Confidence Intervals for Nonparametric Empirical Bayes Analysis

Οικονομικό Πανεπιστήμιο Αθηνών Τμήμα Στατιστικής

Οκτώβριος 7, 2022

Nikos Ignatiadis Statistics Department, Columbia University

What is empirical Bayes?

- Regularization/Shrinkage estimation
- Deconvolution
- Inverse problem
- Hierarchical Bayes
- A Mixed Model with (potentially) some nonparametric components.
- Compound decision theory



Erfahrungs-Tarifierung in der Motorfahrzeughaftpflicht-Versicherung

Von Fritz Bichsel, Muri bei Bern

 $Z_i(1961)$: number of claims of *i*-th policy holder in 1961

$$z = \#\{Z_i(1961) = z\}$$

- 0 103704
- 1 14075
- 2 1766
- 3 255
- 4 45> 5 8



Erfahrungs-Tarifierung in der Motorfahrzeughaftpflicht-Versicherung

Von Fritz Bichsel, Muri bei Bern

 $Z_i(1961)$: number of claims of *i*-th policy holder in 1961

- $\begin{array}{ll}z & \#\{Z_i(1961) = z\}\\0 & 103704\\1 & 14075\end{array}$ Goal: Assign premium for the next year as a function of $Z_i(1961)$.
- 2 1766
- 3 255
- 4 45> 5 8



Erfahrungs-Tarifierung in der Motorfahrzeughaftpflicht-Versicherung

Von Fritz Bichsel, Muri bei Bern

 $Z_i(1961)$: number of claims of *i*-th policy holder in 1961

- $z = \#\{Z_i(1961) = z\}$
- 0 103704
- 1 14075
- 2 1766
- 3 255
- 4 45
- ≥ 5 8

- Goal: Assign premium for the next year as a function of $Z_i(1961)$.
 - Idea: Assess expected number of
 - claims in next year,

 $\mathbb{E}[Z_i(1962) \mid Z_i(1961) = z]$

i: policy holder, $Z_i(t)$: number of claims in year *t*.

$Z_i(1961), Z_i(1962) \stackrel{iid}{\sim} \text{Poisson}(\mu_i)$

i: policy holder, $Z_i(t)$: number of claims in year t.

$Z_i(1961), Z_i(1962) \stackrel{iid}{\sim} \text{Poisson}(\mu_i)$



i: policy holder, $Z_i(t)$: number of claims in year *t*.

$Z_i(1961), Z_i(1962) \stackrel{iid}{\sim} \text{Poisson}(\mu_i)$

 $\mu_i \sim G$ "Structural function representing heterogeneity of the portfolio"

Best guess for the number of claims in 1962 for i is:

 $\mathbb{E}[Z_i(1962) \mid Z_i(1961) = z] \\ = \mathbb{E}[\mu_i \mid Z_i(1961) = z] \\ = \theta_G(z)$

i: policy holder, $Z_i(t)$: number of claims in year *t*.

$$Z_i(1961), Z_i(1962) \stackrel{iid}{\sim} \text{Poisson}(\mu_i)$$

Best guess for the number of claims in 1962 for i is:

$$\mathbb{E}[Z_i(1962) \mid Z_i(1961) = z] \\ = \mathbb{E}[\mu_i \mid Z_i(1961) = z] \\ = \theta_G(z)$$

Bichsel did not know G... But had data.

Nonparametric Maximum Likelihood (NPMLE)



 $G \in \mathscr{G} = \{ all \text{ distributions supported on } [0, 5] \}$



Kiefer-Wolfowitz (1956), Simar (1976), Laird (1978), Lindsay (1983), Walhin and Paris (1999), Koenker and Mizera (2014), Koenker and Gu (2017)

What about confidence intervals?

Estimated posterior mean: $\hat{\theta}_G(3) = \mathbb{E}_{\widehat{G}}[\mu \mid Z = 3] = 0.69$

What about confidence intervals?

