

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Regularized Targeted Learning under Analytic Intractability: Projection-Guided Targeting and Grouped V-Fold Jackknife Inference

Permalink

<https://escholarship.org/uc/item/5v92x36g>

ISBN

9798189277368

Author

Li, Yi

Publication Date

2026-05-01

Peer reviewed|Thesis/dissertation

Regularized Targeted Learning under Analytic Intractability: Projection-Guided Targeting
and Grouped V-Fold Jackknife Inference

By

Yi Li

A dissertation submitted in partial satisfaction of the
requirements for the degree of

Doctor of Philosophy

in

Biostatistics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Mark van der Laan, Chair

Professor Alan Hubbard

Professor Adityanand Guntuboyina

Spring 2026

Copyright 2026
by
Yi Li

Abstract

Regularized Targeted Learning under Analytic Intractability: Projection-Guided Targeting
and Grouped V-Fold Jackknife Inference

by

Yi Li

Doctor of Philosophy in Biostatistics

University of California, Berkeley

Professor Mark van der Laan, Chair

Modern biostatistical problems increasingly require estimators that are flexible enough to accommodate complex nuisance structure and yet stable enough to support reliable inference under analytic intractability. Targeted minimum loss-based estimation (TMLE), adaptive TMLE (A-TMLE), and related debiasing methods provide a principled semiparametric framework for this task, but their most direct implementations often depend on computational objects that are difficult to use in practice. Exact least favorable fluctuations may be numerically unstable, information operators may be ill-conditioned, and direct influence-curve-based variance formulas may be analytically cumbersome or operationally unavailable once the nuisance fit is highly adaptive, the working model is overparameterized, or the target parameter is functional. This dissertation studies how targeted learning can be carried out under those computational and analytic constraints. Its central claim is that stable semiparametric estimation and reliable inference can often be realized in such settings without requiring a tractable closed-form efficient influence curve, least favorable path, or direct variance formula as an operational input, provided the estimator and the inferential procedure are both organized around the same first-order structure.

The dissertation develops a regularized targeted-learning program for stable estimation and reliable inference in such settings. On the estimation side, it shows that the targeting step can be stabilized by introducing data-adaptive working models and by replacing exact but impractical updates with projection-guided surrogates that retain the relevant first-order bias correction. On the inference side, it develops grouped V-fold jackknife procedures that remain practical when closed-form variance calculations are unreliable and when the full estimation pipeline is too expensive for heavy resampling. The common principle is to preserve first-order semiparametric reasoning while replacing fragile operational objects by computational surrogates that remain compatible with the estimation pipeline.

These ideas are developed in four connected methodological chapters, culminating in bivariate right-censored survival as the dissertation's most demanding setting. The methodological core of the dissertation begins in Chapter 3 with regularized A-TMLE in highly adaptive lasso working models, where the geometry of the targeting problem is especially transparent, and then extends the same score-space logic to neural-network working models.

It next turns to inference through the grouped V-fold jackknife for regular and generalized asymptotically linear estimators, emphasizing both statistical validity and the computational advantage of requiring only V grouped refits of the full estimation procedure; because the resulting Studentized procedure is scale-invariant, it also adapts naturally to non-standard convergence rates. The final methodological chapter studies bivariate right-censored survival under coarsened data, where EM-HAL initial estimation, projection-guided targeting, and grouped V-fold jackknife inference are combined in a setting for which a classical full TMLE is not operationally tractable. Taken together, the chapters present regularized targeted learning as a coherent semiparametric approach to modern estimation and inference under instability, analytic intractability, and computational cost. Empirically, the resulting pipeline improves finite-sample stability, lowers mean squared error, and delivers inference closer to nominal coverage in settings where unregularized targeted updates or plug-in variance calculations become unreliable.

To my family

Preface

This dissertation follows a manuscript-based format. Chapters 3 through 6 are adapted from collaborative manuscripts, while Chapters 1, 2, and 7 were written for the dissertation in order to provide a general introduction, a common semiparametric framework, and a concluding discussion. The adapted chapters have been revised for continuity of notation, bibliography, and appendix structure.

Supplementary proofs, extended empirical specifications, and algorithmic details are consolidated in Appendices A through D.

Chapter 3 is adapted from joint work with Sky Qiu, Zeyi Wang, and Mark van der Laan on regularized A-TMLE in highly adaptive lasso working models. Chapter 4 is adapted from joint work with David McCoy, Nolan Gunter, Kaitlyn Lee, Alejandro Schuler, and Mark van der Laan on targeted deep architectures. Chapter 5 is adapted from joint work with Ashkan Ertefaie and Mark van der Laan on grouped V-fold jackknife methods for semiparametric inference. Chapter 6 is adapted from joint work with Ashkan Ertefaie and Mark van der Laan on targeted learning for bivariate right-censored survival.

Throughout the methodological chapters, the first-person plural is retained to reflect the collaborative origin of the underlying manuscripts.

Unless a distinction is mathematically important, the dissertation uses “efficient influence curve” as the default term, “regularized A-TMLE” as the preferred prose form, and “working model” for the lower-dimensional model used to organize targeting. The dissertation uses a single consolidated bibliography. Responsibility for the dissertation-level exposition, transitions, and any remaining errors rests with the author.

Acknowledgments

Above all, I thank my advisor, Professor Mark van der Laan. No one has shaped this thesis, or my development as a researcher, more profoundly than he has. Meeting with Mark has always been one of the most rewarding parts of my graduate experience. Every conversation pushed me toward a deeper understanding, and I often left his office thinking more clearly than when I arrived. Mark gave me the freedom to pursue many ideas; although only a few survived in their original form, those explorations ultimately made this thesis possible. He was unfailingly generous with his time and never once complained about my pace, even when a direction took months to bear fruit—or quietly did not at all.

Through our discussions, I came to truly appreciate research. From Mark, I learned how to think about problems: how to reframe them, how to probe what they are really about, and how to persist in asking the right questions. His perspective is distinctive and often daring, and he taught me not to shy away from bold—even unconventional—ideas. At the same time, his standards are exacting. He can quickly identify where an argument is weak and requires reconstruction, and he is unflinching about the work that entails. He invested tremendous energy in the projects in this thesis, especially the bivariate survival problem, and when progress was difficult, he remained steadfast. His optimism and perseverance sustained us through long stretches where I might otherwise have given up. I feel extraordinarily fortunate to have been his student and look forward to continuing to work with him on the most challenging and meaningful problems in science.

I am also deeply grateful to Professor Maya Petersen. Working with Maya is a privilege; her leadership, clarity, and insight make even complex problems feel approachable. From her, I learned to be more structured, responsive, and optimistic in my research. The 2020 summer COVID project remains especially memorable, and I thank her for including me. I am also grateful for the early-morning meetings with our Novo collaborators on the data-adaptive target parameter project, which revealed the intricacies of real-world applied work and how to navigate them thoughtfully. She encouraged me to share my ideas directly in those discussions—a trust from which I learned a great deal and continue to benefit.

I owe special thanks to Professor Ashkan Ertefaie, my collaborator on the bivariate survival project. He was remarkably patient as I worked through a genuinely challenging problem and consistently encouraged me to voice ideas, even when they were still developing. He would consider them carefully and return with insights that advanced the work. Without his patience and guidance, this project would have been far more difficult. His mathematical rigor, statistical judgment, and professionalism are qualities I greatly admire and hope to emulate.

I am grateful to Professor Alejandro Schuler, whose lab meetings were a constant reminder of how creative and engaging research can be. He provided me with many early opportunities, including reviewing papers, improving my presentation skills, and learning how to communicate statistical ideas more effectively. He was also among the first to hear many of my initial ideas. I was fortunate to coauthor my first PhD statistics paper with him, and I deeply appreciate the thoughtful and constructive way he has engaged with both my successes and my setbacks. His often visual perspective has enriched my own understanding in lasting ways, and I look forward to many future collaborations.

I would also like to thank Professor Alan Hubbard for his steady support throughout my time at Berkeley and for his consistently warm and encouraging conversations. I am grateful to Professor Adityanand Guntuboyina and Professor James Brown, whose courses have taught me a lot. I am grateful to Professor John Marshall, who gave me the opportunity to work in his lab in the summer of 2018; that experience shaped my understanding of biostatistics, introduced me to the impact of research on human health, and ultimately set me on the path to UC Berkeley. I also thank Professor Jingshen Wang, from whom I learned a great deal about rigorous research practice; her advice on persistence and consistency has been invaluable during challenging moments.

I am grateful to the UC Berkeley Department of Biostatistics and the CTML community. The seminars, courses, discussions, and informal gatherings made these years both intellectually stimulating and genuinely enjoyable. I am especially thankful to my peers—Tianyue Zhou, Junming Shi, David McCoy, Sky Qiu, Hao Wang, Yunwen Ji, Lei Shi, Waverly Wei, Lauren Liao, Marie Charpignon, Haodong Li, Toru Shirakawa, Kirsten Landsiedel, Kaitlyn Lee, Yilong Hou, Nolan Hunter, Alissa Gordon, Zhongming Xie, Junzhe Zhao, Yunzhe Zhou, Yichen Xu, and Abhijith Varma Mudunuri among others—whose support, humor, and companionship made even the most difficult periods manageable. I have valued our conversations about each other’s work, our shared celebrations, and our collective belief that we could tackle problems that initially seemed insurmountable.

I also thank the staff of the Division of Biostatistics and CTML who kept everything running smoothly—Lauren Krupa, Jordan Elmowitz, Jessica Angell, and Christina Da Silva—allowing me to focus on my research with confidence and support.

Finally, and most importantly, I thank my parents, my grandparents, and my two younger sisters. Your love sustained me through the most difficult moments. Everything I have accomplished rests on that foundation, and I carry it with me into whatever comes next.

Contents

Preface	ii
Acknowledgments	iii
1 Introduction	1
1.1 Scientific motivation	1
1.2 Targeted learning under computational and analytic constraints	2
1.3 Contribution of the dissertation	3
1.4 Organization of the dissertation	3
2 Semiparametric Framework and Preliminaries	4
2.1 Semiparametric setting and notation	4
2.2 Targeted learning, regularization, and adaptive working models	5
2.3 Projection parameters and working-model score spaces	6
2.4 Exact objects, computational surrogates, and inference	7
2.5 Terminology and notation conventions	8
3 Stable Regularized A-TMLE in HAL Working Models	10
3.1 Local setup and scope	10
3.1.1 Related work and positioning	11
3.2 Problem Setup and Working-Model Efficient Influence Curve Approximation	11
3.2.1 Hessian-Based Working-Model Efficient Influence Curve Approximation	12
3.2.2 Projection-Based Working-Model Efficient Influence Curve Approximation	13
3.2.3 Theoretical Unification of the Two Approximations	15
3.2.4 regHAL-ATMLE Algorithms	17
3.3 Theoretical Guarantees	19
3.3.1 TMLE for working model induced target parameter	19
3.3.2 TMLE for original target parameter	20
3.3.3 Verification of Regularity Conditions under HAL for efficiency	20
3.3.4 Positioning: Regularized TMLE, A-TMLE, and our contribution	21
3.4 Data-Adaptive Working Model Selection via Plateau Selector	23
3.4.1 Plateau Selector Algorithm on Nested Working Models	24

3.5	regHAL-ATMLE for Average Treatment Effects	25
3.5.1	ATE Set up	25
3.5.2	Simulation Design	25
3.5.3	Results with Fixed Working Models: Standard Conditions (DGP 1) .	26
3.5.4	Results with Fixed Working Models: Positivity Violations (DGP 2) .	28
3.5.5	Results with Plateau-Selected Working Models	29
3.6	regHAL-ATMLE for Marginal Survival Curves	32
3.6.1	Right Censored Data Set up	33
3.6.2	Simulation Results	33
3.7	Discussion	35
3.7.1	Broader perspective	36
3.7.2	Limitations	36
4	Regularized A-TMLE Beyond HAL: Neural Network Working Models	37
4.1	Related work and positioning	37
4.2	Targeting inside neural-network working models	38
4.2.1	The Neural Network as a Parametric Working Model	38
4.2.2	Regularized Projection-Based Targeting	40
4.2.3	Score Space Richness in Neural Network Working Models	42
4.2.4	Multi-Parameter Targeting	44
4.3	Theoretical Guarantees for Targeted Deep Architectures	46
4.3.1	Setup and Working Model Construction	46
4.3.2	Regularized TMLE Theorem	46
4.3.3	A-TMLE Theoretical Results	47
4.3.4	Verification of Regularity Conditions for Neural Networks	48
4.4	Experiments	49
4.4.1	ATE Estimation on IHDP: High-Dimensional Confounding	50
4.4.2	ATE Estimation Under Positivity Violations	51
4.4.3	Survival Curve Estimation: Regularizing Functional Updates	53
4.5	Discussion	56
5	Inference for Targeted Estimators via the V-fold Jackknife	58
5.1	Inference problem and chapter scope	58
5.1.1	Related work and positioning	59
5.2	Grouped V-fold jackknife construction	60
5.3	Properties of V-fold Jackknife	61
5.3.1	Fixed V results	62
5.3.2	Diverging V	64
5.4	Simultaneous V-fold Jackknife Confidence Bands	64
5.4.1	Vector-valued V-fold jackknife	65
5.4.2	Joint distribution of the Studentized vector and simultaneous critical values	66
5.4.3	Monte Carlo computation of the critical value	67

5.4.4	Validity of the eigendecomposition and the role of V in simultaneous inference	68
5.4.5	Remark on links to high-dimensional Gaussian approximation	69
5.4.6	Bias-corrected simultaneous confidence bands	69
5.4.7	Algorithm	70
5.5	Extension to Generalized Asymptotically Linear Estimators and Slower Convergence Rates	70
5.5.1	Fixed- V inference for slower rates	72
5.5.2	Diverging- V consistency for slower rates	72
5.5.3	Generalized simultaneous confidence bands	73
5.5.4	Application: HAL dose-response curve	73
5.6	Simulation Studies	75
5.6.1	Scalar Inference: Average Treatment Effect	75
5.6.2	Survival Curve: Kaplan–Meier Estimator	77
5.6.3	Dose-Response Curve: HAL Estimator	80
5.7	Discussion	88
5.7.1	Practical guidance on choosing V	88
5.7.2	Scope and limitations	89
5.7.3	Connections and extensions	90
5.7.4	Concluding remarks	90
6	Culminating Application: Bivariate Right-Censored Survival	92
6.1	Local problem formulation	92
6.1.1	Position relative to existing work	93
6.2	Problem setup and notation	94
6.2.1	Observed data and the CAR model	94
6.2.2	Target parameters	95
6.3	EM-HAL estimation framework	95
6.3.1	HAL working model for the full-data density	95
6.3.2	The EM algorithm for censored data	95
6.3.3	Plug-in survival and density estimation	98
6.3.4	Computational details	98
6.4	Targeting for survival	98
6.4.1	TMLE, Regularized TMLE, and Regularized A-TMLE	98
6.4.2	Projection-based reg-ATMLE update	99
6.5	Inference via V -fold jackknife	105
6.5.1	Motivation and theoretical validity	105
6.5.2	Variance estimation and confidence intervals	106
6.5.3	Implementation	106
6.5.4	Choice of V	106
6.6	Simultaneous confidence bands	107
6.6.1	Monte Carlo calibration	107
6.6.2	Filtered simultaneous confidence bands	108

6.6.3	Bias correction	108
6.7	Simulation study	108
6.7.1	Data-generating process	109
6.7.2	Summary metrics	109
6.7.3	Density empirical performance (EM-HAL)	110
6.7.4	Empirical performance of survival point estimation	110
6.7.5	Pointwise confidence intervals	112
6.7.6	Simultaneous confidence bands	115
6.8	Discussion	116
7	Conclusion and Future Directions	119
7.1	Principal conclusions	119
7.2	A unified methodological perspective	120
7.3	Limitations	120
7.4	Future directions	120
7.5	Concluding remarks	122
	Bibliography	123
A	Supporting Material for Chapter 3	130
A.1	regHAL-ATMLE Algorithm for Multi Dimensional Target Parameter	131
A.2	regHAL-ATMLE for ATE -Details	132
A.2.1	Estimating the ATE using one-step regHAL-ATMLE	132
A.2.2	Simulation Set Up for ATE	133
A.2.3	Standard TMLE with Propensity Score Truncation	134
A.3	regHAL-ATMLE for Survival Curve - Details	134
A.3.1	Estimating the survival curve using one-step regHAL-ATMLE	135
A.3.2	Simulation Set Up for Right-Censored Survival Data	138
A.3.3	Simulation Summary Statistics	138
A.4	Heteroscedasticity and Non-Likelihood Losses: Detailed Development	139
A.5	Proofs	139
A.5.1	Proof of Theorem 3.8 (Regularized TMLE)	140
A.5.2	Theoretical Framework for Hessian-Method: A-TMLE for Projection Parameters	140
A.5.3	Verification of Regularity Conditions under HAL	141
A.6	Alternative plateau selection via <i>bridged</i> targeting	144
A.6.1	Motivation	144
A.6.2	Algorithm	145
A.6.3	Discussion	145
B	Supporting Material for Chapter 4	146
B.1	Appendix: Proofs of Theoretical Results	146
B.1.1	Proof of Theorem 4.6 (A-TMLE for Working Parameter)	146
B.1.2	Proof of Theorem 4.7 (A-TMLE for Original Parameter)	147

B.1.3	Proof of Theorem 4.8 (Verification of Regularity Conditions)	148
B.1.4	Proof of Theorem 4.4 (Regularized TMLE)	149
B.2	Appendix: Experimental Details	150
B.2.1	ATE Experiments: IHDP and Positivity Violations	150
B.2.2	Survival Curve Experiments	152
B.3	TDA Extensions	156
B.3.1	Nested Working Models and Plateau Selection in TDA	156
B.3.2	TDA-Direct: A Streamlined Targeting Step for Special Cases (e.g., the ATE)	159
C	Supporting Material for Chapter 5	161
C.1	Proofs for the main grouped-jackknife results	161
C.1.1	Proof of Lemma 5.1	161
C.1.2	Proof of Theorem 5.2	162
C.1.3	Proof of Lemma 5.3	162
C.1.4	Proof of Corollary 5.4	162
C.1.5	Proof of Theorem 5.6	163
C.1.6	Proof of Theorem 5.7	163
C.1.7	Proof of Theorem 5.8	164
C.1.8	Proof of Theorem 5.9	164
C.1.9	Proof of Theorem 5.12	164
C.1.10	Proof of Corollary 5.13	165
C.1.11	Proof of Theorem 5.14	165
C.2	Connection to High-Dimensional Gaussian Approximation	165
C.3	True Dose-Response Curves for the HAL Simulation	167
C.4	Pointwise Coverage of Simultaneous Confidence Bands	167
C.5	Jackknife-Centered versus Bias-Corrected Confidence Intervals	171
C.6	Diagnostic: Highlighting Unreliable Grid Points in a Simultaneous Band	174
C.7	Isolating the Effect of Correlation Matrix Estimation	175
C.8	Derivation of the Growth-Rate Condition on V	176
D	Supporting Material for Chapter 6	179
D.1	Implementation details for the EM-HAL and targeting pipeline	179
D.2	Pointwise Asymptotic Normality of the EM-HAL Density Sieve MLE	180
D.3	Self-consistency targeting via regularized TMLE	183
D.3.1	Self-consistency equations for survival	183
D.3.2	Indicator-based offset targeting submodel	183
D.3.3	Solving the finite self-consistency system via EM	183
D.3.4	Regularized TMLE theory for the finite self-consistency system	184
D.4	Additional survival point-estimation diagnostics	187
D.5	Additional survival confidence-interval diagnostics	188
D.6	Density estimation under survival targeting	189
D.6.1	Density performance after survival targeting	189

D.6.2	Density V-fold jackknife inference	193
D.6.3	Density simultaneous confidence bands	193
D.7	Proofs of Theoretical Results	194

List of Figures

5.1	Coverage (top) and CI width (bottom) of V -fold jackknife confidence intervals as a function of the number of folds V , for four scenarios and two sample sizes. Horizontal dashed line in the top panels: nominal 95% level. Faded horizontal lines: coverage of the EIC-based and bootstrap methods for each scenario. The V -fold jackknife maintains coverage above the nominal level for moderate positivity and substantially outperforms the EIC-based method under severe positivity, with intervals narrowing as V increases.	78
5.2	Pointwise coverage of 95% confidence intervals for the Kaplan–Meier survival function across evaluation times ($n = 200, 500$ replications). The V -fold jackknife with t_9 critical values ($V = 10$) closely tracks the Greenwood formula, and the bootstrap Wald method performs comparably.	79
5.3	Pointwise coverage of 95% confidence intervals across the dose grid (20 equally spaced points on $[0, 5]$, i.e. $A \in \{0, 0.263, 0.526, \dots, 5\}$) for HAL dose-response curve estimation ($n = 500, 500$ replications). Each point on the horizontal axis represents a dose level at which the dose-response curve $\mathbb{E}[Y \mid A = a, W]$ is evaluated and coverage is computed separately. Rows: DGP 1 (smooth), DGP 2 (oscillating), DGP 3 (piecewise), DGP 4 (discontinuous). Columns: CV-HAL and undersmoothed HAL. The V -fold jackknife with t_{19} critical values ($V = 20$) provides the highest coverage across all DGPs. The delta method fails most notably for DGP 2 (oscillating), where its failure rate also exceeds 5% for undersmoothed HAL.	82
5.4	Bias diagnosis for HAL dose-response curve estimation ($n = 500, 500$ replications). Solid black line: absolute bias $ \hat{\psi}(a) - \psi_0(a) $ at each of the 20 evaluation points, where $\hat{\psi}(a)$ is the simulation average of the point estimate. Dashed blue line: mean estimated standard error from the V -fold jackknife ($V = 20$). When the bias exceeds the standard error, the asymptotic linearity condition is violated and coverage loss is expected. The spatial pattern of bias peaks corresponds to the coverage dips in Figure 5.3: for DGP 3, bias rises for $A > 2$ where oscillations begin; for DGP 4, bias spikes at $A = 2$ and $A = 4$ where the true curve is discontinuous.	84

5.5	Eigenvalue spectrum of the estimated correlation matrix R_0 for the Kaplan–Meier survival curve ($m = 91$ time points) and the HAL dose-response curve ($m = 20$ evaluation points, four DGPs, CV-HAL and undersmoothed variants). The correlation matrix is estimated from the $500 \times m$ matrix of point estimates across simulation replications. (a) Normalized eigenvalue spectrum (fraction of trace m), truncated at index 20 for visibility. The KM curve exhibits the steepest decay, with a single dominant eigenvalue accounting for 59% of the total variance. (b) Cumulative fraction of variance explained. For the KM curve, two eigenvalues suffice for 77% and nine for 95%; for the HAL settings, 6–16 eigenvalues are needed for 95%. In all cases, the effective rank $r(R_0) = m/\lambda_{\max}$ is much smaller than m , confirming that the eigendecomposition of the rank-deficient $\hat{R}_{\text{Jack}}$ accurately captures the dominant correlation structure even when $V \ll m$	87
6.1	EM-HAL density estimation: true bivariate density (left), mean estimate across $R = 213$ replications (center), and pointwise difference (right).	110
6.2	reg-ATMLE: MSE heatmaps for EM-HAL (left), targeted (center), and $\log_{10}$ ratio (right). Blue indicates improvement.	111
6.3	reg-ATMLE: coverage heatmaps comparing CI methods for the targeted survival estimate.	113
B.1	Median survival curves with 95% percentile bands under the concordant late acceleration DGP. Solid curves show the median survival estimate across 500 replications for the initial model (blue), TDA (green), and Kaplan–Meier (cyan). Shaded ribbons depict the 2.5th–97.5th percentiles. The black line is the true survival function. The KM estimator shows severe systematic bias due to concordant informative censoring.	156
B.2	How often does targeting help? Percentage of 500 replications in which TDA has a smaller absolute error than the initial model (blue) and the KM estimator (teal) at each time point. The dashed line indicates parity (50%); values above mean TDA is more accurate in the majority of runs.	157
C.1	True dose-response curves $\psi_0(a) = \mathbb{E}[Y \mid \text{do}(A = a)]$ for the four data generating processes used in the HAL simulation study. DGP 1: smooth monotone increase; DGP 2: oscillating; DGP 3: smooth for $a \leq 2$, oscillating for $a > 2$; DGP 4: jump discontinuities at $a = 2$ and $a = 4$. True values computed by Monte Carlo integration (10^6 samples per evaluation point).	168
C.2	Pointwise coverage (%) of simultaneous 95% confidence bands for the Kaplan–Meier survival function across evaluation times ($n = 200, 500$ replications). Methods: V -fold simultaneous band with $V \in \{5, 10, 20\}$ and the studentized bootstrap simultaneous band. Horizontal dashed line: nominal 95% level. All methods achieve 98–100% pointwise coverage uniformly across time points, with no systematic differences across V	169

C.3	Pointwise coverage (%) of simultaneous 95% confidence bands for the HAL dose-response curve across the dose grid ($n = 500$, 500 replications). Rows: DGP 1 (smooth), DGP 2 (oscillating), DGP 3 (piecewise), DGP 4 (discontinuous). Columns: CV-HAL and undersmoothed HAL. Solid lines: standard simultaneous bands; dashed lines: bias-corrected simultaneous bands. Colors indicate $V \in \{5, 10, 20\}$. Horizontal dashed line: nominal 95% level.	170
C.4	Pointwise coverage (%) of 95% confidence intervals for the HAL dose-response curve, comparing three centering strategies ($V = 20$, $n = 500$, 500 replications). Standard: centered at $\hat{\Psi}_j(P_n)$; Jack-centered (JC): centered at $\hat{\Psi}_{Jack}^{(j)}(P_n) = \hat{\Psi}_j(P_n) - \hat{b}_j$; Bias-corrected (BC): centered at $\hat{\Psi}_j(P_n)$ with interval widened by $ \hat{b}_j $. The JC intervals have lower coverage than the standard intervals in all settings, while BC achieves the highest coverage.	172
C.5	Pointwise coverage (%) of simultaneous 95% confidence bands for the HAL dose-response curve, comparing three centering strategies ($V = 20$, $n = 500$, 500 replications). Same conventions as Figure C.4 but using the simultaneous critical value $\hat{q}_\alpha$ in place of $t_{V-1, 0.975}$. The coverage degradation from JC centering is even more pronounced for simultaneous bands.	173
D.1	Self-consistency targeting: $ \text{Bias} /\text{SD}$ heatmaps for EM-HAL (left), targeted (center), and the change (right). Blue indicates improvement.	187
D.2	Self-consistency targeting: MSE heatmaps for EM-HAL (left), targeted (center), and $\log_{10}$ ratio (right). Blue indicates improvement.	187
D.3	Self-consistency targeting: oracle-coverage heatmaps (95% nominal) for the EM-HAL (left) and targeted (right) stages.	188
D.4	reg-ATMLE: $ \text{Bias} /\text{SD}$ heatmaps for EM-HAL (left), targeted (center), and the change (right). Blue indicates improvement.	188
D.5	reg-ATMLE: oracle-coverage heatmaps (95% nominal) for the EM-HAL (left) and targeted (right) stages.	189
D.6	Self-consistency targeting: coverage heatmaps comparing CI methods for the targeted survival estimate.	190
D.7	Survival coverage and CI length: central $[0.25, 0.75]^2$ vs. tail region for self-consistency targeting (top) and reg-ATMLE (bottom). Metrics are averaged per evaluation point across 213 simulations.	191
D.8	Survival coverage and CI length: 9×9 targeting grid (81 points) vs. non-target points for self-consistency targeting (top) and reg-ATMLE (bottom). Metrics are averaged per evaluation point across 213 simulations.	192

List of Tables

2.1	Shared notation used across the dissertation.	5
2.2	How the shared framework is instantiated in the later chapters.	9
3.1	Simulation results for regHAL-ATMLE under standard conditions (DGP 1) with CV-selected working model.	27
3.2	Simulation results for regHAL-ATMLE under standard conditions (DGP 1) with undersmoothed working model.	27
3.3	Simulation results for regHAL-ATMLE under practical positivity violations (DGP 2) with CV-selected working model.	28
3.4	Simulation results for regHAL-ATMLE under practical positivity violations (DGP 2) with undersmoothed working model.	29
3.5	Simulation results for regHAL-ATMLE with plateau-selected working model complexity under standard conditions (DGP 1). A-TMLE methods adaptively select j via the plateau selector; CV regHAL and US regHAL use fixed CV-selected and undersmoothed working models, respectively, both with projection-based targeting.	31
3.6	Simulation results for regHAL-ATMLE with plateau-selected working model complexity under practical positivity violations (DGP 2).	32
3.7	Simulation results for intensity-based regHAL-ATMLE for marginal survival curve estimation with CV-selected working model. Kaplan-Meier is included as the efficient nonparametric benchmark.	34
3.8	Simulation results for intensity-based regHAL-ATMLE for marginal survival curve estimation with undersmoothed working model. Kaplan-Meier is included as the efficient nonparametric benchmark.	35
4.1	ATE estimation on IHDP (500 replications). Estimators are grouped by targeting strategy. Metrics: Bias, $ \text{Bias} /\text{SD}$ (centering quality), SD (Monte Carlo standard deviation), RMSE ($\sqrt{\text{Bias}^2 + \text{SD}^2}$), Mean SE (average estimated standard error), Nom Cov (95% CI coverage using estimated SE), Orc Cov (coverage using true MC SD), $E P_n D^* $ (mean absolute EIC residual).	52
4.2	ATE estimation under positivity violations ($n = 500$, 500 replications). True ATE = 1.5.	54
4.3	ATE estimation under positivity violations ($n = 1000$, 500 replications). True ATE = 1.5.	54

4.4	Survival curve estimation under concordant informative censoring (500 replications, $n = 1,000$). Each metric is computed at each of 50 time points; we report mean $\pm$ SD across time points. Bold indicates best performance. . . .	56
5.1	Coverage (%) and mean CI width for the average treatment effect via TMLE. Scenarios A–D have the propensity score model g correctly specified. A: Q correct, moderate positivity; B: Q correct, severe positivity; C: Q misspecified, moderate positivity; D: Q misspecified, severe positivity. 500 replications. . .	76
5.2	Pointwise coverage (%) and mean CI width for the Kaplan–Meier survival function ($n = 200$, $m = 91$ evaluation points, 500 replications, averaged over all evaluation times).	77
5.3	Simultaneous coverage (%) and mean band width for the Kaplan–Meier survival function ($n = 200$, $m = 91$ evaluation points, 500 replications). “BC” denotes the bias-corrected band from section 5.4.6.	80
5.4	Pointwise coverage (%), mean CI width, and delta method failure rate for HAL dose-response curve estimation ($n = 500$, 500 replications, averaged over 20 evaluation points).	83
5.5	Simultaneous coverage (%) and mean band width for HAL dose-response curve estimation ($n = 500$, 500 replications, 20 evaluation points). “BC” denotes the bias-corrected band from section 5.4.6.	85
6.1	Simulation parameters.	109
6.2	Density empirical performance: EM-HAL estimate on the interior of the 209×209 evaluation grid, excluding boundary points ($R = 213$, $n = 500$).	110
6.3	Survival empirical performance of self-consistency targeting and the reg-ATMLE update. Metrics are averaged over the 209×209 evaluation grid ($R = 213$, $n = 500$).	111
6.4	Survival CI comparison ($R = 213$, $V = 20$). Coverage and mean CI length are averaged over the 209×209 evaluation grid.	112
6.5	Oracle vs. jackknife CI comparison for targeted estimates ($R = 213$, $V = 20$). Averaged over the 209×209 evaluation grid.	114
6.6	Simultaneous coverage on the 9×9 targeting grid (81 points, $R = 213$). . . .	115
6.7	Partial simultaneous coverage (bias-corrected): fraction of simulations covering $\geq k\%$ of 81 points.	115
6.8	Filtered simultaneous coverage for survival on the targeting grid, using the jackknife bias proxy ($R = 213$).	116
6.9	Head-to-head comparison of self-consistency targeting and reg-ATMLE ($R = 213$). Pointwise metrics are averaged over the 209×209 evaluation grid; simultaneous metrics are computed on the 9×9 targeting grid.	117
B.1	TTR hyperparameter grid search for survival (50 replications each, last-layer targeting). The initial model has $ \text{Bias} /\text{SD} \approx 0.17$. Conservative settings produce coin-flip results; aggressive settings cause numerical divergence. . . .	154

B.2	Plateau selection across nested outcome-head subsets (IHDP, 500 replications). Bias : absolute bias; SD: Monte Carlo standard deviation; SE: average estimated standard error; MSE: mean squared error; Nom Cov: 95% CI coverage using estimated SE; Orc Cov: oracle coverage using true MC SD; $P_n D_{\text{NP}}^*$: non-parametric EIC mean; $P_n D_{\text{proj}}^*$: projected EIC mean.	159
C.1	Decomposing the effect of V on simultaneous coverage for the Kaplan–Meier survival function ($n = 200$, $m = 91$, 500 replications). “Standard” rows use V folds for all components (correlation matrix $\hat{R}_{\text{Jack}}$, degrees of freedom $\nu = V - 1$, and standard errors). “Isolation” rows fix $\nu = 19$ and SE from $V = 20$, varying only the correlation matrix source $\hat{R}_{\text{Jack}}$ estimated from V_{corr} folds.	176
D.1	Density empirical performance before and after survival targeting ($R = 213$, $n = 500$). Metrics averaged over the 209×209 evaluation grid.	189
D.2	Density jackknife CI performance ($R = 213$, $V = 20$). Averaged over the 209×209 evaluation grid.	193
D.3	Density simultaneous coverage on the 9×9 targeting grid (81 points, $R = 213$).194	

Chapter 1

Introduction

1.1 Scientific motivation

Modern biostatistics increasingly depends on estimators that must do two things at once. They must adapt to complex nuisance structure arising from rich covariate information, censoring, missingness, and nonlinear response surfaces, and they must still support valid scientific inference after that flexibility has been used. Semiparametric efficiency theory provides the natural benchmark for this task [Bickel et al., 1993, van der Vaart, 1998]. It identifies the first-order object that should govern estimation and inference and supplies an ideal notion of efficiency against which practical procedures may be judged.

Targeted minimum loss-based estimation (TMLE), adaptive TMLE (A-TMLE), and related debiasing methods were developed to translate that benchmark into practice [van der Laan and Rubin, 2006, Van der Laan et al., 2011, van der Laan and Rose, 2018, Chernozhukov et al., 2018]. Their appeal is precisely that they combine flexible nuisance estimation with explicit control of the first-order bias that matters for the target parameter. This combination is especially important in causal inference and survival analysis, where scientific conclusions often depend on nuisance estimators that are too complex to be handled adequately by rigid parametric models.

These methodological pressures arise from concrete biomedical problems. Comparative effectiveness studies built from high-dimensional electronic health records must adjust for rich patterns of comorbidity, medication history, laboratory measurements, and utilization, making flexible nuisance learning essential for confounding control. Precision-oncology and genomic studies create similar demands when large molecular feature sets are used to model treatment response or prognosis. The bivariate survival problems studied later in the dissertation are also clinically motivated: correlated event times arise naturally in paired-organ or paired-eye studies, in disease-concordance analyses, and in other settings where two survival outcomes must be modeled jointly rather than through independent margins.

These pressures are most visible in the applied settings this dissertation targets. In treatment-effect inference under near-positivity violations, the relevant clever covariate becomes erratic in parts of the covariate space where overlap is weakest. In survival analysis under coarsened data, a literal full TMLE for an entire bivariate survival surface is not

operationally tractable. And in settings in which the nuisance fit is chosen from flexible machine-learning procedures, the information operator used for variance calculation may be ill-conditioned. In each case the theory still identifies the right first-order object, but direct implementation is unstable, analytically difficult, or computationally expensive.

This dissertation studies how targeted learning should be organized once that gap appears. The difficulty is not whether the ideal efficient object exists in principle, but how to turn it into an estimator and an inferential procedure that remain stable and computationally feasible in practice.

1.2 Targeted learning under computational and analytic constraints

The dissertation is concerned with this gap between theoretical availability and operational tractability, which I refer to here as analytic intractability. In the settings studied here, at least one component of the ideal construction becomes difficult to use directly: the efficient influence curve may not have a convenient closed form, the corresponding least favorable path may be cumbersome to construct, the targeted update may be numerically unstable in a data-adaptive working model, or the variance layer may be too delicate or too expensive to evaluate repeatedly inside a full estimation pipeline. These are not isolated technical inconveniences. They arise naturally in treatment-effect estimation under near-positivity violations, in functional survival targets under coarsening, and in adaptive procedures whose score structure is high-dimensional or overparameterized.

The methodological perspective developed here is to distinguish the exact semiparametric object from the *computational surrogate* that makes estimation and inference possible. That distinction preserves the logic of targeted learning without insisting that every component of the ideal construction be available in a tractable closed form. On the estimator side, the dissertation develops projection-guided regularized targeting inside adaptive working models, so that the bias-relevant direction can be represented through a controlled score-space approximation. On the inferential side, it develops grouped V-fold jackknife procedures that remain practical when direct variance calculations are fragile and when bootstrap-style recomputation is too costly; the method is especially attractive because it requires only V grouped refits of the full estimator pipeline.

The central claim of the dissertation is therefore methodological rather than chapter-specific. Under analytic intractability, stable semiparametric point estimation can often be achieved by regularizing the targeting step inside data-adaptive working models, and reliable inference can often be paired with those estimators without requiring a tractable closed-form analytic efficient influence curve, least favorable path, or direct variance formula as an operational input. The point is not to abandon influence-curve reasoning, but to use it in a form that remains workable in the settings where modern data analysis actually creates the most pressure.

1.3 Contribution of the dissertation

This claim is developed through four connected methodological studies. Chapter 3 begins in highly adaptive lasso (HAL) working models, where the geometry of regularized targeting can be stated most explicitly. There the main contribution is to make adaptive targeted learning operational by replacing an unstable exact update with a projection-based approximation to the working-model efficient influence curve. Chapter 4 extends the same idea to neural-network working models, showing that the estimator-side construction is tied to score-space geometry rather than to HAL alone. Chapter 5 develops the inferential complement through grouped V-fold jackknife methods for regular and generalized asymptotically linear estimators. Chapter 6 serves as the culminating synthesis chapter, bringing the estimation and inference programs together in bivariate right-censored survival analysis, where coarsened data, functional targets, and an expensive estimation pipeline make both exact targeting and exact variance calculations difficult to use directly.

The chapters are connected by a common sequence of questions: which working model or score space organizes the relevant bias correction, which exact semiparametric object becomes unstable or analytically difficult in that setting, which projected or regularized surrogate can replace it without losing the first-order logic of the method, and which inferential procedure remains both valid and computationally feasible once that surrogate is adopted. Seen from this perspective, the dissertation contributes not only a set of estimators and asymptotic results, but also a general way of organizing targeted learning when direct implementation of the ideal theory is impractical.

1.4 Organization of the dissertation

Chapter 2 introduces the semiparametric notation and conceptual framework used throughout the dissertation. It also fixes the terminology for working models, projection parameters, efficient influence curves, and computational surrogates that recurs in the later methodological chapters. Chapter 3 then studies regularized A-TMLE in HAL working models and provides the foundational estimator-side analysis. Chapter 4 broadens that analysis to neural-network working models and clarifies the extent to which projection-guided targeting depends on score-space structure rather than on a specific model class. Chapter 5 turns from estimation to inference and develops the grouped V-fold jackknife for pointwise and simultaneous uncertainty quantification. Chapter 6 applies the full program to bivariate right-censored survival under coarsened data. Chapter 7 concludes with the principal methodological lessons of the dissertation, its limitations, and the directions that follow most naturally from the work.

The next chapter establishes the semiparametric framework that underlies the rest of the dissertation. Its purpose is to make precise the common objects that recur across the later chapters and to clarify how the theoretical and computational parts of the argument are related.

Chapter 2

Semiparametric Framework and Preliminaries

The methodological chapters that follow all concern semiparametric estimation, but they do so in different statistical settings and under different computational pressures. The purpose of this chapter is to record the common framework that links them. It introduces the basic observed-data notation used throughout the dissertation and explains how exact semiparametric objects are related to the computational surrogates used in the algorithms.

Two ideas are central. First, the later chapters are all driven by first-order semiparametric structure: target parameters, efficient influence curves, and asymptotic linear expansions. Second, those chapters differ because the exact objects identified by the theory are not always available in an operational form that is stable, explicit, or inexpensive to use. The recurring task is therefore to retain the relevant first-order logic while replacing fragile or intractable constructions by computational surrogates that can actually be implemented.

2.1 Semiparametric setting and notation

Throughout the dissertation, the observed-data unit is denoted by $O \sim P_0$, where P_0 is the true distribution in a statistical model $\mathcal{M}$. The parameter of interest is written as $\Psi(P)$, often through a representation $\Psi^F(Q(P))$ in which $Q(P)$ denotes the nuisance component entering the target parameter directly and $g(P)$ denotes additional nuisance structure entering the efficient influence curve. This notation is broad enough to cover the treatment-effect, survival, and density-based problems that appear later.

Semiparametric efficiency theory supplies the common language in which those problems are analyzed: tangent spaces, efficient influence curves, pathwise differentiability, asymptotic linearity, and efficiency bounds [Bickel et al., 1993, van der Vaart, 1998]. For the purposes of this dissertation, its importance is not only asymptotic. It identifies the first-order correction that a targeted estimator should capture and thereby determines the object against which a computational approximation must be judged.

The notation used most often is summarized in Table 2.1. Fixing these conventions once

at the outset makes it easier to read the later chapters as parts of a single methodological development rather than as isolated applications of similar ideas.

Table 2.1: Shared notation used across the dissertation.

Notation	Meaning
O	Observed-data unit.
P_0	True observed-data distribution.
$\mathcal{M}$	Full statistical model.
$Q(P)$	Nuisance component directly entering the target parameter.
$g(P)$	Additional nuisance structure entering the efficient influence curve.
$\Psi(P)$	Target parameter on the observed-data model.
$\Psi^F(Q)$	Functional form of the target when expressed through Q .
D_P^*	Efficient influence curve of Ψ at P .
$\mathcal{M}_n$	Data-adaptive working model used for targeting.
$\Pi_{\mathcal{M}_n}$	Loss-based projection onto the working model.
$\Psi_{\mathcal{M}_n}(P)$	Projection parameter induced by the working model.
$D_{\mathcal{M}_n, P}^*$	Efficient influence curve of $\Psi_{\mathcal{M}_n}$ at P (the “working-model EIC”).
$\widetilde{D}_{\mathcal{M}_n, P}^*$	Regularized or projected approximation to $D_{\mathcal{M}_n, P}^*$.

2.2 Targeted learning, regularization, and adaptive working models

The dissertation sits at the intersection of several related literatures. Classical TMLE provides a substitution-based route from flexible nuisance estimation to efficient inference [van der Laan and Rubin, 2006, Van der Laan et al., 2011, van der Laan and Rose, 2018]. Closely related debiasing methods, including double machine learning and automated orthogonalization procedures, emphasize the same first-order bias correction in high-dimensional settings [Chernozhukov et al., 2018, 2022b]. The present work is most directly connected to the line of research that studies adaptive targeted learning, in which the working model used for targeting is itself selected from the data [van der Laan et al., 2023, 2026].

Within this perspective, classical TMLE updates an initial estimator along a least favorable path so that the empirical mean of the efficient influence curve is solved to second order. Regularized TMLE retains the same target parameter but replaces an exact fluctuation or exact score equation by a controlled approximation when the direct update is unstable or unavailable in practice. Adaptive TMLE introduces a working-model layer, so that the targeting step is carried out for a projection parameter defined by a lower-dimensional model selected from the data. Regularized A-TMLE combines these ideas by stabilizing the approximation to the efficient influence curve inside that working model.

These distinctions are central to the later chapters. In Chapter 3, the working model is a HAL sieve and regularization enters through an approximation to the working-model efficient

influence curve. In Chapter 4, the same logic is transported to neural-network score spaces. In Chapter 6, the targeted update is built in a coarsened-data problem for which a direct classical full TMLE is not operationally tractable. Chapter 5 addresses the corresponding inferential question by replacing fragile direct variance calculations with grouped V-fold jackknife procedures.

2.3 Projection parameters and working-model score spaces

The full model $\mathcal{M}$ is typically nonparametric or semiparametric. A working model $\mathcal{M}_n \subseteq \mathcal{M}$ is a lower-dimensional, data-adaptive model used to organize the targeting step. The working-model projection parameter is

$$\Psi_{\mathcal{M}_n}(P) = \Psi^F(\Pi_{\mathcal{M}_n}Q(P)),$$

where $\Pi_{\mathcal{M}_n}$ denotes the loss-based projection of the relevant nuisance component onto the working model. This object matters because adaptive targeting is carried out for the projection parameter, while the dissertation’s efficiency arguments, adapted from A-TMLE [van der Laan et al., 2023, 2026], show when the projection-parameter estimator is still first-order valid for the original parameter.

The role of the working model differs across chapters. In Chapter 3 it is a HAL sieve determined by basis selection and sectional variation control. In Chapter 4 it is a neural-network submodel determined by a chosen set of unfrozen weights. In Chapter 6 it is a larger HAL targeting model built on top of an EM-HAL initial estimator. Despite those differences, the working model serves the same purpose throughout: it identifies the score space in which the first-order bias correction will be approximated and regularized.

Anchoring example: ATE under positivity violations. To make the projection-parameter language concrete before it recurs, consider the average treatment effect $\Psi(P) = \mathbb{E}_P[\mathbb{E}_P(Y \mid A = 1, W) - \mathbb{E}_P(Y \mid A = 0, W)]$, with binary treatment A and covariates W . Here $Q(P)(a, w) = \mathbb{E}_P(Y \mid A = a, W = w)$ and $g(P)(w) = P(A = 1 \mid W = w)$. Near-positivity violations drive $g(P)(w)$ close to zero or one in parts of the covariate space, destabilizing the standard efficient-influence-curve evaluation. A working model $\mathcal{M}_n$ that adaptively restricts where the target is identified—for example, through a HAL sieve or a neural-network offset with a chosen subset of unfrozen weights—defines the projection parameter $\Psi_{\mathcal{M}_n}(P)$ on which targeting is actually performed. Chapter 3 returns to this example analytically; Chapter 4 returns to it empirically; Chapter 6 extends the same logic to coarsened survival data.

2.4 Exact objects, computational surrogates, and inference

The later chapters repeatedly distinguish between exact semiparametric objects and the computational constructions used in practice. Four layers are worth separating explicitly:

- **Exact theoretical object.** The efficient influence curve D_P^* of the target parameter, together with a least favorable path when such a path is tractable.
- **Valid influence-curve representation.** A first-order correction direction that identifies the bias-relevant direction even when the exact object is not conveniently available in closed form—for example, an IPCW-based or other valid influence-curve representation.
- **Computational surrogate.** The construction actually used by the algorithm, such as a projected score-space approximation to the working-model EIC or a regularized updating equation.
- **Inferential procedure.** The uncertainty-quantification layer paired with the surrogate, which must remain valid under the same computational constraints rather than relying on a delicate closed-form variance calculation.

This distinction is essential to the methodological program of the dissertation. In Chapter 3, the exact working-model efficient influence curve is replaced by a regularized score-space approximation. In Chapter 4, a valid influence-curve representation is projected into the gradient span of selected neural-network parameters. In Chapter 6, a full exact efficient influence curve and least favorable path for the survival surface are not used as operational inputs, yet valid IPCW and influence-curve representations still identify the bias-relevant direction to be projected into a larger HAL score space.

The inferential layer must respect the same distinction. Asymptotic linearity determines the first-order basis for uncertainty quantification, but it does not require that inference proceed exclusively through a delicate closed-form variance calculation. The grouped V-fold jackknife developed in Chapter 5 is designed precisely for that setting: it uses the first-order expansion while requiring only V grouped refits of the full procedure, making it attractive when each estimator evaluation is itself expensive.

Asymptotic linearity. Most estimators studied in this dissertation admit an asymptotically linear expansion

$$\hat{\Psi}_n - \Psi(P_0) = (P_n - P_0) D_{P_0}^* + o_P(n^{-1/2}),$$

with $D_{P_0}^*$ the efficient influence curve of Ψ at P_0 and $(P_n - P_0)$ the centered empirical measure. Chapter 5 also considers the generalized form in which the influence-function standard deviation σ_n may grow with the sample size, so that the remainder is $o_P(\sigma_n/\sqrt{n})$. This is the setting that arises once the estimator is regularized away from an exact $\sqrt{n}$ -rate ideal, as in the HAL dose-response estimator of Chapter 5. Conceptually, such diverging scales arise when the target is functional and the effective complexity of the approximation

grows with the sample size, so that the first-order stochastic term accumulates over an expanding score or basis representation.

2.5 Terminology and notation conventions

Several terminology choices are fixed here so that the remaining chapters may use them consistently. “Efficient influence curve” is the default term, while “canonical gradient” appears only when mirroring the language of a cited source. “Regularized A-TMLE” is the preferred prose form, with “reg-A-TMLE” used as a shortened label. “Working-model EIC” denotes the efficient influence curve $D_{\mathcal{M}_n, P}^*$ of the projection parameter within a working model. When a tractable exact efficient influence curve or least favorable path is not available, “influence-curve representation” refers to any valid first-order object, such as an IPCW gradient, that identifies the bias-relevant direction to be projected into the working-model score space. Chapter 5, however, treats general regular and generalized asymptotically linear estimators rather than only efficient targeted estimators, and therefore uses φ for a generic influence function; that notation does not imply efficiency and is distinct from the use of D^* in the targeted-estimation chapters.

For reference, Table 2.2 summarizes how the common framework is instantiated in the remaining methodological chapters.

With this framework in place, the dissertation turns next to the most structured setting it studies: regularized adaptive targeting inside HAL working models. Because HAL provides explicit basis functions and nested working models, it is the natural place to begin the estimator-side development.

Table 2.2: How the shared framework is instantiated in the later chapters.

Chapter	Working model	Exact object and computational surrogate	Inferential emphasis
Chapter 3	HAL sieve working model	Exact working-model EIC replaced by a regularized score-space surrogate, with projection-based targeting as the main line	Asymptotic linearity of the targeted estimator and data-adaptive working-model selection
Chapter 4	Neural-network submodel defined by the targeting weights	Bias-relevant influence-function direction projected onto the gradient score span	Generalization of regularized targeting beyond HAL
Chapter 5	Regular and generalized asymptotically linear estimators rather than a single working-model family	Influence-function expansion treated as the theoretical input; grouped pseudo-values replace closed-form variance formulas	Grouped V-fold jackknife for pointwise and simultaneous inference with only V grouped refits
Chapter 6	Larger HAL targeting model layered on top of EM-HAL initial estimation	Valid IPCW and influence-curve representations projected into a larger HAL score span when a full exact EIC and least favorable path are not tractably available	Jackknife pointwise and simultaneous inference under coarsened data with only V grouped refits of the expensive pipeline

Chapter 3

Stable Regularized A-TMLE in HAL Working Models

Highly adaptive lasso working models provide the most transparent setting in which to study regularized adaptive targeting. HAL supplies explicit basis functions, explicit score vectors, and a nested working-model structure, so it makes clear both why adaptive targeting is attractive and why it can become unstable. Once the active score system is highly collinear, an exact update that solves the full working-model efficient-influence-curve equation directly can be numerically fragile even when the working model is statistically well motivated.

The main contribution of the chapter is to replace that unstable exact update by a controlled approximation to the working-model efficient influence curve $D_{\mathcal{M}_n, P}^*$. Projection-based targeting is the primary construction, while Hessian-based targeting is retained as a secondary comparison that helps clarify why the bias-relevant score direction, rather than the full active score system, should determine the update. In this sense, the HAL setting is not simply a first application. It is the setting in which the estimator-side geometry of working models, projection parameters, and regularization can be seen most directly.

3.1 Local setup and scope

This chapter specializes the dissertation-wide framework of Chapter 2 to HAL working models. The target parameter is written as $\Psi(P_0) = \Psi^F(Q_0)$, with nuisance component g_0 , under a full model $\mathcal{M} = \{P \in \mathcal{M}_1 : Q(P) \in \mathcal{Q}\}$. HAL turns the broad càdlàg function class into an explicit finite-dimensional working model whose score geometry can be studied and regularized directly—which is precisely why it is the natural setting in which to develop the dissertation’s targeting mechanism.

The methodological problem. The initial estimator is the cross-validated HAL fit $\hat{Q}_n$, and the targeting step operates inside a possibly larger HAL working model $\mathcal{M}_n = \{P \in \mathcal{M}_1 : Q(P) \in \mathcal{Q}_n\}$. Because HAL basis functions are constructed on a data-adaptive grid of indicators, the induced score vectors are often highly collinear, and exact adaptive targeting becomes numerically unstable even when the working model is statistically attractive. The

exact working-model efficient influence curve $D_{\mathcal{M}_{n,P}}^*$ is therefore the right theoretical object, but a regularized approximation $\widetilde{D}_{\mathcal{M}_{n,P}}^*$ is the right computational object. This chapter’s main text emphasizes stable approximation of the working-model EIC, with the projection-based construction as the primary method and the Hessian-based construction as a secondary contrast that clarifies why regularizing only the bias-relevant direction is the principled choice.

3.1.1 Related work and positioning

Chapter 2 already records the dissertation-wide semiparametric and targeted-learning framework. The present subsection therefore restricts attention to the more specific literature needed to position the HAL instantiation of regularized A-TMLE.

Earlier regularization strategies for TMLE typically stabilized the update while remaining in the full nonparametric model, for example through propensity-score truncation, collaborative TMLE, early-stopping ideas, or score-span approximations [Petersen et al., 2012, van der Laan and Gruber, 2010, van der Laan and Rose, 2018]. Those approaches are important precursors, but they do not by themselves organize the correction through a data-adaptive working model and its projection parameter.

Adaptive TMLE adds that working-model layer by targeting inside a lower-dimensional model selected from the data [van der Laan et al., 2023, 2026]. In the HAL setting this is especially attractive because the basis expansion makes the score geometry explicit, but it also exposes the numerical weakness of relaxed HAL: once selected basis functions are highly collinear, the working-model information operator can become ill-conditioned and exact adaptive targeting becomes unstable.

This chapter is also closely related to automatic debiasing and Riesz-representer constructions [Chernozhukov et al., 2022b, van der Laan et al., 2025]. In a well-conditioned parametric working model, the Hessian-based update direction is closely connected to Hessian-style Riesz learning. The contribution here is to make that connection operational in the collinear HAL regime by using ridge or lasso regularization of the working-model EIC approximation, by separating the initial estimation working model from the targeting working model, and by treating projection-based targeting as the chapter’s main line.

Seen from this literature, the chapter’s specific contribution is not another generic TMLE regularization device. It is a HAL-specific realization of regularized A-TMLE that uses the explicit score geometry of the HAL sieve to compare exact, Hessian-regularized, and projection-regularized targeting, and then pairs that geometry with data-adaptive working-model enlargement through the plateau selector.

3.2 Problem Setup and Working-Model Efficient Influence Curve Approximation

Recall the setup of Section 3.1: we target $\Psi(P_0) = \Psi^F(Q_0)$ with nuisance g_0 , from n i.i.d. observations $O_1, \dots, O_n \sim P_0$.

HAL fits the outcome function as the empirical risk minimizer among $Q_\beta(x) = \beta_0 + \sum_{j=1}^p \beta_j \phi_j(x)$ with $\|\beta\|_1 \leq C$, where each ϕ_j is a zero-order indicator $\mathbf{1}\{x_s \geq t\}$ on a data-adaptive grid [Tibshirani, 1996, Zou and Hastie, 2005]. In practice, a cross-validated lasso over a sequence of penalty parameters λ is used to yield the cross validated HAL fit $\hat{Q}_n$. During the fitting process, it also yields a nested sequence of finite-dimensional working parameter subspaces $\mathcal{Q}_{n,\lambda} = \{Q_\beta : \beta_0 + \sum_{j \in R_{n,\lambda}} \beta_j \phi_j(x)\}$, where $R_{n,\lambda}$ is the index set of non-zero coefficients in the lasso fit with penalty λ . In addition, we obtain a sequence of working models $\mathcal{M}_\lambda = \{P \in \mathcal{M}_1 : Q(P) \in \mathcal{Q}_{n,\lambda}\}$.

Let λ_{cv} be the cross-validated λ that yields $\hat{Q}_n$. Then the working model $\mathcal{M}_{\lambda_{cv}}$ and the sequence of larger working models $\mathcal{M}_\lambda$ induced by undersmoothed HAL fits for $\lambda \leq \lambda_{cv}$ are the playground for our targeting steps. Let M_n be one of those working models, usually $\mathcal{M}_{\lambda_{cv}}$.

We define the projected target parameter as $\Psi_{\mathcal{M}_n}(P) = \Psi^F(\Pi_{\mathcal{M}_n} Q(P))$, where $\Pi_{\mathcal{M}_n}$ is a loss-based dissimilarity projection of $Q(P)$ onto $\mathcal{M}_n$.

We restrict ourselves to the parametric working model $\mathcal{M}_n$ and estimate $\Psi_{\mathcal{M}_n}(P_0)$ by updating the coefficients of the initial fit $\hat{Q}_n$. Besides using the MLE in $\mathcal{M}_n$ (relaxed HAL) to update the coefficients as a simple implementation of A-TMLE, we propose two regularized versions of A-TMLE to update the coefficients of $\hat{Q}_n$ within $\mathcal{M}_n$, based on two different approximations of the EIC of the projection parameter $\Psi_{\mathcal{M}_n}$.

To formalize the scores, let $L_\beta(O_i)$ be the loss function evaluated at observation O_i (e.g., squared error or negative log-likelihood). We define the “score” vectors as the negative gradient of the loss with respect to the coefficients β :

$$S_{i,j}^\beta = -\frac{\partial L_\beta(O_i)}{\partial \beta_j}.$$

For a negative log-likelihood loss, this coincides precisely with the standard statistical score $\frac{\partial \log p_\beta(O_i)}{\partial \beta_j}$. Let S^β be the $n \times p$ score matrix; the j -th column is $S_{*,j}^\beta$ and the i -th row is defined explicitly as a $1 \times p$ row vector $S_{i,*}^\beta$, such that its transpose $(S_{i,*}^\beta)^\top$ yields a valid $p \times 1$ column vector.

3.2.1 Hessian-Based Working-Model Efficient Influence Curve Approximation

Each Q in the parametric space $\mathcal{Q}_n$ can be indexed by coefficients β as Q_β , where β are the coefficients that can vary freely in our previous definitions.

When the working model $\mathcal{M}_n$ is misspecified relative to the true data generating process P_0 , the true EIC of the projection parameter utilizes the expected Hessian matrix rather than the covariance of the scores. The true parametric EIC for the projection parameter $\Psi^F(Q_\beta)$ within $\mathcal{M}_n$ evaluated at a distribution P_0 is explicitly formulated as:

$$D_\beta^*(O_i) = \left\{ \frac{\partial \Psi^F(Q_\beta)}{\partial \beta} I_{P_0}^{-1}(\beta) \right\} (S_{i,*}^\beta)^\top.$$

Remark 3.1 (Fisher information: Expected Hessian vs. Outer-Product). It is a standard result that for a distribution P_β perfectly correctly specified within the working model, Bartlett's second identity holds, and the expected Hessian of the loss and the outer-product of the log likelihood based scores coincide:

$$E_{P_\beta}[\nabla^2 L_\beta(O)] = E_{P_\beta}[S^\beta(O)S^\beta(O)^\top].$$

However, in our setting, the working model $\mathcal{M}_n$ is a restricted sieve that only *approximates* the true data-generating distribution P_0 ; thus, $\mathcal{M}_n$ is generally misspecified for P_0 ($P_0 \notin \mathcal{M}_n$). Because we evaluate the empirical expectation P_n (which approximates P_0 , not P_β), Bartlett's identity fails even under log likelihood scores. To correctly formulate the exact EIC of the projection parameter under P_0 , we must use the **Hessian form** (Huber's sandwich information matrix [White, 1982]), $I_{P_n}(\beta) = P_n[\nabla^2 L_\beta(O)]$.

Since $D_\beta^*(O_i)$ is linear in $S_{i,*}^\beta$, in the TMLE procedure, $\left\{ \frac{\partial \Psi^F(Q_\beta)}{\partial \beta} I_{P_n}^{-1}(\beta) \right\}$ serves as our update direction to solve $P_n D_\beta^* = o_p(n^{-1/2})$ (see Step 8 in Algorithm 1). The Hessian-based method thus analytically formulates the exact EIC of the projection parameter defined for the parametric working model.

While the gradient $\frac{\partial \Psi^F(Q_\beta)}{\partial \beta}$ can be calculated easily, $I_{P_n}^{-1}(\beta)$ is difficult to compute. Because HAL aims to fit Q_0 well in the L_2 sense under a global sectional variation norm constraint, it inherently selects highly collinear basis functions. Consequently, the empirical Hessian $I_{P_n}(\beta) = P_n[\nabla^2 L_\beta(O)]$ is frequently near-singular, with very large norms of the inverse.

To combat this issue, a straightforward regularization blinded from the target parameters can be used. Ridge regularization (adding ηI to the diagonal before taking the inverse) offers a stable solution, replacing $I_{P_n}^{-1}(\beta)$ with $(I_{P_n}(\beta) + \eta I)^{-1}$ [Carroll et al., 1998, Tsiatis, 2006]. Note that ridge regularization equally reduces variability in all directions. For multi-dimensional target parameters (such as a survival curve), such uniform regularization may not adapt well to specific directions requiring different levels of regularization.

3.2.2 Projection-Based Working-Model Efficient Influence Curve Approximation

Our second approximation method achieves the same theoretical goal but takes a different **computational route**. Instead of analytically formulating the EIC of the projection parameter via parameter gradients and Hessian inversion, it approximates the working-model EIC by directly projecting an initial valid nonparametric influence function $D(O)$ onto the linear span of the working model's score functions using an empirical L_2 regression. The key computational advantage is that this projection can be regularized via lasso without explicit matrix inversion, providing a highly stable, target-specific strategy within the working model's tangent space.

Let $D(O)$ be an initial valid non-parametric canonical gradient $D_{Q_n, g_n}^*(O)$ or any other

valid gradient $D_{Q_n, g_n}(O)$ evaluated at the current estimates (Q_n, g_n) . Note that this gradient depends on the distribution through (Q_n, g_n) , but we ignore the subscript for simplicity. Without regularizing the projection, the unpenalized closest $L_2(P_n)$ projection of $D(O)$ onto the span of the working model HAL scores with an intercept γ_0 , which effectively centers the gradient and scores, yields the projected gradient:

$$\hat{D}_{n,\beta}^*(O_i) = \hat{\alpha}_n^\top (S_{i,*}^\beta)^\top, \quad (3.1)$$

where $(\hat{\gamma}_0, \hat{\alpha}_n) = \operatorname{argmin}_{\gamma_0, \alpha} P_n \left(D(O) - \gamma_0 - \sum_j \alpha_j S_j^\beta(O) \right)^2$.

The unregularized $L_2(P_n)$ projection coefficients have the explicit closed form

$$\hat{\alpha}_n = \left(P_n[S^\beta (S^\beta)^\top] \right)^{-1} P_n[S^\beta D],$$

where $P_n[S^\beta D] = (P_n[S_j^\beta D] : j = 1, \dots, p)$ is the $p \times 1$ vector of empirical covariances between each score component and the initial gradient $D(O)$. Using the pathwise derivative identity $P_n[S^\beta D] = \nabla_\beta \Psi^F(Q_\beta)$, the unregularized projection-based EIC approximation takes the form

$$\hat{D}_{n,\beta}^*(O_i) = \nabla_\beta \Psi^F(Q_\beta)^\top \underbrace{I_{cov,n}^{-1}(\beta)}_{\text{score covariance}} (S_{i,*}^\beta)^\top, \quad I_{cov,n}(\beta) = P_n[S^\beta (S^\beta)^\top].$$

Comparing with the Hessian-based approximation from Section 3.2.1,

$$\hat{D}_{\text{hessian},\beta}^*(O_i) = \nabla_\beta \Psi^F(Q_\beta)^\top \underbrace{I_n^{-1}(\beta)}_{\text{Hessian}} (S_{i,*}^\beta)^\top, \quad I_n(\beta) = P_n[-\nabla_\beta S^\beta] = P_n[\nabla^2 L_\beta],$$

we see that both approximations share the identical structure $\nabla_\beta \Psi^F(Q_\beta)^\top \cdot I^{-1} \cdot S^\beta$ and **differ only in the choice of information matrix**: the projection method uses the empirical score covariance $I_{cov,n} = P_n[S^\beta (S^\beta)^\top]$, while the Hessian method uses the empirical expected Hessian $I_n = P_n[\nabla^2 L_\beta]$. The Hessian method directly estimates the exact EIC of the projection parameter $\Psi_{\mathcal{M}_n}$ for any fixed working model—even under misspecification—because the expected Hessian $I_n = P_n[\nabla^2 L_\beta]$ is the natural empirical analog of the population Hessian $I_0(\beta) = P_0[-\nabla_\beta S^\beta]$ that defines the true working-model EIC:

$$D_{\mathcal{M}_n}^*(P_0)(O_i) = \nabla_\beta \Psi^F(Q_\beta)^\top \underbrace{I_0^{-1}(\beta)}_{\text{population}} (S_{i,*}^\beta)^\top, \quad I_0(\beta) = P_0[-\nabla_\beta S^\beta].$$

The Projection method uses the score covariance $I_{cov,n}$ instead, which equals the Hessian only when Bartlett's second identity holds (Section 3.2.3). Thus, for a fixed misspecified working model, the two methods target genuinely different objects. For likelihood-based losses, as the working model $\mathcal{M}_n$ grows toward an oracle model $\mathcal{M}_0 \subseteq \mathcal{M}$ that contains P_0 , Bartlett's identity is restored and the difference between the two approximations vanishes: both converge to $D_{\mathcal{M}_0}^*(P_0)$, the EIC of the projection parameter $\Psi_{\mathcal{M}_0}$. If, moreover, $\mathcal{M}_0 = \mathcal{M}$, both

converge to the nonparametric EIC $D^*(P_0)$ of the original target parameter $\Psi: \mathcal{M} \rightarrow \mathbb{R}$. For non-likelihood losses (e.g. squared error), the two methods remain genuinely different even when $P_0 \in \mathcal{M}_0$; asymptotic equivalence requires $\mathcal{M}_n \rightarrow \mathcal{M}$ (see Remark 3.2).

Collinearity among the basis scores means the unregularized projection coefficients involve a potentially unstable matrix inverse. However, the $L_2(P_n)$ projection itself is a stable object. We add another layer of approximation to directly estimate it stably by imposing a weak ℓ_1 lasso penalty alongside an unpenalized intercept γ_0 :

$$(\hat{\gamma}_0^{reg}, \hat{\alpha}_n^{reg}) = \underset{\gamma_0, \alpha}{\operatorname{argmin}} P_n \left(D(O) - \gamma_0 - \sum_j \alpha_j S_j^\beta(O) \right)^2 + \lambda_n \|\alpha\|_1$$

We then use $(\hat{\alpha}_n^{reg})^\top$ as the direction of update, forming the regularized approximated EIC $\tilde{D}_{n,\beta}^*(O_i) = (\hat{\alpha}_n^{reg})^\top (S_{i,*}^\beta)^\top$. While a small ℓ_1 regularization introduces a minor additional layer of approximation to the exact projection onto the working model tangent space, it prevents instability, effectively controls variance when the working model dimension is large, and acts as a more robust estimator for the EIC.

We emphasize that we separate the issue of selecting an approximation of the EIC from the targeting step itself. While a relaxed HAL fit targets all score directions simultaneously and indiscriminately, our regularized TMLE step explicitly targets only the single score direction given by the approximate EIC, providing a much more stable update.

This lasso-type regularization naturally adapts to different target parameters (specified as different initial gradients) as it fits a separate lasso projection for each target parameter of interest and thus regularizes specific parts of the current fit. In addition, the regularization level can be fully data-dependent through cross validated lasso for example.

3.2.3 Theoretical Unification of the Two Approximations

When evaluated at a distribution $P_\beta \in \mathcal{M}_n$, the EIC of the projection parameter equals the projection of the nonparametric EIC onto the working tangent space, given the loss is likelihood loss.

