Bootstrap Methods and their Application

A. C. Davison

©2021

A short course based on 'Bootstrap Methods and their Application',
by A. C. Davison and D. V. Hinkley, Cambridge University Press, 1997

http://stat.epfl.ch

Bootstrap: Motivation **3**

Bootstrap: Basic Notions **9**

Confidence Intervals **35**

Several Samples **53**

Variance Estimation **60**

Tests **73**

Regression **93**

David Hinkley, 1944–2019

February 2021 – slide 2

Bootstrap: Motivation

slide 3

AIDS data

Diagnosis period		Reporting-delay interval (quarters):									Total reports to end of 1992
Year	Quarter	$0^\dagger$	1	2	3	4	5	6	$\cdots$	$\geq$14	
1988	1	31	80	16	9	3	2	8	$\cdots$	6	174
	2	26	99	27	9	8	11	3	$\cdots$	3	211
	3	31	95	35	13	18	4	6	$\cdots$	3	224
	4	36	77	20	26	11	3	8	$\cdots$	2	205
1989	1	32	92	32	10	12	19	12	$\cdots$	2	224
	2	15	92	14	27	22	21	12	$\cdots$	1	219
	3	34	104	29	31	18	8	6	$\cdots$		253
	4	38	101	34	18	9	15	6	$\cdots$		233
1990	1	31	124	47	24	11	15	8	$\cdots$		281
	2	32	132	36	10	9	7	6	$\cdots$		245
	3	49	107	51	17	15	8	9	$\cdots$		260
	4	44	153	41	16	11	6	5	$\cdots$		285
1991	1	41	137	29	33	7	11	6	$\cdots$		271
	2	56	124	39	14	12	7	10	$\cdots$		263
	3	53	175	35	17	13	11	2			306
	4	63	135	24	23	12	1				258
1992	1	71	161	48	25	5					310
	2	95	178	39	6						318
	3	76	181	16							273
	4	67	66								133

February 2021 – slide 4

AIDS data

- ☐ UK AIDS diagnoses 1988–1992.
- ☐ Reporting delay up to several years!
- ☐ Problem: predict state of epidemic at end 1992, with realistic statement of uncertainty.
- ☐ Simple model: number of reports in row j and column k Poisson, mean

$$\mu_{jk} = \exp(\alpha_j + \beta_k).$$

- ☐ Unreported diagnoses in period j Poisson, mean

$$\sum_{k \text{ unobs}} \mu_{jk} = \exp(\alpha_j) \sum_{k \text{ unobs}} \exp(\beta_k).$$

- ☐ Estimate total unreported diagnoses from period j by replacing α_j and β_k by MLEs.
  - How reliable are these predictions?
  - How sensible is the Poisson model?

AIDS data

- ☐ Data (+), fits of simple model (solid), complex model (dots)
- ☐ Variance formulae could be derived — painful! but useful?
- ☐ Effects of overdispersion, complex model, ...?

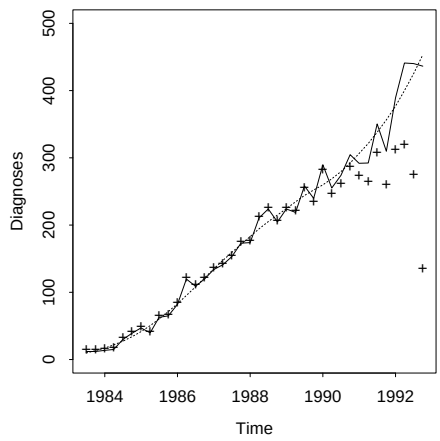

Goal

Find reliable **frequentist** assessment of uncertainty when

- ☐ estimator complex
- ☐ data complex
- ☐ sample size small
- ☐ model non-standard

Reading

- ☐ Chernick (2008) *Bootstrap Methods: A Guide for Practitioner and Researchers*. Second edition. Wiley
- ☐ Davison and Hinkley (1997) *Bootstrap Methods and their Application*. Cambridge University Press
- ☐ Efron and Tibshirani (1993) *An Introduction to the Bootstrap*. Chapman & Hall
- ☐ Hall (1992) *The Bootstrap and Edgeworth Expansion*. Springer
- ☐ Lahiri (2003) *Resampling Methods for Dependent Data*. Springer
- ☐ Lunneborg (2000) *Data Analysis by Resampling: Concepts and Applications*. Duxbury Press
- ☐ Manly and Navarro Alberto (2020) *Randomisation, Bootstrap and Monte Carlo Methods in Biology*. 4th edition. Chapman & Hall
- ☐ Politis, Romano and Wolf (1999) *Subsampling*. Springer
- ☐ Shao and Tu (1995) *The Jackknife and Bootstrap*. Springer
- ☐ Bootstrap anniversary issue of *Statistical Science*, 2003

Bootstrap: Basic Notions

Handedness data

Table 1: Data from a study of handedness; `hand` is an integer measure of handedness, and `dnan` a genetic measure. Data due to Dr Gordon Claridge, University of Oxford.

	`dnan`	`hand`		`dnan`	`hand`		`dnan`	`hand`		`dnan`	`hand`
1	13	1	11	28	1	21	29	2	31	31	1
2	18	1	12	28	2	22	29	1	32	31	2
3	20	3	13	28	1	23	29	1	33	33	6
4	21	1	14	28	4	24	30	1	34	33	1
5	21	1	15	28	1	25	30	1	35	34	1
6	24	1	16	28	1	26	30	2	36	41	4
7	24	1	17	29	1	27	30	1	37	44	8
8	27	1	18	29	1	28	31	1			
9	28	1	19	29	1	29	31	1			
10	28	2	20	29	2	30	31	1			

Handedness data

Scatter plot of handedness data. The numbers show the multiplicities of the observations.

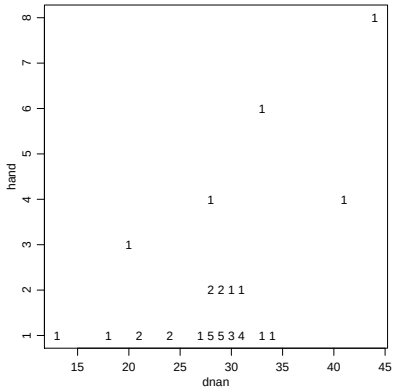

Handedness data

- ☐ Is there dependence between dnan and hand for these $n = 37$ individuals?
- ☐ Sample product-moment correlation coefficient is

$$\begin{aligned}\widehat{\theta} &= \frac{\sum_{j=1}^{n}(\mathtt{dnan}_j - \overline{\mathtt{dnan}})(\mathtt{hand}_j - \overline{\mathtt{hand}})}{\left\{\sum_{j=1}^{n}(\mathtt{dnan}_j - \overline{\mathtt{dnan}})^2 \sum_{j=1}^{n}(\mathtt{hand}_j - \overline{\mathtt{hand}})^2\right\}^{1/2}} \\ &= 0.509.\end{aligned}$$

- ☐ Standard confidence interval (based on bivariate normal population) gives 95% CI $(0.221,\ 0.715)$.
- ☐ Data not bivariate normal!
- ☐ What is the status of the interval? Can we do better?

Frequentist inference

- ☐ Estimator $\widehat{\theta}$ for unknown parameter θ.
- ☐ Statistical model: data $y_1, \ldots, y_n \overset{\text{iid}}{\sim} F$, unknown
- ☐ Handedness data
  - $y = (\mathtt{dnan}, \mathtt{hand})$
  - F puts probability mass on subset of $\mathbb{R}^2$
  - $\widehat{\theta}$ is correlation coefficient
- ☐ Key questions
  - How does $\widehat{\theta}$ behave when samples are repeatedly taken from F?
  - How can we use knowledge of this to learn about θ?
- ☐ Thought experiment: imagine F known — could answer question by
  - analytical (mathematical) calculation
  - simulation

Simulation with F known

- ☐ For $r = 1, \ldots, R$:
  - generate random sample $y_1^*, \ldots, y_n^* \overset{\text{iid}}{\sim} F$;
  - compute $\widehat{\theta}_r^*$ using $y_1^*, \ldots, y_n^*$;
- ☐ Output after R iterations:

$$\widehat{\theta}_1^*, \widehat{\theta}_2^*, \ldots, \widehat{\theta}_R^*$$

- ☐ Use $\widehat{\theta}_1^*, \widehat{\theta}_2^*, \ldots, \widehat{\theta}_R^*$ to estimate sampling distribution of $\widehat{\theta}$ (histogram, density estimate, ...)
- ☐ If $R \to \infty$, then get perfect match to theoretical calculation (if available): Monte Carlo error disappears completely
- ☐ In practice R is finite, so some error remains

Handedness data: Fitted bivariate normal model

Contours of bivariate normal distribution fitted to handedness data; parameter estimates are $\widehat{\mu}_1 = 28.5$, $\widehat{\mu}_2 = 1.7$, $\widehat{\sigma}_1 = 5.4$, $\widehat{\sigma}_2 = 1.5$, $\widehat{\rho} = 0.509$. The data are also shown.

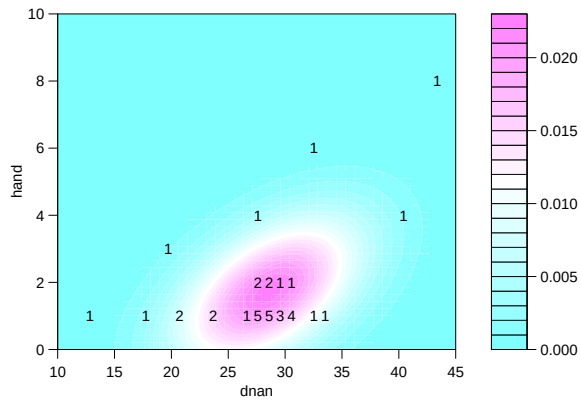

Handedness data: Parametric bootstrap samples

Left: original data, with jittered vertical values. Centre and right: two samples generated from the fitted bivariate normal distribution.

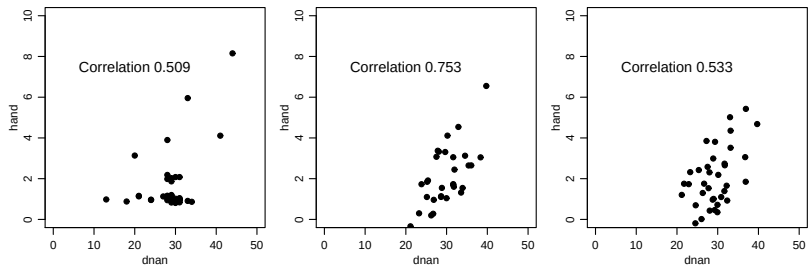

Handedness data: Correlation coefficient

Bootstrap distributions with $R = 10000$. Left: simulation from fitted bivariate normal distribution. Right: simulation from the data by bootstrap resampling. The lines show the theoretical probability density function of the correlation coefficient under sampling from a fitted bivariate normal distribution.

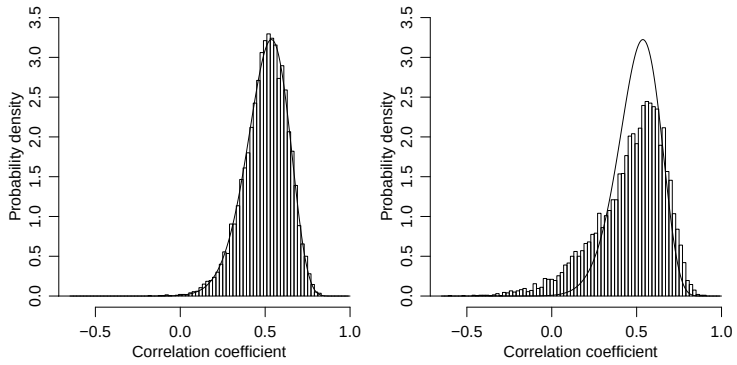

F unknown

- ☐ Replace unknown F by estimate $\widehat{F}$ obtained
  - **parametrically** — e.g., maximum likelihood or robust fit of distribution $F(y) = F(y; \psi)$ (exponential, bivariate normal, ...)
  - **nonparametrically** — using empirical distribution function (EDF) of original data $y_1, \ldots, y_n$, which puts mass $1/n$ on each of the y_j
- ☐ Algorithm: For $r = 1, \ldots, R$:
  - generate random sample $y_1^*, \ldots, y_n^* \overset{\text{iid}}{\sim} \widehat{F}$;
  - compute $\widehat{\theta}_r^*$ using $y_1^*, \ldots, y_n^*$;
- ☐ Output after R iterations:

$$\widehat{\theta}_1^*, \widehat{\theta}_2^*, \ldots, \widehat{\theta}_R^*$$

Nonparametric bootstrap

- ☐ Bootstrap (re)sample $y_1^*, \ldots, y_n^* \overset{\text{iid}}{\sim} \widehat{F}$, where $\widehat{F}$ is EDF of $y_1, \ldots, y_n$
  - If $f_i^* = \#\{y_j^* = y_i\}$, then $(f_1^*, \ldots, f_n^*)$ has the multinomial distribution with denominator n and probability vector $(n^{-1}, \ldots, n^{-1})$.
  - Repetitions will occur!
- ☐ Compute bootstrap $\widehat{\theta}^*$ using $y_1^*, \ldots, y_n^*$.
- ☐ For handedness data take $n = 37$ pairs $y^* = (\texttt{dnan}, \texttt{hand})^*$ with equal probabilities $1/37$ and replacement from original pairs $(\texttt{dnan}, \texttt{hand})$
- ☐ Repeat this R times, to get $\widehat{\theta}_1^*, \ldots, \widehat{\theta}_R^*$
- ☐ See picture
- ☐ Results quite different from parametric simulation — why?

