STATISTICS 568
DESIGN AND ANALYSIS OF EXPERIMENTS

Doug Wiens*

March 16, 2015

Contents

Part I

INTRODUCTION

1. Introduction

- Agricultural experiments – e.g. compare varieties of crops (= 'treatments'), 'plots' of experimental units arranged in homogeneous 'blocks'; within a block each treatment is applied ('randomization' needed here – why?); role of 'replication' to measure and control 'experimental error'.

Sir Ronald Fisher; Frank Yates; Rothamsted Agricultural Experimental Station (1930s)

- Goals

  – treatment comparisons (agriculture, manufacturing, ...)

  - variable screening (response can depend upon many factors – which are important?)

  - response surface exploration (explore effect of the – important – factors on the response variable)

  - system optimization (determine optimal levels of the factors)

  - system robustness (make the system 'robust' against hard-to-control factors)

- Planning an experiment – read §1.2

  - state the objective

  - choose response

  - choose factors and levels

  - choose experimental plan (our focus, in this course)

  - perform the experiment
  - analyze the data
  - draw conclusions, make recommendations

- Fisher's three fundamental principles
  - **replication**
    An *experimental unit* is the basic object of study – machine, crop, time period, ... – to which a treatment is applied. [*Factors* are variables expected to influence the outcome, and each is applied at one of a number of predetermined *levels*. A *treatment* is a combination of factor levels. (So if there is only one factor, treatments and levels are the same thing.)] Typically each treatment is applied to each of several units, thus decreasing the variance of the estimates of the treatment effects. (Replicates – applications to distinct units – not merely repetitions on the same unit.)

- **randomization**
  Randomize the allocation of units to treatments, the order in which the treatments are applied, the order in which the responses are measures ... e.g. clinical trials – physician bias, etc.

- **blocking**
  A block is a group of homogeneous units (fields, stretches of time, ...). Treatments are compared within blocks, so that between-block effects (differing ground fertilities, differences between days, ...) are eliminated.
  Randomization + blocking – 'block what you can, randomize what you cannot'.

- Example: typing/keyboards
  Two keyboards (A/B) to be compared. Six manuscripts given to a typist, who types each. There are considerable differences between the manuscripts, so they are viewed as 'blocks', and the keyboards are the 'treatments'. A possible layout

(rows represent 'time'):

Block:	1	2	3	4	5	6
	A	A	A	A	A	A
	B	B	B	B	B	B

An obvious problem ... randomize:

Block:	1	2	3	4	5	6
	A	B	A	B	A	A
	B	A	B	A	B	B

with some imbalance ... balanced randomization?

- See 'miscellaneous resources' on course web site – (i) review regression, if necessary; (ii) learn the basics of R, if necessary (I will help you get started, but see me <u>now</u> – this is a limited time offer.)

- R code for a regression example on course web site.

Part II

SINGLE FACTOR EXPERIMENTS

2. Single factor experiments

- **One way layout – example**
  Pulp 'reflectance' (brightness) classified by shift operator – each of four operators (A, B, C, D) produce pulp at a pulp mill; the quality is determined by the 'reflectance'. Are there significant differences arising from operators? Each operator is assigned to produce five batches:

A	B	C	D
59.8	59.8	60.7	61.0
⋮	⋮	⋮	⋮
60.8	59.9	60.9	60.5

  Randomization – of the 20 batches, 5 are assigned to each operator; all $\binom{20}{5,5,5,5}$ allocations should be equally likely. Randomly permute $1, \ldots, 20$ etc.

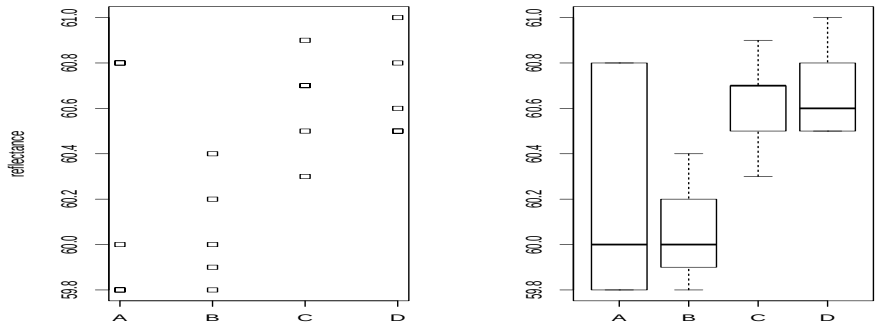

Pulp reflectance data by operator

- **Model**: k treatments, n_i observations on each ($N = \sum n_i$); for $i = 1, ..., k$ we have

$$\begin{aligned} y_{ij} &= \eta + \tau_i + \varepsilon_{ij},\ j = 1, ..., n_i; \\ \varepsilon_{ij} &\sim i.i.d.\ N(0, \sigma^2). \end{aligned}$$

  - Constraints:
    (i) zero-sum is most common:$\sum \tau_i = 0$. With differing group sizes we use instead $\sum n_i \tau_i = 0$. Can be justified as follows: define $\mu_i = E\left[y_{ij}\right]$, $\mu = \frac{1}{N} \sum n_i \mu_i = E\left[\bar{y}_{..}\right]$; now define $\eta = \mu, \tau_i = \mu_i - \mu$. Interpretation: 'overall mean', 'effect specific to the treatment'.
    (ii) 'baseline constraint' sometimes used: $\tau_1 = 0$; then $\eta = E\left[y_{1j}\right]$.

- Decomposition of SS (zero-sum constraint):
  We attempt to 'explain' the differences in the y_{ij} by ascribing them to any measurable effects

(hence the importance of the design). In this case the only possible explanatory factors are 'treatment' (operator) and 'random error':

$$\begin{aligned} y_{ij} &= \bar{y}_{..} + (\bar{y}_{i.} - \bar{y}_{..}) + \left(y_{ij} - \bar{y}_{i.}\right) \\ &= \hat{\eta} + \hat{\tau}_i + e_{ij}. \end{aligned}$$

These $(\hat{\eta}, \hat{\tau}_i)$ are the LSEs - they minimize

$$S(\eta, \tau_1, ..., \tau_k) = \sum_{i,j} \left(y_{ij} - \eta - \tau_i\right)^2.$$

Proof: Under the constraints we have

$$S = \sum_{i,j} \left\{ \left(y_{ij} - \bar{y}_{i.}\right) + (\hat{\tau}_i - \tau_i) + (\hat{\eta} - \eta) \right\}^2$$

$$= \sum_{i,j} \left(y_{ij} - \bar{y}_{i.}\right)^2 + \sum_i n_i (\hat{\tau}_i - \tau_i)^2 + N(\hat{\eta} - \eta)^2,$$

$etc.$

For the analysis of variance we decompose the

'total sum of squares' as

$$\begin{aligned} SST &= \sum_{i,j} \left(y_{ij} - \bar{y}_{..} \right)^2 \\ &= \sum_i n_i \left(\bar{y}_{i.} - \bar{y}_{..} \right)^2 + \sum_{i,j} \left(y_{ij} - \bar{y}_{i.} \right)^2 \\ &= \sum_i n_i \hat{\tau}_i^2 + (n-k)\, \hat{\sigma}^2 \stackrel{def}{=} SSTr + SSE. \end{aligned}$$

The notation implies the statement that

$$MSE = \frac{SSE}{n-k}$$

is an unbiased estimate of σ^2. Reason:

$$\frac{1}{n-k} \sum_{i,j} \left(y_{ij} - \bar{y}_{i.} \right)^2 = \frac{1}{n-k} \sum_i (n_i - 1)\, S_i^2 \text{ etc.}$$

Also $E\left[\hat{\tau}_i\right] = \tau_i$ and

$$E\left[MSTr\right] = \frac{E\left[SSTr\right]}{k-1} = \sigma^2 + \frac{1}{k-1} \sum_i n_i \tau_i^2. \tag{2.1}$$

Calculation in text; will be done later as well.

ANOVA table

Source	df	SS	$E[MS]$
Treatment	$k-1$	SSTr	$\sigma^2 + \frac{1}{k-1}\sum_i n_i\tau_i^2$
Error	$N-k$	SSE	σ^2
Total	$N-1$	SST	

- Likelihood ratio test of $H_0 : \tau_1 = \cdots = \tau_k (= 0)$ is based on

$$F = \frac{MSTr}{MSE} \sim F_{N-k}^{k-1}\left(\lambda^2 = \frac{\sum_i n_i\tau_i^2}{\sigma^2}\right),$$

the 'noncentral' F with noncentrality parameter λ (central if H_0 is true). This can be obtained as an instance of the 'General Linear Hypothesis'

being explored in assignment 1: F is

$$\begin{aligned} F &= \frac{\frac{\min_{H_0} S(\eta,\tau) - \min S(\eta,\tau)}{\nabla df}}{\frac{\min S(\eta,\tau)}{df}} \\ &= \frac{\frac{\min S(\eta,0) - SSE}{\nabla df}}{\frac{SSE}{N-k}} \\ &= \frac{\frac{SST-SSE}{(N-1)-(N-k)}}{\frac{SSE}{N-k}} \\ &= \frac{MSTr}{MSE}, \end{aligned}$$

and $\sigma^2\lambda^2 = \sum_i n_i\tau_i^2$, as in assignment 1.

If $k = 2$ this is the two-sample t-test, with a pooled estimate of the variance.

- If $\mathbf{z} \sim N\left(\mu, \sigma^2 \mathbf{I}_n\right)$ then $\frac{\|\mathbf{z}-\mu\|^2}{\sigma^2} \sim \chi^2_n$ but $\frac{\|\mathbf{z}\|^2}{\sigma^2}$ instead has the noncentral χ^2_n distribution with noncentrality parameter $\lambda^2 = \frac{\|\mu\|^2}{\sigma^2}$.

- If $X \sim \chi^2_n\left(\lambda_1^2\right)$, independently of $Y \sim \chi^2_m\left(\lambda_2^2\right)$, then $F = \frac{X/n}{Y/m}$ has the doubly noncentral $F^n_m\left(\lambda_1^2, \lambda_2^2\right)$ distribution. When $\lambda_2^2 = 0$ (the most common case) this is the singly noncentral F, or merely 'noncentral $F^n_m\left(\lambda_1^2\right)$'.

- The noncentral $F^n_m\left(\lambda^2\right)$ is 'stochastically increasing in λ^2'; this means that for any point f, the probability $P\left(F^n_m\left(\lambda^2\right) > f\right)$ that the F will exceed f is increasing in λ^2. Thus if we carry out a level α test of H_0 as above, then the critical value is f_α defined by

$$P\left(F^{k-1}_{N-k} > f_\alpha\right) = \alpha$$

(the central F) and the power is $P\left(F^{k-1}_{N-k}\left(\lambda^2\right) > f_\alpha\right)$.

- One can determine a common group size n, in order to obtain a particular power, if all other terms are known or can be reliably estimated – there is an R function to carry this out.

```
pwr.anova.test(k = 4, n = NULL, f = .5,
sig.level = .05, power = .8)

Balanced one-way anova power calculation

              k = 4
              n = 11.92613
              f = 0.5
      sig.level = 0.05
          power = 0.8

 NOTE: n is number in each group

# need n=12 for this power against this 'effect'
# function returns whichever argument is 'NULL'
# The effect size f has f^2 =
average(tau1^2,..,tauk^2)/sigmasqd
 (so f = lambda/sqrt(N), where N = kn)
# With n = 5, power = .35 against f = .5
```

3. Multiple comparisons; examples

- Multiple comparisons: suppose that the hypothesis of equal (hence no) treatment effects is rejected, so that at least two of the treatment effects are significantly different. Consider the test of $\tau_i - \tau_j = 0$ based on

$$t_{ij} = \frac{\bar{y}_{j.} - \bar{y}_{i.}}{\hat{\sigma}\sqrt{\frac{1}{n_j} + \frac{1}{n_i}}}.$$

The Bonferroni method declares τ_j to be significantly different than τ_i if $\left|t_{ij}\right| > t_{N-k,\frac{\alpha}{2k'}}$, the upper $100(1 - \frac{\alpha}{2k'})\%$ point of the t_{N-k} distribution. Here $k' = \binom{k}{2}$, the number of comparisons being made. Controls the total probability of a Type I error:

$$\begin{aligned} & P\,(\text{Type I error}) \\ = \;& P\,(\text{at least one false declaration of significance}) \\ \leq \;& \sum_l P\,(\text{declaration '}l\text{' is incorrect}) \\ = \;& k'\frac{\alpha}{k'} = \alpha. \end{aligned}$$

This leads to simultaneous confidence intervals:

$$\bar{y}_{j.} - \bar{y}_{i.} \pm t_{N-k,\frac{\alpha}{2k'}} \hat{\sigma} \sqrt{\frac{1}{n_j} + \frac{1}{n_i}}.$$

- Tukey method: gives <u>exact</u> simultaneous CIs, when all $n_i = n$. First assume this. Determine a value $q_{k,N-k,\alpha}$ as described below, and declare τ_j to be significantly different than τ_i if $\left|t_{ij}\right| > q_{k,N-k,\alpha}/\sqrt{2}$. The probability of a Type I error is

$$\begin{aligned}
& P_{H_0}\left(\begin{array}{c}\text{at least one pair } \bar{y}_{j.}, \bar{y}_{i.} \text{ has} \\ \left|t_{ij}\right| > q_{k,N-k,\alpha}/\sqrt{2}\end{array}\right) \\
= \; & P_{H_0}\left(\frac{\max_i \bar{y}_{i.} - \min_i \bar{y}_{i.}}{\hat{\sigma}\sqrt{\frac{1}{n}+\frac{1}{n}}} > \frac{q_{k,N-k,\alpha}}{\sqrt{2}}\right) \\
= \; & P_{H_0}\left(\frac{\max_i \bar{y}_{i.} - \min_i \bar{y}_{i.}}{\sigma/\sqrt{n}} \Big/ \frac{\hat{\sigma}}{\sigma} > q_{k,N-k,\alpha}\right) \\
= \; & \alpha,
\end{aligned}$$

if $q_{k,N-k,\alpha}$ is the upper α-point in the distribution of the range of k i.i.d. $N(0,1)$ r.v.s divided by an independent $\sqrt{\chi^2_{N-k}/(N-k)}$, and $N = nk$.

It can be shown that in the unbalanced case the probability of a Type I error is <u>at most</u> α. Inverting these tests gives simultaneous CIs

$$\bar{y}_{j.} - \bar{y}_{i.} \pm \frac{q_{k,N-k,\alpha}}{\sqrt{2}} \hat{\sigma} \sqrt{\frac{1}{n_j} + \frac{1}{n_i}}.$$

- Quantitative predictor – orthogonal polynomials: When the predictor is quantitative (perhaps coded as LO, MID, HI for instance) at k levels one can fit a polynomial regression model $E[Y|x] = \beta_0 + \beta_1 P_1(x) + \cdots + \beta_p P_p(x)$, where $P_j(x)$ is a polynomial of degree j and $p < k$ (typically $p = 1, 2$). For instance for levels $x = -1, 0, 1$ and $p = 2$ one puts $P_1(x) = x$, $P_1(x) = 3x^2 - 2$. Then

$$\begin{aligned} E[Y|x = -1] &= \beta_0 - \beta_1 + \beta_2, \\ E[Y|x = 0] &= \beta_0 - 2\beta_2, \\ E[Y|x = 1] &= \beta_0 + \beta_1 + \beta_2, \end{aligned}$$

with

$$\begin{aligned}\beta_1 &= \frac{E[Y|x=1]-E[Y|x=-1]}{2},\\ \beta_2 &= \frac{\begin{array}{c}(E[Y|x=1]-E[Y|x=0])\\ -(E[Y|x=0]-E[Y|x=-1])\end{array}}{6};\end{aligned}$$

these represent the linear and quadratic effects, respectively. For these regressors $\mathbf{X}'\mathbf{X}$ is diagonal, and it is usual to convert the model $E[\mathbf{y}]=\mathbf{X}\boldsymbol{\beta}$ to $E[\mathbf{y}]=\mathbf{Z}\boldsymbol{\gamma}$, where

$$\mathbf{Z}=\mathbf{X}\left(\mathbf{X}'\mathbf{X}\right)^{-1/2},\boldsymbol{\gamma}=\left(\mathbf{X}'\mathbf{X}\right)^{1/2}\boldsymbol{\beta}.$$

The γ_j have interpretations similar to those of the β_j – they are the coefficients of the $P_j(x)/\sqrt{(\mathbf{X}'\mathbf{X})_{jj}}$ – and satisfy $\hat{\boldsymbol{\gamma}}=\mathbf{Z}'\mathbf{y}$, with

$$cov[\hat{\boldsymbol{\gamma}}]=\sigma^2\mathbf{I}.$$

More details in text – read §2.3.

- Pulp data – R code on web site.

```
> fit = aov(y ~operator)
> summary(fit)
            Df Sum Sq Mean Sq F value Pr(>F)
operator     3   1.34  0.4467   4.204 0.0226 *
Residuals   16   1.70  0.1062
```

Differences between operators significant; which are different?

```
> TukeyHSD(fit)
  Tukey multiple comparisons of means
    95% family-wise confidence level


Fit: aov(formula = y ~operator)
```

```
$operator
     diff         lwr       upr     p adj
B-A -0.18 -0.76981435 0.4098143 0.8185430
C-A  0.38 -0.20981435 0.9698143 0.2903038
D-A  0.44 -0.14981435 1.0298143 0.1844794
C-B  0.56 -0.02981435 1.1498143 0.0657945
D-B  0.62  0.03018565 1.2098143 0.0376691
D-C  0.06 -0.52981435 0.6498143 0.9910783
```

D and B significantly different at 5% level.