Estimated posterior mean: $\hat{\theta}_G(3) = \mathbb{E}_{\widehat{G}}[\mu \mid Z = 3] = 0.69$

SCANDINAVIAN ACTUARIAL JOURNAL, 2018 VOL. 2018, NO. 2, 129–144 https://doi.org/10.1080/03461238.2017.1307267





Confidence intervals of the premiums of optimal bonus malus systems

Dimitris Karlis^a, George Tzougas^b and Nicholas Frangos^a

^aDepartment of Statistics, Athens University of Economics and Business, Athens, Greece; ^bDepartment of Statistics, London School of Economics, London, UK

$\theta_G(3) = \mathbb{E}_G[\mu \mid Z = 3] \in [0.55, 0.80].$

Confidence intervals for $\theta_G(z) = \mathbb{E}_G[\mu \mid Z = z]$

1 Let \widehat{G} be the NPMLE of G based on Z_1, \ldots, Z_n . 2 for b = 1 to B do

3 Draw
$$\mu_i^b \sim \widehat{G}, \ Z_i^b \sim p(\cdot \mid \mu_i^b) \text{ for } i = 1, \dots, n \text{ (iid)}.$$

4 Let
$$\widehat{G}^b$$
 be the NPMLE of G based on Z_1^b, \ldots, Z_n^b .

5 Let
$$\hat{\theta}^b(z) = \theta_{\widehat{G}^b}(z)$$
.

6 end

7 Form a percentile bootstrap confidence interval $[\hat{\theta}_{\alpha}^{-}(z), \hat{\theta}_{\alpha}^{+}(z)]$ of $\theta_{G}(z)$ based on $\hat{\theta}^{b}(z), b = 1, \dots, B$.

Confidence intervals for $\theta_G(z) = \mathbb{E}_G[\mu \mid Z = z]$

1 Let \widehat{G} be the NPMLE of G based on Z_1, \ldots, Z_n .

2 for
$$b = 1$$
 to B do

3 Draw
$$\mu_i^b \sim \widehat{G}, \ Z_i^b \sim p(\cdot \mid \mu_i^b) \text{ for } i = 1, \dots, n \text{ (iid)}.$$

4 Let
$$\widehat{G}^b$$
 be the NPMLE of G based on Z_1^b, \ldots, Z_n^b .

5 Let
$$\hat{\theta}^b(z) = \theta_{\widehat{G}^b}(z)$$
.

6 end

7 Form a percentile bootstrap confidence interval $[\hat{\theta}_{\alpha}^{-}(z), \hat{\theta}_{\alpha}^{+}(z)]$ of $\theta_{G}(z)$ based on $\hat{\theta}^{b}(z), b = 1, \ldots, B$.

Theorem [Karlis, Tzougas, and Frangos (2018)]: Assume independence and that $Z_i \mid \mu_i \sim \text{Poisson}(\mu_i), \mu_i \sim G$, with $G \in \mathscr{G} = \left\{ G \text{ supported on compact interval } [0,M], \\ \mathbb{P}_G[\mu \in (0,\varepsilon)] \text{ is sufficiently large for all } \varepsilon > 0 \right\}$ Then: $\lim_{n \to \infty} \left\{ \mathbb{P}_G \left[\theta_G(z) \in [\widehat{\theta}_{\alpha}^-(z), \widehat{\theta}_{\alpha}^+(z)] \right] \right\} \geq 1 - \alpha \,.$

Confidence Intervals for Nonparametric Empirical Bayes Analysis (with Rejoinder)

N.I., and Stefan Wager, JASA T&M (2022)

Empirical Bayes (EB) Robbins (1956), Efron (2010)

We have noisy data Z_i on *n* related units with latent parameter μ_i . Three main ingredients to an EB analysis:

1) Known noise model: $Z_i \mid \mu_i \sim p(\cdot \mid \mu_i)$

 $p(\cdot \mid \mu_i)$ is a density w.r.t. a measure λ , e.g., $Z_i \mid \mu_i \sim \text{Poisson}(\mu_i)$.

2) Class of priors: $\mu_i \sim G, \ G \in \mathcal{G}$

G is unknown.

3) Empirical Bayes estimand: $\theta_G(z) = \mathbb{E}_G[h(\mu_i) \mid Z_i = z]$ for a known function $h(\cdot)$, e.g., for $h(\mu) = \mu$, $\theta_G(z)$ is the posterior mean given $Z_i = z$

Empirical Bayes (EB) Robbins (1956), Efron (2010)

We have noisy data Z_i on *n* related units with latent parameter μ_i . Three main ingredients to an EB analysis:

1) Known noise model: $Z_i \mid \mu_i \sim p(\cdot \mid \mu_i)$

 $p(\cdot \mid \mu_i)$ is a density w.r.t. a measure λ , e.g., $Z_i \mid \mu_i \sim \text{Poisson}(\mu_i)$.

2) Class of priors: $\mu_i \sim G, \ G \in \mathcal{G}$

G is unknown.

3) Empirical Bayes estimand: $\theta_G(z) = \mathbb{E}_G[h(\mu_i) \mid Z_i = z]$ for a known function $h(\cdot)$, e.g., for $h(\mu) = \mu$, $\theta_G(z)$ is the posterior mean given $Z_i = z$

Predominant approach: Point estimate $\widehat{\theta}_G(z)$ of $\theta_G(z)$.