Handedness data: Correlation coefficient

Bootstrap distributions with $R = 10000$. Left: simulation from fitted bivariate normal distribution. Right: simulation from the data by bootstrap resampling. The lines show the theoretical probability density function of the correlation coefficient under sampling from a fitted bivariate normal distribution.

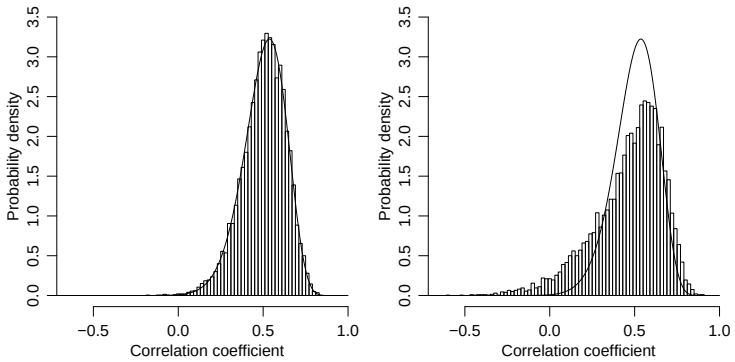

Handedness data: Bootstrap samples

Left: original data, with jittered vertical values. Centre and right: two bootstrap samples, with jittered vertical values.

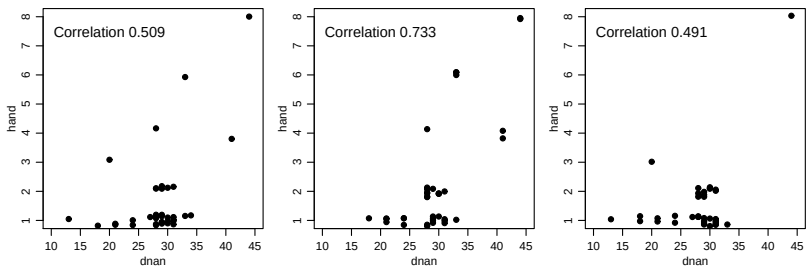

Using the $\widehat{\theta}^*$: Bias

- ☐ **Bootstrap replicates** $\widehat{\theta}_r^*$ used to estimate properties of $\widehat{\theta}$.
- ☐ Write $\theta = \theta(F)$ to emphasize dependence on F
- ☐ **Bias** of $\widehat{\theta}$ as estimator of θ is

$$\beta(F) = \mathrm{E}(\widehat{\theta} \mid y_1, \ldots, y_n \overset{\text{iid}}{\sim} F) - \theta(F)$$

estimated by replacing unknown F by known estimate $\widehat{F}$:

$$\beta(\widehat{F}) = \mathrm{E}(\widehat{\theta} \mid y_1, \ldots, y_n \overset{\text{iid}}{\sim} \widehat{F}) - \theta(\widehat{F})$$

- ☐ Replace theoretical expectation $\mathrm{E}(\cdot)$ by empirical average:

$$\beta(\widehat{F}) \approx b = \overline{\widehat{\theta}^*} - \widehat{\theta} = R^{-1} \sum_{r=1}^{R} \widehat{\theta}_r^* - \widehat{\theta}$$

Using the $\widehat{\theta}^*$: Variance/Quantiles

- ☐ Estimate **variance** $\nu(F) = \mathrm{var}(\widehat{\theta} \mid F)$ by

$$v = \frac{1}{R-1} \sum_{r=1}^{R} \left(\widehat{\theta}_r^* - \overline{\widehat{\theta}^*} \right)^2$$

- ☐ Estimate **quantiles** of $\widehat{\theta}$ by taking empirical quantiles of

$$\widehat{\theta}_1^*, \ldots, \widehat{\theta}_R^*$$

- ☐ For handedness data, 10,000 replicates shown earlier give

$$b = -0.046, \quad v = 0.043 = 0.205^2$$

Handedness data

Summaries of the $\widehat{\theta}^*$. Left: histogram, with vertical line showing $\widehat{\theta}$. Right: normal Q–Q plot of $\widehat{\theta}^*$.

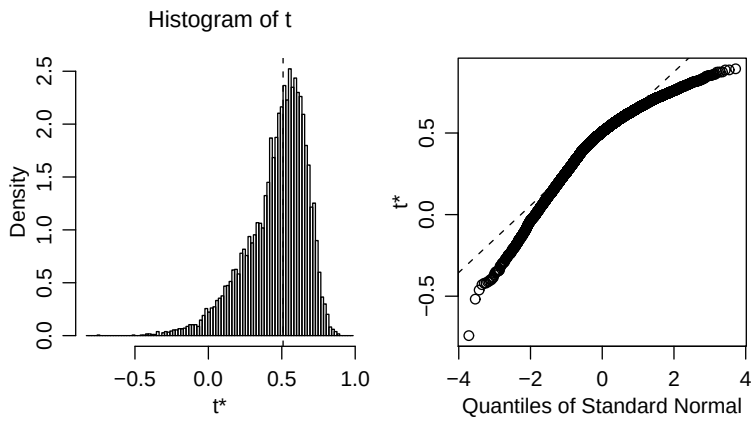

Common questions

- ☐ **How big should n be?** — depends on the context
- ☐ **What if the sample is unrepresentative?** — this is always a potential problem in statistics, not specific to resampling methods.
- ☐ **How big should R be?** — at least 1000 for most purposes
- ☐ **Why take resamples of size n?**
  - We usually want to mimic the sampling properties of samples like the original one, so take resamples of size n,
  - but sometimes we take resamples of size $m \ll n$ in order to achieve validity of the bootstrap—e.g., for extreme quantiles.
- ☐ **Why resample from the EDF?**
  - The EDF is the nonparametric MLE of F, so is a natural choice when there are no restrictions imposed on F, but
  - in some cases (e.g., testing) it may be appropriate to resample from a 'tilted' version of $\widehat{F}$,
  - in other cases it may be useful to smooth $\widehat{F}$

How big should n be?

- For the **average** $\widehat{\theta} = \overline{y}$, the number of distinct samples is

$$m_n = \binom{2n-1}{n},$$

  the most probable of which has probability $p_n = n!/n^n$
- So for $n > 12$, we have $m_n > 10^6$ and $p_n < 6 \times 10^{-5}$.

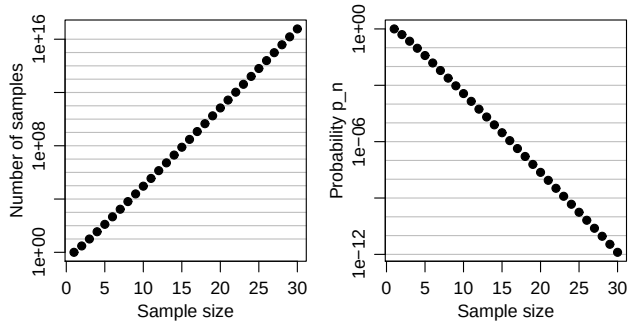

How big should n be?

- For the **median** of a sample of size $n = 2m + 1$, the possible distinct values of $\widehat{\theta}^*$ are $y_{(1)} < \cdots < y_{(n)}$, and

$$\mathrm{P}^*(\widehat{\theta}^* > y_{(l)}) = \sum_{r=0}^{m} \binom{n}{r} \left(\frac{l}{n}\right)^r \left(1 - \frac{l}{n}\right)^{n-r},$$

  so exact calculations of the variance etc. are possible.
- However the median is very vulnerable to bad values in the sample.
- Exact, approximate and mean bootstrap estimates of variance ($\times 10^{-2}$) of sample median, based on 10000 data sets of sizes $n = 11, 21$. The effective degrees of freedom of bootstrap variances uses a χ^2 approximation to their distribution.

	Normal		t_3		Cauchy	
	11	21	11	21	11	21
Theoretical	14.3	7.5	16.8	8.8	22.4	11.7
Exact	13.9	7.3	19.1	9.5	38.3	14.6
Mean bootstrap	17.2	8.8	25.9	11.4	14000	22.8
Effective df	4.3	5.4	3.2	4.9	0.002	0.5

What about outliers?

- Bad sample values not specific to bootstrap methods—need to consider them in any analysis
- Can reduce their effects by using robust estimators, but beware, as outliers may affect bootstrapped estimators even if they don't affect original value
- As example, consider 20% trimmed average with $n = 11$ and two outliers

How many bootstraps?

- ☐ Must estimate moments and quantiles of $\widehat{\theta}$ and derived quantities. Nowadays often feasible to take $R \geq 5000$
- ☐ Need $R \geq 200$ to estimate bias, variance, etc.
- ☐ Need $R \gg 100$, prefer $R \geq 1000$ to estimate quantiles needed for 95% confidence intervals

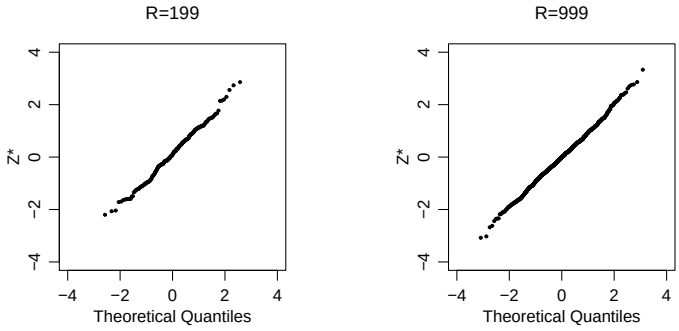

Resamples of size n?

- ☐ Exponential sample of size $n = 1000$
- ☐ Distribution of $n \min(Y_1, \ldots, Y_n)$ is $\exp(1)$
- ☐ Resampling distribution $m \min(Y_1^*, \ldots, Y_m^*)$ using resamples of size $m = 1000$, 100, 50
- ☐ To avoid discreteness must choose $m \ll n$, but how?

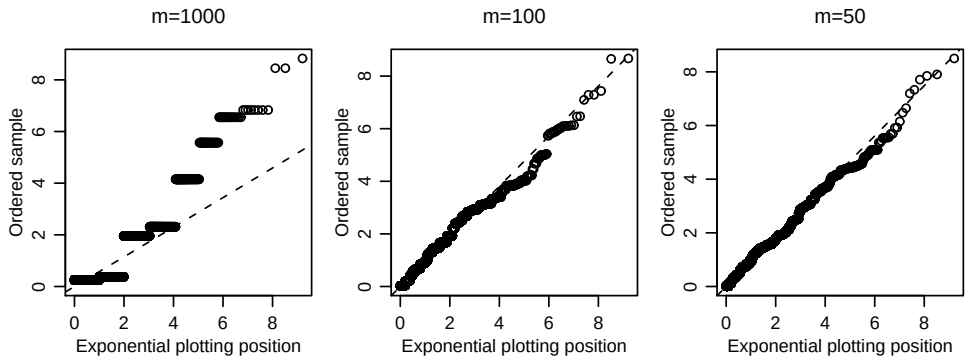

Variants of $\widehat{F}$?