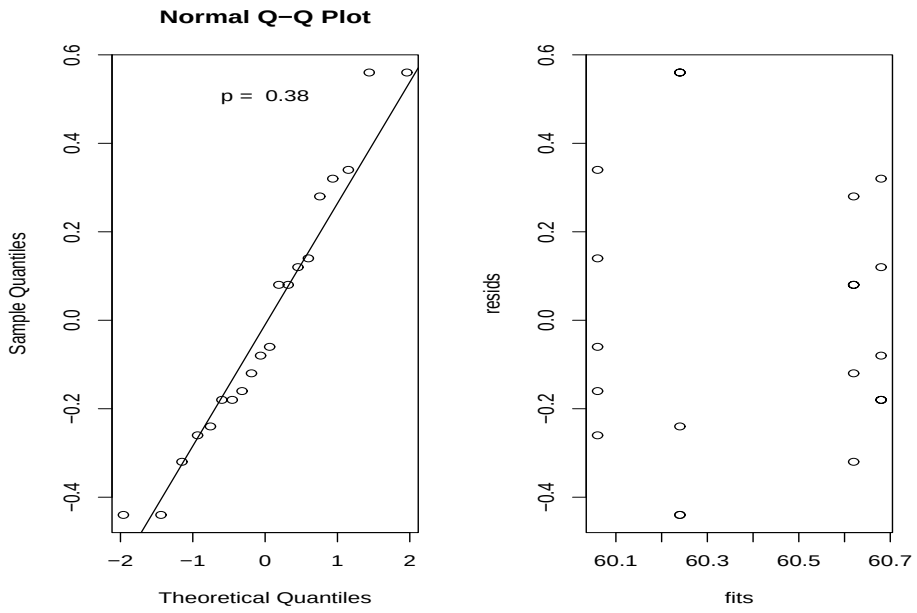

Diagnostic plots; Shapiro-Wilks normality test

- A robust (against non-normality) alternative to the F-test is the Kruskal-Wallis test. It is essentially equivalent to replacing the observations by their ranks, and performing the F-test on the ranked data. The differences are in the manner in which the null distribution is approximated – K-W exploits the fact that $SSTotal$ is fixed (apart from ties) and obtains a χ^2-approximation to $SSTr$.

```
> fit.rank = aov(ranks ~operator)
> summary(fit.rank)
            Df Sum Sq Mean Sq F value Pr(>F)
operator     3  259.3   86.43   3.469 0.0411 *
Residuals   16  398.7   24.92
---

> kruskal.test(y, operator)
        Kruskal-Wallis rank sum test
data:  y and operator
Kruskal-Wallis chi-squared = 7.4874,
df = 3, p-value = 0.05788
```

- **Sampling distributions** via quadratic forms:
  Write $y_{ij} = \eta + \tau_i + \varepsilon_{ij}$; decompose $\sum_{i,j} \varepsilon_{ij}^2 = \boldsymbol{\varepsilon}'\boldsymbol{\varepsilon}$ as

$$\begin{aligned}
\boldsymbol{\varepsilon}'\boldsymbol{\varepsilon} &= \sum_{i,j} \left(y_{ij} - \eta - \tau_i\right)^2 \\
&= \sum_{i,j} \left(y_{ij} - \bar{y}_{i.}\right)^2 + \sum_i n_i \left(\bar{y}_{i.} - \bar{y}_{..} - \tau_i\right)^2 \\
&\quad + N\left(\bar{y}_{..} - \eta\right)^2 \\
&= \sum_{i,j} \left(\varepsilon_{ij} - \bar{\varepsilon}_{i.}\right)^2 + \sum_i n_i \left(\bar{\varepsilon}_{i.} - \bar{\varepsilon}_{..}\right)^2 + N\bar{\varepsilon}_{..}^2 \\
&= \boldsymbol{\varepsilon}'\mathbf{Q}_3\boldsymbol{\varepsilon} + \boldsymbol{\varepsilon}'\mathbf{Q}_2\boldsymbol{\varepsilon} + \boldsymbol{\varepsilon}'\mathbf{Q}_1\boldsymbol{\varepsilon}.
\end{aligned}$$

  Here

$$\begin{aligned}
\mathbf{Q}_1 &= \frac{1}{N} 1_N 1_N' \stackrel{def}{=} \mathbf{J}_N, \\
\mathbf{Q}_3 &= \mathbf{I}_N - \oplus_{i=1}^k \mathbf{J}_{n_i} = \oplus_{i=1}^k \left(\mathbf{I}_{n_i} - \mathbf{J}_{n_i}\right), \\
\mathbf{Q}_2 &= \mathbf{I}_N - Q_3 - Q_1 = \oplus_{i=1}^k \mathbf{J}_{n_i} - \mathbf{J}_N.
\end{aligned}$$

  Note:
  (i) $\mathbf{Q}_1$ is idempotent, with $rank = trace = 1$;
  (ii) $\mathbf{Q}_2$ is idempotent, with $rank = trace = k-1$;
  (iii) $\mathbf{Q}_3$ is idempotent, with $rank = trace = \sum_i (n_i - 1) = N - k$;
  (iv) $\mathbf{Q}_i\mathbf{Q}_j = \mathbf{0}$ for each $i \neq j$.

4. Sampling distributions; random effects model

- From (i) - (iv) at the end of the last lecture, we can write

$$\mathbf{Q}_i = \mathbf{V}_i \mathbf{V}_i',$$

where $\mathbf{V}_i : N \times r_i$ ($r_i = rk(\mathbf{Q}_i)$) and $\mathbf{V}_i' \mathbf{V}_i = \mathbf{I}_{r_i}$ and $\mathbf{V}_i' \mathbf{V}_j = \mathbf{0}$ for each $i \neq j$. Then

$$\mathbf{V}_{N \times N} = (\mathbf{V}_1 \vdots \mathbf{V}_2 \vdots \mathbf{V}_3)$$

is orthogonal, and so $\boldsymbol{\varepsilon}'\boldsymbol{\varepsilon} = \|\mathbf{V}'\boldsymbol{\varepsilon}\|^2$. Now

$$\mathbf{V}'\boldsymbol{\varepsilon} = \begin{pmatrix} \mathbf{V}_1'\boldsymbol{\varepsilon} \\ \mathbf{V}_2'\boldsymbol{\varepsilon} \\ \mathbf{V}_3'\boldsymbol{\varepsilon} \end{pmatrix} \sim N\left(\mathbf{0}, \sigma^2 \mathbf{I}_N\right),$$

so the elements of $\mathbf{V}'\boldsymbol{\varepsilon}$ are i.i.d. $N\left(0, \sigma^2\right)$ – just like those of $\boldsymbol{\varepsilon}$ – but

$$\boldsymbol{\varepsilon}'\mathbf{Q}_i\boldsymbol{\varepsilon} = \left\|\mathbf{V}_i'\boldsymbol{\varepsilon}\right\|^2 \sim \sigma^2 \chi^2_{r_i};$$

all three are independently distributed.

  - **Cochran's Theorem** states that these conclusions hold iff $\sum r_i = N$ (and so they could have been inferred from this). These ranks are called 'degrees of freedom'.

– Note that $\boldsymbol{\varepsilon}'\mathbf{Q}_3\boldsymbol{\varepsilon} = SSE$. Under H_0: all $\tau_i = 0$, $\boldsymbol{\varepsilon}'\mathbf{Q}_2\boldsymbol{\varepsilon} = SSTr$, and so

$$F = \frac{MSTr}{MSE} \sim \frac{\frac{\sigma^2\chi^2_{k-1}}{k-1}}{\frac{\sigma^2\chi^2_{N-k}}{N-k}} \sim F^{k-1}_{N-k}.$$

- Non-null distribution: In general,

$$SSTr = \mathbf{y}'\mathbf{Q}_2\mathbf{y},$$

and

$$\mathbf{y} = \boldsymbol{\varepsilon} + \eta 1_N + \boldsymbol{\tau}, \text{ where } \boldsymbol{\tau} = \begin{pmatrix} \tau_1 1_{n_1} \\ \vdots \\ \tau_k 1_{n_k} \end{pmatrix},$$

with $\mathbf{Q}_2 1_N = \mathbf{Q}_2\mathbf{Q}_1 1_N = \mathbf{0}$. Thus

$$\mathbf{Q}_2\mathbf{y} = \mathbf{Q}_2\left(\boldsymbol{\varepsilon} + \boldsymbol{\tau}\right),$$

and

$$SSTr = \left\|\mathbf{V}_2'\left(\boldsymbol{\varepsilon} + \boldsymbol{\tau}\right)\right\|^2,$$

where $\mathbf{V}_2'\left(\boldsymbol{\varepsilon} + \boldsymbol{\tau}\right) \sim N\left(\mathbf{V}_2'\boldsymbol{\tau}, \sigma^2\mathbf{I}_{k-1}\right)$ (still independent of SSE – why?). Then

$$\frac{SSTr}{\sigma^2} \sim \chi^2_{k-1}\left(\lambda^2\right),$$

the noncentral χ^2_{k-1} with noncentrality parameter

$$\lambda^2 = \frac{\left\|\mathbf{V}_2'\boldsymbol{\tau}\right\|^2}{\sigma^2} = \frac{\boldsymbol{\tau}'\mathbf{Q}_2\boldsymbol{\tau}}{\sigma^2} = \frac{\sum_i n_i \left(\bar{\tau}_{i.} - \bar{\tau}_{..}\right)^2}{\sigma^2} = \frac{\sum_i n_i \tau_i^2}{\sigma^2}.$$

Then $F = \frac{MSTr}{MSE} \sim F^{k-1}_{N-k}\left(\lambda^2 = \frac{\sum_i n_i \tau_i^2}{\sigma^2}\right)$.

- As a general rule, note that λ^2 is obtained by replacing the observations by their means, in the relevant SS, and then dividing by σ^2. The calculation above also yields, using

  $$E\left[\chi^2_{k-1}\left(\lambda^2\right)\right] = (k-1) + \lambda^2,$$

  that

  $$E\left[MSTr\right] = \sigma^2 + \frac{\sum_i n_i \tau_i^2}{k-1},$$

  as claimed at (2.1).

- A useful general formula:
  if $E\left[\mathbf{y}\right] = \boldsymbol{\mu}$, cov$[\mathbf{y}] = \boldsymbol{\Sigma}$, then (how?)

  $$E\left[\mathbf{y}'\mathbf{Q}\mathbf{y}\right] = \boldsymbol{\mu}'\mathbf{Q}\boldsymbol{\mu} + tr\mathbf{Q}\boldsymbol{\Sigma}.$$

- Random effects model

  In the pulp experiment, suppose that the 4 operators were randomly chosen from a population of operators. Then the effects would be random rather than fixed, and an appropriate model might be

$$y_{ij} = \eta + \tau_i + \varepsilon_{ij},\ j = 1, ..., n_i,$$

  as before, but with the assumption that

$$\begin{aligned} \varepsilon_{ij} &\sim i.i.d.\ N(0, \sigma^2), \\ \tau_i &\sim i.i.d.\ N(0, \sigma_\tau^2), \end{aligned}$$

  with ε_{ij} and τ_i independent.

  - σ_τ^2 and σ^2 are the 'variance components' measuring between-treatment and within-treatment variation.

  - A calculation gives that, with SSE and $SSTr$ computed as before, one has

$$\begin{aligned} E[MSE] &= \sigma^2, \\ E[MSTr] &= \sigma^2 + n'\sigma_\tau^2, \end{aligned}$$

where

$$\begin{aligned} n' &= \frac{1}{k-1}\left(N - \frac{\sum_i n_i^2}{N}\right) \\ (&= n \text{ if all } n_i = n). \end{aligned}$$

Now unbiased estimates of the variance components are obtained by equating the observed mean squares to their expectations:

$$\begin{aligned} \hat{\sigma}^2 &= MSE, \\ \hat{\sigma}_\tau^2 &= \frac{MSTr - MSE}{n'}. \end{aligned}$$

Typically we instead take $\hat{\sigma}_\tau^2 = \left(\frac{MSTr-MSE}{n'}\right)^+$.

– Since now $E\left[y_{ij}\right] \equiv \eta$, there is some interest in making inferences about this parameter. With $\tau_\cdot = \sum_i n_i\tau_i/N$ we have

$$\bar{y}_{\cdot\cdot} = \eta + \tau_\cdot + \bar{\varepsilon}_{\cdot\cdot} \sim N\left(\eta, \sigma_{\bar{y}}^2\right),$$

where $\sigma_{\bar{y}}^2 = \frac{\sigma^2}{N} + \frac{\sum_i n_i^2}{N^2}\sigma_\tau^2$. In the balanced case (all $n_i = n$), $\sigma_{\bar{y}}^2 = \frac{\sigma^2 + n\sigma_\tau^2}{N}$ and we estimate

$\sigma^2_{\bar{y}}$ by replacing the variance components by their unbiased estimates, obtaining

$$S^2_{\bar{y}} = \frac{MSTr}{N}.$$

Since

$$\frac{S^2_{\bar{y}}}{\sigma^2_{\bar{y}}} = \frac{MSTr}{\sigma^2 + n\sigma^2_\tau} \sim \frac{\chi^2_{k-1}}{k-1},$$

independently of $\bar{y}_{..}$, we have that $\frac{\bar{y}_{..}-\eta}{S_{\bar{y}}} \sim t_{k-1}$, from which inferences are made – an example using the pulp data is on the course web site.

– Primary interest is on testing $H_0 : \sigma^2_\tau = 0$, implying no treatment effect. This is tested via the same F ratio as is computed in the fixed effects model, and the null distribution remains the same. This is established (almost) exactly as before. Thus in the pulp example, the p-value .0226 obtained in the fixed effects test is now associated with the hypothesis that $\sigma^2_\tau = 0$, leading to the conclusion that $\sigma^2_\tau > 0$. We find $\hat{\sigma}^2_\tau = (MSTr - MSE)/n = .068$.

– Note now that the non-null distribution of the F used to test $H_0 : \sigma_\tau^2 = 0$ is

$$F = \frac{MSTr}{MSE} \sim \frac{\sigma^2 + n\sigma_\tau^2}{\sigma^2} F_{N-k}^{k-1};$$

in this sense the situation (re power calculations) is simpler than in the fixed effects case.

- **Cochran's Theorem**: Let $\mathbf{x} = (X_1, ..., X_n)'$, where the $X_i \sim$ i.i.d. $N\left(0, \sigma^2\right)$. Suppose that

$$\mathbf{x}'\mathbf{x} = \sum_{i=1}^{k} \mathbf{x}'\mathbf{Q}_i\mathbf{x},$$

for symmetric matrices $\mathbf{Q}_i$ with ranks r_i, $i = 1, ..., k$.
The following are equivalent:
(i) the r.v.s $\mathbf{x}'\mathbf{Q}_i\mathbf{x}$ are independently distributed as $\sigma^2\chi_{r_i}^2$, $i = 1, ..., k$;
(ii) $\sum_{i=1}^{k} r_i = n$.

We establish the theorem through a series of preliminary results.

- **Lemma 1**: The r.v. $\mathbf{x}'\mathbf{Q}\mathbf{x}$ is $\sim \sigma^2\chi^2_r$ iff $\mathbf{Q}^2 = \mathbf{Q}$ and $rk(\mathbf{Q}) = r$.

  Proof: Write $\mathbf{Q} = \mathbf{V}\mathbf{D}\mathbf{V}'$, where $\mathbf{V}$ is orthogonal and $\mathbf{D} = diag(\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n)$, the eigenvalues of $\mathbf{Q}$. Then

$$\mathbf{x}'\mathbf{Q}\mathbf{x} = \mathbf{y}'\mathbf{D}\mathbf{y},$$

  where $\mathbf{V}'\mathbf{x} = \mathbf{y} \sim N\left(0, \sigma^2\mathbf{I}_n\right)$. The m.g.f. is $M(t) =$

$$E\left[\exp t\mathbf{x}'\mathbf{Q}\mathbf{x}\right] = E\left[\exp t\mathbf{y}'\mathbf{D}\mathbf{y}\right]$$

$$= \prod_{j=1}^{n} E\left[e^{t\lambda_j Y_j^2}\right] = \prod_{j=1}^{n}\left(1 - 2t\lambda_j\sigma^2\right)^{-1/2} \quad (4.1)$$

  since the $Y_j^2 \sim$ i.i.d. $\sigma^2\chi^2_1$ r.v.s . This holds for all t in some open neighbourhood of 0. This is the m.g.f. of $\sigma^2\chi^2_r$ iff

$$M(t) = \left(1 - 2t\sigma^2\right)^{-r/2}.$$

  Equivalently,

$$\prod_{j=1}^{n}\left(1 - s\lambda_j\right) = (1 - s)^r$$

for all $s = 2t\sigma^2$ in a neighbourhood of 0; by the uniqueness of polynomial factorizations this can hold iff r of the λ_j are ones and the rest are zeros, i.e. $\mathbf{D}^2 = \mathbf{D}$ and $rk(\mathbf{D}) = r$. Equivalently, $\mathbf{Q}^2 = \mathbf{Q}$ and $rk(\mathbf{Q}) = r$. □

At this point we already have the easy half of Cochran's Theorem: If the r.v.s $\mathbf{x}'\mathbf{Q}_i\mathbf{x}$ are distributed as $\sigma^2\chi^2_{r_i}$ then the $\mathbf{Q}_i$ are idempotent and so

$$\begin{aligned}
\mathbf{x}'\mathbf{x} &= \sum_{i=1}^{k} \mathbf{x}'\mathbf{Q}_i\mathbf{x} \\
&\Rightarrow \mathbf{x}'\left[\mathbf{I}_n - \sum_{i=1}^{k}\mathbf{Q}_i\right]\mathbf{x} \equiv \mathbf{0} \\
&\Rightarrow \mathbf{I}_n = \sum_{i=1}^{k}\mathbf{Q}_i \\
&\Rightarrow n = tr\sum_{i=1}^{k}\mathbf{Q}_i = \sum_{i=1}^{k} tr\mathbf{Q}_i = \sum_{i=1}^{k} r_i.
\end{aligned}$$

5. Proof of Cochran's Theorem

Craig and Sakamoto conjectured the following result; Craig's original proof was incorrect. The first correct proof was given by Matusita in 1949, and later that year Ogawa (who had also given an earlier but incorrect proof) gave another. This is Matusita's proof. Some related papers are on the course web site.