EB confidence intervals: statistical task

1) Known noise model: $Z_i \mid \mu_i \sim p(\cdot \mid \mu_i)$ 2) Class of priors: $\mu_i \sim G, \ G \in \mathcal{G}$ i = 1,...,n

G is unknown

3) Empirical Bayes estimand: $\theta_G(z) = \mathbb{E}_G[h(\mu_i) \mid Z_i = z]$

CI: $[\hat{\theta}_{\alpha}^{-}(z), \hat{\theta}_{\alpha}^{+}(z)]$ with pointwise frequentist coverage: $\liminf_{n \to \infty} \left\{ \mathbb{P}_{G} \left[\theta_{G}(z) \in [\hat{\theta}_{\alpha}^{-}(z), \hat{\theta}_{\alpha}^{+}(z)] \right] \right\} \ge 1 - \alpha$

and also with simultaneous coverage.

EB confidence intervals: statistical task

1) Known noise model: $Z_i \mid \mu_i \sim p(\cdot \mid \mu_i)$ 2) Class of priors: $\mu_i \sim G, \ G \in \mathcal{G}$

G is unknown \mathcal{G} is a pre-specified convex class of priors.

3) Empirical Bayes estimand: $\theta_G(z) = \mathbb{E}_G[h(\mu_i) \mid Z_i = z]$

CI: $[\hat{\theta}_{\alpha}^{-}(z), \hat{\theta}_{\alpha}^{+}(z)]$ with pointwise frequentist coverage: $\liminf_{n \to \infty} \left\{ \mathbb{P}_{G} \left[\theta_{G}(z) \in [\hat{\theta}_{\alpha}^{-}(z), \hat{\theta}_{\alpha}^{+}(z)] \right] \right\} \ge 1 - \alpha \quad \text{for all } G \in \mathscr{G}$

and also with simultaneous coverage.

Erik van Zwet¹^(D) | Simon Schwab^{2,3}^(D) | Stephen Senn⁴^(D)

Z-Score from an RCT: $Z_i \mid \mu_i \sim \mathcal{N}(\mu_i, 1)$

Erik van Zwet¹ | Simon Schwab^{2,3} | Stephen Senn⁴

Z-Score from an RCT: For inference:

$$Z_i \mid \mu_i \sim \mathcal{N}(\mu_i, 1)$$
$$\mu_i \sim G$$

Erik van Zwet¹ | Simon Schwab^{2,3} | Stephen Senn⁴

Z-Score from an RCT: For inference:

$$Z_i \mid \mu_i \sim \mathcal{N}(\mu_i, 1)$$
$$\mu_i \sim G$$

n = 23551 Z-scores from RCTs in healthcare (Cochrane Database of Systematic Reviews).



Erik van Zwet¹ | Simon Schwab^{2,3} | Stephen Senn⁴

Z-Score from an RCT: For inference:

$$Z_i \mid \mu_i \sim \mathcal{N}(\mu_i, 1)$$
$$\mu_i \sim G$$

n = 23551 Z-scores from RCTs in healthcare (Cochrane Database of Systematic Reviews).



For an RCT with Z-score $Z_i = -1.81$

Posterior mean: $\mathbb{E}_{\widehat{G}}[\mu \mid Z = -1.81] = -1.12$

Local false sign rate: $\mathbb{P}_{\widehat{G}} [\mu \ge 0 \mid Z = -1.81] = 0.091$

Erik van Zwet¹^(D) | Simon Schwab^{2,3}^(D) | Stephen Senn⁴^(D)

Z-Score from an RCT: For inference:

$$Z_i \mid \mu_i \sim \mathcal{N}(\mu_i, 1)$$
$$\mu_i \sim G$$

n = 23551 Z-scores from RCTs in healthcare (Cochrane Database of Systematic Reviews).



For an RCT with Z-score $Z_i = -1.81$

Posterior mean: $\mathbb{E}_{\widehat{G}}[\mu \mid Z = -1.81] = -1.12$

Local false sign rate: $\mathbb{P}_{\widehat{G}} [\mu \ge 0 \mid Z = -1.81] = 0.091$

What about confidence intervals?

