}&\mathrm{Var}\Big((\mathbb{P}_n - \mathbb{P})\big[\varphi(\cdot; q^*_{1-\alpha}, \widehat{\eta}) - \varphi(\cdot; q^*_{1-\alpha}, \eta_0)\big] \;\Big|\; \widehat{\eta}\Big) \\ &\quad \lesssim \frac{1}{m_{\mathrm{cal}}}\,\mathbb{E}\Big[\big(\varphi(W; q^*_{1-\alpha}, \widehat{\eta}) - \varphi(W; q^*_{1-\alpha}, \eta_0)\big)^2 \,\Big|\widehat{\eta}\Big]\,,\end{aligned}$$

so the term is $o_p(m_{\mathrm{cal}}^{-1/2})$ whenever $\|\widehat{\eta} - \eta_0\|_{L_2(P)} = o_p(1)$, as imposed by

(C4†). This is the standard cross-fitting argument in DML.

Step 6 (CLT and conclusion). Combining Steps 3–5 yields $\sqrt{m_{\mathrm{cal}}}\,\widehat{\Psi}_{\mathrm{cf}}(q^*_{1-\alpha}) =$

$$\sqrt{m_{\mathrm{cal}}}\,(\mathbb{P}_n - \mathbb{P})\varphi(\cdot; q^*_{1-\alpha}, \eta_0) + o_p(1) \Rightarrow \mathcal{N}\big(0, \mathrm{Var}\big(\varphi(W; q^*_{1-\alpha}, \eta_0)\big)\big)\,,$$

by the Lindeberg–Feller CLT under (C4)(i). Plugging this and $\partial_q \widehat{\Psi}_{\mathrm{cf}}(\widetilde{q}) \to_p -f_{\varepsilon^*}(q^*_{1-\alpha})$ into (S1.6) completes the proof. □

Remark 3. One convenient way to estimate $m_0(x, q) = \mathbb{E}[\psi_{1-\alpha}(\varepsilon^* - q) \mid X = x]$ under MAR is via conditional-imputation draws: for each x, draw $Y^{(t)} \sim \widehat{f}(\cdot \mid x)$ from a fitted outcome model, compute $\varepsilon^{(t)} = |Y^{(t)} - \widehat{\mu}(x)|$, and set $\widehat{m}(x, q) = T^{-1}\sum_{t=1}^{T} \psi_{1-\alpha}(\varepsilon^{(t)} - q)$. This preserves the DML requirement that $\widehat{m}(\cdot, q)$ depends only on training data for each evaluation point.

S2 Theoretical Results Proof

Proof of Proposition 1. The weight is derived by the following steps. For the true missing-at-random probability $\mathbb{P}(R = 1|X)$, which we denote by

$\pi(x)$, we obtain, for any function $h(x)$,

$$
\begin{aligned}
0 &= \mathbb{E}\left\{\mathbb{E}\left(\frac{R}{\mathbb{P}(R=1|x)}\left[h(x)-\mathbb{E}h(x)\right] \mid X=x\right)\right\} = \mathbb{E}\left(\frac{R}{\pi(x)}\left[h(X)-\mathbb{E}h(X)\right]\right) \\
&= \mathbb{P}(R=1)\mathbb{E}\left(R\frac{[h(X)-\mathbb{E}h(X)]}{\pi(x)} \mid R=1\right) \\
&\qquad + \mathbb{P}(R=0)\mathbb{E}\left(\frac{R\left[h(X)-\mathbb{E}h(X)\right]}{w(x)} \mid R=0\right) \\
&= \mathbb{P}(R=1)\mathbb{E}\left(R\frac{[h(X)-\mathbb{E}h(X)]}{\pi(x)} \mid R=1\right).
\end{aligned}
$$