- Can be useful to simulate from a smoothed EDF, given by

$$Y^* = y_{j^*} + h\varepsilon^*, \quad \varepsilon^* \sim \mathcal{N}(0,1) \perp\!\!\!\perp j^* \sim U\{1,\ldots,n\},$$

  equivalent to simulating from a kernel density estimate of F'. Below, with $h = 0.1$ (red) and $h = 0.5$ (blue).
- Since $\mathrm{var}^*(Y^*) = \widehat{\sigma}^2 + h^2$, may prefer a shrunk smoothed estimate, given by

$$Y^* = \overline{y} + \frac{(y_{j^*} - \overline{y}) + h\varepsilon^*}{(1 + h^2/\widehat{\sigma}^2)^{1/2}}.$$

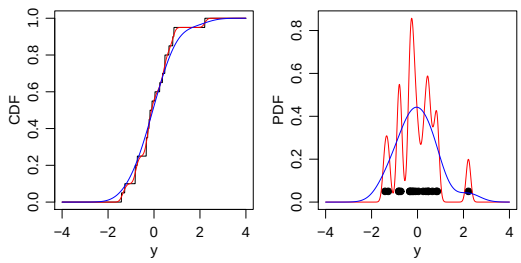

When does the bootstrap work?

- 'Work' might mean the bootstrap gives
  - **reliable** answers when used in practice, or
  - **mathematically correct** answers under 'suitable' regularity conditions.
- For the second of these, suppose we seek to estimate properties of a standardized quantity $Q = q(Y_1,\ldots,Y_n;F)$, maybe $Q = n^{1/2}(\overline{Y} - \theta)$. Let $n \to \infty$ to get limiting results for the distribution function

$$G_{F,n}(q) = \mathrm{P}\left\{Q(Y_1,\ldots,Y_n;F) \le q \mid F\right\},$$

  where the conditioning on F indicates that $Y_1,\ldots,Y_n$ is a random sample from F.
- Bootstrap estimate of this is

$$G_{\widehat{F},n}(q) = \mathrm{P}\left\{Q(Y_1^*,\ldots,Y_n^*;\widehat{F}) \le q \mid \widehat{F}\right\}$$

  where $Q(Y_1^*,\ldots,Y_n^*;\widehat{F}) = n^{1/2}(\overline{Y}^* - \overline{y})$.
- We need conditions under which $G_{\widehat{F},n} \to G_{F,n}$ as $n \to \infty$.

Regularity conditions

- ☐ Suppose that the true distribution F is surrounded by a neighbourhood $\mathcal{N}$ in a suitable space of distributions, and that as $n \to \infty$, $\widehat{F}$ eventually falls into $\mathcal{N}$ with probability one. Then regularity conditions are:
  1. for any $A \in \mathcal{N}$, $G_{A,n}$ must converge weakly to a limit $G_{A,\infty}$;
  2. this convergence must be uniform on $\mathcal{N}$; and
  3. the function mapping A to $G_{A,\infty}$ must be continuous.
- ☐ Weak convergence of $G_{A,n}$ to $G_{A,\infty}$ means that for all integrable $h(\cdot)$,

$$\int h(u)\, dG_{A,n}(u) \quad \to \quad \int h(u)\, dG_{A,\infty}(u) \quad \text{as} \quad n \to \infty.$$

- ☐ Under these conditions the bootstrap is **consistent**: for any q and $\varepsilon > 0$, $\mathrm{P}\{|G_{\widehat{F},n}(q) - G_{F,\infty}(q)| > \varepsilon\} \to 0$ as $n \to \infty$.
- ☐ The first condition ensures that there is a limit for $G_{F,n}$ to converge to, and would be needed even if $\widehat{F} = F$ for every $n \geq n'$, for some n'.
- ☐ As n increases, $\widehat{F}$ changes, so the second and third conditions are needed to ensure that $G_{\widehat{F},n}$ approaches $G_{F,\infty}$ along every possible sequence of $\widehat{F}$s.
- ☐ If any one of these conditions fails, the bootstrap can fail. For the minimum, the convergence is not uniform on suitable neighbourhoods of F.

Key points

- ☐ **Estimator is algorithm**
  - applied to original data $y_1, \ldots, y_n$ gives original $\widehat{\theta}$
  - applied to simulated data $y_1^*, \ldots, y_n^*$ gives $\widehat{\theta}^*$
  - $\widehat{\theta}$ can be of (almost) any complexity
  - for more sophisticated ideas to work, $\widehat{\theta}$ must often be smooth function of data
- ☐ **Sample is used to estimate F**
  - $\widehat{F} \approx F$ — heroic assumption
- ☐ **Simulation replaces theoretical calculation**
  - removes need for mathematical skill
  - does not remove need for thought
  - check code **very** carefully — garbage in, garbage out!
- ☐ **Two sources of error**
  - statistical ($\widehat{F} \neq F$) — reduce by thought
  - simulation ($R \neq \infty$) — reduce by taking R large (enough)

Desiderata

- Recall that a $(1-\alpha)$ **upper confidence limit** for an unknown scalar parameter θ based on data Y is a random variable $\theta_\alpha = \theta_\alpha(Y)$ with the property that

$$\mathrm{P}\left(\theta \le \theta_\alpha\right) = \alpha, \quad 0 < \alpha < 1, \theta \in \Omega_\theta. \tag{1}$$

- We may seek invariance to monotone transformations $\psi = \psi(\theta)$, that is

$$\mathrm{P}\left\{\psi(\theta) \le \psi_\alpha\right\} = \alpha, \quad 0 < \alpha < 1, \theta \in \Omega_\theta.$$

- In practice exact intervals are rarely available, and we seek intervals such that (1) is satisfied as closely as possible. Typically if $Y \equiv Y_1, \ldots, Y_n$, then we have

$$\mathrm{P}\left(\theta \le \theta_\alpha\right) = \alpha + \mathcal{O}(n^{-1/2}), \quad 0 < \alpha < 1, \theta \in \Omega_\theta,$$

  and the corresponding two-sided interval satisfies

$$\mathrm{P}\left(\theta_\alpha < \theta \le \theta_{1-\alpha}\right) = (1-2\alpha) + \mathcal{O}(n^{-1}), \quad 0 < \alpha < \tfrac{1}{2}, \theta \in \Omega_\theta.$$

Normal confidence intervals

- If $\widehat{\theta}$ approximately normal, then $\widehat{\theta} \dot{\sim} N(\theta + \beta, \nu)$, where $\widehat{\theta}$ has bias $\beta = \beta(F)$ and variance $\nu = \nu(F)$
- If β, ν were known, $(1-2\alpha)$ confidence interval for θ would be found as the solution to the equation

$$\mathrm{P}\left(z_\alpha < \frac{\widehat{\theta} - \theta - \beta}{\nu^{1/2}} \le z_{1-\alpha}\right) = 1 - 2\alpha,$$

  giving limits $\widehat{\theta} - \beta \pm z_\alpha \nu^{1/2}$, where $\Phi(z_\alpha) = \alpha$.
- Replace the unknowns β, ν by estimates:

$$\begin{aligned} \beta(F) &\doteq \beta(\widehat{F}) \doteq b = \overline{\widehat{\theta}^*} - \widehat{\theta}, \\ \nu(F) &\doteq \nu(\widehat{F}) \doteq v = (R-1)^{-1} \sum_r (\widehat{\theta}_r^* - \overline{\widehat{\theta}^*})^2, \end{aligned}$$

  giving $(1-2\alpha)$ interval $\widehat{\theta} - b \pm z_\alpha v^{1/2}$.
- Handedness data: $R = 10,000$, $b = -0.046$, $v = 0.205^2$, $\alpha = 0.025$, $z_\alpha = -1.96$, so 95% CI is $(0.147, 0.963)$

Normal confidence intervals

- ☐ Normal approximation reliable? Transformation needed?
- ☐ Here are plots for $\widehat{\psi}^* = \frac{1}{2}\log\{(1+\widehat{\theta}^*)/(1-\widehat{\theta}^*)\}$:

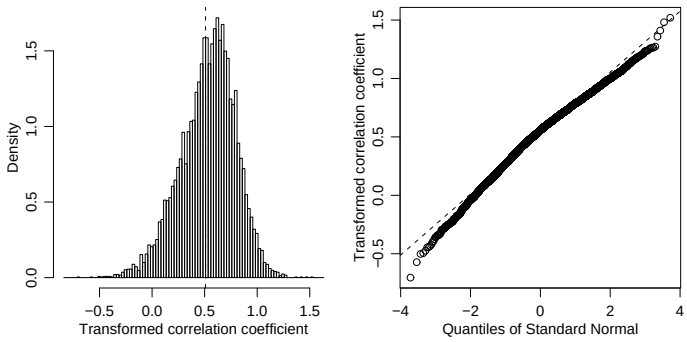

February 2021 – slide 38

Normal confidence intervals

- ☐ Correlation coefficient: try Fisher's z transformation:

$$\widehat{\psi}^* = \psi(\widehat{\theta}^*) = \tfrac{1}{2}\log\{(1+\widehat{\theta}^*)/(1-\widehat{\theta}^*)\}$$

  for which compute

$$b_\psi = R^{-1}\sum_{r=1}^{R}\widehat{\psi}_r^* - \widehat{\psi}, \quad v_\psi = \frac{1}{R-1}\sum_{r=1}^{R}\left(\widehat{\psi}_r^* - \overline{\widehat{\psi}^*}\right)^2,$$

- ☐ After inverse transformation ψ^{-1}, $(1-2\alpha)$ confidence interval for θ is

$$\psi^{-1}\left\{\widehat{\psi} - b_\psi - z_{1-\alpha}v_\psi^{1/2}\right\}, \quad \psi^{-1}\left\{\widehat{\psi} - b_\psi - z_\alpha v_\psi^{1/2}\right\}$$

- ☐ For handedness data, get $(0.074, 0.804)$
- ☐ But how do we choose a transformation in general?

February 2021 – slide 39

Pivots

- ☐ Hope properties of $\widehat{\theta}_1^*, \ldots, \widehat{\theta}_R^*$ mimic effect of sampling from original model.
- ☐ Amounts to faith in **substitution principle** (or **plugin principle**): may replace unknown F with known $\widehat{F}$ — false in general, but often more nearly true for pivots.
- ☐ **Pivot** is combination of data and parameter whose distribution is independent of underlying model.
- ☐ Canonical example: $Y_1, \ldots, Y_n \overset{\text{iid}}{\sim} N(\mu, \sigma^2)$. Then

$$Z = \frac{\overline{Y} - \mu}{(S^2/n)^{1/2}} \quad \sim \quad t_{n-1},$$

  for all μ, σ^2 — so independent of the underlying (normal!) distribution
- ☐ Exact pivot generally unavailable in nonparametric case.

February 2021 – slide 40

Studentized statistic

- ☐ Idea: generalize Student t statistic to bootstrap setting
- ☐ Requires variance V for $\widehat{\theta}$ computed from $y_1, \ldots, y_n$
- ☐ Analogue of Student t statistic:

$$Z = \frac{\widehat{\theta} - \theta}{V^{1/2}}$$

- ☐ If the quantiles z_α of Z known, then

$$\mathrm{P}\left(z_\alpha \leq Z \leq z_{1-\alpha}\right) = \mathrm{P}\left(z_\alpha \leq \frac{\widehat{\theta} - \theta}{V^{1/2}} \leq z_{1-\alpha}\right) = 1 - 2\alpha$$

  (z_α no longer denotes a normal quantile!) implies that

$$\mathrm{P}\left(\widehat{\theta} - V^{1/2} z_{1-\alpha} \leq \theta \leq \widehat{\theta} - V^{1/2} z_\alpha\right) = 1 - 2\alpha$$

  so $(1 - 2\alpha)$ confidence interval is $(\widehat{\theta} - V^{1/2} z_{1-\alpha}, \widehat{\theta} - V^{1/2} z_\alpha)$

February 2021 – slide 41

Studentized statistic

- Bootstrap sample gives $(\widehat{\theta}^*, V^*)$ and hence

$$Z^* = \frac{\widehat{\theta}^* - \widehat{\theta}}{V^{*1/2}}$$

- R bootstrap copies of $(\widehat{\theta}, V)$:

$$(\widehat{\theta}_1^*, V_1^*), \quad (\widehat{\theta}_2^*, V_2^*), \quad \ldots, \quad (\widehat{\theta}_R^*, V_R^*)$$

  and corresponding

$$z_1^* = \frac{\widehat{\theta}_1^* - \widehat{\theta}}{V_1^{*1/2}}, \quad z_2^* = \frac{\widehat{\theta}_2^* - \widehat{\theta}}{V_2^{*1/2}}, \quad \ldots, \quad z_R^* = \frac{\widehat{\theta}_R^* - \widehat{\theta}}{V_R^{*1/2}}.$$

- Use $z_1^*, \ldots, z_R^*$ to estimate distribution of Z — for example, order statistics $z_{(1)}^* < \cdots < z_{(R)}^*$ used to estimate quantiles
- Get $(1 - 2\alpha)$ confidence interval

$$\widehat{\theta} - V^{1/2} z_{((1-\alpha)(R+1))}^*, \quad \widehat{\theta} - V^{1/2} z_{(\alpha(R+1))}^*$$

Mathematical Aside: Why Studentize?