Lemma 2 (The Craig-Sakamoto-Matusita Theorem): The r.v.s $\mathbf{x}'\mathbf{Q}_1\mathbf{x}$, $\mathbf{x}'\mathbf{Q}_2\mathbf{x}$ are independently distributed iff $\mathbf{Q}_1\mathbf{Q}_2 = \mathbf{0}$.

Proof: From (4.1), for any $\mathbf{Q}$,

$$E\left[\exp t\mathbf{x}'\mathbf{Q}\mathbf{x}\right] = \left|\mathbf{I} - 2t\sigma^2\mathbf{D}\right|^{-1/2} = \left|\mathbf{I} - 2t\sigma^2\mathbf{Q}\right|^{-1/2}.$$

Thus the joint m.g.f. of $\mathbf{x}'\mathbf{Q}_1\mathbf{x}$, $\mathbf{x}'\mathbf{Q}_2\mathbf{x}$ is

$$\begin{aligned} M(t_1, t_2) &= E\left[\exp\left(t_1\mathbf{x}'\mathbf{Q}_1\mathbf{x} + t_2\mathbf{x}'\mathbf{Q}_2\mathbf{x}\right)\right] \\ &= \left|\mathbf{I} - 2t_1\sigma^2\mathbf{Q}_1 - 2t_2\sigma^2\mathbf{Q}_2\right|^{-1/2}, \end{aligned}$$

and we have independence iff

$$M(t_1, t_2) = M(t_1, 0)\, M(0, t_2), \qquad (5.1)$$

i.e. iff

$$|\mathbf{I} - s_1\mathbf{Q}_1 - s_2\mathbf{Q}_2| = |\mathbf{I} - s_1\mathbf{Q}_1|\,|\mathbf{I} - s_2\mathbf{Q}_2|\,, \tag{5.2}$$

for all (s_1, s_2) in a neighbourhood of $\mathbf{0}$. If $\mathbf{Q}_1\mathbf{Q}_2 = \mathbf{0}$ then (5.2) holds, hence so does (5.1). For the converse, first diagonalize $\mathbf{Q}_1$:

$$\mathbf{Q}_1 = \mathbf{V}\begin{pmatrix} \Lambda & 0 \\ 0 & 0 \end{pmatrix}\mathbf{V}',$$

where Λ is the diagonal matrix containing the r_1 nonzero eigenvalues of $\mathbf{Q}_1$. Define

$$\mathbf{B} = \mathbf{V}'\mathbf{Q}_2\mathbf{V} = \begin{pmatrix} \mathbf{B}_{11} & \mathbf{B}_{12} \\ \mathbf{B}_{21} & \mathbf{B}_{22} \end{pmatrix}.$$

Then

$$\mathbf{Q}_2 = \mathbf{V}\mathbf{B}\mathbf{V}'$$

and (5.2) becomes

$$\begin{aligned} & \left|\mathbf{I} - s_1\begin{pmatrix} \Lambda & 0 \\ 0 & 0 \end{pmatrix} - s_2\begin{pmatrix} \mathbf{B}_{11} & \mathbf{B}_{12} \\ \mathbf{B}_{21} & \mathbf{B}_{22} \end{pmatrix}\right| \\ = & \left|\mathbf{I} - s_1\begin{pmatrix} \Lambda & 0 \\ 0 & 0 \end{pmatrix}\right|\left|\mathbf{I} - s_2\begin{pmatrix} \mathbf{B}_{11} & \mathbf{B}_{12} \\ \mathbf{B}_{21} & \mathbf{B}_{22} \end{pmatrix}\right|, \end{aligned}$$

i.e.

$$\begin{aligned} & \left| \begin{pmatrix} \mathbf{I}_{r_1} - s_1\Lambda - s_2\mathbf{B}_{11} & -s_2\mathbf{B}_{12} \\ -s_2\mathbf{B}_{21} & \mathbf{I}_{n-r_1} - s_2\mathbf{B}_{22} \end{pmatrix} \right| \\ = \quad & |\mathbf{I}_{r_1} - s_1\Lambda| \, |\mathbf{I} - s_2\mathbf{B}| \, . \end{aligned}$$

The coefficient of $(-s_1)^{r_1}$ – the highest power of s_1 – on the rhs is $\left(\prod_{j=1}^{r_1} \lambda_j\right) |\mathbf{I} - s_2\mathbf{B}|$. On the lhs it arises from the highest power in the expansion of the minor $|\mathbf{I}_{r_1} - s_1\Lambda - s_2\mathbf{B}_{11}|$, times $|\mathbf{I}_{n-r_1} - s_2\mathbf{B}_{22}|$, and is $\left(\prod_{j=1}^{r_1} \lambda_j\right) |\mathbf{I}_{n-r_1} - s_2\mathbf{B}_{22}|$. Thus

$$|\mathbf{I}_{n-r_1} - s_2\mathbf{B}_{22}| = |\mathbf{I} - s_2\mathbf{B}| \, ,$$

and so $\mathbf{B}_{22}$ and $\mathbf{B}$ have the same eigenvalues. The sum of squares of the eigenvalues of a symmetric matrix is the sum of squares of the elements of the matrix; thus

$$\mathbf{B} = \begin{pmatrix} 0 & 0 \\ 0 & \mathbf{B}_{22} \end{pmatrix}$$

and so $\mathbf{Q}_1\mathbf{Q}_2 = \mathbf{0}$. □

- **Lemma 3**: Suppose that $\mathbf{x}'\mathbf{Q}\mathbf{x} = \sum_{i=1}^{k} \mathbf{x}'\mathbf{Q}_i\mathbf{x}$ for symmetric matrices $\mathbf{Q}, \mathbf{Q}_1, ..., \mathbf{Q}_k$, that $\mathbf{x}'\mathbf{Q}\mathbf{x} \sim \sigma^2\chi^2_r$, that $\mathbf{x}'\mathbf{Q}_i\mathbf{x} \sim \sigma^2\chi^2_{r_i}$ for $i = 1, ..., k-1$, and that $\mathbf{x}'\mathbf{Q}_k\mathbf{x} \geq 0$ for all $\mathbf{x}$. Then $\{\mathbf{x}'\mathbf{Q}_i\mathbf{x}\}_{i=1}^{k}$ are independently distributed, and hence $\mathbf{x}'\mathbf{Q}_k\mathbf{x} \sim \sigma^2\chi^2_{r_k}$ with $r_k = r - \sum_{i=1}^{k-1} r_i$.
  Proof: Using Lemmas 1 and 2, the conditions of this lemma are equivalent to
  (i) $\mathbf{Q} = \sum_{i=1}^{k} \mathbf{Q}_i$, all are p.s.d.,
  (ii) $\mathbf{Q}^2 = \mathbf{Q}$ and $rk(\mathbf{Q}) = r$,
  (iii) $\mathbf{Q}_i^2 = \mathbf{Q}_i$ and $rk(\mathbf{Q}_i) = r_i$ for $i = 1, ..., k-1$.
  We are to show that $\mathbf{Q}_1, ..., \mathbf{Q}_k$ are mutually orthogonal. This gives the independence of all $\{\mathbf{x}'\mathbf{Q}_i\mathbf{x}\}_{i=1}^{k}$. It also follows that

$$\sum_{i=1}^{k} \mathbf{Q}_i = \mathbf{Q} = \mathbf{Q}^2 = \sum_{i=1}^{k} \mathbf{Q}_i^2 = \sum_{i=1}^{k-1} \mathbf{Q}_i + \mathbf{Q}_k^2,$$

  hence $\mathbf{Q}_k^2 = \mathbf{Q}_k$ and so $\mathbf{x}'\mathbf{Q}_k\mathbf{x} \sim \sigma^2\chi^2_{r_k}$ with $r_k = rk(\mathbf{Q}_k)$. As well, since $\mathbf{Q}$ is idempotent,

$$r = tr \sum_{i=1}^{k} \mathbf{Q}_i = \sum_{i=1}^{k} tr\mathbf{Q}_i = \sum_{i=1}^{k} r_i, \tag{5.3}$$

  i.e. $r_k = r - \sum_{i=1}^{k-1} r_i$.

To establish the orthogonality first take $k = 2$. Then $\mathbf{Q} = \mathbf{Q}_1 + \mathbf{Q}_2$, all are p.s.d., and $\mathbf{Q}, \mathbf{Q}_1$ are idempotent with ranks r, r_1 respectively. Diagonalize $\mathbf{Q}$, and apply the same transformation to $\mathbf{Q}_1$ and $\mathbf{Q}_2$:

$$\mathbf{V}'\mathbf{Q}\mathbf{V} = \begin{pmatrix} \mathbf{I}_r & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix} = \mathbf{V}'\mathbf{Q}_1\mathbf{V} + \mathbf{V}'\mathbf{Q}_2\mathbf{V}.$$

The diagonal elements of a p.s.d. must be non-negative, thus the last $n - r$ diagonal elements of each of $\mathbf{V}'\mathbf{Q}_1\mathbf{V}$, $\mathbf{V}'\mathbf{Q}_2\mathbf{V}$ must be zero. If p_{ii} is any one of these, then since any submatrix

$$\begin{pmatrix} p_{jj} & p_{ji} \\ p_{ij} & p_{ii} \end{pmatrix}$$

of a p.s.d. matrix is p.s.d., we must have

$$0 \leq p_{ii}p_{jj} - p_{ji}^2 = -p_{ji}^2,$$

i.e. all other elements of the i^{th} row and column must vanish. Thus

$$\mathbf{V}'\mathbf{Q}_1\mathbf{V} = \begin{pmatrix} \mathbf{G}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix}, \ \mathbf{V}'\mathbf{Q}_2\mathbf{V} = \begin{pmatrix} \mathbf{G}_2 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix},$$

and so $\mathbf{I}_r = \mathbf{G}_1 + \mathbf{G}_2$ and $\mathbf{G}_1^2 = \mathbf{G}_1$. This yields

$$\mathbf{G}_1 = \mathbf{G}_1\mathbf{I}_r = \mathbf{G}_1(\mathbf{G}_1 + \mathbf{G}_2) = \mathbf{G}_1 + \mathbf{G}_1\mathbf{G}_2;$$

thus $\mathbf{G}_1\mathbf{G}_2 = \mathbf{0}$ and so $\mathbf{Q}_1\mathbf{Q}_2 = \mathbf{0}$. Also $\mathbf{G}_2^2 = \mathbf{G}_2$ and so $\mathbf{Q}_2$ is idempotent.
We have shown that

$$\begin{aligned} \mathbf{Q} &= \mathbf{Q}_1 + \mathbf{Q}_2, \text{ all p.s.d., } \mathbf{Q}, \mathbf{Q}_1 \text{ idempotent} \\ &\Rightarrow \mathbf{Q}_1\mathbf{Q}_2 = \mathbf{0} \text{ and } \mathbf{Q}_2 \text{ is idempotent.} \end{aligned}$$

If $k = 3$ then write

$$\begin{aligned} \mathbf{Q} &= \mathbf{Q}_1 + \mathbf{Q}_2 + \mathbf{Q}_3 \\ &= \mathbf{Q}_1 + \mathbf{Q}_{23}, \end{aligned}$$

where $\mathbf{Q}, \mathbf{Q}_1$ are idempotent and $\mathbf{Q}_{23} = \mathbf{Q}_2 + \mathbf{Q}_3$ is p.s.d. As above, $\mathbf{Q}_1\mathbf{Q}_{23} = \mathbf{0}$ and so $\mathbf{Q}_{23}$ is idempotent. But '$\mathbf{Q}_{23} = \mathbf{Q}_2 + \mathbf{Q}_3$ and $\mathbf{Q}_{23}, \mathbf{Q}_2$ idempotent' implies (again as above) that $\mathbf{Q}_2\mathbf{Q}_3 = \mathbf{0}$ and $\mathbf{Q}_3$ is idempotent. Now all three are idempotent, so we can relabel them and show in the same way that also $\mathbf{Q}_1\mathbf{Q}_2 = \mathbf{Q}_1\mathbf{Q}_3 = \mathbf{0}$. Continue ... . □

- **Proof of (the remaining half of) Cochran's Theorem.**
  If $\mathbf{x}'\mathbf{x} = \sum_{i=1}^{k} \mathbf{x}'\mathbf{Q}_i\mathbf{x}$ then $\sum_{i=1}^{k} \mathbf{Q}_i = \mathbf{I}_n$. Suppose that $\sum_{i=1}^{k} r_i = n$. Define $\mathbf{P}_i = \mathbf{I}_n - \mathbf{Q}_i$, and let $rk(\mathbf{P}_i) = p_i$. Using that the rank of a sum is bounded above by the sum of ranks, we get

$$\begin{aligned} \mathbf{P}_i &= \sum_{j \neq i} \mathbf{Q}_j \Rightarrow p_i \leq \sum_{j \neq i} r_j = n - r_i; \text{ also} \\ \mathbf{I}_n &= \mathbf{P}_i + \mathbf{Q}_i \Rightarrow n \leq p_i + r_i. \end{aligned}$$

  Thus $p_i = n - r_i$ and so $\mathbf{P}_i$ has r_i eigenvalues $= 0$. It follows that $\mathbf{Q}_i$ has r_i eigenvalues $= 1$. Since this is the rank of $\mathbf{Q}_i$, the other eigenvalues must $= 0$, and so $\mathbf{Q}_i$ is idempotent. Thus

$$\sum_{i=1}^{k} \mathbf{Q}_i = \mathbf{I}_n \text{ and all } \mathbf{Q}_i \text{ are idempotent.}$$

  These are the conditions of Lemma 3 (with $\mathbf{Q} = \mathbf{I}_n$) and so these matrices are mutually orthogonal as well. By Lemmas 1 and 2 the r.v.s $\mathbf{x}'\mathbf{Q}_i\mathbf{x}$ are independently distributed as $\sigma^2\chi^2_{r_i}$, $i = 1, ..., k$.
  □

Part III

MORE THAN ONE FACTOR

6. Randomized blocks

- **Paired comparisons** – read §3.1. **Outline**: Experimental units arise as matched pairs (twins, left and right feet, etc.). View as blocks of size 2 (homogeneity within blocks), with each member of a pair receiving one of the two treatments. If y_{i1} and y_{i2} are the responses within the i^{th} block then the analysis is based on $d_i = y_{i1} - y_{i2}$. One sample t-test. If the blocking were ignored one would carry out a two-sample t-test, and the power would typically be lower – in this example the mse must now account for the differences between the twins, and so would tend to be larger.

- **Randomized block designs** – continues the paired comparison scenario, but now there are more treatments. Basic idea continues to be that units within the same block should be more homogeneous than units between blocks. Each of the k treatments appears once in each of the b blocks ('complete'), and are randomly assigned to units within the block ('randomized' blocks) – hence CRBD.

– Example: Four methods are used to predict the strength of steel girders. Each method is applied on each of nine girders. So $k = 4$, $b = 9$ and within a block the same girder is being used (and the methods are applied in random order), thus reducing the heterogeneity.

– Model: For the j^{th} treatment in the i^{th} block

$$\begin{aligned} y_{ij} &= \eta + \alpha_i + \tau_j + \varepsilon_{ij} \\ &= \text{overall mean} + \text{block effect} \\ &\quad + \text{treatment effect} + \text{error}, \end{aligned}$$

with the constraints $\sum \alpha_i = \sum \tau_j = 0$.

– Decomposition of SS:

$$\begin{aligned} y_{ij} &= \bar{y}_{\cdot\cdot} + (\bar{y}_{i\cdot} - \bar{y}_{\cdot\cdot}) + \left(\bar{y}_{\cdot j} - \bar{y}_{\cdot\cdot}\right) \\ &\quad + \left(y_{ij} - \bar{y}_{i\cdot} - \bar{y}_{\cdot j} + \bar{y}_{\cdot\cdot}\right) \\ &= \hat{\eta} + \hat{\alpha}_i + \hat{\tau}_j + e_{ij}. \end{aligned}$$

These are the lse's (see Stat 368 notes if necessary) and the sums of squares are orthogonal, in that all cross-products vanish when

summed:

$$\begin{aligned}\sum_{i,j}\left(y_{ij}-\bar{y}_{\cdot\cdot}\right)^2 &= \sum_{i,j}\hat{\alpha}_i^2+\sum_{i,j}\hat{\tau}_j^2+\sum_{i,j}e_{ij}^2\\ &= k\sum_i\hat{\alpha}_i^2+b\sum_j\hat{\tau}_j^2+\sum_{i,j}e_{ij}^2;\\ SST &= SSBl+SSTr+SSE.\end{aligned}$$

Expressing these as quadratic forms and determining the ranks (= df) gives the following ANOVA table:

ANOVA table

Source	df	SS	$E[MS]$
Block	$b-1$	$SSBl$	$\sigma^2+\frac{1}{b-1}k\sum_i\alpha_i^2$
Treatment	$k-1$	$SSTr$	$\sigma^2+\frac{1}{k-1}b\sum_j\tau_j^2$
Error	$\substack{(b-1)\cdot\\(k-1)}$	SSE	σ^2
Total	$bk-1$	SST	

Example of calculation of an $E[MS]$: ...