To see this algebraically, note that because $P_\beta \in \mathcal{M}_n$, the working model is correctly specified from the perspective of P_β . Thus, Bartlett's second identity holds, and the Fisher information matrix can be written interchangeably as the expected Hessian or the covariance of the scores:

$$I_{P_\beta}(\beta) = E_{P_\beta}[\nabla^2 L_\beta] = E_{P_\beta}[S^\beta (S^\beta)^\top].$$

The Hessian-based method analytically computes the coefficient

$$\alpha_{\text{hessian}} = I_{P_\beta}(\beta)^{-1} \nabla_\beta \Psi^F(Q_\beta).$$

The unregularized Projection method computes the $L_2(P_\beta)$ projection of the true non-

parametric EIC $D^*(O)$ onto the score space, yielding coefficients

$$\alpha_{\text{proj}} = \left(E_{P_\beta} [S^\beta (S^\beta)^\top] \right)^{-1} E_{P_\beta} [S^\beta D^*].$$

By the definition of the pathwise derivative, given the loss is likelihood loss, the covariance of the canonical gradient with the score is exactly the analytical derivative of the parameter:

$$E_{P_\beta} [S^\beta D^*] = \nabla_\beta \Psi^F(Q_\beta).$$

Substituting these relations, we find that:

$$\alpha_{\text{proj}} = (E_{P_\beta} [S^\beta (S^\beta)^\top])^{-1} \nabla_\beta \Psi^F(Q_\beta) = I_{P_\beta}(\beta)^{-1} \nabla_\beta \Psi^F(Q_\beta) = \alpha_{\text{hessian}}.$$

Thus, when $P_0 \in \mathcal{M}_n$ (i.e. the working model is correctly specified), both methods (evaluated under expectation P_0) yield the exact same object $D_{\mathcal{M}_n}^*(P_0)$. While P_0 is generally not exactly in a finite working model, this demonstrates that as $\mathcal{M}_n$ grows to better approximate P_0 , the difference between the two approximations (using P_n) converges to zero. When $\mathcal{M}_n$ grows to an oracle model $\mathcal{M}_0$ containing P_0 , both converge to $D_{\mathcal{M}_0}^*(P_0)$, the EIC of the projection parameter $\Psi_{\mathcal{M}_0}$. When $\mathcal{M}_0 = \mathcal{M}$, both further converge to the nonparametric EIC $D^*(P_0)$ of the original target parameter. Since any HAL sieve eventually includes P_0 , both methods are asymptotically valid, preserving the A-TMLE interpretation (Section 3.3).

Remark 3.2 (Heteroscedasticity and non-likelihood losses). The unification above relies on Bartlett’s second identity, which is specific to likelihood-based loss functions. For non-likelihood losses such as squared error, the two EIC approximations are **fundamentally different**, not merely different regularizations of the same object. The expected Hessian depends only on covariates ($E_P[2\phi\phi^\top]$), while the score outer-product includes residual variance weights ($E_P[4(Y - Q_\beta)^2\phi\phi^\top]$); these are proportional only under homoscedasticity. Crucially, this distinction persists even when the working model contains P_0 : for any oracle model $\mathcal{M}_0 \subsetneq \mathcal{M}$ that is a proper submodel of the full model, the two methods target genuinely different objects. Asymptotic equivalence requires $\mathcal{M}_n \rightarrow \mathcal{M}$ (the full model), at which point both converge to the nonparametric EIC $D^*(P_0)$. See Appendix A.4 for the detailed algebraic development.

3.2.3.1 Connection to Automatic Debiased ML (autoDML) and Riesz Representer

The concept to directly learn the Riesz representer and bypass analytical gradients was pioneered by Chernozhukov et al. [2022b] in the context of Automatic Debiased Machine Learning. Our projection-based method theoretically aligns with their Riesz representer framework and builds upon the recent autoDML framework of van der Laan et al. [2025], which elegantly establishes that for smooth functionals of M-estimands, the exact update direction is given by a Hessian-weighted Riesz representer. As established above, for likelihood-based losses, the unregularized Hessian-based direction and unregularized Projection direction both coincide with the Hessian Riesz representer whenever the working model grows to an oracle

model $\mathcal{M}_0$ containing P_0 ; when $\mathcal{M}_n \rightarrow \mathcal{M}$, this recovers the nonparametric Riesz representer. For non-likelihood losses, all three agree only in the latter limit ($\mathcal{M}_n \rightarrow \mathcal{M}$). In finite samples, both methods additionally differ due to the use of empirical quantities in place of population expectations. Our contribution lies in adapting these ideas via explicit ℓ_1 and ℓ_2 regularization strategies to provide concrete algorithmic implementations that address the specific challenges of HAL:

1. **Regularization under collinearity.** The autoDML framework assumes a positive-definite Hessian operator, which frequently fails for HAL due to near-collinear basis functions. Our explicit ℓ_1 -regularized projection (Section 3.2.2) and ridge-penalized inversion (Section 3.2.1) provide stable computation when the Fisher information matrix is singular or ill-conditioned.
2. **Iterative targeting via universal least favorable paths.** Rather than a one-shot debiasing correction, we employ iterative TMLE updates along universal least favorable submodel paths with early stopping (Algorithm 1), which can improve finite-sample performance for multi-dimensional target parameters such as survival curves.
3. **Adaptive plateau working model selection.** The regHAL-ATMLE extension (Section 3.4) selects the working model complexity in a data-driven manner via a plateau criterion over nested HAL models.

3.2.4 regHAL-ATMLE Algorithms

Once the approximated EIC within the model $\mathcal{M}_n$ is obtained ($\tilde{D}_n^*$), we use the following algorithm to construct targeted updates using universal least favorable paths via small gradient steps.

Algorithm 1 HAL Updating Algorithm Based on Approximated EIC (One Dimensional Target Parameter)

- 1: **Cross-Validated HAL Fit:** Compute cross-validated or undersmoothed initial HAL fit $\hat{Q}_n$.
- 2: **Working Model Definition:** Define the HAL working model Q_β indexed by β (which is not necessarily the working model of the CV fit).
- 3: **Parameter Initialization:** Initialize $\beta^{(0)}$ based on $\hat{Q}_n$, incorporating zero-padding if using an undersmoothed HAL model.
- 4: **if** using the projection-based method **then**
- 5: **Initial Gradient and Nuisances:** Specify the form of the initial non-parametric gradient $D(O)$ and fit any extra nuisance parameters needed (e.g., propensity score g_n for IPTW). Note that during iteration, one can ignore the dynamic centering at the updated $Q_{\beta^{(k)}}$ since we have empirical centering in lasso step through intercept.
- 6: **end if**
- 7: **EIC Computation:** Dynamically derive the approximated regularized EIC $\tilde{D}_n^*$ at the current $\beta^{(k)}$:

$$\tilde{D}_{n,\beta^{(k)}}^*(O_i) = (\hat{\alpha}_n^{(k)})^\top \text{scores}_{\beta^{(k)}}(O_i)$$

where $\text{scores}_{\beta^{(k)}}$ are the scores (i.e., negative gradients of the loss L_β) w.r.t. β .

- 8: **Update Direction:** Determine the update direction for β as:

$$\text{Update}^{(k)} = \text{sign}(P_n \tilde{D}_{n,\beta^{(k)}}^*) \hat{\alpha}_n^{(k)}.$$

- 9: **Parameter Update:** Update $\beta^{(k)}$ to $\beta^{(k+1)}$ using gradient descent:

$$\beta^{(k+1)} = \beta^{(k)} + \epsilon \cdot \text{Update}^{(k)}$$

using a small step size ϵ .

- 10: **Stopping Criterion:** Repeat steps 5–8 until $|P_n \tilde{D}_{n,\beta^{(k+1)}}^*| < \frac{\text{se}(\tilde{D}_{n,\beta^{(k+1)}}^*)}{\sqrt{n} \log(n)}$.
-

Remark on Step Size and Oscillation: In Step 7, utilizing a fixed step size ϵ combined with the **sign** operator functions similarly to coordinate descent, ensuring stable, incremental movement along the universal least favorable submodel. However, to strictly prevent infinite oscillation (chattering) around the tight $o_p(n^{-1/2})$ tolerance boundary in practice, one must employ a dynamically decaying step-size schedule (e.g., halving ϵ upon sign flips) or implement standard early stopping checks. Alternatively, standard iterative TMLE uses a one-dimensional fluctuation regression (e.g., fitting an MLE submodel $Q_\varepsilon = Q_{\beta^{(k)}} + \varepsilon \hat{\alpha}^\top \phi$) to naturally find an exact optimal step size ε at each iteration, which inherently averts chattering and accelerates convergence. For multi-dimensional target parameter cases, careful scaling across dimensions is maintained. The multi-dimensional algorithm is presented in Appendix A.1.

3.3 Theoretical Guarantees

The next results specialize the general working-model A-TMLE theory of van der Laan et al. [2023] to the HAL constructions developed in this chapter. In particular, the theorem for the projection parameter and the theorem transferring that result to the original target parameter are chapter-specific restatements of that A-TMLE theory in HAL notation. They provide a common proof system for both the Hessian-based and projection-based updates and clarify when the targeted estimator remains first-order valid for the projection parameter and, ultimately, for the original target parameter.

As established in Section 3.2.3, both the regularized Hessian-based method and regularized projection method target the same working-model EIC $D_{\mathcal{M}_n}^*$. Provided the finite-sample regularizations vanish suitably, both algorithms therefore plug directly into the single proof system below. We focus on the case where the oracle model is the full model, i.e., $\mathcal{M}_0 = \mathcal{M}$, so that the theory covers the general setting in which the loss need not be likelihood-based; for likelihood-based losses, the A-TMLE framework further provides super-efficiency relative to a restricted oracle model $\mathcal{M}_0 \subsetneq \mathcal{M}$. In this setting the working model grows toward the full model from below as $n \rightarrow \infty$. The related ADML / autoDML line [van der Laan et al., 2025] remains relevant for the smooth M-estimand and Riesz-representer interpretation discussed in Section 3.2.3.1, but the theorem statements below are inherited from the more general A-TMLE working-model results.

3.3.1 TMLE for working model induced target parameter

Let $\Psi_{\mathcal{M}_n} = \Psi^F \circ \Pi_{\mathcal{M}_n} : \mathcal{M} \rightarrow \mathbb{R}$ to represent the projection parameter and $D_{\mathcal{M}_n}^*(P)$ is the efficient influence curve of the projected parameter in the full model $\mathcal{M}$.

Define $R_{\mathcal{M}_n}(P, P_0) = \Psi_{\mathcal{M}_n}(P) - \Psi_{\mathcal{M}_n}(P_0) + P_0 D_{\mathcal{M}_n}^*(P)$ as the exact remainder.

Our algorithm outputs an estimator $P_{\mathcal{M}_n,n}^*$ that sets the empirical mean of the *regularized* approximation $\tilde{D}_n^*$ to $o_p(n^{-1/2})$. For this to satisfy the theoretical A-TMLE requirements, the algorithmic regularization bias must vanish asymptotically, or more formally stated as assumption 3.3.

Assumption 3.3 (Approximation of the working model EIC). *Let $\tilde{D}_n^*$ denote the regularized EIC estimate used for targeting (either the Ridge-penalized Hessian-based method or the Lasso-penalized Projection method). The targeted fit $P_{\mathcal{M}_n,n}^*$ solves the approximate score equation $P_n \tilde{D}_n^*(P_{\mathcal{M}_n,n}^*) = o_p(n^{-1/2})$. We assume this implies that the true working-model EIC score equation or the EIC for the projected parameter is also solved at the same rate at $P_{\mathcal{M}_n,n}^*$:*

$$P_n D_{\mathcal{M}_n}^*(P_{\mathcal{M}_n,n}^*) = o_p(n^{-1/2}).$$

Theorem 3.4 (TMLE for working model induced target parameter, Restated from van der Laan et al. [2023]). *Let $P_{\mathcal{M}_n,n}^*$ denote our regHAL-ATMLE satisfying Assumption 3.3, such that $P_n D_{\mathcal{M}_n}^*(P_{\mathcal{M}_n,n}^*) = o_p(n^{-1/2})$. Assume:*

$$(A.1) \quad D_{\mathcal{M}_n}^*(P_{\mathcal{M}_n,n}^*) \text{ and } D_{\mathcal{M}_n}^*(P_0) \text{ lie in a } P_0\text{-Donsker class with probability tending to 1.}$$

$$(A.2) \quad P_0 \left(D_{\mathcal{M}_n}^*(P_{\mathcal{M}_{n,n}}^*) - D_{\mathcal{M}_n}^*(P_0) \right)^2 = o_p(1).$$

$$(A.3) \quad R_{\mathcal{M}_n}(P_{\mathcal{M}_{n,n}}^*, P_0) = o_p(n^{-1/2}).$$

Condition (A.3) on the exact remainder is usually satisfied under HAL estimated nuisance parameters, since cross-validated HAL converges in loss-based dissimilarity at rate $O_P(n^{-1/3})$ (up to $\log n$ factors), and the exact remainder is generally bounded by products of $L_2(P_0)$ estimation errors of nuisance parameters, which are controlled under positivity. Then $\Psi_{\mathcal{M}_n}(P_{\mathcal{M}_{n,n}}^*) - \Psi_{\mathcal{M}_n}(P_0) = P_n D_{\mathcal{M}_n}^*(P_0) + o_p(n^{-1/2})$.

See detailed proof in Appendix A.5.2.

3.3.2 TMLE for original target parameter

Theorem 3.5 (TMLE for original target parameter, Restated from van der Laan et al. [2023]). *Assuming the oracle model $\mathcal{M}_0$ is the full model $\mathcal{M}_0 = \mathcal{M}$ (so $D_{\mathcal{M}_0}^*(P_0) = D^*(P_0)$). Let $P_{0,\mathcal{M}_n}$ be the Kullback-Leibler (KL) projection of P_0 on $\mathcal{M}_n$. If we have:*

$$(B.1) \quad P_0 \left(D_{\mathcal{M}_n}^*(P_0) - D^*(P_0) \right)^2 = o_p(1).$$

and

$$(C.2) \quad (P_{0,\mathcal{M}_n} - P_0)(D^*(P_{0,\mathcal{M}_n}) - D_{\mathcal{M}_n}^*(P_{0,\mathcal{M}_n})) = o_p(n^{-1/2}).$$

$$(C.3) \quad P_0 D_{\mathcal{M}_n}^*(P_{0,\mathcal{M}_n}) = o_p(n^{-1/2}).$$

Note that under $\mathcal{M}_0 = \mathcal{M}$, condition (C.1) from the general A-TMLE framework—namely $\Psi(P_{0,\mathcal{M}_n}) - \Psi_{\mathcal{M}_0}(P_{0,\mathcal{M}_n}) = o_p(n^{-1/2})$ —holds trivially as an identity. Then, $\Psi_{\mathcal{M}_n}(P_{\mathcal{M}_{n,n}}^*) - \Psi(P_0) = P_n D^*(P_0) + o_p(n^{-1/2})$.

3.3.3 Verification of Regularity Conditions under HAL for efficiency

Assumption 3.6 (Nuisance parameters HAL estimation). *Let Q_0 and other nuisance parameters g_0 lie in the class of càdlàg functions on $[0, 1]^d$ with sectional-variation norm bounded by M . Let $\mathcal{Q}_n$ be the finite-dimensional HAL sieve selected by cross-validated or undersmoothed k -th order HAL fit, which induces working model $\mathcal{M}_n$. Furthermore, we assume the HAL sieve is sufficiently rich to linearly approximate both Q_0 and the clever covariate H_0 .*

In practice, however, our current software implementations frequently rely on the empirical assumption that the adaptively undersmoothed Q_n basis (via regHAL-ATMLE) is sufficiently rich to span H_0 independently without requiring explicit double-selection unions.

Theorem 3.7. *Under Assumption 3.6, setting the oracle model to the full model $\mathcal{M}_0 = \mathcal{M}$, as $\mathcal{M}_n \rightarrow \mathcal{M}$ from below, conditions (A.1)-(A.3), (B.1), (C.1)-(C.3) hold, and the regHAL-ATMLE achieves standard asymptotic efficiency in the full non-parametric model.*

See detailed proof in Appendix A.5.3, where we explicitly demonstrate via Cauchy-Schwarz bounds that fulfilling Condition C.2 relies inherently on the HAL working model being sufficiently rich to span the true clever covariate.

3.3.4 Positioning: Regularized TMLE, A-TMLE, and our contribution

Before the A-TMLE framework, various approaches regularized the nonparametric TMLE targeting step without invoking working models or projection parameters. These include truncation of clever covariates [Petersen et al., 2012], Collaborative TMLE (C-TMLE) which builds the propensity score model collaboratively to stabilize targeting [van der Laan and Gruber, 2010, Ju et al., 2019], early stopping of iterative updates, and approximating the EIC with linear combinations of HAL scores as in van der Laan and Rose [2018]. Such approaches—which we broadly term *regularized TMLE* (regTMLE)—do not require a specific working model structure: one can use super learner for the initial offset, truncate propensity scores, or approximate the EIC in any suitable manner, as long as the approximation converges fast enough to the true EIC. Similarly, in higher-order TMLE for second-order bias correction [van der Laan et al., 2021], second-order and first-order EICs are approximated using linear spans of HAL scores, possibly adaptively selected.

The A-TMLE framework [van der Laan et al., 2023] introduced a more structured approach and provided the additional benefit of also estimating a projection parameter that is itself of interest. A-TMLE requires: (i) selecting a data-adaptive working model $\mathcal{M}_n$, (ii) ensuring the initial estimator belongs to that model, and (iii) targeting the projection parameter $\Psi_{\mathcal{M}_n}$ within the working model. This yields super-efficiency guarantees when the oracle model $\mathcal{M}_0$ is a proper submodel, and semiparametric efficiency when $\mathcal{M}_0 = \mathcal{M}$. Key implementations include relaxed HAL—refitting coefficients for selected basis functions—and autoTMLE [van der Laan et al., 2025]. However, existing A-TMLE implementations face stability issues when the working model’s Fisher information is ill-conditioned due to collinearity among selected HAL basis functions. Notably, A-TMLE can itself be understood as a more restrictive instance of regularized TMLE: requiring the estimator to remain within the working model inherently provides regularization by restricting the targeting update to the finite-dimensional tangent space of $\mathcal{M}_n$, preventing the TMLE from exploiting potentially unstable directions in the full nonparametric model. Thus, the working model constraint serves a dual role—defining the projection parameter for potential super-efficiency guarantees while simultaneously regularizing the targeting step.

Our regHAL-ATMLE operates within the A-TMLE framework, adding two contributions: (i) *separation of estimation and targeting working models*—the CV-selected HAL provides the initial fit as offset, while targeting operates in a possibly larger working model (under-smoothed or plateau-selected), providing an additional regularization dimension; (ii) *regularized targeting within the working model*—instead of requiring well-conditioned information matrices (as relaxed HAL does), we use ridge (Hessian-based) or lasso (projection-based) regularization to solve the EIC equation stably, even in very large working models with built-in co-linearity. In addition, we employ the iterative targeting with early stopping to enhance the finite-sample performance. The principle of regularized targeting within the A-TMLE framework means we can target in much larger working models than was previously feasible. More broadly, the regularized targeting machinery developed here is not limited to HAL-only initial estimators: by treating any initial estimator (e.g., a super learner) as an offset within

a HAL working model, one obtains a more general regularized A-TMLE. Chapter 7 returns to this broader program.

While our methods are A-TMLEs by construction, they also satisfy the conditions of a more general result. If one is only interested in the original target parameter (rather than the projection parameter), one could directly approximate the EIC $D_{Q,g}^*$ with a score in the working model tangent space without explicitly defining a projection parameter, as long as the approximation converges fast enough (e.g., at rate $n^{-1/4}$ in L_2). The following theorem formalizes this observation.

Theorem 3.8 (Regularized TMLE for the original target parameter). *Let $\tilde{D}_n^*$ be an approximation of the full model EIC $D_{Q,g}^*$. Let P_n^* denote a TMLE that solves the approximate score equation $P_n \tilde{D}_n^*(P_n^*) = o_P(n^{-1/2})$. Define the exact remainder $R(P, P_0) = \Psi(P) - \Psi(P_0) + P_0 D_{Q(P),g(P)}^*$. Assume:*

$$(R.1) \quad D^*(P_n^*) \text{ and } D^*(P_0) \text{ lie in a } P_0\text{-Donsker class with probability tending to 1.}$$

$$(R.2) \quad P_0 \left(D^*(P_n^*) - D^*(P_0) \right)^2 = o_P(1).$$

$$(R.3) \quad R(P_n^*, P_0) = o_P(n^{-1/2}).$$

$$(R.4) \quad P_n \left(D_{Q_n^*, g_n}^*(P_n^*) - \tilde{D}_n^*(P_n^*) \right) = o_P(n^{-1/2}), \text{ where } Q_n^* = Q(P_n^*).$$

Then

$$\Psi(P_n^*) - \Psi(P_0) = P_n D^*(P_0) + o_P(n^{-1/2}),$$

$$\text{and } \sqrt{n} \left(\Psi(P_n^*) - \Psi(P_0) \right) \xrightarrow{d} N \left(0, P_0 D^*(P_0)^2 \right).$$

The proof follows a standard empirical process argument: expand $\Psi(P_n^*) - \Psi(P_0)$ using pathwise differentiability, control the empirical process term via the Donsker property (R.1–R.2), and bound the score equation approximation error via (R.4). See Appendix A.5.1 for details.

Remark 3.9 (Sufficient conditions for (R.4)). Condition (R.4) requires that the regularized score $\tilde{D}_n^*$ approximates the true EIC D^* well enough in an empirical-mean sense at the targeted fit. Decomposing $P_n(D^* - \tilde{D}_n^*) = (P_n - P_0)(D^* - \tilde{D}_n^*) + P_0(D^* - \tilde{D}_n^*)$: the empirical process part is $o_P(n^{-1/2})$ whenever $\|D^*(P_n^*) - \tilde{D}_n^*(P_n^*)\|_{L_2(P_0)} \rightarrow 0$ in probability and the function class is Donsker. The population part $P_0(D^*(P_n^*) - \tilde{D}_n^*(P_n^*))$ requires further analysis depending on the parameter structure. For double-robust target parameters such as the ATE, whose exact remainders are bounded by products of nuisance estimation errors (i.e., $|R| \lesssim \|Q_n - Q_0\| \cdot \|g_n - g_0\|$), an $L_2(P_0)$ approximation rate of $o_P(n^{-1/4})$ for $\tilde{D}_n^* \rightarrow D^*$ is sufficient.

Remark 3.10 (Dual justification of regHAL-ATMLE). Our regHAL-ATMLEs are A-TMLEs that also satisfy Theorem 3.8's conditions, providing a dual theoretical justification. This dual nature reflects the fact that A-TMLE is itself a structured form of regularized TMLE—one where the working model constraint provides implicit regularization of the targeting

step. Specifically, when the oracle model is the full model ($\mathcal{M}_0 = \mathcal{M}$), the A-TMLE theory of Theorems 3.4–3.5 reduces to a special case of Theorem 3.8: conditions (R.1)–(R.4) reduce to the conjunction of (A.1)–(A.3), (B.1), and (C.1)–(C.3). However, when the goal is super-efficiency with respect to an oracle model $\mathcal{M}_0 \subsetneq \mathcal{M}$, the A-TMLE framework provides formal guarantees for the projection parameter $\Psi_{\mathcal{M}_0}$ that Theorem 3.8 does not address. Theorem 3.8 is more general in a different direction: $\tilde{D}_n^*$ need not be the approximation for any exact EIC of well-defined projection parameter. In particular, any regularized linear combination of working-model scores that converges to D^* sufficiently fast will yield a valid and efficient TMLE for the original target parameter. This broader regTMLE perspective encompasses earlier approaches such as truncation, C-TMLE, and HAL score projections [van der Laan and Rose, 2018] that do not invoke the working model or projection parameter framework.

Both of our proposed methods—Hessian-based and Projection-based—are A-TMLEs using different approximations of the EIC of the projection parameter, as established in Sections 3.2.2 and 3.2.3. Since the HAL sieve guarantees that $\mathcal{M}_n$ eventually includes P_0 , both methods solve approximate EIC equations for the oracle projection parameter $\Psi_{\mathcal{M}_0}(P_0)$ with asymptotically equivalent approximation strategies.

3.4 Data-Adaptive Working Model Selection via Plateau Selector

A key question within the regHAL-ATMLE framework is how to choose the working model complexity for targeting. We consider three strategies: (i) using the cross-validated HAL fit as the working model, (ii) using a fixed undersmoothed model, and (iii) data-adaptive selection via a plateau method. The first two use a pre-specified working model, while the third adapts the working model complexity to the specific target parameter. This section develops the plateau approach.

Too simple working models may leave residual bias in the EIC approximation; too complex models inflate variance through collinearity. Traditional cross-validation is computationally expensive and may not optimize the bias-variance trade-off for the specific target parameter of interest. Motivated by Lepski’s method and its applications to CV-TMLE tuning [Cai and van der Laan, 2019, van der Laan et al., 2023], we develop data-driven working model complexity selection by constructing nested working models of increasing complexity, conducting regularized targeting within each, and using a “plateau selector” to identify when additional complexity harms inference quality.

The Plateau Principle Our selector embodies a simple but powerful idea: stop adding complexity when the estimate-confidence interval relationship becomes inconsistent. Specifically, we evaluate a sequence of nested HAL working models created from decreasing penalty parameters λ_j . The algorithm stops when:

- Estimates increase but lower confidence bounds decrease, **OR**

- Estimates decrease but upper confidence bounds increase

Algebraically, the condition where the estimate increases ($\hat{\psi}_j > \hat{\psi}_{j-1}$) but the lower bound decreases ($L_j < L_{j-1}$) implies $\hat{\psi}_j - z_{1-\alpha/2}SE_j < \hat{\psi}_{j-1} - z_{1-\alpha/2}SE_{j-1}$. Rearranging yields $\hat{\psi}_j - \hat{\psi}_{j-1} < z_{1-\alpha/2}(SE_j - SE_{j-1})$, explicitly showing that the algorithm stops exactly when the variance inflation strictly outpaces the magnitude of the bias correction—the hallmark of overfitting.

3.4.1 Plateau Selector Algorithm on Nested Working Models

Consider a sequence of nested working models with increasing complexity:

$$\mathcal{M}_1 = \mathcal{M}_{CV} \subset \mathcal{M}_2 \subset \mathcal{M}_3 \subset \cdots \subset \mathcal{M}_K,$$

where $\mathcal{M}_1$ corresponds to the cross-validated HAL fit and subsequent models incorporate additional HAL basis functions from undersmoothed fits. For each working model $\mathcal{M}_j$, we derive point estimates and confidence intervals for the projected target parameter $\Psi(P_0^{\mathcal{M}_j})$ using our three targeting approaches: relaxed MLE, projection-based, and Hessian-based targeting.

Our plateau selector identifies optimal model complexity j^* by monitoring the behavior of estimates and confidence intervals across increasing complexity levels.

Algorithm 2 Plateau Selector for Optimal Model Complexity

```

1: Input: Point estimates  $\{\hat{\psi}_j\}_{j=1}^K$  and confidence intervals  $\{[L_j, U_j]\}_{j=1}^K$ 
2: Initialize:  $j^* = 1$ 
3: for  $j = 2$  to  $K$  do
4:   if  $\hat{\psi}_j > \hat{\psi}_{j-1}$  and  $L_j < L_{j-1}$  then
5:     Return  $j^* = j - 1$  {Increasing estimate, decreasing lower bound}
6:   else if  $\hat{\psi}_j < \hat{\psi}_{j-1}$  and  $U_j > U_{j-1}$  then
7:     Return  $j^* = j - 1$  {Decreasing estimate, increasing upper bound}
8:   else
9:      $j^* \leftarrow j$  {Update optimal complexity}
10:  end if
11: end for
12: Return  $j^*$ 

```

Intuition: This provides a principled, data-adaptive approach to balance the bias-variance trade-off without computationally intensive re-evaluations. As rigorously proven in Appendix A.5.3 via exact remainder bounds, adaptively enriching the working model fundamentally guarantees that the active functional basis becomes sufficiently rich to efficiently approximate the clever covariate, thereby securing $o_P(n^{-1/2})$ rates even under positivity violations. An alternative bridged targeting approach that warm-starts each model from the previous solution is described in Appendix A.6.

3.5 regHAL-ATMLE for Average Treatment Effects

We evaluate our regularized targeting approaches during A-TMLE’s targeting step through comprehensive simulations for average treatment effect (ATE) estimation. The ATE setting is particularly valuable for demonstrating our methodology because it allows us to test robustness under practical positivity violations—a common real-world challenge where propensity scores approach zero or one.

3.5.1 ATE Set up

Suppose our data is n i.i.d copies of $O_i = (W_i, A_i, Y_i)$. The ATE target parameter is:

$$\Psi^F(Q_0) = E_0[Q_0(1, W)] - E_0[Q_0(0, W)],$$

where $Q_0(A, W) = E_0[Y|A, W]$. We estimate Q_0 with loss $L_\beta(O_i) = (Y_i - Q_\beta(A_i, W_i))^2$ by cross-validated or undersmoothed HAL fit:

$$\hat{Q}_n(A, W) = \beta_0^{(0)} + \sum_{j=1}^p \beta_j^{(0)} \phi_j(A, W),$$

where $\phi_j(A, W)$ are HAL bases with non-zero coefficients. This induces a parametric working model $\mathcal{Q}_n = \{Q_\beta(A, W) = \beta_0 + \sum_{j=1}^p \beta_j \phi_j(A, W) : \beta\}$. At each β , the score matrix entries are the negative loss gradients:

$$S_{i,j}^\beta = -\frac{\partial L_\beta(O_i)}{\partial \beta_j} = 2(Y_i - Q_\beta(A_i, W_i))\phi_j(A_i, W_i).$$

Implementation of the regularized targeting steps and statistical details appear in Appendix A.2.

3.5.2 Simulation Design

We implement both proposed methods using weak regularization: projection-based targeting with $\lambda = 10^{-5}$ for Lasso projection, and Hessian-based targeting with ridge parameter $\eta = 10^{-6}$. Even this minimal regularization provides crucial numerical stability over unregularized projection. Gradient step size is 10^{-4} for both approaches. We evaluate performance within CV-selected HAL working models and undersmoothed variants (using the 10th largest $\lambda \leq \lambda_{cv}$ or smallest available).

We use two data generating processes: DGP 1 (standard conditions) and DGP 2 (practical positivity violations via extreme propensity scores, details in Appendix A.2.2.2–A.2.2.3). As a benchmark, we also include the standard TMLE update, which uses the initial HAL Q fit as offset and targets with the nonparametric clever covariate $H(A, W) = A/g(W) - (1-A)/(1-g(W))$, with propensity scores truncated to $[0.01, 0.99]$ (details in Appendix A.2.3). We use 500 Monte Carlo replications for fixed working model experiments and 200 replications

for plateau-selected model experiments due to the higher computational cost of evaluating multiple nested models.

It is important to separate the statistical performance of the regularized TMLE point estimators (Abs Bias, Std Err, MSE) from the performance of their corresponding confidence interval constructions. Each regularized TMLE provides its own EIC approximation to construct its CI. **Cov NP** refers to the confidence interval coverage using the standard non-parametric EIC variance estimator evaluated at the targeted fits. Note that **Cov NP** consistently achieves near-nominal coverage because the targeted point estimator successfully removes the first-order post-selection bias inherent in the initial HAL fit, allowing the standard non-parametric EIC sample variance to validly quantify uncertainty. **Cov Proj** and **Cov Hess** refer to coverage using the variance of the projection-based and Hessian-based approximated EICs, respectively. **Cov Proj CV** uses cross-validation to select the regularization parameter for the projection variance estimator. **Oracle Cov** uses the true empirical variability of the estimator. ‘Cov Proj’ under-covers structurally because it uses the variance of the Lasso-projected EIC.

3.5.3 Results with Fixed Working Models: Standard Conditions (DGP 1)

Tables 3.1 and 3.2 present results under DGP 1 using 500 replications across sample sizes $n \in \{500, 1000, 1500, 2000\}$.

Superior bias reduction: Projection-based targeting achieves 50–85% bias reduction versus relaxed HAL across all scenarios. At $n = 500$, absolute bias drops from 0.0421 (relaxed) to 0.0058 (projection).

MSE improvements: Projection targeting reduces MSE by 30–40% compared to relaxed HAL and 15–25% versus Hessian-based targeting. Performance advantages persist under undersmoothing, indicating effective regularization of expanded basis sets.

Comparison with standard TMLE: We also include the standard TMLE update as a benchmark, which uses the nonparametric clever covariate $H(A, W) = A/g(W) - (1 - A)/(1 - g(W))$ with propensity score truncation at $[0.01, 0.99]$ (implementation details in Appendix A.2.3). Under standard conditions, the standard TMLE performs comparably to projection-based regHAL-ATMLE: both achieve low bias (0.001–0.009) and near-nominal Cov_{NP} coverage (93–97%) across all sample sizes, confirming that our regularized methods match the performance of standard TMLE under well-behaved conditions.

Estimator	n	Targeting	Bias	SE	MSE	Cov _{NP} ^a	Cov _{Proj} ^a	Cov _{PCV} ^a	Cov _{Hess} ^a	Cov _{Orc} ^a
HAL-TMLE	500	Standard	0.007	0.116	0.0136	93.8 (0.447)	—	—	—	95.0 (0.456)
	500	Relaxed	0.042	0.160	0.0271	80.0 (0.439)	65.0 (0.343)	61.5 (0.309)	63.5 (0.320)	97.0 (0.625)
	500	Projection	0.006	0.131	0.0171	89.5 (0.445)	85.5 (0.368)	80.5 (0.343)	76.5 (0.306)	94.5 (0.513)
	500	Hessian-based	0.053	0.139	0.0220	87.5 (0.445)	77.0 (0.367)	74.0 (0.341)	66.5 (0.307)	93.5 (0.544)
	1000	Standard	0.001	0.083	0.0068	94.0 (0.315)	—	—	—	95.2 (0.324)
	1000	Relaxed	0.007	0.086	0.0074	91.5 (0.312)	87.0 (0.249)	85.0 (0.233)	86.0 (0.229)	93.0 (0.337)
	1000	Projection	0.007	0.078	0.0060	95.0 (0.314)	92.5 (0.265)	91.5 (0.256)	85.5 (0.220)	94.5 (0.304)
	1000	Hessian-based	0.030	0.083	0.0079	91.5 (0.314)	88.0 (0.265)	87.5 (0.255)	80.0 (0.220)	93.0 (0.327)
	1500	Standard	0.006	0.068	0.0046	94.6 (0.257)	—	—	—	96.0 (0.265)
	1500	Relaxed	0.002	0.067	0.0044	96.5 (0.256)	87.0 (0.206)	81.5 (0.198)	80.5 (0.186)	97.5 (0.262)
	1500	Projection	0.003	0.066	0.0043	97.5 (0.258)	89.5 (0.220)	89.0 (0.214)	78.0 (0.179)	97.5 (0.258)
	1500	Hessian-based	0.018	0.068	0.0050	94.5 (0.258)	86.0 (0.218)	85.5 (0.213)	76.5 (0.179)	95.0 (0.268)
	2000	Standard	0.002	0.057	0.0032	94.4 (0.222)	—	—	—	94.2 (0.223)
	2000	Relaxed	0.004	0.058	0.0033	97.5 (0.222)	86.5 (0.177)	85.5 (0.171)	83.5 (0.162)	98.0 (0.226)
	2000	Projection	0.009	0.057	0.0033	96.5 (0.222)	89.5 (0.185)	88.5 (0.180)	82.5 (0.158)	96.5 (0.223)
	2000	Hessian-based	0.020	0.058	0.0037	94.0 (0.222)	87.0 (0.185)	86.0 (0.179)	80.5 (0.158)	94.5 (0.227)

Note: Based on 500 Monte Carlo replications per sample size.

^a Coverage (%) of nominal 95% confidence intervals; mean interval width in parentheses.

Cov_{NP}: nonparametric EIC variance; Cov_{Proj}: projection-based EIC variance;

Cov_{PCV}: projection with CV-selected penalty; Cov_{Hess}: Hessian-based EIC variance;

Cov_{Orc}: oracle (true empirical) variance.

“—” indicates columns not applicable to the standard TMLE update (see Appendix A.2.3).

Table 3.1: Simulation results for regHAL-ATMLE under standard conditions (DGP 1) with CV-selected working model.

Estimator	n	Targeting	Bias	SE	MSE	Cov _{NP} ^a	Cov _{Proj} ^a	Cov _{PCV} ^a	Cov _{Hess} ^a	Cov _{Orc} ^a
HAL-TMLE	500	Standard	0.007	0.116	0.0136	93.8 (0.447)	—	—	—	95.0 (0.456)
	500	Relaxed	0.030	0.163	0.0271	79.5 (0.436)	66.3 (0.343)	61.5 (0.301)	66.3 (0.330)	96.4 (0.637)
	500	Projection	0.010	0.136	0.0186	88.6 (0.445)	85.5 (0.379)	80.7 (0.349)	74.1 (0.307)	94.6 (0.534)
	500	Hessian-based	0.049	0.147	0.0240	84.9 (0.445)	77.1 (0.378)	71.1 (0.350)	65.7 (0.309)	94.6 (0.578)
	1000	Standard	0.001	0.083	0.0068	94.0 (0.315)	—	—	—	95.2 (0.324)
	1000	Relaxed	0.015	0.085	0.0074	92.9 (0.312)	90.0 (0.250)	90.0 (0.231)	90.0 (0.229)	94.3 (0.334)
	1000	Projection	0.017	0.081	0.0068	94.3 (0.315)	92.9 (0.270)	92.9 (0.262)	84.3 (0.218)	94.3 (0.319)
	1000	Hessian-based	0.042	0.087	0.0091	92.9 (0.315)	87.1 (0.270)	82.9 (0.262)	74.3 (0.218)	92.9 (0.339)
	1500	Standard	0.006	0.068	0.0046	94.6 (0.257)	—	—	—	96.0 (0.265)
	1500	Relaxed	0.004	0.067	0.0045	95.0 (0.254)	83.5 (0.203)	81.5 (0.187)	83.0 (0.191)	97.0 (0.263)
	1500	Projection	0.003	0.066	0.0043	97.5 (0.258)	89.5 (0.220)	88.0 (0.209)	78.0 (0.182)	97.0 (0.259)
	1500	Hessian-based	0.019	0.069	0.0052	94.0 (0.258)	85.5 (0.220)	81.5 (0.209)	76.0 (0.181)	95.0 (0.272)
	2000	Standard	0.002	0.057	0.0032	94.4 (0.222)	—	—	—	94.2 (0.223)
	2000	Relaxed	0.001	0.057	0.0033	96.5 (0.219)	87.5 (0.174)	84.5 (0.163)	86.0 (0.166)	96.5 (0.225)
	2000	Projection	0.009	0.057	0.0033	96.5 (0.222)	90.0 (0.186)	88.0 (0.177)	84.0 (0.160)	96.5 (0.223)
	2000	Hessian-based	0.020	0.058	0.0037	94.5 (0.222)	87.0 (0.186)	85.0 (0.176)	80.0 (0.160)	95.0 (0.228)

Note: Based on 500 Monte Carlo replications per sample size.

All column definitions are identical to Table 3.1.

Table 3.2: Simulation results for regHAL-ATMLE under standard conditions (DGP 1) with undersmoothed working model.

3.5.4 Results with Fixed Working Models: Positivity Violations (DGP 2)

Tables 3.3 and 3.4 examine performance under DGP 2’s challenging positivity violation scenario.

Dramatic improvements: Projection-based targeting shows 80% bias reduction under severe positivity violations ($n = 500$: $0.2056 \rightarrow 0.0416$) while maintaining reasonable coverage (Cov NP: 92–95%). Relaxed HAL and Hessian-based targeting suffer substantial coverage degradation.

Methodological Insights: While both projection-based and Hessian-based targeting regularize the canonical gradient within the linear span of scores that matter for fitting Q , projection-based targeting’s consistent superiority stems from the robust approximation of the non-parametric EIC and the additional L_1 regularization of the projection. By intrinsically acting as a variance-weighted regularized regression, it selectively shrinks HAL basis components contributing minimally to bias reduction. This targeted approach proves especially valuable under collinearity—common in HAL implementations—and challenging inference conditions.

Standard TMLE under positivity violations: While the standard TMLE with propensity score truncation maintains low bias under positivity violations (0.000–0.018), it suffers from substantially inflated variance ($SE = 0.137$ – 0.286 versus 0.118 – 0.240 for projection) and degraded Cov_{NP} coverage (84–90% versus 92–96% for projection-based regHAL-ATMLE). The gap is largest at $n \leq 1000$, where the standard TMLE’s reliance on inverse-propensity weighting amplifies variability in the presence of near-zero propensity scores. Note that the standard TMLE’s oracle coverage remains near-nominal ($\approx 95\%$), indicating that the variance estimator rather than the point estimator is the primary source of under-coverage.

Estimator	n	Targeting	Bias	SE	MSE	Cov_{NP}^a	Cov_{Proj}^a	Cov_{PCV}^a	Cov_{Hess}^a	Cov_{Orc}^a
HAL-TMLE	500	Standard	0.012	0.286	0.0820	84.2 (0.883)	—	—	—	94.4 (1.123)
	500	Relaxed	0.206	0.166	0.0697	85.0 (0.999)	68.0 (0.642)	57.0 (0.500)	52.0 (0.400)	74.0 (0.650)
	500	Projection	0.042	0.240	0.0591	92.5 (0.958)	73.0 (0.629)	63.0 (0.537)	53.5 (0.381)	95.0 (0.941)
	500	Hessian-based	0.195	0.215	0.0839	88.0 (1.062)	74.0 (0.697)	61.0 (0.541)	52.0 (0.392)	90.0 (0.843)
	1000	Standard	0.003	0.206	0.0422	86.4 (0.675)	—	—	—	95.6 (0.806)
	1000	Relaxed	0.198	0.124	0.0545	83.0 (0.715)	58.0 (0.462)	53.5 (0.423)	32.5 (0.289)	62.0 (0.488)
	1000	Projection	0.031	0.180	0.0330	95.5 (0.713)	77.0 (0.472)	74.0 (0.435)	49.0 (0.272)	96.0 (0.704)
	1000	Hessian-based	0.191	0.133	0.0541	88.0 (0.746)	62.0 (0.489)	55.5 (0.436)	28.0 (0.289)	70.0 (0.521)
	1500	Standard	0.018	0.163	0.0269	88.2 (0.569)	—	—	—	96.8 (0.640)
	1500	Relaxed	0.203	0.109	0.0528	73.5 (0.585)	43.0 (0.378)	40.0 (0.357)	21.0 (0.231)	55.0 (0.426)
	1500	Projection	0.030	0.148	0.0227	93.0 (0.593)	77.5 (0.398)	76.5 (0.380)	56.0 (0.218)	94.0 (0.580)
	1500	Hessian-based	0.203	0.107	0.0527	79.0 (0.608)	48.5 (0.399)	43.5 (0.373)	21.5 (0.227)	56.5 (0.420)
	2000	Standard	0.000	0.137	0.0188	90.0 (0.475)	—	—	—	96.0 (0.538)
	2000	Relaxed	0.069	0.118	0.0186	90.0 (0.466)	77.0 (0.342)	75.5 (0.326)	64.5 (0.247)	90.0 (0.464)
	2000	Projection	0.027	0.118	0.0145	95.0 (0.479)	82.0 (0.362)	80.0 (0.346)	65.0 (0.230)	94.5 (0.462)
	2000	Hessian-based	0.071	0.215	0.0508	86.0 (0.491)	70.5 (0.366)	68.5 (0.351)	54.0 (0.240)	98.5 (0.841)

Note: Based on 500 Monte Carlo replications per sample size.
All column definitions are identical to Table 3.1.

Table 3.3: Simulation results for regHAL-ATMLE under practical positivity violations (DGP 2) with CV-selected working model.

Estimator	n	Targeting	Bias	SE	MSE	Cov _{NP} ^a	Cov _{Proj} ^a	Cov _{PCV} ^a	Cov _{Hess} ^a	Cov _{Orc} ^a
HAL-TMLE	500	Standard	0.012	0.286	0.0820	84.2 (0.883)	—	—	—	94.4 (1.123)
	500	Relaxed	0.199	0.174	0.0694	87.7 (0.982)	72.5 (0.644)	57.3 (0.484)	53.2 (0.427)	76.0 (0.681)
	500	Projection	0.034	0.235	0.0561	92.4 (0.967)	77.8 (0.655)	65.5 (0.531)	54.4 (0.396)	94.7 (0.921)
	500	Hessian-based	0.212	0.265	0.1146	89.5 (1.071)	78.4 (0.725)	63.2 (0.522)	52.1 (0.419)	95.3 (1.039)
	1000	Standard	0.003	0.206	0.0422	86.4 (0.675)	—	—	—	95.6 (0.806)
	1000	Relaxed	0.212	0.118	0.0587	79.8 (0.728)	56.9 (0.470)	50.5 (0.428)	28.4 (0.294)	53.2 (0.463)
	1000	Projection	0.045	0.182	0.0349	95.4 (0.730)	76.2 (0.487)	73.4 (0.447)	49.5 (0.273)	95.4 (0.715)
	1000	Hessian-based	0.204	0.134	0.0595	84.4 (0.756)	59.6 (0.498)	52.3 (0.450)	26.6 (0.285)	66.1 (0.527)
	1500	Standard	0.018	0.163	0.0269	88.2 (0.569)	—	—	—	96.8 (0.640)
	1500	Relaxed	0.069	0.116	0.0181	91.0 (0.535)	83.5 (0.395)	78.0 (0.367)	68.0 (0.275)	91.0 (0.453)
	1500	Projection	0.035	0.144	0.0219	93.0 (0.582)	83.0 (0.451)	79.5 (0.418)	62.0 (0.252)	94.0 (0.565)
	1500	Hessian-based	0.066	0.260	0.0716	88.5 (0.610)	76.0 (0.466)	72.5 (0.430)	53.5 (0.266)	96.5 (1.019)
	2000	Standard	0.000	0.137	0.0188	90.0 (0.475)	—	—	—	96.0 (0.538)
	2000	Relaxed	0.032	0.117	0.0146	91.5 (0.433)	82.5 (0.352)	79.5 (0.332)	76.0 (0.304)	94.0 (0.457)
	2000	Projection	0.025	0.119	0.0147	94.0 (0.474)	85.5 (0.400)	83.5 (0.383)	70.0 (0.267)	95.0 (0.467)
	2000	Hessian-based	0.031	0.296	0.0881	87.0 (0.493)	79.5 (0.411)	75.5 (0.390)	62.0 (0.283)	96.0 (1.160)

Note: Based on 500 Monte Carlo replications per sample size.
All column definitions are identical to Table 3.1.

Table 3.4: Simulation results for regHAL-ATMLE under practical positivity violations (DGP 2) with undersmoothed working model.

3.5.5 Results with Plateau-Selected Working Models

We now evaluate the plateau-based adaptive working model selection strategy (Section 3.4). We construct nested models using an index $j \in \{1, \dots, 20\}$ corresponding to the 20 largest penalty parameters $\lambda \leq \lambda_{cv}$ when available. Thus, $j = 1$ represents the strict CV-selected model, and $j = 20$ represents a highly undersmoothed model. Since HAL basis sets may not be naturally nested as λ decreases, we create increasing working models by gradually incorporating additional unique active basis functions. Simulations employ 200 replications with projection-based ($\lambda_n = 10^{-5}$), Hessian-based ($\eta_n = 10^{-6}$), and relaxed MLE targeting approaches.

Standard conditions (DGP 1). Table 3.5 compares the three A-TMLE targeting approaches with plateau-selected working models against the fixed working model baselines (CV regHAL and US regHAL, both with projection-based targeting). Under DGP 1, all methods converge to comparable performance at moderate-to-large sample sizes: for $n \geq 1500$, bias is below 0.009, MSE is ≈ 0.003 – 0.004 , and nonparametric coverage exceeds 95% across all projection-based methods.

The plateau selector adapts working model complexity to the available information: it selects intermediate complexity ($\hat{j} = 14$) at $n = 500$, the CV-selected model ($\hat{j} = 1$) at $n = 1000$, and highly undersmoothed models ($\hat{j} = 20$) at $n \geq 1500$. This adaptation is most consequential at $n = 500$, where A-TMLE Projection achieves lower MSE (0.013 vs. 0.017–0.019) and higher nonparametric coverage (95.6% vs. 88.6–89.5%) than both fixed baselines by selecting a working model that controls variance without sacrificing efficiency. Across all sample sizes, projection-based targeting within the A-TMLE framework yields the lowest bias and highest coverage among the three targeting approaches.

Positivity violations (DGP 2). Table 3.6 demonstrates the particular value of adaptive model selection under practical positivity violations. The fixed baselines maintain reasonable performance (coverage 92–95%), but A-TMLE with projection-based targeting is consistently competitive and often superior. At $n = 1000$, A-TMLE Projection achieves markedly lower bias than both the baselines (0.015 vs. 0.031–0.045) and the other A-TMLE variants (0.172–0.184), with correspondingly lower MSE (0.030 vs. 0.033–0.035). Similar advantages persist at $n = 1500$ and $n = 2000$. At $n = 500$, where all methods exhibit higher variance ($SE \approx 0.24$), the plateau selector conservatively retains the CV-selected model ($\hat{j} = 1$), yielding comparable coverage (93.2%) at lower MSE (0.042 vs. 0.056–0.059).

Within the A-TMLE framework, projection-based targeting substantially outperforms both alternatives under positivity violations, with 73–92% bias reductions relative to relaxed targeting across sample sizes and consistently superior coverage (88–94% vs. 79–91%). Hessian-based targeting exhibits numerical instability at $n = 1500$ ($SE = 0.498$, $MSE = 0.253$), further underscoring the practical importance of projection-based regularization. The key advantage of the plateau selector is that it achieves these gains *without requiring the analyst to pre-specify* the working model complexity.

n	Method	Bias	SE	MSE	Cov _{NP} ^a	$\hat{j}$ ^b	Cov _{Orc} ^a
500	A-TMLE (Relaxed)	0.027	0.112	0.0131	93.0 (0.435)	14	93.9 (0.438)
	A-TMLE (Hessian)	0.053	0.115	0.0159	92.2 (0.445)	14	92.2 (0.451)
	A-TMLE (Projection)	0.021	0.111	0.0127	95.6 (0.445)	14	96.5 (0.436)
	CV regHAL (Proj)	0.006	0.131	0.0171	89.5 (0.445)	–	94.5 (0.513)
	US regHAL (Proj)	0.010	0.136	0.0186	88.6 (0.445)	–	94.6 (0.534)
1000	A-TMLE (Relaxed)	0.005	0.086	0.0073	92.4 (0.312)	1	94.5 (0.336)
	A-TMLE (Hessian)	0.031	0.087	0.0084	92.4 (0.315)	1	94.5 (0.339)
	A-TMLE (Projection)	0.012	0.085	0.0072	92.4 (0.315)	1	93.1 (0.331)
	CV regHAL (Proj)	0.007	0.078	0.0060	95.0 (0.314)	–	94.5 (0.304)
	US regHAL (Proj)	0.017	0.081	0.0068	94.3 (0.315)	–	94.3 (0.319)
1500	A-TMLE (Relaxed)	0.003	0.060	0.0036	96.6 (0.249)	20	95.4 (0.235)
	A-TMLE (Hessian)	0.024	0.061	0.0042	95.4 (0.257)	20	93.1 (0.238)
	A-TMLE (Projection)	0.004	0.058	0.0033	97.7 (0.257)	20	95.3 (0.227)
	CV regHAL (Proj)	0.003	0.066	0.0043	97.5 (0.258)	–	97.5 (0.258)
	US regHAL (Proj)	0.003	0.066	0.0043	97.5 (0.258)	–	97.0 (0.259)
2000	A-TMLE (Relaxed)	0.001	0.058	0.0034	92.6 (0.215)	20	95.8 (0.229)
	A-TMLE (Hessian)	0.018	0.058	0.0037	95.3 (0.222)	20	95.3 (0.227)
	A-TMLE (Projection)	0.007	0.058	0.0034	95.2 (0.222)	20	95.7 (0.227)
	CV regHAL (Proj)	0.009	0.057	0.0033	96.5 (0.222)	–	96.5 (0.223)
	US regHAL (Proj)	0.009	0.057	0.0033	96.5 (0.222)	–	96.5 (0.223)

Note: Based on 200 Monte Carlo replications per sample size.

^a Coverage (%) of nominal 95% confidence intervals; mean interval width in parentheses. Cov_{NP}: nonparametric EIC variance estimator; Cov_{Orc}: oracle (true empirical) variance.

^b Mode of the plateau-selected working model index j across replications ($j=1$: CV-selected; $j=20$: most undersmoothed).

“–” indicates fixed working model methods without adaptive selection.

Table 3.5: Simulation results for regHAL-ATMLE with plateau-selected working model complexity under standard conditions (DGP 1). A-TMLE methods adaptively select j via the plateau selector; CV regHAL and US regHAL use fixed CV-selected and undersmoothed working models, respectively, both with projection-based targeting.

n	Method	Bias	SE	MSE	Cov _{NP} ^a	$\hat{j}$ ^b	Cov _{Orc} ^a
500	A-TMLE (Relaxed)	0.193	0.161	0.0630	90.6 (1.017)	1	77.7 (0.631)
	A-TMLE (Hessian)	0.195	0.209	0.0814	94.1 (1.094)	1	91.5 (0.819)
	A-TMLE (Projection)	0.053	0.199	0.0422	93.2 (0.974)	1	94.6 (0.781)
	CV regHAL (Proj)	0.042	0.240	0.0591	92.5 (0.958)	–	95.0 (0.941)
	US regHAL (Proj)	0.034	0.235	0.0561	92.4 (0.967)	–	94.7 (0.921)
1000	A-TMLE (Relaxed)	0.184	0.136	0.0522	79.3 (0.681)	1	72.3 (0.531)
	A-TMLE (Hessian)	0.172	0.178	0.0610	82.2 (0.722)	1	88.5 (0.698)
	A-TMLE (Projection)	0.015	0.173	0.0299	88.1 (0.690)	1	93.7 (0.678)
	CV regHAL (Proj)	0.031	0.180	0.0330	95.5 (0.713)	–	96.0 (0.704)
	US regHAL (Proj)	0.045	0.182	0.0349	95.4 (0.730)	–	95.4 (0.715)
1500	A-TMLE (Relaxed)	0.124	0.136	0.0338	84.5 (0.558)	1	86.6 (0.534)
	A-TMLE (Hessian)	0.081	0.498	0.2530	84.2 (0.615)	1	97.8 (1.951)
	A-TMLE (Projection)	0.023	0.144	0.0211	94.3 (0.578)	20	94.3 (0.564)
	CV regHAL (Proj)	0.030	0.148	0.0227	93.0 (0.593)	–	94.0 (0.580)
	US regHAL (Proj)	0.035	0.144	0.0219	93.0 (0.582)	–	94.0 (0.565)
2000	A-TMLE (Relaxed)	0.033	0.132	0.0184	84.6 (0.428)	20	95.4 (0.517)
	A-TMLE (Hessian)	0.027	0.245	0.0603	91.9 (0.495)	20	95.9 (0.959)
	A-TMLE (Projection)	0.021	0.123	0.0155	93.2 (0.474)	20	94.2 (0.483)
	CV regHAL (Proj)	0.027	0.118	0.0145	95.0 (0.479)	–	94.5 (0.462)
	US regHAL (Proj)	0.025	0.119	0.0147	94.0 (0.474)	–	95.0 (0.467)

Note: Based on 200 Monte Carlo replications per sample size.
All column definitions are identical to Table 3.5.

Table 3.6: Simulation results for regHAL-ATMLE with plateau-selected working model complexity under practical positivity violations (DGP 2).

3.6 regHAL-ATMLE for Marginal Survival Curves

We demonstrate regHAL-ATMLE’s scalability and robustness through marginal survival curve estimation from right-censored data. This covariate-free setting is chosen deliberately as a proof-of-concept where the Kaplan-Meier (KM) estimator—the semiparametrically efficient nonparametric maximum likelihood estimator—provides a gold-standard benchmark for comparison. By validating that our hazard-based HAL targeting machinery matches KM’s statistical efficiency in this simple setting, we establish algorithmic credibility before deployment in the more complex settings that motivate our methodology: survival curve estimation with baseline covariates, informative censoring conditional on covariates, and treatment-specific survival curves under causal frameworks—settings where KM is inapplicable and HAL-based hazard models with regularized targeting become essential.

Evaluating a marginal survival curve over a dense time grid using hazard-based HAL models also serves as an excellent computational stress test, isolating and illustrating the

algorithmic stability of high-dimensional targeting. Standard TMLE methods require recursive clever covariate integrations across time bins, causing the computational burden to scale exponentially and risking numerical fragility. Our method solves this gracefully within the working model without recursive Monte Carlo integrations.

3.6.1 Right Censored Data Set up

Suppose we have $O = (\tilde{T} = \min(T, C), \Delta = \mathbf{1}(T \leq C))$, with realization $o_i = (\tilde{t}_i, \delta_i)$, where $T \in (0, 1)$ and $T \perp C$ under the data generating process described in Appendix A.3.2. Let λ_0^T denote the failure time hazard.

Our target parameters of interest are $S_0(t) = P_0(T > t) = e^{-\int_0^t \lambda_0^T(s)ds}$ for a grid of t .

We estimate the hazard for T using HAL based on minimizing the penalized negative log likelihood:

$$\hat{\beta} = \underset{\beta}{\operatorname{argmin}} -L_{\beta}(O_i) + \lambda \|\beta\|_1 \quad (3.2)$$

$$= \underset{\beta}{\operatorname{argmin}} \int_0^{\tilde{t}_i} \lambda_{\beta}^T(s)ds - \delta_i \log(\lambda_{\beta}^T(\tilde{t}_i)) + \lambda \|\beta\|_1, \quad (3.3)$$

where $\lambda^T(s) = e^{\beta_0^{(0)} + \sum_j \beta_j^{(0)} * \phi_j(s)}$ and $\phi_j(s)$ are zero-order HAL basis taking the form of $\mathbf{1}(s \geq q_j)$. This HAL fit induces a parametric working model $\{\lambda_{\beta}^T(s) = e^{\beta_0 + \sum_j \beta_j * \phi_j(s)} : \beta\}$ for λ_0^T . Implementation of the regularized targeting steps and statistical details appear in Appendix A.3.

3.6.2 Simulation Results

We employ intensity-based HAL models using Poisson regression. Failure time hazard $\lambda^T(t)$ uses 100 percentiles of observed event times as knot points; censoring hazard $\lambda^C(t)$ uses 50 percentiles for efficiency. Undersmoothed models multiply the L_1 constraint by 1.61 to systematically expand working model complexity.

For survival curves evaluated at time grid G , we employ universal least favorable path updates (Algorithm 11, Appendix A.1) with step size 10^{-3} for both targeting approaches.

Tables 3.7–3.8 present results using 200 replications across $n \in \{500, 1000, 1500, 2000\}$. We include the Kaplan-Meier estimator as the semiparametrically efficient benchmark, since it is the NPMLE in this covariate-free, independent-censoring setting.

KM benchmark: The Kaplan-Meier estimator achieves absolute bias 0.007–0.014 (decreasing with n), with oracle coverage consistently near the nominal 95% level (94.9–95.2%), confirming its role as the efficiency bound in this setting.

Relaxed HAL failure: Standard relaxed targeting indiscriminately targets all directions in the working model, becoming highly unstable. It exhibits severe bias persistence (absolute bias 0.30–0.44) and MSE above 0.13 across all sample sizes, with confidence interval coverage of only 3–23%.

Regularized methods match KM: Both projection-based and Hessian-based targeting achieve performance essentially indistinguishable from the KM benchmark (absolute bias 0.007–0.014, MSE 0.0001–0.0004), representing $\approx 98\%$ bias reduction relative to relaxed HAL. Coverage using regularized EIC constructions reaches 87–93%, with oracle coverage 91–93%—close to KM’s 95%. This demonstrates that our regularized targeting within HAL working models recovers the full semiparametric efficiency of the problem, validating the approach for more complex settings where KM is unavailable.

Estimator	n	Targeting	Bias	SE	MSE	Cov _{NP} ^a	Cov _{Proj} ^a	Cov _{PCV} ^a	Cov _{Hess} ^a	Cov _{Orc} ^a
Intensity-HAL-TMLE	500	Kaplan-Meier	0.014	0.018	0.0004	—	—	—	—	94.9 (0.069)
	500	Relaxed	0.359	0.055	0.1895	—	23.8 (0.150)	23.7 (0.148)	6.3 (0.113)	16.0 (0.216)
	500	Projection	0.014	0.017	0.0004	—	86.9 (0.063)	67.8 (0.050)	90.3 (0.071)	93.2 (0.065)
	500	Hessian-based	0.014	0.017	0.0004	—	87.6 (0.065)	68.7 (0.050)	91.1 (0.070)	93.3 (0.066)
	1000	Kaplan-Meier	0.010	0.012	0.0002	—	—	—	—	95.2 (0.048)
	1000	Relaxed	0.402	0.047	0.2391	—	20.2 (0.111)	20.2 (0.110)	4.1 (0.056)	12.8 (0.184)
	1000	Projection	0.010	0.012	0.0002	—	88.5 (0.045)	72.7 (0.038)	92.0 (0.050)	92.8 (0.046)
	1000	Hessian-based	0.010	0.012	0.0002	—	88.9 (0.046)	73.0 (0.038)	92.9 (0.050)	93.2 (0.046)
	1500	Kaplan-Meier	0.008	0.010	0.0001	—	—	—	—	95.1 (0.041)
	1500	Relaxed	0.421	0.056	0.2639	—	18.3 (0.092)	18.3 (0.091)	3.9 (0.040)	9.9 (0.221)
	1500	Projection	0.008	0.010	0.0001	—	87.5 (0.037)	73.7 (0.032)	91.8 (0.041)	91.9 (0.039)
	1500	Hessian-based	0.009	0.010	0.0001	—	87.5 (0.038)	73.8 (0.032)	92.2 (0.041)	92.3 (0.040)
	2000	Kaplan-Meier	0.007	0.009	0.0001	—	—	—	—	94.9 (0.035)
	2000	Relaxed	0.438	0.033	0.2855	—	17.1 (0.080)	17.1 (0.080)	4.3 (0.029)	11.1 (0.130)
	2000	Projection	0.007	0.008	0.0001	—	88.4 (0.033)	76.1 (0.028)	92.4 (0.036)	91.1 (0.033)
	2000	Hessian-based	0.007	0.009	0.0001	—	88.3 (0.033)	76.0 (0.028)	93.1 (0.036)	92.0 (0.034)

Note: Based on 200 Monte Carlo replications per sample size.

All column definitions are identical to Table 3.1.

Kaplan-Meier serves as the semiparametrically efficient benchmark (NPMLE)

in this covariate-free setting; “—” indicates that model-based CI constructions are not applicable.

Table 3.7: Simulation results for intensity-based regHAL-ATMLE for marginal survival curve estimation with CV-selected working model. Kaplan-Meier is included as the efficient non-parametric benchmark.

Estimator	n	Targeting	Bias	SE	MSE	Cov _{NP} ^a	Cov _{Proj} ^a	Cov _{PCV} ^a	Cov _{Hess} ^a	Cov _{Orc} ^a
Intensity-HAL-TMLE	500	Kaplan-Meier	0.014	0.018	0.0004	—	—	—	—	94.9 (0.069)
	500	Relaxed	0.305	0.060	0.1372	—	24.8 (0.147)	24.5 (0.144)	9.7 (0.142)	14.4 (0.235)
	500	Projection	0.014	0.017	0.0004	—	86.2 (0.065)	66.5 (0.050)	90.7 (0.075)	92.2 (0.066)
	500	Hessian-based	0.015	0.018	0.0004	—	87.0 (0.068)	67.4 (0.050)	91.7 (0.074)	92.3 (0.069)
	1000	Kaplan-Meier	0.010	0.012	0.0002	—	—	—	—	95.2 (0.048)
	1000	Relaxed	0.374	0.037	0.2048	—	20.4 (0.111)	20.4 (0.110)	4.5 (0.066)	9.3 (0.145)
	1000	Projection	0.010	0.012	0.0002	—	87.5 (0.046)	71.1 (0.038)	91.6 (0.052)	92.1 (0.047)
	1000	Hessian-based	0.011	0.012	0.0002	—	87.8 (0.047)	71.5 (0.038)	92.6 (0.053)	92.5 (0.049)
	1500	Kaplan-Meier	0.008	0.010	0.0001	—	—	—	—	95.1 (0.041)
	1500	Relaxed	0.404	0.036	0.2413	—	18.5 (0.092)	18.5 (0.091)	4.3 (0.044)	8.0 (0.142)
	1500	Projection	0.009	0.010	0.0001	—	86.7 (0.038)	72.0 (0.032)	91.2 (0.043)	91.0 (0.039)
	1500	Hessian-based	0.009	0.011	0.0002	—	86.3 (0.038)	72.1 (0.032)	91.9 (0.043)	91.5 (0.042)
	2000	Kaplan-Meier	0.007	0.009	0.0001	—	—	—	—	94.9 (0.035)
	2000	Relaxed	0.423	0.030	0.2650	—	16.8 (0.080)	16.8 (0.080)	3.8 (0.030)	7.7 (0.117)
	2000	Projection	0.007	0.009	0.0001	—	87.6 (0.033)	74.4 (0.028)	91.8 (0.037)	90.8 (0.034)
	2000	Hessian-based	0.008	0.009	0.0001	—	87.3 (0.033)	74.2 (0.029)	93.1 (0.037)	91.7 (0.036)

Note: Based on 200 Monte Carlo replications per sample size.

All column definitions are identical to Table 3.7.

Table 3.8: Simulation results for intensity-based regHAL-ATMLE for marginal survival curve estimation with undersmoothed working model. Kaplan-Meier is included as the efficient nonparametric benchmark.

3.7 Discussion

This chapter develops two regularized A-TMLE implementations that stabilize the targeting step inside a HAL-induced working model: Hessian-based regHAL-ATMLE and projection-based regHAL-ATMLE. Both replace unstable exact adaptive targeting by regularized approximations to the working-model efficient influence curve, making the HAL setting the clearest place to see the dissertation’s estimator-side mechanism.

The empirical results show that projection-based targeting is the chapter’s main line. Across the treatment-effect and marginal-survival studies, it reduces bias substantially relative to relaxed HAL, remains stable under practical positivity stress, and recovers near-nominal confidence-interval coverage with modest computational overhead. The Hessian-based construction remains useful as a comparison because it clarifies what is gained by regularizing only the bias-relevant direction rather than the full information operator.

The substantive lesson is that the working model is not a hidden implementation detail. In HAL, it determines the score geometry in which the bias correction is approximated, regularized, and selected. The plateau selector complements this geometry by choosing a working-model size large enough to approximate the clever covariate while avoiding unnecessary variance inflation from collinear basis expansions.

Projection-based regHAL-ATMLE is the main line of the chapter and the methodological ancestor of the survival-specific targeting procedure studied later in Chapter 6. Hessian-based regHAL-ATMLE is retained as a useful contrast—valuable when the influence function is unclear or difficult to evaluate precisely, such as in complex longitudinal or network settings—and the empirical gap between the two constructions clarifies why regularizing only the bias-relevant direction is the principled choice.

3.7.1 Broader perspective

Although developed in the HAL setting, the main estimator-side lesson is broader. HAL is useful here because its basis expansion makes the score geometry visible, but the central idea is to regularize the bias-relevant correction inside a data-adaptive working model rather than insist on an exact and potentially unstable fluctuation. In that sense regHAL-ATMLE is best read as the clearest concrete instantiation of a wider regularized targeting program whose geometry is explicit enough to compare projection-based and Hessian-based corrections directly.

A natural next step is to separate the initial estimator from the targeting model even further by using a strong external learner, such as super learner, as a fixed offset and then carrying out the regularized correction in a HAL targeting model [Van der Laan et al., 2007]. This perspective also clarifies the local relation to collaborative TMLE: both approaches use the targeting step to organize regularization around the bias-relevant part of the problem, but regHAL-ATMLE does so through an explicit working-model score geometry rather than by iteratively rebuilding the nuisance model during targeting [van der Laan and Gruber, 2010, Ju et al., 2019]. Broader questions about offset-based regularized TMLE, cross-fitting, and adaptive inferential regularization are deferred to Chapter 7.

3.7.2 Limitations

The current implementation uses a small numerical floor for the projection penalty λ_n ; in more complex settings, regularization-parameter tuning may require more systematic attention, and theoretical guarantees for a fully data-driven penalty rule remain incomplete. Formal finite-sample theory for regularized EIC approximations in the inference stage is also incomplete. Extensions to longitudinal and time-varying treatments are natural next steps. *This chapter establishes the basic estimator-side mechanism of the dissertation in the most structured setting available. The plateau selector and the undersmoothing-based path of nested HAL working models are specific to that structured geometry. The underlying mechanism—treating the initial fit as an offset, choosing a working-model score space, and projecting onto the bias-relevant direction—is not. The next chapter asks whether that mechanism survives when the working model is no longer a sieve of explicit basis functions but an overparameterized neural network.*

Chapter 4

Regularized A-TMLE Beyond HAL: Neural Network Working Models

The previous chapter develops regularized targeting in the structured setting of HAL working models. This chapter asks whether that idea is specific to HAL’s sieve geometry or whether it is more properly understood as a consequence of score-space geometry in general. It answers that question by constructing regularized A-TMLE in neural-network working models, where the targeted update is organized through a selected subset of weights and the corresponding gradient span.