- Studentize, so $Z \xrightarrow{D} N(0,1)$ as $n \to \infty$. Edgeworth series:

$$\mathrm{P}(Z \le z \mid F) = \Phi(z) + n^{-1/2} a(z)\phi(z) + O(n^{-1});$$

  $a(\cdot)$ even quadratic polynomial.
- Example: if we use $\widehat{\theta} = \overline{Y}$ and $V = n^{-1}S^2$ to compute Z for the average of a sample with skewness γ, then $a(x) = \gamma(2x^2 + 1)/6$ and (next slide) $a'(x) = -\gamma(x^2 - 1)/6$.
- Corresponding expansion for Z^* is

$$\mathrm{P}(Z^* \le z \mid \widehat{F}) = \Phi(z) + n^{-1/2} \widehat{a}(z)\phi(z) + O_p(n^{-1}).$$

  Typically $\widehat{a}(z) = a(z) + O_p(n^{-1/2})$, so

$$\mathrm{P}(Z^* \le z \mid \widehat{F}) - \mathrm{P}(Z \le z \mid F) = O_p(n^{-1}).$$

Aside continued ...

- If don't studentize, $Z = n^{1/2}(\widehat{\theta} - \theta) \xrightarrow{D} N(0, \nu')$. Then

$$\mathrm{P}(Z \le z \mid F) = \Phi\left(\frac{z}{\nu'^{1/2}}\right) + n^{-1/2} a' \left(\frac{z}{\nu'^{1/2}}\right) \phi\left(\frac{z}{\nu'^{1/2}}\right) + O(n^{-1})$$

  and

$$\mathrm{P}(Z^* \le z \mid \widehat{F}) = \Phi\left(\frac{z}{\widehat{\nu}'^{1/2}}\right) + n^{-1/2} \widehat{a}' \left(\frac{z}{\widehat{\nu}'^{1/2}}\right) \phi\left(\frac{z}{\widehat{\nu}'^{1/2}}\right) + O_p(n^{-1}).$$

  Typically $\widehat{\nu}' = \nu' + O_p(n^{-1/2})$, giving

$$\mathrm{P}(Z^* \le z \mid \widehat{F}) - \mathrm{P}(Z \le z \mid F) = O_p(n^{-1/2}).$$

- Thus use of Studentized Z reduces error from $O_p(n^{-1/2})$ to $O_p(n^{-1})$: better than using large-sample asymptotics, for which error is usually $O_p(n^{-1/2})$.

Other confidence intervals

- Problem for studentized intervals: must obtain V, intervals not transformation-invariant
- Simpler approaches:
  - Basic bootstrap interval: treat $\widehat{\theta} - \theta$ as pivot, get

$$\widehat{\theta} - (\widehat{\theta}^*_{((R+1)(1-\alpha))} - \widehat{\theta}), \quad \widehat{\theta} - (\widehat{\theta}^*_{((R+1)\alpha)} - \widehat{\theta}).$$

  - Percentile interval: use empirical quantiles of $\widehat{\theta}^*_1, \ldots, \widehat{\theta}^*_R$:

$$\widehat{\theta}^*_{((R+1)\alpha)}, \quad \widehat{\theta}^*_{((R+1)(1-\alpha))}.$$

- More complex: improved percentile intervals (BC_a, ABC, ...)
  - Replace percentile interval with

$$\widehat{\theta}^*_{((R+1)\alpha')}, \quad \widehat{\theta}^*_{((R+1)(1-\alpha''))},$$

    where α', α'' chosen to improve properties.
  - Transformation-invariant.

BC_a intervals

- ☐ **Bias-corrected and accelerated** confidence intervals
- ☐ Determine α' by

$$\alpha' = \Phi\left\{w + \frac{w + z_\alpha}{1 - a(w + z_\alpha)}\right\},$$

where

$$w = \Phi^{-1}\left\{\widehat{G}^*(\widehat{\theta})\right\}, \quad a = \frac{1}{6}\frac{\sum_{j=1}^n l_j^3}{\left(\sum_{j=1}^n l_j^2\right)^{3/2}},$$

with $\widehat{G}^*$ the EDF of the $\widehat{\theta}_1^*, \ldots, \widehat{\theta}_R^*$, and $l_1, \ldots, l_n$ the empirical influence values, defined below

- ☐ If the **bias** $w = 0$, then $\widehat{G}^*(\widehat{\theta}) = \frac{1}{2}$, so $\widehat{\theta}$ is at the median of the EDF of $\widehat{\theta}^*$
- ☐ If the **acceleration** $a = 0$, then the effect of the data $y_1, \ldots, y_n$ on $\widehat{\theta}$ is symmetric (see below)

Handedness data

- ☐ 95% confidence intervals for correlation coefficient θ, $R = 10,000$:

Normal	0.147	0.963
Percentile	-0.047	0.758
Basic	0.262	1.043
BC_a $(\alpha' = 0.0485, \alpha'' = 0.0085)$	0.053	0.792
Student	0.030	1.206
Basic (transformed)	0.131	0.824
Student (transformed)	-0.277	0.868

- ☐ Transformation is essential here!

Comparisons

Table 2: Empirical error rates (%) for nonparametric bootstrap confidence limits in ratio estimation: rates for sample sizes $n_1 = n_2 = 10$ are given above those for sample sizes $n_1 = n_2 = 25$. $R = 999$ for all bootstrap methods. 10,000 data sets generated from Gamma distributions.

Method	Nominal error rate							
	Lower limit				Upper limit			
	1	2.5	5	10	10	5	2.5	1
Exact	1.0	2.8	5.5	10.5	9.8	4.8	2.6	1.0
	1.0	2.3	4.8	9.9	10.2	4.9	2.5	1.1
Normal approximation	0.1	0.5	1.7	6.3	20.6	15.7	12.5	9.6
	0.1	0.5	2.1	6.4	16.3	11.5	8.2	5.5
Basic bootstrap	0.0	0.0	0.2	1.8	24.4	21.0	18.6	16.4
	0.0	0.1	0.4	3.0	19.2	15.0	12.5	10.3
Basic bootstrap, log scale	2.6	4.9	8.1	12.9	13.1	7.5	4.8	2.5
	1.6	3.2	6.0	11.4	11.5	6.3	3.3	1.7
Studentized bootstrap	0.6	2.1	4.6	9.9	11.9	6.7	4.0	2.0
	0.8	2.3	4.6	9.9	10.9	5.9	3.0	1.4
Studentized bootstrap, log scale	1.1	2.8	5.6	10.7	11.6	6.3	3.5	1.7
	1.1	2.5	5.0	10.1	10.8	5.7	2.9	1.3
Bootstrap percentile	1.8	3.6	6.5	11.6	14.6	8.9	5.9	3.3
	1.2	2.6	5.1	10.1	12.6	7.1	4.2	2.1
BC_a	1.9	4.0	6.9	12.3	14.0	8.3	5.3	3.0
	1.4	3.0	5.6	10.9	11.8	6.8	3.8	1.9
ABC	1.9	4.2	7.4	12.7	14.6	8.7	5.5	3.1
	1.3	3.0	5.7	11.0	12.1	6.8	3.7	1.9

Confidence interval lengths

Lengths of 95% confidence intervals for the first 1000 simulated samples in the numerical experiment with Gamma data.

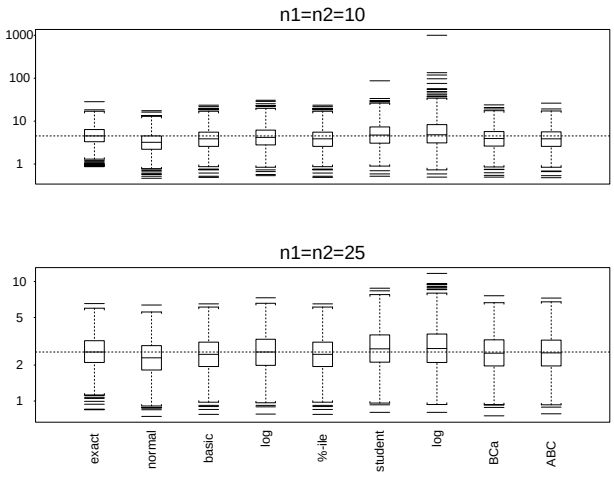

General comparison

- ☐ Bootstrap confidence intervals usually under-cover
- ☐ Normal, basic, and studentized intervals depend on scale.
- ☐ Percentile interval often too short but is transformation-invariant.
- ☐ Studentized intervals give best coverage overall, but
  - depend on scale, can be sensitive to V
  - length can be very variable
  - best on transformed scale, where V is approximately constant
- ☐ Improved percentile intervals have same error in principle as studentized intervals, but often shorter — so lower coverage

February 2021 – slide 50

Caution

- ☐ Edgeworth theory OK for smooth statistics—beware rough statistics: must check output.
- ☐ Bootstrap of median theoretically OK, but very sensitive to sample values in practice.
- ☐ Figure below shows results for bootstrapping a difference of medians—use smoothed bootstap, or sample m-out-of-n observations?

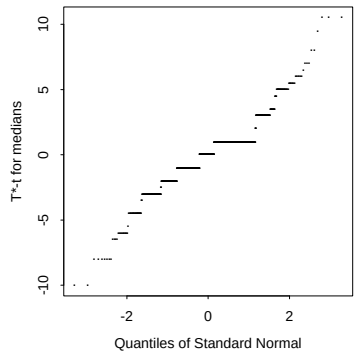

February 2021 – slide 51

Key points

- ☐ Numerous procedures available for 'automatic' construction of confidence intervals
- ☐ Computer does the work
- ☐ Need $R \geq 1000$ in most cases
- ☐ Generally such intervals are a bit too short
- ☐ Must examine output to check if assumptions (e.g. smoothness of statistic) satisfied
- ☐ May need variance estimate V — see later

February 2021 – slide 52

Gravity data

Table 3: Measurements of the acceleration due to gravity, g, given as deviations from 980,000 $\times 10^{-3}$ cms^{-2}, in units of cms$^{-2} \times 10^{-3}$.

Series							
1	2	3	4	5	6	7	8
76	87	105	95	76	78	82	84
82	95	83	90	76	78	79	86
83	98	76	76	78	78	81	85
54	100	75	76	79	86	79	82
35	109	51	87	72	87	77	77
46	109	76	79	68	81	79	76
87	100	93	77	75	73	79	77
68	81	75	71	78	67	78	80
	75	62			75	79	83
	68				82	82	81
	67				83	76	78
						73	78
						64	78

Gravity data

Figure 1: Gravity series boxplots, showing a reduction in variance, a shift in location, and possible outliers.