Let $w(x) = \pi^{-1}(x)$. Thus, for any function $h(x)$ whose expectation exists, $\mathbb{E}(w(X)[h(X) - \mathbb{E}\{h(X)\}] \mid R=1) = 0$. To prove the consistency of $\hat{q}_{\mathrm{MR}}$ in Section 2.2, we can begin by letting $h(x)$ be $\pi^j(x, \hat{a}^j)$ and $\mathbb{E}\psi_{1-\alpha}({\hat{\varepsilon}_i}^k, q_0|X; \hat{b}^k) = 1 - \alpha - \mathbb{P}^k((Y - \mu_{\mathrm{MR}}(X)) - q^*|X, \hat{a}^k)$, where $\mathbb{E}^k\left(\cdot \mid X; b^k\right)$ and $\mathbb{P}^k(\cdot \mid X; b^k)$ are the conditional expectation and the conditional probability under the density $f^k_\varepsilon(\varepsilon \mid X, b^k)$. Using these particular functions as $h(X)$ and the $\hat{a}^j$, $\hat{b}^k$, and $\hat{q}^k$ from Section 2.1, respectively, we construct the following constrained optimization problem:

$$
\begin{aligned}
&\max_{w_i} \sum_{i=1}^{m} \log w_i \quad \text{such that} \\
&\quad w_i \geq 0 \quad (i = 1, \ldots, m), \\
&\quad \sum_{i=1}^{m} w_i = 1, \\
&\quad \sum_{i=1}^{m} w_i \hat{\boldsymbol{v}}_i(\hat{\boldsymbol{a}}, \hat{\boldsymbol{b}}, \hat{\boldsymbol{q}}) = \mathbf{0}.
\end{aligned}
$$

By introducing Lagrange multipliers, we obtain

$$l(w) = \sum_{i=1}^{m} \log w_i - \boldsymbol{\lambda}_1 \sum_{i=1}^{m} w_i \hat{\boldsymbol{v}}_i(\hat{\boldsymbol{a}}, \hat{\boldsymbol{b}}, \hat{\boldsymbol{q}}) - \lambda_2 (\sum_{i=1}^{m} w_i - 1).$$

By taking derivatives with respect to w_i, $\boldsymbol{\lambda}_1$, and λ_2, we obtain

$$\frac{\partial l(w)}{\partial w_i} = \frac{1}{w_i} - \boldsymbol{\lambda}_1 \hat{\boldsymbol{v}}_i(\hat{\boldsymbol{a}}, \hat{\boldsymbol{b}}, \hat{\boldsymbol{q}}) - \lambda_2 = 0 \quad i = 1, \dots, m \tag{S2.7}$$

$$\frac{\partial l(w)}{\partial \lambda_{1l}} = \sum_{i=1}^{m} w_i \hat{\boldsymbol{v}}_i(\hat{\boldsymbol{a}}, \hat{\boldsymbol{b}}, \hat{\boldsymbol{q}}) = 0 \quad l = 1, \dots, J + Kp \tag{S2.8}$$

$$\frac{\partial l(w)}{\partial \lambda_2} = \sum_{i=1}^{m} w_i - 1 = 0. \quad i = 1, \dots, m \tag{S2.9}$$

Combining (S2.7), (S2.8), and (S2.9), we obtain $\lambda_2 = 1$ and

$$w_i = \frac{1}{\boldsymbol{\lambda}_1 \hat{\boldsymbol{v}}_i(\hat{\boldsymbol{a}}, \hat{\boldsymbol{b}}, \hat{\boldsymbol{q}}) + 1}. \tag{S2.10}$$

Combining (S2.8) and (S2.10), we obtain that $\boldsymbol{\lambda}_1$ must satisfy

$$\sum_{i=1}^{m} \frac{\hat{\boldsymbol{v}}_i(\hat{\boldsymbol{a}}, \hat{\boldsymbol{b}}, \hat{\boldsymbol{q}})}{\boldsymbol{\lambda}_1 \hat{\boldsymbol{v}}_i(\hat{\boldsymbol{a}}, \hat{\boldsymbol{b}}, \hat{\boldsymbol{q}}) + 1} = 0.$$

This is also equivalent to solving the optimization problem with respect to $\boldsymbol{\lambda}_1$:

$$\hat{\boldsymbol{\lambda}}_1 = \arg\min_{\boldsymbol{\lambda}_1} \sum_{i=1}^{m} \log[\boldsymbol{\lambda}_1 \hat{\boldsymbol{v}}_i(\hat{\boldsymbol{a}}, \hat{\boldsymbol{b}}, \hat{\boldsymbol{q}}) + 1].$$