Although the survival application studied later returns to a HAL-based construction, the neural-network setting plays an important conceptual role in the dissertation. It shows that projection-guided targeting is not merely a repair for a particular sieve geometry. Rather, it is a more general way of organizing targeted updates in adaptive working models, including settings in which the score structure is nonlinear, high-dimensional, and overparameterized.

4.1 Related work and positioning

Chapter 2 gives the shared targeted-learning framework, and Chapter 3 gives its HAL instantiation. The present section records only the neural-network and causal-inference literatures needed to position TDA as the deep-learning instantiation of that same regularized A-TMLE program.

A first comparison class is the neural-network causal-inference literature built around representation learning and training-time regularization. Balanced-representation methods such as CFR and related architectures learn features that reduce covariate imbalance while preserving predictive performance [Johansson et al., 2016, Shalit et al., 2017]. DragonNet and related neural causal estimators push further by combining nuisance estimation with targeted regularization or orthogonal penalties inside training [Shi et al., 2019, Hatt et al., 2022, Nie et al., 2021]. These methods often improve empirical treatment-effect estimation, but their targeting enters as a penalty or auxiliary loss and therefore does not guarantee that the empirical efficient-influence-curve equation is solved.

A second comparison class keeps the targeting step post hoc. Classical TMLE, one-step TMLE, and later survival-targeting constructions provide the semiparametric template [van der Laan and Rubin, 2006, Van der Laan et al., 2011, Cai and van der Laan, 2020, Rytgaard and van der Laan, 2024]. More recent work has shown how to append debiasing or fluctuation steps to flexible learners, including neural-network-based Riesz or TMLE-style updates [Chernozhukov et al., 2022a, Shirakawa et al., 2024]. The difficulty for the present chapter is that once the target is high-dimensional or functional, explicit post-hoc fluctuations can become analytically cumbersome and computationally heavy.

TDA is positioned between these two lines. Like double machine learning and related orthogonalization approaches, it uses first-order bias correction to obtain debiased estimation [Chernozhukov et al., 2018, Künzel et al., 2019, Athey et al., 2018]. Like Chapter 3, however, it works inside a data-adaptive working model and organizes the correction through a regularized approximation to the working-model EIC [van der Laan et al., 2023, 2025]. The difference from the HAL case is geometric rather than conceptual: the score space is now generated by selected neural-network gradients rather than explicit basis vectors.

Seen from this literature, TDA’s specific contribution is to carry the Chapter 3 projection-targeting logic into overparameterized neural-network working models. The chapter therefore does not argue that deep learning replaces targeted learning, or that training-time regularization is sufficient on its own. It argues that neural architectures can serve as the working-model geometry inside which a post-training regularized A-TMLE step is constructed.

4.2 Targeting inside neural-network working models

We work within the semiparametric notation of Chapter 2. The target parameter $\Psi(P_0)$ depends on a loss-based nuisance component $Q(P)$ and, when needed, additional nuisances such as treatment or censoring mechanisms. TDA keeps the targeted-learning logic but replaces an explicit least favorable fluctuation with a regularized projection of an influence-curve representation onto neural-network score directions.

This section therefore emphasizes the pieces that are specific to the neural-network working model: the offset construction, the choice of targeting subset, the regularized score-space projection, and the multi-parameter update used for functional targets. General definitions of pathwise differentiability and efficient influence curves are taken as given from Chapter 2.

4.2.1 The Neural Network as a Parametric Working Model

Working model and offset structure. Let $f_\theta(x)$ be our neural net with parameters $\theta \in \Theta$. Once trained, the full neural network defines a big parametric working model $\mathcal{M}_{\text{full}} = \{P_\theta : \theta \in \Theta\}$ —a finite-dimensional submodel of the full model $\mathcal{M}$. The trained network $f_{\hat{\theta}}$ serves as the *offset*, with the initial fit $\hat{\theta}$ having nearly all non-zero coefficients (in contrast to HAL’s sparse cross-validated fit).

We then choose a targeting working model within $\mathcal{M}_{\text{full}}$ by partitioning:

$$\theta = (\theta_{\text{fix}}, \theta_{\text{targ}}),$$

where θ_{fix} is kept at its trained value $\hat{\theta}_{\text{fix}}$ and $\theta_{\text{targ}} \in \mathbb{R}^k$ is free to vary from its initial value $\hat{\theta}_{\text{targ}}$. This defines the targeting working model:

$$\mathcal{M}_n = \left\{ P_{\hat{\theta}_{\text{fix}}, \theta_{\text{targ}}} \mid \theta_{\text{targ}} \in \Theta_{\text{targ}} \subseteq \mathbb{R}^k \right\} \subseteq \mathcal{M}_{\text{full}}.$$

The initial fit resides in $\mathcal{M}_n$ at $\theta_{\text{targ}} = \hat{\theta}_{\text{targ}}$. This is the same structure as in regHAL: the big model is $\mathcal{M}_{\text{full}}$, the initial estimator $\hat{\theta}$ lies in it, and we target in a submodel $\mathcal{M}_n \subseteq \mathcal{M}_{\text{full}}$ defined by which parameters are allowed to change from their initial values. Targeting proceeds by fluctuating θ_{targ} away from $\hat{\theta}_{\text{targ}}$ via targeting gradient descent.

As in Chapter 3, we define the *loss-based dissimilarity projection* of $Q(P)$ onto the working model:

$$\Pi_{\mathcal{M}_n} Q(P) = \arg \min_{Q \in \mathcal{Q}_n} P L(Q),$$

where $\mathcal{Q}_n = \{Q_{\hat{\theta}_{\text{fix}}, \theta_{\text{targ}}} : \theta_{\text{targ}} \in \Theta_{\text{targ}}\}$ is the set of outcome functions reachable within the targeting working model. The *projected target parameter* is then:

$$\Psi_{\mathcal{M}_n}(P) = \Psi^F(\Pi_{\mathcal{M}_n} Q(P)).$$

We aim to make the empirical mean of the working-model EIC for $\Psi_{\mathcal{M}_n}$ approximate zero, removing first-order bias. As we show in Section 3.3, with appropriate choice of the targeting subset, this yields asymptotically efficient estimates for the original parameter $\Psi(P_0)$.

Nested working models via subset selection. By unfreezing different subsets of parameters, we obtain nested working models of increasing dimension:

$$\mathcal{M}_n^{(1)} \subset \mathcal{M}_n^{(2)} \subset \dots \subset \mathcal{M}_n^{(K)},$$

where, for example, $\mathcal{M}_n^{(1)}$ varies only the last output layer of each outcome head, $\mathcal{M}_n^{(2)}$ additionally unfreezes a fraction of the penultimate hidden layer, and $\mathcal{M}_n^{(K)}$ varies the full outcome-head subnetwork. Note that all parameters already exist with trained values—unlike HAL where augmentation adds new basis functions with initially-zero coefficients, here we simply allow more existing parameters to move from their trained values. The propensity head and shared trunk remain frozen throughout, serving as the offset. Each $\mathcal{M}_n^{(k)}$ induces a score subspace $\mathcal{T}^{(k)} \equiv \{\alpha^\top S_{\theta_{\text{targ}}^{(k)}}(O) : \alpha \in \mathbb{R}^{|\theta_{\text{targ}}^{(k)}|}\} \subseteq \mathcal{T}_{\mathcal{M}_{\text{full}}}$, where $\theta_{\text{targ}}^{(k)}$ denotes the targeting parameters of $\mathcal{M}_n^{(k)}$; the regularized projection (4.3) applied within each subspace yields a sequence of increasingly accurate EIC approximations, selected via the plateau criterion (Appendix B.3.1). The nested structure serves as a *regularization mechanism* for the targeting step (finite-sample bias–variance tradeoff), since the full network is typically already expressive enough (Section 4.2.3).

The key difference from HAL is that TDA is *nonlinear* in θ_{targ} —the frozen weights $\hat{\theta}_{\text{fix}}$ provide a rich nonlinear feature map—while HAL is linear in β . Because the procedure only requires access to $S_{\theta_{\text{targ}}}(O_i) = -\nabla_{\theta_{\text{targ}}} \ell(\theta, O_i)$, it applies to any architecture (CNNs, RNNs, transformers, etc.).

Example: ATE within the neural network working model. To make the working model concrete, consider the ATE setting with $O = (X, A, Y)$ and the outcome regression $Q(P) = \mathbb{E}_P[Y \mid A, X]$. The neural network models this as $Q_{\hat{\theta}_{\text{fix}}, \theta_{\text{targ}}}(A, X) = f_{\hat{\theta}_{\text{fix}}, \theta_{\text{targ}}}(A, X)$, with loss function $L(Q)(O) = (Y - Q(A, X))^2$ (squared error) or $L(Q)(O) = -\log p(Y \mid Q(A, X))$ (negative log-likelihood). For example, under a Gaussian likelihood:

$$Y \mid A, X \sim \mathcal{N}(f_{\hat{\theta}_{\text{fix}}, \theta_{\text{targ}}}(A, X), \sigma^2).$$

In either case, $\mathcal{Q}_n = \{f_{\hat{\theta}_{\text{fix}}, \theta_{\text{targ}}}(\cdot) : \theta_{\text{targ}} \in \Theta_{\text{targ}}\}$ defines the set of outcome regressions reachable by varying the targeting subset, and the projected target parameter is

$$\Psi_{\mathcal{M}_n}(P) = \Psi^F(\Pi_{\mathcal{M}_n} Q(P)) = \mathbb{E}_{P_{X,n}}[f_{\hat{\theta}_{\text{fix}}, \theta_{\text{targ}}^*}(1, X) - f_{\hat{\theta}_{\text{fix}}, \theta_{\text{targ}}^*}(0, X)],$$

where $\theta_{\text{targ}}^* = \arg \min_{\theta_{\text{targ}}} P L(Q_{\hat{\theta}_{\text{fix}}, \theta_{\text{targ}}})$ realizes the loss-based projection $\Pi_{\mathcal{M}_n}$.

4.2.2 Regularized Projection-Based Targeting

Following the reg-ATMLE principle, and in particular the projection-based construction developed in Chapter 3, TDA targets the projected parameter $\Psi_{\mathcal{M}_n}(P_0)$ within the neural network working model $\mathcal{M}_n$. To motivate our projection-based approach, we first examine the structure of the working-model EIC.

Connection to the Parametric EIC in the Neural Network Working Model. To understand what the projection step targets mathematically, we first define the true parametric EIC for the projected parameter $\Psi_{\mathcal{M}_n}$ within the neural network working model. Let $\theta_{0, \mathcal{M}_n}^* = \arg \min_{\theta_{\text{targ}}} P_0 \ell(Q_{\hat{\theta}_{\text{fix}}, \theta_{\text{targ}}}, O)$ denote the loss-based projection parameter—the value of θ_{targ} at which $Q_{\hat{\theta}_{\text{fix}}, \theta_{\text{targ}}}$ best approximates Q_0 within the working model under P_0 . Let $S_{\theta_{\text{targ}}}(O) = -\nabla_{\theta_{\text{targ}}} \ell(\theta, O)$ denote the score vector (negative gradient of the loss with respect to θ_{targ}). The true parametric EIC for the projected parameter under P_0 involves the population Hessian $I_0(\theta_{0, \mathcal{M}_n}^*) = P_0[\nabla_{\theta_{\text{targ}}}^2 \ell(\theta, O)]$. In practice, we approximate it by replacing population quantities with their empirical counterparts at the current estimate $\hat{\theta}_{\text{targ}}$. The *Hessian-based* approximation is:

$$\hat{D}_{\text{hessian}}^*(O_i) = \nabla_{\theta_{\text{targ}}} \Psi^F(Q_{\theta_{\text{targ}}}) \Big|_{\hat{\theta}_{\text{targ}}}^\top I_n^{-1}(\hat{\theta}_{\text{targ}}) S_{\hat{\theta}_{\text{targ}}}(O_i), \quad (4.1)$$

where $I_n(\theta_{\text{targ}}) = P_n[\nabla_{\theta_{\text{targ}}}^2 \ell(\theta, O)]$ is the empirical Hessian—the natural empirical analog of the population Hessian I_0 that defines the true working-model EIC.

When we perform the unregularized $L_2(P_n)$ projection of a nonparametric influence function $D(O)$ onto the network’s score space, the resulting coefficients are $\alpha^* = (P_n[S_{\theta_{\text{targ}}} S_{\theta_{\text{targ}}}^\top])^{-1} P_n[S_{\theta_{\text{targ}}} D]$. By the *empirical analogue* of the pathwise derivative identity, $P_n[S_{\theta_{\text{targ}}} D] \approx \nabla_{\theta_{\text{targ}}} \Psi^F(Q_{\theta_{\text{targ}}})$. Thus, the projection-based working-model EIC approxima-

tion evaluated at the current estimate $\hat{\theta}_{\text{targ}}$ under P_n is:

$$\tilde{D}_{\text{proj}}^*(O_i) \approx \nabla_{\theta_{\text{targ}}} \Psi^F(Q_{\theta_{\text{targ}}}) \Big|_{\hat{\theta}_{\text{targ}}}^\top I_{\text{cov},n}^{-1}(\hat{\theta}_{\text{targ}}) S_{\hat{\theta}_{\text{targ}}}(O_i), \quad \text{where } I_{\text{cov},n}(\theta_{\text{targ}}) = P_n[S_{\theta_{\text{targ}}} S_{\theta_{\text{targ}}}^\top]. \quad (4.2)$$

Exactly as in the regHAL framework developed in Chapter 3, both approximations share the identical structure $\nabla_{\theta_{\text{targ}}} \Psi^F(Q_{\theta_{\text{targ}}})^\top \cdot I^{-1} \cdot S_{\theta_{\text{targ}}}$, evaluated at the current estimate $\hat{\theta}_{\text{targ}}$ under P_n , and **differ only in the choice of information matrix**: the Hessian method uses $I_n = P_n[\nabla^2 \ell]$, while the projection method uses the score covariance $I_{\text{cov},n} = P_n[S_{\theta_{\text{targ}}} S_{\theta_{\text{targ}}}^\top]$. The Hessian-based method directly estimates the exact EIC of the projection parameter $\Psi_{\mathcal{M}_n}$ —even under misspecification—because I_n is the natural empirical analog of the population Hessian $I_0(\theta_{0,\mathcal{M}_n}^*) = P_0[\nabla^2 \ell]$ that defines the true working-model EIC.

Remark 4.1 (Hessian vs. score covariance and the role of the loss function). For likelihood-based losses, when the working model $\mathcal{M}_n$ is correctly specified (i.e., contains P_0), Bartlett’s second identity holds: $E_{P_0}[\nabla^2 \ell] = E_{P_0}[SS^\top]$, and the Hessian and score covariance coincide at P_0 . Thus, as the neural network working model grows toward an oracle model $\mathcal{M}_0 \subseteq \mathcal{M}$ containing P_0 , both approximations converge to the same EIC of the projection parameter $D_{\mathcal{M}_0}^*(P_0)$, and if $\mathcal{M}_0 = \mathcal{M}$, both converge to the full nonparametric EIC $D^*(P_0)$.

However, for **non-likelihood losses** such as squared error, the two EIC approximations are **fundamentally different**, not merely different regularizations of the same object. The expected Hessian depends only on covariates (e.g., $E_P[2\phi\phi^\top]$), while the score outer-product includes residual variance weights (e.g., $E_P[4(Y - Q)^2\phi\phi^\top]$); these are proportional only under homoscedasticity. Crucially, this distinction persists even when the working model contains P_0 : for any oracle model $\mathcal{M}_0 \subsetneq \mathcal{M}$ that is a proper submodel of the full model, the two methods target genuinely different objects. Asymptotic equivalence under non-likelihood losses requires $\mathcal{M}_n \rightarrow \mathcal{M}$ (the full model), at which point both converge to the nonparametric EIC $D^*(P_0)$.

This distinction is inherited directly from the analysis of Chapter 3 and applies identically to the neural network setting.

To combat the inherent collinearity and potential singularity of the score covariance matrix in overparameterized neural networks, we impose L_2 regularization (ridge) during the projection, obtaining our stable targeting gradient.

Regularized projection onto the score space. Given a (possibly uncentered) influence-curve estimate $\tilde{D}_{\theta_{\text{targ}}}(O)$ —e.g., from AIPW, IPTW, or the EIC itself—we project it onto the network’s score space with L_2 regularization and an unpenalized intercept γ_0 for empirical centering:

$$(\gamma_0^*, \alpha_{\theta_{\text{targ}}}^*) = \arg \min_{\gamma_0, \alpha} \sum_{i=1}^n \left[\tilde{D}_{\theta_{\text{targ}}}(O_i) - \gamma_0 - \alpha^\top S_{\theta_{\text{targ}}}(O_i) \right]^2 + \lambda \|\alpha\|_2^2. \quad (4.3)$$

The intercept γ_0 centers the gradient and scores, matching the unpenalized projection intercept used in Chapter 3. This yields the regularized projected EIC approximation:

$$\tilde{D}_{\text{proj}}^*(O) = \alpha_{\theta_{\text{targ}}}^{*\top} S_{\theta_{\text{targ}}}(O).$$

Remark on regularization: While our asymptotic theory only requires that the regularized projection EIC eventually approximates the true working-model EIC well enough (Assumption 4.3), finite-sample stability in highly overparameterized neural networks often necessitates a non-negligible penalty. In our experiments, we use a small but fixed λ , acting as a variance-stabilizing numerical floor, in the same spirit as Chapter 3’s regularized targeting construction (see Limitation *vii* for discussion of targeting hyperparameter sensitivity).

Iterative Update with Targeting Gradient. We denote $\alpha_{\theta_{\text{targ}}}^*$ as our *targeting gradient* at θ_{targ} for $\Psi(P_0)$. By repeatedly stepping along this gradient direction, we systematically reduce the empirical mean of the approximated EIC. This eliminates the need to derive an analytic least-favorable submodel or compute a closed-form parametric EIC inversion.

Crucially, because we are optimizing the working model’s empirical loss along the 1D direction parameterized by α^* , **any local stationary point along this path mathematically guarantees that the directional derivative is zero.** Since the directional derivative along the $\alpha_{\theta_{\text{targ}}}^*$ of the loss is exactly $\frac{1}{n} \sum \alpha^{*\top} \nabla_{\theta} \ell = -P_n[\tilde{D}_{\text{proj}}^*]$ (recalling $S = -\nabla_{\theta} \ell$), reaching *any* stationary point inherently ensures the empirical EIC equation is solved perfectly ($P_n[\tilde{D}_{\text{proj}}^*] = 0$), completely bypassing the non-convexity challenges typically associated with deep learning optimization.

In practice, we iteratively update θ_{targ} :

$$\theta_{\text{targ}}^{(t+1)} = \theta_{\text{targ}}^{(t)} + \gamma_t \text{sign}\left[P_n(\tilde{D}_{\text{proj}}^{*,(t)})\right] \alpha_{\theta_{\text{targ}}^{(t)}}^*,$$

or we use a simple 1D line search over γ_t . The use of $\text{sign}\left[P_n(\tilde{D}_{\text{proj}}^{*,(t)})\right]$ guarantees we are minimizing the loss as we update. Convergence is checked by verifying $\left|P_n[\tilde{D}_{\text{proj}}^*]\right| \leq \eta_n$ for the tolerance $\eta_n = \frac{sd(\tilde{D}_{\text{proj}}^*)}{\sqrt{n \log n}}$ as suggested in Rytgaard and van der Laan [2024] to guarantee that $P_n(\tilde{D}_{\text{proj}}^*) = o_p(n^{-\frac{1}{2}})$, or upon reaching a max tolerance with no further loss improvement.

Algorithm 3 outlines the full procedure.

4.2.3 Score Space Richness in Neural Network Working Models

A central condition for the reg-ATMLE framework is that the working model’s score space must be rich enough to well-approximate the targeted influence function direction. Both regHAL-ATMLE and TDA operate within a big parametric working model, and the quality of the EIC approximation depends on the richness of the available score space.

In regHAL-ATMLE, the cross-validated HAL fit may not retain enough basis directions to approximate the efficient influence curve well, so targeting proceeds inside a larger working model with additional score directions. TDA faces the same approximation problem in a

Algorithm 3 TDA Procedure (Projection-Based EIC Approximation)

- 1: **Input:** Data $\{O_i\}_{i=1}^n$, initial $\theta_{\text{targ}}^{(0)} = \hat{\theta}_{\text{targ}}$, penalty λ , step size γ_t or line search.
 - 2: **for** $t = 0, 1, 2, \dots$ **until convergence do**
 - 3: Compute $S_{\theta^{(t)}}(O_i) = -\nabla_{\theta_{\text{targ}}^{(t)}} \ell(f_{\theta^{(t)}}, O_i)$ for each i .
 - 4: Solve

$$(\gamma_0^{(t)}, \alpha^{(t)}) = \arg \min_{\gamma_0, \alpha} \sum_i \left[\tilde{D}(O_i) - \gamma_0 - \alpha^\top S_{\theta^{(t)}}(O_i) \right]^2 + \lambda \|\alpha\|_2^2.$$
 - 5: Define

$$\tilde{D}_{\text{proj}}^{*,(t)}(O_i) = \alpha^{(t)\top} S_{\theta^{(t)}}(O_i).$$
 - 6: Update

$$\theta_{\text{targ}}^{(t+1)} = \theta_{\text{targ}}^{(t)} + \gamma_t \text{sign} \left[P_n \left(\tilde{D}_{\text{proj}}^{*,(t)} \right) \right] \alpha^{(t)}.$$
 - 7: Check stopping: if $|P_n[\tilde{D}_{\text{proj}}^{*,(t)}]| \leq \eta_n$ or $t \geq T_{\text{max}}$, break.
 - 8: **end for**
 - 9: **Output:** $\hat{\theta}_{\text{targ}}^{(T)}$, and $\hat{\Psi} = \Psi(\hat{\theta}_{\text{fix}}, \hat{\theta}_{\text{targ}}^{(T)})$.
-

different geometry: the available score space is generated by neural-network gradients rather than HAL basis scores.

In TDA, the big working model is the full neural network $\mathcal{M}_{\text{full}}$. The neural network score space has several properties that aid EIC approximation:

1. **Rich score space from nonlinearity.** Neural network gradients capture complex interactions through the chain rule. Even last-layer gradients involve the full representation learned by earlier layers, making the score space much richer relative to the parameter count than HAL’s linear basis scores.
2. **Joint architecture.** When the propensity score is modeled as a separate head of the same network (as in DragonNet Shi et al. [2019]), the shared representation naturally captures variation relevant to both the outcome and the EIC direction, further aiding EIC approximation within the score space.

Formally, the quality of the working-model EIC approximation depends on how well the score space captures the relevant influence-function direction. We require that the regularized projection of $D(O)$ onto the score space $\{S_{\theta_{\text{targ}}}(\theta) = -\nabla_{\theta_{\text{targ}}} \ell(\theta, O)\}$ yields a sufficiently accurate approximation of the working-model EIC $D_{\mathcal{M}_n}^*$. Letting $P_{\mathcal{M}_n, n}^*$ denote the distribution induced by the converged targeting parameters (formalized in Assumption 4.3), we minimally require consistency under P_0 :

$$\inf_{\alpha \in \mathbb{R}^{|\theta_{\text{targ}}|}} \|D_{\mathcal{M}_n}^*(P_{\mathcal{M}_n, n}^*) - \gamma_0 - \alpha^\top S_{\theta_{\text{targ}}}(O)\|_{L^2(P_0)} = o_P(1). \quad (4.4)$$

To upgrade this consistency statement into the empirical score equation $P_n D_{\mathcal{M}_n}^*(P_{\mathcal{M}_n,n}^*) = o_P(n^{-1/2})$, the approximation error must vanish fast enough. The algorithm enforces $P_n \tilde{D}_{\text{proj}}^*(P_{\mathcal{M}_n,n}^*) = o_P(n^{-1/2})$, so it suffices that

$$P_n \{D_{\mathcal{M}_n}^*(P_{\mathcal{M}_n,n}^*) - \tilde{D}_{\text{proj}}^*(P_{\mathcal{M}_n,n}^*)\} = o_P(n^{-1/2}).$$

Writing $E(P) = D_{\mathcal{M}_n}^*(P) - \tilde{D}_{\text{proj}}^*(P)$, the term $P_n E(P_{\mathcal{M}_n,n}^*)$ decomposes into an empirical-process component and a drift component. Under (4.4), the empirical-process term is negligible for a Donsker approximation class, while the drift term is second order because $P_0 E(P_0) = 0$. A sufficient condition is therefore that the $L^2(P_0)$ approximation error of the projected score decays quickly enough that its rate, multiplied by the initial estimator's convergence rate, is $o_P(n^{-1/2})$.

For the full-network targeting subset, the high dimensionality and nonlinearity of the score space make it plausible that the projection closely approximates the working-model EIC. In practice, this can be assessed either by direct approximation diagnostics or indirectly through the plateau selector in Appendix B.3.1, which expands the targeting subset when narrower score spaces fail to stabilize the update. This remains an empirical property of the chosen architecture rather than an automatic guarantee; severely underparameterized networks can still violate (4.4).

Consequently, the nested working models $\mathcal{M}_n^{(1)} \subset \dots \subset \mathcal{M}_n^{(K)}$ serve primarily as a *regularization mechanism*: starting with a smaller targeting subset (e.g., last layer) controls variance during targeting, while the full-network model is available as a fallback.

Remark 4.2 (Super learner offset extension). In the current presentation, the initial neural-network fit $\hat{\theta}$ serves as the offset within its own big working model $\mathcal{M}_{\text{full}}$. However, the offset need not come from the deep-learning model itself. More generally, one can train a super learner $\hat{Q}_{\text{SL}}(x)$ (which may include HAL, random forests, gradient boosting, and related learners) and define the neural-network model as $Q_{\theta}(x) = \hat{Q}_{\text{SL}}(x) + f_{\theta}(x)$, with the network output f_{θ} initialized near zero. This is the neural analogue of regHAL's offset construction: the external learner carries the convergence rate, while the neural network supplies a rich score space for the targeting correction.

4.2.4 Multi-Parameter Targeting

Many causal estimands are multi-dimensional, such as a *survival curve* $\{S(t)\}_{t \in \mathcal{T}}$ or a vector of treatment effects under different doses. Suppose each parameter Ψ_k has an influence function $D_k(O)$. In TDA, we simply project each D_k separately and then *combine* the resulting directions to create a universal canonical one-dimensional submodel as in van der Laan and Gruber [2016]. Concretely:

$$(\gamma_{0,k}^*, \alpha_k^*) = \arg \min_{\gamma_0, \alpha} \sum_i \left[D_k(O_i) - \gamma_0 - \alpha^\top S_{\theta}(O_i) \right]^2 + \lambda \|\alpha\|_2^2, \quad d_k = \frac{1}{n} \sum_i \alpha_k^{*\top} S_{\theta}(O_i),$$

Algorithm 4 TDA Multi-Parameter Targeting Procedure (Universal Least Favorable Path)

- 1: **Input:** Data $\{O_i\}_{i=1}^n$, target parameters $\{\Psi_s\}_{s \in S}$ with influence functions $\{D_s(O)\}_{s \in S}$, initial $\theta_{\text{targ}}^{(0)} = \hat{\theta}_{\text{targ}}$, penalty λ , step size γ .
- 2: **for** $t = 0, 1, 2, \dots$ **until convergence do**
- 3: Compute scores: $S_{\theta^{(t)}}(O_i) = -\nabla_{\theta_{\text{targ}}^{(t)}} \ell(f_{\theta^{(t)}}, O_i)$ for each i .
- 4: **for** each target parameter $s \in S$ **do**
- 5: Solve the regularized projection:

$$(\gamma_{0,s}^{(t)}, \alpha_s^{(t)}) = \arg \min_{\gamma_0, \alpha} \sum_i \left[D_s(O_i) - \gamma_0 - \alpha^\top S_{\theta^{(t)}}(O_i) \right]^2 + \lambda \|\alpha\|_2^2.$$

- 6: Compute the empirical mean of the projected EIC:

$$d_s^{(t)} = \frac{1}{n} \sum_{i=1}^n \alpha_s^{(t)\top} S_{\theta^{(t)}}(O_i).$$

- 7: **end for**
- 8: Form the universal least favorable path direction:

$$\alpha^{(t)} = \sum_{s \in S} \frac{d_s^{(t)}}{\|d^{(t)}\|_2} \alpha_s^{(t)}, \quad \text{where } d^{(t)} = (d_s^{(t)})_{s \in S}.$$

- 9: Update:

$$\theta_{\text{targ}}^{(t+1)} = \theta_{\text{targ}}^{(t)} + \gamma \text{sign} \left[P_n \left(\alpha^{(t)\top} S_{\theta^{(t)}} \right) \right] \alpha^{(t)}.$$

- 10: Check stopping: if $\frac{\|d^{(t)}\|_2}{\sqrt{|S|}} \leq \frac{\text{Median}_s \left\{ \text{se} \left(\alpha_s^{(t)\top} S_{\theta^{(t)}}(O) \right) \right\}}{\sqrt{n} \log n}$ or $t \geq T_{\text{max}}$, break.

- 11: **end for**

- 12: **Output:** $\hat{\theta}_{\text{targ}}^{(T)}$, and $\{\hat{\Psi}_s = \Psi_s(\hat{\theta}_{\text{fix}}, \hat{\theta}_{\text{targ}}^{(T)})\}_{s \in S}$.
-

then aggregate via

$$w_k = \frac{d_k}{\sqrt{\sum_{j=1}^K d_j^2}}, \quad \alpha^* = \sum_{k=1}^K w_k \alpha_k^*.$$

A single step in direction α^* reduces the empirical mean of each D_k . Combining individual targeting gradients is particularly powerful for *functional* targets like survival curves, preserving coherence across time points while still pushing all marginal biases toward zero. Algorithm 4 formalizes this multi-parameter procedure.

4.3 Theoretical Guarantees for Targeted Deep Architectures

We establish that TDA inherits the statistical guarantees of the Adaptive TMLE (A-TMLE) framework van der Laan et al. [2023]. We present two complementary justifications: (1) a general Regularized TMLE theorem for the original target parameter $\Psi(P_0)$ that applies broadly, and (2) an A-TMLE theorem that additionally provides inference for the data-adaptive projection parameter $\Psi_{\mathcal{M}_n}(P_0)$. We focus on the case where the oracle model is the full model, i.e., $\mathcal{M}_0 = \mathcal{M}$; for likelihood-based losses, the A-TMLE framework further provides super-efficiency relative to a restricted oracle model $\mathcal{M}_0 \subsetneq \mathcal{M}$.

4.3.1 Setup and Working Model Construction

Let $\mathcal{M}$ denote the full (typically nonparametric) statistical model. Given a neural network f_θ with parameters $\theta = (\hat{\theta}_{\text{fix}}, \theta_{\text{targ}})$, TDA induces a parametric working model $\mathcal{M}_n = \{P_{\hat{\theta}_{\text{fix}}, \theta_{\text{targ}}} : \theta_{\text{targ}} \in \Theta_{\text{targ}}\} \subseteq \mathcal{M}_{\text{full}}$. We define the projected target parameter as $\Psi_{\mathcal{M}_n}(P) = \Psi^F(\Pi_{\mathcal{M}_n} Q(P))$, where $\Pi_{\mathcal{M}_n}$ is the loss-based dissimilarity projection onto $\mathcal{M}_n$. Let $D_{\mathcal{M}_n}^*(P)$ denote the efficient influence curve (EIC) of the projected parameter evaluated in the full model $\mathcal{M}$.

Assumption 4.3 (Approximation of the working model EIC). *Let $\tilde{D}_{\text{proj}}^*$ denote the regularized EIC estimate derived from the L_2 -regularized score-covariance projection used for TDA targeting. The targeted fit $P_{\mathcal{M}_n, n}^*$ solves the approximate score equation $P_n \tilde{D}_{\text{proj}}^*(P_{\mathcal{M}_n, n}^*) = o_p(n^{-1/2})$ by the algorithm's stopping criterion. We assume this implies that the true working-model EIC score equation is also solved at the same rate:*

$$P_n D_{\mathcal{M}_n}^*(P_{\mathcal{M}_n, n}^*) = o_p(n^{-1/2}).$$

Assumption 4.3 is the exact analogue of Assumption 1 in Chapter 3: we do not impose an explicit convergence rate on the regularization parameter λ , but instead assume directly that solving the approximate score equation with the regularized projection implies that the true working-model EIC equation is also solved. As discussed in Section 4.2.3, a sufficient condition is that the approximation error $E_n = D_{\mathcal{M}_n}^* - \tilde{D}_{\text{proj}}^*$ at $P_{\mathcal{M}_n, n}^*$ is consistent to zero and converges quickly enough that, multiplied by the initial estimator's convergence rate, it is $o_p(n^{-1/2})$.

4.3.2 Regularized TMLE Theorem

We first present the general Regularized TMLE theorem, which applies broadly to any estimator that approximately solves an approximation of the full-model EIC equation. This result does not require defining a projection parameter or working model—it applies to the original fixed target parameter $\Psi(P_0)$ directly.

Let $\tilde{D}^*$ denote a full-model EIC approximation in which the Q -dependent component $\tilde{D}_Q^*$ is approximated via regularized score projection (Section 4.2.2), and any remaining components of D^* not depending on Q are estimated by other means—for instance, by their empirical analogues. (In the ATE, $D_Q^* = H_0(A, X)(Y - Q(A, X))$ is approximated by the score projection, and $D_{\text{rem}}^* = Q_n(1, X) - Q_n(0, X) - \hat{\Psi}$ is simply solved by the sample average of the marginal terms.)

Theorem 4.4 (Regularized TMLE for the original target parameter). *Let $\tilde{D}^*$ be an approximation of the full-model EIC D^* . Let P_n^* denote the TDA fit that solves $P_n \tilde{D}^*(P_n^*) = o_p(n^{-1/2})$. Define the exact remainder $R(P, P_0) = \Psi(P) - \Psi(P_0) + P_0 D^*(P)$. Assume:*

(R.1) $D^*(P_n^*)$ and $D^*(P_0)$ lie in a P_0 -Donsker class with probability tending to 1.

(R.2) $P_0(D^*(P_n^*) - D^*(P_0))^2 = o_p(1)$.

(R.3) $R(P_n^*, P_0) = o_p(n^{-1/2})$.

(R.4) $P_n(D^*(P_n^*) - \tilde{D}^*(P_n^*)) = o_p(n^{-1/2})$.

Then,

$$\Psi(P_n^*) - \Psi(P_0) = P_n D^*(P_0) + o_p(n^{-1/2}),$$

and

$$\sqrt{n}(\Psi(P_n^*) - \Psi(P_0)) \xrightarrow{d} \mathcal{N}(0, \text{Var}_{P_0}(D^*(P_0))).$$

Remark 4.5 (Sufficient conditions for (R.4)). Condition (R.4) requires that the regularized score $\tilde{D}^*$ approximates the true EIC D^* well enough in an empirical-mean sense at the targeted fit. Decomposing $P_n(D^* - \tilde{D}^*) = (P_n - P_0)(D^* - \tilde{D}^*) + P_0(D^* - \tilde{D}^*)$: the empirical process part is $o_p(n^{-1/2})$ whenever $\|D^*(P_n^*) - \tilde{D}^*(P_n^*)\|_{L_2(P_0)} \rightarrow 0$ in probability and the function class is Donsker. For doubly-robust parameters such as the ATE, whose exact remainders are bounded by products of nuisance estimation errors, an $L_2(P_0)$ approximation rate of $o_p(n^{-1/4})$ for $\tilde{D}^* \rightarrow D^*$ is sufficient. More generally, the required approximation rate depends on the structure of the exact remainder $R(P, P_0)$ for the specific target parameter.

4.3.3 A-TMLE Theoretical Results

Because TDA is a specific algorithmic implementation of the general A-TMLE framework, its theoretical validity relies on the foundational A-TMLE theorems established by van der Laan et al. [2023, 2025]. The next results record those statements in the notation of this chapter and specialize them to the neural-network working model used by TDA.

The A-TMLE framework provides additional structure by working with the projection parameter $\Psi_{\mathcal{M}_n}(P_0)$. TDA solves $P_n D_{\mathcal{M}_n}^*(P_{\mathcal{M}_n, n}^*) = o_p(n^{-1/2})$ —a regularized version of the TMLE for the projection parameter. This yields inference for the data-adaptive target $\Psi_{\mathcal{M}_n}(P_0)$, and if $\mathcal{M}_n \rightarrow \mathcal{M}_0$ (with $\mathcal{M}_0 = \mathcal{M}$ in the general case), also for $\Psi(P_0)$.

Theorem 4.6 (TDA as TMLE for working model target parameter, adapted from van der Laan et al. [2023]). *Let $P_{\mathcal{M}_n,n}^*$ denote the TDA estimator satisfying Assumption 4.3. Define the exact remainder $R_{\mathcal{M}_n}(P, P_0) = \Psi_{\mathcal{M}_n}(P) - \Psi_{\mathcal{M}_n}(P_0) + P_0 D_{\mathcal{M}_n}^*(P)$. Assume:*

(A.1) $D_{\mathcal{M}_n}^*(P_{\mathcal{M}_n,n}^*)$ and $D_{\mathcal{M}_n}^*(P_0)$ lie in a P_0 -Donsker class with probability tending to 1.

(A.2) $P_0 \left(D_{\mathcal{M}_n}^*(P_{\mathcal{M}_n,n}^*) - D_{\mathcal{M}_n}^*(P_0) \right)^2 = o_p(1)$.

(A.3) $R_{\mathcal{M}_n}(P_{\mathcal{M}_n,n}^*, P_0) = o_p(n^{-1/2})$.

Then, $\Psi_{\mathcal{M}_n}(P_{\mathcal{M}_n,n}^*) - \Psi_{\mathcal{M}_n}(P_0) = P_n D_{\mathcal{M}_n}^*(P_0) + o_p(n^{-1/2})$.

Theorem 4.7 (TDA as TMLE for original target parameter, adapted from van der Laan et al. [2023]). *Assume the oracle model is the full model $\mathcal{M}_0 = \mathcal{M}$ (so $D_{\mathcal{M}_0}^*(P_0) = D^*(P_0)$). Let $P_{0,\mathcal{M}_n}$ be the KL projection of P_0 onto $\mathcal{M}_n$. Assume:*

(B.1) $\Psi_{\mathcal{M}_n}(P_0) - \Psi(P_0) = o_p(n^{-1/2})$ (the projection parameter approximates the original target at second order).

(B.2) $D_{\mathcal{M}_n}^*(P_0) \rightarrow D^*(P_0)$ in $L_2(P_0)$, i.e., $P_0 \left(D_{\mathcal{M}_n}^*(P_0) - D^*(P_0) \right)^2 = o_p(1)$.

and

(C.1) $R(P_{0,\mathcal{M}_n}, P_0) = o_p(n^{-1/2})$.

(C.2) $(P_{0,\mathcal{M}_n} - P_0) \left(D^*(P_{0,\mathcal{M}_n}) - D_{\mathcal{M}_n}^*(P_{0,\mathcal{M}_n}) \right) = o_p(n^{-1/2})$.

(C.3) $P_0 D_{\mathcal{M}_n}^*(P_{0,\mathcal{M}_n}) = o_p(n^{-1/2})$.

Note that under $\mathcal{M}_0 = \mathcal{M}$, condition (B.1) holds trivially since $\Psi(P_{0,\mathcal{M}_n}) - \Psi_{\mathcal{M}_0}(P_{0,\mathcal{M}_n})$ is an identity, and (B.2) holds as $\mathcal{M}_n \rightarrow \mathcal{M}$. Then, $\Psi(P_{\mathcal{M}_n,n}^*) - \Psi(P_0) = P_n D^*(P_0) + o_p(n^{-1/2})$, and TDA achieves the semiparametric efficiency bound in $\mathcal{M}$.

TDA estimators are A-TMLEs that can also satisfy Theorem 4.4's conditions, providing a dual theoretical justification in such cases. When the oracle model is the full model ($\mathcal{M}_0 = \mathcal{M}$), the Regularized TMLE theorem (Theorem 4.4) directly applies. The A-TMLE framework (Theorems 4.6–4.7) additionally provides inference for the data-adaptive projection parameter $\Psi_{\mathcal{M}_n}(P_0)$ and, for likelihood-based losses, formal super-efficiency guarantees relative to an oracle model $\mathcal{M}_0 \subsetneq \mathcal{M}$, by changing the $D^*(P_0)$ to $D_{\mathcal{M}_0}^*(P_0)$ in theorem 4.7.

4.3.4 Verification of Regularity Conditions for Neural Networks

Theorem 4.8 (TDA satisfies regularity conditions). *Suppose the neural network working model $\mathcal{M}_n$ is sufficiently rich such that:*

1. The initial estimator satisfies $\|Q_{0,\mathcal{M}_n} - Q_0\|_{L_2(P_0)} = O_p(n^{-r})$ for the relevant nuisance $Q(P)$, and any additional nuisance estimator satisfies $\|g_n - g_0\|_{L_2(P_0)} = O_p(n^{-s})$ with $r + s > 1/2$.

2. The targeting subset θ_{targ} generates a score space rich enough to approximate the Q -component of D^* , such that the $L_2(P_0)$ residual from projecting D_Q^* onto the score space satisfies the rate condition needed for condition (C.2) to hold (see proof below).
3. Assumption 4.3 holds: the regularized projection EIC solving the approximate score equation implies the true working-model EIC equation is also solved at rate $o_p(n^{-1/2})$.

Then, as $\mathcal{M}_n \rightarrow \mathcal{M}$, conditions (A.1)-(A.3), (B.1)-(B.2), and (C.1)-(C.3) hold, and TDA is $\sqrt{n}$ -consistent and asymptotically normal.

As detailed in Appendix B.1, evaluating the exact remainder associated with (C.2) follows the same Cauchy–Schwarz argument used in Chapter 3. The working model $\mathcal{M}_n$ restricts only $Q(P)$ (the part of P relevant for the target parameter), leaving the remaining components of P —including the nuisance $g(P)$ and the marginal distribution—unconstrained. Denoting the Q -dependent component of D^* as D_Q^* (the component lying in the Q -tangent space), the difference $D^*(P_{0,\mathcal{M}_n}) - D_{\mathcal{M}_n}^*(P_{0,\mathcal{M}_n})$ reduces to the residual of projecting D_Q^* onto the working model’s score space. Applying Cauchy–Schwarz yields $|R_n| \lesssim \|D_Q^* - \Pi_S D_Q^*\|_{L_2(P_0)} \times \|Q_{0,\mathcal{M}_n} - Q_0\|_{L_2(P_0)}$, where Π_S denotes projection onto the score space. For TDA, the network’s rich score space causes $\Pi_S D_Q^*$ to well-approximate D_Q^* , making this product $o_p(n^{-1/2})$. Condition (C.3) holds identically via the first-order optimality properties of the loss-based projection.

Remark 4.9 (Convergence rates for neural networks). Unlike HAL, which provides formal convergence rate guarantees under sectional variation norm constraint, general neural networks lack such guarantees. Condition 1 of Theorem 4.8 requires a convergence rate for the initial estimator, which is not guaranteed for arbitrary neural network architectures and training procedures. Two approaches address this limitation. First, variation-constrained neural networks Li and Ding [2022] with arbitrary width can achieve near-parametric rate $n^{-1/2+\delta}$, providing formal guarantees under appropriate architectural constraints. Second, the super learner offset construction (Remark 4.2) resolves this entirely: the initial estimator convergence rate comes from the super learner (which can include HAL or other estimators with known rates), while the neural network only needs to provide a rich score space for the targeting step.

4.4 Experiments

This section retains the empirical evidence needed for the dissertation claim that regularized A-TMLE extends beyond HAL. Detailed training specifications, auxiliary subset-selection results, and implementation specifics are deferred to the appendices.

We evaluate the finite-sample performance of *Targeted Deep Architectures (TDA)* against both standard neural estimators and classical semiparametric causal methods. Beyond the core post-training TDA procedure (Algorithm 3), we investigate *training-time targeting* (TTR), which periodically applies the TDA targeting step during network training, and *Combined estimation*, which chains TTR with post-training TDA. Our primary metrics are

bias, mean squared error (MSE), $E|P_n D^*|$ (whether the EIC equation is solved), and *oracle coverage* (95% CIs from the true Monte Carlo SD, isolating point estimation quality from variance estimation). We investigate three settings:

- (a) **High-dimensional confounding:** ATE estimation on the semi-synthetic IHDP benchmark.
- (b) **Robustness to positivity violations:** ATE estimation under severe near-positivity violations.
- (c) **Functional parameter estimation:** Marginal survival curve estimation under concordant informative censoring.

4.4.1 ATE Estimation on IHDP: High-Dimensional Confounding

Dataset and Setup. The Infant Health and Development Program (IHDP) dataset is a standard causal inference benchmark Hill [2011], Shalit et al. [2017], Shi et al. [2019] comprising 747 units, 25 covariates (X), a binary treatment (A), and semi-synthetic continuous outcomes (Y). The true treatment effect is known, enabling exact assessment of bias and inference calibration. We use an 80%–20% train-validation split for initial model fitting, subsequently applying the targeting procedure on the full sample.

Compared Approaches. We employ DragonNet Shi et al. [2019] as the base architecture and compare 13 estimators organized into three families (full definitions in Appendix B.2.1):

- **Post-training debiasing (standard initialization):** All five estimators in this group start from the *same* trained DragonNet model (outcome regression $\hat{Q}$ and propensity score $\hat{g}$) and differ only in the debiasing step applied post-training. The **Naive Plug-in** applies no debiasing; **TDA-Last** and **TDA-Full** apply Algorithm 3 to the last layer or all outcome-head layers; **Classical TMLE** applies a logistic ε -fluctuation using the same $\hat{Q}$ and $\hat{g}$; and **AIPTW** computes the one-step doubly robust estimator from the same nuisance estimates (not a targeting method).
- **Training-time targeting (TTR):** **TTR-Last** and **TTR-Full** periodically apply the TDA targeting step every 10 training epochs, directly attempting to solve the EIC equation during training. **Combined-Last** and **Combined-Full** chain TTR during training with post-training TDA, using the former to steer the model toward a better initial fit and the latter to finish debiasing.
- **T-regularization family:** **T-reg** Shi et al. [2019] adds an EIC-based loss penalty during training. **T-reg+TDA-Last** and **T-reg+TDA-Full** follow T-reg initialization with post-training TDA. **TMLE Update** follows T-reg with a classical TMLE step.

Standard errors for all methods are computed via the empirical variance of the estimated doubly robust EIC.

Results. Table 4.1 reports performance across 500 Monte Carlo replications. We include the one-step AIPTW estimator as a non-targeting benchmark that uses the same nuisance estimates. The key diagnostic for whether targeting succeeds is $E|P_n D^*|$ (whether the EIC equation is solved); the key metric for point estimation quality is *oracle coverage* (95% CIs constructed from the true Monte Carlo SD), which isolates bias control from variance estimation.

Post-training TDA matches classical targeting methods. TDA-Last and TDA-Full perform comparably to Classical TMLE: all three solve the EIC equation ($E|P_n D^*| \leq 0.011$) and achieve oracle coverages of 95–97%, confirming that TDA is a legitimate TMLE realization. AIPTW, while not a targeting method, also achieves 95.0% oracle coverage by virtue of its one-step EIC correction. The $|\text{Bias}|/\text{SD}$ ratios for TDA-Full (0.565), TDA-Last (0.600), and Classical TMLE (0.600) are all well below 1, confirming proper centering.

TTR alone fails to solve the EIC equation; Combined fixes this. The critical diagnostic is $E|P_n D^*|$: TTR-Last leaves a large EIC residual (0.296) and TTR-Full leaves 0.146, confirming that training-time targeting alone does not solve the score equation—the targeting updates are overwhelmed by the competing predictive loss. The consequence is degraded oracle coverage: TTR-Full achieves only 91.2% (vs. 96.8% for TDA-Full), reflecting residual bias from the unsolved EIC equation. Adding post-training TDA solves the equation and restores oracle coverage: Combined-Last reduces $E|P_n D^*|$ from 0.296 to 0.006 with oracle coverage 95.4% (vs. TTR-Last’s 93.8%), and Combined-Full reduces $E|P_n D^*|$ from 0.146 to 0.006 with oracle coverage 94.8% (vs. TTR-Full’s 91.2%) and the lowest RMSE among all 13 methods (0.171). The key insight is that TTR provides a better *initialization* for post-training targeting—Combined-Full’s SD (0.146) is substantially lower than TDA-Full’s (0.238)—but cannot reliably debias on its own because the EIC equation remains unsolved.

T-reg does not solve the EIC equation. T-reg’s loss penalty reduces but does not eliminate the EIC residual ($E|P_n D^*| = 0.225$), and the consequence is visible in oracle coverage: 93.2%, the lowest among all debiased methods—confirming that a loss penalty alone is insufficient for strict first-order bias removal. Adding post-training TDA solves the equation and restores oracle coverage: T-reg+TDA-Last achieves 95.6% and T-reg+TDA-Full achieves 95.4%. The T-reg family also produces inflated SE estimates (Mean SE ≈ 0.31 vs. ≈ 0.25 for other methods), but this is a secondary concern: the primary problem is that the EIC equation is not solved, leaving residual bias that degrades point estimation.

We additionally validate our data-adaptive plateau selector (Appendix B.3.1, Table B.2). In this well-specified setting, the selector predominantly chooses the largest subset ($k^*=6$ in 88% of replications) because variance inflation is negligible across subsets—confirming the designed conservative behavior of defaulting to the richest working model when no overfitting is detected, while serving as a safety net that would halt expansion if variance inflation were to outpace bias correction.

4.4.2 ATE Estimation Under Positivity Violations

DGP and Setup. To investigate the stability of targeting mechanisms under adversarial conditions, we design a synthetic DGP with 5 covariates $X \sim \text{Uniform}(-1, 1)^5$, a simple

Table 4.1: ATE estimation on IHDP (500 replications). Estimators are grouped by targeting strategy. Metrics: Bias, $|\text{Bias}|/\text{SD}$ (centering quality), SD (Monte Carlo standard deviation), RMSE ($\sqrt{\text{Bias}^2 + \text{SD}^2}$), Mean SE (average estimated standard error), Nom Cov (95% CI coverage using estimated SE), Orc Cov (coverage using true MC SD), $E|P_n D^*|$ (mean absolute EIC residual).

Estimator	Bias	$ \text{B} /\text{SD}$	SD	RMSE	Mean SE	Nom Cov	Orc Cov	$E P_n D^* $
<i>Post-training debiasing (standard initialization)</i>								
Naive Plug-in	−0.197	0.512	0.385	0.432	0.246	71.0%	93.2%	0.209
TDA-Last	−0.135	0.600	0.224	0.261	0.246	91.8%	95.4%	0.006
TDA-Full	−0.135	0.565	0.238	0.273	0.246	93.0%	96.8%	0.011
Classical TMLE	−0.134	0.600	0.224	0.261	0.246	91.2%	95.8%	0.000
AIPTW	−0.139	0.616	0.225	0.264	0.246	91.4%	95.0%	0.000
<i>Training-time targeting (TTR)</i>								
TTR-Last	−0.203	0.262	0.778	0.803	0.247	72.4%	93.8%	0.296
Combined-Last	−0.154	0.420	0.365	0.396	0.249	90.2%	95.4%	0.006
TTR-Full	−0.113	0.481	0.234	0.260	0.246	81.4%	91.2%	0.146
Combined-Full	−0.090	0.614	0.146	0.171	0.246	92.4%	94.8%	0.006
<i>T-regularization family</i>								
T-reg	−0.128	0.553	0.232	0.265	0.326	83.2%	93.2%	0.225
T-reg+TDA-Last	−0.144	0.493	0.292	0.325	0.312	92.0%	95.6%	0.024
T-reg+TDA-Full	−0.129	0.538	0.239	0.272	0.315	91.6%	95.4%	0.055
TMLE Update	−0.158	0.468	0.339	0.374	0.309	93.2%	95.2%	0.003

propensity model $\logit \pi(X) = 3.5X_1 + 2.5X_2$ (yielding propensity scores in $[0.002, 0.998]$), and a complex outcome mechanism where the conditional average treatment effect (CATE) shares drivers with propensity extremity: $\tau(X) = 1 + 2.5X_1^2 + 1.5\cos(2X_2) + 3X_1X_2 + 2\max(0, -X_1)X_3$. The true ATE is approximately 1.5. The design creates systematic plug-in errors in the low-data region—subjects with extreme propensity scores also exhibit the most complex treatment effect heterogeneity, including a “hidden” CATE component $\max(0, -X_1)X_3$ active only in the low-propensity-of-treatment region. Full DGP specification is in Appendix B.2.1.2. We execute $B = 500$ replications at $n \in \{500, 1000\}$.

Results. Tables 4.2 and 4.3 report results at $n = 500$ and $n = 1000$, respectively. The $n = 500$ regime is most discriminating, as smaller samples amplify the impact of near-positivity violations.

Post-training targeting under positivity stress. Under positivity violations, oracle coverage—which uses the true Monte Carlo SD and thus isolates point estimation quality from variance estimation—discriminates sharply among methods. TDA-Full and Classical TMLE achieve comparable oracle coverages (89.6% and 90.0%) with solved EIC equations ($E|P_n D^*| = 0.008$ and 0.000). AIPTW, as a non-targeting one-step estimator, is more sensitive to extreme propensity weights: its oracle coverage drops to 86.0% ($|\text{Bias}|/\text{SD} = 0.866$),

reflecting the heavier-tailed noise from inverse probability weighting.

TTR alone does not solve the EIC equation; Combined restores valid targeting. TTR-Last achieves near-zero raw bias (+0.015) and 95.0% oracle coverage at $n = 500$, but its EIC residual remains large ($E|P_n D^*| = 0.157$): the score equation is not solved, so the theoretical guarantees of TMLE do not apply. Combined-Last achieves comparable oracle coverage (95.2%) while actually solving the EIC equation ($E|P_n D^*| = 0.032$)—and the solved equation enables more reliable variance estimation (nominal coverage 77.0% vs. 66.8%). The same pattern holds at the full-network level: TTR-Full has 92.4% oracle coverage with unsolved EIC ($E|P_n D^*| = 0.127$), while Combined-Full achieves 92.8% oracle coverage with solved EIC ($E|P_n D^*| = 0.008$) and the lowest RMSE (0.292).

TTR provides value as initialization, not as a standalone estimator. Despite failing to solve the EIC equation alone, TTR produces representations that are more robust to positivity violations, giving post-training TDA a better starting point. This is why Combined-Full (RMSE = 0.292, oracle coverage 92.8%) substantially outperforms TDA-Full (RMSE = 0.360, oracle coverage 89.6%) despite both solving the EIC equation. At $n = 1000$ (Table 4.3), the top methods converge—Combined-Full, TMLE Update, and T-reg+TDA variants all achieve oracle coverage $\geq 91.6\%$ and RMSE ≈ 0.186 – 0.190 —confirming that the differentiation is most visible at smaller sample sizes where robustness matters most.

T-reg fails to solve the EIC equation. T-reg alone leaves a large EIC residual ($E|P_n D^*| = 0.205$), and the consequence is severe: oracle coverage is only 77.8%, confirming substantial residual bias in the point estimates. Adding post-training targeting solves the equation and restores oracle coverage: T-reg+TDA-Full achieves $E|P_n D^*| = 0.009$ with oracle coverage 91.0%, and TMLE Update achieves $E|P_n D^*| = 0.000$ with oracle coverage 91.8%.

4.4.3 Survival Curve Estimation: Regularizing Functional Updates

We extend our evaluation to a strictly functional parameter: the marginal counterfactual survival curve $S(t)$ evaluated across a dense grid of 50 time points. The data-generating process features *concordant late acceleration*: a Weibull baseline hazard with nonlinear covariate effects and a late-time acceleration component that activates after $t = 7$, combined with concordant informative censoring that systematically censors the subjects who would experience this acceleration ($n = 1,000$, $p = 10$, $\approx 73\%$ censoring rate; full DGP specification in Appendix B.2.2). This design creates a fundamental challenge: the late hazard acceleration is precisely hidden in the region where data is most scarce, producing systematic bias that cannot be resolved without proper adjustment for informative censoring. We compare TDA against the initial Neural Hazard model, the classical Kaplan-Meier (KM) estimator, and **survtmle** Benkeser and Hejazi [2017], a discrete-time TMLE benchmark whose discretization grid matches our evaluation grid exactly (both use 50 equally-spaced points on $[0, 14]$).

Results and Functional Regularization. Table 4.4 presents performance averaged across the 50 time points (500 replications). As in the ATE setting, oracle coverage—95%

Table 4.2: ATE estimation under positivity violations ($n = 500$, 500 replications). True ATE = 1.5.

Estimator	Bias	B /SD	SD	RMSE	Mean SE	Nom Cov	Orc Cov	$E P_n D^* $
<i>Post-training debiasing (standard initialization)</i>								
Naive Plug-in	-0.481	1.506	0.319	0.577	0.215	42.2%	68.8%	0.251
TDA-Last	-0.227	0.746	0.304	0.379	0.200	70.2%	87.6%	0.040
TDA-Full	-0.175	0.555	0.315	0.360	0.202	72.8%	89.6%	0.008
Classical TMLE	-0.166	0.529	0.314	0.354	0.199	72.8%	90.0%	0.000
AIPTW	-0.254	0.866	0.294	0.388	0.215	70.4%	86.0%	0.000
<i>Training-time targeting (TTR)</i>								
TTR-Last	+0.015	0.043	0.357	0.357	0.163	66.8%	95.0%	0.157
Combined-Last	+0.033	0.124	0.266	0.268	0.161	77.0%	95.2%	0.032
TTR-Full	-0.128	0.418	0.308	0.333	0.186	73.6%	92.4%	0.127
Combined-Full	-0.084	0.301	0.279	0.292	0.178	77.6%	92.8%	0.008
<i>T-regularization family</i>								
T-reg	-0.376	1.196	0.315	0.490	0.204	56.6%	77.8%	0.205
T-reg+TDA-Last	-0.192	0.660	0.291	0.348	0.191	72.2%	89.0%	0.034
T-reg+TDA-Full	-0.154	0.516	0.299	0.336	0.190	73.6%	91.0%	0.009
TMLE Update	-0.143	0.472	0.303	0.335	0.194	75.2%	91.8%	0.000

Table 4.3: ATE estimation under positivity violations ($n = 1000$, 500 replications). True ATE = 1.5.

Estimator	Bias	B /SD	SD	RMSE	Mean SE	Nom Cov	Orc Cov	$E P_n D^* $
<i>Post-training debiasing (standard initialization)</i>								
Naive Plug-in	-0.282	1.282	0.220	0.357	0.150	52.8%	76.6%	0.179
TDA-Last	-0.107	0.578	0.185	0.214	0.141	78.4%	91.4%	0.027
TDA-Full	-0.079	0.416	0.189	0.204	0.142	78.4%	93.2%	0.007
Classical TMLE	-0.072	0.366	0.195	0.208	0.142	81.0%	92.6%	0.000
AIPTW	-0.122	0.653	0.186	0.222	0.150	78.2%	89.8%	0.000
<i>Training-time targeting (TTR)</i>								
TTR-Last	+0.068	0.193	0.352	0.358	0.137	58.0%	95.0%	0.199
Combined-Last	+0.065	0.311	0.210	0.220	0.133	77.6%	94.4%	0.039
TTR-Full	-0.105	0.539	0.195	0.221	0.134	75.0%	91.6%	0.089
Combined-Full	-0.045	0.245	0.185	0.190	0.132	81.2%	95.6%	0.006
<i>T-regularization family</i>								
T-reg	-0.205	1.073	0.191	0.280	0.143	62.8%	83.0%	0.134
T-reg+TDA-Last	-0.079	0.464	0.171	0.189	0.135	81.2%	91.6%	0.025
T-reg+TDA-Full	-0.058	0.324	0.179	0.188	0.134	80.4%	93.2%	0.007
TMLE Update	-0.052	0.293	0.178	0.186	0.134	81.2%	94.6%	0.000

CIs constructed from the true Monte Carlo SD—isolates point estimation quality from variance estimation. The concordant censoring mechanism produces dramatic separation among methods on this metric.

Oracle coverage reveals point estimation quality. KM’s oracle coverage is only 5.4%, confirming that its point estimates are severely biased relative to sampling variability ($|\text{Bias}|/\text{SD} = 13.22$): ignoring informative censoring is a structural failure, not a finite-sample artifact. The initial neural hazard model improves to 73.1% oracle coverage ($|\text{Bias}|/\text{SD} = 1.34$), capturing the complex DGP structure but not fully removing censoring-induced bias without solving the EIC equation. The discrete-time TMLE benchmark `survtmle` (hazard method) achieves 78.7% oracle coverage; like TDA, `survtmle`’s hazard-based TMLE pools information across time bins via a single logistic fluctuation model with clever covariates that couple all time points—both methods share the principle of a combined targeting direction rather than independent per-bin updates. TDA achieves the highest oracle coverage (83.7%) and is the *only* method with $|\text{Bias}|/\text{SD} < 1$ (0.96), confirming that bias is better controlled relative to sampling variability. TDA’s structural advantage is *adaptive regularization through the projection*: both methods iteratively update and share information across time, but they differ in how targeting directions are constructed. TDA projects the EIC equations for all 50 time points onto the neural network’s score space with L_2 regularization, then combines per-time-point targeting gradients into a single universal update direction following the universal least favorable submodel principle van der Laan and Gruber [2016]—with the projection’s adaptive regularization controlling the bias–variance tradeoff across the combined update. In contrast, `survtmle`’s iterative locally least favorable submodel must re-estimate clever covariates after each fluctuation step (up to 25 iterations), without the adaptive regularization that the L_2 -penalized projection provides, and its discrete-hazard parametrization may introduce discretization artifacts.

Nominal coverage reflects the inference problem. Beyond point estimation, nominal coverage depends on whether EIC-based inference is justified. The initial neural hazard model does not solve the EIC equation, so EIC-based inference carries no theoretical guarantee; its 64.6% nominal coverage (8.5-point gap from oracle) reflects this lack of targeting. KM’s Greenwood formula ignores informative censoring entirely, producing erroneously narrow CIs (width 0.075) and only 6.1% nominal coverage. Both `survtmle` and TDA solve a version of the EIC equation, yet nominal coverage remains below 95% for both: `survtmle` achieves 37.4% (a 41.3-point gap from oracle) and TDA achieves 80.1% (a 3.6-point gap). For TDA, the targeting step solves the *projected* EIC—an L_2 -regularized projection of the EIC onto the score space—rather than the full nonparametric EIC, which introduces a gap between the solved equation and the one required for exact semiparametric inference. The built-in complexity of targeting a 50-dimensional functional parameter under 73% informative censoring further challenges EIC-based inference for both methods: each time point’s EIC involves censoring-weighted inverse-probability terms that amplify estimation noise near the boundary, and the high dimensionality of the targeting problem makes variance estimation inherently harder. Closing this gap—through improved variance estimation, cross-fitting, or higher-order corrections—remains an important open problem (see Limitation *viii*).

Training-Time Targeting Does Not Improve Survival in This Setting. In contrast to the ATE setting where TTR improved the Combined estimator’s initialization (Section 4.4.1), training-time targeting provides no improvement for survival-curve estimation in our experiments. Across all 500 replications, TTR performs at approximately chance level: about 51% of replications produce worse $|\text{Bias}|/\text{SD}$ than the untouched initial model. Hyperparameter tuning (7 configurations, 50 replications each; see Appendix B.2.2) did not resolve this: conservative settings remain near chance level while aggressive settings cause numerical divergence. A plausible explanation is the dimensionality of the targeting problem. ATE targeting solves a single scalar EIC equation—the gradient signal is strong and unidirectional. Survival targeting simultaneously solves 50 EIC equations (one per time point) whose corrections can pull network parameters in conflicting directions. Combined with the low per-time-point signal-to-noise ratio, the targeting gradient may be drowned out by noise during training. Consequently, the Combined estimator (TTR + post-training TDA) reduces to TDA alone for survival in this setting. This motivates post-training TDA—which operates on a converged model and solves the system of equations directly without competing with the training loss—as the preferred approach for high-dimensional functional targets, while alternative training-time strategies (e.g., curriculum-based scheduling, adaptive weighting across time points, or architectures specifically designed for multi-target targeting) remain promising directions for future work.

Table 4.4: Survival curve estimation under concordant informative censoring (500 replications, $n = 1,000$). Each metric is computed at each of 50 time points; we report mean $\pm$ SD across time points. Bold indicates best performance.

Metric	Initial Model	Kaplan-Meier	survtmle	TDA-Last
MSE	0.0043 (± 0.0016)	0.0799 (± 0.0414)	0.0023 (± 0.0037)	0.0016 (± 0.0006)
Abs. Bias	0.0494 (± 0.0161)	0.2614 (± 0.1069)	0.0292 (± 0.0259)	0.0253 (± 0.0090)
Variance	0.0016 (± 0.0006)	0.0004 (± 0.0002)	0.0008 (± 0.0004)	0.0008 (± 0.0003)
$ \text{Bias} /\text{SD}$	1.34 (± 0.40)	13.22 (± 3.65)	1.14 (± 1.19)	0.96 (± 0.43)
Coverage (%)	64.6 (± 21.2)	6.1 (± 22.3)	37.4 (± 21.2)	80.1 (± 16.2)
Orc. Cov (%)	73.1 (± 13.1)	5.4 (± 20.8)	78.7 (± 21.9)	83.7 (± 12.1)
CI Width	0.135 (± 0.066)	0.075 (± 0.025)	0.035 (± 0.009)	0.135 (± 0.068)

4.5 Discussion

Targeted Deep Architectures realize the regularized A-TMLE framework of Chapter 3 in a neural-network working model. The trained network serves as the offset, a chosen subset of weights defines the targeting working model, and the working-model efficient influence curve is approximated by regularized projection onto the corresponding gradient score span. Because the projected EIC is by construction a linear combination of score vectors, the empirical score equation is solved exactly at any local stationary point of the targeted risk,

giving post-training TDA a valid targeting step under the A-TMLE theory of this chapter. The multi-parameter extension via a universal least favorable path carries the same mechanism to functional targets, which is the form in which it reappears in Chapter 6’s survival-grid setting. The empirical evaluation supports this conceptual claim: TDA-Full matches Classical TMLE on the IHDP benchmark and behaves well under positivity violations, and a combined training-time/post-training variant gives a better initialization for scalar targets. As expected, training-time targeting does not help in the 50-dimensional survival grid—an early sign that naive training-time blending faces dimensionality challenges, which the dissertation’s main survival pipeline addresses differently in Chapter 6.

Limitations.

- (i) *Non-convex working-model risk.* Unlike a HAL sieve, the neural-network working model is nonlinear in θ_{targ} , so the working-model risk landscape is non-convex. Non-convexity does *not* prevent TDA from solving the empirical score equation, because the projected EIC is by construction linear in the score vectors and $P_n[\tilde{D}_{\text{proj}}^*] = 0$ at any local stationary point of the targeted risk. The practical implication is that TDA may converge to a suboptimal local minimum, which can degrade the predictive fit and leave a larger exact remainder R_n . Verifying the structural regularity conditions in Theorem 4.8 (score-space approximation of the EIC direction and oracle rates) from finite data remains difficult.
- (ii) *Score-space approximation is not guaranteed.* The argument in Section 4.2.3 is that the full network score space is *typically* rich enough to well-approximate the EIC direction, but that is an empirical property that depends on architecture and problem. In pathological cases the plateau selector should detect degradation through widening bands.
- (iii) *Targeting hyperparameters.* The projection penalty λ , the step size γ , and the targeting-subset size are the three knobs TDA introduces that classical TMLE does not. The chapter selects them through small pilot studies; a fully data-driven rule that integrates with inferential validity is still open.
- (iv) *Nominal coverage and the projected EIC.* TDA solves a projected EIC rather than the full nonparametric EIC. The resulting inferential gap is chapter-specific in its exact form but is a general theme across the dissertation; Chapter 5 studies one practical route around it, and Chapter 6 uses that route.

Within the scope of this chapter, the most immediate extensions are more principled selection of the targeting submodel and better strategies for making training-time targeting effective for functional parameters, for example through curriculum-based scheduling or adaptive weighting across equations.

Taken together, Chapters 3 and 4 complete the estimator-side development of the dissertation. The next step is to turn to inference and to ask how uncertainty quantification should proceed once targeted estimators are available but direct variance calculations are unstable, unavailable, or too expensive to use routinely.

Chapter 5

Inference for Targeted Estimators via the V-fold Jackknife

The preceding chapters develop regularized targeting on the estimator side. This chapter turns to the corresponding inferential problem: once a targeted estimator has been constructed, how should uncertainty quantification proceed when direct influence-curve variance calculations are analytically cumbersome, numerically unstable, or impractical for an expensive adaptive pipeline? The grouped V-fold jackknife developed here provides one answer to that question.