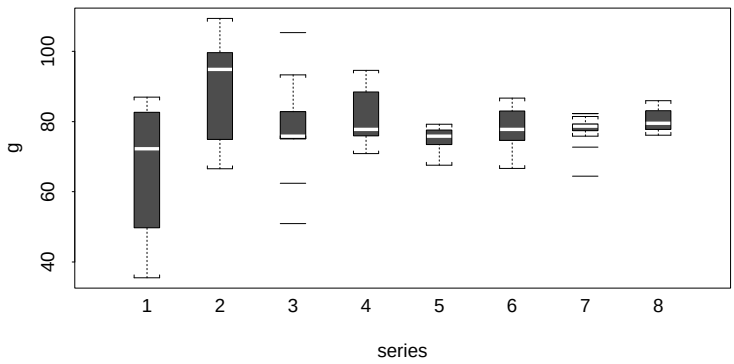

Gravity data

- ☐ Eight series of measurements of gravitational acceleration g made May 1934 – July 1935 in Washington DC
- ☐ Data are deviations from 9.8 m/s^2 in units of 10^{-3} cm/s^2
- ☐ Goal: Estimate g and provide confidence interval
- ☐ Weighted combination of series averages and its variance estimate

$$\widehat{\theta} = \frac{\sum_{i=1}^{8} \overline{y}_i \times n_i/s_i^2}{\sum_{i=1}^{8} n_i/s_i^2}, \quad V = \left(\sum_{i=1}^{8} n_i/s_i^2\right)^{-1},$$

giving

$$\widehat{\theta} = 78.54, \quad V = 0.59^2$$

and 95% confidence interval of $\widehat{\theta} \pm 1.96V^{1/2} = (77.5, 79.8)$

February 2021 – slide 56

Gravity data: Bootstrap

- ☐ Apply stratified (re)sampling to series, taking each series as a separate stratum. Compute $\widehat{\theta}^*$, V^* for simulated data
- ☐ Confidence interval based on

$$Z^* = \frac{\widehat{\theta}^* - \widehat{\theta}}{V^{*1/2}},$$

whose distribution is approximated by simulations

$$z_1^* = \frac{\widehat{\theta}_1^* - \widehat{\theta}}{V_1^{1/2}}, \ldots, z_R^* = \frac{\widehat{\theta}_R^* - \widehat{\theta}}{V_R^{1/2}},$$

giving

$$(\widehat{\theta} - V^{1/2} z^*_{(R(1-\alpha))}, \widehat{\theta} - V^{1/2} z^*_{(R\alpha)})$$

- ☐ For 95% limits set $\alpha = 0.025$, so with $R = 1000$ use $z^*_{(25)}, z^*_{(975)}$, giving interval $(77.1, 80.3)$.

February 2021 – slide 57

Results

Figure 2: Summary plots for 1000 nonparametric bootstrap simulations. Top: normal probability plots of $\widehat{\theta}^*$ and $z^* = (\widehat{\theta}^* - \widehat{\theta})/v^{*1/2}$. Line on the top left has intercept t and slope $v^{1/2}$, line on the top right has intercept zero and unit slope. Bottom: the smallest $\widehat{\theta}_r^*$ also has the smallest v^*, leading to an outlying value of z^*.

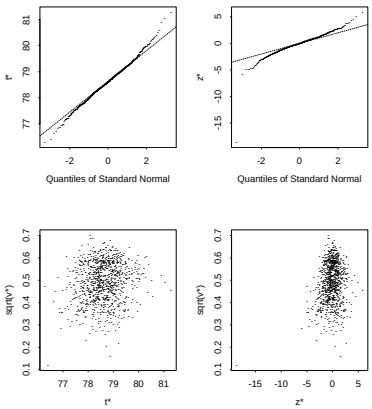

February 2021 – slide 58

Key points

- ☐ For several independent samples, implement bootstrap by stratified sampling independently from each
- ☐ Same basic ideas apply for confidence intervals

February 2021 – slide 59

Variance Estimation

slide 60

Variance estimation

- ☐ Variance estimate V needed for certain types of confidence interval (esp. studentized bootstrap)
- ☐ Ways to compute this:
  - double bootstrap
  - delta method
  - nonparametric delta method
  - jackknife

February 2021 – slide 61

Double bootstrap

- ☐ Bootstrap sample $y_1^*, \ldots, y_n^*$ and corresponding estimate $\widehat{\theta}^*$
- ☐ Take Q second-level bootstrap samples $y_1^{**}, \ldots, y_n^{**}$ from $y_1^*, \ldots, y_n^*$, giving corresponding bootstrap estimates $\widehat{\theta}_1^{**}, \ldots, \widehat{\theta}_Q^{**}$,
- ☐ Compute variance estimate V^* as sample variance of $\widehat{\theta}_1^{**}, \ldots, \widehat{\theta}_Q^{**}$
- ☐ Requires total $R(Q+1)$ resamples, so could be expensive
- ☐ Often reasonable to take $Q \doteq 50$ for variance estimation, so need $O(50 \times 1000)$ resamples — nowadays not infeasible
- ☐ Aside: can use double bootstrap to diagnose problems with statistic, and to suggest changes. Example: next slide shows smoothed estimation of standard error curve

$$\widehat{\sigma}(\theta) = \sqrt{\mathrm{var}(\widehat{\theta})},$$

which can be used to estimate variance-stabilising transformation

$$h(\theta) = \int^{\theta} \frac{du}{\widehat{\sigma}(u)}.$$

Aside: Variance function estimation

Figure 3: Bias and variance estimates for ratio applied to city population data, $n = 10$. For each of $R = 999$ bootstrap samples from the data, $M = 50$ second-level samples were drawn, and the resulting bias and standard error estimates b^* and $v^{*1/2}$ plotted against the bootstrapped ratio $\widehat{\theta}^*$. The lines are from a robust nonparametric curve fit to the simulations.

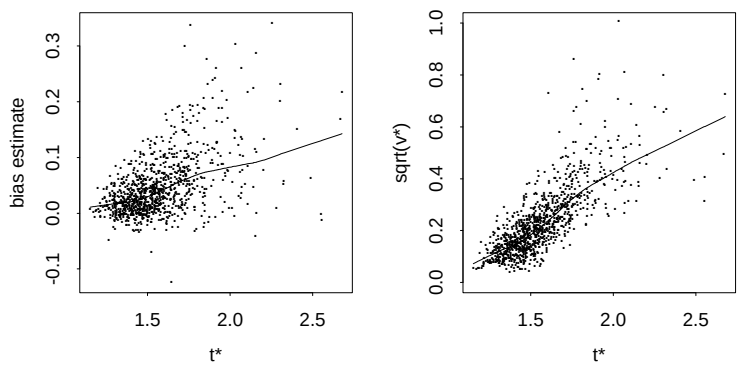

Delta method

- ☐ Computation of variance formulae for functions of averages and other estimators
- ☐ Suppose $\widehat{\psi} = g(\widehat{\theta})$ estimates $\psi = g(\theta)$, and $\widehat{\theta} \dot{\sim} N(\theta, \sigma^2/n)$
- ☐ Then under mild conditions and provided $g'(\theta) \neq 0$, Taylor expansion gives

$$\begin{aligned} \mathrm{E}(\widehat{\psi}) &= g(\theta) + O(n^{-1}), \\ \mathrm{var}(\widehat{\psi}) &= \sigma^2 g'(\theta)^2/n + O(n^{-3/2}). \end{aligned}$$

- ☐ Then $\mathrm{var}(\widehat{\psi}) \doteq \widehat{\sigma}^2 g'(\widehat{\theta})^2/n = V$
- ☐ Example: $\widehat{\theta} = \overline{Y}$, $\widehat{\psi} = \log \widehat{\theta}$
- ☐ Variance stabilisation: if $\mathrm{var}(\widehat{\theta}) \doteq S(\theta)^2/n$, find transformation h such that $\mathrm{var}\{h(\widehat{\theta})\} \doteq$constant
- ☐ Extends to multivariate estimators, and to $\widehat{\psi} = g(\widehat{\theta}_1, \ldots, \widehat{\theta}_d)$

February 2021 – slide 64

Nonparametric delta method

- ☐ Write parameter $\theta = \theta(F)$ as function of distribution F
- ☐ General approximation:

$$V \doteq V_L = \frac{1}{n^2} \sum_{j=1}^{n} L(Y_j; F)^2.$$

- ☐ $L(y; F)$ is **influence function value** for θ for observation at y when distribution is F:

$$L(y; F) = \lim_{\varepsilon \to 0} \frac{\theta\{(1-\varepsilon)F + \varepsilon H_y\} - \theta(F)}{\varepsilon},$$

  where H_y puts unit mass at y. Close link to robustness.
- ☐ Empirical versions of $L(y; F)$ and V_L are

$$l_j = L(y_j; \widehat{F}), \quad v_L = n^{-2} \sum l_j^2,$$

  usually obtained by analytic/numerical differentiation.

February 2021 – slide 65

Nonparametric delta method: Average

- Population mean and its empirical version

$$\theta = t(F) = \int x\, dF(x), \quad \widehat{\theta} = t(\widehat{F}) = \int x\, d\widehat{F}(x) = n^{-1}\sum_{j=1}^{n} Y_j = \overline{Y}.$$

- If H_y puts unit mass at y, its 'density' is a Dirac delta function $\delta_y(x)$, and

$$\begin{aligned}\theta\left\{(1-\varepsilon)F + \varepsilon H_y\right\} &= \int x\, d\left\{(1-\varepsilon)F + \varepsilon H_y\right\}(x) \\ &= (1-\varepsilon)\int x\, dF(x) + \varepsilon \int x\, dH_y(x) = (1-\varepsilon)\theta(F) + \varepsilon y\end{aligned}$$

  and therefore

$$L(y;F) = \lim_{\varepsilon\to 0}\frac{\theta\left\{(1-\varepsilon)F + \varepsilon H_y\right\} - \theta(F)}{\varepsilon} = \lim_{\varepsilon\to 0}\frac{(1-\varepsilon)\theta(F) + \varepsilon y - \theta(F)}{\varepsilon} = y - \theta(F),$$

- Hence empirical influence values and variance estimate are

$$l_j = L(y_j;\widehat{F}) = y_j - \overline{y}, \qquad v_L = \frac{1}{n^2}\sum(y_j - \overline{y})^2 = \frac{n-1}{n}n^{-1}s^2.$$

February 2021 – slide 66

Computation of l_j

- Write $\widehat{\theta}$ in weighted form, differentiate with respect to ε
- Sample average:

$$\widehat{\theta} = \overline{y} = \frac{1}{n}\sum y_j = \sum w_j y_j\Big|_{w_j \equiv 1/n}$$

  Change weights:

$$w_j \mapsto \varepsilon + (1-\varepsilon)\tfrac{1}{n}, \quad w_i \mapsto (1-\varepsilon)\tfrac{1}{n}, \quad i \neq j$$

  so

$$\overline{y} \mapsto \overline{y}_\varepsilon = \varepsilon y_j + (1-\varepsilon)\overline{y} = \varepsilon(y_j - \overline{y}) + \overline{y},$$

  giving $l_j = y_j - \overline{y}$ and $v_L = \frac{1}{n^2}\sum(y_j - \overline{y})^2 = \frac{n-1}{n}n^{-1}s^2$
- Interpretation: l_j is standardized change in $\overline{y}$ when increase mass on y_j

February 2021 – slide 67

Nonparametric delta method: Ratio

- Population $F(u,x)$ with $y=(u,x)$ and

$$\theta = \theta(F) = \int x\,dF(u,x) / \int u\,dF(u,x),$$

sample version is

$$\widehat{\theta} = \theta(\widehat{F}) = \int x\,d\widehat{F}(u,x) / \int u\,d\widehat{F}(u,x) = \overline{x}/\overline{u}$$

- Then using result for averages and chain rule,

$$l_j = (x_j - \widehat{\theta} u_j)/\overline{u},$$

giving

$$v_L = \frac{1}{n^2} \sum \left(\frac{x_j - \widehat{\theta} u_j}{\overline{u}} \right)^2 .$$

February 2021 – slide 68

Handedness data: Correlation coefficient

- Correlation coefficient may be written as a function of averages $\overline{xu} = n^{-1} \sum x_j u_j$ etc.:

$$\widehat{\theta} = \frac{\overline{xu} - \overline{x}\,\overline{u}}{\left\{ (\overline{x^2} - \overline{x}^2)(\overline{u^2} - \overline{u}^2) \right\}^{1/2}},$$

from which empirical influence values l_j can be derived

- In this example (and for others involving only averages), nonparametric delta method is equivalent to delta method
- Get

$$v_L = 0.029$$

for comparison with $v = 0.043$ obtained by bootstrapping.

- v_L typically underestimates $\mathrm{var}(\widehat{\theta})$ — as here!