– If the hypothesis of no treatment effects is rejected ($F=MSTr/MSE\sim F_{dfSSE}^{dfSSTr}$ under the null hypothesis that all $\tau_j=0$) then one can go on to compare pairs of treatments. To

compare treatments i and j, assuming i.i.d. $N\left(0, \sigma^2\right)$ errors:

$$\bar{y}_{\cdot j} - \bar{y}_{\cdot i} \sim N\left(\tau_j - \tau_i, \frac{\sigma^2}{b} + \frac{\sigma^2}{b}\right),$$

leading to

$$t_{ij} = \frac{\left(\bar{y}_{\cdot j} - \bar{y}_{\cdot i}\right) - \left(\tau_j - \tau_i\right)}{\sqrt{MSE\left(\frac{1}{b} + \frac{1}{b}\right)}} \sim t_{dfSSE.}$$

This is used to test that $\tau_j = \tau_i$, or to get CIs. Tukey's method applies exactly as before, leading to simultaneous CIs

$$\bar{y}_{\cdot j} - \bar{y}_{\cdot i} \pm q_{k,dfSSE,\alpha}\sqrt{\frac{MSE}{b}}.$$

- **Two-way layout with fixed factor levels** – rather than blocks $\times$ treatments we might plan an experiment with two factors, each at several levels. Again a 'two-way' layout, but now the notion of 'homogeneity within blocks' no longer pertains. We make $n > 1$ observations at each of I levels of factor A and J levels of Factor B.

– Within each of the n replicates the IJ experimental units are randomly assigned to the IJ treatments.

- Model:

$$\begin{aligned} y_{ijl} &= \eta + \alpha_i + \beta_j + \omega_{ij} + \varepsilon_{ijl} \\ &= \text{overall mean} + \text{effect of level } i, \text{ Factor } A \\ &\quad + \text{effect of level } j, \text{ Factor } B \\ &\quad + \text{interaction} + \text{error}, \end{aligned}$$

with the constraints that all effects vanish when summed over either index. The lse's are obtained as before:

$$\begin{aligned} \hat{\alpha}_i &= \bar{y}_{i\cdot\cdot} - \bar{y}_{\cdot\cdot\cdot}, \\ \hat{\beta}_j &= \bar{y}_{\cdot j\cdot} - \bar{y}_{\cdot\cdot\cdot}, \\ \hat{\omega}_{ij} &= \left(\bar{y}_{ij\cdot} - \bar{y}_{\cdot\cdot\cdot}\right) - \left(\bar{y}_{i\cdot\cdot} - \bar{y}_{\cdot\cdot\cdot}\right) - \left(\bar{y}_{\cdot j\cdot} - \bar{y}_{\cdot\cdot\cdot}\right) \\ &= \bar{y}_{ij\cdot} - \bar{y}_{i\cdot\cdot} - \bar{y}_{\cdot j\cdot} + \bar{y}_{\cdot\cdot\cdot}, \\ e_{ijl} &= y_{ijl} - \bar{y}_{ij\cdot}\ . \end{aligned}$$

ANOVA table

Source	df	SS	$E[MS]$
A	$I-1$	$SSA = nJ\sum_i \hat{\alpha}_i^2$	$\sigma^2 + \frac{nJ\sum_i \alpha_i^2}{df\,A}$
B	$J-1$	$SSB = nI\sum_j \hat{\beta}_j^2$	$\sigma^2 + \frac{nI\sum_j \beta_j^2}{df\,B}$
A×B	$\substack{(I-1)\cdot \\ (J-1)}$	$SSAB = n\sum_{i,j} \hat{\omega}_{ij}^2$	$\sigma^2 + \frac{n\sum_{i,j} \omega_{ij}^2}{df\,AB}$
Error	$\substack{IJ\cdot \\ (n-1)}$	$SSE = \sum_{i,j,l} e_{ijl}^2$	σ^2
Total	$IJn-1$	SST	

- The hypothesis that all $\omega_{ij} = 0$ is tested by $F = MSAB/MSE$; if accepted then the response is additive and one can go on to test the equality of the effects of factor A ($F = MSA/MSE$) or of factor B ($F = MSB/MSE$), and make inferences. In the presence of significant interactions these latter hypotheses are not very meaningful.

- **Regression approach** – define indicators

  $$U_i = I\left(\text{Factor A set at level } i+1\right), i = 1, ..., I-1,$$
  $$V_j = I\left(\text{Factor B set at level } j+1\right), j = 1, ..., J-1;$$

  fit a regression model with independent variables $\{U_i\}$, $\left\{V_j\right\}$, $\left\{U_i V_j\right\}$:

  $$Y = \eta + \sum_i \beta_i U_i + \sum_j \gamma_j V_j + \sum_{i,j} (\beta\gamma)_{ij} U_i V_j + \varepsilon.$$

  Then the effects $\alpha_i, \beta_j, \omega_{ij}$ can be expressed in terms of the regression parameters. For instance the hypothesis of no interactions becomes the hypothesis that all $(\beta\gamma)_{ij} = 0$. If accepted one can go on to test the equality of the effects of factor A: all $\beta_i = 0$, etc.

- If 'n' varies from cell to cell then the various sums of squares are no longer additive, and a regression approach is necessary – the ANOVA table above is no longer correct.

- If $n = 1$ then interactions cannot be estimated – they must be assumed to vanish, and then $SSAB$ is used in place of SSE. There is a 'one-degree of freedom for non-additivity' test, in which one fits the model

$$y_{ij} = \eta + \alpha_i + \beta_j + \omega\alpha_i\beta_j + \varepsilon_{ij}$$

and tests that $\omega = 0$ (Stat 368, lecture 16).

- Read §3.5 on three-way and higher layouts.

7. Latin & Graeco-Latin squares, BIBD

- R example of 2-factor layout – data in Table 3.8. Response is work required to tighten a locknut; the units (locknuts) are subject to factors $A =$ test medium, at $I = 2$ levels, and $B =$ type of plating process, at $J = 3$ levels. There are $n = 10$ replicates. Industry standard is $Y \leq 45$ (so look for treatments with $\hat{\mu}_{ij} + 2\hat{\sigma}_{ij} \leq 45$?).

  - Transformation needed to validate the normality assumption?

  - Testing equality of variances when there are replicates – the methods mentioned in the text are all quite non-robust against non-normality. A robust alternative is Levene's test. First calculate absolute deviations from the group medians $\tilde{y}_i$:

    $$d_{ij} = \left| y_{ij} - \tilde{y}_i \right| .$$

    Equality of the variances is indicated by equality of the group means of the absolute deviations, which is tested by the usual F-test. An R function is available.

- **Latin squares** – two blocking variables. Example: one tests abrasion resistance of four types of material (the treatments: A,B,C,D). The testing equipment accommodates the 4 materials in 4 positions, but there might be differences between the positions. In each group of 4 applications there might be an 'application' effect (arising from the setup of the equipment). A $k \times k$ ($k = 4$) Latin square layout:

	Position			
App.	1	2	3	4
1	C	D	B	A
2	A	B	D	C
3	D	C	A	B
4	B	A	C	D

Each material is tested in each position, and in each application (the two blocking variables).

  - There are tables of Latin squares (and R functions) available, from which one can choose a design at random.

- Linear model: if i runs over one blocking variable ('rows') and j over the other ('columns), and l denotes the treatment, then

$$E\left[y_{ijl}\right] = \eta + \alpha_i + \beta_j + \tau_l + \varepsilon_{ijl},$$

and

$$\begin{aligned} y_{ijl} &= \bar{y}... + (\bar{y}_{i}.. - \bar{y}...) + \left(\bar{y}_{.j}. - \bar{y}...\right) + (\bar{y}_{..l} - \bar{y}...) \\ &\quad + \text{ by subtraction} \\ &= \hat{\eta} + \hat{\alpha}_i + \hat{\beta}_j + \hat{\tau}_l + e_{ijl} \end{aligned}$$

  - Anova: There are only k^2 observations, so with $\sum'_{i,j,l}$ being shorthand for the sum over all k^2 cells:

$$SS_{rows} = \sum{}'_{i,j,l} \hat{\alpha}_i^2 = k \sum_{i=1}^{k} \hat{\alpha}_i^2, \text{ etc.}$$

ANOVA table

Source	df	SS
Rows	$k-1$	$k\sum_{i=1}^{k}\hat{\alpha}_i^2$
Columns	$k-1$	$k\sum_{j=1}^{k}\hat{\beta}_j^2$
Treatments	$k-1$	$k\sum_{l=1}^{k}\hat{\tau}_l^2$
Error	by subt'n	$\sum{}'_{i,j,l} e_{ijl}^2$
Total	k^2-1	

- **Graeco-Latin squares** – three blocking variables; superimpose a Greek square (a third blocking variable – day?) on a Latin square in such a way that each Greek-Latin pair appears once:

	Position			
App.	1	2	3	4
1	Cα	Dβ	Bγ	Aδ
2	Aβ	Bα	Dδ	Cγ
3	Dγ	Cδ	Aα	Bβ
4	Bδ	Aγ	Cβ	Dα

  - Anova is similar

  - These do not exist for all values of k. Sometimes one can even add another blocking variable ('Hyper-Graeco-Latin squares') by superimposing a third Latin square which is orthogonal to the other two (so that each triple of letters appears exactly once).

- **Balanced Incomplete Block Designs (BIBD)**: We compare t treatments in b blocks of size $k < t$ ('incomplete'); each treatment is replicated r times. The design is 'balanced' in that each pair of treatment appears together in the same number (λ) of blocks.

- **Example**: Each of 4 tires is divided into three sections and each section is treated with one of 4 compounds (so a tire cannot accommodate all 4 treatments). Then they are road tested to determine which compound results in the least wear. The tires are the blocks, and so $b = t = 4$, $k = 3$. The design, with $r = 3$, $\lambda = 2$ is:

	Compound			
Tire	A	B	C	D
1	x	x	x	
2	x	x		x
3	x		x	x
4		x	x	x

- The parameters b, t, r, k, λ are positive integers satisfying:

$$\begin{aligned} bk &= rt, \\ r\,(k-1) &= \lambda\,(t-1)\,. \end{aligned}$$

The first of these is clear (is it?); the second follows from the two ways of enumerating the number of units that appear with (say) treatment 1 in the same block. Some consequences: as above $t > k$, so $b > r > \lambda$:

$$\begin{aligned} r &= \frac{\lambda\,(t-1)}{(k-1)} > \lambda \text{ and} \\ b &= \frac{rt}{k} > r; \end{aligned}$$

also $rk = \lambda t + (r - \lambda) > \lambda t$. Thus BIBDs don't exist for all values of the parameters.

8. BIBD continued; ANCOVA; Split plots

- As for the RCB design the model is ($i = 1, ..., b$, $j = 1, ..., t$)

$$y_{ij} = \eta + \alpha_i + \tau_j + \varepsilon_{ij}$$

and the total sum of squares is $\sum' \left(y_{ij} - \bar{y}_{\cdot\cdot}\right)^2$, where $\sum'$ denotes summation over the bk occupied cells in the $b \times t$ table, and $\bar{y}_{\cdot\cdot} = \sum' y_{ij}/bk$. Define

$$n_{ij} = I\,(\text{treatment } j \text{ appears in block } i)\,;$$

note that $\sum_{i=1}^{b} n_{ij} = r$ and $\sum_{j=1}^{t} n_{ij} = k$. Then the total ss is

$$\sum_{j=1}^{t} \sum_{i=1}^{b} n_{ij} \left(y_{ij} - \bar{y}_{\cdot\cdot}\right)^2$$

and minimizing this under the zero-sum constraints on $\{\alpha_i\}$ and $\left\{\tau_j\right\}$ gives the LSEs

$$\begin{aligned} \hat{\eta} &= \bar{y}_{\cdot\cdot} \\ \hat{\alpha}_i &= \bar{y}_{i\cdot} - \bar{y}_{\cdot\cdot}, \\ \hat{\tau}_j &= \frac{k}{\lambda t} Q_j, \end{aligned}$$

where

$$Q_j = r\bar{y}_{\cdot j} - \sum_{i=1}^{b} n_{ij}\bar{y}_{i\cdot},$$

the 'adjusted treatment total' for the j^{th} treatment. Details are in the book by Peter Johns, on the course web site. If $k = t$ (completeness) then $r = b = \lambda$ and $\hat{\tau}_j = \bar{y}_{\cdot j} - \bar{y}_{\cdot\cdot}$.

Note:

(i) $\bar{y}_{\cdot j} = \sum_{i=1}^{b} n_{ij}y_{ij} / \sum_{i=1}^{b} n_{ij} = \sum_{i=1}^{b} n_{ij}y_{ij}/r$; similarly

(ii) $\bar{y}_{i\cdot} = \sum_{j=1}^{t} n_{ij}y_{ij}/k$; thus

$$r\sum_{j=1}^{t} \bar{y}_{\cdot j} = k\sum_{i=1}^{b} \bar{y}_{i\cdot} = \sum_{j=1}^{t}\sum_{i=1}^{b} n_{ij}y_{ij} = bk\bar{y}_{\cdot\cdot}.$$

From this, $\sum_{i=1}^{b} \hat{\alpha}_i = 0$; also $\sum_{j=1}^{t} Q_j = 0$, so that $\sum_{j=1}^{t} \hat{\tau}_j = 0$. The incompleteness complicates the anova since the estimates of the effects are no longer orthogonal – the estimates of treatment effects change if block effects are not estimated. The SS due to treatments is the change

in the SS under the hypothesis that all $\tau_j = 0$:

ANOVA table

Source	df	SS
Blocks	$b-1$	$k\sum_{i=1}^{b}\hat{\alpha}_i^2$
Treatments	$t-1$	$\frac{k}{\lambda t}\sum_{j=1}^{t}Q_j^2$
Error	by subt'n	by subt'n
Total	$bk-1$	

- The order in which the factors are entered is important, since they are not orthogonal. In R, use `aov(y ~blocks + treatments)` so that the effect of the blocking will be accounted for first. The remaining SS is that of treatments 'adjusted for blocks'.

- The balance ensures that all differences $\hat{\tau}_j - \hat{\tau}_i$ are equally varied (var $= \frac{2k}{\lambda t}\sigma_\varepsilon^2$). Simultaneous CIs based on Tukey's method (see reference in text) are

$$\hat{\tau}_j - \hat{\tau}_i \pm \frac{q_{t,dfSSE}}{\sqrt{2}}\hat{\sigma}\sqrt{\frac{2k}{\lambda t}}.$$

- Analysis of covariance (ANCOVA) – In addition to a treatment classification, the values of an auxiliary covariate might be known. Example: breaking strength (Y) of three types (CA, CO, PO) of starch film is analyzed; the thickness (x) of the film is also observed and possibly significant. Each of the $I = 3$ films is tested at n_i levels of x:

$$y_{ij} = \eta + \tau_i + \gamma x_{ij} + \varepsilon_{ij}.$$

  Generally done by fitting a regression model; one is then interested in comparing the τ_i. R code on web site gives output

```
             Estimate Std. Error t value Pr(>|t|)
(Intercept)   158.26     179.78   0.880 0.383360
thick          62.50      17.06   3.664 0.000653***
typeCO        -83.67      86.10  -0.972 0.336351
typePO         70.36      67.78   1.038 0.304795
```

- Interpretation (note baseline constraint):

$$\begin{aligned} CA &: \hat{y} = 158.26 + 62.50x, \\ CO &: \hat{y} = (158.26 - 83.67) + 62.50x, \\ PO &: \hat{y} = (158.26 + 70.36) + 62.50x. \end{aligned}$$

  The model implies that the regression lines have the same slope; one can instead fit separate regressions (different intercepts, different slopes) and first test for the equality of the slopes.

- Split plot designs: see Stat 368 lecture 30. A review: In an agricultural experiment, $I = 3$ fields are planted, with a different crop in each field (the 'whole plots'; 'crop' is the 'whole plot treatment'). Each field is then divided into $J = 4$ 'subplots', or 'split-plots', and each subplot is treated with a different fertilizer (the 'subplot treatment'). Then the whole experiment is replicated $n = 3$ times. Sounds like a two factor, replicated CRD, but *the randomization is only within each whole plot*.

- The hard to change factor (crops) is the whole plot treatment, and the main factor of interest (fertilizer) is the subplot treatment. Note that the whole plot treatments (crops) are confounded with the whole plots (fields) – if there is a systematic difference between the fields (soil quality?) it will show up as a difference between the crops.

- If this experiment is replicated n times, then we view the replicates as blocks. The simplest model, and one that can be easily fitted on R, includes (random) effects of $\{$replicates $+$ replicates $\times$ whole plot interactions$\}$, and (fixed) effects of $\{$wholeplot and subplot treatments and their interaction$\}$:

$$\begin{aligned} y_{ijk} &= \eta + \underbrace{\tau_k + (\tau\alpha)_{ki}}_{\text{replicate effects}} \\ &\quad + \underbrace{\alpha_i + \beta_j + (\alpha\beta)_{ij}}_{\text{treatment effects}} + \varepsilon_{kij} \\ i &= 1, ..., I, \ j = 1, ..., J, \ k = 1, ..., n. \end{aligned}$$

These random effects are viewed as $N\left(0, \sigma^2_\tau\right)$ and $N\left(0, \sigma^2_{\tau\alpha}\right)$ respectively.

- If we call these effects A (whole plot treatments, with effects α_i), B (subplot treatments, with effects β_j) and AB (interactions, with effects $(\alpha\beta)_{ij}$) then the sums of squares are computed by fitting a linear model aov$(y \sim (R + RA) + (A + B)^2)$. Note $R \times B$ and $R \times A \times B$ are not fitted, these form the 'subplot error' SS_{Sub} and used to estimate σ^2_ε. Apart from the names used, these are the SS for an unreplicated 3-way model, as in §3.5.