Its importance in the dissertation is complementary to the targeting chapters. The grouped jackknife exploits the asymptotic-linearity expansion displayed in Chapter 2, including the generalized form in which the influence-function variance may diverge with the sample size, without relying on delicate inverse-information calculations. It is also computationally attractive because it requires only V grouped refits of the full estimator rather than bootstrap-scale recomputation. Those features make it especially well suited to the expensive adaptive procedures considered later, where inference must respect the same computational constraints that shaped the estimator itself. Because this chapter treats regular and generalized asymptotically linear estimators beyond the efficient targeted estimators of the earlier chapters, its technical development uses φ for a generic influence function rather than D^* for an efficient influence curve.

5.1 Inference problem and chapter scope

Chapters 3 and 4 stabilize the estimator; this chapter stabilizes inference. The earlier chapters construct targeted estimators whose bias is controlled through efficient-influence-curve equations inside adaptive working models. Once such an estimator is in hand, the next problem is how to carry out inference when the direct variance layer is numerically fragile, analytically cumbersome, or unavailable after regularization. The grouped V-fold jackknife answers this question by replacing analytic variance calculations with fold-level refits while preserving the asymptotic-linear logic established in Chapter 2.

Three features of the procedure matter for the dissertation. First, valid asymptotic inference already holds for any fixed $V \geq 2$, which makes it practical even when each estimator refit is expensive—as it is for targeted-learning pipelines and for HAL-based estimators. Second, the method requires no knowledge of the influence function or the convergence rate; it extends to generalized asymptotically linear estimators whose influence function variance σ_n^2 diverges with n , including HAL functionals converging at rates slower than $\sqrt{n}$. Third, a vector-valued extension yields simultaneous confidence bands for functional parameters, which is the form in which the jackknife is used in Chapter 6.

Observed-data notation, the statistical model $\mathcal{M}$, and the target parameter $\Psi(P_0)$ follow Chapter 2, Table 2.1. Throughout this chapter we consider an estimator $\hat{\Psi}(P_n)$ admitting the asymptotically linear expansion

$$\hat{\Psi}(P_n) - \Psi(P_0) = P_n\varphi + R(P_n, P_0),$$

where P_n is the empirical measure, $\varphi = \varphi_{P_0}$ is a mean-zero influence function with $\mathbb{E}_{P_0}[\varphi^2] < \infty$, and $R(P_n, P_0) = o_P(n^{-1/2})$ is a remainder term. We use φ rather than D^* here because Chapter 5 treats general regular and generalized asymptotically linear estimators, not only efficient ones. In Section 5.5, this is generalized to allow the influence-function variance $\sigma_n^2 = \mathbb{E}_{P_0}[\varphi_n^2]$ to diverge with n , covering estimators that converge at rates slower than $\sqrt{n}$ such as HAL functional estimators. For the jackknife construction, $\{I_1, \dots, I_V\}$ denotes a partition of $\{1, \dots, n\}$ into disjoint folds of approximately equal size $|I_v| \approx n/V$, with fold-level empirical measure $P_{n,v}^1 f = |I_v|^{-1} \sum_{i \in I_v} f(O_i)$ and leave-fold-out measure $P_{n,-v}$. The v th jackknife pseudo-value is

$$IC_{\text{Jack}}(v) = V\hat{\Psi}(P_n) - (V-1)\hat{\Psi}(P_{n,-v}),$$

and the jackknife variance estimator of $\sigma^2 = \text{Var}(\varphi(O))$ is the (scaled) empirical variance of these pseudo-values.

5.1.1 Related work and positioning

Chapter 2 already gives the dissertation-wide first-order semiparametric framework. The present subsection records the more specific resampling and simultaneous-inference literatures needed to position grouped V-fold jackknife inference.

The grouped jackknife belongs to the pseudo-value tradition beginning with Quenouille and Tukey, with classical fixed- V Studentization results established in parametric settings by Brillinger and reviewed by Miller; related grouped constructions also appear in the U -statistics literature [Quenouille, 1949, Tukey, 1958, Brillinger, 1964, Miller, 1974, Arvesen, 1969]. The present chapter keeps that grouped pseudo-value construction but extends it from classical parametric estimators to regular and generalized asymptotically linear estimators of pathwise differentiable targets.

A complementary comparison class is delete-a-group jackknife methodology in survey sampling. Under that name, grouped jackknife methods are used with replicate weights in large-sample stratified survey settings, where Kott [1998, 2001] develop practical con-

structions together with asymptotic design-based justification, and Kott and Garren [2011] study small-sample bias concerns. That line is close in spirit because it also treats grouped deletions as a computationally manageable alternative to more cumbersome variance calculations. The contribution here is different: this chapter is not a survey-replicate-weight procedure, but a first-order theory for fixed- V Studentized inference, diverging- V consistency, and simultaneous bands for regular and generalized asymptotically linear estimators.

Another comparison class is the delete- d jackknife and its consistency theory. Deleting one fold in a V -fold partition is a structured delete- d scheme with $d = n/V$, so the diverging- V results here are closest in spirit to the literature developed by Wu [1986] and Shao and Wu [1989]. The difference is that the analysis is stated directly in terms of asymptotic linearity and fold-stable remainder control, which is the natural language for targeted and other semiparametric estimators.

The main practical alternative remains the bootstrap and its cheaper variants. Bootstrap-style intervals are broadly applicable but can be computationally prohibitive when each estimator refit is expensive, and recent cheaper bootstrap or subsampling variants still require an extra Monte Carlo layer and subsample-fraction choices [Lam, 2022, Ohlendorff et al., 2025]. The grouped jackknife instead uses one deterministic partition, keeps a large fraction of the sample in each refit, and yields fixed- V Studentized inference without requiring thousands of replicate fits.

For functional targets, the closest comparison class is simultaneous-band construction via stacked influence curves or bootstrap sup- t methods [van der Laan, 2019, Díaz and van der Laan, 2013, Montiel Olea and Plagborg-Møller, 2019]. The contribution here is to give a grouped-jackknife alternative that uses only V refits, extends to the generalized asymptotically linear setting, and matches the dissertation’s broader strategy of replacing fragile analytic variance calculations by computationally feasible first-order surrogates.

5.2 Grouped V -fold jackknife construction

For an asymptotically linear estimator $\hat{\Psi}(P_n)$, we have

$$\hat{\Psi}(P_n) - \Psi(P_0) = \frac{1}{V} \sum_{v=1}^V P_{n,v}^1 \varphi + R(P_n, P_0),$$

and

$$\hat{\Psi}(P_{n,-v}) - \Psi(P_0) = \frac{1}{V-1} \sum_{\substack{v'=1 \\ v' \neq v}}^V P_{n,v'}^1 \varphi + R(P_{n,-v}, P_0).$$

Let $IC_{Jack}(v) = V\hat{\Psi}(P_n) - (V-1)\hat{\Psi}(P_{n,-v})$ denote the v th *jackknife pseudo-value* [Tukey, 1958], which can also be viewed as a fold-level empirical influence curve. Then we have,

$$IC_{Jack}(v) - \Psi(P_0) = P_{n,v}^1 \varphi + R(P_{n,-v}, P_0) + V\{R(P_n, P_0) - R(P_{n,-v}, P_0)\}.$$

Algorithm 5 Univariate V -fold jackknife confidence interval

- 1: **Input:** Estimator $\hat{\Psi}$, observations $O_1, \dots, O_n$, number of folds $V \geq 2$, confidence level $1 - \alpha$.
 - 2: **Output:** $(1 - \alpha)$ -level confidence interval for $\Psi(P_0)$.
 - 3: Compute $\hat{\Psi}(P_n)$ using all n observations.
 - 4: Partition $\{1, \dots, n\}$ into V disjoint folds $I_1, \dots, I_V$ of approximately equal size.
 - 5: **for** $v = 1$ **to** V **do**
 - 6: Compute the leave-fold-out estimate $\hat{\Psi}(P_{n,-v})$ using $\{1, \dots, n\} \setminus I_v$.
 - 7: Compute the pseudo-value $IC_{Jack}(v) = V\hat{\Psi}(P_n) - (V - 1)\hat{\Psi}(P_{n,-v})$.
 - 8: **end for**
 - 9: Compute $\widehat{SE} = \hat{S}_{Jack}/\sqrt{V}$, where $\hat{S}_{Jack}^2 = (V - 1)^{-1} \sum_{v=1}^V (IC_{Jack}(v) - \bar{IC}_{Jack})^2$.
 - 10: **Return** the confidence interval $\hat{\Psi}(P_n) \pm t_{V-1, 1-\alpha/2} \cdot \widehat{SE}$.
-

Let $\hat{\Psi}_{Jack}(P_n) = \frac{1}{V} \sum_{v=1}^V IC_{Jack}(v)$ denote the *jackknife estimator* of $\Psi(P_0)$, i.e., the mean of the pseudo-values. We define the V -fold jackknife variance estimator of $\sigma^2 = \text{Var}(\varphi(O))$ as

$$\hat{S}_{Jack}^2 = (V - 1)^{-1} \sum_{v=1}^V (IC_{Jack}(v) - \hat{\Psi}_{Jack}(P_n))^2. \quad (5.1)$$

In this formulation, the true variance is $\text{Var}(\hat{\Psi}(P_n)) \approx \frac{\sigma^2}{n}$. By the independence of folds, $\text{Var}(P_{n,v}^1 \varphi) = \left(\frac{1}{n/V}\right)^2 \sum_{i \in I_v} \sigma^2 = \frac{V}{n} \sigma^2$. Hence, the sample variance of the fold-level averages $S_W^2 = (V - 1)^{-1} \sum_v (P_{n,v}^1 \varphi - \bar{W})^2$ naturally estimates $\frac{V}{n} \sigma^2$. Consequently, the scaled V -fold Jackknife variance estimator $\frac{n}{V} \hat{S}_{Jack}^2$ estimates σ^2 .

Define the jackknife standard error $\widehat{SE} = \hat{S}_{Jack}/\sqrt{V}$. The V -fold jackknife $(1 - \alpha)$ -level confidence interval for $\Psi(P_0)$ is

$$\hat{\Psi}(P_n) \pm t_{V-1, 1-\alpha/2} \cdot \widehat{SE},$$

where $t_{V-1, 1-\alpha/2}$ denotes the $(1 - \alpha/2)$ -quantile of the Student t -distribution with $V - 1$ degrees of freedom. The use of t_{V-1} critical values rather than normal quantiles provides a finite-sample correction. Its theoretical justification is that, under the conditions stated below, the Studentized statistic converges in distribution to t_{V-1} when V is fixed. The complete procedure is summarized in Algorithm 5.

5.3 Properties of V -fold Jackknife

We study a single Studentized jackknife procedure for constructing confidence intervals for $\Psi(P_0)$, but analyze it in two asymptotic regimes: one with a fixed number of folds V , and one with $V \rightarrow \infty$. When V is fixed, the t_{V-1} correction reflects the finite number of fold-

level contributions; when V grows, the same procedure remains valid and the t_{V-1} critical values approach the corresponding normal quantiles. This makes the trade-off transparent: small fixed V typically improves the finite-sample t approximation but yields wider intervals, whereas slowly increasing V produces narrower intervals while preserving asymptotic validity. To clarify this behavior, we also derive the exact algebraic decomposition of the jackknife variance estimator, which isolates the dominant linear component and makes the role of the remainder explicit. In particular, it shows why the t correction is central: for linear estimators it acts directly on the fold-level means, and for asymptotically linear estimators the same logic applies to the leading linear term.

5.3.1 Fixed V results

We now consider the regime in which V is fixed. With V fixed, the variance estimator does not concentrate to the asymptotic variance in probability as $n \rightarrow \infty$, because the effective number of independent components remains V . Nevertheless, the fold-level fluctuations still support valid asymptotic t -type inference. Specifically, we prove that

$$\frac{\sqrt{V}(\hat{\Psi}_{Jack}(P_n) - \Psi(P_0))}{\hat{S}_{Jack}} \Rightarrow_d t_{V-1}.$$

The intuition is straightforward: if the remainder terms are negligible, then $IC_{Jack}(v) - \Psi(P_0) = P_{n,v}^1 \varphi$ for each v , so each fold mean is asymptotically normal and Studentization across the V independent fold means yields the t_{V-1} limit.

Lemma 5.1 (Equivalence of $\hat{S}_{Jack}$ and S_W for fixed V). *Fix $V \geq 2$ and let $|I_v| = n/V \rightarrow \infty$. Denote $W_v = P_{n,v}^1 \varphi$, $\bar{W} = V^{-1} \sum_{v=1}^V W_v$, and*

$$S_W^2 = \frac{1}{V-1} \sum_{v=1}^V (W_v - \bar{W})^2, \quad \hat{S}_{Jack}^2 = \frac{1}{V-1} \sum_{v=1}^V \left(IC_{Jack}(v) - \hat{\Psi}_{Jack}(P_n) \right)^2.$$

Assume that $\max_{1 \leq v \leq V} |R(P_{n,-v}, P_0)| = o_p(n^{-1/2})$, and that $\text{Var}(\varphi(O)) \in (0, \infty)$. Then, $\hat{S}_{Jack}^2 = S_W^2 + o_p(n^{-1})$, which implies $\hat{S}_{Jack}^2 = S_W^2 \{1 + o_p(1)\}$.

Proof location. The proof is given in Appendix C.1.1.

Theorem 5.2. *Let $W_v = P_{n,v}^1 \varphi$ and $S_W^2 = (V-1)^{-1} \sum_{v=1}^V (W_v - \bar{W})^2$. Assuming $R(P_n, P_0) = o_p(n^{-1/2})$ and $\max_{1 \leq v \leq V} |R(P_{n,-v}, P_0)| = o_p(n^{-1/2})$, we have*

$$\frac{\sqrt{V}(\hat{\Psi}_{Jack}(P_n) - \Psi(P_0))}{\hat{S}_{Jack}} \Rightarrow_d t_{V-1}.$$

Proof location. The proof is given in Appendix C.1.2.

The Studentized statistic in Theorem 5.2 is centered at $\hat{\Psi}_{Jack}(P_n)$, whereas the confidence intervals we construct are centered at the full-sample estimator $\hat{\Psi}(P_n)$. The following lemma

shows that these two centerings are asymptotically equivalent, connecting the centering mismatch to the jackknife bias estimate and to the remainder structure already established in Theorem 5.1.

Lemma 5.3 (Centering equivalence for fixed V). *Under the conditions of Theorem 5.2, define the jackknife bias estimate*

$$\hat{b} = (V - 1)(\bar{\Psi}_{(-v)} - \hat{\Psi}(P_n)), \quad \bar{\Psi}_{(-v)} = \frac{1}{V} \sum_{v=1}^V \hat{\Psi}(P_{n,-v}). \quad (5.2)$$

Then $\hat{\Psi}_{Jack}(P_n) - \hat{\Psi}(P_n) = -\hat{b}$ and $\hat{b} = o_p(\widehat{SE})$, where $\widehat{SE} = \hat{S}_{Jack}/\sqrt{V}$.

Proof location. The proof is given in Appendix C.1.3.

As an immediate consequence, we can center the confidence interval at the full-sample estimator $\hat{\Psi}(P_n)$.

Corollary 5.4 (Confidence interval centered at $\hat{\Psi}(P_n)$). *Under the conditions of Theorem 5.2,*

$$\frac{\sqrt{V}(\hat{\Psi}(P_n) - \Psi(P_0))}{\hat{S}_{Jack}} \Rightarrow_d t_{V-1}.$$

Hence a $(1 - \alpha)$ -level confidence interval for $\Psi(P_0)$ is

$$\hat{\Psi}(P_n) \pm t_{V-1, 1-\alpha/2} \cdot \widehat{SE}, \quad \widehat{SE} = \frac{\hat{S}_{Jack}}{\sqrt{V}}. \quad (5.3)$$

Proof location. The proof is given in Appendix C.1.4.

In the simulation studies section we show that the confidence intervals (5.3) have valid coverage even when $V = 5$. One disadvantage of choosing smaller values of V is that the resulting confidence intervals can be wide. As we increase the number of folds, the quantiles of the t -distribution get closer to those from a normal distribution and thus the confidence intervals shrink.

Remark 5.5 (Finite-sample behavior and the role of V). The fixed- V result of Theorem 5.2 also has a useful finite-sample interpretation. Consider first *linear estimators* — those for which $\hat{\Psi}(P_n) - \Psi(P_0) = P_n\varphi$ exactly, with no remainder ($R(P_n, P_0) = 0$). Then the fold-level averages $W_v = P_{n,v}^1\varphi$ are exact sample means of n/V i.i.d. random variables. Choosing V small enough that n/V is moderately large improves the normal approximation for each W_v by the central limit theorem and therefore improves the t_{V-1} approximation.

For *asymptotically linear estimators* with a nonzero remainder $R(P_{n,-v}, P_0) = o_p(n^{-1/2})$, the t_{V-1} Studentization still acts directly on the dominant linear component, namely the fold-level averages W_v . Only the negligibility of the remainder relies on asymptotics, and smaller V helps here as well: larger fold sizes n/V make $\hat{\Psi}(P_{n,-v})$ closer to its population target and thereby shrink the remainder terms. In practice, one may therefore choose a

small V to improve the t_{V-1} approximation, accepting wider confidence intervals (because of the heavier tails of t_{V-1}) in exchange for more reliable coverage.

5.3.2 Diverging V

We study the asymptotic behavior of the V -fold jackknife variance estimator in the regime where the number of folds V diverges with n . Our objective is to characterize its consistency, convergence rate, and asymptotic linear representation. In this setting, the primary question is whether the fold-based construction consistently estimates the asymptotic variance $\sigma^2 = \mathbb{E}_{P_0}[\varphi^2]$ of the underlying estimator. We show that consistency holds under mild remainder conditions whenever $V \rightarrow \infty$, without requiring a specific growth rate (e.g., $V = \log n$). Thus, the usual Wald type confidence intervals for the asymptotic linear estimators can be constructed based on the estimated $\hat{\sigma}$

Theorem 5.6 (Consistency under $V \rightarrow \infty$). *Let $\{O_i\}_{i=1}^n$ be i.i.d. from P_0 . Suppose the estimator $\hat{\Psi}(P_n)$ satisfies*

$$\hat{\Psi}(P_n) - \Psi(P_0) = P_n\varphi + R(P_n, P_0),$$

where $\mathbb{E}_{P_0}[\varphi] = 0$, $\mathbb{E}_{P_0}[\varphi^2] = \sigma^2 < \infty$, and $R(P_n, P_0) = o_p(n^{-1/2})$. Partition the sample into V disjoint folds of equal size n/V , and define the V -fold jackknife variance estimator

$$\hat{\sigma}^2 = \frac{n}{V} \hat{S}_{Jack}^2, \quad \hat{S}_{Jack}^2 = \frac{1}{V-1} \sum_{v=1}^V (IC_{Jack}(v) - \hat{\Psi}_{Jack}(P_n))^2.$$

Assume also that $V \rightarrow \infty$ with $V = o(n)$. Under the standard second-order remainder expansion shared by TMLE, one-step, and AIPW estimators, Section C.8 shows that this rate implies the stability condition $\|R(P_n, P_0) - R(P_{n,-v}, P_0)\|_2 = o_p(n^{-1/2}V^{-1})$. Then

$$\hat{\sigma}^2 - \sigma^2 = \left(\frac{1}{n} \sum_{i=1}^n \varphi_i^2 - \sigma^2 \right) + T_n + o_p(V^{-1/2}),$$

where $T_n = \frac{1}{n} \sum_{v=1}^V \sum_{\substack{i,j \in I_v \\ i \neq j}} \varphi_i \varphi_j$. If $V \rightarrow \infty$, then $\hat{\sigma}^2 \xrightarrow{p} \sigma^2$. Moreover, if $V = o(n)$,

$$T_n = o_p(1/\sqrt{V}), \quad \sqrt{V}(\hat{\sigma}^2 - \sigma^2) = O_p(1).$$

Proof location. The proof is given in Appendix C.1.5.

5.4 Simultaneous V -fold Jackknife Confidence Bands

We now extend the V -fold jackknife framework from a scalar parameter $\Psi(P_0) \in \mathbb{R}$ to a vector-valued parameter $\Psi(P_0) = (\Psi_1(P_0), \dots, \Psi_m(P_0))' \in \mathbb{R}^m$, where m is a fixed positive

integer (not growing with n). Throughout this section, φ_j denotes the influence function for the j th component, consistent with the notation φ used in Section 2 for the scalar case. In Section 5.5, we show that all results in this section extend to the generalized asymptotic linearity framework of Assumption 5.10, where the influence function variance may diverge. A leading example is the bivariate survival function evaluated on a grid: $\Psi_j(P_0) = S(t_{1,j}, t_{2,j}) = P_0(T_1 > t_{1,j}, T_2 > t_{2,j})$ for $j = 1, \dots, m$. The goal is to construct a simultaneous confidence band $[\hat{\Psi}_j \pm q_\alpha \cdot \widehat{SE}_j]$ that covers all m components simultaneously with probability at least $1 - \alpha$.

5.4.1 Vector-valued V-fold jackknife

Suppose that for each $j = 1, \dots, m$, the estimator $\hat{\Psi}_j(P_n)$ admits an asymptotically linear expansion

$$\hat{\Psi}_j(P_n) - \Psi_j(P_0) = P_n \varphi_j + R_{n,j},$$

where $\varphi_j = \varphi_{j,P_0}$ is a mean-zero influence function with $\mathbb{E}_{P_0}[\varphi_j^2] = \sigma_j^2 < \infty$ and $R_{n,j} = o_p(n^{-1/2})$. Define the $m \times m$ true covariance matrix $\Sigma = (\sigma_{jk})$ with $\sigma_{jk} = \mathbb{E}_{P_0}[\varphi_j \varphi_k]$.

For each fold $v = 1, \dots, V$ and component $j = 1, \dots, m$, the V -fold jackknife pseudo-value is

$$IC_{Jack}^{(j)}(v) = V \hat{\Psi}_j(P_n) - (V-1) \hat{\Psi}_j(P_{n,-v}).$$

Let $\hat{\Psi}_{Jack}^{(j)}(P_n) = V^{-1} \sum_{v=1}^V IC_{Jack}^{(j)}(v)$ denote the componentwise jackknife estimator, and define the V -fold jackknife covariance matrix estimator

$$\hat{\Sigma}_{Jack}(j, k) = \frac{1}{V-1} \sum_{v=1}^V \left(IC_{Jack}^{(j)}(v) - \hat{\Psi}_{Jack}^{(j)}(P_n) \right) \left(IC_{Jack}^{(k)}(v) - \hat{\Psi}_{Jack}^{(k)}(P_n) \right). \quad (5.4)$$

The componentwise jackknife variance and standard error are

$$\hat{\sigma}_j^2 = \hat{\Sigma}_{Jack}(j, j), \quad \widehat{SE}_j = \frac{\hat{\sigma}_j}{\sqrt{V}},$$

and the jackknife correlation matrix is

$$\hat{R}_{Jack}(j, k) = \frac{\hat{\Sigma}_{Jack}(j, k)}{\sqrt{\hat{\Sigma}_{Jack}(j, j) \hat{\Sigma}_{Jack}(k, k)}}. \quad (5.5)$$

Theorem 5.7 (Multivariate asymptotic equivalence, fixed V). *Let $\text{Var}(\varphi_j(O)) = \sigma_j^2$ and $\mathbb{E}_{P_0}[\varphi_j \varphi_k] = \sigma_{jk}$ for each $j, k = 1, \dots, m$, where m is fixed. Suppose that for each component j , the remainder terms satisfy $\max_{1 \leq v \leq V} |R_j(P_{n,-v}, P_0)| = o_p(n^{-1/2})$. Define the oracle covariance estimator $\hat{S}_W^2(j, k) = (V-1)^{-1} \sum_{v=1}^V (W_v^{(j)} - \bar{W}^{(j)})(W_v^{(k)} - \bar{W}^{(k)})$ where $W_v^{(j)} = P_{n,v}^1 \varphi_j$. Then, for V fixed:*

$$\frac{n}{V} \hat{\Sigma}_{Jack}(j, k) = \frac{n}{V} \hat{S}_W^2(j, k) + o_p(1).$$

Furthermore, when folds are of exactly equal size n/V , $\frac{n}{V}\hat{S}_W^2(j, k)$ is an unbiased estimator of σ_{jk} with $\mathbb{E}[\frac{n}{V}\hat{S}_W^2(j, k)] = \sigma_{jk}$.

Proof location. The proof is given in Appendix C.1.6.

5.4.2 Joint distribution of the Studentized vector and simultaneous critical values

When V is fixed, Theorem 5.2 established that the scalar Studentized statistic converges in distribution to t_{V-1} . For simultaneous inference over m components, we characterize the joint distribution of the componentwise-Studentized vector. Define

$$\mathbf{T} = \left(\frac{\sqrt{V}(\hat{\Psi}_{Jack}^{(1)}(P_n) - \Psi_1(P_0))}{\hat{\sigma}_1}, \dots, \frac{\sqrt{V}(\hat{\Psi}_{Jack}^{(m)}(P_n) - \Psi_m(P_0))}{\hat{\sigma}_m} \right)'.$$

Theorem 5.8 (Componentwise-Studentized distribution, fixed V). *Suppose the conditions of Theorem 5.7 hold. Let $R_0 = D^{-1}\Sigma D^{-1}$ denote the true correlation matrix, where $D = \text{diag}(\sigma_1, \dots, \sigma_m)$. As $n \rightarrow \infty$ with V fixed, the vector $\mathbf{T}$ converges in distribution: $\mathbf{T} \Rightarrow_d \mathbf{T}^\infty$, where*

$$\mathbf{T}^\infty = \left(\frac{\sqrt{V} \bar{Z}^{(1)}}{S_Z^{(1)}}, \dots, \frac{\sqrt{V} \bar{Z}^{(m)}}{S_Z^{(m)}} \right)', \quad (5.6)$$

with $\mathbf{Z}_1, \dots, \mathbf{Z}_V \stackrel{i.i.d.}{\sim} \mathcal{N}_m(\mathbf{0}, R_0)$, $\bar{\mathbf{Z}} = V^{-1} \sum_{v=1}^V \mathbf{Z}_v$, and $S_Z^{(j)} = [(V-1)^{-1} \sum_{v=1}^V (Z_v^{(j)} - \bar{Z}^{(j)})^2]^{1/2}$. Each marginal T_j^∞ has an exact t_{V-1} distribution.

Proof location. The proof is given in Appendix C.1.7.

The asymptotic normality of the fold-level vectors $\mathbf{W}_v$ places the problem in the classical framework of simultaneous inference for the mean of a multivariate normal distribution. However, a crucial distinction arises for the componentwise-Studentized vector $\mathbf{T}$. While a standard multivariate t -distribution models a vector scaled by a single shared χ^2 denominator [Type I in the classification of Kotz and Nadarajah, 2004], the limit distribution $\mathbf{T}^\infty$ features m dependent denominators corresponding to the diagonal elements of a sample covariance matrix (which follows a Wishart distribution). Because the sample variances $(S_Z^{(1)})^2, \dots, (S_Z^{(m)})^2$ are correlated but not identically equal, approximating this joint distribution with a standard multivariate t -distribution using a single shared denominator will lead to an incorrect limit and substantial undercoverage, particularly when V is small. Thus, simultaneous critical values must be simulated directly from the correct joint law of $\mathbf{T}^\infty$, which accurately preserves the dependency structure of both the numerators and the specific denominators.

The confidence band (5.7) below is centered at $\hat{\Psi}_j(P_n)$, while the Studentized statistic uses $\hat{\Psi}_{Jack}^{(j)}(P_n)$. By the same argument as in Theorem 5.3, applied componentwise, $\hat{\Psi}_{Jack}^{(j)}(P_n) -$

Algorithm 6 Monte Carlo computation of the simultaneous critical value $\hat{q}_\alpha$

- 1: **Input:** Estimated correlation matrix $\hat{R}_{Jack} \in \mathbb{R}^{m \times m}$, number of folds V , confidence level $1 - \alpha$, Monte Carlo size B (e.g., $B = 10,000$).
 - 2: **Output:** Critical value $\hat{q}_\alpha$ for the simultaneous confidence band.
 - 3: Compute eigendecomposition $\hat{R}_{Jack} = U\Lambda U'$.
 - 4: **for** $b = 1$ **to** B **do**
 - 5: **for** $v = 1$ **to** V **do**
 - 6: Draw $\mathbf{Z}_v^{(b)} \sim \mathcal{N}_m(\mathbf{0}, \hat{R}_{Jack})$ via $\mathbf{Z}_v^{(b)} = U\Lambda^{1/2}\boldsymbol{\eta}_v^{(b)}$, where $\boldsymbol{\eta}_v^{(b)} \sim \mathcal{N}_m(\mathbf{0}, I_m)$. (This procedure remains valid when $m > V - 1$ and $\hat{R}_{Jack}$ is singular; see section 5.4.4.)
 - 7: **end for**
 - 8: Compute the sample mean vector $\bar{\mathbf{Z}}^{(b)} = \frac{1}{V} \sum_{v=1}^V \mathbf{Z}_v^{(b)}$.
 - 9: **for** $j = 1$ **to** m **do**
 - 10: Compute the sample standard deviation $S_Z^{(j,b)} = \left[\frac{1}{V-1} \sum_{v=1}^V \left(Z_{v,j}^{(b)} - \bar{Z}_j^{(b)} \right)^2 \right]^{1/2}$.
 - 11: Form the Studentized statistic $T_j^{(b)} = \sqrt{V} \bar{Z}_j^{(b)} / S_Z^{(j,b)}$.
 - 12: **end for**
 - 13: Record $M^{(b)} = \max_{1 \leq j \leq m} |T_j^{(b)}|$.
 - 14: **end for**
 - 15: **Return** $\hat{q}_\alpha = Q_{1-\alpha}(\{M^{(b)}\}_{b=1}^B)$, where Q_p denotes the p -th sample quantile.
-

$\hat{\Psi}_j(P_n) = -\hat{b}_j$ where $\hat{b}_j = (V - 1)(\bar{\Psi}_{j,(-v)} - \hat{\Psi}_j(P_n))$ is the componentwise jackknife bias estimate, and $\hat{b}_j = o_p(\widehat{SE}_j)$. Hence the centering mismatch does not affect coverage.

A $(1 - \alpha)$ -level simultaneous confidence band is

$$\hat{\Psi}_j \pm q_\alpha \cdot \widehat{SE}_j, \quad j = 1, \dots, m, \quad (5.7)$$

where q_α satisfies

$$P\left(\max_{1 \leq j \leq m} |T_j^\infty| \leq q_\alpha\right) = 1 - \alpha. \quad (5.8)$$

Since the distribution of $\max_j |T_j^\infty|$ does not admit a closed form for general R_0 and m , the critical value q_α is computed by Monte Carlo simulation.

5.4.3 Monte Carlo computation of the critical value

Given the estimated correlation matrix $\hat{R}_{Jack}$, the critical value q_α is approximated by simulating directly from the joint componentwise-Studentized limit distribution $\mathbf{T}^\infty$. This relies on V independent Gaussian draws per Monte Carlo step to ensure the correct Wishart-diagonal denominator structure is enforced. The procedure is given in Algorithm 6.

Under the asymptotic normality of the fold-level vectors, the sample correlation $\hat{R}_{Jack}$ is the maximum likelihood estimate of R_0 based on V independent $\mathcal{N}_m$ observations. As $V \rightarrow \infty$, $\hat{R}_{Jack} \xrightarrow{p} R_0$, $S_Z^{(j)} \xrightarrow{p} 1$, and the joint componentwise-Studentized distribution

converges to the supremum of $\mathcal{N}_m(\mathbf{0}, R_0)$.

5.4.4 Validity of the eigendecomposition and the role of V in simultaneous inference

When $m > V - 1$, the estimated correlation matrix $\hat{R}_{Jack}$ has rank at most $V - 1 < m$, and the Monte Carlo draws in Step 1 are generated through the eigendecomposition $\mathbf{Z}_v^{(b)} = U\Lambda^{1/2}\boldsymbol{\eta}_v^{(b)}$. This is exact sampling from the degenerate plug-in distribution $\mathcal{N}_m(\mathbf{0}, \hat{R}_{Jack})$: the zero-eigenvalue directions carry no variation and are therefore correctly excluded from the simulation.

The real inferential question is not whether this eigendecomposition is legitimate, but whether $\hat{R}_{Jack}$ is close enough to the true correlation matrix R_0 . When V is fixed, Theorem 5.8 identifies the correct componentwise-Studentized limit distribution conditional on the true correlation matrix R_0 , and Theorem 5.7 shows that the jackknife covariance estimator is asymptotically unbiased. This still falls short of a formal simultaneous-coverage theorem for the plug-in procedure, because consistency of $\hat{R}_{Jack}$ is missing. Formal plug-in simultaneous validity therefore requires consistent estimation of R_0 and hence $V \rightarrow \infty$. For the variance results developed here, no particular divergence rate is needed, so one may let V increase arbitrarily slowly, for example $V \asymp \log n$.

To understand why the plug-in procedure can nevertheless work well when V is small, it is useful to introduce the notion of *effective rank*. For a positive semidefinite matrix A , define

$$r(A) = \frac{\text{tr}(A)}{\|A\|_{\text{op}}}, \quad (5.9)$$

where $\|A\|_{\text{op}}$ denotes the operator norm. For a correlation matrix R_0 , the trace is exactly m , so $r(R_0) = m/\lambda_{\max}(R_0)$. When the eigenvalues of R_0 decay rapidly, the effective rank can be much smaller than the ambient dimension m .

Concentration inequalities for sample covariance operators show that convergence rates depend on this effective rank rather than directly on m . Let $\hat{\Sigma}_{Jack}$ denote the sample covariance matrix of the fold-level vectors and let Σ_0 denote the corresponding population covariance. Results of Vershynin [2012] and Koltchinskii and Lounici [2017] imply

$$\mathbb{E}[\|\hat{\Sigma}_{Jack} - \Sigma_0\|_{\text{op}}] \leq C \cdot \|\Sigma_0\|_{\text{op}} \left(\sqrt{\frac{r(\Sigma_0)}{V}} + \frac{r(\Sigma_0)}{V} \right), \quad (5.10)$$

for a universal constant C . Because the sample correlation matrix is obtained by symmetric diagonal scaling, $\hat{R}_{Jack} = \hat{D}^{-1}\hat{\Sigma}_{Jack}\hat{D}^{-1}$ with $\hat{D} = \text{diag}(\hat{\Sigma}_{Jack})^{1/2}$, standard algebraic decompositions show that $\|\hat{R}_{Jack} - R_0\|_{\text{op}}$ is bounded by a constant multiple of $\|\hat{\Sigma}_{Jack} - \Sigma_0\|_{\text{op}}$ whenever the marginal variances are bounded away from zero. Thus the same effective-rank rate carries over to the correlation scale.

Consequently, even when $V \ll m$, the sample correlation matrix $\hat{R}_{Jack}$ can still be close to R_0 whenever V is large relative to the effective rank $r(R_0)$. For smooth functional parameters

such as survival curves or dose-response curves evaluated on a grid, adjacent evaluation points are highly correlated and the eigenvalues of R_0 often decay quickly. In that case, the dominant eigendirections of R_0 are captured by the rank-deficient $\hat{R}_{Jack}$ even when V is moderate, which helps explain the robust finite-sample performance seen in our simulations. This bound-based explanation supports the practical use of small V , but it does not replace the formal asymptotic requirement that $V \rightarrow \infty$ for plug-in simultaneous validity.

This perspective also suggests a practical finite-sample tuning rule. Along a candidate sequence such as $V \in \{5, 10, 20, 40, \dots\}$, compute the corresponding estimated correlation matrix $\hat{R}_{Jack}$ and its effective rank

$$\hat{r}_V := r(\hat{R}_{Jack}) = \frac{\text{tr}(\hat{R}_{Jack})}{\|\hat{R}_{Jack}\|_{\text{op}}} = \frac{m}{\lambda_{\max}(\hat{R}_{Jack})}, \quad (5.11)$$

where the final equality uses that $\hat{R}_{Jack}$ is a correlation matrix and therefore has trace m . One may then choose the smallest V such that

$$V \geq 2\hat{r}_V. \quad (5.12)$$

This rule should be viewed as a heuristic implementation of the effective-rank argument, not as a substitute for the formal asymptotic requirement above. When such tuning is not carried out, the simulations in Sections 5.6.2 and 5.6.3 suggest that $V = 20$ is a reasonable default in the settings we study.

For functional estimators with growing dimension, a natural next step would be to establish uniform convergence of the jackknife correlations to the true correlation function, which would then support Gaussian-process-style simultaneous inference.

5.4.5 Remark on links to high-dimensional Gaussian approximation

Our simultaneous-inference problem is related to the literature on Gaussian and multiplier-bootstrap approximation for maxima of high-dimensional means. However, those results are not needed for the development of the main theory here and require additional growth conditions linking V and m . Since the key message is that formal simultaneous validity follows once $V \rightarrow \infty$ —possibly at an arbitrarily slow rate such as $V \asymp \log n$ —we defer this connection to Appendix C.2.

5.4.6 Bias-corrected simultaneous confidence bands

When the asymptotic linear expansion is not uniformly negligible across all m grid points, the simultaneous coverage of (5.7) can be eroded by finite-sample bias. Following earlier work on bootstrapping HAL estimators [e.g., Cai and van der Laan, 2020], we address this by *enlarging* the confidence intervals to account for the estimated bias. Extending the scalar

jackknife bias estimate (5.2) to the j th component, define

$$\hat{b}_j = (V - 1) \left(\bar{\Psi}_{j,(-v)} - \hat{\Psi}_j(P_n) \right), \quad \bar{\Psi}_{j,(-v)} = \frac{1}{V} \sum_{v=1}^V \hat{\Psi}_j(P_{n,-v}). \quad (5.13)$$

This estimate is computable from the same fold-level estimates used for variance estimation and requires no additional computation; it is reused, as a post-hoc reliability diagnostic, in Appendix C.6. The bias-corrected simultaneous confidence band is

$$\left[\hat{\Psi}_j(P_n) - q_\alpha \cdot \widehat{SE}_j - |\hat{b}_j|, \quad \hat{\Psi}_j(P_n) + q_\alpha \cdot \widehat{SE}_j + |\hat{b}_j| \right], \quad j = 1, \dots, m, \quad (5.14)$$

where $\hat{b}_j$ is the jackknife bias estimate defined in (5.13).

Theorem 5.9 (Bias-corrected simultaneous coverage). *Suppose the conditions of Theorem 5.8 hold. Additionally, suppose the jackknife bias estimate satisfies $\max_{1 \leq j \leq m} |\hat{b}_j - b_j| = o_p(\widehat{SE}_j)$, where $b_j = \mathbb{E}[\hat{\Psi}_j(P_n)] - \Psi_j(P_0)$ is the true bias. Then*

$$P\left(\Psi_j(P_0) \in \left[\hat{\Psi}_j(P_n) - q_\alpha \widehat{SE}_j - |\hat{b}_j|, \quad \hat{\Psi}_j(P_n) + q_\alpha \widehat{SE}_j + |\hat{b}_j| \right] \right. \\ \left. \forall j = 1, \dots, m \right) \geq 1 - \alpha + o(1).$$

Proof location. The proof is given in Appendix C.1.8.

The bias correction may also be paired with the post-hoc diagnostic of Appendix C.6, which uses the same jackknife bias estimate $\hat{b}_j$ to flag grid points whose linearization is under stress, complementing the widening provided by the bias-corrected band.

5.4.7 Algorithm

We summarize the complete procedure for constructing simultaneous V -fold jackknife confidence bands. The univariate procedure (Algorithm 5) was presented in Section 4; the simultaneous extension is given in Algorithm 7.

5.5 Extension to Generalized Asymptotically Linear Estimators and Slower Convergence Rates

The theory developed in the preceding sections assumes the target parameter $\Psi(P_0)$ is pathwise differentiable, meaning the estimator converges at the parametric $n^{-1/2}$ rate with an influence function φ having finite variance $\sigma^2 = \mathbb{E}_{P_0}[\varphi^2] < \infty$. However, in many modern semiparametric and nonparametric applications, researchers target non-pathwise differentiable parameters. A prominent example is the causal dose-response curve for a continuous treatment evaluated at a specific point, estimated via the highly adaptive lasso (HAL) [van der Laan, 2023, Shi et al., 2024].

Algorithm 7 Simultaneous V -fold jackknife confidence band

- 1: **Input:** Estimator $\hat{\Psi} = (\hat{\Psi}_1, \dots, \hat{\Psi}_m)$, observations $O_1, \dots, O_n$, number of folds $V \geq 2$, confidence level $1 - \alpha$, Monte Carlo size B (e.g., $B = 10,000$).
 - 2: **Output:** $(1 - \alpha)$ -level simultaneous confidence band for $(\Psi_1(P_0), \dots, \Psi_m(P_0))$.
 - 3: Compute $\hat{\Psi}_j(P_n)$ for $j = 1, \dots, m$ using all n observations.
 - 4: Partition $\{1, \dots, n\}$ into V disjoint folds $I_1, \dots, I_V$ of approximately equal size.
 - 5: **for** $v = 1$ **to** V **do**
 - 6: Compute $\hat{\Psi}_j(P_{n,-v})$ for all $j = 1, \dots, m$.
 - 7: **for** $j = 1$ **to** m **do**
 - 8: Compute the pseudo-value $IC_{Jack}^{(j)}(v) = V\hat{\Psi}_j(P_n) - (V - 1)\hat{\Psi}_j(P_{n,-v})$.
 - 9: **end for**
 - 10: **end for**
 - 11: Compute the jackknife covariance matrix $\hat{\Sigma}_{Jack}$ via (5.4), the standard errors $\widehat{SE}_j = \hat{\sigma}_j/\sqrt{V}$, and the correlation matrix $\hat{R}_{Jack}$ via (5.5).
 - 12: Compute the critical value $\hat{q}_\alpha$ via Algorithm 6.
 - 13: **Return** the simultaneous confidence band: $\hat{\Psi}_j(P_n) \pm \hat{q}_\alpha \cdot \widehat{SE}_j$ for $j = 1, \dots, m$. Optionally, apply the bias correction (5.14). See Appendix C.6 for a diagnostic that highlights grid points with large $|\hat{b}_j|/\widehat{SE}_j$.
-

In such settings, the estimator may still admit a generalized asymptotic linear representation, but the influence curve depends on the sample size n , and its variance diverges. For instance, the variance of the influence function for the first-order HAL dose-response curve scales with the effective dimension J_n of the working model (i.e., $\sigma_n^2 = \text{Var}_{P_0}(\varphi_n(O)) \asymp J_n \rightarrow \infty$), yielding a slower stabilization rate of $\sqrt{n/J_n}$.

A profound mathematical and operational advantage of the V -fold jackknife is its inherent *scale invariance*. Because the jackknife pseudo-values are computed directly from the estimator, their empirical dispersion automatically tracks the inflating variance of the influence curve. Consequently, the unknown diverging rate $\sigma_n \asymp \sqrt{J_n}$ perfectly cancels out in the Studentized t -statistic. This allows practitioners to obtain valid confidence intervals using the exact same jackknife procedure without needing to analytically derive, estimate, or explicitly divide by the effective dimension J_n .

We formalize this by extending our fixed- V and diverging- V results to a sequence of sample-size-dependent influence functions.

Assumption 5.10 (Generalized Asymptotic Linearity). *Suppose the estimator admits the expansion*

$$\hat{\Psi}(P_n) - \Psi(P_0) = P_n \varphi_n + R_n(P_n, P_0), \quad (5.15)$$

where $\varphi_n = \varphi_{n, P_0}$ is a sequence of mean-zero influence functions with diverging variance $\sigma_n^2 = \mathbb{E}_{P_0}[\varphi_n^2(O)] \rightarrow \infty$ as $n \rightarrow \infty$. We assume:

1. **Scaled Remainder:** The remainder terms are asymptotically negligible relative to the standard error: $R_n(P_n, P_0) = o_p(\sigma_n n^{-1/2})$ and $\max_{1 \leq v \leq V} |R_n(P_{n,-v}, P_0)| =$

$o_p(\sigma_n n^{-1/2})$.

2. **Lindeberg Condition:** The standardized influence functions satisfy the Lindeberg condition: for any $\epsilon > 0$,

$$\lim_{n \rightarrow \infty} \frac{1}{\sigma_n^2} \mathbb{E}_{P_0} \left[\varphi_n^2(O) \mathbb{1} \left\{ |\varphi_n(O)| > \epsilon \sigma_n \sqrt{n} \right\} \right] = 0.$$

Remark 5.11. Because $\sigma_n \rightarrow \infty$, the generalized remainder condition $o_p(\sigma_n n^{-1/2})$ is strictly weaker than the standard $o_p(n^{-1/2})$ requirement. This mathematical relaxation aligns with the intuition that when the target parameter converges at a slower rate (i.e., has a larger standard error), a proportionally larger finite-sample bias is statistically tolerable for valid coverage. This perfectly mirrors the logic in the uniform HAL results of van der Laan [2023]: slower convergence rates tolerate slightly larger finite-sample biases.

5.5.1 Fixed- V inference for slower rates

We first generalize the fixed- V framework.

Theorem 5.12 (Fixed- V Inference for Slower Rates). *Under Assumption 5.10, for any fixed $V \geq 2$, the Studentized V -fold jackknife statistic adapts to the unknown convergence rate σ_n , yielding*

$$\frac{\sqrt{V} (\hat{\Psi}_{Jack}(P_n) - \Psi(P_0))}{\hat{S}_{Jack}} \implies_d t_{V-1}.$$

Proof location. The proof is given in Appendix C.1.9.

Corollary 5.13 (Confidence interval centered at $\hat{\Psi}(P_n)$). *Under the conditions of Theorem 5.12,*

$$\frac{\sqrt{V} (\hat{\Psi}(P_n) - \Psi(P_0))}{\hat{S}_{Jack}} \implies_d t_{V-1}.$$

Hence, a valid $(1-\alpha)$ -level confidence interval for $\Psi(P_0)$ is given by $\hat{\Psi}(P_n) \pm t_{V-1, 1-\alpha/2} \cdot \frac{\hat{S}_{Jack}}{\sqrt{V}}$.

Proof location. The proof is given in Appendix C.1.10.

5.5.2 Diverging- V consistency for slower rates

The diverging- V consistency also naturally extends to this regime. We show that the relative ratio of the jackknife variance estimator to the true diverging variance converges to one.

Theorem 5.14 (Diverging- V Consistency for Slower Rates). *Under Assumption 5.10, define the scaled variance estimator $\hat{s}_n^2 = \frac{n}{V} \hat{S}_{Jack}^2$. Suppose that a Weak Law of Large Numbers holds*

for the squared influence curve: $\frac{1}{n\sigma_n^2} \sum_{i=1}^n \varphi_n^2(O_i) \xrightarrow{p} 1$ (which is satisfied if $\mathbb{E}_{P_0}[\varphi_n^4(O)] = o(n\sigma_n^4)$). Assume $V \rightarrow \infty$ with $V = o(n/\sigma_n^4)$, so that the remainder stability condition

$$\|R_n(P_n, P_0) - R_n(P_{n,-v}, P_0)\|_2 = o_p\left(\frac{\sigma_n}{\sqrt{n}V}\right)$$

holds. For the HAL dose-response curve this yields $V = o(n^{1/3})$ for 0-order HAL ($\sigma_n \asymp n^{1/6}$) and $V = o(n^{3/5})$ for first-order HAL ($\sigma_n \asymp n^{1/10}$); see Section C.8 for the derivation. If $V \rightarrow \infty$ with $V = o(n)$ as $n \rightarrow \infty$, then

$$\frac{\hat{s}_n^2}{\sigma_n^2} \xrightarrow{p} 1.$$

Proof location. The proof is given in Appendix C.1.11.

5.5.3 Generalized simultaneous confidence bands

The simultaneous confidence band results developed in the preceding section also extend to the generalized asymptotic linearity framework of Assumption 5.10. Consider a vector-valued parameter $\Psi(P_0) = (\Psi_1(P_0), \dots, \Psi_m(P_0))'$ where each component admits

$$\hat{\Psi}_j(P_n) - \Psi_j(P_0) = P_n \varphi_{n,j} + R_{n,j}(P_n, P_0),$$

with $\sigma_{n,j}^2 = \mathbb{E}_{P_0}[\varphi_{n,j}^2] \rightarrow \infty$, possibly at different rates for different j . The componentwise Studentization eliminates $\sigma_{n,j}$ from each component independently:

$$T_j = \frac{\sqrt{V}(\hat{\Psi}_{jack}^{(j)}(P_n) - \Psi_j(P_0))}{\hat{\sigma}_j} \Rightarrow_d T_j^\infty \sim t_{V-1},$$

with the joint distribution determined by the correlation matrix $R_{0,n} = D_n^{-1} \Sigma_n D_n^{-1}$, where $D_n = \text{diag}(\sigma_{n,1}, \dots, \sigma_{n,m})$ and $\Sigma_n(j, k) = \mathbb{E}_{P_0}[\varphi_{n,j} \varphi_{n,k}]$. This correlation matrix has entries in $[-1, 1]$ regardless of the scale of the variances. Consequently, all results in the preceding section—including the componentwise-Studentized distribution (Theorem 5.8), the Monte Carlo critical value algorithm, and the bias-corrected bands (Theorem 5.9)—hold verbatim under Assumption 5.10 applied componentwise, since all proofs are scale-invariant.

5.5.4 Application: HAL dose-response curve

The causal dose-response curve $\Psi(a) = \mathbb{E}_{P_0}[\mathbb{E}_{P_0}(Y \mid A = a, W)]$ estimated via the HAL plug-in $\hat{\Psi}(a) = \frac{1}{n} \sum_{i=1}^n \hat{Q}_n(W_i, a)$ provides the primary example of the generalized framework, where $\hat{Q}_n$ denotes the highly adaptive lasso estimator for the conditional mean outcome $Q_0(W, A) = \mathbb{E}_{P_0}[Y \mid A, W]$.

To establish the pointwise asymptotic normality of the HAL estimator without delving into the complex proof details of L_1 -penalization, we outline the logical progression developed by van der Laan [2023]. The analysis centers on the construction of a designed oracle working model. Let $\mathcal{R}_{0,n}$ denote a fixed, oracle set of spline basis functions. This set is designed to be as sparse as possible—meaning its dimension $d_n = |\mathcal{R}_{0,n}|$ grows slowly with n —while still being sufficiently rich to maintain the optimal approximation error rate for the true function Q_0 .

Let $D(\mathcal{R}_{0,n})$ denote the finite-dimensional linear working model spanned by these d_n basis functions, and let $Q_{0,n} = \operatorname{argmin}_{Q \in D(\mathcal{R}_{0,n})} P_0 L(Q)$ be the theoretical oracle maximum likelihood estimator (MLE) on this model under a loss function L . A standard analysis of the unpenalized empirical MLE on this fixed working model readily establishes its asymptotic normality.

The key observation in van der Laan [2023] is that the actual HAL estimator $\hat{Q}_n$, despite selecting its basis adaptively via L_1 -regularization, effectively solves the unpenalized score equations associated with the basis functions in the oracle model $D(\mathcal{R}_{0,n})$. Because $\hat{Q}_n$ solves the scores of this oracle model, it is asymptotically equivalent to the oracle MLE and seamlessly inherits its asymptotic normality. This yields the generalized asymptotic linear representation:

$$(\hat{Q}_n - Q_{0,n})(x) \approx (P_n - P_0)D_{Q_{0,n},x},$$

where $x = (W, a)$, and the influence-curve-like object $D_{Q_{0,n},x} = \sum_{j \in \mathcal{R}_{0,n}} S_{Q_{0,n}}(\phi_j^*) \phi_j^*(x)$ is a sum over d_n terms. Here, $\{\phi_j^*\}_{j \in \mathcal{R}_{0,n}}$ forms an orthonormalized basis of $D(\mathcal{R}_{0,n})$ in $L^2(P_0)$, and $S_{Q_{0,n}}(\phi) = \left. \frac{d}{d\epsilon} L(Q_{0,n} + \epsilon\phi) \right|_{\epsilon=0}$ is the corresponding score operator evaluated at $Q_{0,n}$.

The variance of this expansion, $\sigma_n^2(x) = P_0\{D_{Q_{0,n},x}\}^2 = d_n \cdot \tilde{\sigma}_n^2(x)$, diverges linearly with the dimension d_n of the oracle model, while the normalized variance $\tilde{\sigma}_n^2(x) = d_n^{-1} \sum_{j \in \mathcal{R}_{0,n}} \{\phi_j^*(x)\}^2$ remains bounded. Standardizing by this diverging standard error establishes the pointwise asymptotic normality of the HAL estimator at a slower convergence rate: $(n/d_n)^{1/2}(\hat{Q}_n - Q_{0,n})(x)/\tilde{\sigma}_n(x) \Rightarrow_d \mathcal{N}(0, 1)$. Integrating over the empirical distribution of W extends this pointwise linearization to the dose-response plug-in estimator $\hat{\Psi}(a)$.

Assumption 5.10 is verified for the HAL dose-response curve precisely through this oracle model formulation. First, the scaled remainder condition $R_n = o_p(\sigma_n n^{-1/2})$ holds because the bias $Q_{0,n} - Q_0$ of the working model projection vanishes at a rate controlled by the approximation error of the sparse oracle sieve, which is $o_p(\sqrt{d_n/n})$ for an appropriate choice of d_n [van der Laan, 2023]. Second, the Lindeberg condition is satisfied because $D_{Q_{0,n},x}$ is a sum of d_n terms involving bounded basis function contributions, which perfectly aligns with the setting where the Lindeberg–Feller central limit theorem for triangular arrays applies.

The V -fold jackknife confidence interval $\hat{\Psi}(a) \pm t_{V-1, 1-\alpha/2} \cdot \hat{S}_{Jack}(a)/\sqrt{V}$ therefore provides asymptotically valid pointwise inference for the dose-response curve. Crucially, the jackknife standard error $\hat{S}_{Jack}(a)/\sqrt{V} \asymp_p \sqrt{d_n/n}$ automatically adapts to the correct diverging rate without requiring the practitioner to know the effective dimension d_n . This scale invariance explains why the delta method—which requires explicitly extracting the active basis functions, computing the growing d_n -dimensional influence curve, and calculating its diverging

variance—can fail numerically (producing NAs at rates up to 6.2% in our simulations). In contrast, the V -fold jackknife reliably captures the diverging variance automatically through the fold-level variability of the pseudo-values. The simulation study in Section 5.6.3 provides empirical validation of this theoretical guarantee.

5.6 Simulation Studies

We evaluate the finite-sample performance of the V -fold jackknife through three simulation studies of increasing complexity. Section 5.6.1 validates the fixed- V t_{V-1} inference from Theorem 5.2 for a scalar parameter (the average treatment effect). Section 5.6.2 extends the evaluation to a functional parameter (the Kaplan–Meier survival curve), comparing pointwise and simultaneous confidence bands from Theorem 5.8. Section 5.6.3 demonstrates the practical value of the approach for HAL-based dose-response curve estimation, where analytic standard errors may be unreliable.

5.6.1 Scalar Inference: Average Treatment Effect

We first examine pointwise confidence intervals for the average treatment effect $\Psi(P_0) = \mathbb{E}_{P_0}[\mathbb{E}_{P_0}(Y \mid A = 1, W) - \mathbb{E}_{P_0}(Y \mid A = 0, W)]$ estimated by targeted minimum loss-based estimation (TMLE) with Super Learner. The data generating process draws eight binary and continuous covariates $W = (W_1, \dots, W_8)$, a binary treatment A , and a continuous outcome Y . We consider four scenarios in which the propensity score model g is correctly specified, varying whether the outcome regression Q is correctly specified or misspecified, and whether positivity violations are moderate or severe. Nuisance functions are estimated via Super Learner with library (`SL.glm`, `SL.step`, `SL.glm.interaction`). We report results for $n \in \{200, 1000\}$ based on 500 replications.

We compare four approaches to constructing 95% confidence intervals:

- (i) *EIC-based*: the plug-in influence function variance estimator with normal critical values, as implemented in the `tmle` package [Gruber and van der Laan, 2012];
- (ii) *Targeted bootstrap*: bootstrap standard errors ($B = 10,000$) with normal critical values;
- (iii) *Robust variance*: the robust variance estimator of Tran et al. [2023], as implemented in `ltmle`;
- (iv) *V-fold jackknife* ($V \in \{5, 10, 20\}$): the proposed estimator with t_{V-1} critical values.

All four methods use the same TMLE point estimate; only the variance estimation and critical values differ.

Table 5.1 summarizes coverage and CI width across all four scenarios. Under moderate positivity (Scenarios A and C), the EIC-based method achieves only 84–87% coverage at $n = 200$ because the plug-in variance estimator underestimates the true sampling variability when nuisance estimation is imprecise. Under severe positivity (Scenarios B and D), the

Table 5.1: Coverage (%) and mean CI width for the average treatment effect via TMLE. Scenarios A–D have the propensity score model g correctly specified. A: Q correct, moderate positivity; B: Q correct, severe positivity; C: Q misspecified, moderate positivity; D: Q misspecified, severe positivity. 500 replications.

Method	Scenario A		Scenario B		Scenario C		Scenario D	
	Cov	Width	Cov	Width	Cov	Width	Cov	Width
$n = 200$								
EIC-based	86.6	0.426	73.2	0.390	84.0	0.412	69.8	0.362
Targeted bootstrap	94.8	0.542	95.2	0.642	96.9	0.565	95.8	0.688
Robust variance	92.8	0.580	88.8	0.785	94.6	0.587	90.2	0.809
V-fold ($V=5$)	96.8	0.753	92.2	1.091	97.4	0.831	94.4	1.262
V-fold ($V=10$)	96.8	0.726	93.8	1.025	98.0	0.754	95.2	1.138
V-fold ($V=20$)	98.2	0.750	94.8	1.107	98.4	0.786	96.0	1.253
$n = 1000$								
EIC-based	90.8	0.222	81.8	0.234	91.2	0.224	79.6	0.237
Targeted bootstrap	93.8	0.252	96.4	0.395	95.6	0.258	96.6	0.402
Robust variance	93.6	0.241	95.8	0.377	93.8	0.247	94.8	0.379
V-fold ($V=5$)	94.0	0.339	95.0	0.492	94.2	0.340	94.2	0.501
V-fold ($V=10$)	96.2	0.287	96.2	0.443	97.6	0.301	95.6	0.444
V-fold ($V=20$)	95.6	0.282	96.0	0.432	96.0	0.286	96.0	0.440

EIC-based undercoverage worsens to 70–73%. In contrast, the V -fold jackknife with t_{V-1} critical values achieves 93–98% coverage across all four scenarios at $n = 200$, confirming the finite-sample validity of Theorem 5.2. The targeted bootstrap achieves near-nominal coverage uniformly, while the robust variance estimator undercovers slightly (89–93%) in the severe positivity settings. At $n = 1000$, the EIC-based method remains substantially below nominal under severe positivity (80–82%), whereas all other methods converge toward the 95% level.

Figure 5.1 illustrates the coverage–width tradeoff as a function of V . The top row shows that coverage remains above the nominal 95% level for all $V \geq 2$ at $n = 200$ in the moderate-positivity scenarios, while under severe positivity the V -fold jackknife still substantially outperforms the EIC-based method. The bottom row confirms that CI width decreases with V as the t_{V-1} quantiles approach the normal quantiles. This tradeoff is consistent with the theory: for fixed V , the variance estimator does not concentrate, and the t -distribution accounts for this additional uncertainty.

5.6.2 Survival Curve: Kaplan–Meier Estimator

We next evaluate both pointwise and simultaneous confidence bands for the Kaplan–Meier estimator of the survival function $S(t) = P(T > t)$. This setting provides a natural testbed because the Greenwood formula yields an analytic standard error for comparison. We generate right-censored survival data from exponential distributions with event rate $\lambda = 0.1$ and censoring rate $\lambda_c = 0.05$, yielding approximately 33% censoring. The sample size is $n = 200$, and we evaluate the survival function at $m = 91$ equally spaced time points on $[2, 20]$. Results are based on 500 replications.

Pointwise inference. Figure 5.2 and Table 5.2 compare pointwise coverage for the Greenwood formula with normal critical values, the V -fold jackknife with t_{V-1} critical values ($V \in \{5, 10, 20\}$), and the bootstrap Wald method ($B = 1000$). All methods achieve approximately 95% coverage, averaged over evaluation times, with the V -fold jackknife tracking the Greenwood formula closely. Table 5.2 confirms that the V -fold jackknife variance estimator is well calibrated for the Kaplan–Meier estimator even at moderate sample sizes: the V -fold t achieves 94.9% ($V = 5$), 94.8% ($V = 10$), and 94.7% ($V = 20$), comparable to the Greenwood formula (94.6%) and bootstrap (94.7%).

Table 5.2: Pointwise coverage (%) and mean CI width for the Kaplan–Meier survival function ($n = 200$, $m = 91$ evaluation points, 500 replications, averaged over all evaluation times).

Method	Coverage (%)	Mean Width
Greenwood (z)	94.6	0.145
V -fold t ($V=5$)	94.9	0.197
V -fold t ($V=10$)	94.8	0.164
V -fold t ($V=20$)	94.7	0.153
Bootstrap Wald	94.7	0.145




Figure 5.1: Coverage (top) and CI width (bottom) of V -fold jackknife confidence intervals as a function of the number of folds V , for four scenarios and two sample sizes. Horizontal dashed line in the top panels: nominal 95% level. Faded horizontal lines: coverage of the EIC-based and bootstrap methods for each scenario. The V -fold jackknife maintains coverage above the nominal level for moderate positivity and substantially outperforms the EIC-based method under severe positivity, with intervals narrowing as V increases.




Figure 5.2: Pointwise coverage of 95% confidence intervals for the Kaplan–Meier survival function across evaluation times ($n = 200$, 500 replications). The V -fold jackknife with t_9 critical values ($V = 10$) closely tracks the Greenwood formula, and the bootstrap Wald method performs comparably.

Simultaneous inference. Table 5.3 reports simultaneous coverage for six methods: the V -fold simultaneous band from Theorem 5.8 (using the componentwise-Studentized critical value from section 5.4.3), the Greenwood normal band (Greenwood SE with multivariate normal critical values), the classical Hall–Wellner and Equal Precision (EP) bands, the studentized bootstrap band [Davison and Hinkley, 1997], the bias-corrected V -fold band from Theorem 5.9, and a bias-corrected bootstrap (BC Bootstrap) band that widens the studentized bootstrap band by the absolute bootstrap bias $|\bar{S}_b^*(t) - \hat{S}(t)|$ at each evaluation point, where $\bar{S}_b^*(t) = B^{-1} \sum_{b=1}^B \hat{S}_b^*(t)$ is the mean of the bootstrap Kaplan–Meier estimates.

The V -fold simultaneous band achieves 93.4–94.6% coverage across all three choices of V , approaching the classical EP (96.6%) and Hall–Wellner (96.6%) bands. For $V = 20$ the band achieves 94.6% coverage with width 0.239, comparable to these established methods, while smaller V produces wider but slightly less precise bands: $V = 5$ achieves 93.6% coverage with mean width 0.454, and $V = 10$ achieves 93.4% with width 0.282. Larger V benefits from a more precise correlation matrix estimate $\hat{R}_{\text{Jack}}$ (since V is closer to the $m = 91$ evaluation points), which outweighs the effect of the lighter-tailed t_{V-1} critical values. The net effect is that all three choices of V achieve reasonable simultaneous coverage, with $V = 20$ offering the best coverage and narrowest bands. The bias-corrected bands yield coverage nearly identical to the standard bands (e.g., 94.8% versus 94.6% for $V = 20$) with negligible width increase (0.240 versus 0.239), reflecting the fact that the Kaplan–Meier estimator has minimal bias in this setting and the jackknife bias estimates $|\hat{b}_j|$ are effectively zero. Similarly, the BC

Table 5.3: Simultaneous coverage (%) and mean band width for the Kaplan–Meier survival function ($n = 200$, $m = 91$ evaluation points, 500 replications). “BC” denotes the bias-corrected band from section 5.4.6.

Method	Simult. Cov. (%)	Mean Width
Equal Precision	96.6	0.239
Hall–Wellner	96.6	0.251
Bootstrap	95.8	0.226
V-fold simultaneous ($V=20$)	94.6	0.239
V-fold simultaneous ($V=5$)	93.6	0.454
V-fold simultaneous ($V=10$)	93.4	0.282
Greenwood Normal	92.4	0.210
BC simultaneous ($V=20$)	94.8	0.240
BC simultaneous ($V=5$)	93.8	0.456
BC simultaneous ($V=10$)	93.4	0.283
BC Bootstrap	95.8	0.228

Bootstrap band achieves the same 95.8% coverage as the standard bootstrap with only marginally wider bands (0.228 versus 0.226), confirming that the bootstrap bias correction $|\bar{S}_b^*(t) - \hat{S}(t)|$ is also effectively zero for the Kaplan–Meier estimator.

Effective rank of the correlation matrix. Since $m = 91$ exceeds $V - 1$ for all choices of V considered here, the estimated correlation matrix $\hat{R}_{\text{Jack}}$ is singular and the simultaneous critical value is computed via the eigendecomposition described in section 5.4.4. To verify that the rank deficiency is benign, we estimate the true correlation matrix R_0 from the 500 simulation replications: the 500×91 matrix of Kaplan–Meier estimates across replications yields a sample correlation matrix whose eigendecomposition reveals the intrinsic dimensionality of the problem. The estimated effective rank is $\hat{r}(R_0) = m/\hat{\lambda}_{\max} = 91/53.8 \approx 1.7$, with the first two eigenvalues alone capturing 77% of the total variance and only 9 eigenvalues needed to explain 95% (Figure 5.5). This extremely low effective rank reflects the strong correlation among adjacent time points on a smooth survival curve: the Kaplan–Meier estimator at nearby time points is driven by essentially the same at-risk set, producing a nearly one-dimensional correlation structure. Even $V = 5$ folds far exceed the effective rank, confirming that the eigendecomposition accurately captures the dominant correlation structure of R_0 .

5.6.3 Dose-Response Curve: HAL Estimator

Our final study demonstrates the V -fold jackknife for inference on a dose-response curve estimated by the highly adaptive lasso (HAL). Unlike the pathwise differentiable parameters in Sections 5.6.1 and 5.6.2, the dose-response curve is a non-pathwise differentiable functional whose HAL plug-in estimator converges at the slower rate $\sqrt{n/J_n}$, where J_n is the number of non-zero HAL basis functions. The theoretical framework for this setting is developed in

Section 5.5, where we show that the V -fold jackknife naturally adapts to this non-standard rate through the scale invariance of Studentization. Unlike the Kaplan–Meier estimator, deriving and implementing analytic standard errors for HAL estimates is non-trivial: the delta method requires computing the influence curve of the HAL-based plug-in estimator, which can fail numerically. The V -fold jackknife bypasses this difficulty entirely, requiring only the ability to refit the estimator on leave-fold-out subsamples.

We consider four data generating processes of increasing complexity from Shi et al. [2024]: DGP 1 (smooth, monotonically increasing dose-response), DGP 2 (oscillating), DGP 3 (piecewise smooth with a transition at $A = 2$), and DGP 4 (discontinuous). All have binary outcomes Y , continuous treatment $A \in [0, 5]$, and a single confounder $W \sim N(0, 1)$. The dose-response curve is evaluated at 20 equally spaced points on $[0, 5]$ with $n = 500$ and 500 replications. All HAL fits use first-order (piecewise linear) basis functions, corresponding to smoothness order 1 in the framework of Shi et al. [2024], who also consider zero-order and adaptive choices. We consider two HAL fitting strategies: cross-validated HAL (CV-HAL) and undersmoothed HAL ($\lambda = \lambda_{CV}/\sqrt{\log n}$).

We compare three methods: the delta method using the HAL influence curve with normal critical values, the V -fold jackknife with t_{V-1} critical values ($V = 20$), and the bootstrap Wald method ($B = 200$).

Pointwise inference. Table 5.4 reports pointwise coverage, mean CI width, and the delta method failure rate (percentage of replications producing NA standard errors) across all four DGPs. For DGP 1 (smooth), all methods achieve near-nominal coverage (93–95%), with the V -fold jackknife at 93.7% (CV-HAL) and 95.0% (undersmoothed). For DGP 2 (oscillating), the delta method’s coverage drops sharply to 77.8% (CV-HAL) and 72.2% (undersmoothed), with failure rates of 0.8% and 6.2%, respectively. In contrast, the V -fold jackknife achieves 94.3% and 95.1%, demonstrating its robustness when the delta method is unreliable. For DGPs 3 (piecewise) and 4 (discontinuous), coverage is lower across all methods due to estimator bias, but the V -fold jackknife consistently provides the highest coverage: for undersmoothed HAL, 89.8% and 90.0% compared to 81.8% and 85.5% for the delta method, and 85.4% and 86.3% for the bootstrap.