This is how we obtain the multiple robust weight $\hat{w}_i$, which estimates $w(x) = \mathbb{P}(R = 1|X)$ on complete cases. □

Proof of Theorem 2. Let $\mathcal{I}_{\text{cal}}$ be the calibration set, $S = \{i \in \mathcal{I}_{\text{cal}} : R_i = 1\}$ the complete calibration cases, and $m_{\text{cal}} = |S|$. Define the observed conformity scores $\hat{\varepsilon}_i = \big|Y_i - \hat{\mu}_{\mathrm{MR}}(X_i)\big|$ for $i \in S$ and the EL weights

$$\hat{d}_i \;=\; \frac{1}{m_{\mathrm{cal}}}\,\frac{1}{1+\hat{\lambda}^\top \hat{v}_i}, \qquad \hat{\lambda} \in \arg\min_{\lambda}\left\{-\frac{1}{m_{\mathrm{cal}}}\sum_{i\in S}\log\big(1+\lambda^\top \hat{v}_i\big)\right\},$$

where $\hat{v}_i$ stacks centered propensity and centered $\psi_{1-\alpha}$-moments defined in (2.14). The double-calibrated quantile $\hat{q}_{\mathrm{MR}}$ solves the weighted score equation

$$\Psi_m(q) \;:=\; \sum_{i\in S}\hat{d}_i\,\psi_{1-\alpha}\big(\hat{\varepsilon}_i - q\big) \;=\; 0. \tag{S2.11}$$

Write the population target map

$$\Psi(q) \;=\; \mathbb{E}\big[\psi_{1-\alpha}\big(\varepsilon^* - q\big)\big] \;=\; 1-\alpha-\mathbb{P}(\varepsilon^* \le q)\,, \qquad \varepsilon^* := \big|Y - \mu_0(X)\big|.$$

By (C2), $\Psi(\cdot)$ is continuous and strictly decreasing at $q^*_{1-\alpha}$ with derivative $\Psi'(q^*_{1-\alpha}) = -f_{\varepsilon^*}(q^*_{1-\alpha}) < 0$.

Step 1 Uniform stochastic equicontinuity. We show $\sup_{q\in\mathcal{Q}}|\Psi_m(q)-\Psi(q)| = o_p(1)$ for any compact $\mathcal{Q}$ around $q^*_{1-\alpha}$. Decompose

$$\begin{aligned}\Psi_m(q)-\Psi(q) = &\underbrace{\sum_{i\in S}\hat{d}_i\big[\psi_{1-\alpha}\big(\hat{\varepsilon}_i - q\big)-\psi_{1-\alpha}\big(\varepsilon_i^* - q\big)\big]}_{(A)}\\ &+\underbrace{\left\{\sum_{i\in S}\hat{d}_i\,\psi_{1-\alpha}\big(\varepsilon_i^*-q\big)-\mathbb{E}\big[\psi_{1-\alpha}\big(\varepsilon^*-q\big)\big]\right\}}_{(B)},\end{aligned}$$

where $\varepsilon_i^* = |Y_i - \mu_0(X_i)|$. For (A), by sample splitting (C1), the training estimate $\hat{\mu}_{\mathrm{MR}}$ is independent of calibration data; by Theorem 2 and continuity of the absolute value, $\sup_{i\in S}|\hat{\varepsilon}_i - \varepsilon_i^*| = o_p(1)$ in L_1 or in probability.

Since $q \mapsto \psi_{1-\alpha}(u-q)$ has bounded variation and a single jump at $q = u$, the difference (A) is $o_p(1)$ uniformly on compact $\mathcal{Q}$ via a standard argument for indicator classes, Glivenko–Cantelli, and continuity of the cdf near $q^*_{1-\alpha}$; see Van der Vaart (2000).