February 2021 – slide 69

Delta methods: Comments

- ☐ Can be applied to many complex statistics
- ☐ For small n, delta method often underestimates true variances:

$$v_L < \mathrm{var}(\widehat{\theta})$$

- ☐ Can be applied automatically (numerical differentation) if algorithm for $\widehat{\theta}$ written in weighted form, e.g.

$$\overline{x}_w = \sum w_j x_j, \quad w_j \equiv 1/n \text{ for } \overline{x}$$

and vary weights successively for $j = 1, \ldots, n$, setting

$$w_j = w_i + \varepsilon, \quad i \neq j, \quad \sum w_i = 1$$

for $\varepsilon = 1/(100n)$ and using the definition as derivative

February 2021 – slide 70

Jackknife

- ☐ Approximation to empirical influence values given by

$$l_j \approx l_{\text{jack},j} = (n-1)(\widehat{\theta} - \widehat{\theta}_{-j}),$$

where $\widehat{\theta}_{-j}$ is value of $\widehat{\theta}$ computed from sample

$$y_1, \ldots, y_{j-1}, \quad y_{j+1}, \ldots, y_n$$

- ☐ Jackknife bias and variance estimates are

$$b_{\text{jack}} = -\frac{1}{n} \sum l_{\text{jack},j}, \quad v_{\text{jack}} = \frac{1}{n(n-1)} \left(\sum l_{\text{jack},j}^2 - n b_{\text{jack}}^2 \right)$$

- ☐ Requires $n+1$ calculations of $\widehat{\theta}$
- ☐ Corresponds to numerical differentiation of $\widehat{\theta}$, with $\varepsilon = -1/(n-1)$

February 2021 – slide 71

Key points

- ☐ Several methods available for estimation of variances
- ☐ Needed for some types of confidence interval
- ☐ Most general method is double bootstrap: can be expensive
- ☐ Delta methods rely on linear expansion, can be applied numerically or analytically
- ☐ Jackknife gives approximation to delta method, can fail for rough statistics

February 2021 – slide 72

Ingredients

- ☐ Ingredients for testing problems:
  - data $y_1, \ldots, y_n$;
  - hypothesis H_0 (equivalently model M_0) to be tested;
  - test statistic $t = t(y_1, \ldots, y_n)$, with large values giving evidence against H_0, and observed value t_{obs}
- ☐ P-value

$$p_{\text{obs}} = \mathrm{P}_0(T \geq t_{\text{obs}}) = \mathrm{P}(T \geq t_{\text{obs}} \mid M_0)$$

  measures evidence against M_0 — small p_{obs} indicates evidence against M_0.
- ☐ Difficulties:
  - p_{obs} may depend upon 'nuisance' parameters, those of M_0;
  - p_{obs} often hard to calculate.

February 2021 – slide 74

Examples

- ☐ Balsam-fir seedlings in 5×5 quadrats — Poisson sample?

0	1	2	3	4	3	4	2	2	1
0	2	0	2	4	2	3	3	4	2
1	1	1	1	4	1	5	2	2	3
4	1	2	5	2	0	3	2	1	1
3	1	4	3	1	0	0	2	7	0

- ☐ Two-way layout: row-column independence?

1	2	2	1	1	0	1
2	0	0	2	3	0	0
0	1	1	1	2	7	3
1	1	2	0	0	0	1
0	1	1	1	1	0	0

February 2021 – slide 75

Estimation of p_{obs}

- ☐ Estimate p_{obs} by simulation from fitted null hypothesis model $\widehat{M_0}$.
- ☐ Algorithm: for $r = 1, \ldots, R$:
  - simulate data set $y_1^*, \ldots, y_n^*$ from $\widehat{M_0}$;
  - calculate test statistic t_r^* from $y_1^*, \ldots, y_n^*$.
- ☐ Calculate simulation estimate

$$\widehat{p} = \frac{\#\{t_r^* \geq t_{\text{obs}}\}}{R}$$

of

$$\widehat{p}_{\text{obs}} = \mathrm{P}(T \geq t_{\text{obs}} \mid \widehat{M_0}).$$

- ☐ Simulation and statistical errors:

$$\widehat{p} \approx \widehat{p}_{\text{obs}} \approx p_{\text{obs}}$$

Handedness data: Test of independence

- ☐ Are `dnan` and `hand` positively associated?
- ☐ Take $T = \widehat{\theta}$ (correlation coefficient), with $\widehat{\theta}_{\text{obs}} = 0.509$; this is large in case of positive association (one-sided test)
- ☐ Null hypothesis of independence: $F(u, x) = F_1(u)F_2(x)$
- ☐ Take bootstrap samples independently from $\widehat{F_1} \equiv (\texttt{dnan}_1, \ldots, \texttt{dnan}_n)$ and from $\widehat{F_2} \equiv (\texttt{hand}_1, \ldots, \texttt{hand}_n)$, then put them together to get bootstrap data $(\texttt{dnan}_1^*, \texttt{hand}_1^*), \ldots, (\texttt{dnan}_n^*, \texttt{hand}_n^*)$.
- ☐ With $R = 10,000$ get 19 values of $\widehat{\theta}^* \geq \widehat{\theta}_{\text{obs}}$, so

$$\widehat{p} = \frac{19}{10000} = 0.0019 :$$

`hand` and `dnan` seem to be positively associated

- ☐ To test positive or negative association (two-sided test), take $T = |\widehat{\theta}|$: gives $\widehat{p} = 0.004$.

Handedness data: Bootstrap from $\widehat{M_0}$

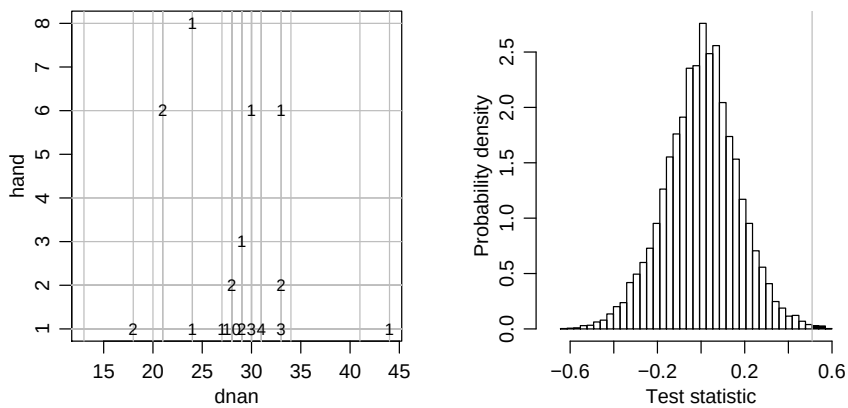

Choice of R

- ☐ Take R big enough to get small standard error for $\widehat{p}$, typically ≥ 100, using binomial calculation:

$$\begin{aligned} \mathrm{var}(\widehat{p}) &= \mathrm{var}\left(\frac{\#\{t_r^* \geq t_{\mathrm{obs}}\}}{R}\right) \\ &= \frac{1}{R^2}\, R p_{\mathrm{obs}}(1-p_{\mathrm{obs}}) = \frac{p_{\mathrm{obs}}(1-p_{\mathrm{obs}})}{R} \end{aligned}$$

  so if $p_{\mathrm{obs}} \doteq 0.05$ need $R \geq 1900$ for 10% relative error
- ☐ Can choose R sequentially: e.g. if $\widehat{p} \doteq 0.06$ and $R = 99$, can augment R enough to diminish standard error.
- ☐ Taking R too small lowers power of test, because the randomness of the t_r^* means that the rejection region for the test is random.

Two situations

- ☐ Tests about parameters:
  - Is correlation coefficient equal to zero for handedness data?
- ☐ Pure significance tests:
  - Are the balsam-fir seedling data Poisson?
  - Are the row and column classifications independent for the two-way layout data?
- ☐ In the first case, can exploit duality of test and confidence interval to obtain significance levels, or can use pivot test.

Duality with confidence interval

- ☐ For test concerning the value of a parameter θ
- ☐ General approach based on duality between confidence interval $\mathcal{I}_{1-\alpha} = (\theta_\alpha, \infty)$ and test of null hypothesis $\theta = \theta_0$
- ☐ Reject null hypothesis at level α in favour of alternative $\theta > \theta_0$, if $\theta_0 < \theta_\alpha$
- ☐ Handedness data: $\theta_0 = 0 \notin \mathcal{I}_{0.95}$, but $\theta_0 = 0 \in \mathcal{I}_{0.99}$, so estimated significance level $0.01 < \widehat{p} < 0.05$: weaker evidence than before
- ☐ Extends to tests of $\theta = \theta_0$ against other alternatives:
  - if $\theta_0 \notin \mathcal{I}^{1-\alpha} = (-\infty, \theta^\alpha)$, have evidence that $\theta < \theta_0$
  - if $\theta_0 \notin \mathcal{I}_{1-2\alpha} = (\theta_\alpha, \theta^\alpha)$, have evidence that $\theta \neq \theta_0$

Pivot tests

- ☐ Equivalent to use of confidence intervals
- ☐ Idea: use (approximate) pivot such as $Z = (\widehat{\theta} - \theta)/V^{1/2}$ as statistic to test $\theta = \theta_0$
- ☐ Observed value of pivot is $z_{\text{obs}} = (\widehat{\theta} - \theta_0)/V^{1/2}$
- ☐ Significance level is

$$\begin{aligned} \mathrm{P}\left(\frac{\widehat{\theta} - \theta}{V^{1/2}} \geq z_{\text{obs}} \mid M_0\right) &= \mathrm{P}(Z \geq z_{\text{obs}} \mid M_0) \\ &= \mathrm{P}(Z \geq z_{\text{obs}} \mid F) \\ &\doteq \mathrm{P}(Z \geq z_{\text{obs}} \mid \widehat{F}) \end{aligned}$$

- ☐ Compare observed z_{obs} with simulated distribution of $Z^* = (\widehat{\theta}^* - \widehat{\theta})/V^{*1/2}$, without needing to construct null hypothesis model $\widehat{M_0}$
- ☐ Use of (approximate) pivot is essential for success

Example: Handedness data

- ☐ Test zero correlation ($\theta_0 = 0$), not independence; $\widehat{\theta} = 0.509$, $V = 0.170^2$:

$$z_{\text{obs}} = \frac{\widehat{\theta} - \theta_0}{V^{1/2}} = \frac{0.509 - 0}{0.170} = 2.99$$

- ☐ Observed significance level is

$$\widehat{p} = \frac{\#\{z_r^* \geq z_{\text{obs}}\}}{R} = \frac{216}{10000} = 0.0216$$

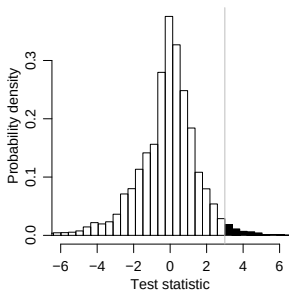

Exact tests

- ☐ Problem: bootstrap estimate is

$$\widehat{p}_{\text{obs}} = \mathrm{P}(T \geq t_{\text{obs}} \mid \widehat{M_0}) \neq \mathrm{P}(T \geq t \mid M_0) = p_{\text{obs}},$$

  so estimate the wrong thing
- ☐ In some cases can eliminate parameters from null hypothesis distribution by conditioning on sufficient statistic
- ☐ Then simulate from conditional distribution
- ☐ More generally, can use Metropolis–Hastings algorithm to simulate from conditional distribution (below)

Example: Fir data

- ☐ Data $Y_1, \ldots, Y_n \stackrel{\text{iid}}{\sim} \text{Pois}(\lambda)$, with λ unknown
- ☐ Poisson model has $\text{E}(Y) = \text{var}(Y) = \lambda$: base test of overdispersion on

$$T = \sum (Y_j - \overline{Y})^2 / \overline{Y} \dot{\sim} \chi^2_{n-1};$$

  observed value is $t_{\text{obs}} = 55.15$
- ☐ Unconditional significance level:

$$\text{P}(T \geq t_{\text{obs}} \mid \widehat{M}_0, \lambda)$$

- ☐ Condition on value w of sufficient statistic $W = \sum Y_j$:

$$p_{\text{obs}} = \text{P}(T \geq t_{\text{obs}} \mid \widehat{M}_0, W = w),$$

  independent of λ, owing to sufficiency of W
- ☐ Exact test: simulate from multinomial distribution of $Y_1, \ldots, Y_n$ given $W = \sum Y_j = 107$.

Example: Fir data

Simulation results for dispersion test. Left panel: $R = 1000$ values of the dispersion statistic t^* obtained under multinomial sampling: the data value is $t_{\text{obs}} = 55.15$ and $\widehat{p} = 0.25$. Right panel: chi-squared plot of ordered values of t^*, dotted line shows χ^2_{49} approximation to null conditional distribution.