Figure 5.3 shows pointwise coverage across the dose grid for all four DGPs. For DGP 1 with undersmoothed HAL, all methods maintain coverage near 95% uniformly over the evaluation points. For DGP 2, the delta method exhibits severe undercoverage throughout the dose range, while the V -fold jackknife remains near nominal. For DGPs 3 and 4, coverage deteriorates at specific dose levels where the true dose-response curve is non-smooth: for DGP 3, coverage dips appear for $A > 2$ (where oscillations begin), and for DGP 4, coverage dips concentrate at $A \approx 2$ and $A \approx 4$ (the discontinuity locations). As the bias diagnosis in Figure 5.4 confirms, these are precisely the points where the absolute bias of the first-order HAL estimator exceeds the estimated standard error, indicating that coverage loss reflects estimator bias rather than variance estimation failure. In all settings, the V -fold jackknife provides the highest coverage, reflecting the conservatism of the t_{V-1} critical values.

To diagnose the source of coverage loss in DGPs 3 and 4, Figure 5.4 displays the absolute bias $|\hat{\psi}(a) - \psi_0(a)|$ (averaged over replications) alongside the mean estimated standard error at each dose point, using the V -fold jackknife ($V = 20$). When the absolute bias exceeds




Figure 5.3: Pointwise coverage of 95% confidence intervals across the dose grid (20 equally spaced points on $[0, 5]$, i.e. $A \in \{0, 0.263, 0.526, \dots, 5\}$) for HAL dose-response curve estimation ($n = 500$, 500 replications). Each point on the horizontal axis represents a dose level at which the dose-response curve $\mathbb{E}[Y \mid A = a, W]$ is evaluated and coverage is computed separately. Rows: DGP 1 (smooth), DGP 2 (oscillating), DGP 3 (piecewise), DGP 4 (discontinuous). Columns: CV-HAL and undersmoothed HAL. The V-fold jackknife with t_{19} critical values ($V = 20$) provides the highest coverage across all DGPs. The delta method fails most notably for DGP 2 (oscillating), where its failure rate also exceeds 5% for undersmoothed HAL.

Table 5.4: Pointwise coverage (%), mean CI width, and delta method failure rate for HAL dose-response curve estimation ($n = 500$, 500 replications, averaged over 20 evaluation points).

DGP	Method	CV-HAL			Undersmoothed HAL		
		Cov	Width	NA%	Cov	Width	NA%
1 (smooth)	Delta	93.1	0.129	0.0	93.6	0.171	0.4
	V -fold t ($V=20$)	93.7	0.137	0.0	95.0	0.185	0.0
	Bootstrap	92.8	0.126	0.0	93.6	0.161	0.0
2 (oscillating)	Delta	77.8	0.105	0.8	72.2	0.122	6.2
	V -fold t ($V=20$)	94.3	0.177	0.0	95.1	0.214	0.0
	Bootstrap	90.5	0.138	0.0	92.8	0.161	0.0
3 (piecewise)	Delta	75.2	0.161	0.0	81.8	0.208	1.4
	V -fold t ($V=20$)	82.2	0.211	0.0	89.8	0.288	0.0
	Bootstrap	77.5	0.163	0.0	85.4	0.207	0.0
4 (discontinuous)	Delta	76.0	0.219	0.0	85.5	0.279	0.0
	V -fold t ($V=20$)	80.2	0.255	0.0	90.0	0.343	0.0
	Bootstrap	75.3	0.209	0.0	86.3	0.267	0.0

the standard error, the asymptotic linearity assumption underlying the confidence interval is stressed, and coverage loss is expected. For DGP 1 (smooth), the bias is uniformly small relative to the standard error across the entire dose grid, consistent with the near-nominal coverage in Figure 5.3. For DGP 2 (oscillating), the bias is similarly controlled, which explains why the V -fold jackknife achieves excellent coverage despite the complex curve shape—the first-order HAL can approximate the oscillating function adequately at this sample size. For DGP 3 (piecewise), the bias is small for $A \leq 2$ (where the true curve is smooth and nearly linear) but rises sharply for $A > 2$ (where the oscillating pattern begins), producing a bias-to-SE ratio that exceeds 1 at several dose points, particularly for CV-HAL. These are precisely the dose points where Figure 5.3 shows coverage dips. For DGP 4 (discontinuous), the bias spikes at $A \approx 2$ and $A \approx 4$ —the locations of the jump discontinuities—where the piecewise-linear HAL basis cannot represent the true function, again matching the coverage dip locations in Figure 5.3. These patterns confirm that the coverage loss in DGPs 3 and 4 is attributable to estimator bias at non-smooth points, not to a failure of the V -fold jackknife variance estimator.

Simultaneous inference. Table 5.5 reports simultaneous coverage for the delta method simultaneous band (multivariate normal with estimated IC covariance), the V -fold simultaneous band with $V \in \{5, 10, 20\}$ (componentwise-Studentized critical value from Theorem 5.8), the studentized bootstrap band [Davison and Hinkley, 1997], the bias-corrected bootstrap band, and the bias-corrected V -fold band from Theorem 5.9.

For DGP 1 (smooth), the standard V -fold simultaneous band achieves 89.8–91.4% coverage (CV-HAL) and 88.8–91.0% (undersmoothed), the highest among the uncorrected meth-




Figure 5.4: Bias diagnosis for HAL dose-response curve estimation ($n = 500$, 500 replications). Solid black line: absolute bias $|\hat{\psi}(a) - \psi_0(a)|$ at each of the 20 evaluation points, where $\hat{\psi}(a)$ is the simulation average of the point estimate. Dashed blue line: mean estimated standard error from the V -fold jackknife ($V = 20$). When the bias exceeds the standard error, the asymptotic linearity condition is violated and coverage loss is expected. The spatial pattern of bias peaks corresponds to the coverage dips in Figure 5.3: for DGP 3, bias rises for $A > 2$ where oscillations begin; for DGP 4, bias spikes at $A = 2$ and $A = 4$ where the true curve is discontinuous.

Table 5.5: Simultaneous coverage (%) and mean band width for HAL dose-response curve estimation ($n = 500$, 500 replications, 20 evaluation points). “BC” denotes the bias-corrected band from section 5.4.6.

DGP	Method	CV-HAL		Undersmoothed HAL	
		Cov	Width	Cov	Width
1 (smooth)	Delta simult.	84.8	0.179	82.4	0.249
	V-fold simult. ($V=5$)	91.4	0.338	91.0	0.479
	V-fold simult. ($V=10$)	91.2	0.230	88.8	0.325
	V-fold simult. ($V=20$)	89.8	0.199	88.8	0.288
	Bootstrap simult.	90.0	0.186	88.4	0.256
	BC simult. ($V=5$)	94.8	0.355	94.6	0.513
	BC simult. ($V=10$)	96.0	0.255	92.8	0.372
	BC simult. ($V=20$)	94.2	0.239	95.0	0.357
	BC Bootstrap simult.	92.4	0.194	90.0	0.268
2 (oscillating)	Delta simult.	27.6	0.156	19.2	0.181
	V-fold simult. ($V=5$)	90.0	0.447	92.2	0.541
	V-fold simult. ($V=10$)	86.8	0.311	88.0	0.382
	V-fold simult. ($V=20$)	87.8	0.281	89.0	0.351
	Bootstrap simult.	85.6	0.230	89.4	0.281
	BC simult. ($V=5$)	94.2	0.497	95.8	0.603
	BC simult. ($V=10$)	94.8	0.380	91.8	0.470
	BC simult. ($V=20$)	96.6	0.393	95.4	0.485
	BC Bootstrap simult.	87.6	0.247	91.6	0.299
3 (piecewise)	Delta simult.	17.4	0.235	26.2	0.313
	V-fold simult. ($V=5$)	64.2	0.519	79.2	0.712
	V-fold simult. ($V=10$)	47.2	0.366	65.4	0.500
	V-fold simult. ($V=20$)	38.4	0.331	59.8	0.462
	Bootstrap simult.	30.0	0.261	56.4	0.349
	BC simult. ($V=5$)	70.2	0.553	80.0	0.738
	BC simult. ($V=10$)	53.4	0.430	72.0	0.583
	BC simult. ($V=20$)	51.2	0.446	72.2	0.593
	BC Bootstrap simult.	33.4	0.278	60.8	0.372
4 (discontinuous)	Delta simult.	10.8	0.324	31.6	0.423
	V-fold simult. ($V=5$)	52.2	0.656	77.6	0.883
	V-fold simult. ($V=10$)	30.6	0.455	65.0	0.616
	V-fold simult. ($V=20$)	26.4	0.406	59.6	0.556
	Bootstrap simult.	16.0	0.342	56.8	0.474
	BC simult. ($V=5$)	64.4	0.714	84.8	0.977
	BC simult. ($V=10$)	49.6	0.552	78.6	0.741
	BC simult. ($V=20$)	45.4	0.549	76.8	0.736
	BC Bootstrap simult.	21.0	0.363	63.4	0.503

ods. For DGP 2 (oscillating), the advantage is striking: the delta method achieves only 27.6% (CV-HAL) and 19.2% (undersmoothed), while the V -fold band reaches 87.8–90.0% and 88.0–92.2%. For DGPs 3 and 4, all uncorrected methods exhibit substantial under-coverage due to estimator bias at non-smooth points, but the V -fold simultaneous band consistently provides the best coverage. With undersmoothed HAL on DGP 4, the V -fold band ($V = 5$) reaches 77.6% compared to 31.6% for the delta method and 56.8% for the bootstrap.

The bias-corrected bands from Theorem 5.9 provide a substantial improvement. For DGPs 1 and 2, where the estimator bias is moderate, the bias-corrected V -fold band achieves near-nominal coverage: 94.2–96.0% (CV-HAL) and 92.8–95.0% (undersmoothed) for DGP 1, and 94.2–96.6% and 91.8–95.8% for DGP 2. These results validate Theorem 5.9 in finite samples. For DGPs 3 and 4, the bias correction substantially improves coverage relative to the standard band—for example, on DGP 4 with undersmoothed HAL the bias-corrected band ($V = 5$) reaches 84.8% versus 77.6% without correction—but coverage remains below nominal because the estimator bias is too large for the jackknife bias estimate to fully compensate. The bias-corrected bands are wider than the standard bands (e.g., for DGP 2 with $V = 20$: 0.393 versus 0.281 for CV-HAL), reflecting the price of the $|\hat{b}_j|$ correction term.

The bias-corrected bootstrap band, which widens the standard bootstrap band by $|\hat{b}_{\text{boot},j}|$ where $\hat{b}_{\text{boot},j} = \bar{\psi}_{\text{boot}}(a_j) - \hat{\psi}(a_j)$, similarly improves over the standard bootstrap by 2–7 percentage points across all DGPs. For DGPs 1 and 2, it reaches 87.6–92.4% (CV-HAL) and 90.0–91.6% (undersmoothed). For DGPs 3 and 4, the improvement is more modest (e.g., 21.0% versus 16.0% for CV-HAL on DGP 4). However, the bias-corrected bootstrap consistently achieves lower coverage than the bias-corrected V -fold band, indicating that the jackknife bias estimate is more effective than the bootstrap bias estimate for bias correction in this setting.

These results illustrate three complementary points. First, the V -fold jackknife provides valid variance estimation and the most reliable inference among the methods compared, regardless of the HAL tuning strategy. The delta method can fail both numerically (producing NAs at rates up to 6.2%) and statistically (coverage as low as 19.2% for simultaneous bands), whereas the V -fold jackknife never fails and consistently achieves the highest coverage. Second, the bias-corrected simultaneous band from section 5.4.6 recovers near-nominal coverage when the estimator bias is moderate (DGPs 1–2), confirming the theoretical guarantee of Theorem 5.9; the bias-corrected bootstrap similarly improves over the standard bootstrap, though the jackknife bias estimate proves more effective for bias correction than the bootstrap bias estimate. Third, achieving nominal simultaneous coverage ultimately requires the estimator’s bias to be negligible relative to its standard error—a condition that depends on the estimator and data generating process rather than the variance estimation method. The V -fold jackknife correctly quantifies sampling variability; the remaining coverage gap in non-smooth settings is attributable to estimator bias, which the bias correction partially but not fully addresses when the bias is large. The spatial pattern of the coverage dips in Figure C.3 mirrors the bias profile in Figure 5.4: the simultaneous coverage is driven down by the same high-bias dose points identified in the pointwise analysis.

Effective rank of the correlation matrix. For the simultaneous bands with $V = 5$

or $V = 10$, we have $m = 20 > V - 1$, so $\hat{R}_{\text{Jack}}$ is rank-deficient and the eigendecomposition of section 5.4.4 is used. We estimate the effective rank $r(R_0)$ from the 500 simulation replications, as for the KM setting. For CV-HAL, the effective rank ranges from 2.8 (DGP 1, smooth) to 3.8 (DGP 3, piecewise); for undersmoothed HAL, it ranges from 4.1 (DGP 2, oscillating) to 5.3 (DGP 3 and 4). The higher effective rank for undersmoothed HAL reflects the reduced smoothing: adjacent curve points are less correlated, distributing variance across more eigendirections. Nevertheless, in all cases the effective rank remains well below $m = 20$, and in particular below $V - 1 = 4$ for $V = 5$, confirming that even the smallest fold count captures the dominant correlation structure. Figure 5.5 displays the eigenvalue spectrum and cumulative variance explained for all settings, showing rapid eigenvalue decay in every case.




Figure 5.5: Eigenvalue spectrum of the estimated correlation matrix R_0 for the Kaplan–Meier survival curve ($m = 91$ time points) and the HAL dose-response curve ($m = 20$ evaluation points, four DGPs, CV-HAL and undersmoothed variants). The correlation matrix is estimated from the $500 \times m$ matrix of point estimates across simulation replications. (a) Normalized eigenvalue spectrum (fraction of trace m), truncated at index 20 for visibility. The KM curve exhibits the steepest decay, with a single dominant eigenvalue accounting for 59% of the total variance. (b) Cumulative fraction of variance explained. For the KM curve, two eigenvalues suffice for 77% and nine for 95%; for the HAL settings, 6–16 eigenvalues are needed for 95%. In all cases, the effective rank $r(R_0) = m/\lambda_{\max}$ is much smaller than m , confirming that the eigendecomposition of the rank-deficient $\hat{R}_{\text{Jack}}$ accurately captures the dominant correlation structure even when $V \ll m$.

5.7 Discussion

This chapter develops the grouped (V -fold) jackknife as a practical inference tool for regular asymptotically linear estimators of pathwise differentiable parameters. Building on the classical theory of Tukey [1958], Brillinger [1964], and Shao and Wu [1989], it extends fixed- V t_{V-1} Studentization to the regular asymptotically linear class, establishes diverging- V consistency at rate $\sqrt{V}$, and develops simultaneous confidence bands and slower-rate extensions that remain usable when direct variance formulas are unstable or unavailable.

The simulations support both the formal results and the computational motivation. For scalar inference on the average treatment effect, the grouped jackknife with t_{V-1} critical values achieves strong coverage even at $n = 200$, whereas the plug-in influence-function-based variance estimator undercovers substantially. For the Kaplan–Meier survival function, the grouped jackknife yields pointwise and simultaneous bands that are practically comparable to established methods. For HAL dose-response curve estimation, where the delta method can fail numerically and be severely anti-conservative, the grouped jackknife never fails and consistently delivers the strongest empirical coverage among the methods compared. When the estimator bias is moderate, the bias-corrected simultaneous band restores near-nominal coverage.

5.7.1 Practical guidance on choosing V

The choice of V involves a trade-off between computational cost and statistical precision. Fewer folds require fewer refits but produce wider confidence intervals because the t_{V-1} quantiles exceed their normal counterparts; more folds narrow the intervals as the t quantiles approach the normal quantiles and the variance estimator concentrates. For scalar inference, the simulations suggest that $V \in \{10, 20\}$ offers a practical balance: for the ATE at $n = 1000$, increasing V from 5 to 20 reduces the mean CI width by roughly 15% while maintaining coverage above the nominal level. For simultaneous confidence bands, the more relevant issue is the quality of the estimated correlation matrix $\hat{R}_{Jack}$. As discussed in section 5.4.4, a practical approach is to increase V along a sequence such as $\{5, 10, 20, 40, \dots\}$ until it satisfies $V \geq 2\hat{r}_V$; absent such tuning, the simulations suggest that $V = 20$ is a reasonable default in the settings studied. When $V < m$, the estimated correlation matrix is singular and the Monte Carlo draws are confined to a rank- $(V - 1)$ subspace of $\mathbb{R}^m$, but this confinement is exact for the plug-in distribution and is often benign when the dominant eigendirections are already captured. As a fallback rule, $V = 10$ to 20 remains reasonable for scalar and moderate-dimensional problems, with larger V preferred when m is large or when the computational budget allows it. The total cost of the procedure is $V + 1$ fits of the estimator (one full-sample fit and V leave-fold-out fits), compared to thousands for the bootstrap.

5.7.2 Scope and limitations

The theoretical guarantees rest on the estimator admitting an asymptotically linear expansion with a pathwise differentiable influence function. This includes a broad class of semiparametric estimators—TMLE, one-step estimators, augmented IPW estimators—and the generalized asymptotic linearity framework of Section 5.5 extends the scope beyond $\sqrt{n}$ -rate estimators. Parameters fall outside the present theory when the required asymptotic linearity or fold-stable remainder conditions fail, so genuinely non-smooth cases require separate analysis. Efron’s sample-median example shows that the ordinary delete-1 jackknife can fail in such settings, but that example should be read as a cautionary boundary case rather than as a blanket exclusion of all quantile-type parameters. The key insight is that the Studentized jackknife statistic is scale-invariant: σ_n cancels between numerator and denominator, so the t_{V-1} inference remains valid regardless of whether the influence-function variance is bounded or diverges with the number of basis functions. This allows practitioners to obtain valid confidence intervals without knowing J_n —in contrast to the delta method, which requires computing these quantities explicitly. The simultaneous confidence-band theory assumes fixed m ; extending to a growing number of evaluation points $m = m_n \rightarrow \infty$ would require uniform-in- m control of the remainder terms and a different limiting process (e.g., a Gaussian-process approximation), which is beyond the scope of this chapter.

For simultaneous confidence bands, the fixed- m theory identifies the correct component-wise Studentized limit distribution for fixed V conditional on the true correlation matrix R_0 , while consistency of the correlation estimate requires $V \rightarrow \infty$. Formal plug-in simultaneous coverage therefore also requires $V \rightarrow \infty$, although the divergence can be at a reasonably slow rate such as $O(\log n)$. The strong performance observed for small fixed V should thus be interpreted as empirical support, informed by the low-effective-rank argument, rather than as a formal fixed- V coverage theorem.

The bias-corrected bands from Theorem 5.9 improve coverage when the estimator’s bias is moderate relative to its standard error, as demonstrated for DGPs 1 and 2 in the HAL simulations. However, when the bias is large—as in the piecewise and discontinuous dose-response settings (DGPs 3 and 4)—the bias correction provides partial but insufficient improvement. The bias diagnosis in Figure 5.4 makes the mechanism precise: coverage loss occurs at dose points where the absolute bias exceeds the standard error, and these points coincide with features of the true curve (oscillations for $A > 2$ in DGP 3, discontinuities at $A = 2$ and $A = 4$ in DGP 4) that the first-order HAL’s piecewise-linear basis cannot well approximate. This is a limitation of the estimator’s approximation class, not of the variance estimation method. Achieving nominal coverage in such settings requires reducing the estimator’s bias—through higher-order HAL basis functions, adaptive smoothness selection as in Shi et al. [2024], or other bias-reduction techniques; the V -fold jackknife correctly quantifies sampling variability but cannot compensate for systematic bias in the point estimator.

Finally, the results are conditional on a single random partition of the data into V folds. While the theory establishes asymptotic validity for any partition satisfying the balanced-fold condition, finite-sample performance may exhibit partition-to-partition variability, particularly when V is small. Averaging the variance estimator over multiple independent partitions

could reduce this variability at additional computational cost.

5.7.3 Connections and extensions

The V -fold jackknife shares a natural connection with cross-fitting procedures used in modern semiparametric inference. In debiased machine learning [Chernozhukov et al., 2018] and cross-validated TMLE [Zheng and van der Laan, 2011], the nuisance functions are already trained on each leave-fold-out subsample $P_{n,-v}$, and the target parameter is evaluated on the held-out fold I_v . The V -fold jackknife requires the complementary quantity $\hat{\Psi}(P_{n,-v})$ —the estimator evaluated on the $V - 1$ training folds rather than on the single held-out fold. Since the expensive nuisance estimation step has already been performed, computing $\hat{\Psi}(P_{n,-v})$ requires only an additional targeting or evaluation pass on the training folds using the existing nuisance estimates, avoiding the cost of a full refit. The V -fold jackknife thus integrates naturally into cross-fitting workflows at modest additional computational cost.

The V -fold jackknife should also be distinguished from the jackknife+ and CV+ methods of Barber et al. [2021], which use fold-based leave-one-out residuals to construct distribution-free prediction intervals for a new observation. While both methods partition the data into folds, the inferential targets are fundamentally different: the jackknife+ provides marginal coverage for a future observation under exchangeability, whereas the V -fold jackknife developed here provides asymptotic coverage for fixed population parameters.

The extension to non- $\sqrt{n}$ rates developed in Section 5.5 opens several chapter-level directions. The generalized asymptotic linearity framework applies to any sieve MLE whose influence-curve-like representation has growing variance, including not only HAL but also kernel and series estimators. Combined with the vector-valued theory, this provides a principled approach to uniform inference for HAL functional estimators—the form in which the jackknife is used in Chapter 6. Extensions to dependent data through block partitions, to simultaneous bands with $m = m_n \rightarrow \infty$, and to adaptive data-driven selection of V remain natural next steps.

5.7.4 Concluding remarks

The grouped jackknife has a long history in statistics, but its properties for the semiparametric and machine-learning-based estimators used in modern practice have not been systematically examined. The analysis here shows that the grouped V -fold jackknife is a simple and reliable tool for variance estimation and inference when analytic standard errors are unavailable, difficult to derive, or numerically unreliable. Its primary appeal is computational: only V refits are required, compared to the thousands demanded by the bootstrap. Fixed- V t_{V-1} inference is especially attractive for expensive estimators, where even $V = 5$ or 10 folds yield valid confidence intervals. For simultaneous confidence bands, the fixed- V theory identifies the correct componentwise-Studentized limit conditional on R_0 , while the plug-in procedure becomes formally valid once $V \rightarrow \infty$ at a reasonably slow rate such as $O(\log n)$; low-effective-rank correlation structure helps explain why smaller values of V can still perform well. The same scale-invariant Studentization extends the method to estimators

whose influence-function variance diverges with n , eliminating the need to know or estimate the effective convergence rate.

This chapter supplies the inferential complement to the regularized targeting methods developed earlier. The next chapter applies both parts of the dissertation's methodological program at once, using projection-based regularized targeting for point estimation and grouped V-fold jackknife methods for feasible pointwise and simultaneous inference on a functional survival target under coarsened data.

Chapter 6

Culminating Application: Bivariate Right-Censored Survival

Bivariate right-censored survival is the setting in which the methodological program of the dissertation is under the greatest pressure. The full-data density must be estimated under coarsening, the target parameter is functional, and a literal classical full TMLE for the entire survival surface is not operationally tractable. In addition, uncertainty quantification is demanding because the full estimation pipeline is expensive to refit. A problem in which all of these difficulties occur simultaneously is therefore the natural place to assess the scope of the earlier developments.

The chapter addresses these difficulties by combining the central ideas of the dissertation. EM-HAL provides a smooth initial estimator of the full-data density. Projection-based regularized A-TMLE extends the targeting logic developed in Chapter 3 through a larger HAL working model and valid IPCW or influence-curve representations, thereby avoiding the need to rely on a tractable closed-form exact efficient influence curve and least favorable fluctuation as operational inputs for the full survival surface. The grouped V-fold jackknife developed in Chapter 5 then supplies feasible pointwise and simultaneous inference, a particularly important feature here because the EM-HAL-plus-targeting pipeline is too expensive for heavy resampling. Although the implementation is HAL-based, the broader interpretation is informed by Chapter 4, which shows that projection-guided targeting is more naturally understood as a score-space principle than as a feature specific to one model class.

6.1 Local problem formulation

This chapter studies the bivariate right-censored survival surface

$$S(t_1, t_2) = P(T_1 > t_1, T_2 > t_2)$$

under a coarsened-data structure that makes each stage of the targeted-learning pipeline nontrivial. Three simultaneous difficulties drive the chapter's method choices. First, the

initial estimator must respect the censoring geometry, so naive bivariate NPMLE is inconsistent and classical density estimators introduce bandwidth choices that become more difficult in higher dimensions. Second, a literal classical full TMLE for the full survival surface is not computationally available in this coarsened-data model: an exact least favorable fluctuation whose score equals the efficient influence curve is not tractable in closed form. Third, direct plug-in variance formulas are unstable once the initial density estimator is regularized, so an inferential layer that does not rely on inverting a potentially ill-conditioned information operator is needed. This chapter therefore takes the initial-estimator problem from EM-HAL, the projection-based targeting logic from Chapter 3, and the grouped-jackknife inferential layer from Chapter 5, and combines them into a single pipeline.

6.1.1 Position relative to existing work

Chapter 2 already records the shared targeted-learning framework, and Chapters 3 and 5 provide the estimator-side and inferential machinery reused here. The present subsection therefore restricts attention to the chapter-specific literatures needed to locate the bivariate-survival contribution.

Classical bivariate survival estimators provide the first comparison class. Product-limit and hazard-based constructions due to Dabrowska [1988], Prentice and Cai [1992], and Prentice et al. [2004] remain the canonical references, with Lin and Ying [1993] providing weak-convergence theory in a simpler censoring setting. Those estimators define the basic nonparametric problem, but they do not simultaneously deliver a proper smooth density representation, regularized targeting of the survival functional, and a practical inferential layer for the coarsened-data pipeline studied here.

A second comparison class is the NPMLE/SOR-MLE line. The inconsistency of the naive bivariate NPMLE under singly censored half-lines was identified by Tsai et al. [1986] and van der Laan [1996], with the SOR-MLE repairing the problem through smoothing. That line shows why the initial estimator must respect the coarsening geometry, but it leaves a bandwidth-selection problem and does not by itself provide the targeted functional correction or the inference layer developed here. Copula models offer a different route by imposing a parametric dependence structure [Clayton, 1978, Oakes, 1989, Genest et al., 1995, Shih and Louis, 1995]; they can be efficient under correct specification, but their dependence assumptions can be misspecified and they do not supply the nonparametric density-plus-targeting program pursued in this chapter.

The initial-estimator side of the present chapter is instead closest to the HAL/EM line. HAL provides the nonparametric likelihood-based geometry [van der Laan, 2017, Benkeser and van der Laan, 2016, Bibaut and van der Laan, 2019, Ertefaie et al., 2023, van der Laan, 2023], while the EM literature provides the standard computational strategy for incomplete-data likelihoods [Dempster et al., 1977, Turnbull, 1976, Wei and Tanner, 1990, Groeneboom and Wellner, 1992]. The specific contribution here is to combine an exponential-family HAL density parameterization with EM-HAL estimation for bivariate right-censored data, following the same broader sieve-HAL perspective used later in the targeting arguments [Munch et al., 2024, Hou et al., 2026].

On the targeting side, this chapter inherits the projection-based regularized A-TMLE logic of Chapter 3 rather than re-deriving a classical full TMLE. The substantive reason is local to this problem: a tractable exact efficient influence curve and least favorable fluctuation for the full bivariate survival surface are not available as operational inputs. Adaptive targeted learning and related automatic-debiasing ideas provide the relevant working-model perspective [van der Laan et al., 2023, 2025, Li et al., 2025b,a]. The self-consistency construction is retained as a comparator, but the main-line method is the projection-based reg-ATMLE update layered on top of the EM-HAL initial estimator.

On the inferential side, the chapter reuses the grouped V-fold jackknife developed in Chapter 5 because direct influence-curve-based or information-matrix-based variance calculations are too fragile for the full EM-HAL-plus-targeting pipeline. In that sense the chapter's contribution is cumulative: it is not only another survival paper, but the place where the dissertation's estimator-side and inferential developments are simultaneously stress-tested in a coarsened-data functional problem. Chapter 4's neural-network construction is not used algorithmically here, but it remains interpretively relevant because it shows that projection-guided targeting is a score-space principle rather than a HAL accident. The full self-consistency construction is retained as a secondary comparison and moved to Appendix D.3.

6.2 Problem setup and notation

6.2.1 Observed data and the CAR model

Let $X = (T_1, T_2) \in [0, 1]^2$ denote the full-data bivariate event-time vector with density $f = p_X$ with respect to Lebesgue measure. The event times are subject to independent right-censoring:

$$O = (\tilde{T}_1, \Delta_1, \tilde{T}_2, \Delta_2), \quad \tilde{T}_j = \min(T_j, C_j), \quad \Delta_j = I(T_j \leq C_j), \quad j = 1, 2,$$

where $C = (C_1, C_2)$ is independent of X . We observe n i.i.d. copies $O_1, \dots, O_n$ of O .

Under this censoring mechanism, the coarsening at random (CAR) assumption holds [Gill et al., 1997, van der Laan and Robins, 2003], so the observed-data likelihood factorizes as

$$p_{P_X, G}(o) = q_X(C(o)) \cdot h_G(o),$$

where $q_X(C(o))$ depends on P_X only and $h_G(o)$ depends on the censoring distribution G only. Since the coarsening set $C(O)$ is observed, the factorized density contribution $q_X(C(O))$ can be written explicitly for the full data X . The likelihood factor $q_X(C(O))$ evaluates to:

$$q_X(C(O)) = \begin{cases} f(\tilde{T}_1, \tilde{T}_2) & \text{if } \Delta_1 = 1, \Delta_2 = 1 \\ \int_{\tilde{T}_2}^1 f(\tilde{T}_1, u_2) du_2 & \text{if } \Delta_1 = 1, \Delta_2 = 0 \\ \int_{\tilde{T}_1}^1 f(u_1, \tilde{T}_2) du_1 & \text{if } \Delta_1 = 0, \Delta_2 = 1 \\ \int_{\tilde{T}_1}^1 \int_{\tilde{T}_2}^1 f(u_1, u_2) du_2 du_1 & \text{if } \Delta_1 = 0, \Delta_2 = 0. \end{cases}$$

Thus, inference for P_X can correctly ignore the censoring mechanism $h_G(O)$.

Here $C(O) \subset [0, 1]^2$ is the coarsening set implied by O : for uncensored observations ($\Delta_1 = \Delta_2 = 1$), $C(O) = \{(\tilde{T}_1, \tilde{T}_2)\}$; for singly censored, $C(O)$ is a half-line segment; for doubly censored ($\Delta_1 = \Delta_2 = 0$), $C(O) = (\tilde{T}_1, 1] \times (\tilde{T}_2, 1]$.

Since the censoring mechanism can be ignored under CAR, inference for P_X is based on the likelihood factor $q_X(C(O))$.

6.2.2 Target parameters

The primary target parameter is the bivariate survival surface, evaluated on a grid $\mathcal{T} = \{t_\ell = (t_{1\ell}, t_{2\ell}) : \ell = 1, \dots, m\} \subset [0, 1]^2$:

$$\Psi_\ell(P) \equiv S_P(t_{1\ell}, t_{2\ell}) = P(T_1 > t_{1\ell}, T_2 > t_{2\ell}) = \int_{t_{1\ell}}^1 \int_{t_{2\ell}}^1 f(u_1, u_2) du_2 du_1.$$

The bivariate density $f(t_1, t_2)$ is a secondary object of interest; its performance as a byproduct of the EM-HAL estimation is discussed in the main text and Appendix D.6.

6.3 EM-HAL estimation framework

6.3.1 HAL working model for the full-data density

We parameterize the full-data density in an exponential-family form on $[0, 1]^2$:

$$\log f_\theta(t_1, t_2) = \sum_{k=1}^p \theta_k \phi_k(t_1, t_2) - \log Z(\theta), \quad Z(\theta) = \int_{[0,1]^2} \exp \left\{ \sum_{k=1}^p \theta_k \phi_k(u) \right\} du, \quad (6.1)$$

where $\{\phi_k\}_{k=1}^p$ is a HAL basis consisting of first-order hinge functions $\phi_k(t) = (t_j - g_k)_+$ and their pairwise tensor products. (Note that an explicit global intercept θ_0 is mathematically redundant, as it factors out of the exponential and cancels exactly with the normalizing constant $Z(\theta)$.) The coefficients are subject to an ℓ_1 constraint $\|\theta\|_1 \leq \lambda$ to control model complexity.

In this exponential-family parameterization, the score vector for the censored-data likelihood takes the conditional-expectation form

$$S_\theta(O) = E_\theta\{\phi(X) \mid O\} - E_\theta\{\phi(X)\}, \quad (6.2)$$

where $\phi(X) = (\phi_1(X), \dots, \phi_p(X))^\top$ and the conditional expectation $E_\theta(\cdot \mid O)$ is computed under f_θ restricted to $C(O)$. This is precisely what the EM algorithm computes in its E-step.

6.3.2 The EM algorithm for censored data

The observed-data HAL-MLE $\hat{\theta} = \arg \min_\theta P_n L(\theta)$, with $L(\theta)(O) = -\log q_{X,\theta}(C(O))$, is generally non-convex in θ because $q_{X,\theta}$ involves integrals of f_θ over coarsening sets. By

contrast, the unpenalized full-data log-likelihood in the exponential-family model (6.1) is concave in θ , and the ℓ_1 constraint makes the M-step a well-posed convex problem. This makes EM a natural implementation strategy.

We initialize the EM procedure with a warm-start density estimate f_{θ^0} obtained from a HAL fit on multiply imputed full data. In practice, each censored observation is first completed under a simple initial distribution, and the resulting pseudo-complete sample is used to fit an ℓ_1 -penalized exponential-family model with cross-validation selecting λ_{CV} . An IPCW-based initialization is also possible, but throughout this chapter we use the imputation-based warm start.

Algorithm 8 summarizes the EM-HAL implementation. At iteration k , the E-step computes or approximates the conditional law of X_i given $X_i \in C(O_i)$ under the current density f_{θ^k} , while the M-step updates θ through the corresponding weighted full-data HAL likelihood under the relaxed constraint $\|\theta\|_1 \leq \lambda_{EM}$. For uncensored observations, the conditional distribution degenerates at the observed event-time pair, so no imputation is required.

Algorithm 8 EM-HAL estimation under bivariate right-censoring

- 1: **Input:** observed data $\{O_i\}_{i=1}^n$, HAL basis ϕ , cross-validated penalty λ_{CV} , relaxed penalty λ_{EM} , imputation size B
- 2: Obtain a warm-start density estimate f_{θ^0} from multiply imputed full data
- 3: **for** $k = 0, 1, 2, \dots$ until convergence **do**
- 4: **E-step:**
 For each i , compute or approximate the conditional law of X_i given $X_i \in C(O_i)$ under f_{θ^k}
 Draw weighted imputations $\{x_i^{(b)}, p_i^k(b)\}_{b=1}^{B_i}$; for uncensored observations use the observed point mass
- 5: **M-step:**
 Update θ^{k+1} by solving

$$\theta^{k+1} = \arg \max_{\theta, \|\theta\|_1 \leq \lambda_{EM}} \frac{1}{n} \sum_{i=1}^n \sum_{b=1}^{B_i} p_i^k(b) \log f_{\theta}(x_i^{(b)})$$

This is a weighted ℓ_1 -penalized full-data HAL-MLE

- 6: Evaluate the observed-data log-likelihood $\sum_{i=1}^n \log q_{X, \theta^{k+1}}(C(O_i))$
 - 7: **end for**
 - 8: **Output:** EM-HAL density estimate $f_{\hat{\theta}}$
-

The relaxation parameter is taken as a multiple of the cross-validated penalty, typically $\lambda_{EM} = 5 \cdot \lambda_{CV}$, to enlarge the effective support while retaining regularization. Convergence is monitored through the observed-data log-likelihood $\sum_i \log q_{X, \theta^k}(C(O_i))$, and the iterations are terminated once the increment falls below a prespecified tolerance or a maximum number of iterations is reached.

Connection to the HAL sieve MLE framework. To connect the EM-HAL estimator with the sieve-HAL theory of van der Laan [2023], we distinguish the global log-density class

$\mathcal{Q}$, the data-adaptive HAL sieve indexed by the penalty level λ , and the corresponding oracle sieve approximation.

Let $\mathcal{Q}$ denote the class of càdlàg functions on $[0, 1]^2$ with bounded sectional variation norm, and let $Q(P)$ denote the unnormalized log-density associated with P , so that $f_Q = \exp(Q)/Z(Q)$ as parameterized in (6.1). Let $L(Q)(O) = -\log q_X(C(O))$ denote the observed-data log-likelihood loss under CAR, and assume this loss identifies the target element through

$$Q(P) = \arg \min_{Q \in \mathcal{Q}} PL(Q).$$

For the true distribution P_0 , we therefore write

$$Q_0 = Q(P_0) = \arg \min_{Q \in \mathcal{Q}} P_0 L(Q).$$

Thus Q_0 is defined as the loss-minimizing element of the global càdlàg bounded-variation class, not of a finite sieve.

The data-adaptive HAL working model for Q indexed by the ℓ_1 norm constraint λ is the k -th order spline class

$$D_\lambda^{(k)}(\mathcal{R}_n) = \left\{ \sum_{j \in \mathcal{R}_n} \beta_j \phi_j : \|\beta\|_1 \leq \lambda \right\},$$

where $\mathcal{R}_n$ indexes the data-adaptively selected basis functions. The EM-HAL empirical sieve MLE is then the empirical risk minimizer in this working model:

$$Q_n = \arg \min_{Q \in D_\lambda^{(k)}(\mathcal{R}_n)} P_n L(Q).$$

While the full-data log-likelihood is strictly concave in θ , the observed-data marginal likelihood involves integrals over coarsening sets and is generally non-concave. Consequently, while the EM algorithm provides monotonic likelihood increases and converges to a stationary point [Wu, 1983], it is not mathematically guaranteed to find the unique global maximum. In our theoretical development, we proceed under the standard assumption that the EM algorithm—initialized via our multiply imputed warm-start—successfully converges to the global empirical risk minimizer (Sieve MLE) over the working model.

The main sieve-estimation complication is that the active basis set $\mathcal{R}_n$ is selected data-adaptively. The theory therefore introduces a fixed sequence of index sets $\mathcal{R}_{0,n}$ of dimension $d_{0,n} = |\mathcal{R}_{0,n}|$ that approximates the empirical working model. The corresponding oracle working model is $D_\lambda^{(k)}(\mathcal{R}_{0,n})$, and its loss-based projection of Q_0 is

$$Q_{0,n} = \arg \min_{Q \in D_\lambda^{(k)}(\mathcal{R}_{0,n})} P_0 L(Q).$$

The estimation error then decomposes into a stochastic component within the oracle working model and a deterministic working-model approximation error:

$$Q_n - Q_0 = (Q_n - Q_{0,n}) + (Q_{0,n} - Q_0).$$

The technical role of EM-HAL in the dissertation is twofold. First, it supplies a smooth initial density estimator whose loss-based convergence rate is strong enough for the later targeting arguments. Second, because it is an ℓ_1 -constrained likelihood-based sieve estimator, it can be analyzed within the existing HAL theory for oracle approximation and empirical-process control. The detailed regularity transfer and pointwise asymptotic-normality arguments are deferred to Appendix D.2.

6.3.3 Plug-in survival and density estimation

Given the EM-HAL density estimate $f_{\hat{\theta}}$, the bivariate survival function is computed via the inclusion-exclusion formula:

$$\hat{S}(t_1, t_2) = 1 - \hat{F}_1(t_1) - \hat{F}_2(t_2) + \hat{F}(t_1, t_2), \quad (6.3)$$

where $\hat{F}(t_1, t_2) = \int_0^{t_1} \int_0^{t_2} f_{\hat{\theta}}(u_1, u_2) du_2 du_1$ is the joint CDF and $\hat{F}_j$ are the marginal CDFs. In practice, these integrals are approximated on a numerical grid using the trapezoidal rule.

The density is evaluated directly from the exponential-family model (6.1). Both the survival surface and density are evaluated on a fine grid $\mathcal{T}$ of 209×209 points on $[0, 1]^2$, constructed as the union of 200 equally spaced points on $[0, 1]$ and the 9 targeting grid coordinates $\{0.1, 0.2, \dots, 0.9\}$.

6.3.4 Computational details

The main text keeps only the modeling choices needed to understand the estimator. The numerical integration grids, pruning thresholds, imputation counts, EM stopping rules, and working-model relaxation settings used in the simulations are collected in Appendix D.1.

6.4 Targeting for survival

6.4.1 TMLE, Regularized TMLE, and Regularized A-TMLE

The EM-HAL estimator maximizes a regularized observed-data likelihood over the HAL model class. This yields a stable global density estimate, but the resulting plug-in survival estimator $\hat{S}_{\text{EMHAL}}(t_1, t_2) = \int_{t_1}^1 \int_{t_2}^1 f_{\hat{\theta}}(u_1, u_2) du_2 du_1$ is not directly targeted to the survival functional. Because the ℓ_1 penalty is chosen for global density estimation rather than for $\Psi_{\ell}(P) = S_P(t_{1\ell}, t_{2\ell})$, the resulting regularization bias can remain relevant at the $n^{-1/2}$ scale needed for valid inference.

Our theoretical framing follows the hierarchy

$$\text{TMLE} \rightarrow \text{regularized TMLE} \rightarrow \text{A-TMLE} \rightarrow \text{regularized A-TMLE}.$$

TMLE solves the exact score or efficient-influence-curve equation for the target of interest. In this problem, however, a literal classical full TMLE is not computationally tractable because

no tractable expression is available for the exact efficient influence curve of the full bivariate survival surface, nor for a corresponding least favorable fluctuation. Regularized TMLE replaces that infeasible equation with a controlled approximation. A-TMLE is the corresponding working-model formulation, and regularized A-TMLE further replaces the exact working-model EIC with a regularized approximation. Self-consistency targeting plays the regularized-TMLE role through score control for a selected self-consistency system, whereas the projection-based reg-ATMLE update plays the regularized A-TMLE role as the corresponding working-model special case within the broader regularized TMLE framework, by targeting a projected survival parameter inside a larger working model. The main text develops the projection-based reg-ATMLE update; Appendix D.3 gives the full self-consistency construction.

Let $\mathcal{G}_n = \{s_1, \dots, s_{m_n}\} \subset [0, 1]^2$ denote a targeting grid and, for $s = (t_1, t_2)$, write $\Psi_s(P) = S_P(t_1, t_2)$. We denote the corresponding grid-based survival parameter by

$$\Psi_{\mathcal{G}_n}(P) = \left(\Psi_{s_1}(P), \dots, \Psi_{s_{m_n}}(P) \right)^\top.$$

In this chapter the implemented targeting grid is $\{0.1, \dots, 0.9\}^2$, so $m_n = 81$, but conceptually m_n is a precision choice and may grow as the grid is refined.

6.4.2 Projection-based reg-ATMLE update

The projection-based reg-ATMLE update is the primary method developed in the main text. Its target is the finite-dimensional survival parameter $\Psi_{\mathcal{G}_n}(P) = \left(\Psi_{s_1}(P), \dots, \Psi_{s_{m_n}}(P) \right)^\top$. For each grid point $s \in \mathcal{G}_n$, we project a valid IPCW gradient for $\Psi_s(P)$ onto the HAL score span of a larger targeting working model and then aggregate the resulting directions into a joint fluctuation over the grid. Let $\mathcal{M}_{n,\text{init}}$ denote the initial EM-HAL working model induced by the pruned fit and let $\mathcal{M}_{n,\text{targ}}$ denote a larger targeting working model obtained by relaxing the HAL penalty. Because the EM-HAL fit is already the MLE in $\mathcal{M}_{n,\text{init}}$, targeting must occur in $\mathcal{M}_{n,\text{targ}}$. In the implementation, $\mathcal{G}_n$ is the 9×9 targeting grid, and any spillover from that grid to the denser evaluation grid is assessed empirically rather than built into the theorem.

6.4.2.1 IPCW gradient for bivariate survival

The reg-ATMLE update starts from a scalar gradient for $\Psi_s(P)$ at a fixed grid point $s = (t_1, t_2)$. Let $\bar{G}(c_1, c_2) = P(C_1 \geq c_1, C_2 \geq c_2)$ denote the joint censoring survival function. Under independent censoring, a valid mean-zero nonparametric gradient for $\Psi_s(P) = S_P(t_1, t_2)$ is the predictable left-limit IPCW estimating function:

$$D_s(O; P, \bar{G}) = \frac{\Delta_1 \Delta_2 I(T_1 > t_1, T_2 > t_2)}{\bar{G}(T_1-, T_2-)} - S_P(t_1, t_2). \quad (6.4)$$

6.4.2.2 Larger targeting working model

The reg-ATMLE update separates the initial estimation working model from the targeting working model. The initial EM-HAL fit determines $\mathcal{M}_{n,\text{init}}$, while the targeting step is carried out in a larger model $\mathcal{M}_{n,\text{targ}}$ obtained by relaxing the ℓ_1 constraint. Concretely, for a relaxation factor κ , we refit HAL with constraint $\kappa \times \lambda_{\text{CV}}$, collect the additional active basis functions, and append them to the EM-HAL offset with zero initial coefficients. This enlargement is necessary because the score equations in $\mathcal{M}_{n,\text{init}}$ have already been exhausted by the EM-HAL fit. In the simulations we use $\kappa = 6$.

6.4.2.3 Regularized projection onto the HAL score span

Regularized A-TMLE becomes feasible by approximating the working-model EIC for Ψ_s inside the score span of the larger targeting model. Write

$$\mathcal{M}_n = \mathcal{M}_{n,\text{targ}} = \{P_\theta : \theta \in \Theta_n\}$$

for the finite-dimensional empirical targeting model. For a fixed working model, the exact efficient influence curve of the projected parameter uses the population Hessian of the likelihood loss, not the score covariance. The projection approach takes a different computational route. It starts from a valid nonparametric gradient D_s and projects that gradient onto the score span. Without regularization, this yields

$$D_{s,\theta}^*(O) = \alpha_s(\theta)^\top S_\theta(O), \quad \text{where} \quad \alpha_s(\theta) = \left(P_n[S_\theta S_\theta^\top]\right)^{-1} P_n[S_\theta D_s]. \quad (6.5)$$

Thus the projection and Hessian constructions have the same algebraic form but use different information matrices. For a fixed misspecified working model, these are genuinely different objects. Their asymptotic equivalence relies on two ingredients from the regHAL framework: the loss must be likelihood-based, so Bartlett's identity is available, and the targeting working model must grow toward an oracle model containing P_0 .

The practical obstacle is that the theoretical working-model EIC and its empirical coordinate representation behave very differently in this HAL geometry. Because HAL selects highly collinear basis functions to control sectional variation, the empirical score covariance matrix $P_n[S_\theta S_\theta^\top]$ is often nearly singular. The geometric projection of the nonparametric EIC onto the working-model score span remains a well-defined $L_2(P_0)$ object, but its exact unregularized coefficient vector $\alpha_s(\theta)$ can be numerically unstable. This is why delta-method inference based on the raw inverse information matrix performs poorly in the simulations.

Following the projection-based regularized A-TMLE construction established in Chapter 3, we therefore estimate a stable approximation of the projected EIC through a lasso-penalized score regression:

$$\hat{\alpha}_{s,\text{reg}}(\theta) = \arg \min_{\alpha \in \mathbb{R}^p} \left\{ P_n \left[D_s(O; \hat{P}_\theta, \hat{G}) - \alpha^\top S_\theta(O) \right]^2 + \lambda_{\text{proj}} \|\alpha\|_1 \right\}. \quad (6.6)$$

The resulting projected EIC approximation is $\widetilde{D}_{s,\theta}^*(O) = \hat{\alpha}_{s,\text{reg}}(\theta)^\top S_\theta(O)$. The ℓ_1 penalty

bypasses the ill-conditioned inversion by selecting a sparse and stable coordinate representation of the projected direction. For each fixed grid point s , the targeting goal is to make the empirical mean $P_n \widetilde{D}_{s,\theta}^*$ negligible.

6.4.2.4 Multi-parameter targeting via universal direction

For a generic targeting grid $\mathcal{G}_n = \{s_1, \dots, s_{m_n}\}$, write

$$d(\theta) = \left(d_1(\theta), \dots, d_{m_n}(\theta)\right)^\top, \quad d_\ell(\theta) = P_n \widetilde{D}_{s_\ell, \theta}^*,$$

for the vector of pointwise empirical projected scores we want to drive toward zero. As a geometric motivation, one can aggregate these pointwise projected directions through a universal least-favorable fluctuation following van der Laan and Gruber [2016] and the Chapter 3 construction:

$$\hat{\alpha}_{\text{univ}}(\theta) = \sum_{\ell=1}^{m_n} w_\ell(\theta) \hat{\alpha}_{s_\ell, \text{reg}}(\theta), \quad (6.7)$$

where $w_\ell(\theta) = d_\ell(\theta) / \|d(\theta)\|_2$. Define the aggregated projected-EIC fluctuation by

$$\widetilde{D}_{\text{univ}, \theta}^*(O) = \sum_{\ell=1}^{m_n} w_\ell(\theta) \widetilde{D}_{s_\ell, \theta}^*(O) = \hat{\alpha}_{\text{univ}}(\theta)^\top S_\theta(O).$$

Its empirical mean is

$$P_n \widetilde{D}_{\text{univ}, \theta}^* = \sum_{\ell=1}^{m_n} \frac{d_\ell(\theta)^2}{\|d(\theta)\|_2} = \|d(\theta)\|_2. \quad (6.8)$$

Since the derivative of the empirical log-likelihood along a fluctuation submodel equals the empirical mean of its score, (6.8) shows that the universal direction has first-order gain $\|d(\theta)\|_2$. Moreover, for any normalized aggregation $a \in \mathbb{R}^{m_n}$ with $\|a\|_2 = 1$,

$$P_n \sum_{\ell=1}^{m_n} a_\ell \widetilde{D}_{s_\ell, \theta}^* = a^\top d(\theta) \leq \|d(\theta)\|_2$$

by Cauchy–Schwarz, so (6.7) is the normalized aggregate direction with maximal first-order likelihood increase. Thus $\|d(\theta)\|_2$ is not itself the optimization criterion; it is the score diagnostic induced by the universal fluctuation. The censored-data implementation below does not update by a direct small step along (6.7); instead, it uses an EM reparameterization of these projected directions through a derived targeting basis. Finally,

$$|d_\ell(\theta)| \leq \|d(\theta)\|_2,$$

so small joint norm implies small componentwise empirical projected scores. The theoretical justification below remains scalar: the implemented algorithm applies the fixed- s construction across $s_1, \dots, s_{m_n}$ and then combines the resulting directions through (6.7). In the present implementation $m_n = 81$.

6.4.2.5 Targeting EM implementation

The implemented projection-based update follows the censored-data EM reparameterization developed for this setting. At each outer iteration, the lasso projections are recomputed, yielding a coefficient matrix

$$C(\theta) \in \mathbb{R}^{m_n \times p}, \quad C_{\ell,\cdot}(\theta) = \hat{\alpha}_{s_\ell, \text{reg}}(\theta)^\top,$$

and hence the targeting basis vector

$$X_{\text{targ}}(t) = C(\theta) \phi(t), \quad (6.9)$$

whose components are linear combinations of the original HAL basis functions. Holding the current fixed component θ as an offset, we then target within the two-component submodel

$$\log f_{\theta, \alpha}(t) = \phi(t)^\top \theta + X_{\text{targ}}(t)^\top \alpha - \log Z(\theta, \alpha), \quad (6.10)$$

with $\alpha \in \mathbb{R}^{m_n}$ initialized at 0. The inner targeting-EM maximizes the observed-data likelihood over α in (6.10); after convergence, the resulting targeting contribution is absorbed back into the fixed component through $\theta \leftarrow \theta + C(\theta)^\top \hat{\alpha}$, α is reset to zero, and the outer projection step is recomputed. Algorithm 9 summarizes this implementation, adapted to the present survival setting from the regHAL-style projection construction of Chapter 3.

Algorithm 9 Projection-based reg-ATMLE update

- 1: **Input:** EM-HAL event-time fit, targeting grid $\mathcal{G}_n = \{s_\ell\}_{\ell=1}^{m_n}$, censoring model, score basis ϕ
 - 2: Fit the event-time model and the corresponding censoring model to obtain $\hat{G}(t_1-, t_2-)$
 - 3: Let $\mathcal{M}_{n, \text{init}}$ be the initial EM-HAL working model and construct the larger targeting model $\mathcal{M}_{n, \text{targ}}$ by relaxing the HAL penalty and augmenting the active basis set
 - 4: Embed the EM-HAL coefficients in $\mathcal{M}_{n, \text{targ}}$ by treating the fitted log-density as an offset and initializing any newly added coefficients at zero
 - 5: Initialize θ at the resulting parameter value in $\mathcal{M}_{n, \text{targ}}$
 - 6: **for** outer iterations until convergence **do**
 - 7: Compute the IPCW gradients $D_{s_\ell}(O_i)$ over the targeting grid
 - 8: Form the HAL score matrix $S(\theta)$ in the enlarged targeting model $\mathcal{M}_{n, \text{targ}}$
 - 9: Solve the regularized score regressions in (6.6) to obtain $C(\theta)$ and define the targeting basis vector $X_{\text{targ}}(t) = C(\theta)\phi(t)$
 - 10: Run an inner targeting-EM for α in the offset model (6.10), holding θ fixed
 - 11: Update $\theta \leftarrow \theta + C(\theta)^\top \hat{\alpha}$ and reset $\alpha \leftarrow 0$
 - 12: Evaluate the empirical mean of the targeted score vector induced by X_{targ}
 - 13: **end for**
 - 14: **Output:** projection-targeted density and survival estimates
-

In our implementation, $\mathcal{G}_n = \{0.1, \dots, 0.9\}^2$ so $m_n = 81$, the larger targeting model is

built with relaxation factor $\kappa = 6$, and the outer loop terminates once the norm of the empirical targeted score vector induced by X_{targ} satisfies $\|P_n S_{\text{targ}}(O)\|_2 / \sqrt{m_n} < 10^{-3}$ or 50 outer iterations have been reached, thereby operationalizing Assumption 6.1.

6.4.2.6 Regularized A-TMLE theory via the working model

The next results adapt the scalar working-model expansions from the A-TMLE framework [van der Laan et al., 2023], together with the smooth M-estimand / ADML continuation [van der Laan et al., 2025], to fixed survival-grid points in the present coarsened-data problem. They operate inside the larger targeting working model and replace the exact working-model EIC with the regularized projection described above, so that the asymptotic justification comes from the inherited adaptive-targeting theory while the computable targeting step follows the Chapter 3 projection-based reg-ATMLE mechanism. Fix $s \in \mathcal{G}_n$ and write

$$\Psi_{\mathcal{M}_n, s}(P) = \Psi_s(\Pi_{\mathcal{M}_n} P)$$

for the working-model-induced scalar target. Let $D_{\mathcal{M}_n, s}^*(P)$ denote the corresponding efficient influence curve in the full nonparametric CAR model $\mathcal{M}$, and define the exact scalar remainder

$$R_{\mathcal{M}_n, s}(P, P_0) = \Psi_{\mathcal{M}_n, s}(P) - \Psi_{\mathcal{M}_n, s}(P_0) + P_0 D_{\mathcal{M}_n, s}^*(P).$$

The projection-based targeting algorithm outputs an estimator P_n^* that is intended to make the empirical mean of a regularized approximation $\tilde{D}_{n, s}^*$ asymptotically negligible at each fixed targeting point.

Assumption 6.1 (Approximation of the working model EIC). *For a fixed grid point $s \in \mathcal{G}_n$, let $\tilde{D}_{n, s}^*$ denote the regularized EIC approximation used by the projection-based update. We assume the targeted fit P_n^* satisfies the approximate scalar score equation*

$$P_n \tilde{D}_{n, s}^*(P_n^*) = o_P(n^{-1/2}),$$

and that this implies the corresponding exact working-model score equation for the projected parameter is also solved at the same rate:

$$P_n D_{\mathcal{M}_n, s}^*(P_n^*) = o_P(n^{-1/2}).$$

The first theorem below gives the fixed-grid working-model result for the projected target $\Psi_{\mathcal{M}_n, s}$, while the second applies the same A-TMLE logic to transfer that conclusion to the original survival target when the oracle model is the full model. In both results, the computable fluctuation is still the Chapter 3 projection-based reg-ATMLE update, specialized here to the EM-HAL survival construction.

Theorem 6.2 (Regularized A-TMLE for a fixed working-model-induced survival probability, adapted from van der Laan et al. [2023]). *Fix $s \in \mathcal{G}_n$ and let P_n^* denote the targeted estimator satisfying Assumption 6.1, so that $P_n D_{\mathcal{M}_n, s}^*(P_n^*) = o_P(n^{-1/2})$. Assume:*

(A.1) $D_{\mathcal{M}_{n,s}}^*(P_n^*)$ and $D_{\mathcal{M}_{n,s}}^*(P_0)$ lie in a P_0 -Donsker class with probability tending to 1.

(A.2) $P_0\{D_{\mathcal{M}_{n,s}}^*(P_n^*) - D_{\mathcal{M}_{n,s}}^*(P_0)\}^2 = o_P(1)$.

(A.3) $R_{\mathcal{M}_{n,s}}(P_n^*, P_0) = o_P(n^{-1/2})$.

Then

$$\Psi_{\mathcal{M}_{n,s}}(P_n^*) - \Psi_{\mathcal{M}_{n,s}}(P_0) = P_n D_{\mathcal{M}_{n,s}}^*(P_0) + o_P(n^{-1/2}).$$

Theorem 6.3 (Regularized A-TMLE for a fixed survival probability via the working model, adapted from van der Laan et al. [2023]). *Fix $s \in \mathcal{G}_n$ and assume the oracle model $\mathcal{M}_0$ is the full model $\mathcal{M}$, so that $D_{\mathcal{M}_0,s}^*(P_0) = D_{\Psi_s}^*(P_0)$. Let $P_{0,\mathcal{M}_n}$ be the loss-based projection of P_0 onto $\mathcal{M}_n$. Assume the conditions of Theorem 6.2 hold, and suppose:*

(B.1) $P_0\{D_{\mathcal{M}_{n,s}}^*(P_0) - D_{\Psi_s}^*(P_0)\}^2 = o_P(1)$.

and

(C.2) $(P_{0,\mathcal{M}_n} - P_0)(D_{\Psi_s}^*(P_{0,\mathcal{M}_n}) - D_{\mathcal{M}_{n,s}}^*(P_{0,\mathcal{M}_n})) = o_P(n^{-1/2})$.

(C.3) $P_0 D_{\mathcal{M}_{n,s}}^*(P_{0,\mathcal{M}_n}) = o_P(n^{-1/2})$.

Under $\mathcal{M}_0 = \mathcal{M}$, condition (C.1) from the general A-TMLE framework holds trivially:

$$\Psi_s(P_{0,\mathcal{M}_n}) - \Psi_{\mathcal{M}_0,s}(P_{0,\mathcal{M}_n}) = 0.$$

Then the targeted plug-in survival estimate at s

$$\hat{S}^{\text{targ}}(s) = \Psi_{\mathcal{M}_{n,s}}(P_n^*) = \Psi_s(P_n^*)$$

is asymptotically linear with influence curve $D_{\Psi_s}^*(P_0)$:

$$\hat{S}^{\text{targ}}(s) - \Psi_s(P_0) = \frac{1}{n} \sum_{i=1}^n D_{\Psi_s}^*(P_0)(O_i) + o_P(n^{-1/2}), \quad (6.11)$$

for each fixed targeting-grid point s .

Verification of the regularized A-TMLE assumptions.

For the bivariate right-censoring problem, the key structural issue is that the working-model EIC is geometrically well defined but numerically hard to represent without regularization. The targeting working model contains many collinear HAL basis functions, so the unregularized inverse information matrix is unstable even when the projected EIC itself is a perfectly regular $L_2(P_0)$ object. Assumption 6.1 therefore relies on the lasso projection in (6.6) to recover a stable empirical representation of that projected direction. Solving the approximate score equation with this regularized projection is the operational substitute for directly solving the exact working-model score equation.

Efficiency relative to the full model is obtained by allowing the targeting working model $\mathcal{M}_{n,\text{targ}}$ to expand beyond the initial EM-HAL fit. As the relaxed HAL model grows, the score span better approximates the full tangent space, and Bartlett’s identity for the CAR likelihood links the projected IPCW gradient to the full-model efficient influence curve. This is the same oracle-model logic used in the earlier regularized A-TMLE chapters: finite samples may behave as though the estimator targets a smaller oracle object, while asymptotically the working model grows toward the full nonparametric model.

The empirical-process requirements are supported by the regularity of the EM-HAL density and by the dual ℓ_1 constraints on the initial fit and the projected EIC coefficients. In particular, the continuous EM-HAL density smooths the singly censored half-line coarsening sets that make fully discrete NPMLE-based arguments unstable, and the bounded sectional-variation structure keeps the relevant score and projected-EIC classes inside a manageable Donsker envelope. Finally, the exact remainders are controlled by the EM-HAL loss-based convergence rate: the working-model remainder is of squared-density-error order, while the full-model remainder has a double-robust product form involving the event-time and censoring estimators. With the EM-HAL rate recorded in Appendix D.2 and a sufficiently accurate censoring estimator, both contributions are $o_P(n^{-1/2})$ for each fixed targeting-grid point.

These results are specific to the projection-based reg-ATMLE update. For each fixed targeting-grid point s , they justify the update as a regularized A-TMLE and provide the scalar first-order expansion used below to motivate jackknife inference. The implemented algorithm targets the full finite grid by aggregating the pointwise projected directions; any spillover from the targeting grid to the denser evaluation grid is assessed empirically in the simulations.

6.5 Inference via V-fold jackknife

6.5.1 Motivation and theoretical validity

For bivariate right-censoring, neither a tractable closed-form efficient influence curve nor a simple plug-in variance formula is available for the survival surface. As Section 6.7.5 will show, the delta method based on the HAL information matrix is inflated by basis collinearity, while the projected-influence-curve variance from the projection-based update remains too sensitive to regularization in finite samples. A bootstrap is also not practically viable because each refit of the full EM-HAL-plus-targeting pipeline is expensive. The grouped V -fold jackknife is attractive because it reduces that burden to V refits and does not require direct EIC estimation.

The jackknife is justified only to the extent that targeting moves the estimator closer to an asymptotically linear representation for the survival functional. Section 6.4 makes this explicit for the projection-based reg-ATMLE update at any fixed targeting-grid point through Theorem 6.3. Self-consistency targeting is justified more modestly through control of a precision-chosen self-consistency score system rather than through the same working-model expansion.

This motivates the jackknife as a practical inferential device, but not as a consequence of one uniform theorem for both updates over the full 209×209 evaluation grid. We therefore use it operationally for both targeted estimators and let the simulations assess whether the resulting pseudo-values behave as though the leading regularization bias has been removed. Without targeting, jackknife coverage collapses to about 43%; with targeted estimators, it returns to the nominal range.

6.5.2 Variance estimation and confidence intervals

Let $\mathcal{T}$ be a fixed evaluation grid. We partition the sample $\{1, \dots, n\}$ into V disjoint folds $I_1, \dots, I_V$ of approximately equal size n/V . For each $t \in \mathcal{T}$, let $\hat{S}(t)$ denote the full-sample targeted estimate and $\hat{S}_{-v}(t)$ the leave-fold-out targeted estimate computed on the training sample $I_{-v} = \{1, \dots, n\} \setminus I_v$.

Following the grouped V-fold jackknife methodology of Chapter 5, inference is based on the fold-level pseudo-values:

$$IC_v(t) = V \hat{S}(t) - (V - 1) \hat{S}_{-v}(t), \quad v = 1, \dots, V. \quad (6.12)$$

Let $\overline{IC}(t) = V^{-1} \sum_{v=1}^V IC_v(t)$. The V -fold jackknife sample variance is given by:

$$\hat{\sigma}_{\text{JK}}^2(t) = \frac{1}{V-1} \sum_{v=1}^V \{IC_v(t) - \overline{IC}(t)\}^2. \quad (6.13)$$

The standard error of the survival estimator is $\widehat{\text{SE}}_{\text{JK}}(t) = \hat{\sigma}_{\text{JK}}(t)/\sqrt{V}$.

Relying on the fixed- V theoretical justification, we construct pointwise $(1 - \alpha)$ -level confidence intervals using the t_{V-1} critical value:

$$\widehat{\text{CI}}_{1-\alpha}(t) = \left[\hat{S}(t) - t_{V-1, 1-\alpha/2} \widehat{\text{SE}}_{\text{JK}}(t), \hat{S}(t) + t_{V-1, 1-\alpha/2} \widehat{\text{SE}}_{\text{JK}}(t) \right], \quad (6.14)$$

clipped to $[0, 1]$ to respect the bounds of survival probabilities.

6.5.3 Implementation

To accurately capture the full sampling variability, it is essential that every data-adaptive stage be refit within each jackknife training sample rather than recycled from the full-data analysis. Algorithm 10 summarizes the resulting fold-wise implementation. All folds use the same evaluation grid and numerical settings so that the pseudo-values in (6.12) are directly comparable across leave-fold-out fits.

6.5.4 Choice of V

The choice of V involves a tradeoff between statistical precision and computational cost. Standard rate calculations suggest $V = O(\sqrt{n})$ is sufficient for valid Wald confidence intervals. For our sample size of $n = 500$, we use $V = 20$ across all simulations. The corresponding

Algorithm 10 V-fold jackknife refitting and survival evaluation

- 1: **Input:** observed data $\{O_i\}_{i=1}^n$, folds $\{I_v\}_{v=1}^V$, evaluation grid $\mathcal{T}$, chosen targeting method
 - 2: Fit the full-sample targeted estimator and evaluate $\hat{S}(t)$ for all $t \in \mathcal{T}$
 - 3: **for** $v = 1, \dots, V$ **do**
 - 4: Restrict the observed data to the training sample $I_{-v} = \{1, \dots, n\} \setminus I_v$
 - 5: Refit the EM-HAL density estimator on I_{-v} from the common warm-start specification
 - 6: Refit the targeting procedure on I_{-v} ; for projection-based targeting also refit the censoring model on I_{-v}
 - 7: Evaluate the leave-fold-out targeted estimate $\hat{S}_{-v}(t)$ for all $t \in \mathcal{T}$
 - 8: **end for**
 - 9: Form the pseudo-values $IC_v(t)$ in (6.12) and the variance estimate in (6.13)
 - 10: **Output:** jackknife standard errors and confidence intervals over $\mathcal{T}$
-

$t_{19,0.975} = 2.093$ critical value provides a necessary finite-sample correction over the standard normal $z_{0.975} = 1.96$ multiplier.

6.6 Simultaneous confidence bands

The pointwise confidence intervals defined in (6.14) guarantee coverage at individual grid points, but global inference on the survival surface requires simultaneous confidence bands over an evaluation grid $\mathcal{G} = \{s_1, \dots, s_m\}$. A simultaneous band replaces the pointwise critical value with a calibrated multiplier q such that $P(S_{\text{true}}(s) \in [\hat{S}(s) \pm q \cdot \widehat{\text{SE}}_{\text{JK}}(s)] \text{ for all } s \in \mathcal{G}) = 1 - \alpha$.

6.6.1 Monte Carlo calibration

Standard multivariate normal or multivariate t approximations are invalid in this setting because they assume a single shared scaling denominator. Instead, the componentwise-Studentized jackknife pseudo-values limit to a joint distribution with dependent Wishart-diagonal denominators, exactly the setting treated in Chapter 5.

We therefore calibrate q directly from this exact limit distribution using the Monte Carlo procedure developed in Chapter 5. From the $V \times m$ matrix of pseudo-values $IC_v(s)$ evaluated over $\mathcal{G}$, we compute the $m \times m$ sample correlation matrix $\hat{R}$. Per Monte Carlo step, we simulate V independent Gaussian vectors from $\mathcal{N}_m(\mathbf{0}, \hat{R})$, calculate the sample mean and variance for each component, form the componentwise-Studentized statistics, and record their maximum absolute value across the m components. The $(1 - \alpha)$ -quantile of these simulated maxima yields the correctly calibrated simultaneous critical value q .

6.6.2 Filtered simultaneous confidence bands

The simultaneous coverage guarantee relies on the uniform validity of the asymptotic linear expansion across $\mathcal{G}$. In finite samples, elevated remainder bias in boundary regions (e.g., heavy censoring tails) can degrade overall simultaneous coverage. Following Chapter 5, we utilize the classic jackknife bias estimate:

$$\hat{b}(s) = (V - 1)(\bar{S}_{(-v)}(s) - \hat{S}(s)), \quad \bar{S}_{(-v)}(s) = \frac{1}{V} \sum_{v=1}^V \hat{S}_{-v}(s). \quad (6.15)$$

We compute the bias-to-standard-error ratio $\rho(s) = |\hat{b}(s)|/\widehat{\text{SE}}_{\text{JK}}(s)$. By specifying a threshold $\tau > 0$, we restrict the simultaneous band to the adaptively selected, well-targeted subset $\mathcal{G}(\tau) = \{s \in \mathcal{G} : \rho(s) < \tau\}$. The critical value q_τ is then computed specifically from the submatrix $\hat{R}[\mathcal{G}(\tau), \mathcal{G}(\tau)]$ via the Monte Carlo algorithm.