For (B), we use the EL weight expansion (Lemma 2 below) to write

$$\hat{d}_i = \frac{1}{m_{\text{cal}}}\Big\{1-\hat{\lambda}^\top \hat{v}_i\Big\}+r_{i,n}, \qquad \max_{i\in S}|r_{i,n}| = o_p\big(m_{\text{cal}}^{-1}\big), \qquad \|\hat{\lambda}\| = O_p\big(m_{\text{cal}}^{-1/2}\big),$$

and the calibration condition $\sum_{i\in S}\hat{d}_i\hat{v}_i = 0$. Because $\mathbb{E}[\hat{v}_i] = o_p(1)$ by centering (definitions of $\hat{\theta}^{(j)}$ and $\hat{\xi}^{(k)}$) and (C3) guarantees at least one correctly specified moment block, the weighted empirical process for the Donsker class $\{\psi_{1-\alpha}(\varepsilon^* - q) : q \in \mathcal{Q}\}$ is $o_p(1)$ uniformly; see Qin and Lawless (1994). Therefore, $\sup_{q\in\mathcal{Q}}|(B)| = o_p(1)$.

Consistency. By (C2), $\Psi(\cdot)$ has a unique zero at $q^*_{1-\alpha}$ and is strictly decreasing there. Since $\Psi_m \to \Psi$ uniformly on $\mathcal{Q}$, Theorem 5.9 of Van der Vaart (2000) yields $\hat{q}_{\text{MR}} \xrightarrow{p} q^*_{1-\alpha}$.

Step 2 Bahadur representation. Define the weighted cdf of complete-case conformity scores at q,

$$\hat{F}_w(q) = \sum_{i\in S}\hat{d}_i\,\mathbb{1}\{\hat{\varepsilon}_i \le q\}, \qquad Z_i(q) := 1-\alpha-\mathbb{1}\{\varepsilon^*_i \le q\}.$$

Equation (S2.11) is equivalent to $\hat{F}_w(\hat{q}_{\text{MR}}) = 1-\alpha$. A standard weighted-quantile Bahadur expansion (Lemma 3 below) gives, uniformly in a neighborhood of $q^*_{1-\alpha}$,

$$\hat{q}_{\mathrm{MR}} - q^*_{1-\alpha} \;=\; \frac{1-\alpha-\hat{F}_w(q^*_{1-\alpha})}{f_{\varepsilon^*}(q^*_{1-\alpha})} \;+\; o_p\big(m_{\mathrm{cal}}^{-1/2}\big).$$

By Step 1, replacing $\hat{\varepsilon}_i$ with ε^*_i in $\hat{F}_w(\cdot)$ incurs $o_p(m_{\mathrm{cal}}^{-1/2})$. Hence,

$$\sqrt{m_{\mathrm{cal}}}\,\big(\hat{q}_{\mathrm{MR}} - q^*_{1-\alpha}\big) \;=\; \frac{1}{f_{\varepsilon^*}(q^*_{1-\alpha})}\cdot\sqrt{m_{\mathrm{cal}}}\sum_{i\in S}\hat{d}_i\, Z_i(q^*_{1-\alpha}) \;+\; o_p(1). \quad \text{(S2.12)}$$

Step 3 CLT for the calibrated mean. Let $V_i := \hat{v}_i$ evaluated at the limits of the nuisance estimates. By sample splitting and (C3), the plug-in estimation error is $o_p(m_{\mathrm{cal}}^{-1/2})$. The EL optimality conditions imply $\sum_{i\in S}\hat{d}_i V_i = 0$ and the expansion

$$\hat{d}_i \;=\; \frac{1}{m_{\mathrm{cal}}}\Big\{1-\hat{\lambda}^\top V_i\Big\} + o_p\big(m_{\mathrm{cal}}^{-1}\big), \qquad \hat{\lambda} \;=\; \Big(\bar{\Sigma}_{VV}\Big)^{-1}\bar{\Sigma}_{VZ} \;+\; o_p\big(m_{\mathrm{cal}}^{-1/2}\big),$$

where $\bar{\Sigma}_{VV} = \frac{1}{m_{\mathrm{cal}}}\sum_{i\in S} V_i V_i^\top$ and $\bar{\Sigma}_{VZ} = \frac{1}{m_{\mathrm{cal}}}\sum_{i\in S} V_i Z_i(q^*_{1-\alpha})$. Thus,