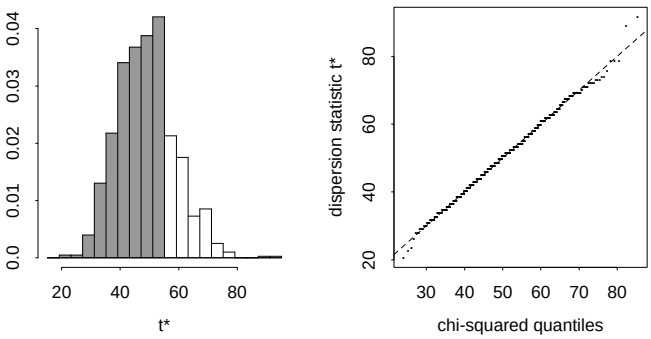

Handedness data: Permutation test

- ☐ Are `dnan` and `hand` related?
- ☐ Take $T = \widehat{\theta}$ (correlation coefficient) again
- ☐ Impose null hypothesis of independence: $F(u, x) = F_1(u)F_2(x)$, but condition so that marginal distributions $\widehat{F}_1$ and $\widehat{F}_2$ are held fixed under resampling plan — permutation test
- ☐ Take resamples of form
$$(\texttt{dnan}_1, \texttt{hand}_{1^*}), \ldots, (\texttt{dnan}_n, \texttt{hand}_{n^*})$$
where $(1^*, \ldots, n^*)$ is random permutation of $(1, \ldots, n)$
- ☐ Doing this with $R = 10,000$ gives one- and two-sided significance probabilities of 0.002, 0.003
- ☐ Typically values of $\widehat{p}$ very similar to those for corresponding bootstrap test

Handedness data: Permutation resample

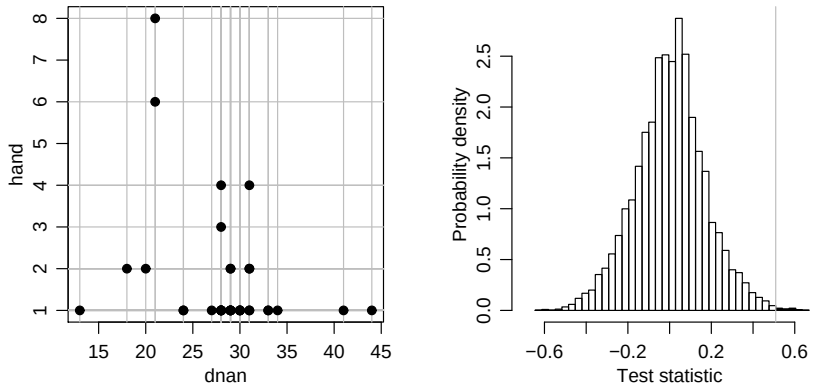

Contingency table

1	2	2	1	1	0	1
2	0	0	2	3	0	0
0	1	1	1	2	7	3
1	1	2	0	0	0	1
0	1	1	1	1	0	0

- ☐ Are row and column classifications independent:
$$\mathrm{P}(\text{row } i,\ \text{column } j) = \mathrm{P}(\text{row } i) \times \mathrm{P}(\text{column } j)?$$
- ☐ Standard test statistic for independence is
$$T = \sum_{i,j} \frac{(y_{ij} - \widehat{y}_{ij})^2}{\widehat{y}_{ij}}, \quad \widehat{y}_{ij} = \frac{y_{i\cdot} y_{\cdot j}}{y_{\cdot\cdot}}$$
- ☐ Get $\mathrm{P}(\chi^2_{24} \geq 38.52) = 0.048$, but is $T \dot{\sim} \chi^2_{24}$?

Exact tests: Contingency table

- ☐ For exact test, need to simulate distribution of T conditional on sufficient statistics — row and column totals
- ☐ Metropolis–Hastings algorithm for conditional simulation:
  1. choose two rows $j_1 < j_2$ and two columns $k_1 < k_2$ at random
  2. generate new values from hypergeometric distribution of $y_{j_1 k_1}$ conditional on margins of 2×2 table
  $$\begin{matrix} y_{j_1 k_1} & y_{j_1 k_2} \\ y_{j_2 k_1} & y_{j_2 k_2} \end{matrix}$$
  3. compute test statistic T^* every $I = 100$ iterations, say
- ☐ Compare observed value $t_{\text{obs}} = 38.52$ with simulated T^* — get $\widehat{p} \doteq 0.08$

Exact tests: Contingency table

Comparison of 10,000 simulated values of T^* with the χ^2_{24} distribution. Monte Carlo 'exact' value is $\widehat{p} \doteq 0.08$, while χ^2_{24} approximation gives 0.048.

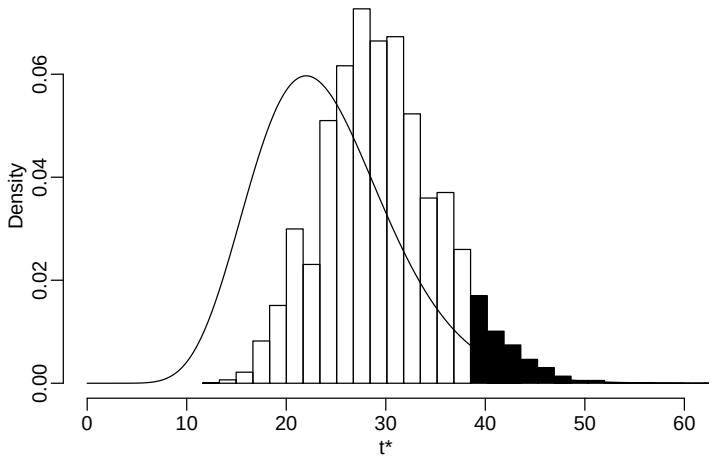

Key points

- ☐ Tests can be performed using resampling/simulation
- ☐ Must take account of null hypothesis, by
  - modifying sampling scheme to satisfy null hypothesis
  - inverting confidence interval (pivot test)
- ☐ Can use Monte Carlo simulation to get approximations to exact tests — simulate from null distribution of data, conditional on observed value of sufficient statistic
- ☐ Sometimes obtain permutation tests — very similar to bootstrap tests

Linear regression

- ☐ Independent data (x_1, y_1), ..., (x_n, y_n) with

$$y_j = x_j^{\mathrm{T}}\beta + \varepsilon_j, \quad \varepsilon_j \sim (0, \sigma^2)$$

- ☐ Least squares estimates $\widehat{\beta}$, leverages h_j, residuals

$$e_j = \frac{y_j - x_j^{\mathrm{T}}\widehat{\beta}}{(1-h_j)^{1/2}} \quad \dot{\sim} \quad (0, \sigma^2)$$

- ☐ Design matrix X should be held fixed if possible, at least for designed experiments, as

$$\mathrm{var}(\widehat{\beta}) = \sigma^2 (X^{\mathrm{T}} X)^{-1}$$

  if model $y = X\beta + \varepsilon$ correct

Linear regression: Resampling schemes

- ☐ Two main resampling schemes
- ☐ Model-based resampling:

$$y_j^* = x_j^{\mathrm{T}}\widehat{\beta} + \varepsilon_j^*, \quad \varepsilon_j^* \sim \mathrm{EDF}(e_1 - \overline{e}, \ldots, e_n - \overline{e})$$

  - Fixes design but not robust to model failure
  - Assumes ε_j sampled from population
- ☐ Case resampling:

$$(x_j, y_j)^* \sim \mathrm{EDF}\{(x_1, y_1), \ldots, (x_n, y_n)\}$$

  - Varying design X but robust
  - Assumes (x_j, y_j) sampled from population
  - Usually design variation no problem; can prove awkward in designed experiments and when design sensitive.

Belgian phone calls

Annual number of calls (millions), with least squares (black), M-estimate (red), and least trimmed squares (blue) fits.

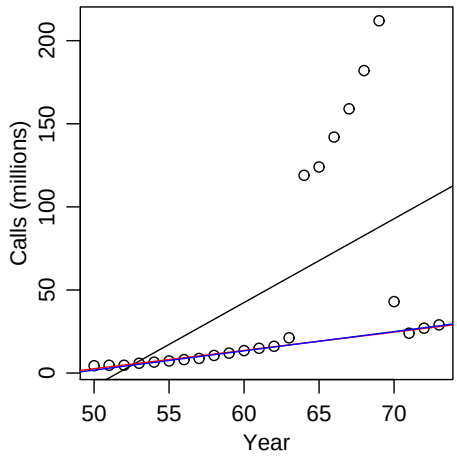

February 2021 – slide 96

Robust estimators

- ☐ Obvious need for robust fits here
- ☐ Two possibilities:
  - **M-estimator** result of minimising

$$\sum_{j=1}^{n} \rho\{(y_j - x_j^{\mathrm{T}}\beta)/\sigma\},$$

    with ρ chosen to downweight or eliminate big residuals
  - **least trimmed squares** result of minimising

$$\sum_{\lfloor (n+p+1)/2 \rfloor \text{ smallest}} (y_j - x_j^{\mathrm{T}}\beta)^2$$

  - Neither has reliable large-sample theory for estimates
- ☐ For phone data, years are fixed, so apply model-based resampling

February 2021 – slide 97

Belgian phone calls

Linear model fits, with nominal and bootstrap standard errors

	Intercept, 1950	Slope
Least squares	-8.0 (22.3)	5.0 (1.66)
M-estimate	2.6 (0.6)	1.1 (0.05)
bootstrap SE	(8.6)	(0.27)
Least trimmed squares	1.8	1.1
bootstrap SE	(4.5)	(0.14)

- ☐ Least squares fit is nonsense, with huge standard errors
- ☐ M-estimate nominal standard errors are much too small, don't take large residuals into account
- ☐ Least trimmed squares uses genetic algorithm, gives no nominal standard errors, so bootstrap is best (only?) option
- ☐ Least trimmed squares seems less perturbed by bootstrap including big residuals than is M-estimation

Cement data

y is the heat (calories per gram of cement) evolved while samples of cement set. The covariates are percentages by weight of four constituents, tricalciumaluminate x_1, tricalcium silicate x_2, tetracalcium alumino ferrite x_3 and dicalcium silicate x_4.

	x_1	x_2	x_3	x_4	y
1	7	26	6	60	78.5
2	1	29	15	52	74.3
3	11	56	8	20	104.3
4	11	31	8	47	87.6
5	7	52	6	33	95.9
6	11	55	9	22	109.2
7	3	71	17	6	102.7
8	1	31	22	44	72.5
9	2	54	18	22	93.1
10	21	47	4	26	115.9
11	1	40	23	34	83.8
12	11	66	9	12	113.3
13	10	68	8	12	109.4

Cement data

- ☐ Fit linear model
$$y = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 + \beta_4 x_4 + \varepsilon$$
and apply case resampling
- ☐ Covariates compositional: $x_1 + \cdots + x_4 \doteq 100\%$ so X almost collinear — smallest eigenvalue of $X^{\mathrm{T}}X$ is $l_5 = 0.0012$
- ☐ Plot of $\widehat{\beta}_1^*$ against smallest eigenvalue of $X^{*\mathrm{T}}X^*$ reveals that $\mathrm{var}^*(\widehat{\beta}_1^*)$ strongly variable
- ☐ Relevant subset for case resampling — post-stratification of output based on l_5^*?

Cement data

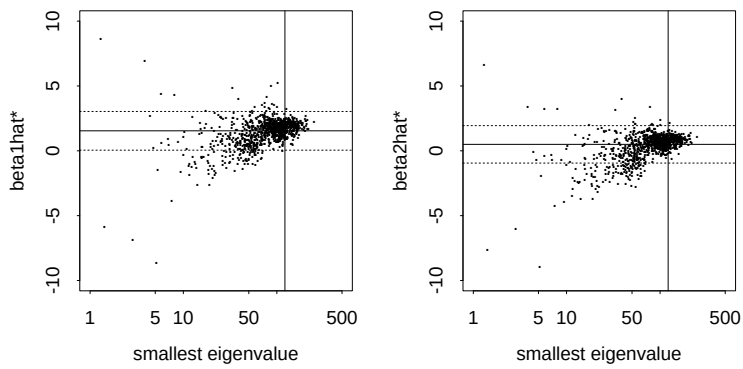

Cement data

Table 4: Standard errors of linear regression coefficients for cement data. Theoretical and error resampling assume homoscedasticity. Resampling results use $R = 1000$ samples, but last two rows are based only on those samples with the middle 500 and the largest 800 values of ℓ_5^*.