Educational Testing Service

n = 12990 students Z_i : Score on multiple choice test with 20 questions (5 choices per question) Lord (1969), Lord and Cressie (1975), Lord and Stocking (1976)



Educational Testing Service

n = 12990 students Z_i : Score on multiple choice test with 20 questions (5 choices per question) Lord (1969), Lord and Cressie (1975), Lord and Stocking (1976)



Empirical Bayes model: $\mu_i \sim G$ $Z_i \mid \mu_i \sim \text{Binomial}(20, \mu_i)$ Posterior mean: $\theta_G(z) = \mathbb{E}_G[\mu_i \mid Z_i = z]$

Educational Testing Service

n = 12990 students Z_i : Score on multiple choice test with 20 questions (5 choices per question) Lord (1969), Lord and Cressie (1975), Lord and Stocking (1976)



Empirical Bayes model: $\mu_i \sim G$ $Z_i \mid \mu_i \sim \text{Binomial}(20, \mu_i)$ Posterior mean: $\theta_G(z) = \mathbb{E}_G[\mu_i \mid Z_i = z]$

Related work: EB confidence intervals

Efron (2014, 2016, 2019), Narasimhan and Efron (2020) Use a flexible parametric model for *G*, and estimate variance of point estimates $\hat{\theta}_G(z)$.

Related work: EB confidence intervals

Efron (2014, 2016, 2019), Narasimhan and Efron (2020)

Use a flexible parametric model for *G*, and estimate variance of point estimates $\hat{\theta}_G(z)$.

Nonparametrics:

Poisson model, Posterior Mean

Robbins (1980, Proceedings of the National Academy of Sciences) Karlis, Tzougas and Frangos (2018, Scandinavian Actuarial Journal)

Binomial model, Posterior Mean

Lord and Cressie (1975, Sankhyā Series B) Lord and Stocking (1976, Psychometrika)

Our work

A unified inference framework that works in the general nonparametric situation described.

Nonparametric EB estimation

Poisson model, Posterior Mean:



It is possible to estimate $\theta_G(z)$ at the parametric rate $1/\sqrt{n}$ [Robbins (1956), Lambert and Tierney (1984)]

Gaussian model, Posterior Mean:

It is possible to estimate $\theta_G(z)$ at the quasi-parametric rate $\log(n)^{3/4}/\sqrt{n}$ [Matias and Taupin (2004)]

Gaussian model, Local False Sign Rate: 🔀



XX

For Sobolev G, minimax rates for estimating $\theta_G(z)$ are polynomial in $1/\log(n)$ [Butucea and Comte (2009), Pensky (2017)]

Binomial model, Posterior Mean:

If we impose no restrictions on G, $\theta_G(z)$ is only partially identified [Robbins (1956)]

Why does the difficulty of estimation vary?

Hierarchical model:	$\mu_i \sim G,$	$Z_i \mid \mu_i \sim p(\cdot \mid \mu_i)$
Marginally:	$Z_i \sim f_G,$	$f_G(z) = \int p(z \mid \mu) dG(\mu)$

Why does the difficulty of estimation vary?

Hierarchical model:	$\mu_i \sim G,$	$Z_i \mid \mu_i \sim p(\cdot \mid \mu_i)$
Marginally:	$Z_i \sim f_G,$	$f_G(z) = \int p(z \mid \mu) dG(\mu)$

 $f_{G_1} \approx f_{G_2} \quad \rightleftharpoons \quad G_1 \approx G_2$

Why does the difficulty of estimation vary?

Hierarchical model:	$\mu_i \sim G,$	$Z_i \mid \mu_i \sim p(\cdot \mid \mu_i)$
Marginally:	$Z_i \sim f_G,$	$f_G(z) = \int p(z \mid \mu) dG(\mu)$

 $f_{G_1} \approx f_{G_2} \quad \rightleftharpoons \quad G_1 \approx G_2$ $f_{G_1} \approx f_{G_2} \quad \Longrightarrow \quad \theta_{G_1}(z) \approx \theta_{G_2}(z)$

Why is a unified inference approach difficult?

 Distributional theory for e.g., NPMLE is notoriously difficult. Some results for discrete distributions by Lambert and Tierney (1984), Böhning and Patilea (2005), Karlis and Patilea (2008).

Why is a unified inference approach difficult?

- Distributional theory for e.g., NPMLE is notoriously difficult. Some results for discrete distributions by Lambert and Tierney (1984), Böhning and Patilea (2005), Karlis and Patilea (2008).
- 2) Suppose we had asymptotic normality: $\widehat{\theta}_G(z) \pm 1.96 \cdot \widehat{se}$

Why is a unified inference approach difficult?

- Distributional theory for e.g., NPMLE is notoriously difficult. Some results for discrete distributions by Lambert and Tierney (1984), Böhning and Patilea (2005), Karlis and Patilea (2008).
- 2) Suppose we had asymptotic normality: $\widehat{\theta}_G(z) \pm 1.96 \cdot \widehat{se}$

Assumes that:



Not true under partial identification or when minimax estimation rates are slow!