6.6.3 Bias correction

Alternatively, rather than excluding biased points, we can retain the full grid and apply a truth-free bias correction that selectively widens the confidence band by the absolute estimated bias:

$$\text{CI}_q^{\text{BC}}(s) = [\hat{S}(s) - |\hat{b}(s)| - q \cdot \widehat{\text{SE}}_{\text{JK}}(s), \hat{S}(s) + |\hat{b}(s)| + q \cdot \widehat{\text{SE}}_{\text{JK}}(s)].$$

This approach safely enlarges intervals at high-bias regions while leaving well-estimated interior points largely unaffected. The filtered and bias-corrected approaches may also be combined by filtering to $\mathcal{G}(\tau)$ and subsequently applying the absolute bias correction to the remaining grid points.

6.7 Simulation study

We evaluate the full estimation-and-inference pipeline through a Monte Carlo study. We first assess the EM-HAL density estimate, then examine survival point estimation before and after targeting, followed by pointwise jackknife intervals and simultaneous confidence bands. Throughout, the projection-based reg-ATMLE update is the primary main-text method, while self-consistency targeting is the comparative alternative detailed in Appendix D.3.

Monte Carlo setup. The grouped $V = 20$ jackknife already requires 20 refits of the full EM-HAL-plus-targeting pipeline per data set, so the simulation study is computationally intensive even without bootstrap resampling. For that reason, the $R = 213$ simulation batches for self-consistency targeting and the projection-based reg-ATMLE update were run independently on a distributed cluster with different random seeds. Consequently, the baseline EM-HAL metrics show minor Monte Carlo variation across later tables. With $R = 213$ replications, the Monte Carlo standard error for a nominal 95% coverage estimate is

approximately $\sqrt{0.95 \times 0.05/213} \approx 0.015$, so coverages in roughly $[0.92, 0.98]$ are statistically indistinguishable from 95%.

The initial HAL fit on imputed data serves only as a warm start for the EM algorithm and is not reported.

6.7.1 Data-generating process

We generate the event-time pair from a Gaussian copula with parameter $\rho = 0.7071$, corresponding to Kendall's $\tau = 0.5$, combined with truncated normal marginals $T_j \sim \text{TruncNorm}(\mu = 0.5, \sigma^2 = 0.05; [0, 1])$ for $j = 1, 2$. The censoring variables satisfy $C_1, C_2 \stackrel{\text{i.i.d.}}{\sim} \text{Unif}(0, 1)$ and are independent of (T_1, T_2) , yielding approximately 50% censoring in each coordinate. The true survival surface $S_{\text{true}}(t_1, t_2)$ and density $f_{\text{true}}(t_1, t_2)$ are computed by numerical evaluation of the Gaussian copula density with truncated normal marginals on the evaluation grid.

Table 6.1: Simulation parameters.

Parameter	Value
Sample size n	500
Monte Carlo replications R	213
V-fold jackknife folds V	20
t critical value (df = 19)	2.093
Evaluation grid $\mathcal{T}$	209×209 on $[0, 1]^2$ (43,681 points; see Section 6.3.4)
Targeting grid $\mathcal{G}$	9×9 on $\{0.1, \dots, 0.9\}^2$ (81 points)
Confidence level	95%

Throughout the simulations, the 9×9 targeting grid is the common implementation choice for both methods. We treat it as a computational precision setting rather than a theoretically fixed dimension.

6.7.2 Summary metrics

For a generic estimand $\psi(s)$ evaluated at grid point $s = (t_1, t_2)$, with estimate $\hat{\psi}^{(r)}(s)$ in replication r , we compute:

$$\begin{aligned} \text{Bias}(s) &= \bar{\psi}(s) - \psi_0(s), & \text{MSE}(s) &= \frac{1}{R} \sum_{r=1}^R \left(\hat{\psi}^{(r)}(s) - \psi_0(s) \right)^2, \\ |\text{Bias}|/\text{SD}(s) &= \frac{|\text{Bias}(s)|}{\text{SD}(s)}, & \text{SD}(s) &= \left[\frac{1}{R-1} \sum_{r=1}^R \left(\hat{\psi}^{(r)}(s) - \bar{\psi}(s) \right)^2 \right]^{1/2}. \end{aligned}$$

The ratio $|\text{Bias}|/\text{SD}$ is a key diagnostic: values near or above 1 indicate that bias dominates sampling variability, making oracle confidence intervals (using the true SD) unable to achieve

nominal coverage. Oracle coverage is defined as the fraction of replications in which $\psi_0(s)$ falls within $\hat{\psi}^{(r)}(s) \pm 1.96 \cdot \text{SD}(s)$. All metrics are averaged over the evaluation grid unless otherwise noted.

6.7.3 Density empirical performance (EM-HAL)

The EM-HAL framework produces both density and survival estimates. We begin with the density, which is the direct output of the EM algorithm before any survival targeting. Table 6.2 reports performance on interior grid points, excluding boundary rows where the truncated normal density is extreme.

Table 6.2: Density empirical performance: EM-HAL estimate on the interior of the 209×209 evaluation grid, excluding boundary points ($R = 213$, $n = 500$).

Mean Bias	Mean Bias	Mean Bias /SD	Mean Var	Mean MSE	Oracle Cov
−0.0019	0.1316	0.566	0.1265	0.1696	0.898

Since targeting has not yet been applied, the density at this stage is identical across pipelines and serves as the common baseline for the rest of the study. The mean $|\text{Bias}|/\text{SD}$ ratio of about 0.57 indicates that sampling variability still dominates bias, and oracle coverage of 0.898 is respectable though below the nominal 95%. Figure 6.1 shows that EM-HAL captures the broad unimodal shape of the Gaussian-copula density. The remaining subsections focus on survival; the effect of survival targeting on density is deferred to Appendix D.6.




Figure 6.1: EM-HAL density estimation: true bivariate density (left), mean estimate across $R = 213$ replications (center), and pointwise difference (right).

6.7.4 Empirical performance of survival point estimation

We now turn to the primary target parameter, the bivariate survival function $S(t_1, t_2)$. For each targeting strategy, we compare the EM-HAL and targeted stages to isolate what

targeting contributes beyond the common baseline. Table 6.3 summarizes the average point-estimation performance of self-consistency targeting and the projection-based reg-ATMLE update side by side.

Table 6.3: Survival empirical performance of self-consistency targeting and the reg-ATMLE update. Metrics are averaged over the 209×209 evaluation grid ($R = 213$, $n = 500$).

Method	Stage	Mean Bias	Bias	Bias /SD	Mean Var	Mean MSE	Oracle Cov
Self-consistency targeting	EM-HAL	0.00417	0.00632	7.518	0.000345	0.000421	0.904
	Targeted	0.00312	0.00480	7.231	0.000427	0.000463	0.925
	Change (%)		-24.1%	-3.8%	+24.0%	+10.1%	+2.1pp
reg-ATMLE	EM-HAL	0.00430	0.00670	7.606	0.000345	0.000470	0.904
	Targeted	0.00230	0.00320	7.140	0.000311	0.000389	0.931
	Change (%)		-52.2%	-6.1%	-9.8%	-17.2%	+2.8pp

Projection-based reg-ATMLE update

The projection-based reg-ATMLE update delivers the strongest overall point-estimation performance in this Gaussian-copula design. Relative to EM-HAL, mean absolute bias drops by 52%, variance drops by about 10%, and MSE drops by 17%. The gains are broad-based as well: 84.0% of grid points show smaller absolute bias and 67.7% show smaller MSE.

Spatial pattern. Figure 6.2 is the main spatial diagnostic for the projection-based reg-ATMLE update. MSE improvement is broad across the interior and lower-left region, while deterioration is largely confined to a thin strip near the heavily censored boundary. Additional bias-to-standard-deviation and oracle-coverage diagnostics appear in Appendix D.4.




Figure 6.2: reg-ATMLE: MSE heatmaps for EM-HAL (left), targeted (center), and log₁₀ ratio (right). Blue indicates improvement.

Self-consistency targeting

Self-consistency targeting also improves centering relative to EM-HAL, but it does so through a rougher correction. Mean absolute bias falls by 24% and oracle coverage rises from 90.4% to 92.5%, but variance increases enough to raise MSE by 10%. That pattern matches the indicator-based geometry of the update. Appendix D.4 provides the spatial diagnostics, and Appendix D.3 gives the full comparative construction.

6.7.5 Pointwise confidence intervals

The next question is whether those point-estimation gains translate into practical inference. Across both methods, the main empirical issue is the same: does targeting remove enough first-order bias for the grouped V -fold jackknife to behave like a useful inferential device on the survival surface?

Table 6.4 compares the available inference methods for self-consistency targeting and the reg-ATMLE update.

Table 6.4: Survival CI comparison ($R = 213$, $V = 20$). Coverage and mean CI length are averaged over the 209×209 evaluation grid.

Method	Inference Method	Coverage	Mean CI Length
reg-ATMLE	Targeted + Delta method	0.988	1.314
	Targeted + Projected IC	0.600	0.104
	Targeted + Jackknife ($V = 20$)	0.947	0.083
	EM-HAL + Jackknife ($V = 20$)	0.431	0.021
Self-consistency targeting	Targeted + Delta method	0.989	2.494
	Targeted + Jackknife ($V = 20$)	0.951	0.090
	EM-HAL + Jackknife ($V = 20$)	0.436	0.023

reg-ATMLE inference

For the projection-based reg-ATMLE update, Table 6.4 supports a clear conclusion: targeting is what makes jackknife inference viable, while the projected-IC variance formula remains unreliable in finite samples. The targeted jackknife reaches 94.7% coverage with mean CI length 0.083, close to the nominal 95% target. By contrast, the projected-IC variance formula undercovers at 60.0%, and the untargeted jackknife fails with only 43.1% coverage.

Figure 6.3 gives the corresponding spatial diagnostic: the targeted jackknife is close to nominal across most of the interior, with the main residual deficit confined to the heavily censored upper-right boundary, whereas the projected-IC method undercovers much more broadly.

Targeted CI Methods: Delta vs Projection vs Jackknife
(213 simulations, V20, gaussian)




Figure 6.3: reg-ATMLE: coverage heatmaps comparing CI methods for the targeted survival estimate.

Self-consistency targeting inference

For self-consistency targeting, the ranking is simpler but qualitatively the same: the targeted jackknife is the only practically useful interval procedure among those considered. It attains near-nominal coverage (95.1%) with mean CI length 0.090, whereas the untargeted jackknife fails badly and the delta intervals are far too wide. Appendix D.5 gives the supporting spatial and regional diagnostics.

6.7.5.1 Oracle vs. jackknife comparison

Table 6.5 shows how close the targeted jackknife comes to oracle inference once targeting has been carried out.

Table 6.5: Oracle vs. jackknife CI comparison for targeted estimates ($R = 213$, $V = 20$). Averaged over the 209×209 evaluation grid.

Method	Oracle Cov.	JK Cov.	Oracle CI Len.	JK CI Len.
Self-consistency				
targeting	0.925	0.951	0.076	0.090
reg-ATMLE	0.931	0.947	0.070	0.083

For both targeted estimators, the jackknife CIs are about 19% wider than the oracle CIs, so the price of practical validity is only mild conservatism. The more important point is that, after targeting, the jackknife closely tracks the true sampling variability rather than missing it by the margins seen for the untargeted estimator.

6.7.5.2 Regional analysis: central vs. tail and target band vs. non-target

We partition the 209×209 evaluation grid in two ways to examine spatial variation in CI performance: (i) a central region $[0.25, 0.75]^2$ (11,025 points) vs. a tail region (32,656 points), and (ii) the 81 points on the 9×9 targeting grid $\{0.1, \dots, 0.9\}^2$ vs. the remaining 43,600 non-target points.

Central vs. tail (see Appendix Figure D.7). The regional breakdown shows that targeted-jackknife validity is not confined to a small subset of the surface. In the central region, coverage is 97.1% for self-consistency targeting and 96.7% for the projection-based reg-ATMLE update; in the tail region it drops only slightly to 94.4% and 94.1%. By contrast, the untargeted EM-HAL jackknife coverage remains near 43–44% in both regions.

Target band vs. non-target (see Appendix Figure D.8). On the 81 targeting-grid points, targeted jackknife coverage reaches 96.6% for self-consistency targeting and 96.2% for the projection-based reg-ATMLE update. On the 43,600 non-target points, coverage is only slightly lower at 95.1% and 94.7%. This modest gap indicates that the benefit of targeting is not confined to the grid itself.

6.7.6 Simultaneous confidence bands

The simultaneous-band results extend the pointwise story: targeting makes jackknife-based surface inference feasible, but multiplicity still requires additional bias control. Table 6.6 summarizes simultaneous coverage, per-point average coverage, band width, and the calibrated critical value on the 81-point targeting grid using the componentwise-Studentized Monte Carlo calibration from Section 6.6.1.

Table 6.6: Simultaneous coverage on the 9×9 targeting grid (81 points, $R = 213$).

Metric	Self-consistency	
	targeting	reg-ATMLE
<i>Without bias correction</i>		
Simultaneous coverage	0.897	0.925
Per-point average coverage	0.966	0.962
Mean band width	0.192	0.155
<i>With bias correction</i>		
Simultaneous coverage	0.953	0.977
Per-point average coverage	0.994	0.995
Mean band width	0.291	0.233
<i>Calibration</i>		
Mean critical value q	3.68	3.53
$q/t_{19, 0.975}$ ratio	1.76	1.69

Without bias correction, simultaneous coverage falls short of nominal for both strategies despite high per-point coverage. The bias correction from Section 6.6.3 raises simultaneous coverage to 95.3% for self-consistency targeting and 97.7% for the projection-based reg-ATMLE update, at the cost of wider bands.

Table 6.7 reports partial simultaneous coverage—the fraction of simulations in which at least $k\%$ of the 81 target points are simultaneously covered.

Table 6.7: Partial simultaneous coverage (bias-corrected): fraction of simulations covering $\geq k\%$ of 81 points.

Threshold $k\%$	Self-consistency	
	targeting	reg-ATMLE
90% ($\geq 73/81$ points)	1.000	0.991
95% ($\geq 77/81$ points)	0.991	0.991
100% (all 81 points)	0.953	0.977

Table 6.7 shows that the remaining failures are localized rather than catastrophic. Nearly all simulations cover at least 95% of the 81 target points, so when full simultaneous coverage fails the miss is usually limited to one or two tail points.

Table 6.8 makes the same point from the opposite direction. Filtering to grid points with small bias-to-SE ratio $\rho(s) < \tau$ recovers nominal coverage on the surviving subset, and combining that filter with bias correction raises coverage above 95% even for moderate thresholds. Once multiplicity is layered on top of the targeted jackknife, the remaining problem is concentrated at higher-bias locations rather than spread uniformly across the grid.

Table 6.8: Filtered simultaneous coverage for survival on the targeting grid, using the jackknife bias proxy ($R = 213$).

τ	Avg. $ \mathcal{G}(\tau) $	Without bias correction			With bias correction		
		Self-consistency targeting	reg-ATMLE	Width	Self-consistency targeting	reg-ATMLE	Width
0.50	14–16	0.967	0.953	0.16	0.977	0.967	0.17
1.00	27–30	0.962	0.967	0.17	0.977	0.972	0.20
1.50	40–42	0.944	0.962	0.18	0.967	0.977	0.22
∞ (all 81)	81	0.897	0.925	0.19	0.953	0.977	0.29

Taken together, the simulation study supports three conclusions in this Gaussian-copula setting. First, EM-HAL is a useful starting estimator but not by itself enough for reliable survival inference. Second, both targeting strategies improve the survival estimator relative to untargeted EM-HAL, with the projection-based reg-ATMLE update providing the stronger overall bias-variance tradeoff in this study. Third, grouped V -fold jackknife becomes credible only after targeting, and simultaneous inference requires an additional bias-control step on top of the targeted fit.

6.8 Discussion

This chapter develops a smooth density-based approach to bivariate failure-time data under right censoring and studies how estimation, targeting, and inference interact in a setting where both the statistical geometry and the computation are nontrivial. Its main conceptual contribution is a regularized TMLE strategy for bivariate right-censored survival that operationalizes the Chapter 3 projection-based targeting logic and the Chapter 5 jackknife inference logic in a setting where a literal classical full TMLE is not computationally tractable.

The main empirical message is straightforward. EM-HAL provides a useful initial non-parametric representation of the joint density, but untargeted plug-in survival estimation is not reliable enough for inference on the survival functional. Survival-specific targeting is therefore not optional in this setting: it is the step that materially improves survival estimation and makes jackknife-based inference plausible. Table 6.9 summarizes the cross-method comparison.

The contrast between self-consistency targeting and the projection-based reg-ATMLE update is clearest when the simulation evidence is viewed as a whole. Self-consistency tar-

Table 6.9: Head-to-head comparison of self-consistency targeting and reg-ATMLE ($R = 213$). Pointwise metrics are averaged over the 209×209 evaluation grid; simultaneous metrics are computed on the 9×9 targeting grid.

Aspect	Self-consistency	
	targeting	reg-ATMLE
Survival MSE (targeted)	0.000463	0.000389
Oracle coverage (targeted)	92.5%	93.1%
MSE change after targeting	+10.1%	−17.2%
Grid pts with MSE improvement	29.2%	67.7%
Jackknife pointwise coverage	95.1%	94.7%
Jackknife CI width	0.090	0.083
Simultaneous coverage (base)	89.7%	92.5%
Simultaneous coverage (BC)	95.3%	97.7%

geting is the more aggressive centering correction: it reduces survival bias and improves coverage, but it does so through a rougher update that pays in variance and local instability. The projection-based reg-ATMLE update is smoother because it stays within the HAL score geometry; in this Gaussian-copula design, that translated into lower targeted MSE, slightly better oracle and simultaneous coverage, and narrower jackknife intervals. The practical conclusion is comparative and design-specific: both strategies show that targeting helps, while the projection-based reg-ATMLE update provides the better overall bias-variance-inference tradeoff in this study.

The inferential lesson is equally important. The simulation does not show that grouped V -fold jackknife is automatically valid for this problem. Rather, it shows that the jackknife becomes effective only after targeting has removed enough first-order bias from the survival estimator. Without targeting, jackknife coverage collapses to about 43% for both strategies. With targeting, pointwise jackknife coverage returns to the nominal range. By contrast, the delta-method intervals remain overly conservative because of information-matrix instability, and the projected-IC variance formula remains too optimistic in finite samples.

The practical workflow suggested by these results is straightforward for the setting studied here: use EM-HAL to obtain the initial density estimate, apply survival-specific targeting, and then use grouped V -fold jackknife for uncertainty quantification. In our simulations, the projection-based reg-ATMLE update is the better default within that workflow, though that recommendation should be read relative to the targeting grid, sample size, and data-generating mechanism considered here. The common 9×9 targeting grid is best viewed as a computational precision choice adopted for feasibility and comparability, not as a canonical dimension. For simultaneous bands, the truth-free jackknife bias correction remained important: targeting made simultaneous inference feasible, but multiplicity still required a further bias-control step.

Limitations. The study considered one Gaussian-copula data-generating mechanism and one sample size, so the ranking of the methods should be checked under alternative dependence structures, censoring levels, and sample-size regimes. The upper-right tail remains difficult for both procedures. The computational burden is substantial because valid inference requires refitting EM-HAL and targeting inside each jackknife training sample; those costs are manageable in two dimensions with parallelization but will be more serious in higher dimensions. A chapter-local open problem is how large the targeting grid should be and how it should be selected: for self-consistency targeting, larger grids more closely approximate the full self-consistency system; for the projection-based reg-ATMLE update, they more directly approximate the survival surface. Data-adaptive grid selection and principled growth rules for m_n are therefore methodological questions in their own right at the chapter level.

Overall, this chapter treats smooth density estimation, survival-specific targeting, and resampling-based inference as distinct but coordinated parts of the bivariate survival problem. EM-HAL provides the initial representation, self-consistency targeting offers an aggressive centering correction retained in the appendix as a comparator, the projection-based reg-ATMLE update inherited from Chapter 3 provides the smoother main-line alternative with better overall stability, and the V -fold jackknife of Chapter 5 supplies practical uncertainty quantification once targeting has done the work needed to support it. Together these components form a coherent nonparametric alternative to classical bivariate survival estimators under right-censoring and show why the estimation and inference sides of the regularized targeted-learning program belong together.

This chapter shows how the estimator-side and inferential ideas developed earlier combine in the setting where exact semiparametric constructions are hardest to deploy directly. The concluding chapter draws together the main methodological lessons of the dissertation, clarifies the limitations that remain, and outlines the future directions suggested by these results.

Chapter 7

Conclusion and Future Directions

7.1 Principal conclusions

The dissertation has studied regularized targeted learning under analytic intractability, that is, settings in which the formal semiparametric ideal is clear but its direct implementation is unstable, analytically difficult, or computationally impractical. The main conclusion is that targeted learning remains viable in such settings when the estimator and the inferential procedure are both organized around the same first-order structure. On the estimation side, regularizing the targeting step inside a data-adaptive working model can preserve the relevant bias correction while stabilizing computation. On the inference side, grouped V-fold jackknife methods can provide reliable uncertainty quantification without depending on a fragile closed-form variance calculation or bootstrap-scale recomputation.

Three conclusions follow from the chapters taken together. First, the working model is a substantive part of the method rather than a hidden implementation detail, because it determines which score directions are treated as bias-relevant and therefore how regularization should be organized. Second, projection-guided approximations to influence-function directions provide a principled route to targeted estimation when exact updates or least favorable paths are not available in a tractable operational form. Third, inferential procedures should be designed under the same computational constraints as the estimator itself, which is why grouped V-fold jackknife methods play such a central role in the overall program developed here.

The empirical studies sharpen these conclusions in practice. Across the HAL, neural-network, and bivariate-survival settings, the dissertation shows that unregularized targeted updates and plug-in variance calculations can become numerically unstable or markedly anti-conservative in finite samples when score systems are collinear, overlap is weak, or the target is functional. By contrast, the projection-guided estimation and grouped V-fold jackknife pipeline improves stability, reduces mean squared error, and moves inference substantially closer to nominal coverage in precisely the settings that place the greatest pressure on classical implementations.

7.2 A unified methodological perspective

The methodological chapters are most naturally read as successive parts of one argument. Chapter 3 establishes the regularized targeting mechanism in the setting where its geometry is easiest to describe, namely HAL working models with explicit basis functions and nested score spaces. Chapter 4 then shows that the same projection logic is not confined to HAL, but extends to overparameterized neural-network working models once the relevant score space is identified through selected gradients. Taken together, these chapters show that the core estimator-side idea concerns score-space organization and bias-relevant regularization rather than any single model class.

Chapter 5 supplies the inferential counterpart by showing how uncertainty quantification can be built from the same first-order structure without relying on fragile closed-form variance calculations. Chapter 6 brings the estimator-side and inferential developments together in the setting where the pressure on the theory is greatest: coarsened-data survival with a functional target, an expensive estimation pipeline, and no operational classical full TMLE for the entire survival surface. Across these chapters, the dissertation repeatedly replaces exact semiparametric constructions that are unstable or intractable in practice by valid first-order surrogates inside adaptive working models, and then pairs the resulting estimators with an inference procedure designed for the same computational setting.

7.3 Limitations

Several limitations remain. The first is theoretical. Parts of the analysis still rely on Donsker-type control or related arguments that would ideally be replaced by a more systematic cross-fitted treatment. A fully developed cross-fitted version of the regularized targeting procedures in Chapters 3, 4, and 6 would sharpen the point-estimation theory and clarify its interaction with grouped V-fold jackknife inference.

The second limitation concerns the selection of regularization itself. Projection penalties, working-model size, and stopping rules can materially affect finite-sample behavior, yet they are not determined by a single theory-driven principle across all chapters. The plateau selector studied in Chapter 3 is an important step, but a broader theory connecting regularization choice to inferential performance remains to be developed.

The third limitation is dimensional. Functional targets require simultaneous control of many first-order equations, and both the computational and inferential problems become more difficult as the target dimension grows. This issue appears already in Chapters 4 and 6, and it is likely to become more pronounced in richer coarsened-data and event-time settings.

7.4 Future directions

Offset-based regularized TMLE with grouped V-fold jackknife inference. The most immediate methodological direction is to formalize a broader offset-based regularized

TMLE program suggested by the dissertation. The completed chapters are written primarily in regularized A-TMLE language because their formal targeting steps are defined within adaptive working models. At the same time, the dissertation has already implemented external-offset ideas in practice: Chapter 4 treats the trained neural network as an offset, and Chapter 6 targets a fixed EM-HAL initial density estimator. The genuine future direction is therefore not the introduction of the idea itself, but the development of a unified general theory for this broader construction. A canonical example would write

$$Q_\beta(x) = \hat{Q}_{n,\text{init}}(x) + \sum_j \beta_j \phi_j(x),$$

with $\hat{Q}_{n,\text{init}}$ supplied by a strong external learner and the coefficients β used only for the targeting correction. The super-learner-plus-offset construction should be read as one motivating instance of the broader theory rather than as a new algorithmic prescription introduced at the end of the dissertation.

Seen this way, projection-guided targeting becomes a regularized analogue of the classical TMLE targeting step. Rather than fluctuating along an exactly tractable least favorable path that solves the efficient-influence-curve equation directly, one projects a valid influence-curve representation onto a chosen score space and solves the corresponding regularized approximation. The working model therefore serves as the geometric space in which the correction is constructed, not necessarily as the model that generates the entire initial estimator. This is the sense in which the future program moves beyond the specific A-TMLE restriction that the initial fit remain within the working model while still preserving the targeted-learning logic developed in Chapter 3.

The natural inferential partner for this program is the grouped V-fold jackknife of Chapter 5. It uses the same first-order structure as the targeting step, avoids dependence on a tractable closed-form variance formula, and remains feasible for expensive pipelines because it requires only V grouped refits rather than heavy bootstrap recomputation. Chapter 3 supplies the regularized targeting logic, Chapter 4 shows that the targeting principle is not tied to HAL alone, Chapter 5 provides the computationally feasible inferential layer, and Chapter 6 shows that the combination is useful when exact efficient-influence-curve and least-favorable-path machinery is difficult to deploy directly. Formalizing this offset-based regularized TMLE program, and establishing general conditions under which the combined point-estimation and inference pipeline is valid, is the most immediate next step suggested by the dissertation.

Cross-fitting. Parts of the analysis still rely on Donsker-type control that would ideally be replaced by a fully cross-fitted construction. Developing a cross-fitted version of the regularized targeting procedures in Chapters 3, 4, and 6 would sharpen the point-estimation theory and clarify its interaction with the fold-based structure of Chapter 5, whose fold geometry already offers a natural cross-fitting analogue.

Adaptive regularization selection. Projection penalties, working-model size, and stopping rules each matter, and they currently rely on pilot experimentation or external tuning rather than on a unified rule. A theory of adaptive regularization selection that is coordinated with inferential validity—an analogue of the plateau selector of Chapter 3 tied directly to grouped V-fold jackknife coverage—is a concrete next target.

Growing target dimension and richer coarsening. Functional parameters require joint control of many first-order equations, and both computational and inferential difficulty grow with the target dimension. Chapter 4 and Chapter 6 already show this pressure. Extending the grouped V-fold jackknife to settings in which the evaluation grid grows with the sample size, and carrying the bivariate-survival program into higher-dimensional event-time problems, competing-risks structures, or longitudinal treatment-effect settings, are natural follow-ups.

Practical guidance. The chapters also suggest a few practical defaults for applied use. When grouped V-fold jackknife inference is used as the primary uncertainty layer, a moderate fixed choice such as $V = 20$ is a reasonable default because it remains computationally feasible while performing well across the simulation settings studied here. HAL working models are especially attractive when the score space is structured and interpretable enough that one wants direct control over the geometry of the targeting step, whereas neural-network working models are more natural when representation richness matters more than explicit basis interpretation and the targeting correction must be regularized through score-space selection or penalization.

7.5 Concluding remarks

The contribution of this dissertation is therefore not only a collection of estimators, algorithms, and asymptotic results, but a methodological perspective on targeted learning under instability, analytic intractability, and computational cost. The chapters developed here show that semiparametric rigor and computational pragmatism need not pull in opposite directions: the relevant first-order structure can be preserved even when exact updates, exact least favorable paths, or closed-form variance formulas are difficult to use directly. That perspective proves productive in HAL working models, in neural-network working models, in grouped V-fold jackknife inference, and in bivariate right-censored survival under coarsened data.

More broadly, the dissertation lowers the barrier to robust semiparametric inference in modern applied work. By replacing fragile derivations and unstable matrix inversions with stable projection-based computation and grouped V-fold jackknife inference, it creates a more accessible path for applied statisticians to combine modern machine learning with principled semiparametric reasoning. That broader accessibility is part of the dissertation’s practical impact and a natural foundation for the work that follows.

Bibliography

- James N. Arvesen. Jackknifing u -statistics. *The Annals of Mathematical Statistics*, 40(6): 2076–2100, 1969. doi: 10.1214/aoms/1177697287.
- Susan Athey, Guido W Imbens, and Stefan Wager. Approximate residual balancing: debiased inference of average treatment effects in high dimensions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 80(4):597–623, 2018.
- Rina Foygel Barber, Emmanuel J. Candès, Aaditya Ramdas, and Ryan J. Tibshirani. Predictive inference with the jackknife+. *The Annals of Statistics*, 49(1):486–507, 2021. doi: 10.1214/20-AOS1965.
- Alexandre Belloni, Victor Chernozhukov, and Christian Hansen. Inference on treatment effects after selection among high-dimensional controls. *The Review of Economic Studies*, 81(2):608–650, 2014.
- D. Benkeser and M. J. van der Laan. The highly adaptive lasso estimator. In *Proceedings of the IEEE International Conference on Data Science and Advanced Analytics*, pages 689–696, 2016.
- David C Benkeser and Nima S Hejazi. *survtmle: Targeted Minimum Loss-Based Estimation for Survival Analysis in R*, 2017. URL <http://dx.doi.org/10.5281/zenodo.835868>.
- A. F. Bibaut and M. J. van der Laan. Fast rates for empirical risk minimization over càdlàg functions with bounded sectional variation norm. *arXiv preprint arXiv:1907.09244*, 2019.
- Peter J Bickel, Chris AJ Klaassen, Ya’acov Ritov, and Jon A Wellner. *Efficient and Adaptive Estimation for Semiparametric Models*. Johns Hopkins University Press, 1993.
- David R. Brillinger. The asymptotic behaviour of Tukey’s general method of setting approximate confidence limits (the jackknife) when applied to maximum likelihood estimates. *Review of the International Statistical Institute*, 32(3):202–206, 1964.
- Weixin Cai and Mark van der Laan. Nonparametric bootstrap inference for the targeted highly adaptive lasso estimator. *arXiv preprint arXiv:1905.10299*, 2019.
- Weixin Cai and Mark J. van der Laan. One-step targeted maximum likelihood estimation for time-to-event outcomes. *Biometrics*, 76(3):722–733, 2020. doi: 10.1111/biom.13172.

- RJ Carroll, Suojin Wang, DG Simpson, AJ Stromberg, and D Ruppert. The sandwich (robust covariance matrix) estimator. *Unpublished manuscript*, 1998.
- Xiaohui Chen and Kengo Kato. Jackknife multiplier bootstrap: Finite-sample approximations to the U -process supremum with applications. *Probability Theory and Related Fields*, 176(3–4):1097–1163, 2020. doi: 10.1007/s00440-019-00936-y.
- Victor Chernozhukov, Denis Chetverikov, and Kengo Kato. Gaussian approximations and multiplier bootstrap for maxima of sums of high-dimensional random vectors. *The Annals of Statistics*, 41(6):2786–2819, 2013. doi: 10.1214/13-AOS1161.
- Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 2018.
- Victor Chernozhukov, Whitney K. Newey, Victor Quintas-Martinez, and Vasilis Syrgkanis. RieszNet and ForestRiesz: Automatic Debiased Machine Learning with Neural Nets and Random Forests, June 2022a. URL <http://arxiv.org/abs/2110.03031>. arXiv:2110.03031 [cs, econ, stat].
- Victor Chernozhukov, Whitney K Newey, and Rahul Singh. Automatic debiased machine learning of causal and structural effects. *Econometrica*, 90(3):967–1027, 2022b.
- D. G. Clayton. A model for association in bivariate life tables and its application in epidemiological studies of familial tendency in chronic disease incidence. *Biometrika*, 65(1):141–151, 1978.
- D. M. Dabrowska. Kaplan–Meier estimate on the plane. *Annals of Statistics*, 16(4):1475–1489, 1988.
- A. C. Davison and D. V. Hinkley. *Bootstrap Methods and their Application*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, 1997.
- A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B*, 39(1):1–38, 1977.
- Iván Díaz and Mark J. van der Laan. Assessing the causal effect of policies: An example using stochastic interventions. *The International Journal of Biostatistics*, 9(2):161–174, 2013. doi: 10.1515/ijb-2013-0014.
- A. Ertefaie, N. S. Hejazi, and M. J. van der Laan. Nonparametric inverse-probability-weighted estimators based on the highly adaptive lasso. *Biometrics*, 79(2):1029–1041, 2023.
- C. Genest, K. Ghoudi, and L.-P. Rivest. A semiparametric estimation procedure of dependence parameters in multivariate families of distributions. *Biometrika*, 82(3):543–552, 1995.

- R. D. Gill, M. J. van der Laan, and J. M. Robins. Coarsening at random: Characterizations, conjectures, counter-examples. In *Proceedings of the First Seattle Symposium in Biostatistics*, pages 255–294. Springer, 1997.
- P. Groeneboom and J. A. Wellner. *Information Bounds and Nonparametric Maximum Likelihood Estimation*. DMV Seminar Band 19. Birkhäuser, 1992.
- Susan Gruber and Mark J. van der Laan. tmle: An R package for targeted maximum likelihood estimation. *Journal of Statistical Software*, 51(13):1–35, 2012. doi: 10.18637/jss.v051.i13.
- Tobias Hatt, Jeroen Berrevoets, Alicia Curth, Stefan Feuerriegel, and Mihaela van der Schaar. Combining observational and randomized data for estimating heterogeneous treatment effects. *arXiv preprint arXiv:2202.12891*, 2022.
- Jennifer L Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240, 2011.
- Yilong Hou, Zhengpu Zhao, Yi Li, and Mark van der Laan. Hal-mle log-splines density estimation (part i: Univariate). *arXiv preprint arXiv:2602.16259*, 2026.
- Fredrik D Johansson, Uri Shalit, and David Sontag. Learning representations for counterfactual inference. In *Proceedings of the 33rd International Conference on Machine Learning*, pages 3020–3029, 2016.
- Cheng Ju, Susan Gruber, Samuel D Lendle, Antoine Chambaz, Jessica M Franklin, Richard Wyss, Sebastian Schneeweiss, and Mark J van Der Laan. Scalable collaborative targeted learning for high-dimensional data. *Statistical methods in medical research*, 28(2):532–554, 2019.
- Vladimir Koltchinskii and Karim Lounici. Concentration inequalities and moment bounds for sample covariance operators. *Bernoulli*, 23(1):110–133, 2017. doi: 10.3150/15-BEJ730.
- Phillip S. Kott. Using the delete-a-group jackknife variance estimator in practice. In *Proceedings of the Section on Survey Research Methods*, pages 763–768, Alexandria, VA, 1998. American Statistical Association.
- Phillip S. Kott. The delete-a-group jackknife. *Journal of Official Statistics*, 17(4):521–526, 2001.
- Phillip S. Kott and Steven T. Garren. Evaluating the small-sample bias of the delete-a-group jackknife for model analyses. *Journal of Official Statistics*, 27(1):121–134, 2011.
- Samuel Kotz and Saralees Nadarajah. *Multivariate t Distributions and Their Applications*. Cambridge University Press, Cambridge, 2004. ISBN 978-0-521-82654-9.

- Sören R. Künnel, Jasjeet S. Sekhon, Peter J. Bickel, and Bin Yu. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences of the United States of America*, 116(10):4156–4165, March 2019. ISSN 1091-6490. doi: 10.1073/pnas.1804597116.
- Henry Lam. A cheap bootstrap method for fast inference. 2022. arXiv preprint arXiv:2202.00090.
- Gen Li and Jie Ding. The rate of convergence of variation-constrained deep neural networks. *arXiv preprint arXiv:2106.12068*, 2022.
- Y. Li, D. McCoy, N. Gunter, K. Lee, A. Schuler, and M. J. van der Laan. Targeted deep architectures: A TMLE-based framework for robust causal inference in neural networks. *arXiv preprint arXiv:2507.12435*, 2025a.
- Yi Li, Sky Qiu, Zeyi Wang, and Mark J. van der Laan. Regularized targeted maximum likelihood estimation in highly adaptive lasso implied working models. *arXiv preprint arXiv:2506.17214*, 2025b.
- D. Y. Lin and Z. Ying. A simple nonparametric estimator of the bivariate survival function under univariate censoring. *Biometrika*, 80(3):573–581, 1993.
- Rupert G. Miller. The jackknife—a review. *Biometrika*, 61(1):1–15, 1974. doi: 10.1093/biomet/61.1.1.
- José Luis Montiel Olea and Mikkel Plagborg-Møller. Simultaneous confidence bands: Theory, implementation, and an application to SVARs. *Journal of Applied Econometrics*, 34(1):1–17, 2019. doi: 10.1002/jae.2656.
- Anders Munch, Thomas A Gerds, Mark J van der Laan, and Helene CW Rytgaard. Estimating conditional hazard functions and densities with the highly-adaptive lasso. *arXiv preprint arXiv:2404.11083*, 2024.
- Lizhen Nie, Mao Ye, Qiang Liu, and Dan Nicolae. VCNet and functional targeted regularization for learning causal effects of continuous treatments. In *International Conference on Learning Representations*, 2021.
- D. Oakes. Bivariate survival models induced by frailties. *Journal of the American Statistical Association*, 84(406):487–493, 1989.
- Johan Sebastian Ohlendorff, Anders Munch, Kathrine Kold Sørensen, and Thomas Alexander Gerds. Cheap subsampling bootstrap confidence intervals for fast and robust inference. 2025. arXiv preprint arXiv:2501.10289.
- Maya L Petersen, Kristin E Porter, Susan Gruber, Yue Wang, and Mark J van der Laan. Diagnosing and responding to violations in the positivity assumption. *Statistical Methods in Medical Research*, 21(1):31–54, 2012.

- R. L. Prentice and J. Cai. Covariance and survivor function estimation using censored multivariate failure time data. *Biometrika*, 79(3):495–512, 1992.
- R. L. Prentice, F. Moodie, and J. Wu. Hazard-based nonparametric survivor function estimation. *Journal of the Royal Statistical Society: Series B*, 66(2):305–319, 2004.
- M. H. Quenouille. Approximate tests of correlation in time-series. *Journal of the Royal Statistical Society, Series B*, 11(1):68–84, 1949.
- Helene CW Rytgaard and Mark J van der Laan. One-step targeted maximum likelihood estimation for targeting cause-specific absolute risks and survival curves. *Biometrika*, 111(1):129–145, 2024.
- Uri Shalit, Fredrik D Johansson, and David Sontag. Estimating individual treatment effect: generalization bounds and algorithms. In *International conference on machine learning*, pages 3076–3085. PMLR, 2017.
- Jun Shao and C. F. Jeff Wu. A general theory for jackknife variance estimation. *The Annals of Statistics*, 17(3):1176–1197, 1989. doi: 10.1214/aos/1176347263.
- Claudia Shi, David Blei, and Victor Veitch. Adapting neural networks for the estimation of treatment effects. *Advances in neural information processing systems*, 32, 2019.
- Junming Shi, Wenxin Zhang, Alan E. Hubbard, and Mark J. van der Laan. HAL-based plug-in estimation with pointwise asymptotic normality of the causal dose-response curve. *arXiv preprint arXiv:2406.05607*, 2024.
- J. H. Shih and T. A. Louis. Inferences on the association parameter in copula models for bivariate survival data. *Biometrics*, 51(4):1384–1399, 1995.
- Toru Shirakawa, Yi Li, Yulun Wu, Sky Qiu, Yuxuan Li, Mingduo Zhao, Hiroyasu Iso, and Mark Van Der Laan. Longitudinal targeted minimum loss-based estimation with temporal-difference heterogeneous transformer. In *Proceedings of the 41st International Conference on Machine Learning*, pages 45097–45113. PMLR, 2024.
- Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 58(1):267–288, 1996.
- Linh Tran, Maya Petersen, Joshua Schwab, and Mark J. van der Laan. Robust variance estimation and inference for causal effect estimation. *Journal of Causal Inference*, 11(1): 20210067, 2023. doi: 10.1515/jci-2021-0067.
- W.-Y. Tsai, S. Leurgans, and J. Crowley. Nonparametric estimation of a bivariate survival function in the presence of censoring. *Annals of Statistics*, 14(4):1351–1365, 1986.
- Anastasios A Tsiatis. *Semiparametric theory and missing data*, volume 4. Springer, 2006.

- John W. Tukey. Bias and confidence in not-quite large samples. In Jerzy Neyman, editor, *Proceedings of the Symposium on Mathematical Statistics and Probability*, volume 1, page 614. University of California Press, Berkeley, 1958.
- B. W. Turnbull. The empirical distribution function with arbitrarily grouped, censored and truncated data. *Journal of the Royal Statistical Society: Series B*, 38(3):290–295, 1976.
- Lars van der Laan, Marco Carone, Alex Luedtke, and Mark van der Laan. Adaptive de-biased machine learning using data-driven model selection techniques. *arXiv preprint arXiv:2307.12544*, 2023.
- Lars van der Laan, Aurelien Bibaut, Nathan Kallus, and Alex Luedtke. Automatic debiased machine learning for smooth functionals of nonparametric m-estimands. *arXiv preprint arXiv:2501.11868*, 2025.
- M. J. van der Laan. Efficient estimation in the bivariate censoring model and repairing NPMLE. *Annals of Statistics*, 24(2):596–627, 1996.
- M. J. van der Laan. A generally efficient targeted minimum loss based estimator based on the highly adaptive lasso. *The International Journal of Biostatistics*, 13(2), 2017.
- M. J. van der Laan. Higher order spline highly adaptive lasso estimators of functional parameters: Pointwise asymptotic normality and uniform convergence rates. *arXiv preprint arXiv:2301.13354*, 2023.
- M. J. van der Laan and J. M. Robins. *Unified Methods for Censored Longitudinal Data and Causality*. Springer, 2003.
- Mark van der Laan and Susan Gruber. One-step targeted minimum loss-based estimation based on universal least favorable one-dimensional submodels. *The international journal of biostatistics*, 12(1):351–378, 2016.
- Mark van der Laan, Zeyi Wang, and Lars van der Laan. Higher order targeted maximum likelihood estimation. *arXiv preprint arXiv:2101.06290*, 2021.
- Mark van der Laan, Sky Qiu, Jens Magelund Tarp, and Lars van der Laan. Adaptive-tmle for the average treatment effect based on randomized controlled trial augmented with real-world data. *Journal of Causal Inference*, 14(1):20240025, 2026.
- Mark J. van der Laan. Simultaneous inference with the Kaplan–Meier estimator of survival. <https://vanderlaan-lab.org/2019/12/29/simultaneous-inference-with-the-kaplan-meier-estimator-of-survival/>, 2019. Blog post.
- Mark J van der Laan and Susan Gruber. Collaborative double robust targeted maximum likelihood estimation. *The International Journal of Biostatistics*, 6(1), 2010.

- Mark J van der Laan and Sherri Rose. *Targeted learning in data science: causal inference for complex longitudinal studies*. Springer International Publishing, 2018.
- Mark J van der Laan and Daniel Rubin. Targeted maximum likelihood learning. *The International Journal of Biostatistics*, 2(1), 2006.
- Mark J Van der Laan, Eric C Polley, and Alan E Hubbard. Super learner. *Statistical applications in genetics and molecular biology*, 6(1), 2007.
- Mark J Van der Laan, Sherri Rose, et al. *Targeted learning: causal inference for observational and experimental data*, volume 4. Springer, 2011.
- A. W. van der Vaart. *Asymptotic Statistics*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, 1998.
- Aad W Van Der Vaart and Jon A Wellner. *Weak Convergence and Empirical Processes: With Applications to Statistics*. Springer, 1996.
- Roman Vershynin. How close is the sample covariance matrix to the actual one? *Journal of Theoretical Probability*, 25(3):655–686, 2012. doi: 10.1007/s10959-010-0338-z.
- G. C. G. Wei and M. A. Tanner. A Monte Carlo implementation of the EM algorithm and the poor man’s data augmentation algorithms. *Journal of the American Statistical Association*, 85(411):699–704, 1990.
- Halbert White. Maximum likelihood estimation of misspecified models. *Econometrica*, 50(1):1–25, 1982.
- C. F. J. Wu. On the convergence properties of the EM algorithm. *The Annals of Statistics*, 11(1):95–103, 1983.
- C. F. Jeff Wu. Jackknife, bootstrap and other resampling methods in regression analysis. *The Annals of Statistics*, 14(4):1261–1295, 1986. doi: 10.1214/aos/1176350142.
- Wenjing Zheng and Mark J van der Laan. Cross-validated targeted minimum-loss-based estimation. In *Targeted Learning: Causal Inference for Observational and Experimental Data*, pages 459–474. Springer, 2011.
- Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 67(2):301–320, 2005.

Appendix A

Supporting Material for Chapter 3

This appendix collects the algorithmic details, extended empirical specifications, and formal proofs for Chapter 3. The main chapter keeps the estimator construction, theorem statements, and headline applications in the body; longer derivations and proofs are recorded here.

A.1 regHAL-ATMLE Algorithm for Multi Dimensional Target Parameter

Algorithm 11 HAL Updating Algorithm Based on Approximated EIC (Multi Dimensional Target Parameter)

- 1: **Target Parameter Indexed By** s Ψ_s for $s \in S$.
- 2: **Cross-Validated HAL Fit:** Compute cross validated $\hat{Q}_n$.
- 3: **Working Model Definition:** Define the HAL working model $\mathcal{Q}_\beta$ (which is not necessarily the working model of the CV fit, e.g., it can be undersmoothed).
- 4: **Parameter Initialization:** Initialize $\beta^{(0)}$ based on $\hat{Q}_n$, incorporating zero-padding if using an undersmoothed HAL model.
- 5: **if** using the projection-based method **then**
- 6: **Initial Gradient and Nuisances:** Specify the form of the initial non-parametric gradient $D(O)$ and fit any extra nuisance parameters required.
- 7: **end if**
- 8: **EIC Computation:** Dynamically derive the approximated regularized EIC $\tilde{D}_n^*$ for Ψ_s as

$$\tilde{D}_{s,n,\beta^{(k)}}^*(O_i) = (\hat{\alpha}_{s,n}^{(k)})^\top \text{scores}_{\beta^{(k)}}(O_i)$$

at the current $\beta^{(k)}$ (re-evaluating $D(O)$ and associated components using updated $Q_{\beta^{(k)}}$ residuals), where scores_β are the negative gradients of the loss function L_β w.r.t. β .

- 9: **Universal Least Favorable Path Update Direction:**

$$\begin{aligned} d(s) &= P_n \tilde{D}_{s,n,\beta^{(k)}}^* \quad \forall s \in S \\ d &= (d(1), d(2), \dots, d(|S|)) \\ \text{Update}^{(k)} &= \left(\sum_{s \in S} \frac{d(s)}{\|d\|_2} \hat{\alpha}_{s,n}^{(k)} \right) \end{aligned}$$

- 10: **Parameter Update:** Update $\beta^{(k)}$ to $\beta^{(k+1)}$ using gradient descent:

$$\beta^{(k+1)} = \beta^{(k)} + \epsilon \cdot \text{Update}^{(k)}$$

where ϵ is a small step size.

- 11: **Stopping Criterion:** Repeat steps 6–8 until

$$\frac{\|d\|_2}{\sqrt{|S|}} < \frac{\text{Median}_s \{se(\tilde{D}_{s,n,\beta^{(k+1)}}^*)\}}{\sqrt{n} \log n}$$

A.2 regHAL-ATMLE for ATE -Details

We discuss such regularized targeting in estimation of the average treatment effect parameter through regHAL-ATMLE.

A.2.1 Estimating the ATE using one-step regHAL-ATMLE

A.2.1.1 Projection Based Parametric EIC Approximation

Let's first choose an initial non-parametric gradient $D(O)$ that we are going to project. As strictly recommended in Section 3.2.2, we use the centered gradient $D = \left(\frac{A}{\hat{g}_n(1|W)} - \frac{1-A}{\hat{g}_n(0|W)} \right) * (Y - \hat{Q}_n(A, W))$.

Now, we can do a one step TMLE update in the parametric working model using the approximated non-parametric EIC based on projection. At each Q_β , we have

$$\tilde{D}_{n,\beta,proj}^*(O_i) = \hat{\gamma}_{0,n}^{\beta,reg} + (\hat{\alpha}_n^{\beta,reg})^\top (S_{i,*}^\beta)^\top$$

where $\hat{\alpha}_n^{\beta,reg}$ and $\hat{\gamma}_{0,n}^{\beta,reg}$ are the coefficients and unpenalized intercept gained by running a regularized least square regression of $D(O_i)$ on $S_{i,1}^\beta, S_{i,2}^\beta, \dots, S_{i,p}^\beta$. Here, we choose a weakly penalized lasso problem (e.g., using a very small penalization factor such as 10^{-5}) as a numerical floor to maintain stability.

A.2.1.2 Hessian-Based Parametric EIC Approximation

In the Hessian-based approach, we first need to compute:

$$\begin{aligned} \frac{\partial \Psi^F(Q_\beta)}{\partial \beta_j} &= \frac{\partial [E_0(Q_\beta(1, W)) - E_0(Q_\beta(0, W))]}{\partial \beta_j} \\ &= P_0 \left(\phi_j(1, W) - \phi_j(0, W) \right) \end{aligned}$$

We can approximate it by

$$P_n \left(\phi_j(1, W) - \phi_j(0, W) \right)$$

In addition, we use the empirical expected Hessian $I_{P_n}(\beta) = P_n [\nabla^2 L_\beta(O)]$, which for a linear model with squared error loss $L_\beta(O) = (Y - Q_\beta(A, W))^2$ mathematically coincides with $2P_n [\phi(A, W)\phi(A, W)^\top]$. Adding a ridge parameter η_{weak} to stabilize the inversion, we have:

$$\tilde{D}_{n,\beta,hessian}^*(O_i) = \left(\left(P_n(\phi_j(1, W) - \phi_j(0, W)) \right)_{j=1}^p \right)^\top * \left(I_{P_n}(\beta) + \eta_{\text{weak}} * I \right)^{-1} (S_{i,*}^\beta)^\top$$

A.2.1.3 Targeting

Choosing $\tilde{D}_{n,\beta}^* = \tilde{D}_{n,\beta,proj}^*$ or $\tilde{D}_{n,\beta,hessian}^*$.

We have, for a chosen small step size ϵ , starting from $\beta^{(0)} = \beta^{initial}$, we update it as $\beta^{(k+1)} = \beta^{(k)} + \epsilon \cdot \text{sign}(P_n \tilde{D}_{n,\beta^{(k)}}^* \hat{\alpha}_n^{\beta^{(k)}})$ until $|P_n \tilde{D}_{n,\beta^{(k+1)}}^*|$ is smaller than the desired threshold δ_n . Note that each time, we need to recalculate the score matrix $S_{i,j}^{\beta^{(k+1)}}$ since it is involved in the $\tilde{D}_{n,\beta^{(k+1)}}^*(O_i)$.

Suppose we stop at $\beta^{(K)}$, then our final point estimate for the ATE is $\hat{\Psi}^F(Q_{\beta^{(K)}}) = \frac{1}{n} \sum_{i=1}^n (Q_{\beta^{(K)}}(1, W_i) - Q_{\beta^{(K)}}(0, W_i))$.

Moreover, since we have the approximated non-parametric EIC, we can estimate the standard error of this estimator and construct a Wald-type 95% confidence interval as:

$$\hat{\Psi}^F(Q_{\beta^{(K)}}) \pm 1.96 * \sqrt{\frac{P_n[\tilde{D}_{n,\beta^{(K)}}^*(O)]^2}{n}} \quad (\text{A.1})$$

A.2.2 Simulation Set Up for ATE

A.2.2.1 HAL basis choice

We use the zero-order HAL basis up to two-way interactions. In addition, the knot points are selected to be the $\frac{n}{20}$ number of percentiles for each continuous predictor.

A.2.2.2 DGP 1

First simulation is the regular setting with no positivity violations for the average treatment effect.

Consider a point-treatment causal inference setting where researchers collect three baseline covariates W_1 , W_2 , and W_3 , each following $\text{Uniform}(-1, 1)$. The treatment assignment A is drawn from a Bernoulli distribution with probability $\text{expit}(-0.25W_1 + 0.7W_2)$. The outcome is given by $Y = 1.9 + 1.5A + (2.5W_1 + 0.7W_2)A + 1.5\sin(W_1 + W_2) + 0.3|W_1| + 0.9W_1^2 + 1.4W_2 + 2.1W_3 + U_Y$, where U_Y follows a standard normal distribution.

A.2.2.3 DGP 2

The second simulation is about ATE estimation under practical positivity violations.

We consider a point-treatment causal inference setting with three baseline covariates W_1 , W_2 , and W_3 . These covariates are generated as

$$W_i \sim \text{Uniform}(-1, 1) \quad (\text{rounded to one decimal})$$

for $i = 1, 2, 3$.

The treatment assignment A is generated as follows. If no counterfactual treatment is specified, the probability of treatment is given by

$$\mathbb{P}(A = 1 \mid W_1, W_2) = \text{expit}(-0.25 W_1 + 5 W_2) = \frac{\exp(-0.25 W_1 + 5 W_2)}{1 + \exp(-0.25 W_1 + 5 W_2)}.$$

Otherwise, if a counterfactual treatment a is fixed, we set

$$A = a \quad \text{for all observations.}$$

The outcome Y is generated according to the model

$$Y = 1.9 + 1.5 A + (2.5 W_1 + 0.7 W_2)A + 1.5 \sin(W_1 + W_2) + 0.3 |W_1| + 0.9 W_1^2 + 1.4 W_2 + 2.1 W_3 + U_Y,$$

where the error term is distributed as

$$U_Y \sim \mathcal{N}(0, 1).$$

The true average treatment effect (ATE) is then estimated by computing

$$\text{ATE} = \mathbb{E}[Y(1) - Y(0)]$$

using a large sample.

A.2.3 Standard TMLE with Propensity Score Truncation

As a benchmark, we implemented the standard TMLE update. We use the initial HAL Q fit as a fixed offset and fit a linear regression with the clever covariate $H(A, W) = \frac{A}{g(W)} - \frac{1-A}{1-g(W)}$ as the sole covariate, yielding the updated $Q^*(A, W) = \hat{Q}(A, W) + \hat{\epsilon} \cdot H(A, W)$. Propensity scores are truncated to $[0.01, 0.99]$.

The standard TMLE results are included directly in the main text tables (Tables 3.1–3.4) as the “Standard” targeting row.

A.3 regHAL-ATMLE for Survival Curve - Details

In survival set ups, one usually cares about targeting multiple time points survivals and even the whole survival curve. In order to keep the monotonicity of the survival fit, the hazard based models are being used and the targeting happens on the hazard level. As illustrated in Helene’s work [Rytgaard and van der Laan, 2024], the targeting can be done using the universal least favorable path which in practice needs the iterative one step TMLE with very small step sizes. As one can see, the clever co-variate in such case take a very complicated form and due to the iterative nature, it recursively depends on the initial fit

of hazard in a very complicated way, making the targeting harder and harder as targeting rounds increases. Computation burden increases exponentially. In addition, the simulation from the final updated hazard to get the Monte Carlo based estimated survival curve is computationally heavy if not impossible.

Nevertheless, our targeting within the HAL working model, possibly under-smoothed, solves this very naturally since each update only means updating the coefficients and thus the historical path is all stored within the current values of the coefficients of the HAL model. Thus, the need for recursive computations is gone and the computation burden is relieved.

We will demonstrate here how to target the whole survival curve in the classical right-censored setup.

A.3.1 Estimating the survival curve using one-step regHAL-ATMLE

A.3.1.1 Right Censored Data Set up

Suppose we have $O = (\tilde{T} = \min(T, C), \Delta = \mathbf{1}(T \leq C))$, where $T \in (0, 1)$ and $T \perp C$. For a realization o_i , we denote it as $(\tilde{t}_i, \delta_i)$.

Now, we can estimate the hazard for T using HAL based on the minimizing the penalized negative log likelihood as

$$\hat{\beta} = \underset{\beta}{\operatorname{argmin}} -L_{\beta}(O) + \lambda \|\beta\|_1 \quad (\text{A.2})$$

$$= \underset{\beta}{\operatorname{argmin}} \int_0^{\tilde{t}_i} \lambda_{\beta}^T(s) ds - \delta_i \log(\lambda_{\beta}^T(\tilde{t}_i)) + \lambda \|\beta\|_1, \quad (\text{A.3})$$

where $\lambda^T(s) = e^{\beta_0 + \sum_j \beta_j \phi_j(s)}$ and $\phi_j(s)$ are zero-order HAL basis taking the form of $\mathbf{1}(s > q_j)$ for some selected knot points q_j .

In R, we can run Poisson glmnet with the repeated data when we have 0 order HAL basis.

In addition, in order to select λ from a sequence of provided λ 's through cross validation as in usual HAL, we will fit the Poisson HAL using the Poisson cvglmnet on the repeated data.

Now, after we have the $\hat{\beta}_{cv}$, we can do the pruning to have the HAL working model for the λ^T by only retaining those basis with non zero coefficients in the cv selected model. Moreover, one can also choose to work with the under-smoothed model.

For simplicity, we now denote the β as the vector of coefficients in this new model.

A.3.1.2 Projection Based regHAL-ATMLE

Objective: Refine the initial hazard estimator so that its score aligns with the efficient influence function, thereby yielding an estimator with improved asymptotic efficiency.

1. Score Function Calculation:

For subject i with observed time $\tilde{t}_i$ and event indicator δ_i , the score is approximated by

$$(S_{i,*}^\beta)^\top = \delta_i X(\tilde{t}_i) - \int_0^{\tilde{t}_i} \exp\{X(u)^\top \beta\} X(u) du,$$

where $X(u)$ is the design vector (including an intercept and indicator basis functions).

2. Initial Gradient Computation:

An initial non-parametric gradient is computed on a grid of time points s as G . For example,

$$g_i(s) = \frac{\mathbf{1}\{\delta_i = 1 \text{ and } \tilde{t}_i > s\}}{P(C > \tilde{t}_i)},$$

where $P(C > \tilde{t}_i)$ is estimated from the censoring model. We first fit the hazard of C similar to that of T using Poisson HAL and then use that fitted hazard to calculate this survival probability of C .

3. Projection via LASSO:

For each grid point s , the initial gradient $g(s)$ is projected onto the space spanned by the score vectors $S_i = (S_{i,*}^\beta)^\top$ with unpenalized intercept γ_0 by solving

$$\min_{\gamma_0, \gamma} \sum_i \left(g_i(s) - \gamma_0 - S_i^\top \gamma \right)^2 + \lambda \|\gamma\|_1,$$

where λ is a tuning parameter. This produces a vector of projected coefficients $\gamma(s)$ and intercept $\gamma_0(s)$ for each s .

4. Aggregation of Projected Coefficients:

To form a single update direction for the working model, the candidate projections are aggregated. First, for each grid point s a scalar direction is computed:

$$d(s) = \frac{1}{n} \sum_{i=1}^n \left[S_i^\top \gamma(s) \right].$$

The collection $\{d(s)\}$ is normalized by its Euclidean norm $\|d\|_2$ and then used to weight the corresponding $\gamma(s)$ vectors:

$$\text{update} = \sum_s \frac{d(s)}{\|d\|_2} \gamma(s).$$

This weighted sum yields a single update vector for the non-intercept coefficients.

5. Projection based non-parametric EIC approximation

$$\tilde{D}_{n,\beta,proj}^*(O_i) = \sum_s \frac{d(s)}{\|d\|_2} \gamma_0(s) + \left(\sum_s \frac{d(s)}{\|d\|_2} \gamma(s) \right)^\top (S_{i,*}^\beta)^\top \quad (\text{A.4})$$

In addition, for each time grid point s , we have approximated the EIC of the survival probability at this point as

$$\tilde{D}_{s,n,\beta,proj}^*(O_i) = \gamma_0(s) + \gamma(s)^\top (S_{i,*}^\beta)^\top \quad (\text{A.5})$$

6. Coefficient Update:

The working model's coefficients are updated via

$$\beta^{\text{new}} = \beta^{\text{old}} + \text{step factor} \times \text{update.}$$

with $\text{sign}(P_n \tilde{D}_{n,\beta,proj}^*) = \text{sign}\|d\|_2 = 1$ whenever not converged.

7. Iteration and Convergence:

Steps 1–5 are repeated until

$$\frac{\|d\|_2}{\sqrt{|G|}} < \frac{\text{Median}_s \{se(\tilde{D}_{s,n,\beta,proj}^*)\}}{\sqrt{n} \log n}$$

A.3.1.3 Hessian-Based regHAL-ATMLE

As discussed in general before, we only need to change the $\gamma(s)$ in step 3 using the analytical form using the Hessian method instead of doing the lasso as

$$\gamma_{\text{Hessian}}(s)^\top = e^{-\int_0^s \lambda_\beta(u) du} \left(- \int_0^s \exp\{X(u)^\top \beta\} X(u) du \right)^\top \left(I_{P_n}(\beta) + \eta_{\text{weak}} * I \right)^{-1}$$

In addition, in step 5, we have Hessian-based parametric EIC approximations as

$$\tilde{D}_{n,\beta,hessian}^*(O_i) = \left(\sum_s \frac{d(s)}{\|d\|_2} \gamma_{\text{Hessian}}(s) \right)^\top (S_{i,*}^\beta)^\top$$

A.3.1.4 Inference for survival curve

Choosing $\tilde{D}_{s,n,\beta}^* = \tilde{D}_{s,n,\beta,proj}^*$ (Projection method) or $\tilde{D}_{s,n,\beta}^* = \tilde{D}_{s,n,\beta,hessian}^*$ (Hessian method).

After updating the β , we can now construct the Wald type 95% confidence interval for $P(T > s)$ at any given grid point s as

$$e^{-\int_0^s \lambda_\beta^T(u) du} \pm 1.96 * \sqrt{\frac{P_n[\tilde{D}_{s,n,\beta}^*(O)]^2}{n}} \quad (\text{A.6})$$

In addition, we can construct the simultaneous confidence intervals for the survival curve at all time grid point selected as mentioned in [Rytgaard and van der Laan, 2024] by

$$e^{-\int_0^s \lambda_\beta^T(u) du} \pm z_{0.95}^G * \sqrt{\frac{P_n[\tilde{D}_{s,n,\beta}^*(O)]^2}{n}} \quad \forall s \in G \quad (\text{A.7})$$

A.3.2 Simulation Set Up for Right-Censored Survival Data

We generate n independent observations. The data generation proceeds as follows:

True Event Times: For each individual $i = 1, \dots, n$, the true event time t_i^* is drawn from a Beta distribution with parameters $(\alpha, \beta) = (2, 2)$:

$$t_i^* \sim \text{Beta}(2, 2).$$

Censoring Times: Independently, a censoring time c_i is generated from a Uniform distribution on the interval $[0, 1.2]$:

$$c_i \sim \text{Uniform}(0, 1.2).$$

Observed Times and Event Indicator: The observed time $\tilde{t}_i$ is defined as the minimum of the true event time and the censoring time:

$$\tilde{t}_i = \min\{t_i^*, c_i\}.$$

An event indicator δ_i is then defined to indicate whether the event is observed:

$$\delta_i = \mathbf{1}\{t_i^* \leq c_i\},$$

where $\mathbf{1}\{\cdot\}$ is the indicator function, equal to 1 if the condition holds (i.e., the event is observed) and 0 otherwise (i.e., the observation is censored).

A.3.3 Simulation Summary Statistics

For each targeting method and for each point s in common grid of time points, we computed the following performance metrics:

- **Absolute Bias:** The average absolute difference between the estimated survival probability, $\hat{S}(s)$, and the true survival function, $S(s)$.
- **Mean Squared Error (MSE):** The average of the squared differences between $\hat{S}(s)$ and $S(s)$, reflecting overall estimation accuracy.
- **Standard Error:** The average variability (sample standard deviation) of the estimation errors across replications.
- **Coverage Probability:** The percentage of time points where the 95% confidence interval for $\hat{S}(s)$ contains the true survival value, $S(s)$. In addition, an *oracle coverage* is computed by constructing confidence intervals using the empirical variability of the estimates.

These metrics are first calculated at each time point and then averaged across the grid to provide summary performance measures for each targeting method and sample size. Table 3.8 presents the resulting metrics, facilitating a direct comparison of the estimator performance in terms of bias, error, and confidence interval coverage.

A.4 Heteroscedasticity and Non-Likelihood Losses: Detailed Development

For squared error loss $L_\beta(O) = (Y - Q_\beta(A, W))^2$, the expected Hessian is $E_P[2\phi\phi^\top]$, depending only on the covariates, while the expected outer-product of the scores is $E_P[4(Y - Q_\beta)^2\phi\phi^\top]$. These two matrices are proportional only under homoscedasticity and diverge in the presence of heteroscedasticity. Our Hessian-based method uses the unweighted empirical covariance $P_n[\phi\phi^\top]$, while the Projection method (Section 3.2.2) implicitly uses the variance-weighted empirical score covariance $P_n[4(Y - Q_\beta)^2\phi\phi^\top]$. For the ATE centered gradient $D = H(Y - Q_\beta)$, the empirical L_2 projection expands into a weighted least squares objective: $\min_\alpha P_n[(Y - Q_\beta)^2(H - 2\alpha^\top\phi)^2]$.

If the two definitions were changed to use E_P everywhere (population expectations), they would be identical. It is precisely the use of empirical quantities that makes them truly different. They become identical asymptotically only if $P_n \rightarrow P_0$ and the working model Q_n approaches the full model, since then the weighted and unweighted projections both recover the true $E[Y \mid W, A]$. If the working model converges to a proper submodel, the distinction between the weighted (Projection) and unweighted (Hessian) approaches persists asymptotically.

A.5 Proofs

As highlighted in Section 3.3, the Hessian-method and Projection-based method rely on two distinct conceptual frameworks. We present two parallel systems of proofs to guarantee

their respective asymptotic validities, assuming for simplicity that the oracle model is the full nonparametric model ($\mathcal{M}_0 = \mathcal{M}$).

A.5.1 Proof of Theorem 3.8 (Regularized TMLE)

Proof. By pathwise differentiability of Ψ at P_0 , we have the exact expansion

$$\Psi(P_n^*) - \Psi(P_0) = -P_0 D^*(P_n^*) + R(P_n^*, P_0).$$

Adding and subtracting $P_n D^*(P_n^*)$:

$$\Psi(P_n^*) - \Psi(P_0) = (P_n - P_0) D^*(P_n^*) - P_n D^*(P_n^*) + R(P_n^*, P_0).$$

Empirical process term. Write

$$(P_n - P_0) D^*(P_n^*) = (P_n - P_0) (D^*(P_n^*) - D^*(P_0)) + (P_n - P_0) D^*(P_0).$$

By conditions (R.1) and (R.2), $D^*(P_n^*)$ and $D^*(P_0)$ lie in a P_0 -Donsker class and $\|D^*(P_n^*) - D^*(P_0)\|_{L_2(P_0)} \rightarrow 0$ in probability. By asymptotic equicontinuity of the Donsker class, the first term is $o_P(n^{-1/2})$.

Score equation term. Decompose

$$P_n D^*(P_n^*) = P_n \tilde{D}_n^*(P_n^*) + P_n (D^*(P_n^*) - \tilde{D}_n^*(P_n^*)).$$

The first term is $o_P(n^{-1/2})$ by the targeting step (the algorithm solves the approximate score equation). The second term is $o_P(n^{-1/2})$ by condition (R.4).

Collecting all terms and using $P_0 D^*(P_0) = 0$:

$$\Psi(P_n^*) - \Psi(P_0) = (P_n - P_0) D^*(P_0) + o_P(n^{-1/2}) + R(P_n^*, P_0) = P_n D^*(P_0) + o_P(n^{-1/2}),$$

where the last equality uses condition (R.3). The CLT applied to

$$P_n D^*(P_0) = n^{-1} \sum_{i=1}^n D^*(P_0)(O_i)$$

yields the stated normal limit. □

A.5.2 Theoretical Framework for Hessian-Method: A-TMLE for Projection Parameters

The proofs below are direct adaptations of the proofs developed by van der Laan et al. [2023], restated here using our notation for completeness.