	$\widehat{\beta}_0$	$\widehat{\beta}_1$	$\widehat{\beta}_2$
Normal theory	70.1	0.74	0.72
Model-based resampling, $R = 1000$	66.3	0.70	0.69
Case resampling, all $R = 1000$	108.5	1.13	1.12
Case resampling, largest 500	68.4	0.76	0.71
Case resampling, largest 800	67.3	0.77	0.69

Survival data

dose x	117.5	235.0	470.0	705.0	940.0	1410
survival % y	44.000	16.000	4.000	0.500	0.110	0.700
	55.000	13.000	1.960	0.320	0.015	0.006
			6.120		0.019	

- ☐ Data on survival % for rats at different doses
- ☐ Linear model:

$$\log(\text{survival}) = \beta_0 + \beta_1 \text{dose}$$

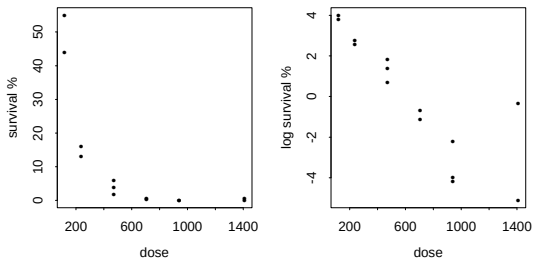

February 2021 – slide 103

Survival data

- ☐ Case resampling
- ☐ Replication of outlier: none (0), once (1), two or more (•).
- ☐ Model-based sampling including residual would lead to change in intercept but not slope.

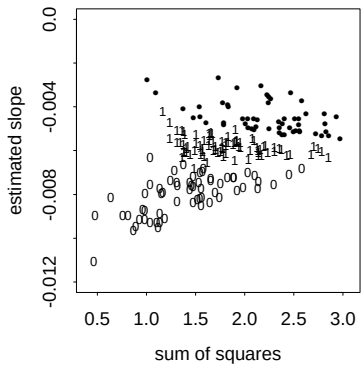

February 2021 – slide 104

Generalized linear model

- Response may be binomial, Poisson, gamma, normal, ...

$$y_j \quad \sim \quad \text{mean } \mu_j, \quad \text{variance } \phi V(\mu_j),$$

where $g(\mu_j) = x_j^{\mathrm{T}}\beta$ is linear predictor; $g(\cdot)$ is link function.

- MLE $\widehat{\beta}$, fitted values $\widehat{\mu}_j$, Pearson residuals

$$r_{Pj} = \frac{y_j - \widehat{\mu}_j}{\{V(\widehat{\mu}_j)(1-h_j)\}^{1/2}} \quad \dot{\sim} \quad (0, \phi).$$

- Bootstrapped responses

$$y_j^* = \widehat{\mu}_j + V(\widehat{\mu}_j)^{1/2}\varepsilon_j^*$$

where $\varepsilon_j^* \sim \mathrm{EDF}(r_{P1} - \overline{r}_P, \ldots, r_{Pn} - \overline{r}_P)$. However
  - possible that $y_j^* \notin \{0, 1, 2, \ldots, \}$
  - r_{Pj} not exchangeable, so may need stratified resampling

February 2021 – slide 105

Remission data

27 observations on how remission of patient depends on measure of cell activity, LI, with fit of logistic regression model

$$\mathrm{P}(\text{ remission }) = \frac{\exp(\beta_0 + \beta_1 LI)}{1 + \exp(\beta_0 + \beta_1 LI)}.$$

with $\widehat{\beta}_0 = -3.8\ (1.4)$, $\widehat{\beta}_1 = 2.9\ (1.2)$.

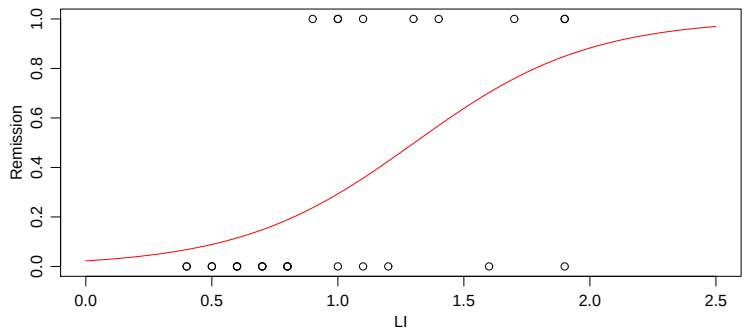

February 2021 – slide 106

Remission data

Case resampling, 95% confidence intervals for slope, $\widehat{\beta}_1 = 2.90$

GLM output	$(0.57, 5.22)$
Normal	$(-27.9, 30.7)$
Basic	$(-5.50, 4.62)$
Studentized	$(0.92, 5.00)$
Percentile	$(1.18, 11.29)$
BC_a	$(0.74, 7.57)$

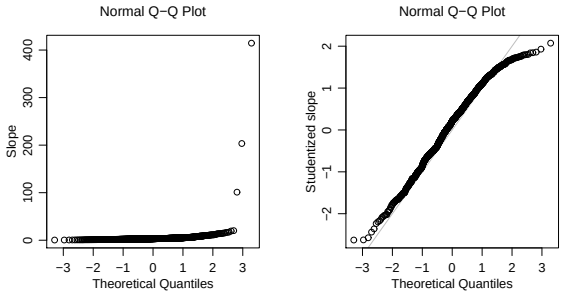

AIDS data

- ☐ Log-linear model: number of reports in row j and column k follows Poisson distribution with mean

$$\mu_{jk} = \exp(\alpha_j + \beta_k)$$

- ☐ Log link function

$$g(\mu_{jk}) = \log \mu_{jk} = \alpha_j + \beta_k$$

and variance function

$$\mathrm{var}(Y_{jk}) = \phi \times V(\mu_{jk}) = 1 \times \mu_{jk}$$

- ☐ Pearson residuals:

$$r_{jk} = \frac{Y_{jk} - \widehat{\mu}_{jk}}{\{\widehat{\mu}_{jk}(1 - h_{jk})\}^{1/2}}$$

- ☐ Model-based simulation:

$$Y^*_{jk} = \widehat{\mu}_{jk} + \widehat{\mu}_{jk}^{1/2} \varepsilon^*_{jk}$$

AIDS data

Diagnosis period		Reporting-delay interval (quarters):									Total reports to end
Year	Quarter	$0^{\dagger}$	1	2	3	4	5	6	$\cdots$	$\geq$14	of 1992
1988	1	31	80	16	9	3	2	8	$\cdots$	6	174
	2	26	99	27	9	8	11	3	$\cdots$	3	211
	3	31	95	35	13	18	4	6	$\cdots$	3	224
	4	36	77	20	26	11	3	8	$\cdots$	2	205
1989	1	32	92	32	10	12	19	12	$\cdots$	2	224
	2	15	92	14	27	22	21	12	$\cdots$	1	219
	3	34	104	29	31	18	8	6	$\cdots$		253
	4	38	101	34	18	9	15	6	$\cdots$		233
1990	1	31	124	47	24	11	15	8	$\cdots$		281
	2	32	132	36	10	9	7	6	$\cdots$		245
	3	49	107	51	17	15	8	9	$\cdots$		260
	4	44	153	41	16	11	6	5	$\cdots$		285
1991	1	41	137	29	33	7	11	6	$\cdots$		271
	2	56	124	39	14	12	7	10	$\cdots$		263
	3	53	175	35	17	13	11	2			306
	4	63	135	24	23	12	1				258
1992	1	71	161	48	25	5					310
	2	95	178	39	6						318
	3	76	181	16							273
	4	67	66								133

AIDS data

- ☐ Poisson two-way model deviance 716.5 on 413 df — indicates strong overdispersion: $\phi > 1$, so Poisson model implausible
- ☐ Residuals highly inhomogeneous — exchangeability doubtful

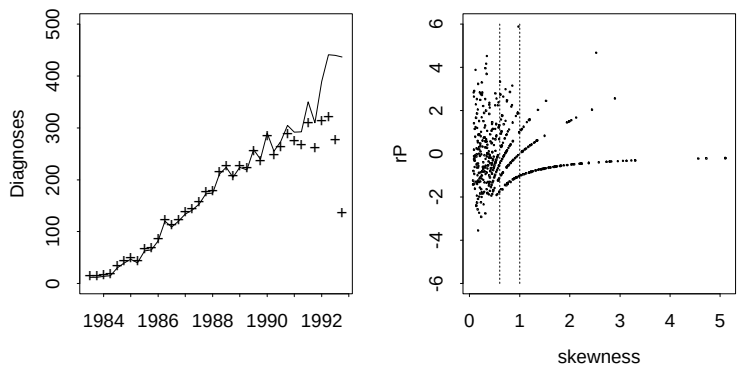

AIDS data: Prediction intervals

- To estimate prediction error:
  - simulate complete table y^*_{jk};
  - estimate parameters from incomplete y^*_{jk}
  - get estimated row totals and 'truth'

$$\widehat{\mu}^*_{+,j} = e^{\widehat{\alpha}^*_j} \sum_{k \text{ unobs}} e^{\widehat{\beta}^*_k}, \quad y^*_{+,j} = \sum_{k \text{ unobs}} y^*_{jk}.$$

  - Prediction error

$$\frac{y^*_{+,j} - \widehat{\mu}^*_{+,j}}{\widehat{\mu}^{*1/2}_{+,j}}$$

    studentized so more nearly pivotal.
- Form prediction intervals from R replicates.

February 2021 – slide 111

AIDS data: Resampling plans

- Resampling schemes:
  - parametric simulation, fitted Poisson model
  - parametric simulation, fitted negative binomial model
  - nonparametric resampling of r_P
  - stratified nonparametric resampling of r_P
- Stratification based on skewness of residuals, equivalent to stratifying original data by values of fitted means
- Take strata for which

$$\widehat{\mu}_{jk} < 1, \quad 1 \le \widehat{\mu}_{jk} < 2, \quad \widehat{\mu}_{jk} \ge 2$$

February 2021 – slide 112

AIDS data: Results

- Deviance/df ratios for the sampling schemes, $R = 1000$.
- Poisson variation inadequate.
- 95% prediction limits.

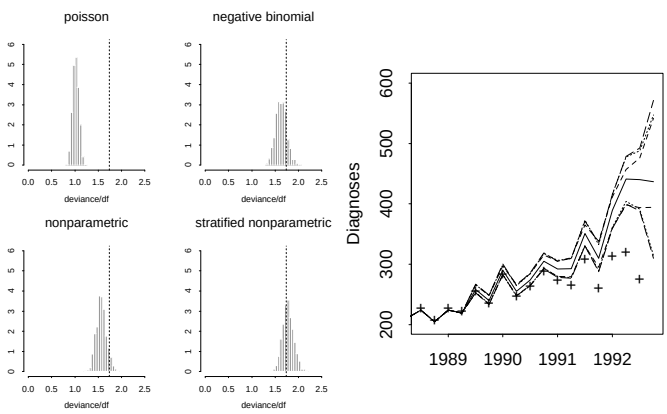

February 2021 – slide 113

AIDS data: Semiparametric model

- ☐ More realistic: generalized additive model

$$\mu_{jk} = \exp\{\alpha(j) + \beta_k\},$$

  where $\alpha(j)$ is locally-fitted smooth.
- ☐ Same resampling plans as before
- ☐ 95% intervals now generally narrower and shifted upwards

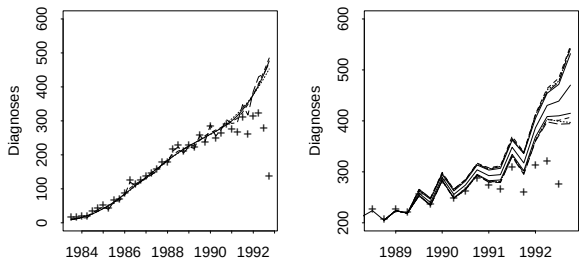

February 2021 – slide 114

Key points

- ☐ Key assumption: independence of cases
- ☐ Two main resampling schemes for regression settings:
  - Model-based
  - Case resampling
- ☐ Intermediate schemes possible
- ☐ Can help to reduce dependence on assumptions needed for regression model
- ☐ These two basic approaches also used for more complex settings (time series, ...), where data are dependent

February 2021 – slide 115