Source	MS=SS/df	E(MS)
τ: Replicates (blocks)	$\frac{SS_{Rep}}{n-1}$	$\left[\begin{array}{c}\sigma^2_\varepsilon + J\sigma^2_{\tau\alpha} \\ +IJ\sigma^2_\tau\end{array}\right]$
α: Whole plot treatment	$\frac{SS_A}{I-1}$	$\left[\begin{array}{c}\sigma^2_\varepsilon + J\sigma^2_{\tau\alpha} \\ +\frac{nJ\sum\alpha_i^2}{I-1}\end{array}\right]$
$(\tau\alpha)$: RA inter'n	$\frac{SS_{Whole}}{(n-1)(I-1)}$	$\sigma^2_\varepsilon + J\sigma^2_{\tau\alpha}$
β: Subplot treatment	$\frac{SS_B}{J-1}$	$\sigma^2_\varepsilon + \frac{nI\sum\beta_j^2}{J-1}$
$(\alpha\beta)$: AB inter'n	$\frac{SS_{AB}}{(I-1)(J-1)}$	$\sigma^2_\varepsilon + \frac{n\sum_{i,j}(\alpha\beta)_{ij}^2}{(I-1)(J-1)}$
ε: Subplot error	$\frac{SS_{Sub}}{I(n-1)(J-1)}$	σ^2_ε

- Testing is carried out by examining the expected mean squares – for instance the hypothesis of no whole plot effect is tested by $F = MS_A/MS_{Whole}$.

- If one fits only the treatment effects – $(A+B)^2$ – then the residuals from this can be used to then fit the (random) replicate effects. The R command is `aov(y ~(A + B)^2 + Error(R/A))`.

  - Example: Two panels of wood are each given a 'pretreatment' (A, at 2 levels), then split into 4 sections and a stain (B, at 4 levels) is applied to each section. This is replicated three times. Data in Table 3.29; R code on web site.

- Transformations – read §3.11; see also Lecture 20 of Stat 312 (Delta method).

Part IV

FULL FACTORIALS

9. Full factorials at two levels

- A 2^k factorial design has runs at all combinations of levels (each 'HI' or 'LO', $+/-$, etc.) of each of k factors – thus 2^k runs in one replicate.

- **Epitaxial layer growth experiment** – description in text; response is $Y =$ thickness of an 'epitaxial layer' on an integrated circuit device. Vapors are sprayed on silicon wafers being spun in a jar. There are $k = 4$ factors: $A =$ rotation methods, $B =$ nozzle position, $C =$ temperature, $D =$ length of time allowed for the depositing of particles on the wafer. Goal is to find combinations of factors achieving $Y = 14.5 \pm .5\mu m$. Factors C and D are quantitative. The 16 runs are carried out in random order (lurking variables – for instance room humidity might change over the course of the day and adversely affect runs carried out in the afternoon. Why does this necessitate randomization?).

- A variable might be hard to change (furnace temperature?); when this is the case a split-plot design might be employed, with the hard-to-change factor used as the whole plot factor.

- A 2^k factorial design is 'balanced' – each factor level appears in the same number (2^{k-1}) of runs, and 'orthogonal' – all level combinations appear in the same number of runs.

- Representation via model matrix (one replicate here; this display becomes the X-matrix – with orthogonal columns – in a regression representation, upon adding a column of $+$'s):

1	2	3	12	13	23	123	
A	B	C	AB	AC	BC	ABC	run #:
$-$	$-$	$-$	$+$	$+$	$+$	$-$	5
$-$	$-$	$+$	$+$	$-$	$-$	$+$	1
$-$	$+$	$-$	$-$	$+$	$-$	$+$	8
$-$	$+$	$+$	$-$	$-$	$+$	$-$	7
$+$	$-$	$-$	$-$	$-$	$+$	$+$	2
$+$	$-$	$+$	$-$	$+$	$-$	$-$	4
$+$	$+$	$-$	$+$	$-$	$-$	$-$	6
$+$	$+$	$+$	$+$	$+$	$+$	$+$	3

- Let $\bar{y}(A+)$ denote the average of all $m2^{k-1}$ of all runs (assuming m replicates) with A at the high level, ... The main effect of A is $ME(A) = \bar{y}(A+) - \bar{y}(A-)$, etc. The 'conditional main effect of B when A is HI' is $ME(B|A+) = \bar{y}(B+|A+) - \bar{y}(B-|A+)$, etc. Then

$$INT(A,B) \stackrel{def}{=} \frac{1}{2}[ME(A|B+) - ME(A|B-)]$$
$$= \frac{1}{2}[ME(B|A+) - ME(B|A-)],$$

the effect of changing A with B HI, minus the effect of changing A with B LO. Or the effect of changing B with A HI, minus the effect of changing B with A LO: $INT(A,B) = INT(B,A)$. Main effect plots and interaction plots are valuable displays; see R example for the epitaxial experiment.

- If $ME(B|A+)$ and $ME(B|A-)$ have the same sign we say the interaction plot is 'synergistic' (changing B when A is HI results in the same

type of effect as changing B when A is LO), otherwise 'antagonistic'. In the epitaxial experiment the C against D plot is synergistic (thickness Y increases with time D regardless of the other factors such as temperature C); the D against C plot is antagonistic.

- Three-factor (and higher order) interactions are defined recursively:

$$INT(A,B,C) = \frac{1}{2}INT\,(A,B|C+) - \frac{1}{2}INT(A,B|C-)$$

where

$$\begin{aligned} & INT\,(A,B|C+) \\ = \; & \frac{1}{2}\left[ME\,(A|B+,C+) - ME\,(A|B-,C+)\right], \end{aligned}$$

the two-factor interaction computed from that half of the data with C HI. Then

$$\begin{aligned} & INT(A,B,C,D) \\ = \; & \frac{1}{2}INT\,(A,B,C|D+) - \frac{1}{2}INT(A,B,C|D-), \end{aligned}$$

and so on. It follows that this is independent of the labelling: $INT(A,B,C) = INT(B,A,C) = INT(B,C,A)$ etc.

- The $2^k - 1$ main effects and interaction effects are the 'factorial effects'; if θ is any such effect then $\hat{\theta} = \bar{y}_+ - \bar{y}_-$, where $\bar{y}_+$ is the average of all y_{ij} at which the product of all levels (denoted $\pm$) of the factors is $+$, etc. For instance

$$\begin{aligned} & INT\,(A,B) \\ = \; & \frac{\bar{y}\,[(A,B)=(+,+)] + \bar{y}\,[(A,B)=(-,-)]}{2} \\ & -\frac{\bar{y}\,[(A,B)=(+,-)] + \bar{y}\,[(A,B)=(-,+)]}{2}. \end{aligned}$$

The general case is assigned; it justifies the arrangement of the signs in the model matrix above. It also follows that, if there are m replicates, so $N = m2^k$ runs, then

$$\begin{aligned} \text{var}\left[\hat{\theta}\right] &= \text{var}\,[\bar{y}_+] + \text{var}\,[\bar{y}_-] \\ &= \frac{\sigma_\varepsilon^2}{N/2} + \frac{\sigma_\varepsilon^2}{N/2} = \frac{4\sigma_\varepsilon^2}{N}. \end{aligned}$$

- Regression model – these effects can be obtained directly by fitting a regression model – read §4.4; the derivation of the LSEs is assigned. It follows from this that the regression estimates are half the factorial effects. For instance using the model matrix as above,

$$\hat{\beta}_1 = \frac{1}{2}\bar{y}(A+) - \frac{1}{2}\bar{y}(A-).$$

With $x_A, x_B, x_C = \pm 1$ the underlying regression model is

$$\begin{aligned} Y &= \beta_0 + \beta_1 x_A + \beta_2 x_B + \beta_3 x_C + \beta_{12} x_A x_B \\ &\quad + \beta_{13} x_A x_C + \beta_{23} x_B x_C + \beta_{123} x_A x_B x_C + \varepsilon, \end{aligned}$$

with

$$\begin{aligned} & E\left[Y | x_A = \pm 1\right] \\ = \; & E_{x_B, x_C} E\left[Y | x_A = \pm 1, x_B, x_C\right] \\ = \; & E_{x_B, x_C} \left[\begin{array}{l} \beta_0 \pm \beta_1 + \beta_2 x_B + \beta_3 x_C \pm \beta_{12} x_B \\ \pm \beta_{13} x_C + \beta_{23} x_B x_C \pm \beta_{123} x_B x_C \end{array} \right] \\ = \; & \beta_0 \pm \beta_1 \end{aligned}$$

and so

$$\beta_1 = \frac{1}{2} E\left[Y | x_A = 1\right] - \frac{1}{2} E\left[Y | x_A = -1\right].$$

- Some guiding principles:

  - Effect Hierarchy Principle:

    (i) Lower order effects are more likely to be important than higher order effects.

    (ii) Effects of the same order are equally likely to be important.

  - Effect Sparsity principle (Box-Meyer)

    The number of relatively important effects in a factorial experiment is small.

  - Effect Heredity Principle (Hamada-Wu)

    In order for an interaction to be significant, at least one of its parent factors should be significant.

- 'One-factor at a time' approach – read. Roughly, the procedure is to investigate what one feels is the most important factor, with the others fixed at a particular set of levels. Decide on the best

level of that factor, fix it there and move on to study another factor. Requires many more runs to achieve the same precision.

- **Half-normal plots**: Graphical way to judge significance of effects. Plot the ordered absolute values $\left|\hat{\theta}\right|_{(i)}$ of the effects against the quantiles $F^{-1}\left((i-.5)/\left(2^k-1\right)\right)$, where F is the df of $|Z|$ (the positive normal distribution). It follows that

$$F^{-1}\left(\frac{i-.5}{2^k-1}\right)=\Phi^{-1}\left(\frac{1}{2}\left(\frac{i-.5}{2^k-1}+1\right)\right).$$

  Significantly large absolute effects show up as exceeding their (large) values of $F^{-1}\left(i/\left(2^k-1\right)\right)$. There is an R function available.

  - Not all $I=2^k-1$ effects need be compared; a smaller set might be plotted.

10. Full factorials – inferences

- With replicates, the anova table yields an estimate of the error variance, and F-ratios which can be used to judge the significance of individuals effects (under the assumption of homoscedasticity). There are other methods of testing significance, even in the absence of replicates. We have seen that the use of half-normal plots gives a graphical indication of significance. Another approach is 'Lenth's method'.

- **Lenth's method**: Suppose estimates $\left\{\hat{\theta}_i\right\}_{i=1}^{I}$ are i.i.d. $N\left(0, \sigma^2\right)$. A consistent estimate of σ is

$$s_0 = \frac{\text{med}\left\{\left|\hat{\theta}_i\right|\right\}_{i=1}^{I}}{\Phi^{-1}\left(.75\right))}.$$

Reason: Put $Z_i = \hat{\theta}_i/\sigma$. Let F be the distribution function of $|Z|$. Then

$$\sqrt{I}\left(\text{med}\left\{|Z_i|\right\} - F^{-1}\left(.5\right)\right)$$

is asymptotically normal with mean 0 (Stat 665 lecture 3); thus

$$\begin{aligned} s_0 - \sigma &= \left(\frac{\text{med}\left\{\left|\hat{\theta}_i\right|\right\}}{\sigma} - \Phi^{-1}(.75)) \right) \frac{\sigma}{\Phi^{-1}(.75))} \\ &= \left(\text{med}\left\{|Z_i|\right\} - F^{-1}(.5) \right) \frac{\sigma}{\Phi^{-1}(.75))}, \end{aligned}$$

and so

$$s_0 - \sigma \xrightarrow{pr} 0.$$

A robust, trimmed (about 1%) estimate is

$$\text{pse} = \frac{\text{med}\left\{\left|\hat{\theta}_i\right| \mid \left|\hat{\theta}_i\right| < 2.5 s_0\right\}}{\Phi^{-1}(.75)},$$

the 'pseudo standard error'. The studentized values

$$t_{\text{pse},i} = \frac{\hat{\theta}_i}{\text{pse}}$$

can be used to test significance when, as here, the effects $\hat{\theta}_i$ are equally varied. There is an R function available; it plots 'margin of error' for individual tests and 'simultaneous ME' for simultaneous tests.

- When there are replicates and, as here (the epitaxial layer growth experiment), there is no evidence of variance inequality at the 2^k different levels of the factors, then the variance estimates can be pooled:

$$s^2 = \sum_{i=1}^{2^k} w_i s_i^2,$$

where $w_i = (n_i - 1) \,/\, \sum_{i=1}^{2^k} (n_i - 1)$ ($= 2^{-k}$ if all levels are replicated the same number of times). This is the same as the mean square of the residuals, returned by anova.

- **Simultaneous comparisons**. If the factorial is replicated, then to obtain simultaneous intervals on the mean responses at differing levels of the factors, Tukey's HSD may be used. The same theory extends to simultaneous testing of effects. Consider testing the hypothesis that all of the effects $\{\theta_i\}_{i=1}^{I}$ are zero, by rejecting for large values of

$$R = \max_i \frac{\left|\hat{\theta}_i\right|}{\mathsf{se}\left[\hat{\theta}_i\right]}.$$

We suppose that under the hypothesis, we have $\hat{\theta}_i \sim N\left(0, \nu_i^2\sigma_\varepsilon^2\right)$. (In a factorial with N observations we have $\nu_i^2 = 4/N$, derived below.) Suppose also that the variance σ_ε^2 is estimated by $S^2 \sim \sigma_\varepsilon^2\chi_{df}^2/df$ (independently of the $\hat{\theta}_i$). Then $\mathsf{se}\left[\hat{\theta}_i\right] = \nu_i S$. Put

$$Z_i = \frac{\hat{\theta}_i}{\nu_i\sigma_\varepsilon} \overset{i.i.d.}{\sim} N(0,1).$$

Then

$$\begin{aligned} R &= \max_i \frac{\left|\hat{\theta}_i\right|}{\mathsf{se}\left[\hat{\theta}_i\right]} = \max_i \frac{\nu_i\sigma_\varepsilon\left|Z_i\right|}{\nu_i S} = \frac{\max_i\left|Z_i\right|}{\sqrt{S^2/\sigma_\varepsilon^2}} \\ &\sim \frac{\max_i\left|Z_i\right|}{\sqrt{\chi_{df}^2/df}} \overset{def}{=} M_{I,df}, \end{aligned}$$

the 'studentized maximum modulus' statistic. Tables are in the text. An effect is declared significant, with experimentwise error rate α, if

$$\left|\hat{\theta}_i\right| \Big/ \mathsf{se}\left[\hat{\theta}_i\right] > M_{I,df;\alpha}.$$

This is because then

$$\begin{aligned} & P_H\,(\text{at least one effect is declared significant}) \\ = \;& P\left(R > M_{I,df;\alpha}\right) = \alpha. \end{aligned}$$

- Extending a previous formula gives

$$\text{var}\left[\hat{\theta}_i\right] = \frac{\sigma_\varepsilon^2}{N_{i+}} + \frac{\sigma_\varepsilon^2}{N_{i-}} = \sigma_\varepsilon^2 \left(\frac{1}{N_{i+}} + \frac{1}{N_{i-}}\right),$$

where N_{i+} and N_{i-} are the numbers of observations at the $+$ and $-$ values of the product of terms in $\hat{\theta}_i$. Then in the above

$$\nu_i^2 = \frac{1}{N_{i+}} + \frac{1}{N_{i-}} = \frac{N}{N_{i+}N_{i-}}.$$

Generally $N_{i+} = N_{i-} = N/2$ and then $\nu_i^2 = 4/N$.

- In the epitaxial data $M_{15,80;.05} = 3.01$ with $M^2 = 9.06$. This can be compared with the F-values in the anova output, which are the squares of the $\left|\hat{\theta}_i\right| / \text{se}\left[\hat{\theta}_i\right]$. By this measure only D and CD are significant. For full factorials, $I = 2^k - 1$ can be very large and the test lacks power. But suppose we know (before looking at the data!) that the interactions of order higher than 2 are negligible, and we seek inferences only on the $I =$

$k+k(k-1)/2 = k(k+1)/2$ main effects and two-factor interactions. Then in the epitaxial example we use $M_{10,80;.05} = 2.87$, with $M^2 = 8.24$; still only D and CD are significant.

- **Nominal-the-best** analysis of epitaxial data. We aim for a response as close to 14.5 as possible. A plausible, two-step procedure for finding levels minimizing the mse $E\left[(Y-14.5)^2\right]$ is to first look for levels of factors minimizing the variance $E\left[(Y-\mu_Y)^2\right]$, and to then choose, among the factors which have not been set in the first step, levels which reduce the (squared) bias $(\mu_Y - 14.5)^2$. Here μ_Y is estimated by the fitted value.

- In this example the original epitaxial data are used; for these the variances depend very significantly on the factor levels (Levene's test).