Proof for theorem 3.4. We have the exact expansion:

$$\begin{aligned}
& \Psi_{\mathcal{M}_n}(P_{\mathcal{M}_n,n}^*) - \Psi_{\mathcal{M}_n}(P_0) \\
&= -P_0 D_{\mathcal{M}_n}^*(P_{\mathcal{M}_n,n}^*) + R_{\mathcal{M}_n}(P_{\mathcal{M}_n,n}^*, P_0) \\
&= P_n \tilde{D}_n^*(P_{\mathcal{M}_n,n}^*) - P_0 D_{\mathcal{M}_n}^*(P_{\mathcal{M}_n,n}^*) - P_n \tilde{D}_n^*(P_{\mathcal{M}_n,n}^*) + R_{\mathcal{M}_n}(P_{\mathcal{M}_n,n}^*, P_0) \\
&= P_n D_{\mathcal{M}_n}^*(P_0) - P_0 D_{\mathcal{M}_n}^*(P_0) + (P_n - P_0)(D_{\mathcal{M}_n}^*(P_{\mathcal{M}_n,n}^*) - D_{\mathcal{M}_n}^*(P_0)) \\
&\quad - P_n \tilde{D}_n^*(P_{\mathcal{M}_n,n}^*) + P_n(D_{\mathcal{M}_n}^*(P_{\mathcal{M}_n,n}^*) - \tilde{D}_n^*(P_{\mathcal{M}_n,n}^*)) + R_{\mathcal{M}_n}(P_{\mathcal{M}_n,n}^*, P_0)
\end{aligned}$$

By definition of the non-parametric influence curve, $P_0 D_{\mathcal{M}_n}^*(P_0) = 0$. The targeting step solves $P_n \tilde{D}_n^*(P_{\mathcal{M}_n,n}^*) = o_p(n^{-1/2})$ by the algorithm stopping criterion. By Assumption 3.3 (the algorithmic regularization bias vanishes), this implies $P_n(D_{\mathcal{M}_n}^*(P_{\mathcal{M}_n,n}^*) - \tilde{D}_n^*(P_{\mathcal{M}_n,n}^*)) = o_p(n^{-1/2})$.

By condition (A.1) (Donsker class) and condition (A.2) (L_2 convergence of the working EIC), the asymptotic equicontinuity in empirical process theory implies $(P_n - P_0)(D_{\mathcal{M}_n}^*(P_{\mathcal{M}_n,n}^*) - D_{\mathcal{M}_n}^*(P_0)) = o_p(n^{-1/2})$.

By condition (A.3), $R_{\mathcal{M}_n}(P_{\mathcal{M}_n,n}^*, P_0) = o_p(n^{-1/2})$. Thus, we conclude:

$$\Psi_{\mathcal{M}_n}(P_{\mathcal{M}_n,n}^*) - \Psi_{\mathcal{M}_n}(P_0) = P_n D_{\mathcal{M}_n}^*(P_0) + o_p(n^{-1/2}).$$

□

Proof for theorem 3.5. Under $\mathcal{M}_0 = \mathcal{M}$, $\Psi(P) = \Psi_{\mathcal{M}_0}(P)$ and $D_{\mathcal{M}_0}^*(P) = D^*(P)$. Since $\mathcal{M}_0 = \mathcal{M}$, the general A-TMLE condition (C.1)— $\Psi(P_{0,\mathcal{M}_n}) - \Psi_{\mathcal{M}_0}(P_{0,\mathcal{M}_n})$ —holds trivially as an identity. Thus:

$$\begin{aligned}
\Psi_{\mathcal{M}_n}(P_0) - \Psi(P_0) &= \Psi(P_{0,\mathcal{M}_n}) - \Psi(P_0) \\
&= -P_0 D^*(P_{0,\mathcal{M}_n}) + R(P_{0,\mathcal{M}_n}, P_0) \\
&= (P_{0,\mathcal{M}_n} - P_0)(D^*(P_{0,\mathcal{M}_n}) - D_{\mathcal{M}_n}^*(P_{0,\mathcal{M}_n})) \\
&\quad + R(P_{0,\mathcal{M}_n}, P_0) - P_0 D_{\mathcal{M}_n}^*(P_{0,\mathcal{M}_n})
\end{aligned}$$

Now, $R(P_{0,\mathcal{M}_n}, P_0) = o_p(n^{-1/2})$ due to the HAL sieve's universal approximation property in that $Q_{0,\mathcal{M}_n} = \Pi_{\mathcal{M}_n}(Q_0)$ converges to Q_0 in sup-norm [van der Laan, 2023]. Combining assumptions (C.2) and (C.3), $\Psi_{\mathcal{M}_n}(P_0) - \Psi(P_0) = o_p(n^{-1/2})$.

So we have $\Psi_{\mathcal{M}_n}(P_{\mathcal{M}_n,n}^*) - \Psi(P_0) = \Psi_{\mathcal{M}_n}(P_{\mathcal{M}_n,n}^*) - \Psi_{\mathcal{M}_n}(P_0) + \Psi_{\mathcal{M}_n}(P_0) - \Psi(P_0) = P_n D_{\mathcal{M}_n}^*(P_0) + o_p(n^{-1/2})$. By asymptotic equicontinuity and assumption (B.1), $(P_n - P_0)(D_{\mathcal{M}_n}^*(P_0) - D^*(P_0)) = o_p(n^{-1/2})$. Thus, $\Psi_{\mathcal{M}_n}(P_{\mathcal{M}_n,n}^*) - \Psi(P_0) = P_n D^*(P_0) + o_p(n^{-1/2})$. □

A.5.3 Verification of Regularity Conditions under HAL

Proof for theorem 3.7. Let $\mathcal{Q}_M$ be the càdlàg class on $[0, 1]^d$ with variation norm $\leq M$. By Theorem 6 of [Bibaut and van der Laan, 2019], $\mathcal{Q}_M$ has polynomial-order bracketing entropy

and so is P_0 -Donsker.

(A.1). Each working-model EIC

$$D_{\mathcal{M}_n}^*(P) = \gamma(P)^\top S^\beta, \quad P \in \mathcal{M}_n,$$

is a finite linear combination of HAL-basis scores $S^\beta \in \mathcal{Q}_M$. A finite-dimensional span of a Donsker class remains Donsker (van der Vaart & Wellner, 1996, Ch. 2)[Van Der Vaart and Wellner, 1996]. Hence $\{D_{\mathcal{M}_n}^*(P) : P \in \mathcal{M}_n\}$ is P_0 -Donsker, and in particular $D_{\mathcal{M}_n}^*(P_{\mathcal{M}_n,n}^*)$ and $D_{\mathcal{M}_n}^*(P_0)$ lie in this class w.p. $\rightarrow 1$.

(A.2). The working-model EIC has the form $D_{\mathcal{M}_n}^*(P) = \gamma(P)^\top S^\beta$, where the coefficients $\gamma(P)$ depend smoothly on $Q(P)$ through the parameter gradient and Fisher information. Under Assumption 3.6, HAL converges in $L_2(P_0)$ at rate $O_P(n^{-1/3}(\log n)^{d/2})$, so $\|Q(P_{\mathcal{M}_n,n}^*) - Q_0\|_{L_2(P_0)} = o_P(1)$. Since the scores S^β are bounded (they are formed from HAL basis functions with bounded variation norm), the $L_2(P_0)$ convergence of Q implies

$$\|D_{\mathcal{M}_n}^*(P_{\mathcal{M}_n,n}^*) - D_{\mathcal{M}_n}^*(P_0)\|_{L_2(P_0)} = o_P(1).$$

(A.3). The exact remainder for the working-model projection parameter is

$$R_{\mathcal{M}_n}(P_{\mathcal{M}_n,n}^*, P_0) = \Psi_{\mathcal{M}_n}(P_{\mathcal{M}_n,n}^*) - \Psi_{\mathcal{M}_n}(P_0) + P_0 D_{\mathcal{M}_n}^*(P_{\mathcal{M}_n,n}^*).$$

For doubly robust target parameters such as the ATE, this remainder is bounded by cross-products of nuisance estimation errors:

$$|R_{\mathcal{M}_n}(P_{\mathcal{M}_n,n}^*, P_0)| \lesssim \|Q(P_{\mathcal{M}_n,n}^*) - Q_0\|_{L_2(P_0)} \|g_n - g_0\|_{L_2(P_0)}.$$

Under Assumption 3.6, cross-validated HAL converges in loss-based dissimilarity at rate $O_P(n^{-1/3})$ (up to $\log n$ factors), which controls each factor at rate $O_P(n^{-1/3}(\log n)^{d/2})$. Hence $|R_{\mathcal{M}_n}| = O_P(n^{-2/3}(\log n)^d) = o_P(n^{-1/2})$.

(B.1). Since $\mathcal{M}_n \rightarrow \mathcal{M}$ from below, we have

$$\|D_{\mathcal{M}_n}^*(P_0) - D^*(P_0)\|_{L^2} = o(1).$$

(C.1). Because $\mathcal{M}_0 = \mathcal{M}$, $\Psi_{\mathcal{M}_0} = \Psi$ and thus $\Psi(P_{0,\mathcal{M}_n}) - \Psi_{\mathcal{M}_0}(P_{0,\mathcal{M}_n}) = 0$.

(C.2). We must establish that

$$R_n = (P_{0,\mathcal{M}_n} - P_0)(D^*(P_{0,\mathcal{M}_n}) - D_{\mathcal{M}_n}^*(P_{0,\mathcal{M}_n})) = o_p(n^{-1/2}).$$

We prove this utilizing the general exact remainder bounds for doubly robust parameters. The working model $\mathcal{M}_n$ restricts only the outcome regression Q , leaving the marginal density and propensity score g unspecified. Therefore, the projected distribution $P_{0,\mathcal{M}_n}$ shares the exact same true nuisance component g_0 as P_0 .

Let $H_0 = H(A, W; g_0)$ denote the true clever covariate evaluated under the shared true nuisance g_0 . The non-parametric EIC evaluated at $P_{0, \mathcal{M}_n}$ takes the general form:

$$D^*(P_{0, \mathcal{M}_n}) = H_0(Y - Q_{0, \mathcal{M}_n}) + D_{\text{marg}}(P_0),$$

where D_{marg} consists only of marginal terms independent of Y .

As established in Section 3.2.3, the working-model EIC $D_{\mathcal{M}_n}^*(P_{0, \mathcal{M}_n})$ exactly corresponds to the L_2 projection of the nonparametric EIC onto the span of the active working model score space. Since the working model scores exclusively update Q , this projection implicitly projects the clever covariate H_0 onto the linear span of the basis functions used for Q . Denoting this score space projection operator as Π_S , we have:

$$D_{\mathcal{M}_n}^*(P_{0, \mathcal{M}_n}) = (\Pi_S H_0)(Y - Q_{0, \mathcal{M}_n}) + D'_{\text{marg}}(P_0).$$

Thus, taking the difference completely cancels the marginal components, and the difference is strictly given by the residual of the clever covariate projection:

$$D_{\perp} := D^*(P_{0, \mathcal{M}_n}) - D_{\mathcal{M}_n}^*(P_{0, \mathcal{M}_n}) = (H_0 - \Pi_S H_0)(Y - Q_{0, \mathcal{M}_n}).$$

Because $Y - Q_0$ is a mean-zero residual under the true data generating distribution P_0 , evaluating the integral with respect to P_0 cancels Y , and the exact remainder simplifies as:

$$\begin{aligned} R_n &= (P_{0, \mathcal{M}_n} - P_0)D_{\perp} \\ &= 0 - P_0[D_{\perp}] \\ &= - \int (H_0 - \Pi_S H_0)(Q_0 - Q_{0, \mathcal{M}_n}) dP_0. \end{aligned}$$

Applying the Cauchy-Schwarz inequality, we bound the magnitude of this integral:

$$|R_n| \leq \|H_0 - \Pi_S H_0\|_{L_2(P_0)} \times \|Q_0 - Q_{0, \mathcal{M}_n}\|_{L_2(P_0)}.$$

Under Assumption 3.6, the HAL sieve converges to the true function such that $\|Q_0 - Q_{0, \mathcal{M}_n}\|_{L_2(P_0)} = O_P(n^{-r})$ with $r > 1/4$. Therefore, the exact remainder is strictly $o_p(n^{-1/2})$ if and only if the active HAL basis selected for Q is sufficiently rich to approximate the clever covariate H_0 such that $\|H_0 - \Pi_S H_0\|_{L_2(P_0)} = O_P(n^{-r'})$ where $r + r' \geq 1/2$.

Because Π_S projects only onto the active subset $\mathcal{M}_n$ selected by the Lasso to fit Q , the penalty may drop terms vital for H_0 . To mathematically guarantee this $o_p(n^{-1/2})$ remainder bound strictly under post-selection sparsity, one must assume $\mathcal{M}_n$ is sufficiently rich to span H_0 . As explicitly noted in Assumption 3.6, this is rigorously achieved in theory by defining $\mathcal{M}_n$ as the *union* of the active bases selected from preliminary regularized fits of both Q and g —a standard post-double selection strategy [Belloni et al., 2014] that theoretically preserves the required $O_p(n^{-1/4})$ approximation rate for both the outcome and the clever covariate components. Our practical algorithmic implementations instead intentionally rely on the empirical richness of an adaptively undersmoothed HAL sieve for Q_n to flexibly adapt to capture H_0 . Thus, when satisfied, $|R_n| = O_p(n^{-1/4}) \times O_p(n^{-1/4}) = o_p(n^{-1/2})$.

(C.3) We must show $P_0[D_{\mathcal{M}_n}^*(P_{0,\mathcal{M}_n})] = o_p(n^{-1/2})$.

By definition, the projected distribution $P_{0,\mathcal{M}_n}$ minimizes the expected loss within the working model $\mathcal{M}_n$ under the true distribution P_0 . The first-order optimality condition dictates that the expected scores evaluated at this projection parameter are exactly zero: $P_0[S^{\beta_{0,\mathcal{M}_n}}] = 0$.

Since the working-model EIC $D_{\mathcal{M}_n}^*(P_{0,\mathcal{M}_n})$ is purely a linear combination of these exact score functions (i.e., $D_{\mathcal{M}_n}^* = \gamma^\top S^{\beta_{0,\mathcal{M}_n}}$), it is a strict mathematical identity that:

$$P_0[D_{\mathcal{M}_n}^*(P_{0,\mathcal{M}_n})] = \gamma^\top P_0[S^{\beta_{0,\mathcal{M}_n}}] = 0$$

This algebraic identity satisfies the condition immediately, establishing $P_0[D_{\mathcal{M}_n}^*(P_{0,\mathcal{M}_n})] = 0 = o_p(n^{-1/2})$ without relying on asymptotic rate bounds. $\square$

A.6 Alternative plateau selection via *bridged* targeting

The *plateau selector* in §3.4 evaluates a *separate* targeting path inside each nested working model

$$\mathcal{M}_1 = \mathcal{M}_{\text{cv}} \subset \mathcal{M}_2 \subset \dots \subset \mathcal{M}_K,$$

always *re-initialising* from the cross-validated HAL estimate $\hat{Q}_n$. Here we describe a computationally lighter alternative in which the TMLE update in one model *feeds directly* into the next, so that information gathered in $\mathcal{M}_j$ is *bridged* to $\mathcal{M}_{j+1}$.

A.6.1 Motivation

Let $\hat{\beta}_j$ denote the targeted coefficient vector in $\mathcal{M}_j$, obtained by the regHAL-ATMLE of §3.2.2. Because $\mathcal{M}_j \subset \mathcal{M}_{j+1}$, we can embed $Q_{\hat{\beta}_j}$ in the larger model by padding zeros for the additional basis functions. Targeting then continues in $\mathcal{M}_{j+1}$ *from this padded point* instead of from $\hat{Q}_n$. Formally,

$$\tilde{\beta}_{j \rightarrow j+1} = \left(\hat{\beta}_j^\top, \mathbf{0}_{p_{j+1}-p_j}^\top \right)^\top, \quad Q_{\tilde{\beta}_{j \rightarrow j+1}} \in \mathcal{M}_{j+1},$$

and the projection-based TMLE update is applied to $\tilde{\beta}_{j \rightarrow j+1}$ with the *larger* score matrix $S^{(j+1)}$.

Intuitively, the canonical gradient is approximated ever more accurately as additional scores enter, so the bias-reduction step becomes “finer-grained” while retaining the variance control already achieved in the smaller model.

A.6.2 Algorithm

Algorithm 12 BRIDGED TARGETING PLATEAU SELECTOR

Require: Initial HAL fit $\hat{Q}_n$; nested models $\{\mathcal{M}_j\}_{j=1}^K$; stopping tolerance δ_n (§3.2.2).

- 1: **Initial step** ($j = 1$): run regHAL-ATMLE in $\mathcal{M}_1$, giving $\hat{\beta}_1$.
 - 2: **for** $j = 2, \dots, K$ **do**
 - 3: Embed $\hat{\beta}_{j-1}$ in $\mathcal{M}_j$ to get $\tilde{\beta}_{j-1 \rightarrow j}$.
 - 4: Continue projection-based targeting in $\mathcal{M}_j$ starting from $\tilde{\beta}_{j-1 \rightarrow j}$ until $|P_n D_{\mathcal{M}_j}^*| < \delta_n$, producing $\hat{\beta}_j$.
 - 5: Record point estimate $\hat{\psi}_j = \Psi^F(Q_{\hat{\beta}_j})$ and $(1 - \alpha)$ CI from $P_n D_{\mathcal{M}_j}^*(\hat{\beta}_j)$.
 - 6: **Check plateau rule:** if $\hat{\psi}_j$ and its CI violate the monotone-width pattern of §3.4, break and set $j^* = j - 1$.
 - 7: **end for**
 - 8: **return** $\hat{\psi}_{j^*}$ and its CI.
-

Stopping criterion. Within each model the update stops when the empirical mean of the *projected* EIC has magnitude below $\delta_n = \text{se}[D_{\mathcal{M}_j}^*(\hat{\beta}_j)] / \{\sqrt{n} \log n\}$ (same as Algorithm 1).

A.6.3 Discussion

- **Efficiency gain.** Only $K - 1$ additional targeting runs are needed, each warm-started close to optimality. Empirically we may observe significant reduction in wall-clock time relative to cold starts.
- **Bias-variance trade-off.** Because every subsequent model begins with a *zero-bias* estimator for $\Psi^F(Q_{\mathcal{M}_{j-1}})$, the additional bias removed in $\mathcal{M}_j$ is purely due to scores absent from $\mathcal{M}_{j-1}$. Variance increases more gradually, often yielding a clearer plateau.
- **Link to sieve TMLE.** Bridged targeting resembles a sieve estimator with increasing basis dimension, but retains HAL’s automatic knot placement and ℓ_1 regularization.
- **When to prefer.** Use bridged targeting when K is large or when the working models differ only by a handful of basis functions (typical for HAL sequences along λ).

This alternative rule offers a pragmatic middle ground: it exploits the nesting of HAL models to recycle information, yet retains the intuitive “stop-when-CI-widens” principle that makes the original plateau selector attractive in practice.

Appendix B

Supporting Material for Chapter 4

This appendix gathers the proof arguments, training details, subset-selection material, and supplemental empirical specifications for Chapter 4. The main chapter keeps the construction, theorem statements, and headline empirical evidence in the body.

B.1 Appendix: Proofs of Theoretical Results

Our theoretical analysis seamlessly adapts the proof systems developed for A-TMLE and, in the dissertation setting, the regHAL construction of Chapter 3 to neural network working models.

B.1.1 Proof of Theorem 4.6 (A-TMLE for Working Parameter)

Proof. We have the exact pathwise expansion:

$$\Psi_{\mathcal{M}_n}(P_{\mathcal{M}_{n,n}}^*) - \Psi_{\mathcal{M}_n}(P_0) = -P_0 D_{\mathcal{M}_n}^*(P_{\mathcal{M}_{n,n}}^*) + R_{\mathcal{M}_n}(P_{\mathcal{M}_{n,n}}^*, P_0)$$

Adding and subtracting $P_n D_{\mathcal{M}_n}^*(P_{\mathcal{M}_{n,n}}^*)$:

$$\Psi_{\mathcal{M}_n}(P_{\mathcal{M}_{n,n}}^*) - \Psi_{\mathcal{M}_n}(P_0) = (P_n - P_0) D_{\mathcal{M}_n}^*(P_{\mathcal{M}_{n,n}}^*) - P_n D_{\mathcal{M}_n}^*(P_{\mathcal{M}_{n,n}}^*) + R_{\mathcal{M}_n}(P_{\mathcal{M}_{n,n}}^*, P_0)$$

We split the empirical process term $(P_n - P_0) D_{\mathcal{M}_n}^*(P_{\mathcal{M}_{n,n}}^*)$ into:

$$(P_n - P_0) \left(D_{\mathcal{M}_n}^*(P_{\mathcal{M}_{n,n}}^*) - D_{\mathcal{M}_n}^*(P_0) \right) + (P_n - P_0) D_{\mathcal{M}_n}^*(P_0)$$

By condition (A.1) (Donsker class) and condition (A.2) (L_2 convergence of the working EIC), asymptotic equicontinuity ensures that $(P_n - P_0) \left(D_{\mathcal{M}_n}^*(P_{\mathcal{M}_{n,n}}^*) - D_{\mathcal{M}_n}^*(P_0) \right) = o_p(n^{-1/2})$.

Furthermore, because $P_0 D_{\mathcal{M}_n}^*(P_0) = 0$ by definition of the EIC, we have:

$$(P_n - P_0) D_{\mathcal{M}_n}^*(P_0) = P_n D_{\mathcal{M}_n}^*(P_0)$$

For the score equation term, Assumption 4.3 directly guarantees $-P_n D_{\mathcal{M}_n}^*(P_{\mathcal{M}_n,n}^*) = o_p(n^{-1/2})$. Finally, condition (A.3) provides $R_{\mathcal{M}_n}(P_{\mathcal{M}_n,n}^*, P_0) = o_p(n^{-1/2})$.

Substituting these bounds into the expansion yields:

$$\Psi_{\mathcal{M}_n}(P_{\mathcal{M}_n,n}^*) - \Psi_{\mathcal{M}_n}(P_0) = P_n D_{\mathcal{M}_n}^*(P_0) + o_p(n^{-1/2}).$$

□

B.1.2 Proof of Theorem 4.7 (A-TMLE for Original Parameter)

Proof. Under $\mathcal{M}_0 = \mathcal{M}$, we have $\Psi_{\mathcal{M}_0} = \Psi$ and $D_{\mathcal{M}_0}^* = D^*$. The condition $\Psi(P_{0,\mathcal{M}_n}) - \Psi_{\mathcal{M}_0}(P_{0,\mathcal{M}_n}) = 0$ holds trivially as an identity. Thus, by pathwise expansion around the working model projection $P_{0,\mathcal{M}_n}$:

$$\begin{aligned} \Psi_{\mathcal{M}_n}(P_0) - \Psi(P_0) &= \Psi(P_{0,\mathcal{M}_n}) - \Psi(P_0) \\ &= -P_0 D^*(P_{0,\mathcal{M}_n}) + R(P_{0,\mathcal{M}_n}, P_0) \end{aligned}$$

Because the full non-parametric EIC has mean zero evaluated at its own underlying distribution, we strictly have $P_{0,\mathcal{M}_n} D^*(P_{0,\mathcal{M}_n}) = 0$. Similarly, the working-model EIC inherently has mean zero evaluated under the projected distribution, giving $P_{0,\mathcal{M}_n} D_{\mathcal{M}_n}^*(P_{0,\mathcal{M}_n}) = 0$.

Leveraging these identities, we can rewrite the primary linear term as:

$$\begin{aligned} -P_0 D^*(P_{0,\mathcal{M}_n}) &= (P_{0,\mathcal{M}_n} - P_0) D^*(P_{0,\mathcal{M}_n}) \\ &= (P_{0,\mathcal{M}_n} - P_0) \left(D^*(P_{0,\mathcal{M}_n}) - D_{\mathcal{M}_n}^*(P_{0,\mathcal{M}_n}) \right) - P_0 D_{\mathcal{M}_n}^*(P_{0,\mathcal{M}_n}). \end{aligned}$$

Substituting this identity into our original expansion yields:

$$\begin{aligned} \Psi_{\mathcal{M}_n}(P_0) - \Psi(P_0) &= (P_{0,\mathcal{M}_n} - P_0) \left(D^*(P_{0,\mathcal{M}_n}) - D_{\mathcal{M}_n}^*(P_{0,\mathcal{M}_n}) \right) \\ &\quad - P_0 D_{\mathcal{M}_n}^*(P_{0,\mathcal{M}_n}) + R(P_{0,\mathcal{M}_n}, P_0) \end{aligned}$$

For neural networks, $R(P_{0,\mathcal{M}_n}, P_0) = o_p(n^{-1/2})$ holds under condition (C.1). Condition (C.2) bounds the first term to $o_p(n^{-1/2})$, and condition (C.3) bounds the second term to $o_p(n^{-1/2})$. Thus, $\Psi_{\mathcal{M}_n}(P_0) - \Psi(P_0) = o_p(n^{-1/2})$.

We can now write:

$$\Psi_{\mathcal{M}_n}(P_{\mathcal{M}_n,n}^*) - \Psi(P_0) = \left(\Psi_{\mathcal{M}_n}(P_{\mathcal{M}_n,n}^*) - \Psi_{\mathcal{M}_n}(P_0) \right) + \left(\Psi_{\mathcal{M}_n}(P_0) - \Psi(P_0) \right)$$

From Theorem 4.6, the first term is $P_n D_{\mathcal{M}_n}^*(P_0) + o_p(n^{-1/2})$, giving:

$$\Psi_{\mathcal{M}_n}(P_{\mathcal{M}_n,n}^*) - \Psi(P_0) = P_n D_{\mathcal{M}_n}^*(P_0) + o_p(n^{-1/2})$$

Since $P_{\mathcal{M}_n,n}^* \in \mathcal{M}_n$ (which holds by construction since the TDA estimator updates the network weights and thus strictly remains within the working model), the projected parameter evaluated at the estimator equals the original parameter evaluated at the estimator:

$$\Psi_{\mathcal{M}_n}(P_{\mathcal{M}_n,n}^*) = \Psi(P_{\mathcal{M}_n,n}^*).$$

Finally, we decompose $P_n D_{\mathcal{M}_n}^*(P_0)$ as:

$$P_n D^*(P_0) + (P_n - P_0)(D_{\mathcal{M}_n}^*(P_0) - D^*(P_0)) + P_0(D_{\mathcal{M}_n}^*(P_0) - D^*(P_0))$$

By asymptotic equicontinuity and assumption (B.1), the empirical process term is $o_p(n^{-1/2})$. Because both EICs have mean zero under P_0 , the third term is exactly zero. Thus, $\Psi(P_{\mathcal{M}_n,n}^*) - \Psi(P_0) = P_n D^*(P_0) + o_p(n^{-1/2})$. $\square$

B.1.3 Proof of Theorem 4.8 (Verification of Regularity Conditions)

Proof. Conditions (A.1), (A.2), and (B.1) hold by arguments analogous to those used in Chapter 3. Specifically, neural network gradients form a Donsker class under standard boundedness and weight decay regularizations, and $D_{\mathcal{M}_n}^*$ converges to D^* as the network's capacity grows. Assumption 4.3 provides the bridge from the algorithmic stopping criterion to the true working-model EIC equation being solved.

To establish (A.3), we note that for doubly robust target parameters, the exact remainder $R_{\mathcal{M}_n}(P_{\mathcal{M}_n,n}^*, P_0)$ is bounded by the product of $L_2(P_0)$ estimation errors of the relevant nuisance parameters. By Condition 1, this product is bounded by $O_p(n^{-r})O_p(n^{-s}) = o_p(n^{-1/2})$ since $r + s > 1/2$.

For (C.2), we establish that $(P_{0,\mathcal{M}_n} - P_0)(D^*(P_{0,\mathcal{M}_n}) - D_{\mathcal{M}_n}^*(P_{0,\mathcal{M}_n})) = o_p(n^{-1/2})$. The working model $\mathcal{M}_n$ constrains only $Q(P)$ (the part of P relevant for the target parameter), leaving the remaining components of P —including the nuisance $g(P)$ and the marginal distribution—unconstrained. Consequently, the projected distribution $P_{0,\mathcal{M}_n}$ shares the exact same nuisance component g_0 and marginal distribution as P_0 ; only Q is replaced by its working-model projection $Q_{0,\mathcal{M}_n}$.

The nonparametric EIC can generally be decomposed as

$$D^*(Q, g_0)(O) = D_Q^*(Q, g_0)(O) + D_{\text{rem}}^*(g_0)(O),$$

where D_Q^* is the Q -tangent component that depends on Q and D_{rem}^* collects terms that do not depend on Q . (For the ATE, $D_Q^*(Q, g_0)(O) = H_0(A, X)(Y - Q(A, X))$ where H_0 is the clever covariate, and $D_{\text{rem}}^* = Q(1, X) - Q(0, X) - \Psi(P)$.)

Because $\mathcal{M}_n$ leaves g and the marginals unconstrained, its tangent space at $P_{0,\mathcal{M}_n}$ includes the full nonparametric nuisance tangent space. Thus, the $L_2(P_{0,\mathcal{M}_n})$ projection of D^* onto the working model tangent space perfectly preserves D_{rem}^* , while projecting the Q -component D_Q^* onto the network's score space. Denoting this score space projection operator as Π_S , the difference is:

$$D_{\perp} := D^*(P_{0,\mathcal{M}_n}) - D_{\mathcal{M}_n}^*(P_{0,\mathcal{M}_n}) = D_Q^*(Q_{0,\mathcal{M}_n}, g_0) - \Pi_S D_Q^*(Q_{0,\mathcal{M}_n}, g_0).$$

Because $Q_{0,\mathcal{M}_n}$ minimizes the expected loss $P_0 L(Q)$ within the working model, the first-

order optimality condition ensures $P_{0,\mathcal{M}_n}[D_\perp] = 0$. Thus:

$$(P_{0,\mathcal{M}_n} - P_0)D_\perp = -P_0[D_\perp].$$

The loss-based structure of Q ensures that $P_0[D_Q^*(Q_{0,\mathcal{M}_n}, g_0)]$ involves the discrepancy $Q_0 - Q_{0,\mathcal{M}_n}$. (For the ATE, this is explicit: $P_0[H_0(Y - Q_{0,\mathcal{M}_n})] = P_0[H_0(Q_0 - Q_{0,\mathcal{M}_n})]$ since $E_{P_0}[Y | A, X] = Q_0$.) Applying Cauchy–Schwarz:

$$|(P_{0,\mathcal{M}_n} - P_0)D_\perp| \leq \|D_Q^* - \Pi_S D_Q^*\|_{L_2(P_0)} \times \|Q_0 - Q_{0,\mathcal{M}_n}\|_{L_2(P_0)}.$$

Under Condition 1, $\|Q_0 - Q_{0,\mathcal{M}_n}\|_{L_2(P_0)} = O_p(n^{-r})$. Under Condition 2, the neural network score space well-approximates the Q -component of the EIC such that the projection residual $\|D_Q^* - \Pi_S D_Q^*\|_{L_2(P_0)}$ is small enough for the product to be $o_p(n^{-1/2})$.

Furthermore, for (C.1), the exact remainder $R(P_{0,\mathcal{M}_n}, P_0)$ is bounded by products of nuisance estimation errors. Because $P_{0,\mathcal{M}_n}$ and P_0 share the same g_0 , the g -error factor vanishes identically, giving $R(P_{0,\mathcal{M}_n}, P_0) = 0$ for doubly robust parameters. (For non-doubly-robust parameters, $R(P_{0,\mathcal{M}_n}, P_0) = o_p(n^{-1/2})$ follows from the working model approximation rate for Q .)

For (C.3), the projected distribution $P_{0,\mathcal{M}_n}$ minimizes the expected loss within the working model under P_0 . The first-order optimality condition yields $P_0[S_{\theta_{\text{targ}}}(P_{0,\mathcal{M}_n})] = 0$. Since $D_{\mathcal{M}_n}^*(P_{0,\mathcal{M}_n})$ is a linear combination of these scores, $P_0[D_{\mathcal{M}_n}^*(P_{0,\mathcal{M}_n})] = 0$, satisfying the condition algebraically. $\square$

B.1.4 Proof of Theorem 4.4 (Regularized TMLE)

Proof. By pathwise differentiability of Ψ at P_0 , we have the exact expansion:

$$\Psi(P_n^*) - \Psi(P_0) = -P_0 D^*(P_n^*) + R(P_n^*, P_0).$$

Adding and subtracting $P_n D^*(P_n^*)$:

$$\Psi(P_n^*) - \Psi(P_0) = (P_n - P_0) D^*(P_n^*) - P_n D^*(P_n^*) + R(P_n^*, P_0).$$

Empirical process term. Write

$$(P_n - P_0) D^*(P_n^*) = (P_n - P_0) (D^*(P_n^*) - D^*(P_0)) + (P_n - P_0) D^*(P_0).$$

By conditions (R.1) and (R.2), $D^*(P_n^*)$ and $D^*(P_0)$ lie in a P_0 -Donsker class and their $L_2(P_0)$ difference converges to zero in probability. By asymptotic equicontinuity, the first term is $o_p(n^{-1/2})$.

Score equation term. Decompose

$$-P_n D^*(P_n^*) = -P_n \tilde{D}^*(P_n^*) - P_n (D^*(P_n^*) - \tilde{D}^*(P_n^*)).$$

The first term is $o_p(n^{-1/2})$ by the TDA stopping criterion. The second term is $o_p(n^{-1/2})$ by

condition (R.4).

Collecting all terms and using $P_0 D^*(P_0) = 0$:

$$\Psi(P_n^*) - \Psi(P_0) = P_n D^*(P_0) + o_p(n^{-1/2}),$$

where the last equality uses condition (R.3). The CLT applied to $P_n D^*(P_0)$ yields the stated normal limit. $\square$

B.2 Appendix: Experimental Details

This appendix provides additional information on dataset characteristics, implementation details, hyperparameter choices, and supplementary results for the experiments described in Section 4.4 of the main text.

B.2.1 ATE Experiments: IHDP and Positivity Violations

B.2.1.1 IHDP Dataset

Dataset and Setup. The IHDP dataset combines covariates from a real randomized experiment studying educational interventions for premature infants with semi-synthetic outcomes following the “Response Surface B” setting from Hill [2011]. It contains 747 units with 25 covariates (6 continuous, 19 binary), a binary treatment indicator, and an approximate true ATE of 4.0. We run 500 replications (seed offset 2000), each using a different synthetic outcome realization. Data splitting: 80% training, 20% validation (with early stopping), and full data for the post-training targeting step.

DragonNet Architecture. We implement DragonNet Shi et al. [2019] in PyTorch with: 2 trunk layers (64 units each, ELU activation); 1 propensity head (64 units, ELU $\rightarrow$ sigmoid); 2-layer outcome heads for μ_0 and μ_1 (64 hidden units, ELU, linear output). Training uses Adam (lr = 0.001, weight decay 10^{-5} , batch size 128) with early stopping (patience 10 epochs) on the combined loss (binary cross-entropy for propensity + MSE for outcomes).

Estimator Definitions. We compare 13 estimators organized into three families.

Group 1: Post-training debiasing (standard initialization).

- **Naive Plug-in:** $\hat{\psi} = n^{-1} \sum_i [\hat{\mu}_1(X_i) - \hat{\mu}_0(X_i)]$, no targeting.
- **TDA-Last:** Post-training TDA targeting only the final linear layers of the outcome heads (≈ 128 parameters). Uses Algorithm 1 with $\lambda = 0.01$, patience 10, max 100 iterations.
- **TDA-Full:** Post-training TDA targeting both layers of each outcome head ($\approx 8,320$ parameters). Same hyperparameters as TDA-Last.
- **Classical TMLE:** Single ε -fluctuation via the clever covariate $H(X) = A/\hat{p}(X) - (1 - A)/(1 - \hat{p}(X))$ on the logit scale.

- **AIPTW**: Augmented Inverse Probability of Treatment Weighting, a one-step doubly robust estimator (not a targeting method): $\hat{\psi}_{\text{AIPTW}} = n^{-1} \sum_i [A_i(Y_i - \hat{\mu}_1)/\hat{p} - (1 - A_i)(Y_i - \hat{\mu}_0)/(1 - \hat{p}) + \hat{\mu}_1 - \hat{\mu}_0]$.

Group 2: Training-time targeting (TTR). TTR periodically applies the TDA targeting step *during* network training, every $f_{\text{targ}} = 10$ training epochs. At each targeting epoch, the current model parameters are updated via Algorithm 1 (with $\lambda_{\text{targ}} = 0.01$), and training then resumes from the updated parameters. This steers the model toward regions of parameter space where the EIC equation is approximately satisfied, potentially producing a better initialization for post-training targeting.

- **TTR-Last**: TTR applied to last layer only; no post-training TDA.
- **TTR-Full**: TTR applied to both outcome head layers; no post-training TDA.
- **Combined-Last**: Two-stage—TTR-Last during training, then post-training TDA-Last.
- **Combined-Full**: Two-stage—TTR-Full during training, then post-training TDA-Full.

Group 3: T-regularization family. T-regularization Shi et al. [2019] adds an EIC-based loss penalty $\alpha \cdot |P_n[D^*(\hat{Q}, \hat{g})]|$ during training. Unlike TTR, which directly applies the targeting *step*, T-reg penalizes the loss function to encourage EIC residual reduction without explicitly solving the EIC equation.

- **T-reg**: T-regularization only; no post-training targeting.
- **T-reg + TDA-Last**: T-regularization training, then post-training TDA-Last.
- **T-reg + TDA-Full**: T-regularization training, then post-training TDA-Full.
- **TMLE Update**: T-regularization training, then Classical TMLE update.

Propensity Truncation. All methods use Gruber truncation: $\hat{p}(X) \in [g_{\text{bound}}, 1 - g_{\text{bound}}]$ with $g_{\text{bound}} = 5/(\sqrt{n} \log n)$.

Confidence Intervals. 95% CIs are constructed as $\hat{\psi} \pm 1.96 \cdot \hat{\sigma}/\sqrt{n}$, where $\hat{\sigma}^2$ is the sample variance of the estimated efficient influence curve values.

B.2.1.2 Positivity Violation DGP

Data Generation. We design a DGP where the propensity model creates near-positivity violations and the CATE shares drivers with the propensity, so that plug-in errors are systematic in the low-data region. Specifically, $X \sim \text{Uniform}(-1, 1)^5$, with propensity:

$$\text{logit } \pi(X) = 3.5 X_1 + 2.5 X_2,$$

yielding propensity scores ranging from 0.002 to 0.998. The control outcome is:

$$\mu_0(X) = 2 \sin(2X_1 X_2) + X_3^2 + 1.5|X_1 - X_4| + X_2 X_5,$$

and the CATE is:

$$\tau(X) = 1 + 2.5 X_1^2 + 1.5 \cos(2X_2) + 3 X_1 X_2 + 2 \max(0, -X_1) X_3,$$

with $Y = \mu_0(X) + \tau(X) A + \varepsilon$, $\varepsilon \sim \mathcal{N}(0, 1)$, and true ATE ≈ 1.5 (computed via Monte Carlo with 500,000 samples). The last CATE term, $2 \max(0, -X_1) X_3$, is a “hidden” effect active only in the low-propensity region ($X_1 < 0$), where treated units are scarce and plug-in estimators have limited information.

We run 500 replications at each of $n \in \{500, 1,000\}$ (seed offset 3000). All other hyperparameters (architecture, training, targeting) are identical to the IHDP setup.

B.2.2 Survival Curve Experiments

B.2.2.1 Concordant Late Acceleration DGP

We simulate survival data from a DGP designed to stress-test censoring adjustment. The key feature is a *late-time treatment effect acceleration* that is systematically hidden by *concordant informative censoring*: subjects most affected by the acceleration are preferentially censored before it manifests.

Covariates. $X \sim \mathcal{N}(0, \Sigma)$ with $\Sigma_{ij} = 0.5^{|i-j|}$ ($n = 1,000$, $p = 10$).

Event time model. Event times are generated via numerical inversion of the cumulative hazard:

$$\lambda(t \mid X) = \lambda_0(t; X) \cdot \exp\left(\eta_{\text{smooth}}(t, X) + \eta_{\text{late}}(t, X)\right),$$

where $\lambda_0(t; X)$ is a Weibull baseline with shape $\alpha = 1.5$ and subject-specific scale:

$$\eta_0(X) = 8 \cdot \exp\left(-0.5 \underbrace{\left[0.4(e^{0.5X_4} + e^{0.3X_5}) - 0.8 + 0.3 \sin(\pi X_6)\right]}_{\text{nonlinear effect}} + \underbrace{0.3 X_1 X_4 \cdot \mathbf{1}(X_7 > 0)}_{\text{three-way interaction}}\right).$$

The time-varying component combines smooth and threshold effects:

$$\eta_{\text{smooth}}(t, X) = \left[0.8(0.8X_1 + 0.6X_2 + 0.4X_3) + 0.2 X_2 X_8\right] \sqrt{t} + 0.5 \cdot \mathbf{1}(X_9 > 0.5) \cdot \log(1 + t).$$

The late acceleration activates after $t = 7$:

$$\eta_{\text{late}}(t, X) = 2 \sigma(3(X_1 + 0.5X_2)) \cdot \log(1 + \max(0, t - 7)),$$

where $\sigma(\cdot)$ is the sigmoid function. This creates a hidden effect: subjects with large X_1 (and positive X_2) experience dramatically increasing hazard after $t = 7$.

Concordant censoring. Censoring times are generated from an exponential model with all-positive covariate effects: $\gamma = (1.0, 0.8, 0.5, 0.4, 0.3, 0.2, 0.15, 0.15, 0.1, 0.05)$, with baseline rate $\eta_{\text{base}} = 3.0$. Because $\gamma_j > 0$ for all j and the late acceleration depends on X_1, X_2 with

positive coefficients, subjects who would experience the strongest acceleration are also the most heavily censored—precisely when they are most informative. The overall censoring rate is approximately 73%.

Observed data. $\tilde{T} = \min(T, C)$, $\Delta = \mathbf{1}(T \leq C)$, evaluated on a grid of 50 equally-spaced time points $s_{\text{fine}} = \text{linspace}(0, 14, 50)$.

B.2.2.2 Neural Hazard Model

Architecture. Input: covariates X concatenated with time t ($p + 1$ dimensions). Three hidden layers with dimensions $[64, 32, 16]$, ReLU activations, dropout (rate 0.2). Final linear layer outputs $\log \lambda(t|X)$.

Training. Poisson log-likelihood: $\ell(\lambda, \Delta, \tau) = \tau\lambda - \Delta \log(\lambda\tau)$. Adam optimizer (lr = 0.001), mini-batch size 64, early stopping on validation loss (patience 10 epochs). Survival functions computed by numerical integration: $\hat{S}(t|X) = \exp(-\int_0^t \hat{\lambda}(u|X) du)$.

Censoring model. A separate neural network with identical architecture is trained to model the conditional censoring process $G(t|X) = P(C > t|X)$, used in EIC construction.

B.2.2.3 TDA Implementation for Survival

Post-training targeting for survival proceeds via Algorithm 1 with several practical safeguards:

Regularization scaling. For last-layer targeting (TDA-Last, 16 parameters), the base regularization is $\lambda = 10^{-5}$. When targeting additional layers, we scale $\lambda_k = \lambda_{\text{base}} \cdot (d_k/d_{\text{base}})$, where d_k is the number of targeted parameters and $d_{\text{base}} = 16$, ensuring that larger parameter spaces receive proportionally more regularization.

Step factor scaling. The effective step size is $\varepsilon_{\text{eff}} = \varepsilon_{\text{base}} \cdot (16/d_k)$, with $\varepsilon_{\text{base}} = 10^{-2}$ for last-layer targeting. When targeting additional layers, smaller base step sizes are needed (e.g., 10^{-3}) to account for the amplification of perturbations through nonlinear hidden layers.

NLL tolerance. Targeting is expected to slightly increase the negative log-likelihood while improving the survival curve (solving the EIC equation). We allow NLL to increase up to 1% from its initial value; if it exceeds this threshold, we restore the best model state and terminate. This prevents catastrophic divergence while permitting the controlled trade-off between predictive fit and EIC satisfaction.

Convergence criterion. Targeting terminates when the normalized EIC residual satisfies $\|P_n[\tilde{D}^*]\|_2/\sqrt{m} < \hat{\sigma}_{\text{med}}/(\sqrt{n} \log n)$, where $m = 50$ is the number of time points and $\hat{\sigma}_{\text{med}}$ is the median standard deviation across time points. Maximum 100 iterations.

Universal targeting. For $m = 50$ time points, we compute the EIC at each time point, project onto the score space via regularized least squares, and form a single universal update direction weighted by the per-time-point EIC residuals. This forces all 50 targeting equations to share a common parameter update direction, globally smoothing the survival curve.

Confidence intervals. 95% CIs: $\hat{S}(t) \pm 1.96 \cdot \hat{\sigma}_t / \sqrt{n}$, where $\hat{\sigma}_t^2$ is the empirical variance of the projected influence function at time t . For KM: standard Greenwood formula. For survtmle: its built-in influence-function-based variance estimator.

survtmle comparison. The survtmle estimator Benkeser and Hejazi [2017] uses the same discretization grid as our evaluation grid: 50 equally-spaced points on $[0, 14]$ (i.e., `seq(0, 14, length.out = 50)` in R = `linspace(0, 14, 50)` in Python), ensuring a fair comparison. survtmle is run with SuperLearner including GLM and XGBoost learners for both the outcome and censoring models.

B.2.2.4 TTR Failure Analysis for Survival

Training-time targeting (TTR) does not improve survival-curve estimation. Across all 500 replications with last-layer targeting, TTR performs at approximately chance level: about 51% of replications produce *worse* $|\text{Bias}|/\text{SD}$ than the untouched initial model, with a correlation of only 0.38 between Initial and TTR errors.

Hyperparameter grid search. To verify this is not merely a tuning problem, we conducted a systematic grid search over 7 configurations (50 replications each, last-layer targeting). Table B.1 shows that no configuration reliably improves over the initial model.

Table B.1: TTR hyperparameter grid search for survival (50 replications each, last-layer targeting). The initial model has $|\text{Bias}|/\text{SD} \approx 0.17$. Conservative settings produce coin-flip results; aggressive settings cause numerical divergence.

λ	Step factor	Frequency	TTR $ \text{Bias} /\text{SD}$	Notes
10^{-4}	10^{-2}	10	0.225	Default (baseline)
10^{-5}	10^{-2}	10	0.196	Best config (marginal)
10^{-6}	10^{-2}	10	0.229	No improvement
10^{-5}	5×10^{-2}	5	0.308	Worse
10^{-6}	5×10^{-2}	5	0.306	NaN in 2 reps
10^{-5}	10^{-1}	1	0.406	NaN in 12 reps
10^{-6}	10^{-1}	1	0.390	NaN in 11 reps

The best configuration ($\lambda = 10^{-5}$, step = 10^{-2} , frequency = 10) improves $|\text{Bias}|/\text{SD}$ from 0.225 to 0.196, but this improvement comes primarily from SD inflation (+8%), not bias reduction; MSE actually worsens.

Structural explanation. The failure reflects three structural differences from ATE targeting:

1. **Dimensionality:** ATE targets one scalar EIC equation; survival targets 50 simultaneous equations whose corrections can pull parameters in conflicting directions.
2. **Signal-to-noise ratio:** ATE bias is large relative to variability ($|\text{Bias}|/\text{SD} = 0.51$ for IHDP), giving TTR a strong, consistent gradient. Survival per-time-point biases are small relative to noise, drowning out the targeting signal.
3. **Gradient competition:** During training, the Poisson log-likelihood gradient competes with 50 targeting gradient directions. For ATE, the single targeting direction can coexist with the MSE gradient; for survival, the competition is 50-fold worse.

As a result, the Combined estimator (TTR + post-training TDA) reduces to TDA alone for survival: the TTR stage contributes no information that the post-training TDA step cannot already achieve. This motivates post-training TDA as the preferred approach for high-dimensional functional targets in the current setting.

Time grids. TTR during training uses a coarser grid of 25 time points for computational efficiency, while post-training TDA uses the full 50-point evaluation grid. The Combined estimator uses 25 points during Stage 1 (TTR) and 50 points during Stage 2 (TDA). All methods are evaluated on the full 50-point grid.

B.2.2.5 Visualization of Survival-Curve Accuracy

We aggregate results from the 500 Monte Carlo replications in two complementary visualizations.

Uncertainty bands for the survival function. For each of the 50 evaluation times $t \in [0, 14]$, we collect estimated survival probabilities, compute their empirical median, and derive a 95% percentile band (2.5th–97.5th percentiles) across replications. Figure B.1 juxtaposes the true survival curve with the median trajectories of the initial neural network, TDA, and Kaplan–Meier.

How often does the targeted estimator win? Figure B.2 reports the proportion of replications in which TDA’s absolute error is smaller than that of the initial neural network or the KM estimator at each time point. Beyond the early time points where all estimators are nearly unbiased, TDA wins in approximately 60–70% of replications.




Figure B.1: **Median survival curves with 95% percentile bands** under the concordant late acceleration DGP. Solid curves show the median survival estimate across 500 replications for the initial model (blue), TDA (green), and Kaplan–Meier (cyan). Shaded ribbons depict the 2.5th–97.5th percentiles. The black line is the true survival function. The KM estimator shows severe systematic bias due to concordant informative censoring.

B.3 TDA Extensions

B.3.1 Nested Working Models and Plateau Selection in TDA

In Chapter 3, nested working models arise from decreasing the lasso penalty λ : $\mathcal{M}_1 = \mathcal{M}_{\text{CV}} \subset \mathcal{M}_2 \subset \dots \subset \mathcal{M}_K$, where each model includes more HAL basis functions. In TDA, nested working models arise from expanding the targeting subset of weights:

$$\mathcal{M}_n^{(1)} \subset \mathcal{M}_n^{(2)} \subset \dots \subset \mathcal{M}_n^{(K)},$$

where each model unfreezes more layers. The plateau selector from Chapter 3 applies identically.

The EIC approximation condition (4.4) requires that the projection of $D(O)$ onto the score space yields a sufficiently accurate approximation. As argued in Section 4.2.3, the full network typically satisfies this condition without augmentation. The nested structure serves as a *regularization mechanism* for the targeting step, not as basis enrichment.

From Last-Layer to Partial Unfreezing. A common practical choice is to freeze all but the *final* layer. However, in settings with richer structure (multiple treatments, survival




Figure B.2: **How often does targeting help?** Percentage of 500 replications in which TDA has a smaller absolute error than the initial model (blue) and the KM estimator (teal) at each time point. The dashed line indicates parity (50%); values above mean TDA is more accurate in the majority of runs.

analysis, complex interactions), the final layer alone may not provide sufficient flexibility. In such cases, TDA can *selectively* unfreeze additional blocks of weights:

- **Gradient-based selection:** Unfreeze blocks θ_j with large $\|\nabla_{\theta_j} \Psi(\hat{\theta})\|$.
- **Second-order criteria:** Unfreeze blocks that yield large expected reduction in $|P_n[\tilde{D}^*(\theta)]|$ per unit norm.

Plateau-Based Selection. Following the plateau selector of Chapter 3, we select the working model complexity k^* data-adaptively. For each candidate $\mathcal{M}_n^{(k)}$, $k = 1, \dots, K$, run TDA to obtain the point estimate $\hat{\Psi}_k$, standard error $\widehat{se}_k$, and 95% CI $[L_k, U_k] = [\hat{\Psi}_k \pm z_{1-\alpha/2} \widehat{se}_k]$. The stopping rule is:

IF $\hat{\Psi}_k > \hat{\Psi}_{k-1}$ **and** $L_k < L_{k-1}$: **stop**, set $k^* = k - 1$.

ELSIF $\hat{\Psi}_k < \hat{\Psi}_{k-1}$ **and** $U_k > U_{k-1}$: **stop**, set $k^* = k - 1$.

ELSE continue to $k + 1$.

Algebraic justification. When the estimate increases ($\hat{\Psi}_k > \hat{\Psi}_{k-1}$) but the lower bound decreases ($L_k < L_{k-1}$), expanding the first condition gives $\hat{\Psi}_k - z_{1-\alpha/2} \widehat{se}_k < \hat{\Psi}_{k-1} - z_{1-\alpha/2} \widehat{se}_{k-1}$, which rearranges to

$$\hat{\Psi}_k - \hat{\Psi}_{k-1} < z_{1-\alpha/2} (\widehat{se}_k - \widehat{se}_{k-1}).$$

This shows explicitly that the algorithm stops when the standard-error increase (variance inflation) strictly exceeds the magnitude of the point-estimate shift (bias correction)—the hallmark of overfitting the targeting step.

Concrete Example: DragonNet Nested Subsets for ATE. For the DragonNet architecture with `hidden_outcome` = 64, we construct $K = 6$ nested subsets by progressively unfreezing rows of the penultimate outcome-head layer in 20% increments. The propensity head and shared trunk remain frozen:

k	Targeting parameters	Dim
1	Last output layer of each outcome head	128
2	$k=1$ + first 20% of penultimate layer	1,792
3	$k=1$ + first 40% of penultimate layer	3,456
4	$k=1$ + first 60% of penultimate layer	5,120
5	$k=1$ + first 80% of penultimate layer	6,784
6	Full outcome heads (all layers)	8,320

Note that regHAL traverses ~ 20 nested models via decreasing λ , whereas TDA typically has $K = 3\text{--}6$ natural layer-based subsets. For each subset, we also report two EIC diagnostics: the *non-parametric EIC mean* $P_n[D_{\text{NP}}^*]$ (empirical mean of the full doubly-robust influence function residual) and the *projected EIC mean* $P_n[D_{\text{proj}}^*]$ (empirical mean of the EIC projected onto the working model’s score space). After targeting converges, $P_n[D_{\text{proj}}^*] \approx 0$ by construction; the gap $|P_n[D_{\text{NP}}^*]| - |P_n[D_{\text{proj}}^*]|$ diagnoses whether the score space is too narrow to adequately approximate the full non-parametric EIC.

Empirical Demonstration (IHDP). Table B.2 reports the plateau selector applied to the DragonNet ATE experiment (Section 4.4.1) across 500 IHDP replications. For each replication, the plateau selector evaluates all $K = 6$ nested subsets; columns report per-subset averages across replications. The regularization is scaled as $\lambda_k = 0.01 \cdot \sqrt{d_k/d_{\text{base}}}$, where d_k is the number of targeted parameters in subset k and $d_{\text{base}} = 128$ is the base dimension; this ensures larger subsets receive proportionally more regularization (see Appendix B.2.1).

The plateau selector chose the largest subset $k^* = 6$ in 88.0% of replications, with the remaining selections spread across $k^* = 1$ (3.4%), $k^* = 2$ (3.6%), $k^* = 3$ (2.2%), $k^* = 4$ (1.6%), and $k^* = 5$ (1.2%). This outcome reflects the near-constant SE estimates (0.246–0.247) across all subsets: because expanding the targeting set introduces negligible variance inflation, the stopping rule—which triggers when variance growth outpaces bias correction—is never activated. The selector thus defaults to the largest working model, which is the designed conservative behavior.

Oracle coverage—95% CIs constructed from the true Monte Carlo SD—reveals that point estimation quality is uniformly excellent across all subsets: oracle coverage ranges from 95.4% ($k=1$) to 96.4% ($k=6$), confirming that every working model adequately controls bias. Among fixed subsets, $k=2$ (+20% penultimate) achieves the lowest MSE (0.057) and SD (0.205), while the full model $k=6$ has MSE 0.074 and SD 0.223. The modest differences confirm that the DragonNet outcome heads are well-specified for this problem: even the smallest subset ($k=1$, 128 parameters) already approximates the EIC direction adequately, and expanding does not degrade inference. Nominal coverage (91.6–92.0%) lags oracle coverage (95.4–96.4%) by 3–4 percentage points uniformly across subsets, reflecting the variance

Table B.2: Plateau selection across nested outcome-head subsets (IHDP, 500 replications). |Bias|: absolute bias; SD: Monte Carlo standard deviation; SE: average estimated standard error; MSE: mean squared error; Nom Cov: 95% CI coverage using estimated SE; Orc Cov: oracle coverage using true MC SD; $P_n D_{\text{NP}}^*$: non-parametric EIC mean; $P_n D_{\text{proj}}^*$: projected EIC mean.

k	Parameters	Dim	Bias	SD	SE	MSE	Nom Cov	Orc Cov	$P_n D_{\text{NP}}^*$	$P_n D_{\text{proj}}^*$
1	Last output	128	0.134	0.224	0.246	0.068	92.0%	95.4%	−0.001	−0.000
2	+20% penult.	1,792	0.127	0.205	0.246	0.057	92.0%	96.0%	−0.001	−0.000
3	+40% penult.	3,456	0.127	0.208	0.246	0.059	91.6%	96.2%	−0.000	−0.000
4	+60% penult.	5,120	0.129	0.217	0.247	0.062	92.0%	96.2%	−0.000	−0.001
5	+80% penult.	6,784	0.132	0.220	0.247	0.066	91.8%	96.0%	+0.001	+0.001
6	Full outcome	8,320	0.133	0.223	0.247	0.074	92.0%	96.4%	+0.001	+0.001
$\hat{k}^*$	Plateau-selected	—	0.133	0.218	0.246	0.074	92.0%	96.2%	—	—

estimation gap rather than any deficiency in the targeting step. This stability is precisely the regime where the plateau selector is unnecessary by design—it serves as a safety net that detects overfitting when present, rather than an oracle selector that finds the MSE-optimal subset.

The EIC diagnostics are informative: $|P_n D_{\text{proj}}^*| < 0.001$ at all levels, confirming that the model score equation is solved within each working model after targeting. The non-parametric EIC mean $P_n D_{\text{NP}}^*$ remains small throughout (magnitude ≤ 0.001), indicating that the score space is rich enough to well-approximate the full non-parametric EIC even at the smallest subset size.

B.3.2 TDA-Direct: A Streamlined Targeting Step for Special Cases (e.g., the ATE)

The general TDA algorithm (Sections 4.2–4.2.4) projects an estimated influence function onto the span of the score vectors $S_{\theta_{\text{targ}}} = -\nabla_{\theta_{\text{targ}}} \ell$. For many causal parameters, a closed-form *linear universal* least-favourable submodel is already known. A canonical example is the average treatment effect (ATE), whose update is driven by the *clever covariate* $H(A, X) = \frac{A}{g(X)} - \frac{1-A}{1-g(X)}$. In such settings we can bypass the projection step entirely and embed the fluctuation directly in the *last layer* of the outcome network. We denote this simplified variant **TDA-Direct**.

Algorithm. Assume the final layer of Q is linear: $Q_{\theta}(a, x) = \sum_{j=1}^M \theta_j^{\text{last}} \phi_j^{(\text{early})}(a, x)$.

- (1) **Initial fits.** Train a network for Q and a separate model for the propensity score g .
- (2) **Compute clever covariate.** Evaluate $\hat{H}_i = \hat{H}(A_i, X_i)$ for all observations.

- (3) **Linear regression in weight space.** Freeze all weights *except* θ^{last} and solve $\hat{\alpha} = \arg \min_{\alpha \in \mathbb{R}^M} \sum_{i=1}^n \left(\hat{H}_i - \alpha^\top \phi^{(\text{early})}(A_i, X_i) \right)^2$. The solution $\hat{\alpha}$ is our *targeting gradient* and we hope to overfit $\hat{H}_i$ as much as we can.
- (4) **One-step update.** Replace the last-layer weights by $\theta^{\text{last}} \leftarrow \theta^{\text{last}} + \varepsilon \hat{\alpha}$, where ε is found by a one-dimensional line search.

Expressiveness. The key requirement is that $\hat{H}$ lies (or at least approximately) in the span of the final-layer features. When this is doubtful, one can add *skip connections* from deeper hidden units to the output so that the final linear layer spans a richer function class while remaining linear in θ^{last} during targeting.

Such special cases with a universal path that is linear in clever covariates which only depends on the orthogonal nuisance parameters are prevalent (e.g., the ATE in the longitudinal case). **TDA-Direct** thus serves as a simpler and faster alternative to the general TDA discussed in the main paper. The general score-projection-based TDA algorithm applies broadly to any target parameter and loss function.

Appendix C

Supporting Material for Chapter 5

This appendix collects the proofs and supplementary diagnostics for Chapter 5. The main chapter keeps the grouped-jackknife construction, theorem statements, simultaneous-band methodology, and the simulations that most directly support the dissertation's targeted-estimation storyline.

C.1 Proofs for the main grouped-jackknife results

C.1.1 Proof of Lemma 5.1

Proof. Let $W_v = P_{n,v}^1 \varphi$ and $\bar{W} = V^{-1} \sum_{v=1}^V W_v$. Substituting the asymptotically linear expansion into the pseudo-value definition gives

$$IC_{Jack}(v) - \hat{\Psi}_{Jack}(P_n) = (W_v - \bar{W}) + \Delta_v,$$

where

$$\Delta_v = -(V-1) \left\{ R(P_{n,-v}, P_0) - \bar{R}_{n,-v} \right\}, \quad \bar{R}_{n,-v} = \frac{1}{V} \sum_{v'=1}^V R(P_{n,-v'}, P_0).$$

Under the stated assumptions, $\max_v |\Delta_v| = o_p(n^{-1/2})$ for fixed V . Expanding the centered empirical variance gives

$$\hat{S}_{Jack}^2 - S_W^2 = \frac{1}{V-1} \sum_{v=1}^V \left\{ \Delta_v^2 + 2(W_v - \bar{W})\Delta_v \right\}.$$

Because $W_v - \bar{W} = O_p(n^{-1/2})$ and $\max_v |\Delta_v| = o_p(n^{-1/2})$, the difference is $o_p(n^{-1})$. Since $S_W^2 \asymp_p V/n$ with fixed V , dividing by S_W^2 yields the claimed relative equivalence. $\square$

C.1.2 Proof of Theorem 5.2

Proof. Using the asymptotically linear expansion, the numerator satisfies

$$\sqrt{V}(\hat{\Psi}_{Jack}(P_n) - \Psi(P_0)) = \sqrt{V}\bar{W} + \sqrt{V}\left(VR(P_n, P_0) - (V-1)\bar{R}_{n,-v}\right).$$

Because V is fixed, the remainder assumptions imply that the second term is $o_p(n^{-1/2})$ and therefore becomes $o_p(1)$ after division by $\hat{S}_{Jack} \asymp_p \sqrt{V/n}$. By Lemma 5.1, $\hat{S}_{Jack} = S_W\{1 + o_p(1)\}$, so

$$\frac{\sqrt{V}(\hat{\Psi}_{Jack}(P_n) - \Psi(P_0))}{\hat{S}_{Jack}} = \frac{\sqrt{V}\bar{W}}{S_W} + o_p(1).$$

The disjoint folds imply that $(\sqrt{n/V}W_1, \dots, \sqrt{n/V}W_V)$ converges in distribution to V independent centered normal variables. Studentization across those V fold means therefore yields the t_{V-1} limit. $\square$

C.1.3 Proof of Lemma 5.3

Proof. By definition,

$$\hat{\Psi}_{Jack}(P_n) = V\hat{\Psi}(P_n) - (V-1)\bar{\Psi}_{(-v)},$$

so $\hat{\Psi}_{Jack}(P_n) - \hat{\Psi}(P_n) = -\hat{b}$. Substituting the asymptotically linear expansions for the full-sample and leave-fold-out estimators, the influence-function terms cancel exactly because each observation appears in exactly $V-1$ leave-fold-out samples. Hence

$$-\hat{b} = (V-1) \left\{ R(P_n, P_0) - \frac{1}{V} \sum_{v=1}^V R(P_{n,-v}, P_0) \right\}.$$

Under the fixed- V assumptions, each remainder term is $o_p(n^{-1/2})$, so $\hat{b} = o_p(n^{-1/2})$. Since $\widehat{SE} = \hat{S}_{Jack}/\sqrt{V} \asymp_p n^{-1/2}$ for fixed V , the result follows. $\square$

C.1.4 Proof of Corollary 5.4

Proof. Write

$$\frac{\sqrt{V}(\hat{\Psi}(P_n) - \Psi(P_0))}{\hat{S}_{Jack}} = \frac{\sqrt{V}(\hat{\Psi}_{Jack}(P_n) - \Psi(P_0))}{\hat{S}_{Jack}} + \frac{\sqrt{V}(\hat{\Psi}(P_n) - \hat{\Psi}_{Jack}(P_n))}{\hat{S}_{Jack}}.$$

The first term converges to t_{V-1} by Theorem 5.2. The second term equals $\hat{b}/\widehat{SE} = o_p(1)$ by Lemma 5.3. Slutsky's theorem gives the stated limit and confidence interval. $\square$

C.1.5 Proof of Theorem 5.6

Proof. Let $W_v = P_{n,v}^1 \varphi$ and $\bar{W} = V^{-1} \sum_{v=1}^V W_v$. As in the fixed- V argument,

$$IC_{Jack}(v) - \hat{\Psi}_{Jack}(P_n) = (W_v - \bar{W}) + \Delta_v,$$

with Δ_v determined by the fold-level remainder differences. Under the stated stability condition,

$$\sum_{v=1}^V \Delta_v^2 \leq (V-1)^2 \|R(P_n, P_0) - R(P_{n,-v}, P_0)\|_2^2 = o_p(n^{-1}).$$

Multiplying by n/V shows that the remainder contribution to $\hat{\sigma}^2 = (n/V) \hat{S}_{Jack}^2$ is $o_p(V^{-1})$; the cross-term is also $o_p(V^{-1})$ by Cauchy–Schwarz. Thus

$$\hat{\sigma}^2 = \frac{n}{V} \hat{S}_W^2 + o_p(V^{-1}).$$

Expanding $\hat{S}_W^2$ exactly yields

$$\frac{n}{V} \hat{S}_W^2 = \frac{1}{n} \sum_{i=1}^n \varphi_i^2 + T_n + \frac{1}{V-1} T_n - \frac{1}{V-1} U_n,$$

where $T_n = \frac{1}{n} \sum_{v=1}^V \sum_{i \neq j \in I_v} \varphi_i \varphi_j$ and $U_n = \frac{1}{n} \sum_{i \neq j} \varphi_i \varphi_j$. The diagonal term converges to σ^2 by the law of large numbers, $U_n = O_p(1)$, and $\text{Var}(T_n) = O(V^{-1})$, so $T_n = O_p(V^{-1/2})$. Therefore

$$\hat{\sigma}^2 - \sigma^2 = \left(\frac{1}{n} \sum_{i=1}^n \varphi_i^2 - \sigma^2 \right) + T_n + o_p(V^{-1/2}),$$

which implies consistency when $V \rightarrow \infty$ and root- V consistency under $V = o(n)$. $\square$

C.1.6 Proof of Theorem 5.7

Proof. For component j , write

$$IC_{Jack}^{(j)}(v) - \hat{\Psi}_{Jack}^{(j)}(P_n) = (W_v^{(j)} - \bar{W}^{(j)}) + \Delta_v^{(j)},$$

where $\Delta_v^{(j)} = -(V-1) \{R_j(P_{n,-v}, P_0) - V^{-1} \sum_{v'} R_j(P_{n,-v'}, P_0)\}$. Under the fixed- V remainder assumption, $\max_v |\Delta_v^{(j)}| = o_p(n^{-1/2})$. Expanding the cross-products that define $\hat{\Sigma}_{Jack}(j, k)$ gives

$$\hat{\Sigma}_{Jack}(j, k) = \hat{S}_W^2(j, k) + A_{jk} + A_{kj} + B_{jk},$$

where A_{jk} is a cross-term between $W^{(j)}$ and $\Delta^{(k)}$, and B_{jk} is quadratic in the remainders. The quadratic term is $o_p(n^{-1})$ immediately. The cross-term is also $o_p(n^{-1})$ by Cauchy–Schwarz because $\hat{S}_W(j, j)^{1/2} = O_p(n^{-1/2})$ and the remainder factor is $o_p(n^{-1/2})$. Multiplying by n/V yields the claimed asymptotic equivalence. For exactly equal folds, the oracle covariance estimator is unbiased because the fold means are i.i.d. with covariance $(V/n)\sigma_{jk}$. $\square$

C.1.7 Proof of Theorem 5.8

Proof. Let $\mathbf{W}_v = (P_{n,v}^1 \varphi_1, \dots, P_{n,v}^1 \varphi_m)'$. Because the folds are disjoint, the $\mathbf{W}_v$ are independent, and each is an average over n/V observations. The multivariate central limit theorem therefore implies

$$\sqrt{n/V} (\mathbf{W}_1, \dots, \mathbf{W}_V) \implies_d (\tilde{\mathbf{Z}}_1, \dots, \tilde{\mathbf{Z}}_V),$$

with $\tilde{\mathbf{Z}}_v \stackrel{i.i.d.}{\sim} \mathcal{N}_m(0, \Sigma)$. For each component, the same remainder argument used in Theorem 5.2 shows that the jackknife-centered estimator equals $\bar{W}^{(j)} + o_p(n^{-1/2})$, and Theorem 5.7 gives $\hat{\sigma}_j^2 = \hat{S}_W^2(j, j) + o_p(n^{-1})$. After scaling by $\sqrt{n/V}$ and applying Slutsky's theorem, the componentwise-Studentized vector converges to the version built from $\tilde{\mathbf{Z}}_1, \dots, \tilde{\mathbf{Z}}_V$. Rescaling each component by its marginal standard deviation converts that limit to the correlation-matrix representation in (5.6). Each marginal is therefore exactly t_{V-1} . $\square$

C.1.8 Proof of Theorem 5.9

Proof. Write $\hat{\Psi}_j(P_n) - \Psi_j(P_0) = (\hat{\Psi}_j(P_n) - \mathbb{E}[\hat{\Psi}_j(P_n)]) + b_j$. By the triangle inequality,

$$|\hat{\Psi}_j(P_n) - \Psi_j(P_0)| \leq |\hat{\Psi}_j(P_n) - \mathbb{E}[\hat{\Psi}_j(P_n)]| + |b_j|.$$

The reverse triangle inequality gives $|b_j| \leq |\hat{b}_j| + |\hat{b}_j - b_j|$, so it suffices that

$$|\hat{\Psi}_j(P_n) - \mathbb{E}[\hat{\Psi}_j(P_n)]| + |\hat{b}_j - b_j| \leq q_\alpha \widehat{SE}_j \quad \forall j.$$

Let $\varepsilon_n = \max_j |\hat{b}_j - b_j| / \widehat{SE}_j = o_p(1)$. Then the sufficient condition above is implied by $\max_j |T_j^*| \leq q_\alpha - \varepsilon_n$, where $T_j^* = (\hat{\Psi}_j(P_n) - \mathbb{E}[\hat{\Psi}_j(P_n)]) / \widehat{SE}_j$. Theorem 5.8, together with the identity $\hat{\Psi}_{Jack}^{(j)}(P_n) - \hat{\Psi}_j(P_n) = -\hat{b}_j$, shows that $T_j^* = T_j + o_p(1)$ uniformly over j . Therefore $\max_j |T_j^*| + \varepsilon_n$ converges to $\max_j |T_j^\infty|$, whose $(1 - \alpha)$ quantile is q_α . Continuity of the limit distribution gives the stated simultaneous coverage. $\square$

C.1.9 Proof of Theorem 5.12

Proof. Define $W_{n,v} = P_{n,v}^1 \varphi_n$ and $\bar{W}_n = V^{-1} \sum_{v=1}^V W_{n,v}$. Exact algebraic cancellation again yields

$$IC_{Jack}(v) - \hat{\Psi}_{Jack}(P_n) = (W_{n,v} - \bar{W}_n) + \Delta_{n,v},$$

where $\max_v |\Delta_{n,v}| = o_p(\sigma_n n^{-1/2})$ by Assumption 5.10. Since $\text{Var}(W_{n,v}) = V\sigma_n^2/n$, the same variance expansion as before gives $\hat{S}_{Jack}^2 = S_{W_n}^2 \{1 + o_p(1)\}$. Writing the Studentized statistic and dividing numerator and denominator by the unknown scale $\sigma_n \sqrt{V/n}$ leaves only the normalized fold averages

$$Z_{n,v} = \frac{\sqrt{n/V}}{\sigma_n} W_{n,v}.$$

The Lindeberg condition implies $(Z_{n,1}, \dots, Z_{n,V})$ converges to V independent standard normals. Studentization therefore again yields the t_{V-1} limit, with the diverging rate canceled by scale invariance. $\square$

C.1.10 Proof of Corollary 5.13

Proof. As in the fixed-rate case, $\hat{\Psi}_{Jack}(P_n) - \hat{\Psi}(P_n) = -\hat{b}_n$ with

$$\hat{b}_n = (V-1) \left\{ R_n(P_n, P_0) - \frac{1}{V} \sum_{v=1}^V R_n(P_{n,-v}, P_0) \right\}.$$

Assumption 5.10 makes this difference $o_p(\sigma_n n^{-1/2})$, while $\widehat{SE} = \hat{S}_{Jack}/\sqrt{V} \asymp_p \sigma_n n^{-1/2}$. Hence the centering mismatch is $o_p(\widehat{SE})$, and Slutsky's theorem transfers the t_{V-1} limit from Theorem 5.12 to the full-sample estimator. $\square$

C.1.11 Proof of Theorem 5.14

Proof. Following the exact decomposition from Theorem 5.6,

$$\frac{\hat{S}_n^2}{\sigma_n^2} = \frac{n}{V\sigma_n^2} \hat{S}_{W_n}^2 + \frac{n}{V(V-1)\sigma_n^2} \sum_{v=1}^V \Delta_{n,v}^2 + \text{cross-term}.$$

The stated remainder-stability condition makes the quadratic remainder term $o_p(V^{-1})$ after scaling, and the cross-term is $o_p(V^{-1/2})$ by Cauchy–Schwarz. Thus it is enough to analyze $(n/(V\sigma_n^2))\hat{S}_{W_n}^2$. Expanding that sample variance yields

$$\frac{n}{V\sigma_n^2} \hat{S}_{W_n}^2 = \frac{1}{n\sigma_n^2} \sum_{i=1}^n \varphi_n^2(O_i) + \frac{T_n}{\sigma_n^2} + \frac{1}{V-1} \frac{T_n}{\sigma_n^2} - \frac{1}{V-1} \frac{U_n}{\sigma_n^2},$$

where T_n and U_n are the same foldwise and global cross-product sums as in the fixed-rate theorem. The weak law of large numbers gives convergence of the first term to 1. The variance calculation gives $\mathbb{V}\text{ar}(T_n/\sigma_n^2) \sim 2/V$, so $T_n/\sigma_n^2 = o_p(1)$ when $V \rightarrow \infty$, and the $(V-1)^{-1}$ terms vanish as well. Hence $\hat{S}_n^2/\sigma_n^2 \xrightarrow{p} 1$. $\square$

C.2 Connection to High-Dimensional Gaussian Approximation

This appendix records a broader connection between the simultaneous-inference problem studied here and the literature on Gaussian and multiplier-bootstrap approximation for maxima of high-dimensional means. These results are not needed for the main theorems of the chapter, but they help situate the grouped-jackknife construction relative to existing high-dimensional approximation methods.