"when the minimax convergence rate is slower than any algebraic rate, the optimal linear estimator must have maximum squared bias completely dominating the variance" [Cai and Low (2003)]

Instead: Bias-aware inference

[Armstrong and Kolesár (2018), Imbens and Manski (2004), Imbens and Wager (2018)]
$\mathcal{F}_{n}(\alpha) : \text{a set of distributions, such that} \\ \liminf_{n \to \infty} \left\{ \mathbb{P}_{G}[F_{G} \in \mathcal{F}_{n}(\alpha)] - (1 - \alpha) \right\} \geq 0.$

$$\mathcal{F}_{n}(\alpha) : \text{a set of distributions, such that} \\ \liminf_{n \to \infty} \left\{ \mathbb{P}_{G}[F_{G} \in \mathcal{F}_{n}(\alpha)] - (1 - \alpha) \right\} \geq 0.$$

Set of all priors consistent with the F-Localization

$$\mathcal{G}(\mathcal{F}_n(\alpha)) = \{ G \in \mathcal{G} : F_G \in \mathcal{F}_n(\alpha) \}$$

Confidence intervals for empirical Bayes estimand $\hat{\theta}_{\alpha}^{+}(z) = \sup\{\theta_{G}(z) : G \in \mathcal{G}(\mathcal{F}_{n}(\alpha))\}, \ \mathcal{F}_{\alpha}(z) = \left[\hat{\theta}_{\alpha}^{-}(z), \hat{\theta}_{\alpha}^{+}(z)\right]$

$$\mathcal{F}_{n}(\alpha) : \text{a set of distributions, such that} \\ \liminf_{n \to \infty} \left\{ \mathbb{P}_{G}[F_{G} \in \mathcal{F}_{n}(\alpha)] - (1 - \alpha) \right\} \geq 0.$$

Set of all priors consistent with the F-Localization

$$\mathcal{G}(\mathcal{F}_n(\alpha)) = \{ G \in \mathcal{G} : F_G \in \mathcal{F}_n(\alpha) \}$$

Confidence intervals for empirical Bayes estimand $\hat{\theta}_{\alpha}^{+}(z) = \sup\{\theta_{G}(z) : G \in \mathcal{G}(\mathcal{F}_{n}(\alpha))\}, \ \mathcal{F}_{\alpha}(z) = \left[\hat{\theta}_{\alpha}^{-}(z), \hat{\theta}_{\alpha}^{+}(z)\right]$

Proposition (I., Wager):
$$\liminf_{n \to \infty} \left\{ \mathbb{P}_{G} \left[\theta_{G}(z) \in [\widehat{\theta}_{\alpha}^{-}(z), \widehat{\theta}_{\alpha}^{+}(z)] \right] \right\} \ge 1 - \alpha$$

Robbins (1956), Anderson (1964), Deely and Kruse (1968), Romano and Wolf (2000), Stark (1992), Donoho and Reeves (2013), Kuusela and Stark (2017), Greenshtein and Itskov (2018), Brennan et al. (2020),

Dvoretzky–Kiefer–Wolfowitz

Theorem [Massart (1990)]: $\mathbb{P}_G[F_G \in \mathcal{F}_n(\alpha)] \ge 1 - \alpha$

Dvoretzky–Kiefer–Wolfowitz

Theorem [Massart (1990)]: $\mathbb{P}_G[F_G \in \mathcal{F}_n(\alpha)] \ge 1 - \alpha$

n = 12990 students Z_i : Score on multiple choice test with 20 questions (5 choices per question) $Z_i \mid \mu_i \sim \text{Binomial}(20, \mu_i)$















Choice of F-localization

DKW F-Localization: Universal and finite-sample χ^2 -F-Localization: When $Z_i \mid \mu_i \sim \text{Binomial}(N, \mu_i)$, $\mathscr{F}_n(\alpha) = \left\{ F \text{ with pmf} f \text{ on } \{0, ..., N\} : \sum_{z=0}^N \frac{(n\hat{f}_n(z) - nf(z))^2}{nf(z)} \le \chi^2_{N,1-\alpha} \right\}$ $\frac{1}{n} \sum_{i=1}^n \mathbf{1}(Z_i = z)$

Choice of F-localization





Choice of F-localization





Lord and Cressie (1975), Lord and Stocking (1976)



Beyond F-localization



Easy to construct

"Common sense approach"

Simultaneous coverage



What about power? Power strong only if F-localization gives tight characterization of uncertainty. Not true in general, because of simultaneous coverage.