- To analyze the dispersion, it is helpful to study $\log S^2$ since, by virtue of the delta method, $S^2 \sim$

$\left(\sigma^2, k\sigma^4\right)$ so that $\log S^2 \sim \left(\log \sigma^2, k\right)$, approximately. At each of the 16 combinations of levels, a value $\log S^2$ is computed and used as the response variable in an unreplicated 2^4 factorial. As well, the thickness values (Y) are used to estimate the 'location' effects, as above. These effects are estimated by regression on the ± 1 variables x_A, x_B,... .The half-normal plots identify A as the most important factor affecting the dispersion, with the variance being lowest at $A-$, and D as the most important factor affecting the thickness. Regressing $\log S^2$ on x_A gives the regression equation

$$\widehat{\log S^2} = -3.772 + 1.917 x_A;$$

this is minimized when $x_A = -1$ and then the estimated variance is $\hat{\sigma}^2 = \exp(-3.772 - 1.917) = .0034$. The regression of Y on x_D gives

$$\hat{Y} = 14.3889 + 0.4181 x_D;$$

now $\hat{Y} = 14.5$ when $x_D = .2657$. Assume (for reasons of confidentiality this is not known exactly) that $x_D = -1/+1$ correspond to $D =$

30/40 respectively; then by interpolation we obtain that $x_D = .2657$ corresponds to

$$D = 30 + \frac{1.2657}{2}(40 - 30) = 36.3286.$$

Note from Table 4.1 that the smallest values of S^2 all occur at $A-$, and are comparable to .0034. So the best settings to obtain the target value of $Y = 14.5 \pm .5\mu m$ (or less) appear to be A at the LO level, and D set at 36.33 (seconds).

- Minimize mse directly? Substitute the (full) fits for both $\hat{Y}$ and $\widehat{\log S^2}$ into

$$\begin{aligned} \text{mse} &= \left(\hat{Y} - 14.5\right)^2 + \exp\left(\widehat{\log S^2}\right) \\ &= \text{mse}\left(x_A, x_B, x_C, x_D\right) \end{aligned}$$

and minimize this (nonlinear) function over

$$\begin{aligned} x_A, x_B &\in \{-1, 1\}, \\ x_C, x_D &\in [-1, 1]. \end{aligned}$$

11. Full factorials – blocking

- One might wish to run a 2^k factorial in blocks – 2^q blocks, each of size 2^{k-q}. As usual, units within a block should be homogeneous. For instance there might be two blocks, corresponding to running the experiment over two days (with 'days' not interacting with any of the effects being modelled).

- Example: 2^3 design in 2 blocks I and II. The block effect estimate will be $\bar{y}(II) - \bar{y}(I)$, the difference in the averages of the observations in the two blocks. There are only 7 df, so only 6 other effects can be estimated. It is reasonable to sacrifice the highest order interaction. Label the factors 123 rather than ABC. Since the three factor interaction is estimated by $\bar{y}(123+) - \bar{y}(123-)$, we put the $123+$ runs in one block and the $123-$ runs in another block. We say $B = 123$ – blocks are 'confounded' with the 123 effect.

1	2	3	12	13	23	123	B
-	-	-	+	+	+	-	I
-	-	+	+	-	-	+	II
-	+	-	-	+	-	+	II
-	+	+	-	-	+	-	I
+	-	-	-	-	+	+	II
+	-	+	-	+	-	-	I
+	+	-	+	-	-	-	I
+	+	+	+	+	+	+	II

- To run the experiment in 4 blocks, we might choose blocking variables $B_1 = 12$ and $B_2 = 13$ (days and suppliers, for instance), and then the blocks are the 4 combinations of signs of B_1 and B_2:

1	2	3	B_1	B_2	23	123	B
-	-	-	+	+	+	-	IV
-	-	+	+	-	-	+	III
-	+	-	-	+	-	+	II
-	+	+	-	-	+	-	I
+	-	-	-	-	+	+	I
+	-	+	-	+	-	-	II
+	+	-	+	-	-	-	III
+	+	+	+	+	+	+	IV

- With 4 blocks there are 3 df for estimating block effects; the third effect is the interaction $B_1B_2 = 23$. We have $B_1B_2 = +$ in blocks IV and I, and $B_1B_2 = -$ in blocks III and II, so that a systematic difference between these pairs of blocks is confounded with the 23 interaction. (Just as B_1 and 12 have the same signs in blocks III/IV and I/II, and B_2 and 13 have the same signs in blocks II/IV and I/III.)

- Why not $B_1 = 12$ and $B_2 = 123$? Then $B_1B_2 = 12{*}123 = 1^2 2^2 3 = 3$ (any square is the product of two equal signs, hence is $+$: we say $1^2 = 2^2 = I$, the identity) and the corresponding block effect is confounded with a main effect. Similarly $B_1 = 13$ and $B_2 = 123$ results in the main effect 2 being confounded with the B_1B_2 interaction.

- There is an assumption that there are no block-treatment interactions. For instance the relationship $B_1 = 12$ implies that $B_1 1 = 2$ and $B_1 2 = 1$.

If there is a $B_1 1$ interaction it is confounded with 2, etc. To check the assumption one might plot the residuals separately from each block, looking for differences. (Example of a block-treatment interaction: blocks are separate blends of raw materials, one of which is contaminated in such a way as to negatively impact only one of the treatments.)

- In general, to run a 2^k factorial in 2^q blocks, each of size 2^{k-q}, denote the q blocking variables by $B_1, ..., B_q$ and choose q factorial effects $v_1, ..., v_q$ to be confounded with them:

$$B_1 = v_1, ..., B_q = v_q.$$

There are $\sum_{i=1}^{q} \binom{q}{i} = 2^q - 1$ possible products of the B's; these and I form a multiplicative group called the block 'defining contrast subgroup' (DCS).

- Example: 2^5 in $2^3 = 8$ blocks. Try the 'generators'

$$B_1 = 135, B_2 = 235, B_3 = 1234,$$

implying

$$\begin{aligned} B_1B_2 &= 12, B_1B_3 = 245, \\ B_2B_3 &= 145, B_1B_2B_3 = 34. \end{aligned}$$

There are no confounded main effects and two confounded two-factor interactions. For this scheme 'b' we say $g_1(b) = 0$, $g_2(b) = 2$, $g_3(b) = 4$, $g_4(b) = 1$. For any two schemes b_1 and b_2, if r is the smallest value with $g_r(b_1) \neq g_r(b_2)$, the scheme b_1 is preferable if $g_r(b_1) < g_r(b_2)$ (recall the effect hierarchy principle) – we say that b_1 has 'less aberration' than b_2. In this example another possible scheme has generators

$$B_1 = 12, B_2 = 13, B_3 = 45,$$

implying

$$\begin{aligned} B_1B_2 &= 23, B_1B_3 = 1245, \\ B_2B_3 &= 1345, B_1B_2B_3 = 2345, \end{aligned}$$

with $g_1(b) = 0$, $g_2(b) = 4$. It can be shown that the first scheme is a 'minimum aberration' design – no scheme has smaller aberration. Tables of minimum aberration designs are available

(and are in the text) for small values of k. Lots of interesting mathematics involved in the constructions.

- The 'order of estimability' of a blocking scheme is $e =$ one less than the lowest order interaction confounded with block effects ($=$ highest order not confounded). The schemes above have $e = 1$.

- Different sets of generators can generate the same blocking scheme. Those presented in the tables have generators with the longest 'words'. For instance if $k = 4$ and $q = 2$ these generators

$$(B_1, B_2) = (134, 234), (12, 134), (12, 234)$$

generate the same scheme. But the first, with longer words (on average) is preferable, based on the principle that 'main' block effects B_1 and B_2 might be more important than their products (interactions).

Part V

FRACTIONAL FACTORIALS

12. Fractional factorials – construction

- **Overview**: One might run only a fraction of a 2^k factorial, in order to save resources. As with blocking, this results in some effects being 'aliased' with others – they have the same estimate, and so the effect is estimable only if its aliases can be assumed to be negligible. If they cannot, there are 'fold-over' techniques for adding more runs in such a way as to resolve this difficulty. As well, fractional factorials can be blocked.

- **Leaf spring experiment**: five factors affect heat treatment process on the leaf springs of trucks. The response variable is the height of an unloaded spring, and the factors are $B =$ furnace temperature, $C =$ heating time, $D =$ transport time (from the furnace to the forming machine), $E =$ time spent in a press and $Q =$ quench oil temperature. A half-fraction is run, with each of the 2^4 runs being replicated three times.

- Were a full 2^5 factorial to be run, then the three- and higher order interactions would use up $\sum_{i=3}^{5} \binom{5}{i} = 10 + 5 + 1 = 16$ df. These interactions are often of little interest and can perhaps be assumed to vanish. Then the remaining 15 effects can (perhaps) be estimated in the half-fraction.

B	C	D	$E = BCD$	Q	$E = BCDQ$
-	+	+	-	-	+
+	+	+	+	-	-
-	-	+	+	-	-
+	-	+	-	-	+
-	-	-	+	-	+
+	-	-	-	-	-
-	+	-	-	-	-
+	+	-	+	-	+
-	+	+	-	+	-
+	+	+	+	+	+
-	-	+	+	+	+
+	-	+	-	+	-
-	-	-	+	+	-
+	-	-	-	+	+
-	+	-	-	+	+
+	+	-	+	+	-

- The design used in the leaf spring experiment is given by the first four columns on the last slide. Note $E = BCD$ – these effects are estimated by the same linear combination of the Ys, and so are indistinguishable from each other without further assumptions. The relationship $E = BCD$ is equivalent to $I = BCDE$, the 'defining relation' of this 2^{5-1} 'fractional factorial'. Since the 'word' in the defining relationship is of length 4, we call this a resolution IV design. From this defining relationship we obtain

$$\begin{aligned} B &= CDE, C = BDE, D = BCE, E = BCD, \\ BC &= DE, BD = CE, BE = CD, \\ Q &= BCDEQ, BQ = CDEQ, CQ = CDEQ, \\ DQ &= BCEQ, EQ = BCDQ, BCQ = DEQ, \\ BDQ &= CEQ, BEQ = CDQ \end{aligned} \tag{12.1}$$

  - A main effect or two-factor interaction is 'clear' if none of its aliases are main effects or two-factor interactions, and 'strongly clear' if as well none of its aliases are three factor interactions. Here, B, C, D, and E are clear, Q is strongly clear as are BQ, CQ, DQ, EQ.

- The design with $E = BCDQ$, i.e. $I = BCDEQ$ is resolution V. Each main effect is aliased with a four-factor interaction, and so is strongly clear. All ten two-factor interactions are clear but not strongly clear. Is this layout preferable? – depends on how important it is that BQ, CQ, DQ, EQ be strongly clear. How important are the two-factor interactions with Q? If not important, the second plan seems preferable.

- Note that the 2^{5-1} design can be constructed by writing down the 16 runs of a full 2^4 factorial, i.e. 4 columns with all their $\pm$ signs, and then calculating the fifth column by applying the defining relation. This extends to 2^{k-p} fractional factorials.

- In general, a 2^{k-p} fractional factorial is a 2^{-p} fraction of a full 2^k factorial.

– It is generated by equating p words to I (the defining relations); since then all products of these words also $= I$ there are $2^p - 1$ aliased effects which together with I form a multiplicative group – the (treatment) defining contrast subgroup (DCS).

– The number of letters in a word is the 'wordlength'. We always avoid words of length 1 or 2, since then main effects are aliased with each other. Denote by A_i the number of words of length i in the DCS. Then $W = (A_3, ..., A_k)$ is the 'wordlength pattern' of the design.

– The 'resolution' is the smallest r with $A_r > 0$, i.e. the length of the shortest word in the DCS. The resolution is generally denoted by Roman numerals III,IV, ... and a 2^{k-p} fractional factorial of resolution R is denoted by 2^{k-p}_R.

- In a resolution R design, no effect involving i factors is aliased with effects involving fewer than $R-i$ factors. (**Proof**: If

$$A_1 \cdots A_i = A_{i+1} \cdots A_{R-1}$$

then $I = A_1 \cdots A_i A_{i+1} \cdots A_{R-1}$, a contradiction.) Thus in a resolution IV design all main effects are clear. In a resolution V design the main effects are strongly clear and the two-factor interactions are clear.

- Generally, for given k and p a higher resolution is better.

- **Example**: A 2^{6-2}_{III}design. It is convenient now to denote factors by $1, 2, 3$ ... rather than A, B, C ... . Start with $I = 125 = 1346$, then also $= 23456$ ($p = 2$, $2^p - 1 = 3$). The DCS $= \{I, 125, 1346, 23456\}$. The wordlength pattern is $W = (1, 1, 1, 0)$ (so resolution III) and the aliased

effects are

$$
\begin{aligned}
I &= 125 = 1346 = 23456, \\
1 &= 25 = 346 = 123456, \\
2 &= 15 = 12346 = 3456, \\
3 &= 1235 = 146 = 2456, \text{(clear)} \\
4 &= 1245 = 136 = 2356, \text{(clear)} \\
5 &= 12 = 13456 = 2346, \\
6 &= 1256 = 134 = 2346, \text{(clear)} \\
13 &= 235 = 46 = 12456, \\
14 &= 245 = 36 = 12356, \\
16 &= 256 = 34 = 12345, \\
23 &= 135 = 1246 = 456, \text{(clear)} \\
&\vdots \\
56 &= 126 = 1345 = 234 \text{ (clear)}.
\end{aligned}
$$

- **Projection property** – a resolution R design can be 'projected' so that it is a full factorial in $R-1$ factors (assigned). For instance the 2^{3-1}_{III} design with generator $I = 123$ is obtained by writing out all 8 runs of a 2^3 factorial and then using only the 4 with $123 = +$ (if $I = -123$ use the others).

1	2	3	123
+	+	+	+
+	+	-	-
+	-	+	-
+	-	-	+
-	+	+	-
-	+	-	+
-	-	+	+
-	-	-	-

Then the half fraction with $123 = +$ is

1	2	3	123
+	+	+	+
+	-	-	+
-	+	-	+
-	-	+	+

and this is a full 2^2 factorial in any two of the factors.

- If two designs d_1 and d_2 have wordlength patterns

$$\begin{aligned} W(d_1) &= (A_3(d_1), ..., A_k(d_1)), \\ W(d_2) &= (A_3(d_2), ..., A_k(d_2)) \end{aligned}$$

let r be the smallest integer with $A_r(d_1) \neq A_r(d_2)$. We say d_1 has 'less aberration' than d_2 if $A_r(d_1) < A_r(d_2)$ (i.e. a lexicographic ordering of wordlength patterns), and is a 'minimum aberration' design if no design has less aberration. We typically seek a minimum aberration design, which always exists.

13. Fractional factorials – analysis

Leaf spring experiment. Recall Y = height of unloaded spring (inches); nominal ($Y = 8$) is best.

	Factor	$-$	$+$
$B = 1$	Temperature	1840	1880
$C = 2$	Heating time	23	25
$D = 3$	Transfer time	10	12
$E = 4$	Hold down time	2	3
$Q = 5$	Quench temp.	$130 - 150$	$150 - 170$

- 2^{5-1}_{IV} design – a half-fraction with $4 = 123$, i.e. $I = 1234 = BCDE$. All main effects clear.

- R reports the effects with the shortest words, these include the contributions of aliased effects. But also, e.g., $BCQ = DEQ$ (reported lexicographically?). Max modulus statistic is $|M_{15,32,.05}| = 3.14 = \sqrt{9.86}$ (or $|M_{12,32,.05}| = 3.06 = \sqrt{9.36}$ if three-factor interactions assumed negligible); by

this measure B, C, Q, CQ are significant. This is borne out by the half-normal plot. Note E and BQ are marginally significant.

```
Response: y
           Df  Sum Sq Mean Sq F value    Pr(>F)
xB          1 0.58742 0.58742 35.4846 1.230e-06 ***
xC          1 0.37277 0.37277 22.5181 4.160e-05 ***
xD          1 0.00992 0.00992  0.5992  0.444577
xE          1 0.12917 0.12917  7.8028  0.008735 **
xQ          1 0.80860 0.80860 48.8458 6.416e-08 ***
xB:xC       1 0.00350 0.00350  0.2116  0.648663
xB:xD       1 0.00460 0.00460  0.2780  0.601651
xB:xE       1 0.01505 0.01505  0.9093  0.347458
xB:xQ       1 0.08585 0.08585  5.1861  0.029592 *
xC:xQ       1 0.32835 0.32835 19.8350 9.645e-05 ***
xD:xQ       1 0.03467 0.03467  2.0943  0.157579
xE:xQ       1 0.00880 0.00880  0.5317  0.471191
xB:xC:xQ    1 0.00130 0.00130  0.0787  0.780932
xB:xD:xQ    1 0.01960 0.01960  1.1841  0.284647
xB:xE:xQ    1 0.02660 0.02660  1.6070  0.214065
Residuals  32 0.52973 0.01655
```

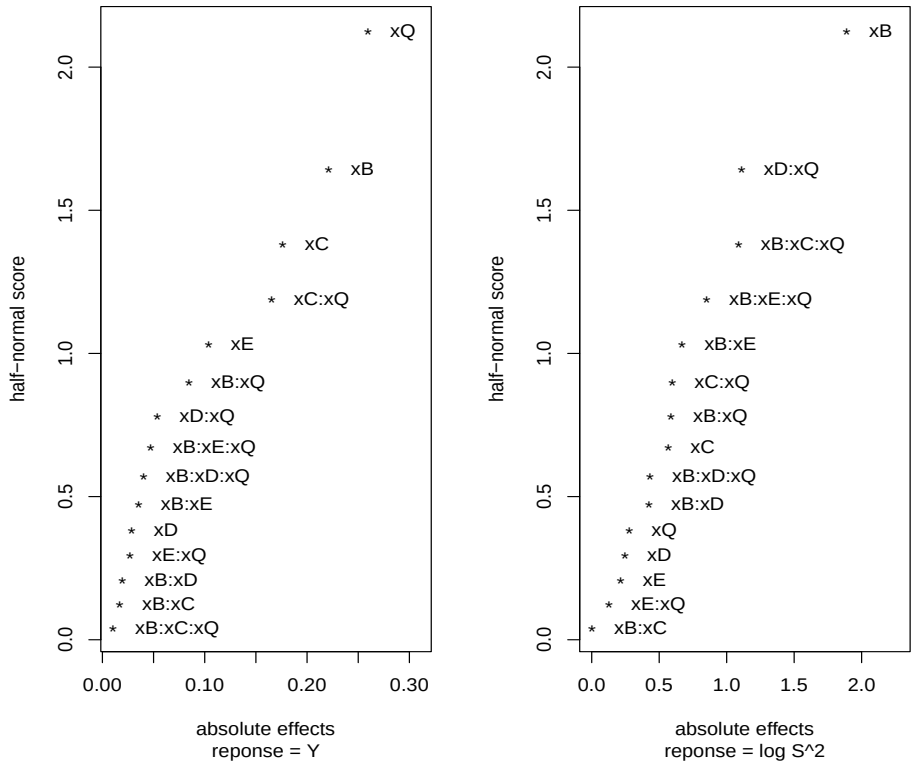

Half-normal plots of absolute location and dispersion effects

- Dispersion – B, possibly DQ, BCQ significant. Include Q as well (principle of hierarchy; also it makes the following analysis easier). Fitting

$$logs2 \sim xB + xD : xQ + xB : xC : xQ + xQ \tag{13.1}$$

and looking at model.tables(aov(fit),type="effects") gives the conclusion that the best settings are $(B, C, D, Q) = (-, +, +, -)$. At these levels $\hat{y} = 7.868 + 0.0519xE$, so for $y = 8$ we want $xE = 2.538$. By interpolation $E = 3.767$ – too far out of the range?