$$\sum_{i\in S}\hat{d}_i Z_i(q^*_{1-\alpha}) \;=\; \frac{1}{m_{\mathrm{cal}}}\sum_{i\in S}\Big\{Z_i(q^*_{1-\alpha}) - V_i^\top\gamma_0\Big\} + o_p\big(m_{\mathrm{cal}}^{-1/2}\big), \qquad \gamma_0 := \Sigma_{VV}^{-1}\Sigma_{VZ},$$

with $\Sigma_{VV} = \mathbb{E}[V_i V_i^\top]$ and $\Sigma_{VZ} = \mathbb{E}[V_i Z_i(q^*_{1-\alpha})]$. By the multivariate CLT and Slutsky,

$$\sqrt{m_{\mathrm{cal}}}\sum_{i\in S}\hat{d}_i Z_i(q^*_{1-\alpha}) \;\rightsquigarrow\; \mathcal{N}\Big(0,\ \mathrm{Var}\big\{Z_i(q^*_{1-\alpha}) - V_i^\top\gamma_0\big\}\Big).$$

Observe $\mathrm{Var}\big\{Z_i(q^*_{1-\alpha})\big\} = 1-\alpha(1-1-\alpha)$ and

$$\mathrm{Var}\big\{Z_i(q^*_{1-\alpha}) - V_i^\top\gamma_0\big\} \;=\; 1-\alpha(1-1-\alpha)\cdot\kappa_{\mathrm{MR}}, \qquad \kappa_{\mathrm{MR}} \;=\; 1-\frac{\mathrm{Cov}\big(Z_i,V_i\big)\,\Sigma_{VV}^{-1}\,\mathrm{Cov}\big(V_i,Z_i\big)}{1-\alpha(1-1-\alpha)}.$$

By the projection theorem, $\kappa_{\mathrm{MR}} \in [0,1]$, with $\kappa_{\mathrm{MR}} = 1$ when V_i carries no linear information for Z_i and $\kappa_{\mathrm{MR}} < 1$ when the balancing moments are informative.

Combining (S2.12) with the above CLT yields

$$\sqrt{m_{\mathrm{cal}}}\,\left(\hat{q}_{\mathrm{MR}} - q^*_{1-\alpha}\right) \;\rightsquigarrow\; \mathcal{N}\Big(0,\; \frac{1-\alpha(1-1-\alpha)}{\{f_{\varepsilon^*}(q^*_{1-\alpha})\}^2} \cdot \kappa_{\mathrm{MR}}\Big),$$

and consistency follows from Step 1. This proves the theorem. □

Lemma 2 (EL weight expansion)**.** *Under (C1)–(C5) and (C3), with fixed dimension of V_i, the EL minimizer $\hat{\lambda}$ satisfies $\|\hat{\lambda}\| = O_p(m_{\mathrm{cal}}^{-1/2})$ and*

$$\hat{d}_i \;=\; \frac{1}{m_{\mathrm{cal}}}\Big\{1 - \hat{\lambda}^\top V_i\Big\} + o_p\big(m_{\mathrm{cal}}^{-1}\big), \qquad \hat{\lambda} \;=\; \Big(\bar{\Sigma}_{VV}\Big)^{-1}\bar{\Sigma}_{VZ} + o_p\big(m_{\mathrm{cal}}^{-1/2}\big),$$

where $\bar{\Sigma}_{VV}$ and $\bar{\Sigma}_{VZ}$ are sample analogues. Moreover, $\sum_{i\in S}\hat{d}_i V_i = 0$ holds exactly.

Use strict convexity and KKT conditions of the EL problem; a Taylor expansion of the first-order conditions around $\lambda = 0$, along with $\mathbb{E}[V_i] = 0$ (centering) and $\mathbb{E}[V_i V_i^\top]$ nonsingular, yields the stated rates.