Conservative!

Choice of F-Localization matters.

(Affine Minimax Anderson Rubin Inference)

(Affine Minimax Anderson Rubin Inference)

$$H_0: \theta_G(z) = c$$

(Affine Minimax Anderson Rubin Inference)

$$H_0: \theta_G(z) = c$$

 $\theta_G(z) = \mathbb{E}_G[h(\mu_i) \mid Z_i = z] = \frac{\int h(\mu) p(z \mid \mu) \, dG(\mu)}{\int p(z \mid \mu) \, dG(\mu)} = \frac{a_G(z)}{f_G(z)}$

$$H_0: \theta_G(z) = c \iff H_0: a_G(z) - c \cdot f_G(z) = 0$$

(Affine Minimax Anderson Rubin Inference)

$$H_0: \theta_G(z) = c$$

$$\theta_G(z) = \mathbb{E}_G[h(\mu_i) \mid Z_i = z] = \frac{\int h(\mu) p(z \mid \mu) \, dG(\mu)}{\int p(z \mid \mu) \, dG(\mu)} = \frac{a_G(z)}{f_G(z)}$$

$$H_0: \theta_G(z) = c \iff H_0: a_G(z) - c \cdot f_G(z) = 0$$

Gu Linear in G

Jiaying Gu Noack and Rothe (2019) Anderson and Rubin (1949) Fieller (1954)

Upshot:

Can focus on inference for linear functionals!

 $L(\mathbf{G})$

(Affine Minimax Anderson Rubin Inference)

Seek to conduct inference for the linear functional

L(G)

Ansatz: Consider affine estimators of the form

$$\widehat{L} = \widehat{L}(G) = \frac{1}{n} \sum_{i=1}^{n} Q_n(Z_i).$$

Why?

- Convenient computationally.
- Class of estimators that includes kernel density estimators and minimax optimal estimators of linear functionals in the Gaussian deconvolution problem [Butucea and Comte (2009), Pensky (2017)]

(Affine Minimax Anderson Rubin Inference)



(Affine Minimax Anderson Rubin Inference)

$$\hat{L} = \frac{1}{n} \sum_{i=1}^{n} Q_n(Z_i). \quad \text{How to choose } Q = Q_n?$$
$$\min_{Q:\mathbb{R}\to\mathbb{R}} \left\{ \max_{G\in\mathscr{G}_n} \left\{ \text{Bias}_G[\hat{L}]^2 \right\} + \widehat{\text{Var}}[\hat{L}] \right\}$$

Bias: $\max_{G \in \mathscr{G}_n} \left\{ \text{Bias}_G[\widehat{L}]^2 \right\} = \max_{G \in \mathscr{G}_n} \left\{ \left(\mathbb{E}_G[Q(Z_i)] - L(G) \right)^2 \right\}$ $\mathscr{G}_n = \{ G \in \mathscr{G} : F_G \in \mathscr{F}_n(\alpha_n) \}, \ \alpha_n \to 0$

(Affine Minimax Anderson Rubin Inference)

$$\hat{L} = \frac{1}{n} \sum_{i=1}^{n} Q_n(Z_i). \quad \text{How to choose } Q = Q_n?$$
$$\min_{Q:\mathbb{R}\to\mathbb{R}} \left\{ \max_{G\in\mathscr{G}_n} \left\{ \text{Bias}_G[\hat{L}]^2 \right\} + \widehat{\text{Var}}[\hat{L}] \right\}$$

Bias: $\max_{G \in \mathscr{G}_n} \left\{ \operatorname{Bias}_G[\widehat{L}]^2 \right\} = \max_{G \in \mathscr{G}_n} \left\{ \left(\mathbb{E}_G[Q(Z_i)] - L(G) \right)^2 \right\}$ $\mathscr{G}_n = \{ G \in \mathscr{G} : F_G \in \mathscr{F}_n(\alpha_n) \}, \ \alpha_n \to 0$

Variance: We use a pilot estimator $\bar{f}_n(z)$ of the marginal density. $\widehat{\text{Var}}[\hat{L}] = \frac{1}{n} \left\{ \int Q^2(z) \bar{f}_n(z) d\lambda(z) - \left(\int Q(z) \bar{f}_n(z) d\lambda(z) \right)^2 \right\}$

Minimax optimization

$$\min_{Q:\mathbb{R}\to\mathbb{R}}\left\{\max_{G\in\mathscr{G}_n}\left\{\mathsf{Bias}_G[\widehat{L}]^2\right\}+\widehat{\mathsf{Var}}[\widehat{L}]\right\}\right\}$$

The above optimization problem can be solved by building upon:

Donoho (1994), Armstrong and Kolesár (2018, 2020, 2021, ...), Sacks and Ylvisacker (1978), Ibragimov and Hasminskii (1984), Donoho and Liu (1989, 1991), Low (1995), Zhao (1997), Cai and Low (2003, 2004), ...