- Perhaps better to substitute the models

$$\hat{y} \sim xQ + xB + xC + xC * xQ + xE + xB * xQ,$$
$$logs2 \sim xB + xD * xQ + xB * xC * xQ$$

into

$$mse = (\hat{y} - 8)^2 + \exp(\ln s^2) \tag{13.2}$$

to find the levels ($x \in [-1, 1]$) minimizing the mean squared error (assigned).

- Recall that $BCQ = DEQ$ – the data from this experiment cannot distinguish between these effects. If neither can be assumed negligible then a

follow-up experiment is necessary. Two methods are available. The first is the 'fold-over' technique. Read §5.4.1. Here is a rough description of method, illustrated on a 2^{7-4}_{III} design d_1 with defining relations

$$4 = 12, 5 = 13, 6 = 23, 7 = 123.$$

The DCS has

$$\begin{aligned} I &= 124 = 135 = 236 = 1237 \\ &= 1346 = 2345 = 347 = 1256 = 257 = 167 \\ &= 3567 = 2467 = 1457 = 456 \\ &= 1234567; \end{aligned}$$

with 7 words of even length and 8 of odd length. None of the main effects or two-factor interactions are clear. Consider augmenting d_1 with the 2^{7-4}_{III} 'fold-over' design d_2 obtained by replacing the columns $1, ..., 7$ in d_1 by $-1, ..., -7$. For instance the eight runs in column 1 have signs $(-,-,-,-,+,+,+,+)$; they become $(+,+,+,+,-,-,-,-)$. Now the defining relations for d_2 are

$$4 = -12, 5 = -13, 6 = -23, 7 = 123,$$

and

$$\begin{aligned} I &= -124 = -135 = -236 = 1237 \\ &= 1346 = 2345 = -347 = 1256 = -257 = -167 \\ &= 3567 = 2467 = 1457 = -456 \\ &= -1234567. \end{aligned}$$

Thus, in the combined design, the DCS drops all odd-length words and has

$$\begin{aligned} I &= 1237 \\ &= 1346 = 2345 = 1256 \\ &= 3567 = 2467 = 1457. \end{aligned}$$

Now all main effects are clear.

- It is usual to add a column 8 ($\equiv +$ in d_1, $-$ in d_2) representing blocks.

- The same idea can be used to de-alias a main affect and all of its two-factor interactions. For instance start with d_1 and consider d_3, in which 5 is replaced by -5:

$$4 = 12, 5 = -13, 6 = 23, 7 = 123.$$

In the combined design, the DCS drops all words containing '5':

$$\begin{aligned} I &= 124 = 236 = 1237 \\ &= 1346 = 347 = 167 \\ &= 2467; \end{aligned}$$

with defining relations (can just drop the one involving 5 in d_1)

$$4 = 12, 6 = 23, 7 = 123.$$

The main effect 5 is strongly clear, and all two-factor interactions involving 5 are clear. Again, this would be run in blocks.

- The fold-over technique might use a lot of resources to address what is maybe a small problem, or the wrong problem – the effects to be de-aliased might not be as described in these examples.

14. Fractional factorials – selection and blocking

- Another approach to de-aliasing uses optimal design theory. Here is the general idea, applied in the leaf spring 2^{5-1}_{IV} design to de-alias BCQ and DEQ in the dispersion model (13.1) (recall that Q was not significant). Consider adding $n = 2$ runs to the original design and then fitting the model

$$E\left[\log S^2\right] = \beta_0 + \beta_1 x_{bl} + \beta_2 x_B + \beta_3 x_D x_Q + \beta_4 x_B x_C x_Q + \beta_5 x_D x_E x_Q.$$

  Here x_{bl} (blocks) will $= -1$ for the 16 runs in the original experiment and $= +1$ in the additional runs. Let X_1 denote the 16×6 model matrix consisting of the columns 1 (intercept), $x_{bl} = -1$, x_B,..., $x_D x_E x_Q$. Let X_2 be $n \times 6$, with its rows – values of 1, 1, x_B,..., $x_D x_E x_Q$ – determined by the levels of B, C, D, E, Q used in these additional runs. There are $2^{5n} = 1024$ possible ways

to choose these levels. The 'D-optimality criterion' dictates that we choose them so that the augmented matrix

$$X = \begin{pmatrix} X_1 \\ X_2 \end{pmatrix}$$

has

$$\det\left(X'X\right) = \det\left(X_1'X_1 + X_2'X_2\right) = \max.$$

(if $\mathbf{y} = X\beta + \varepsilon$ then $cov\left[\hat{\beta}\right] = \sigma^2\left(X'X\right)^{-1}$, etc.)
For small n an exhaustive search is quite feasible. Otherwise, specialized software is available.

- It should be obvious (is it?) that a D-optimal design has de-aliased the effects.

- One might instead seek a design giving maximum precision in the estimation of β_4 and β_5. Define $M = X'X$, partitioned as

$$M = \begin{pmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{pmatrix}$$

with $M_{22}: 2 \times 2$ and $M_{11}: 4 \times 4$. Then

$$\text{cov}\left[\begin{pmatrix} \hat{\beta}_4 \\ \hat{\beta}_5 \end{pmatrix}\right] \propto M^{22} = \left(M_{22} - M_{21}M_{11}^{-1}M_{12}\right)^{-1},$$

so that we aim to choose X_2 so that

$$\det\left(M_{22} - M_{21}M_{11}^{-1}M_{12}\right) = \frac{\det M}{\det M_{11}} = \max;$$

this is the 'D_S-optimality criterion'.

- Possible optimality criteria for fractional factorials - minimum aberration might no longer be a reasonable goal. An example: See Table 5A.3 in text, where two 2_{IV}^{9-4} designs are given. Design d_1 has generators $6 = 123, 7 = 124, 8 = 125, 9 = 1345$. The words in the DCS of length 4 are

$$\begin{aligned} d_1\text{: } I &= 1236 = 1247 = 1258 \\ &= 3467 = 4578 = 3568. \end{aligned} \quad (14.1)$$

Design d_2 has generators $6 = 123, 7 = 124, 8 = 134, 9 = 2345$ and 7 words of length 4:

$$\begin{aligned} d_2\text{: } I &= 1236 = 1247 = 1348 = 3467 \\ &= 2378 = 2468 = 1678. \end{aligned} \quad (14.2)$$

Thus

$$\begin{aligned} W(d_1) &= (0, 6, ...), \\ W(d_2) &= (0, 7, ...), \end{aligned}$$

and d_1 has less aberration than d_2. In fact d_1 has minimum aberration. What about clear two-factor interactions? Any such interaction using numbers which appear together in (14.1) or (14.2) is not clear. In d_1 this leaves only 9 – all 8 two-factor interactions involving 9 are clear. In d_2 all $8 + 8 - 1 = 15$ two-factor interactions involving one or both of 5 and 9 are clear. So we might prefer d_2. More on this in §5.5.

- Another possibility is to seek a large number of clear effects. This can be ambiguous as well - the text gives an example of a 2^{6-2}_{IV} design with 6 clear main effects but no clear two-factor interactions, and a 2^{6-2}_{III} design with only three clear main effects but six clear two-factor interactions.

- **Blocking** of fractional factorials is more complicated than with full factorials, due to there now being two defining subgroups - one for treatments (TDCS) and one for blocks (BDCS). An example will suffice. Table 5B.2 gives (the table of 2^{k-p} fractional factorials run in 2^q blocks) gives a 2^{6-2} design in 4 blocks, with TDCS

$$I = 1235 = 1246 = 3456 \tag{14.3}$$

and the BDCS

$$B1 = 134, B2 = 234; \text{ hence } B1B2 = 12. \tag{14.4}$$

First - what is the design? Start by writing out a full $2^{6-2} = 2^4$ factorial. These form columns $1, 2, 3, 4$. Then determine columns 5 and 6 so that they obey $5 = 123, 6 = 124$. Determine the blocks by computing $B1$ and $B2$ from (14.4) and then taking the 4 combinations of $\pm$ signs.

We note that, multiplying (14.3) by $B1 = 134$ gives

$$B1 = 134 = 245 = 236 = 156;$$

similarly

$$\begin{aligned} B2 &= 234 = 145 = 136 = 256, \\ B1B2 &= 12 = 35 = 46 = 123456. \end{aligned}$$

1	2	3	4	5	6	B1	B2	block
+	+	+	+	+	+	+	+	1
+	+	+	-	+	-	-	-	4
+	+	-	+	-	+	-	-	4
+	+	-	-	-	-	+	+	1
+	-	+	+	-	-	+	-	2
+	-	+	-	-	+	-	+	3
+	-	-	+	+	-	-	+	3
+	-	-	-	+	+	+	-	2
-	+	+	+	-	-	-	+	3
-	+	+	-	-	+	+	-	2
-	+	-	+	+	-	+	-	2
-	+	-	-	+	+	-	+	3
-	-	+	+	+	+	-	-	4
-	-	+	-	+	-	+	+	1
-	-	-	+	+	+	+	+	1
-	-	-	-	+	-	-	-	4

The three-factor interactions confounded with blocks are thus aliased with each other. The six main effects are clear - not aliased with another main

effect or two-factor interaction, nor with blocks. Of the 15 df, we can use 6 to estimate the main effects if three- and higher order interactions can be discounted. In that case the two block effects and their interaction can be estimated, but the latter then prevents estimating $12, 35, 46$. This leaves $16 - 1 - 6 - 3 = 6$ df. The remaining 12 two-factor interactions appear – two each – in six groups of aliases:

$$13 = 25 = 2346 = 1456$$
$$\cdots \text{(four others)} \cdots$$
$$36 = 45 = 1256 = 1234.$$

A member of a group can be estimated only if the other members are negligible. (So perhaps use all 6 df to estimate error?)

- Blocked fractional factorial can be ranked by the number of clear effects ($= 6$ in the example above).

- Many (all?) of the tabulated designs are computer generated; this is an active research area for the combinatorially inclined.

15. Full factorials at three levels

- Motivation - investigate quadratic effects in quantitative factors, investigate all naturally occurring levels (rather than only two and then have to add more later), investigate effects of both increasing and decreasing a nominal level.

- A full factorial with k factors, each at 3 levels, with require 3^k runs (replicated, if error is to be estimated).

- Example: 3^3 factorial, factors A, B, C. The main effects are estimated in the usual way; for instance that of level i of A is estimated by $\bar{y}_{i..} - \bar{y}_{...}$, where $\bar{y}_{i..}$ is the average, over replicates and over the other factors, of the response with A at level i.

  - There are 2 df for each main effect; for instance $SS_A = \sum_{i=1}^{3} (\bar{y}_{i..} - \bar{y}_{...})^2$.

– There are 4 df for each two-factor interaction $A \times B$, $B \times C$, $A \times C$. Each of these can be split into orthogonal components with 2 df, as follows. Denote by $x_1 = 0, 1, 2$ and $x_2 = 0, 1, 2$ the levels of A and B, say. Denote by α, β, γ the runs with $x_1 + x_2$ (mod 3) $= 0, 1, 2$ respectively, and i, j, k the runs with $x_1 + 2x_2$ (mod 3) $= 0, 1, 2$, and $y_{x_1 x_2}$ the average (over the n replicates) of the response at these levels. For instance y_{21} corresponds to (α, j), since

$$\begin{aligned} x_1 + x_2 &= 2 + 1 = 0 \text{ (mod 3)}, \\ x_1 + 2x_2 &= 2 + 2 \cdot 1 = 1 \text{ (mod 3)}. \end{aligned}$$

These observations form a 3×3 Graeco-Latin square:

	x_2		
x_1	0	1	2
0	αi	βk	γj
1	βj	γi	αk
2	γk	αj	βi

– Form averages

$$\begin{aligned} \bar{y}_\alpha &= \frac{1}{3}(y_{00} + y_{12} + y_{21}) \\ &\vdots \\ \bar{y}_k &= \frac{1}{3}(y_{01} + y_{12} + y_{20})\,. \end{aligned}$$

We denote by SS_{AB} the component of $SS_{A\times B}$ obtained by viewing the square as the output in factors (α, β, γ), with overall average $\bar{y}_\cdot = \frac{1}{3}\left(\bar{y}_\alpha + \bar{y}_\beta + \bar{y}_\gamma\right)$:

$$SS_{AB} = 3n\left[(\bar{y}_\alpha - \bar{y}_\cdot)^2 + \left(\bar{y}_\beta - \bar{y}_\cdot\right)^2 + (\bar{y}_\gamma - \bar{y}_\cdot)^2\right].$$

Similarly SS_{AB^2} is the component arising from viewing the square as the output for factors (i, j, k).

– Orthogonal contrasts, but otherwise no physical interpretation?

– The $A \times B \times C$ interaction, on 8 df, splits into 4 orthogonal components, each with 2 df. They are denoted ABC, ABC^2, AB^2C and AB^2C^2, and correspond to classifications

of the observations into groups corresponding to

$$\begin{aligned} x_1 + x_2 + x_3 \text{ (mod 3)} &= 0, 1, 2, \\ x_1 + x_2 + 2x_3 \text{ (mod 3)} &= 0, 1, 2, \\ x_1 + 2x_2 + x_3 \text{ (mod 3)} &= 0, 1, 2, \\ x_1 + 2x_2 + 2x_3 \text{ (mod 3)} &= 0, 1, 2, \end{aligned}$$

respectively.

– Note $A^2B^jC^k$ represents contrasts

$$\begin{aligned} & 2x_1 + jx_2 + kx_3 \text{ (mod 3)} = 0, 1, 2 \\ \Leftrightarrow \quad & x_1 + 2jx_2 + 2kx_3 \text{ (mod 3)} = 0, 2, 1. \end{aligned}$$

Thus $A^2BC = AB^2C^2$, $A^2B^2C = ABC^2$, etc. (square and do mod 3 arithmetic on exponents.) By convention, the exponent of A is 1.

– Why can we equate effects to their squares in this way? Just consider one main effect, say that of A. It is estimated by splitting the observations into three groups, corresponding to $x_1 = 0, 1, 2$ respectively, computing the three averages $\bar{y}_0$, $\bar{y}_1, \bar{y}_2$ and then looking at any two orthogonal contrasts among these (so that the estimate has 2 df) – for instance $[(\bar{y}_2 - \bar{y}_1) - (\bar{y}_1 - \bar{y}_0)] / \sqrt{6}$ and $(\bar{y}_2 - \bar{y}_0) / \sqrt{2}$. The SS of these two contrasts is symmetric in $\bar{y}_0$, $\bar{y}_1, \bar{y}_2$ – they can be interchanged. In fact the SS reduces to $\sum_{i=0,1,2} (\bar{y}_i - \bar{y}_\cdot)^2$, which is clearly symmetric. Now if instead we consider A^2, the three groups corresponding to $x_1 = 0, 1, 2$ will have averages $\bar{y}_{2x_1} = \bar{y}_0$, $\bar{y}_2, \bar{y}_1$ respectively – the groupings are the same, but occur in a different order. Because of the symmetry, this order is irrelevant. So the groupings used to estimate A are the same as those used to estimate A^2, and are used symmetrically.

- **Seat belt experiment**. Four factors affect the 'pull strength' of truck seat belts after a 'crimping operation' joining an anchor and cable: $A =$ pressure (psi) used by crimping machine, $B =$ 'die flat' setting (mm), $C =$ length of crimp (mm) and $D =$ anchor lot (a label); each is at three levels. There are two possible aims: maximize the strength (Y_1), or minimize the 'flash' (Y_2) - the excess metal after the crimping process. For the first of these one might first find settings that maximize $E[Y_1]$ (the harder task), and then others that minimize var$[Y_1]$. For the second, minimize $E[Y_2]$ and then minimize var$[Y_2]$; perhaps combine these by minimizing $E\left[Y_2^2\right]$.

- See R code treating the data from only factors A, B, C as a full 3^3 factorial.

  – Outliers causing non-normality?