Lemma 3 (Weighted Bahadur representation)**.** *Let $\hat{F}_w(q) = \sum_{i\in S}\hat{d}_i \mathbb{1}\{\varepsilon^*_i \le q\}$ with $\hat{d}_i$ as in Lemma 2. If F_{ε^*} is continuously differentiable at $q^*_{1-\alpha}$ with density $f_{\varepsilon^*}(q^*_{1-\alpha}) > 0$, then the $1-\alpha$-quantile $\hat{q}$ solving $\hat{F}_w(\hat{q}) = 1-\alpha$ obeys*

$$\hat{q} - q^*_{1-\alpha} = \frac{1-\alpha-\hat{F}_w(q^*_{1-\alpha})}{f_{\varepsilon^*}(q^*_{1-\alpha})} + o_p\big(m_{\mathrm{cal}}^{-1/2}\big).$$

Use a Taylor expansion of $q \mapsto \hat{F}_w(q)$ around $q^*_{1-\alpha}$; the derivative converges in probability to $f_{\varepsilon^*}(q^*_{1-\alpha})$ by the weighted Glivenko–Cantelli/CLT and the fact that the random weights are asymptotically close to constants at rate $m_{\text{cal}}^{-1/2}$.

Proof of Theorem 3. Since (X_{n+1}, Y_{n+1}) is sampled i.i.d. from $f(Y|X)$, we have

$$
\begin{aligned}
&\mathbb{P}\left(Y_{n+1} \in \hat{C}_{\text{MR}}\right)\\
=&\mathbb{P}\left(|Y_{n+1} - \hat{\mu}_{\text{MR}}| \leq \hat{Q}_{\text{MR}}(E_i)\right)\\
=&\mathbb{P}\left(\hat{\varepsilon}_{n+1} \leq q^* + \hat{Q}_{\text{MR}}(E_i) - q^*\right)\\
\geq&\mathbb{P}\left(\hat{\varepsilon}_{n+1} \leq q^*\right) - \mathbb{P}\left(\hat{\varepsilon}_{n+1} \leq |\hat{Q}_{\text{MR}}(E_i) - q^*|\right)\\
\geq&1 - \alpha + o_p(1)
\end{aligned}
$$

The last inequality follows from Theorem 2. The proof of Theorem 3 is completed. □

Proof of Theorem 4. In Feldman et al. (2021), various metrics of conditional coverage are proposed. Adopting the asymptotic conditional coverage notion of Lei et al. (2018), we need to verify that the second moment of

$\hat{\mu}_{\mathrm{MR}} - \mu_0$ is $o_p(1)$.

$$
\begin{aligned}
&\left|\mathbb{P}(Y \in \hat{C}_{\mathrm{MR}}(x) \mid X = x) - (1 - 1 - \alpha)\right| \\
=&\left|\mathbb{P}(Y \in \hat{C}_{\mathrm{MR}}(x) \mid X = x) - \mathbb{P}(Y \in C^*(x) \mid X = x)\right| \\
\leq&\left|\mathbb{P}(Y \in C^*(x) \Delta \hat{C}_{\mathrm{MR}}(x) \mid X = x)\right|.
\end{aligned}
$$

where $C^*(x)\Delta\hat{C}_{\mathrm{MR}}(x)$ is the symmetric difference between the sets $C^*(x)$ and $\hat{C}_{\mathrm{MR}}(x)$. It contains the elements that are either in $C^*(x)$ or in $\hat{C}_{\mathrm{MR}}(x)$ but not in both. Since both sets are confidence intervals, we can bound the difference in length for fixed x,

$$
|C^*(x)\Delta\hat{C}_{\mathrm{MR}}(x)| \leq \underbrace{|\hat{\mu}_{\mathrm{MR}}(x) - \mu_0(x)|}_{(A.1)} + 2\underbrace{|\hat{q}_{\mathrm{MR}} - q^*|}_{(A.2)}.
$$

From Theorem 2 and Theorem 3, we obtain that $\hat{\beta}_{\mathrm{MR}} - \beta_0 = o_P(1)$ and $(A.2) = o_P(1)$. By Condition (C2), X has a bounded subgaussian norm. Thus $|\hat{\mu}_{\mathrm{MR}}(x) - \mu^*(x)| \leq |\beta_0 - \beta_{\mathrm{MR}}| = o_p(1)$ for fixed x. We define the subset

$$
\mathcal{S} = \left\{x : L(C^*(x)\Delta\hat{C}_{\mathrm{MR}}(x)) = o_p(1)\right\}.
$$

Then, we obtain

$$
\sup_{x\in\mathcal{S}}\left|\mathbb{P}(Y \in \hat{C}_{\mathrm{MR}}(x) \mid X = x) - (1 - 1 - \alpha)\right| = o_p(1).
$$

Therefore, conditional coverage is proved under the consistency of the multiple robust estimates.