Bias-aware inference

Armstrong and Kolesár (2018), Imbens and Manski (2004), Imbens and Wager (2018)

Estimate
$$L(G)$$
 by $\widehat{L} = \sum_{i=1}^{n} Q_n(Z_i) / n$
 $\widehat{V} = \widehat{\operatorname{Var}}(Q_n(Z_i)) / n$ $\widehat{B} = \sup_{G \in \mathscr{G}_n} \left| \operatorname{Bias}_G[\widehat{L}] \right|$

 $\widehat{L} \pm t_{\alpha}(\widehat{B}, \widehat{V}) \qquad t_{\alpha}(B, V) = \inf\{t : \mathbb{P}[|b + V^{1/2}W| > t] < \alpha \text{ for all } |b| \le B\} \\ W \sim \mathcal{N}(0, 1)$

Theorem (I., Wager), Informal

Suppose we choose $Q_n(\cdot)$ as piecewise constant outside [-M, M]. If $\mathscr{F}_n(\alpha_n), \bar{f}_n(\cdot)$ are constructed by sample splitting, then in the Binomial, Gaussian and Poisson empirical Bayes models, our intervals have asymptotic coverage $\geq 1 - \alpha$.

Confidence intervals in the RCT example

$$n = 23,551 \text{ RCTs} \qquad \begin{array}{l} \mu_i \sim G, \\ Z_i \mid \mu_i \sim \mathcal{N}(\mu_i, 1) \end{array}$$

- $\mathscr{G} \stackrel{\circ}{=} Scale Mixture of centered Gaussians$
 - $= \left\{ G \text{ with Leb. density } g(\mu) = \int \frac{1}{\tau} \varphi\left(\frac{\mu}{\tau}\right) d\Pi(\tau), \text{ } \Pi \text{ supported on } [0.1, 60] \right\}$

Confidence intervals in the RCT example

$$n = 23,551 \text{ RCTs} \qquad \begin{array}{l} \mu_i \sim G, \\ Z_i \mid \mu_i \sim \mathcal{N}(\mu_i, 1) \end{array}$$

 $\mathscr{G} \stackrel{c}{=} \text{Scale Mixture of centered Gaussians}$ $= \left\{ G \text{ with Leb. density } g(\mu) = \int \frac{1}{\tau} \varphi\left(\frac{\mu}{\tau}\right) d\Pi(\tau), \ \Pi \text{ supported on } [0.1, 60] \right\}$



Confidence intervals in the RCT example

$$n = 23,551 \text{ RCTs} \qquad \begin{array}{l} \mu_i \sim G, \\ Z_i \mid \mu_i \sim \mathcal{N}(\mu_i, 1) \end{array}$$

 $\mathscr{G} \stackrel{c}{=} \text{Scale Mixture of centered Gaussians}$ $= \left\{ G \text{ with Leb. density } g(\mu) = \int \frac{1}{\tau} \varphi\left(\frac{\mu}{\tau}\right) d\Pi(\tau), \ \Pi \text{ supported on } [0.1, 60] \right\}$





 $\mathscr{G}_{loc} \triangleq \text{Location Mixture of Gaussians} \qquad \boxed{\tau = 0.25}$ $= \left\{ G \text{ with Leb. density } g(\mu) = \int \frac{1}{\tau} \varphi \left(\frac{\mu - u}{\tau} \right) d\Pi(u), \ \Pi \text{ supported on } [-3,3] \right\}$ [Magder and Zeger (1996), Cordy and Thomas (1997)] $\mathscr{G}_{sc} \triangleq \text{Scale Mixture of centered Gaussians}$ $= \left\{ G \text{ with Leb. density } g(\mu) = \int \frac{1}{\tau} \varphi \left(\frac{\mu}{\tau} \right) d\Pi(\tau), \ \Pi \text{ supported on } [0.1, 15.6] \right\}$





 $\mathscr{G}_{loc} \doteq$ Location Mixture of Gaussians





 $\mathscr{G}_{loc} \doteq$ Location Mixture of Gaussians





Coverage





Conclusion

Even if n is large, there may be substantial uncertainty in the estimation of empirical Bayes quantities.

Here we describe two approaches to conduct inference for empirical Bayes estimands that can accompany empirical Bayes point estimation in practice.