The fold-level pseudo-value vectors provide a natural bridge to the Gaussian multiplier bootstrap framework of Chernozhukov et al. [2013]. By Theorem 5.8, the m -dimensional centered pseudo-value vectors $(IC_{Jack}^{(1)}(v) - \hat{\Psi}_{Jack}^{(1)}(P_n), \dots, IC_{Jack}^{(m)}(v) - \hat{\Psi}_{Jack}^{(m)}(P_n))'$, $v = 1, \dots, V$, behave asymptotically as V independent draws from a mean-zero distribution with covariance proportional to Σ . This places the problem of constructing simultaneous inference for the mean vector $\Psi(P_0)$ squarely within the framework of the Gaussian multiplier bootstrap for high-dimensional means.

In the standard formulation of Chernozhukov et al. [2013], given n independent centered observations $\mathbf{X}_1, \dots, \mathbf{X}_n \in \mathbb{R}^p$, the Gaussian multiplier bootstrap approximates the distribution of the max statistic $\max_{1 \leq j \leq p} |n^{-1/2} \sum_{i=1}^n X_{ij}|$ as follows: for $b = 1, \dots, B$, generate independent standard normal multipliers $\xi_1^{(b)}, \dots, \xi_n^{(b)} \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1)$, compute the multiplier process $\max_{1 \leq j \leq p} |n^{-1/2} \sum_{i=1}^n \xi_i^{(b)} X_{ij}|$, and use the empirical distribution of these B maxima to estimate quantiles. The key theoretical result is that this bootstrap consistently estimates the distribution of the max statistic even when p greatly exceeds n , with no rank restriction on the covariance matrix, and the approximation error converges to zero uniformly over hyperrectangles even when p grows as fast as $\exp(Cn^c)$ for certain constants $C, c > 0$.

Chen and Kato [2020] develop a “jackknife multiplier bootstrap” (JMB) for U -processes of general order $r \geq 2$ indexed by VC-type function classes. In their setting, the standard Gaussian multiplier bootstrap is infeasible because it requires the Hájek projection $P^{r-1}h$, which depends on the unknown distribution P . Their key idea is to replace each $P^{r-1}h(X_i)$ with its leave-one-out jackknife estimate, and they establish non-asymptotic coupling and Kolmogorov distance bounds showing that the resulting bootstrap consistently approximates the distribution of the U -process supremum with polynomial error rates. While this is the closest existing combination of jackknife ideas with simultaneous inference, it differs from the present approach in several respects: (i) it uses the leave-one-out (not grouped) jackknife, where each deletion removes a single observation; (ii) the jackknife serves to estimate an unknown nuisance quantity (the Hájek projection) that is then fed into an external Gaussian multiplier scheme, rather than producing pseudo-values that are directly Studentized; (iii) it targets suprema over infinite-dimensional function classes rather than finite-dimensional parameter vectors; and (iv) the theoretical guarantees are non-asymptotic finite-sample bounds for U -processes, whereas the results here provide asymptotic distributional limits for general regular asymptotically linear estimators of pathwise differentiable parameters.

In the present setting, V plays the role of n and m (the number of evaluation points) plays the role of p . One could, in principle, apply the Gaussian multiplier bootstrap directly to the V centered pseudo-value vectors to approximate the distribution of $\max_{1 \leq j \leq m} |\sqrt{V}(\hat{\Psi}_{Jack}^{(j)}(P_n) - \Psi_j(P_0))/\hat{\sigma}_j|$. However, the theoretical guarantees of Chernozhukov et al. [2013] require the sample size to satisfy $n \gg (\log p)^c$ for certain $c > 0$, which translates to $V \gg (\log m)^c$ in this framework. Since typical choices of $V \in \{10, 20\}$ are modest, this high-dimensional rate condition may not be satisfied when m is large, and the multiplier-bootstrap approximation may be unreliable in this regime. This observation helps explain why the main text instead focuses on the exact componentwise-Studentized limit together with slow divergence of V , while viewing the high-dimensional Gaussian

approximation as contextual rather than essential.

C.3 True Dose-Response Curves for the HAL Simulation

Figure C.1 displays the true dose-response curves $\psi_0(a) = \mathbb{E}_{P_0}[\mathbb{E}_{P_0}(Y \mid A = a, W)]$ for the four data generating processes used in the HAL simulation study (Section 5.6.3). DGP 1 is a smooth, monotonically increasing curve. DGP 2 exhibits pronounced oscillations driven by a $\sin(1.25a^{3/2})$ term. DGP 3 is smooth and nearly linear for $a \leq 2$, then transitions to an oscillating pattern for $a > 2$ via an indicator-modulated sine component. The transition at $a = 2$ is continuous (since $\sin((0.8 \cdot 2)^2 - 2.56) = \sin(0) = 0$), but the derivative is discontinuous—a kink that the first-order HAL can represent, though the subsequent oscillations remain challenging for a piecewise-linear basis. DGP 4 features jump discontinuities at $a = 2$ and $a = 4$, where the curve jumps from approximately 0.12 to 0.50 and from 0.84 to 0.16, respectively. These four settings represent a progression from easy (smooth) to challenging (discontinuous) for nonparametric curve estimation, and the coverage results in Section 5.6.3 should be interpreted in light of the varying smoothness of the underlying truth.

C.4 Pointwise Coverage of Simultaneous Confidence Bands

The simultaneous confidence bands reported in Tables 5.3 and 5.5 are designed to cover the entire curve simultaneously with probability $1 - \alpha$. A natural diagnostic is to examine the *pointwise* coverage of these simultaneous bands at each individual evaluation point. By construction, pointwise coverage must be at least as large as simultaneous coverage, since the event that the band covers every point implies coverage at each individual point. Here we display pointwise coverage profiles for the Kaplan–Meier and HAL simulation studies to examine how the choice of V affects the uniformity of coverage across evaluation points.

Figure C.2 displays pointwise coverage for the Kaplan–Meier simultaneous bands. All methods— V -fold with $V \in \{5, 10, 20\}$ and the studentized bootstrap—achieve approximately 99–100% pointwise coverage uniformly across the entire time range. The coverage profiles are nearly indistinguishable across V , showing no additional fluctuation for small V . This uniformity reflects the conservatism inherent in simultaneous bands: the simultaneous critical value q_α is substantially larger than the pointwise t_{V-1} quantile, yielding pointwise coverage well above the nominal 95% level at every evaluation point.

Figure C.3 shows the corresponding pointwise coverage profiles for the HAL dose-response simultaneous bands. For DGPs 1 (smooth) and 2 (oscillating), pointwise coverage is uniformly high (approximately 95–100%) across all dose points for both CV-HAL and under-smoothed HAL, and the bias-corrected bands achieve even more uniform coverage. For DGPs 3 (piecewise) and 4 (discontinuous), coverage dips at specific dose points where the




Figure C.1: True dose-response curves $\psi_0(a) = \mathbb{E}[Y \mid \text{do}(A = a)]$ for the four data generating processes used in the HAL simulation study. DGP 1: smooth monotone increase; DGP 2: oscillating; DGP 3: smooth for $a \leq 2$, oscillating for $a > 2$; DGP 4: jump discontinuities at $a = 2$ and $a = 4$. True values computed by Monte Carlo integration (10^6 samples per evaluation point).




Figure C.2: Pointwise coverage (%) of simultaneous 95% confidence bands for the Kaplan–Meier survival function across evaluation times ($n = 200, 500$ replications). Methods: V -fold simultaneous band with $V \in \{5, 10, 20\}$ and the studentized bootstrap simultaneous band. Horizontal dashed line: nominal 95% level. All methods achieve 98–100% pointwise coverage uniformly across time points, with no systematic differences across V .




Figure C.3: Pointwise coverage (%) of simultaneous 95% confidence bands for the HAL dose-response curve across the dose grid ($n = 500$, 500 replications). Rows: DGP 1 (smooth), DGP 2 (oscillating), DGP 3 (piecewise), DGP 4 (discontinuous). Columns: CV-HAL and undersmoothed HAL. Solid lines: standard simultaneous bands; dashed lines: bias-corrected simultaneous bands. Colors indicate $V \in \{5, 10, 20\}$. Horizontal dashed line: nominal 95% level.

true curve is non-smooth, but these dips are consistent across all values of V —the coverage profiles for $V = 5$, $V = 10$, and $V = 20$ are nearly superimposed.

A striking feature of these figures is that small V does not produce more variable pointwise coverage across evaluation points compared to large V . The coverage profiles are remarkably stable across V . This stability arises because the simultaneous critical value q_α compensates for the additional uncertainty associated with fewer folds: smaller V yields larger t_{V-1} -based critical values, which uniformly inflate the band width at every evaluation point, producing stable (if conservative) pointwise coverage. The variation in pointwise coverage across evaluation points is driven primarily by the estimator’s bias profile—which is a property of the estimator and the data generating process—rather than by the choice of V .

This behavior contrasts with pointwise confidence intervals (cf. Figure 5.3), where different methods (delta method, V -fold t , bootstrap) produce visibly different coverage profiles across dose points. The simultaneous bands homogenize coverage across evaluation points because the simultaneous critical value acts as a uniform multiplier of the standard error, so the pointwise coverage at each point is determined primarily by the ratio of the estimator’s bias to its standard error at that point.

C.5 Jackknife-Centered versus Bias-Corrected Confidence Intervals

The standard V -fold jackknife confidence intervals center at the full-sample estimator $\hat{\Psi}_j(P_n)$, while the bias-corrected band from section 5.4.6 widens each interval by $|\hat{b}_j|$ to accommodate bias. A natural alternative is to *shift* the center to the jackknife estimator $\hat{\Psi}_{Jack}^{(j)}(P_n) = \hat{\Psi}_j(P_n) - \hat{b}_j$, which subtracts the estimated bias while keeping the interval width unchanged. The resulting “jackknife-centered” (JC) interval is

$$\hat{\Psi}_{Jack}^{(j)}(P_n) \pm \text{crit} \cdot \widehat{SE}_j = \left(\hat{\Psi}_j(P_n) - \hat{b}_j \right) \pm \text{crit} \cdot \widehat{SE}_j,$$

where $\text{crit} = t_{V-1, 1-\alpha/2}$ for pointwise or $\hat{q}_\alpha$ for simultaneous bands. Because JC has the same width as the standard interval (and narrower than the BC band), it would be more efficient if the bias estimate $\hat{b}_j$ were accurate.

Figure C.4 compares pointwise coverage across the dose grid for the three centering strategies (standard, JC, and BC) using $V = 20$. Contrary to the expectation that bias correction via centering might help, the JC intervals exhibit substantially *lower* coverage than the standard intervals in every DGP and HAL variant. For DGP 1 (smooth), pointwise coverage drops from 93.7% to 86.0% (CV-HAL) and from 95.6% to 85.7% (undersmoothed). For DGP 2 (oscillating), the drop is more severe: from 95.2% to 71.9% (CV-HAL). For DGPs 3 and 4, where bias is already problematic, JC coverage falls to 64–72%. The simultaneous results are even more striking (Figure C.5): for DGP 2, simultaneous coverage drops from 91.8% to 21.0% (CV-HAL), and for DGP 4 from 25.4% to 5.8%. In contrast, the BC approach consistently achieves the highest coverage across all settings, reaching near-nominal levels for DGPs 1–2.




Figure C.4: Pointwise coverage (%) of 95% confidence intervals for the HAL dose-response curve, comparing three centering strategies ($V = 20$, $n = 500$, 500 replications). Standard: centered at $\hat{\Psi}_j(P_n)$; Jack-centered (JC): centered at $\hat{\Psi}_{Jack}^{(j)}(P_n) = \hat{\Psi}_j(P_n) - \hat{b}_j$; Bias-corrected (BC): centered at $\hat{\Psi}_j(P_n)$ with interval widened by $|\hat{b}_j|$. The JC intervals have lower coverage than the standard intervals in all settings, while BC achieves the highest coverage.




Figure C.5: Pointwise coverage (%) of simultaneous 95% confidence bands for the HAL dose-response curve, comparing three centering strategies ($V = 20$, $n = 500$, 500 replications). Same conventions as Figure C.4 but using the simultaneous critical value $\hat{q}_\alpha$ in place of $t_{V-1, 0.975}$. The coverage degradation from JC centering is even more pronounced for simultaneous bands.

The failure of JC centering can be understood through the lens of Theorem 5.3. For pathwise differentiable parameters, the centering mismatch $\hat{b}_j = \hat{\Psi}_j(P_n) - \hat{\Psi}_{Jack}^{(j)}(P_n)$ satisfies $\hat{b}_j = o_p(\widehat{SE}_j)$, so the JC and standard intervals are asymptotically equivalent. However, for the HAL dose-response estimator with slower convergence rate $\sqrt{n/J_n}$, the bias b_j can be of the same order as $\widehat{SE}_j$ or larger, as Figure 5.4 confirms. The jackknife bias estimate $\hat{b}_j$ is computed from only $V = 20$ pseudo-values and is therefore noisy: its estimation error $\hat{b}_j - b_j$ is not negligible relative to $\widehat{SE}_j$. Shifting the center by $\hat{b}_j$ introduces a first-order noise term that inflates the effective variance of the centered estimator without reliably reducing its bias. In contrast, the BC approach widens the interval by $|\hat{b}_j|$, which is valid by the triangle inequality argument in Theorem 5.9: the widening accommodates bias regardless of the sign of the estimation error $\hat{b}_j - b_j$. These results confirm the paper’s recommendation to center confidence intervals at the full-sample estimator $\hat{\Psi}_j(P_n)$ and address bias through widening (BC) rather than shifting (JC).

C.6 Diagnostic: Highlighting Unreliable Grid Points in a Simultaneous Band

The simultaneous confidence band (5.7) presumes that the asymptotically linear expansion is equally reliable at every grid point j . In finite samples this may fail: at some j the remainder $R_{n,j}$ can be large relative to the standard error, so the pointwise interval — and hence its contribution to the simultaneous band — is suspect. This is a property of the underlying *estimator*, not of the V -fold construction; it would be present under any inferential procedure that uses the same estimator.

Because the jackknife already produces a fold-level bias estimate $\hat{b}_j = (V - 1)(\bar{\Psi}_{j,(-v)} - \hat{\Psi}_j(P_n))$ (see (5.13)) and a fold-level standard error $\widehat{SE}_j$, the ratio

$$\rho_j = \frac{|\hat{b}_j|}{\widehat{SE}_j}$$

is computed at no extra cost and serves as a natural point-level reliability diagnostic: large ρ_j indicates grid points where the linearization is under stress, and small ρ_j indicates well-targeted points.

We therefore recommend plotting the simultaneous band with a *color overlay* proportional to ρ_j — for example, shading grid points with $\rho_j \geq \tau$ (e.g. $\tau = 1$) to warn the reader that coverage may be unreliable there. This is a visualization aid, not a modification of the simultaneous band: the band (5.7) remains defined on the full grid $\{1, \dots, m\}$, and the overall $(1 - \alpha)$ -simultaneous coverage guarantee is unaffected. Analysts who prefer to report a confidence band over a restricted set — e.g. to describe the dose-response curve only at well-targeted doses — may use the flagged points as an *informal* data-adaptive restriction, with the understanding that post-selection inference issues then apply and the simultaneous coverage over the selected set is not formally controlled at level $1 - \alpha$.

The bias-corrected bands from section 5.4.6 and this diagnostic serve complementary roles: the former widens intervals by $|\hat{b}_j|$ to provide asymptotic coverage when $\hat{b}_j$ estimates b_j well, while the latter transparently communicates to the user where the linearization assumption is most strained.

C.7 Isolating the Effect of Correlation Matrix Estimation

In the simultaneous confidence band simulations of Section 5.6.2, changing V simultaneously alters three components of the inference procedure: (i) the quality of the estimated correlation matrix $\hat{R}_{\text{Jack}}$ (more folds yield a higher-rank, more precise estimate), (ii) the degrees of freedom $\nu = V - 1$ in the t -distribution used for the Monte Carlo critical value, and (iii) the standard error estimates. This confounding makes it difficult to attribute coverage differences across V to any single factor. To disentangle these effects, we conduct a supplementary simulation that *isolates* the correlation matrix by fixing the other two components.

Specifically, we fix the degrees of freedom at $\nu = 19$ (corresponding to $V = 20$), use the standard errors from a $V = 20$ jackknife, and center all confidence bands at the full-data Kaplan–Meier estimate. The only quantity that varies is the source of the correlation matrix: $\hat{R}_{\text{Jack}}$ is estimated from $V_{\text{corr}} \in \{5, 10, 20\}$ folds. The simultaneous confidence band is then

$$\hat{S}(t) \pm \hat{q}_\alpha(\hat{R}_{V_{\text{corr}}}, \nu = 19) \cdot \widehat{SE}_{20}(t),$$

where $\hat{q}_\alpha(\hat{R}_{V_{\text{corr}}}, \nu = 19)$ denotes the Monte Carlo critical value computed from Algorithm 1 using $\hat{R}_{V_{\text{corr}}}$ as the correlation matrix and $\nu = 19$ degrees of freedom. The data generating process, sample size ($n = 200$), evaluation grid ($m = 91$ points on $[2, 20]$), and number of replications (500) are identical to Section 5.6.2. We note that the “Standard” and “Isolation” rows in Table C.1 come from separate simulation runs with independent fold assignments, though both use the same data-generating mechanism and random seed for dataset generation.

Table C.1 reveals that the correlation matrix quality and the t -distribution tail behavior act in *opposite directions* and partially cancel. In the isolation rows, where only the correlation matrix varies, the coverage gap between $V_{\text{corr}} = 5$ and $V_{\text{corr}} = 20$ is 2.4 percentage points (92.8% versus 95.2%): the rank-deficient $\hat{R}$ from $V_{\text{corr}} = 5$ underestimates the $\max |T_j^\infty|$ quantile, yielding a smaller critical value ($\bar{q}_\alpha = 3.11$ versus 3.28) and narrower bands. In the standard simulation, however, the total gap between $V = 5$ and $V = 20$ is only 1.0 percentage point (93.6% versus 94.6%). The difference is explained by the heavy tails of t_4 : for the standard $V = 5$ procedure, the t_4 -distributed Monte Carlo draws inflate the critical value to $\bar{q}_\alpha = 6.52$ —more than double the isolation value—producing much wider bands (0.454 versus 0.229) that compensate for the imprecise $\hat{R}$. Thus, the heavier t_{V-1} tails for small V serve as a natural safeguard, inflating the critical value enough to partially offset the less precise correlation matrix estimate. The estimated effective rank $\hat{r}(\hat{R}) \approx 1.6$ –1.7 is stable across all settings, confirming the low intrinsic dimensionality of the Kaplan–Meier

Table C.1: Decomposing the effect of V on simultaneous coverage for the Kaplan–Meier survival function ($n = 200$, $m = 91$, 500 replications). “Standard” rows use V folds for all components (correlation matrix $\hat{R}_{\text{Jack}}$, degrees of freedom $\nu = V-1$, and standard errors). “Isolation” rows fix $\nu = 19$ and SE from $V=20$, varying only the correlation matrix source $\hat{R}_{\text{Jack}}$ estimated from V_{corr} folds.

Setting	V	ν	Simult. Cov. (%)	Mean q_α	Mean Width	$\hat{r}(\hat{R})$
Standard	5	4	93.6	6.52	0.454	1.6
Standard	10	9	93.4	3.91	0.282	1.7
Standard	20	19	94.6	3.28	0.239	1.7
Isolation	5	19	92.8	3.11	0.229	1.6
Isolation	10	19	93.8	3.23	0.237	1.7
Isolation	20	19	95.2	3.28	0.240	1.7

correlation structure discussed in section 5.4.4. Even $V = 5$, which yields a correlation matrix of rank at most 4, far exceeds this effective rank, and the resulting simultaneous coverage remains close to nominal.

C.8 Derivation of the Growth-Rate Condition on V

The diverging- V results (Theorems 5.6 and 5.14) require the fold-level remainder difference $R(P_n, P_0) - R(P_{n,-v}, P_0)$ to vanish strictly faster than $1/(\sqrt{n}V)$ in the standard regime, or faster than $\sigma_n/(\sqrt{n}V)$ in the slower-rate regime. This appendix derives an explicit upper growth rate on V under which those conditions hold, assuming the standard second-order von Mises remainder shape $R(P, P_0) \approx \int (p - p_0)^2$ shared by TMLE, one-step, AIPW, and HAL plug-in estimators. The derivation treats the leave-fold-out sample P_{n_v} , consisting of $n_v = n(V-1)/V$ observations, as the full-sample empirical measure with one fold of size n/V removed.

G.1 Bounding the leading linear difference

Writing the difference between the full-sample and leave-fold-out empirical averages of the influence function $\varphi = D_{P_0}^*$ gives

$$\begin{aligned}
P_n \varphi - P_{n_v} \varphi &= \left(\frac{1}{n} - \frac{1}{n_v} \right) \sum_{i=1}^{n_v} \varphi(O_i) + \frac{1}{n} \sum_{i=n_v+1}^n \varphi(O_i) \\
&= \underbrace{\frac{n_v - n}{n} \cdot \frac{1}{n_v} \sum_{i=1}^{n_v} \varphi(O_i)}_A + \underbrace{\frac{1}{V} \cdot \frac{1}{n/V} \sum_{i=n_v+1}^n \varphi(O_i)}_B.
\end{aligned}$$

Because $(n_v - n)/n = -1/V$, direct variance calculations yield

$$\mathbb{V}\text{ar}(A) = \frac{\sigma^2}{nV(V-1)} \asymp \frac{\sigma^2}{nV^2}, \quad \mathbb{V}\text{ar}(B) = \frac{1}{V^2} \cdot \frac{\sigma^2}{n/V} = \frac{\sigma^2}{nV}.$$

Since $\mathbb{V}\text{ar}(B)$ dominates $\mathbb{V}\text{ar}(A)$ by a factor of V , we obtain

Lemma C.1 (Linear-term rate). *Under $\mathbb{E}[\varphi^2] = \sigma^2 < \infty$,*

$$P_n \varphi - P_{n_v} \varphi = O_p \left(\frac{\sigma}{\sqrt{n} \sqrt{V}} \right).$$

G.2 Bounding the remainder difference

Under the standard second-order remainder shape $R(P, P_0) \approx \int (p - p_0)^2$, write the density-ratio difference $p_n - p_0 = P_n D_n + r(n)$ for a score-like term D_n and a higher-order residual $r(n) = o_p(n^{-1/2})$. Expanding $R(P_{n_v}, P_0) - R(P_n, P_0)$ gives

$$\begin{aligned} R(P_{n_v}, P_0) - R(P_n, P_0) &= \int (P_{n_v} D_n)^2 + 2 P_{n_v} D_n r(n_v) + r(n_v)^2 \\ &\quad - (P_n D_n)^2 - 2 P_n D_n r(n) - r(n)^2 \\ &= \underbrace{\int (P_{n_v} D_n - P_n D_n)(P_{n_v} D_n + P_n D_n)}_{A_1} + \underbrace{\int (r(n_v) - r(n))(r(n_v) + r(n))}_{B_1} \\ &\quad + \underbrace{2 \int (r(n_v) - r(n)) P_{n_v} D_n}_{C_1} + \underbrace{2 \int r(n) (P_{n_v} D_n - P_n D_n)}_{C_2}. \end{aligned}$$

The A_1 term dominates. By Lemma C.1 applied to D_n (with population variance scale σ_n^2),

$$P_{n_v} D_n - P_n D_n = O_p \left(\frac{\sigma_n}{\sqrt{n} \sqrt{V}} \right), \quad P_{n_v} D_n + P_n D_n = O_p \left(\frac{\sigma_n}{\sqrt{n}} \right),$$

so $A_1 = O_p(\sigma_n^2/(n\sqrt{V}))$. The terms B_1 , C_1 , and C_2 are of smaller order because each contains at least one factor of the higher-order residual $r(\cdot) = o_p(n^{-1/2})$. We conclude:

Lemma C.2 (Remainder-difference rate). *Under $R(P, P_0) \approx \int (p - p_0)^2$,*

$$\|R(P_n, P_0) - R(P_{n_v}, P_0)\|_2 = O_p \left(\frac{\sigma_n^2}{n\sqrt{V}} \right).$$

G.3 Translating into a growth rate on V

Matching the rate of Lemma C.2 against the stability requirement of Theorems 5.6 and 5.14 gives

$$\frac{\sigma_n^2}{n\sqrt{V}} = o\left(\frac{1}{\sqrt{n}V}\right) \iff \sqrt{V} < \frac{\sqrt{n}}{\sigma_n^2} \iff \boxed{V = o\left(\frac{n}{\sigma_n^4}\right)}.$$

When σ_n diverges, the requirement $o_p(\sigma_n/(\sqrt{n}V))$ in Theorem 5.14 is weaker than $o_p(1/(\sqrt{n}V))$ by a factor of σ_n ; the bound above is therefore a sufficient, slightly conservative condition that guarantees the stability requirements used in both theorems.

G.4 Examples

Substituting specific rates of σ_n yields:

- **Standard RAL estimators** ($\sigma_n = O(1)$): $V = o(n)$. This reproduces the condition already stated in Theorem 5.6, confirming that any slow divergence such as $V = \log n$ suffices.
- **0-order HAL** ($\sigma_n \asymp n^{1/6}$, equivalently $\sigma_n = n^{-1/3} \cdot n^{1/2}$ from a HAL convergence rate of $n^{-1/3}$):

$$V < \frac{n}{\sigma_n^4} = \frac{n}{n^{2/3}} = n^{1/3} \implies \boxed{V = o(n^{1/3})}.$$

- **First-order HAL** ($\sigma_n \asymp n^{1/10}$, equivalently $\sigma_n = n^{-2/5} \cdot n^{1/2}$ from a HAL convergence rate of $n^{-2/5}$):

$$V < \frac{n}{\sigma_n^4} = \frac{n}{n^{2/5}} = n^{3/5} \implies \boxed{V = o(n^{3/5})}.$$

G.5 Remark

The bound $V = o(n/\sigma_n^4)$ is obtained under the canonical $f(p - p_0)^2$ remainder shape; estimators whose remainder has a sharper, for example bilinear or product-form, structure may admit a faster growth rate for V . The derivation above should therefore be read as a sufficient condition that delivers the remainder-stability requirement used in the diverging- V proofs, rather than as a tight characterization. In practice, small fixed V such as 5, 10, or 20 remains the recommended choice in finite samples; the explicit rates above guarantee that the diverging- V theory also applies if one wishes to let V grow slowly with n to obtain consistent variance estimation.

Appendix D

Supporting Material for Chapter 6

This appendix gathers the technical regularity arguments, the alternative self-consistency targeting construction, and the supplemental diagnostics for Chapter 6. The main chapter keeps the EM-HAL construction, the projection-based regularized A-TMLE update, the grouped-jackknife inference layer, and the core simulation results in the body.

D.1 Implementation details for the EM-HAL and targeting pipeline

This appendix records the numerical settings used in the simulation study of Chapter 6. These choices are implementation details rather than methodological claims, so they are kept out of the main text in the dissertation version.

For the bivariate right-censoring simulations in Section 6.7, the HAL basis was constructed from $K_{\text{main}} = 20$ marginal knots placed at empirical quantiles and $K_{\text{interact}} = 5$ interaction knots, yielding approximately 66 active basis functions after pruning. Basis functions with estimated coefficients below $\varepsilon_{\text{prune}} = 10^{-4}$ were removed after the initial fit in order to stabilize the subsequent EM iterations and targeting steps.

In the E-step, each censored observation was represented by $B = 20$ imputations drawn on the numerical integration grid restricted to its coarsening set $C(O_i)$. The EM algorithm was run for at most 50 iterations and terminated earlier whenever the change in observed-data log-likelihood satisfied $|\Delta\ell_{\text{obs}}| < 0.1$. Numerical integration for density normalization and conditional expectation calculations was carried out on an approximately 45×45 fitting grid, while survival and density summaries were reported on the finer 209×209 evaluation grid described in Section 6.3.3; this grid is obtained by merging 200 equally spaced points on $[0, 1]$ with the nine targeting coordinates $\{0.1, 0.2, \dots, 0.9\}$ in each margin, giving 209 unique coordinates per dimension.

D.2 Pointwise Asymptotic Normality of the EM-HAL Density Sieve MLE

This appendix records the sieve-HAL regularity conditions and pointwise asymptotic-normality result used to place the EM-HAL estimator within the framework of van der Laan [2023]. The main-text survival-targeting arguments rely primarily on the loss-based convergence of the initial density estimator established in that technical report; the pointwise result is included here for completeness.

Assumption D.1 (HAL Sieve Regularity; adapted from van der Laan [2023], Section 10). *Let $M_n = \|Q_n\|_{v,k}^*$ denote the k -th order sectional variation norm of the unnormalized log-density estimator Q_n . Assume the following regularity conditions hold for the working models $D_\lambda^{(k)}(\mathcal{R}_{0,n})$:*

(G0) (Adaptive Uniform Approximation.)

$$\sup_{Q \in \{Q_n, Q_0\}} \inf_{g \in D_\lambda^{(k)}(\mathcal{R}_{0,n})} \|Q - g\|_\infty = O_P\left(C(M_n) r(d, d_{0,n})^{k+1}\right).$$

Here $r(d, d_{0,n})$ is the approximation rate (noting that the total sieve dimension $d_{0,n}$ is proportional to the average knot-points per component $J_{0,n}$ used in van der Laan [2023], so $r(d, d_{0,n}) \asymp r(d, J_{0,n})$). This directly aligns with the adaptive uniform approximation condition (Definition 10 of van der Laan). By Theorem 3 therein, this property transfers to the oracle MLE, ensuring that the deterministic sieve bias satisfies

$$\|Q_{0,n} - Q_0\|_\infty = O_P\left(C(M_n) r(d, d_{0,n})^{k+1}\right).$$

(G1) (Empirical Score Residual.) *The empirical score residual on the orthonormal basis satisfies $\tilde{r}_n(x) \equiv \sum_{j \in \mathcal{R}_{0,n}} \{P_n S_{Q_n}(\phi_j^*)\} \phi_j^*(x) = o_P\left((n/d_{0,n})^{-1/2}\right)$.*

(G2) (Second-Order Score Remainder.) *The exact second-order Taylor remainder of the score expansion satisfies $\bar{R}_n(x) \equiv \sum_{m=1}^3 \sum_{j \in \mathcal{R}_{0,n}} R_{m,n}(\phi_j^*) \phi_j^*(x) = o_P\left((n/d_{0,n})^{-1/2}\right)$, where the specific remainder components are defined as:*

$$\begin{aligned} R_{1,n}(\phi) &\equiv P_0\{S_{Q_n}(\phi) - S_{Q_{0,n}}(\phi)\} - P_0 dS_{Q_{0,n}}(\phi)(Q_n - Q_{0,n}), \\ R_{2,n}(\phi) &\equiv P_0 dS_{Q_{0,n}}(\phi)(Q_n - \Pi_{J_{0,n}}(Q_n)), \\ -R_{3,n}(\phi) &\equiv -P_0 dS_{Q_{0,n}}(\phi)(\Pi_{J_{0,n}}(Q_n) - Q_{0,n}) - \langle \phi, \Pi_{J_{0,n}}(Q_n) - Q_{0,n} \rangle_{J_{0,n}}, \end{aligned}$$

with $dS_Q(\phi)(h)$ denoting the directional derivative of the score. These remainders are thoroughly analyzed in van der Laan [2023].

(G3) (Empirical Process Term.) *The empirical process evaluated at the score differences satisfies $E_n(x) \equiv \sum_{j \in \mathcal{R}_{0,n}} (P_n - P_0)\{S_{Q_n}(\phi_j^*) - S_{Q_{0,n}}(\phi_j^*)\} \phi_j^*(x) = o_P\left((n/d_{0,n})^{-1/2}\right)$.*

(G4) (Rate Condition.) *The sieve bias is asymptotically negligible relative to the standard error: $C(M_n) r(d, d_{0,n})^{k+1} = o_P((n/d_{0,n})^{-1/2})$.*

For the EM-HAL MLE, (G1) holds naturally because the algorithm directly optimizes the penalized likelihood, ensuring the score equations evaluated at the empirical distribution are bounded by the subgradient of the ℓ_1 penalty. (G4) is satisfied by selecting an appropriate complexity parameter (e.g., $d_{0,n} \asymp n^{1/5} \log^{m_1} n$).

Theorem D.2 (Pointwise Asymptotic Normality of the EM-HAL Density Sieve MLE; van der Laan [2023], Theorems 4–5, 11–12). *Under Assumption D.1 with $k = 1$ (first-order continuous hinge functions, $k^* = 2$) and the sieve dimension tuned to balance bias and variance as $d_{0,n} \asymp n^{1/5} \log^{m_1} n$:*

(i) (Convergence rate.) *The EM-HAL log-density estimator achieves the dimension-free uniform convergence rate:*

$$\|Q_n - Q_0\|_\infty = O_P(n^{-2/5} \log^m n)$$

for some $m < \infty$.

(ii) (Pointwise asymptotic normality.) *For each fixed $x \in [0, 1]^2$, the estimator within the working model admits the exact linear expansion:*

$$(Q_n - Q_{0,n})(x) = (P_n - P_0)D_{Q_{0,n},x} + R_n(x), \quad (\text{D.1})$$

where the total remainder $R_n(x) = -\tilde{r}_n(x) + \bar{R}_n(x) - E_n(x)$ satisfies $|R_n(x)| = o_P((d_{0,n}/n)^{1/2})$ and the influence curve evaluated along the fixed working model is

$$D_{Q_{0,n},x}(O) \equiv \sum_{j \in \mathcal{R}_{0,n}} S_{Q_{0,n}}(\phi_j^*)(O) \phi_j^*(x).$$

Because the observed-data loss $L(Q)$ is a true log-likelihood, the information equality ensures that the score operators $S_{Q_{0,n}}(\phi_j^*)$ are identically orthonormal in $L^2(P_0)$. Therefore, the pointwise asymptotic variance reduces to:

$$\tilde{\sigma}_{0,n}^2(x) \equiv \frac{1}{d_{0,n}} \text{Var}_{P_0}(D_{Q_{0,n},x}) = \frac{1}{d_{0,n}} \sum_{j \in \mathcal{R}_{0,n}} \{\phi_j^*(x)\}^2. \quad (\text{D.2})$$

Because the deterministic oracle bias $(Q_{0,n} - Q_0)(x) = o_P((n/d_{0,n})^{-1/2})$ by (G4), we obtain weak convergence for the overall error to a standard normal limit:

$$\tilde{\sigma}_{0,n}^{-1}(x) (n/d_{0,n})^{1/2} (Q_n - Q_0)(x) \Rightarrow N(0, 1).$$

Remark D.3 (Rigorous Transfer of Sieve-HAL Theory to the Coarsened-Data Likelihood). While van der Laan [2023] develops the general Sieve-HAL theory primarily through the lens

of full-data empirical risk minimization, the framework rigorously transfers to our coarsened observed-data setting. By analyzing the EM-HAL directly, we can formalize it similarly to regular HAL: (a) we use log-likelihood cross-validation to select a working model, ensuring it is large enough that the loss dissimilarity between its oracle and the truth is negligible; and (b) we analyze our estimator as the MLE over that parametric model under an ℓ_1 -norm constraint (where the EM algorithm with an ℓ_1 penalty is used to implement this estimator). The observed-data likelihood structurally satisfies the necessary conditions (G0) through (G4):

- 1. Working Model Selection and Loss-Based Dissimilarity (Condition G0):** Log-likelihood cross-validation selects a sufficiently rich working model such that the deterministic oracle approximation error vanishes. Under standard positivity assumptions for right-censoring, the observed-data loss-based dissimilarity strictly bounds the full-data squared L_2 distance. Because standard regular HAL theory directly applies to our ℓ_1 -constrained MLE, the optimal rate for the observed-data dissimilarity guarantees the required full-data L_2 convergence rate.
- 2. Solving the Empirical Score (Condition G1):** The EM algorithm optimizes the likelihood over the data-adaptively selected working model subject to the ℓ_1 bound λ_{EM} . This explicit constraint provides the requisite stability against collinearity while ensuring the active empirical scores are solved well enough. By sufficiently reducing the dual Lagrange multiplier γ_n , the empirical score residual on the active orthonormal basis satisfies $\tilde{r}_n(x) = o_P((n/d_{0,n})^{-1/2})$, fulfilling condition (G1).
- 3. Log-Likelihood Behaving Loss & Exact Remainder (Condition G2):** Because our observed-data CAR loss $L(Q)(O) = -\log q_X(C(O))$ is a proper log-likelihood, Bartlett's identity strictly holds natively. This guarantees that the exact second-order Taylor remainder of the score expansion is strictly bounded by the squared loss-based dissimilarity, fulfilling condition (G2).
- 4. Empirical Process Term and Donsker Property (Conditions G3 & G4):** Under coarsening, the observed-data score takes the conditional expectation form $S_Q(\phi_j)(O) = E_Q[\phi_j(X) | C(O)] - E_Q[\phi_j(X)]$. Because the EM algorithm implements the MLE under an ℓ_1 -norm control, it directly enforces a bounded sectional variation norm constraint on the full data density. This strict sectional variation control inherently forces the Donsker condition to hold even in the presence of the inverse information operator. It ensures the observed-data score class maintains bounded sectional variation and operates as a P_0 -Donsker class, allowing the empirical process term to be bounded at the required rate without needing strip-censoring.

Given these verifications, the structural conditions hold for the observed-data likelihood. By the typical proof of regular HAL [van der Laan, 2023], our estimator converges to the working-model oracle at the rate implied by the underlying càdlàg function class, directly yielding the stated full-data rates (e.g., $O_P(n^{-1/3})$ for zero-order or $O_P(n^{-2/5} \log^m n)$ for first-order splines) and pointwise normal limit.

D.3 Self-consistency targeting: NPMLE-type regularized TMLE via self-consistency equations

Self-consistency targeting is the comparative grid-based alternative studied in this chapter. It is motivated by the NPMLE self-consistency system for the survival surface. Because solving the full system is practically intractable, the method works on a finite targeting grid $\mathcal{G}_n$ and treats m_n as a precision parameter. Since it uses global indicators rather than projecting onto the HAL score span, it tends to produce a rougher fit and higher variance.

D.3.1 Self-consistency equations for survival

For a targeting grid $\mathcal{G}_n = \{s_1, \dots, s_{m_n}\} \subset [0, 1]^2$ with $s_\ell = (t_{1\ell}, t_{2\ell})$, define the full-data step function

$$h_\ell(X) \equiv I(T_1 > t_{1\ell}, T_2 > t_{2\ell}) - \Psi_\ell(P),$$

where $\Psi_\ell(P) = S_P(t_{1\ell}, t_{2\ell})$. Under CAR, the NPMLE self-consistency conditions require

$$P_n[E_P\{h_\ell(X) \mid O\}] = 0, \quad \ell = 1, \dots, m_n. \quad (\text{D.3})$$

These grid-indexed equations are the computable subsystem solved by self-consistency targeting. In the notation introduced below, they play the role of B_n , while a larger monitored self-consistency family, evaluated on a finer grid or richer collection of points, plays the role of the target system A_n . Solving (D.3) on a sufficiently fine grid $\mathcal{G}_n$ therefore provides a feasible approximation to the full NPMLE characterization.

D.3.2 Indicator-based offset targeting submodel

For implementation, we treat the initial EM-HAL density $f_{\hat{\theta}}$ as a fixed offset and augment the log-density with indicator-based targeting functions:

$$\log f_{\theta, \alpha}(t_1, t_2) = \underbrace{\log f_{\hat{\theta}}(t_1, t_2)}_{\text{Fixed Offset}} + \sum_{\ell=1}^{m_n} \alpha_\ell \tilde{\phi}_\ell(t_1, t_2) - \log \tilde{Z}(\theta, \alpha), \quad (\text{D.4})$$

where $\tilde{\phi}_\ell(t_1, t_2) = I(t_1 > t_{1\ell}, t_2 > t_{2\ell})$. The parameter θ is fixed at the initial EM-HAL estimate $\hat{\theta}_{\text{EMHAL}}$, and $\alpha \in \mathbb{R}^{m_n}$ is initialized at 0.

D.3.3 Solving the finite self-consistency system via EM

The targeted model (D.4) is fitted by a targeting EM scheme in which θ is held fixed at the initial EM-HAL estimate and only the targeting coefficients α are updated. Algorithm 13 summarizes this offset-based procedure. Because the maximization is carried out directly within the augmented working model $\mathcal{M}_{n, \text{extend}}$, the resulting fit drives the empirical score equations for α_ℓ toward zero and enforces the chosen self-consistency system on $\mathcal{G}_n$. The

added indicator basis functions are chosen to reduce survival bias, not to represent the density efficiently.

Algorithm 13 Self-consistency targeting via indicator-based offset targeting

- 1: **Input:** initial EM-HAL density $f_{\hat{\theta}_{\text{EMHAL}}}$, targeting grid $\mathcal{G}_n = \{s_\ell\}_{\ell=1}^{m_n}$, offset model (D.4)
 - 2: Fix $\theta = \hat{\theta}_{\text{EMHAL}}$ and initialize $\alpha = 0$
 - 3: **for** EM iterations until convergence **do**
 - 4: **E-step:**
 Compute the conditional distribution of X given O under the current offset density $f_{\theta, \alpha}$
 - 5: **M-step:**
 Maximize the corresponding complete-data log-likelihood over α in the offset model (D.4)
 Hold θ fixed and impose only a mild ℓ_1 constraint for numerical stability
 - 6: **end for**
 - 7: **Output:** targeted density $f_{\theta, \hat{\alpha}}$ and survival estimates on $\mathcal{G}_n$
-

Crucially, the score with respect to α_ℓ equals $E_{f_{\theta, \alpha}}\{h_\ell(X) \mid O\}$. Taking the MLE in this augmented working model therefore solves the self-consistency equations (D.3) on the chosen targeting grid. In our implementation, $\mathcal{G}_n = \{0.1, 0.2, \dots, 0.9\}^2$, so $m_n = 81$.

D.3.4 Regularized TMLE theory for the finite self-consistency system

The preceding construction solves the computable proxy subsystem B_n on the targeting grid $\mathcal{G}_n$. The following finite-system regularized-TMLE statement formalizes the logic that solving this proxy system can still control a broader target system A_n .

Theorem D.4 (Regularized TMLE as indirect control of a target score system). *Let $A_n(P)$ denote a target score system we would ideally like to solve (e.g., evaluated on a continuous surface or a dense grid), and let $B_n(P)$ denote a proxy score system that is computationally available and actually solved. Suppose the targeted fit P_n^* satisfies*

$$\|P_n B_n(P_n^*)\| = o_P(n^{-1/2})$$

under a chosen norm $\|\cdot\|$, and suppose that solving B_n controls the target system A_n in the sense that

$$\|P_n A_n(P_n^*)\| \leq C_n \|P_n B_n(P_n^*)\| + r_n,$$

with $C_n = O_P(1)$ and an approximation remainder $r_n = o_P(n^{-1/2})$. Then

$$\|P_n A_n(P_n^*)\| = o_P(n^{-1/2}).$$

Thus, regularized TMLE need not solve the full target score system directly: it is enough to solve an available proxy system B_n to precision faster than $n^{-1/2}$, provided that doing so yields structural control on the original system A_n .

Proof of Theorem D.4. By hypothesis, the targeted estimator solves the proxy system such that $\|P_n B_n(P_n^*)\| = o_P(n^{-1/2})$. Because the scaling factor C_n is stochastically bounded ($C_n = O_P(1)$), their product is $C_n \|P_n B_n(P_n^*)\| = O_P(1) \cdot o_P(n^{-1/2}) = o_P(n^{-1/2})$. Furthermore, the remainder is assumed to satisfy $r_n = o_P(n^{-1/2})$. Applying the assumed structural inequality yields

$$0 \leq \|P_n A_n(P_n^*)\| \leq o_P(n^{-1/2}) + o_P(n^{-1/2}) = o_P(n^{-1/2}),$$

which establishes the result. $\square$

To connect the finite grid-based self-consistency system to the full survival surface, we next show that monotonicity of the indicator and boundedness of the EM-HAL density yield a uniform control bound over the continuous domain.

Theorem D.5 (Uniform control of the continuous self-consistency system via finite-grid self-consistency targeting). *Let $U(t, P) = P_n[E_P\{I(T_1 > t_1, T_2 > t_2) \mid O\}] - S_P(t_1, t_2)$ denote the continuous self-consistency score evaluated at any point $t \in [0, 1]^2$. Let $\mathcal{G}_n \subset [0, 1]^2$ be the finite Cartesian targeting grid, and let $\bar{\mathcal{G}}_n = \mathcal{G}_n \cup \partial[0, 1]^2$ denote the grid extended to the boundaries. Define its maximum mesh size in the supremum norm as δ_n . Suppose the targeted EM-HAL density estimate P_n^* has marginal densities uniformly bounded by a constant $M = O_P(1)$. If self-consistency targeting solves the proxy system B_n on $\mathcal{G}_n$ such that*

$$\|P_n B_n(P_n^*)\|_\infty = \max_{s \in \mathcal{G}_n} |U(s, P_n^*)| = o_P(n^{-1/2}),$$

then the target score system $A_n(P_n^) = U(\cdot, P_n^*)$ is uniformly controlled over the entire continuous domain $[0, 1]^2$:*

$$\|P_n A_n(P_n^*)\|_\infty = \sup_{t \in [0, 1]^2} |U(t, P_n^*)| \leq \|P_n B_n(P_n^*)\|_\infty + 2M\delta_n.$$

Consequently, if the grid is sufficiently dense such that the mesh size satisfies $\delta_n = o_P(n^{-1/2})$, self-consistency targeting achieves uniform parametric control everywhere:

$$\sup_{t \in [0, 1]^2} |U(t, P_n^*)| = o_P(n^{-1/2}).$$

Proof of Theorem D.5. Fix an arbitrary point $t = (t_1, t_2) \in [0, 1]^2$. By the definition of the grid mesh size, there exist points $s_L = (s_{1,L}, s_{2,L})$ and $s_U = (s_{1,U}, s_{2,U})$ in $\bar{\mathcal{G}}_n$ such that $s_L \leq t \leq s_U$ coordinate-wise, and $\|s_U - s_L\|_\infty \leq \delta_n$. (Note that on the upper boundaries where $t_1 = 1$ or $t_2 = 1$, the survival function and indicator are deterministically zero, so the error is identically zero.) Because the indicator function $I(T > \cdot)$ is monotonically non-increasing in each coordinate, we have the sure ordering:

$$I(T_1 > s_{1,U}, T_2 > s_{2,U}) \leq I(T_1 > t_1, T_2 > t_2) \leq I(T_1 > s_{1,L}, T_2 > s_{2,L}).$$

Taking the conditional expectation $E_{P_n^*}[\cdot \mid O]$ and the empirical average P_n preserves this inequality:

$$P_n E_{P_n^*}[I(T > s_U) \mid O] \leq P_n E_{P_n^*}[I(T > t) \mid O] \leq P_n E_{P_n^*}[I(T > s_L) \mid O].$$

Let $\epsilon_n(s) = U(s, P_n^*)$ denote the target score evaluated at s . We can substitute $P_n E_{P_n^*}[I(T > s) \mid O] = S_{P_n^*}(s) + \epsilon_n(s)$ for the grid points. Subtracting the target survival function $S_{P_n^*}(t)$ throughout yields:

$$S_{P_n^*}(s_U) - S_{P_n^*}(t) + \epsilon_n(s_U) \leq U(t, P_n^*) \leq S_{P_n^*}(s_L) - S_{P_n^*}(t) + \epsilon_n(s_L).$$

Because $S_{P_n^*}(\cdot)$ is also monotonically non-increasing, $S_{P_n^*}(s_U) \leq S_{P_n^*}(t) \leq S_{P_n^*}(s_L)$. This allows us to bound the error strictly by:

$$|U(t, P_n^*)| \leq (S_{P_n^*}(s_L) - S_{P_n^*}(s_U)) + \max_{s \in \{s_L, s_U\}} |\epsilon_n(s)|.$$

The difference $S_{P_n^*}(s_L) - S_{P_n^*}(s_U)$ is exactly the probability mass under P_n^* in the region $\{T_1 > s_{1,L}, T_2 > s_{2,L}\} \setminus \{T_1 > s_{1,U}, T_2 > s_{2,U}\}$. This region is fully contained within the union of the two marginal strips $\{s_{1,L} < T_1 \leq s_{1,U}\}$ and $\{s_{2,L} < T_2 \leq s_{2,U}\}$. Because the marginal densities of P_n^* are bounded by M , the probability mass of this union is bounded by:

$$S_{P_n^*}(s_L) - S_{P_n^*}(s_U) \leq M(s_{1,U} - s_{1,L}) + M(s_{2,U} - s_{2,L}) \leq 2M\|s_U - s_L\|_\infty \leq 2M\delta_n.$$

Taking the supremum over all $t \in [0, 1]^2$, and recognizing that

$$\max_{s \in \mathcal{G}_n} |\epsilon_n(s)| \leq \|P_n B_n(P_n^*)\|_\infty,$$

we obtain:

$$\sup_{t \in [0, 1]^2} |U(t, P_n^*)| \leq \|P_n B_n(P_n^*)\|_\infty + 2M\delta_n.$$

This maps directly to the structural requirement of Theorem D.4 under the supremum norm, with $C_n = 1$ and $r_n = 2M\delta_n$. If the grid spacing is set such that $\delta_n = o_P(n^{-1/2})$, then $r_n = o_P(n^{-1/2})$, which strictly bounds the full continuous target system to the desired parametric rate. $\square$

Remark D.6 (Theoretical basis for the bias-variance tradeoff). This structural bound provides a mathematically exact explanation for the empirical bias-variance tradeoff observed for self-consistency targeting in Section 6.7. To ensure the theoretical remainder $r_n = 2M\delta_n$ falls below the $n^{-1/2}$ threshold required for proper debiasing, the mesh size δ_n must shrink, necessitating a dense targeting grid $\mathcal{G}_n$. However, adding more indicator basis functions locally to achieve this resolution creates a rougher targeted fit, which inherently inflates variance. This mathematical reality captures why self-consistency targeting successfully centers the estimator but exhibits weaker variance control than the smoother geometry of

the reg-ATMLE update.

D.4 Additional survival point-estimation diagnostics

This appendix collects the complementary spatial point-estimation diagnostics referenced in Section 6.7.4. For self-consistency targeting, we report the omitted $|\text{Bias}|/\text{SD}$, MSE, and oracle-coverage heatmaps. For the projection-based reg-ATMLE update, Figure 6.2 remains the main-text spatial figure, and the present appendix records the complementary diagnostics.

Self-consistency targeting

Figures D.1–D.3 show the spatial pattern behind the summary in Section 6.7.4: the centering gains are concentrated in the interior and around the targeting grid, while the variance cost is most visible near the heavily censored boundary.




Figure D.1: Self-consistency targeting: $|\text{Bias}|/\text{SD}$ heatmaps for EM-HAL (left), targeted (center), and the change (right). Blue indicates improvement.




Figure D.2: Self-consistency targeting: MSE heatmaps for EM-HAL (left), targeted (center), and $\log_{10}$ ratio (right). Blue indicates improvement.

NPMLE: Oracle Coverage (213 simulations, 95% nominal)




Figure D.3: Self-consistency targeting: oracle-coverage heatmaps (95% nominal) for the EM-HAL (left) and targeted (right) stages.

reg-ATMLE

Figures D.4 and D.5 complement the main-text MSE diagnostic in Figure 6.2 by showing that the reg-ATMLE gains are not limited to one summary metric.

ATMLE: |Bias|/SD Diagnostic (213 simulations, $n=500$)




Figure D.4: reg-ATMLE: |Bias|/SD heatmaps for EM-HAL (left), targeted (center), and the change (right). Blue indicates improvement.

D.5 Additional survival confidence-interval diagnostics

This appendix presents the supplemental spatial and regional diagnostics for the survival confidence-interval comparisons discussed in Section 6.7.5.

ATMLE: Oracle Coverage (213 simulations, 95% nominal)




Figure D.5: reg-ATMLE: oracle-coverage heatmaps (95% nominal) for the EM-HAL (left) and targeted (right) stages.

D.6 Density estimation under survival targeting

The targeting procedures in Section 6.4 are designed for survival, not density, so it is natural to ask how they affect the estimated density. The answer depends on the geometry of the update: self-consistency targeting uses indicator corrections aimed at survival self-consistency, whereas the projection-based reg-ATMLE update remains inside the HAL score span and therefore stays more relevant for density estimation. Here too the 9×9 targeting grid is only the implementation grid; the substantive question is how targeting on that grid propagates to the denser surface summaries reported below.

D.6.1 Density performance after survival targeting

Table D.1 reports the density empirical performance before and after survival targeting.

Table D.1: Density empirical performance before and after survival targeting ($R = 213$, $n = 500$). Metrics averaged over the 209×209 evaluation grid.

Method	Stage	Mean Bias	Mean MSE	Oracle Cov.
Self-consistency targeting	EM-HAL	0.132	0.166	0.898
Self-consistency targeting	Targeted (survival)	0.093	0.699	0.945
reg-ATMLE	EM-HAL	0.131	0.173	0.898
reg-ATMLE	Targeted (survival)	0.065	0.164	0.935

The projection-based reg-ATMLE update modestly improves density performance: mean




Figure D.6: Self-consistency targeting: coverage heatmaps comparing CI methods for the targeted survival estimate.




Figure D.7: Survival coverage and CI length: central $[0.25, 0.75]^2$ vs. tail region for self-consistency targeting (top) and reg-ATMLE (bottom). Metrics are averaged per evaluation point across 213 simulations.




Figure D.8: Survival coverage and CI length: 9×9 targeting grid (81 points) vs. non-target points for self-consistency targeting (top) and reg-ATMLE (bottom). Metrics are averaged per evaluation point across 213 simulations.

absolute bias falls by 50.5% and MSE falls by 4.9%. The gain is smaller than for survival, but it is consistent with the projection-based update staying aligned with the smooth exponential-family geometry of the initial fit.

Self-consistency targeting behaves very differently. Although absolute bias decreases, density MSE rises by more than 300% because the indicator-based correction introduces substantial variance and a visibly rougher fit near the targeting grid. This is the density-side version of the same structural tradeoff seen for survival point estimation: the added basis functions are useful for survival centering, but they are poor basis functions for density representation.

D.6.2 Density V-fold jackknife inference

Table D.2 reports jackknife CI coverage and length for the density.

Table D.2: Density jackknife CI performance ($R = 213$, $V = 20$). Averaged over the 209×209 evaluation grid.

Method	Stage	JK Coverage	Mean CI Length
Self-consistency targeting	EM-HAL	0.874	1.020
Self-consistency targeting	Targeted	0.851	2.984
reg-ATMLE	EM-HAL	0.869	0.997
reg-ATMLE	Targeted	0.968	1.677

The density jackknife results tell a similar story. Targeted reg-ATMLE reaches 96.8% coverage, which suggests that the projection-based update improves the linearization of the density functional as well. By contrast, the untargeted EM-HAL density estimator remains undercovered at about 87%, and self-consistency targeting mainly widens intervals without recovering comparable stability.

D.6.3 Density simultaneous confidence bands

Table D.3 extends the simultaneous confidence band analysis to the density functional on the 9×9 targeting grid, using the same componentwise-Studentized Monte Carlo calibration as for survival (Section 6.6.1).

The simultaneous-band results reinforce this contrast. After bias correction, reg-ATMLE reaches 94.4% simultaneous coverage, close to the nominal 95% level. The untargeted EM-HAL density estimator has only about 31% simultaneous coverage because small pointwise biases accumulate under the 81-point multiplicity correction. Self-consistency targeting fails even more dramatically here: the rougher targeted density has 0% base simultaneous coverage and only 4.7% after bias correction, again reflecting the variance inflation already seen in Table D.1.

Table D.3: Density simultaneous coverage on the 9×9 targeting grid (81 points, $R = 213$).

Method	Stage	Simult. Cov.	Simult. Cov. (BC)	Per-pt Avg	Mean q
Self-consistency targeting	EM-HAL	0.310	0.554	0.967	3.79
Self-consistency targeting	Targeted	0.000	0.047	0.884	4.01
reg-ATMLE	EM-HAL	0.315	0.549	0.965	3.79
reg-ATMLE	Targeted	0.822	0.944	0.994	3.89

D.7 Proofs of Theoretical Results

Proof of Theorem 6.2. Fix a targeting-grid point $s \in \mathcal{G}_n$. We have the exact pathwise Taylor expansion of the working-model-induced scalar target:

$$\begin{aligned}
& \Psi_{\mathcal{M}_n, s}(P_n^*) - \Psi_{\mathcal{M}_n, s}(P_0) \\
&= -P_0 D_{\mathcal{M}_n, s}^*(P_n^*) + R_{\mathcal{M}_n, s}(P_n^*, P_0) \\
&= (P_n - P_0) D_{\mathcal{M}_n, s}^*(P_n^*) - P_n D_{\mathcal{M}_n, s}^*(P_n^*) + R_{\mathcal{M}_n, s}(P_n^*, P_0) \\
&= P_n D_{\mathcal{M}_n, s}^*(P_0) - P_0 D_{\mathcal{M}_n, s}^*(P_0) + (P_n - P_0) (D_{\mathcal{M}_n, s}^*(P_n^*) - D_{\mathcal{M}_n, s}^*(P_0)) \\
&\quad - P_n \tilde{D}_{n, s}^*(P_n^*) - P_n (D_{\mathcal{M}_n, s}^*(P_n^*) - \tilde{D}_{n, s}^*(P_n^*)) + R_{\mathcal{M}_n, s}(P_n^*, P_0).
\end{aligned}$$

By definition of the influence curve, $P_0 D_{\mathcal{M}_n, s}^*(P_0) = 0$. The targeting step solves the approximate scalar score equation $P_n \tilde{D}_{n, s}^*(P_n^*) = o_P(n^{-1/2})$. By Assumption 6.1, the corresponding exact working-model score is also resolved at the same rate, so $P_n \{D_{\mathcal{M}_n, s}^*(P_n^*) - \tilde{D}_{n, s}^*(P_n^*)\} = o_P(n^{-1/2})$.

By conditions (A.1) and (A.2), asymptotic equicontinuity implies

$$(P_n - P_0) (D_{\mathcal{M}_n, s}^*(P_n^*) - D_{\mathcal{M}_n, s}^*(P_0)) = o_P(n^{-1/2}).$$

By condition (A.3), the exact second-order remainder satisfies $R_{\mathcal{M}_n, s}(P_n^*, P_0) = o_P(n^{-1/2})$. Substituting these bounds yields

$$\Psi_{\mathcal{M}_n, s}(P_n^*) - \Psi_{\mathcal{M}_n, s}(P_0) = P_n D_{\mathcal{M}_n, s}^*(P_0) + o_P(n^{-1/2}),$$

which proves the theorem. $\square$

Proof of Theorem 6.3. Fix $s \in \mathcal{G}_n$. Under $\mathcal{M}_0 = \mathcal{M}$, the projected parameter equals the original scalar target $\Psi_s(P) = \Psi_{\mathcal{M}_0, s}(P)$ and the corresponding efficient influence curve satisfies $D_{\mathcal{M}_0, s}^*(P) = D_{\Psi_s}^*(P)$.

Because $\mathcal{M}_0 = \mathcal{M}$, condition (C.1) holds trivially: $\Psi_s(P_{0, \mathcal{M}_n}) - \Psi_{\mathcal{M}_0, s}(P_{0, \mathcal{M}_n}) = 0$. We therefore decompose the deterministic projection bias via an exact pathwise Taylor expansion

of Ψ_s around $P_{0,\mathcal{M}_n}$:

$$\begin{aligned}\Psi_{\mathcal{M}_n,s}(P_0) - \Psi_s(P_0) &= \Psi_s(P_{0,\mathcal{M}_n}) - \Psi_s(P_0) \\ &= -P_0 D_{\Psi_s}^*(P_{0,\mathcal{M}_n}) + R(P_{0,\mathcal{M}_n}, P_0) \\ &= (P_{0,\mathcal{M}_n} - P_0) \left(D_{\Psi_s}^*(P_{0,\mathcal{M}_n}) - D_{\mathcal{M}_n,s}^*(P_{0,\mathcal{M}_n}) \right) \\ &\quad + R(P_{0,\mathcal{M}_n}, P_0) - P_0 D_{\mathcal{M}_n,s}^*(P_{0,\mathcal{M}_n}),\end{aligned}$$

where $R(P, P_0) = \Psi_s(P) - \Psi_s(P_0) + P_0 D_{\Psi_s}^*(P)$ is the exact second-order remainder for the scalar target Ψ_s .

Because the EM-HAL density sieve eventually includes P_0 , $P_{0,\mathcal{M}_n}$ converges to P_0 , ensuring $R(P_{0,\mathcal{M}_n}, P_0) = o_P(n^{-1/2})$. Combining conditions (C.2) and (C.3) yields

$$\Psi_{\mathcal{M}_n,s}(P_0) - \Psi_s(P_0) = o_P(n^{-1/2}).$$

We then combine this deterministic bias bound with Theorem 6.2:

$$\begin{aligned}\Psi_{\mathcal{M}_n,s}(P_n^*) - \Psi_s(P_0) &= \Psi_{\mathcal{M}_n,s}(P_n^*) - \Psi_{\mathcal{M}_n,s}(P_0) + \Psi_{\mathcal{M}_n,s}(P_0) - \Psi_s(P_0) \\ &= P_n D_{\mathcal{M}_n,s}^*(P_0) + o_P(n^{-1/2}).\end{aligned}$$

Finally, add and subtract the true nonparametric influence curve:

$$P_n D_{\mathcal{M}_n,s}^*(P_0) = P_n D_{\Psi_s}^*(P_0) + (P_n - P_0) \left(D_{\mathcal{M}_n,s}^*(P_0) - D_{\Psi_s}^*(P_0) \right) + P_0 \left(D_{\mathcal{M}_n,s}^*(P_0) - D_{\Psi_s}^*(P_0) \right).$$

Because both influence curves have mean zero under P_0 , the last term vanishes. By condition (B.1) and asymptotic equicontinuity,

$$(P_n - P_0) \left(D_{\mathcal{M}_n,s}^*(P_0) - D_{\Psi_s}^*(P_0) \right) = o_P(n^{-1/2}),$$

so

$$P_n D_{\mathcal{M}_n,s}^*(P_0) = P_n D_{\Psi_s}^*(P_0) + o_P(n^{-1/2}).$$

Substituting this back yields

$$\Psi_{\mathcal{M}_n,s}(P_n^*) - \Psi_s(P_0) = P_n D_{\Psi_s}^*(P_0) + o_P(n^{-1/2}).$$

Because $\Psi_s(P_n^*) = \Psi_{\mathcal{M}_n,s}(P_n^*)$ inside the working model $\mathcal{M}_n$, this establishes

$$\hat{S}^{\text{targ}}(s) - \Psi_s(P_0) = \frac{1}{n} \sum_{i=1}^n D_{\Psi_s}^*(P_0)(O_i) + o_P(n^{-1/2}),$$

which proves the theorem. □