When we adopt the form of asymptotic conditional coverage in Sesia and Romano (2021), we need to verify the following assumption for $\hat{q}_{\mathrm{MR}}$:

$$\mathbb{P}\left\{\mathbb{E}\left(\hat{q}_{\mathrm{MR}}-q^{*}\right)^{2} \geq \eta_{n}^{2}\right\} \leq 1-\eta_{n}^{2} \tag{S2.13}$$

Actually, we can prove the asymptotic normality of $\hat{q}_{\mathrm{MR}}$ by

$$\sqrt{n}(\hat{q}_{\mathrm{MR}}-q^{*}) \sim N(0, Var(Z)),$$

where Z is the random variable with the form detailed in Han et al. (2019). Therefore, we obtain $\eta_{n}^{2}=Var(Z)/n^{2}$, which satisfies (S2.13), and $\eta_{n}=o_{p}(1)$. Therefore, $b_{n}=O(\eta_{n})$ and $\xi_{n}=O(\eta_{n})$. The proof is completed. $\square$

Proof of Theorem 5. Assume (C1)–(C5) and (C3) hold, and that the cdf of $\tilde{\varepsilon}$ is continuously differentiable with density $f_{\tilde{\varepsilon}}\big(q_{1-\alpha}(\tilde{\varepsilon})\big)>0$. Then, with $m_{\mathrm{cal}}=\sum_{i\in\mathcal{I}_{\mathrm{cal}}} R_{i}$,

$$\sqrt{m_{\mathrm{cal}}}\left(\hat{q}_{\mathrm{MR}}-q_{1-\alpha}(\tilde{\varepsilon})\right) \rightsquigarrow \mathcal{N}\left(0,\ \frac{1-\alpha(1-1-\alpha)}{f_{\tilde{\varepsilon}}\big(q_{1-\alpha}(\tilde{\varepsilon})\big)^{2}}\,\kappa_{\mathrm{MR}}\right),$$

$$\sqrt{m_{\mathrm{cal}}}\left(\hat{q}_{S}-q_{1-\alpha}(\tilde{\varepsilon})\right) \rightsquigarrow \mathcal{N}\left(0,\ \frac{1-\alpha(1-1-\alpha)}{f_{\tilde{\varepsilon}}\big(q_{1-\alpha}(\tilde{\varepsilon})\big)^{2}}\,\kappa_{S}\right),$$

where

$$\kappa_{\mathrm{MR}}=\frac{\mathrm{Var}\big(Z-\Pi_{\mathcal{V}}Z\big)}{\mathrm{Var}(Z)} \quad \text{and} \quad \kappa_{S}=\frac{\mathrm{Var}\big(Z-\Pi_{\mathcal{V}_{S}}Z\big)}{\mathrm{Var}(Z)},$$

$$Z:=1-\alpha-\mathbb{1}\{\tilde{\varepsilon}\leq q_{1-\alpha}(\tilde{\varepsilon})\}.$$

Here $\Pi_{\mathcal{A}}$ denotes the L^{2}-projection onto the linear span $\mathcal{A}$, $\mathcal{V}$ is the span of *all* centered calibration moments used in (2.4)–(2.5), and $\mathcal{V}_{S}$ is the span of the *single* moment used by method S. Consequently,

$$0 \;\le\; \kappa_{\mathrm{MR}} \;\le\; \kappa_S \;\le\; 1, \tag{S2.14}$$

with equality on the left only if the extra moments add no predictive content for Z.

As a corollary for interval length, we have

$$\mathrm{Var}\big(\sqrt{m_{\mathrm{cal}}}\,(L_{\mathrm{MR}} - 2q_{1-\alpha}(\tilde{\varepsilon}))\big) \;\le\; \mathrm{Var}\big(\sqrt{m_{\mathrm{cal}}}\,(L_S - 2q_{1-\alpha}(\tilde{\varepsilon}))\big), \tag{S2.15}$$

and, for every $\epsilon > 0$ and $\delta \in (0,1)$, there exists N such that for all $n \ge N$,

$$\mathbb{P}\big(\,\mathrm{Len}_{\mathrm{MR}} \le \mathrm{Len}_S + \epsilon\big) \;\ge\; 1 - \delta.$$

Thus, *even when $\tilde{\mu}$ (e.g. a linear $\hat{\beta}^\top x$) is misspecified*, the CM–MRL interval is asymptotically no longer than any single-model interval that honors only one of (2.4) and (2.5).