Our approaches can also be used to assess the sensitivity to the choice of prior class \mathcal{G} .

Thank you for your attention!
AMARI

(Affine Minimax Anderson Rubin Inference)

Seek to conduct inference for the linear functional

L(G)

We consider affine estimators

Butucea and Comte (2009), Pensky (2017)

We choose the affine estimator by minimax optimization

Donoho (1994), Armstrong and Kolesár (2018, 2020, 2021, ...), Sacks and Ylvisacker (1978), Ibragimov and Hasminskii (1984), Donoho and Liu (1989, 1991), Low (1995), Zhao (1997), Cai and Low (2003, 2004), ...

We conduct bias-aware inference

Armstrong and Kolesár (2018), Imbens and Manski (2004), Imbens and Wager (2018)

Suppose $G_1, G_2 \in \mathcal{G}_n$ solve the following optimization problem ($\delta > 0$):

maximize
$$L(G_1) - L(G_2)$$
 s.t. $\int \frac{(f_{G_1}(z) - f_{G_2}(z))^2}{\bar{f}_n(z)} d\lambda(z) \le \frac{\delta^2}{n}$

We call $\text{Conv}(G_1, G_2) = \{\eta G_1 + (1 - \eta)G_2 : \eta \in [0, 1]\}$ a hardest 1-dimensional subfamily.

Suppose $G_1, G_2 \in \mathcal{G}_n$ solve the following optimization problem ($\delta > 0$):

maximize
$$L(G_1) - L(G_2)$$
 s.t. $\int \frac{(f_{G_1}(z) - f_{G_2}(z))^2}{\bar{f}_n(z)} d\lambda(z) \le \frac{\delta^2}{n}$

We call $\text{Conv}(G_1, G_2) = \left\{ \eta G_1 + (1 - \eta)G_2 : \eta \in [0, 1] \right\}$ a hardest 1-dimensional subfamily.

$$\min_{Q:\mathbb{R}\to\mathbb{R}}\max_{G\in\mathscr{G}_n}\left\{\mathsf{Bias}_G[\widehat{L}]^2 + \Gamma_n \cdot \widehat{\mathsf{Var}}[\widehat{L}]\right\}$$

Suppose $G_1, G_2 \in \mathcal{G}_n$ solve the following optimization problem ($\delta > 0$):

maximize
$$L(G_1) - L(G_2)$$
 s.t. $\int \frac{(f_{G_1}(z) - f_{G_2}(z))^2}{\bar{f}_n(z)} d\lambda(z) \le \frac{\delta^2}{n}$

We call $\text{Conv}(G_1, G_2) = \left\{ \eta G_1 + (1 - \eta) G_2 : \eta \in [0, 1] \right\}$ a hardest 1-dimensional subfamily.

$$\min_{\substack{Q:\mathbb{R}\to\mathbb{R}\ G\ \mathcal{G}_n}} \max_{\boldsymbol{G}} \left\{ \operatorname{Bias}_{\boldsymbol{G}}[\widehat{L}]^2 + \Gamma_n(\boldsymbol{\delta}) \cdot \widehat{\operatorname{Var}}[\widehat{L}] \right\}$$
$$\boldsymbol{G} \in \operatorname{Conv}(G_1, G_2)$$

The above optimization problem can be solved analytically!

Suppose $G_1, G_2 \in \mathcal{G}_n$ solve the following optimization problem ($\delta > 0$):

maximize
$$L(G_1) - L(G_2)$$
 s.t. $\int \frac{(f_{G_1}(z) - f_{G_2}(z))^2}{\bar{f}_n(z)} d\lambda(z) \le \frac{\delta^2}{n}$

We call $\text{Conv}(G_1, G_2) = \left\{ \eta G_1 + (1 - \eta) G_2 : \eta \in [0, 1] \right\}$ a hardest 1-dimensional subfamily.

$$\min_{Q:\mathbb{R}\to\mathbb{R}} \max_{G \in \mathcal{G}_n} \left\{ \operatorname{Bias}_G[\widehat{L}]^2 + \Gamma_n(\delta) \cdot \widehat{\operatorname{Var}}[\widehat{L}] \right\}$$
$$G \in \operatorname{Conv}(G_1, G_2)$$

The above optimization problem can be solved analytically!

Donoho (1994), Armstrong and Kolesár (2018, 2020, 2021, ...), Sacks and Ylvisacker (1978), Ibragimov and Hasminskii (1984), Donoho and Liu (1989, 1991), Low (1995), Zhao (1997), Cai and Low (2003, 2004), ...