Work on the calibration split, with $\tilde{\mu}$ treated as fixed by sample splitting (C1). Both $\hat{q}_{\mathrm{MR}}$ and $\hat{q}_S$ are solutions to weighted score equations of the form

$$\sum_{i \in S_{\mathrm{cal}}} \tilde{d}_i\, \psi_{1-\alpha}(\tilde{\varepsilon}_i - q) = 0,$$

where $\tilde{d}_i$ are EL-type weights that satisfy either the full set of balancing constraints (2.4) and (2.5) in CM–MRL or just a single constraint in single-model S. A weighted Bahadur representation yields

$$\hat{q} - q_{1-\alpha}(\tilde{\varepsilon}) = \frac{1 - \alpha - \sum_i \tilde{d}_i\, \mathbb{1}\{\tilde{\varepsilon}_i \le q_{1-\alpha}(\tilde{\varepsilon})\}}{f_{\tilde{\varepsilon}}(q_{1-\alpha}(\tilde{\varepsilon}))} + o_p(m_{\mathrm{cal}}^{-1/2}).$$

Empirical-likelihood expansions for $\tilde{d}_i$ imply that the influence function of $\hat{q}$ is proportional to

$$Z - \Pi_{\mathcal{A}} Z, \qquad Z := 1 - \alpha - \mathbb{1}\{\tilde{\varepsilon} \le q_{1-\alpha}(\tilde{\varepsilon})\},$$

where $\mathcal{A}$ is the L^2-span of the centered moments used in the calibration. Hence the asymptotic variance is $\mathrm{Var}(Z - \Pi_{\mathcal{A}} Z)/\{f_{\tilde{\varepsilon}}(q_{1-\alpha}(\tilde{\varepsilon}))\}^2$. Because $\mathcal{V}_S \subseteq \mathcal{V}$, the projection property in a Hilbert space gives

$$\|Z - \Pi_{\mathcal{V}} Z\|_2 \le \|Z - \Pi_{\mathcal{V}_S} Z\|_2,$$

which yields (S2.14). Multiplying by 4 transfers the ordering to interval lengths, giving (S2.15). The probabilistic dominance follows from asymptotic normality and variance ordering by Slutsky's theorem and CLT tail bounds. The argument is unaffected by misspecification of $\tilde{\mu}$, since $q_{1-\alpha}(\tilde{\varepsilon})$ is defined relative to the fixed and possibly incorrect center used by all methods. $\square$

Bibliography

Chernozhukov, V., D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, and J. Robins (2018, 01). Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal 21*(1), C1–C68.

Feldman, S., S. Bates, and Y. Romano (2021). Improving conditional coverage via orthogonal quantile regression. *Advances in neural information processing systems 34*, 2060–2071.

Han, P. (2016). Intrinsic efficiency and multiple robustness in longitudinal studies with drop-out. *Biometrika 103*(3), 683–700.

Han, P., L. Kong, J. Zhao, and X. Zhou (2019). A general framework for quantile estimation with incomplete data. *Journal of the Royal Statistical Society Series B: Statistical Methodology 81*(2), 305–333.

Lei, J., M. G'Sell, A. Rinaldo, R. J. Tibshirani, and L. Wasserman (2018). Distribution-free predictive inference for regression. *Journal of the American Statistical Association 113*(523), 1094–1111.

Qin, J. and J. Lawless (1994). Empirical likelihood and general estimating equations. *the Annals of Statistics 22*(1), 300–325.

Sesia, M. and Y. Romano (2021). Conformal prediction using conditional histograms. *Advances in Neural Information Processing Systems 34*, 6304–6315.

Van der Vaart, A. W. (2000). *Asymptotic statistics*, Volume 3. Cambridge university press.