16. Fractional factorials at three levels

- In the seat belt experiment only a 3^{4-1} fractional factorial was run; it was defined by $D = ABC$, i.e. $x_4 = x_1 + x_2 + x_3$ (mod 3) or $0 = x_1 + x_2 + x_3 + 2x_4$ (mod 3) or

$$I = ABCD^2. \tag{16.1}$$

Multiply by A^2: $A^2 = BCD^2$. But the levels $2x_1$ range over $\{0, 1, 2\}$ as x_1 does, so

$$A = BCD^2. \tag{16.2}$$

The product of (16.1) and (16.2) gives $A = AB^2C^2D$, so $A = BCD^2 = AB^2C^2D$. The other aliases are obtained in a similar way, yielding 13 sets of aliased effects (p. 277 in text):

$$\begin{aligned} A &= BCD^2 = AB^2C^2D, \\ &\vdots \\ AB &= CD^2 = ABC^2D, \\ AB^2 &= AC^2D = BC^2D, \\ &\vdots \\ CD &= ABC^2 = ABD. \end{aligned}$$

If three-factor interactions are negligible then $A, B, C, D, AB^2, AC^2, AD, BC^2, BD$ and CD can be estimated; these are clear. As well, $AB = CD^2$, $AC = BD^2$ and $BC = AD^2$, so one of each pair must be assumed negligible. (So AB^2 can be estimated even though AB perhaps cannot – a possible advantage of the orthogonal breakdown of the interactions as considered here? Interpretation is a problem still.)

- See main effects and interaction plots for the strength data. It seems that the primary factors affecting the mean strength are A (should be at level 2) and C (should be at level 0 – or 1, from interaction plot, although the difference seems slight).

- Plotting the 2 df effects - various methods are possible, which do not require the fitting of a regression model. If no effects are significant, then the sums of squares are $\sim \sigma^2\chi^2_2$, so a plot of these ordered and scaled effects against the χ^2_2

quantiles gives a graphical display of significance. Alternatively, the p-values are $\sim U(0,1)$, so that $-2\log p \sim \chi^2_2$. See the plots - they tell very much the same story, and bear out the significance of A and C.

- Dispersion analysis: Compute s^2 over the 3 replicates, in each of the 27 runs. Fit $\log\left(s^2\right)$ to the 13 effects listed above. A plot of the ordered SS against the χ^2_2 quantiles indicates that A and AB^2 are the treatments with a significant effect on the dispersion. The main effects plots suggest A at level 2 (in agreement with maximizing the mean) or 1, B at level 1. The interaction plots suggest both at level 1.

- **Blocking**. First consider a 3^2 design in 3 blocks. Recall the Graeco-Latin square used to define the two orthogonal components on an $A \times B$ interaction:

	x_2		
x_1	0	1	2
0	αi	βk	γj
1	βj	γi	αk
2	γk	αj	βi

  Here α, β, γ defined the AB contrast, and i, j, k the AB^2 contrast. If the three blocks correspond to α, β, γ respectively, then block effects are confounded with AB:

$$b_1 = AB.$$

  Blocking schemes for full 3^k factorials ($k > 2$) can be derived in a similar fashion; they are tabulated and classified (in the text) and one can use minimum aberration as a guide.

- When blocking a 3^{k-p} fractional factorial then, as with the 2^{k-p} case, we must deal with both a TDCS and a BDCS. Consider running the 3^{4-1} seat belt experiment (with $I = ABCD^2$) in 9 blocks (of 3 runs each). Take

$$b_1 = AB, b_2 = AC.$$

to generate the blocking scheme. The 9 blocks correspond to the groups of runs with $(b_1, b_2) \in \{0, 1, 2\} \times \{0, 1, 2\}$. There are two other defining relations for blocking:

$$\begin{aligned} b_3 &= b_1 b_2 = A^2BC = AB^2C^2, \\ b_4 &= b_1 b_2^2 = BC^2. \end{aligned}$$

(Recall that relations ranging over all of $\{0, 1, 2\}$ can be equated to their squares.) Any others are repetitions of these; for instance $b_1^2 b_2 = \left(b_1^2 b_2\right)^2 = b_1 b_2^2$.

- Now multiply $I = ABCD^2 = A^2B^2C^2D \left(= I^2\right)$ by b_1, b_2, b_3, b_4:

  $b_1 = (AB)\left(ABCD^2\right)$ & $= (AB)\left(A^2B^2C^2D\right)$

  so $= A^2B^2CD^2$ & $= C^2D$; squaring these gives

  $= ABC^2D$ & $= CD^2$ (and $= AB$).

  See p. 303 – there are 13 sets of aliased and confounded effects. The main effects are clear, as are AB, AC^2, BD and CD.

- **Multiple responses/out-of-spec probabilities**. As in the seat belt experiment we might wish to have one response Y_1 maximized and another response Y_2 minimized. A possibility is to aim to minimize

$$\psi = \lambda P\left(Y_1 \leq c_1\right) + (1-\lambda)\, P\left(Y_2 \geq c_2\right)$$

  for specified constants c_i. In this experiment the factor levels are quantitative and so a possible approach is to fit quadratic regression models to Y_i and to $Z_i = \log\left(s_i^2\right)$ (as in the location/dispersion analyses) with the values of the k factor levels,

and their powers and products, as regressors. These give estimates $\hat{\mu}_i(x)$ and $\hat{\sigma}_i^2(x) = \exp \hat{z}_i(x)$ (for $i = 1, 2$ and $x = (X_1, ..., X_k)$) and then

$$\hat{\psi}(x) = \lambda \Phi\left(\frac{c_1 - \hat{\mu}_1(x)}{\hat{\sigma}_1(x)}\right) + (1-\lambda)\,\Phi\left(\frac{\hat{\mu}_2(x) - c_2}{\hat{\sigma}_2(x)}\right)$$

can be minimized over x.

– The same ideas extend to 'nominal is best'. For each response Y_i, with target value t_i, we might set limits $t_i \pm \Delta_i$. Then estimate the mse by

$$\widehat{\mathsf{mse}}_i(x) = \exp \hat{z}_i(x) + (\hat{y} - t_i)^2.$$

Here $\hat{z}$ and $\hat{y}$ are the predicted values at x. Then choose x to minimize

$$\lambda \frac{\widehat{\mathsf{mse}}_1(x)}{\Delta_1^2} + (1-\lambda)\frac{\widehat{\mathsf{mse}}_2(x)}{\Delta_2^2}.$$

Part VI

RESPONSE SURFACE METHODOLOGY

17. Response surface methodology: introduction and example

- Problem: explore a 'response surface' $f(x_1, ..., x_k)$. Assume variables have been coded to lie in some symmetric interval; for instance if $L \leq X \leq U$ with midpoint $M = (U + L)/2$ then

$$x = 2(X - M)/(U - L) \in [-1, 1].$$

With $\mathbf{x} = (x_1, ..., x_k)'$ we may observe

$$y_i = f(\mathbf{x}_i) + \varepsilon_i$$

at chosen values $\{\mathbf{x}_i\}$. Assume i.i.d. $N\left(0, \sigma^2\right)$ errors. Often the purpose is to find the extrema of f, without a precise knowledge of this function.

- Assume that a small number (k) of important factors have been identified, perhaps by running a preliminary fractional factorial.

- **Example**: The production of Zantac (used to treat ulcers) requires the separation of the active ingredient ranitidine from other substances. Three important factors $x_1 =$ pH level, $x_2 =$ voltage, $x_3 =$ concentration have been identified as affecting $y =$ CEF, a certain quality measure. In the initial screening experiment, and before coding, these variables have possible levels $X_1 \in \{2, 3.42, 5.5, 7.58, 9\}$, $X_2 \in \{9.9, 14, 20, 26, 30.1\}$, $X_3 \in \{0, 2, 5, 8, 10\}$. After coding they have possible levels $x_1, x_2 \in \{-1.68, -1, 0, 1, 1.68\}$, $x_3 \in \{-1.67, -1, 0, 1, 1.67\}$.

- Start by taking a linear approximation of f:

$$f(\mathbf{x}) \approx f(\mathbf{0}) + \mathbf{x}'\nabla_f(\mathbf{0}) = \beta_0 + \sum_{i=1}^{k} \beta_i x_i.$$

Run a first order design, typically a 2^{k-p} factorial, i.e. 2^{k-p} of the 2^k 'corner points' $\mathbf{x} = (\pm 1, ..., \pm 1)$, with ± 1 denoting the HI and LO levels of the coded variables. Denote by n_f the

number of such runs. As well, make n_c runs at the 'centre point' $\mathbf{x} = \mathbf{0}$. These may be used to estimate error and to check for curvature of the response surface, as follows. Let $\bar{y}_f$ and $\bar{y}_c$ be the response averages at the factorial and centre points. Consider the second order model

$$E[y|\mathbf{x}] = \beta_0 + \sum_{i=1}^{k} \beta_i x_i + \sum_{i<j}^{k} \beta_{ij} x_i x_j + \sum_{i=1}^{k} \beta_{ii} x_i^2. \quad (17.1)$$

In this model

$$E[\bar{y}_c] = E[y|\mathbf{x} = \mathbf{0}] = \beta_0,$$

and, since the x_i and $x_i x_j$ sum to zero over the n_f factorial runs (and $= 0$ otherwise),

$$E\left[\bar{y}_f\right] = \beta_0 + \sum_{i=1}^{k} \beta_{ii}.$$

Thus the 'overall curvature' $\sum_{i=1}^{k} \beta_{ii}$ is estimated by $\bar{y}_f - \bar{y}_c \sim N\left(\sum_{i=1}^{k} \beta_{ii}, \sigma^2\left(\frac{1}{n_f} + \frac{1}{n_c}\right)\right)$ and tested by

$$t_{curv} = \frac{\bar{y}_f - \bar{y}_c}{\hat{\sigma}\sqrt{\frac{1}{n_f} + \frac{1}{n_c}}},$$

where $\hat{\sigma}$ is computed from the responses at the centre point (why?), and is on $n_c - 1$ df.

– An equivalent method of testing curvature, which applies even when $n_c = 1$, is as follows. Note that all k squares are equal, so that (17.1) is equivalent to

$$E[y|\mathbf{x}] = \beta_0 + \sum_{i=1}^{k} \beta_i x_i + \sum_{i<j}^{k} \beta_{ij} x_i x_j + x_1^2 \beta_*, \tag{17.2}$$

with the single parameter $\beta_* = \sum_{i=1}^{k} \beta_{ii}$ replacing $\beta_{11}, ..., \beta_{kk}$. The t-test for β_* is then equivalent to that for curvature.

– All this assumes that the β_{ii} all have the same sign, else they could cancel each other out. In practice this is apparently unlikely (except near a saddlepoint), but one could fit (17.1), using a design with axial points, if one were worried about this.

- Suppose that the result from the first order design is that curvature appears to be negligible. If we are seeking a maximum of f, we should move in the direction of steepest ascent:

$$\arg \max_{\|\mathbf{x}\|=1} \left|\mathbf{x}'\widehat{\nabla_f(\mathbf{0})}\right| = \arg \max_{\|\mathbf{x}\|=1} \left|\mathbf{x}'\hat{\boldsymbol{\beta}}\right| = \hat{\boldsymbol{\beta}} \Big/ \left\|\hat{\boldsymbol{\beta}}\right\|,$$

where $\hat{\boldsymbol{\beta}} = \left(\hat{\beta}_1, ..., \hat{\beta}_k\right)$. Then observations are to be made at several points $\mathbf{x}_\lambda = \lambda\hat{\boldsymbol{\beta}} \Big/ \left\|\hat{\boldsymbol{\beta}}\right\|$. Among these, if a maximum is found at $\lambda = \lambda_0$ then $\lambda_0\hat{\boldsymbol{\beta}} \Big/ \left\|\hat{\boldsymbol{\beta}}\right\|$ is taken as the new centre point and the experiment is repeated. Curvature and interaction are tested after fitting (17.2), for which t_{curv} is produced as the t-statistic for testing the significance of β_*.

- See R code on web site. Responses are simulated from $E[y|\mathbf{w}] =$

  $$5+\cos\left(\frac{2w_1+w_2-15}{10}\right)+\cos\left(\frac{w_1-2w_2+5}{10}\right),$$

  $0 \le w_1, w_2 \le 20$, with $N\left(0, .15^2\right)$ errors added.

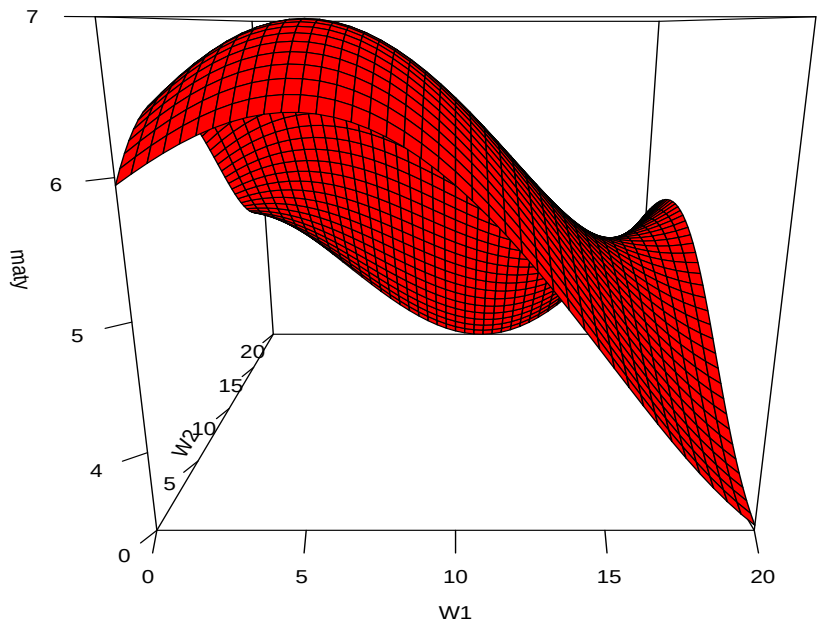

Mean response; maximum at $(w_1, w_2) = (5, 5)$.

A 2^2 factorial augmented by 5 centre points is employed. Three iterations of steepest ascent, starting at $(10, 10)$ and testing curvature and interaction after each yield:

	w1	w2	p.inter	p.curv
[0]	10.000	10.000	1.000	1.000
[1]	7.225	7.119	0.604	0.746
[2]	5.963	5.567	0.084	0.041

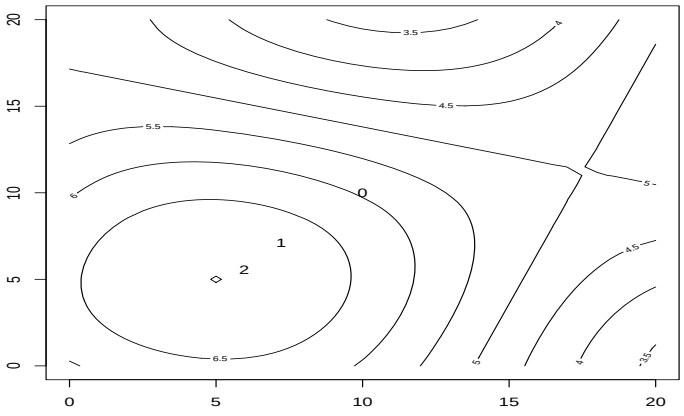

Contour plot; iteratively chosen centre points shown.

- Once curvature becomes an important issue a second order design and model should be employed. The most common second order design augments the first order design with $2k$ 'axial points'$= \pm\alpha$ on each coordinate axis in the x-space. Such 'central composite' designs are studied below. Then a full second order model is fitted. If this model is represented as

$$\begin{aligned} E[y|\mathbf{x}] &= \beta_0 + \sum_{i=1}^{k} \beta_i x_i + \sum_{i<j}^{k} \beta_{ij} x_i x_j + \sum_{i=1}^{k} \beta_{ii} x_i^2 \\ &= \beta_0 + \mathbf{x}'\mathbf{b} + \mathbf{x}'\mathbf{B}\mathbf{x}, \end{aligned} \tag{17.3}$$

where

$$\mathbf{B}_{ij} = \begin{cases} \beta_{ii}, & i = j, \\ \frac{1}{2}\beta_{ij}, & i \neq j, \end{cases}$$

then the stationary point is estimated by

$$\mathbf{x}_s = -\frac{1}{2}\hat{\mathbf{B}}^{-1}\hat{\mathbf{b}},$$

with the eigenvalues of $\hat{\mathbf{B}}$ determining the nature of the extremum.

18. Response surface methodology: designs

- Inserting $\mathbf{x} = \mathbf{x}_s + (\mathbf{x} - \mathbf{x}_s) \stackrel{def}{=} \mathbf{x}_s + \mathbf{z}$ into (17.3) and simplifying gives

$$\begin{aligned} \hat{y}_{|\mathbf{z}} &= \hat{\beta}_0 - \frac{1}{4}\hat{\mathbf{b}}'\hat{\mathbf{B}}^{-1}\hat{\mathbf{b}} + \mathbf{z}'\hat{\mathbf{B}}\mathbf{z} \\ &= \hat{y}_s + \mathbf{z}'\hat{\mathbf{B}}\mathbf{z}, \end{aligned}$$

  since that the estimated stationary point $\mathbf{x}_s$ corresponds to $\mathbf{z} = \mathbf{0}$. We infer that this gives a maximum if $\hat{\mathbf{B}}$ is negative definite.

- In the example above, we start with the new centre point $(w_1, w_2) = (5.593, 5.567)$ and add four axial points $(\pm\alpha, 0)$, $(0, \pm\alpha)$ with $\alpha = \sqrt{2}$ to the original design. Sample output:

```
Estimated stat. point: (w1,w2) =  (5.46, 5.09)
Eigenvalues of Hessian are -0.289 -0.399
```

- Suppose that $\hat{\mathbf{B}}$ was not negative definite, and in fact was indefinite. Then the stationary point is a saddlepoint, and the response increases in some directions and decreases in others as one moves away from the stationary point. Diagonalize $\hat{\mathbf{B}}$:

$$\hat{\mathbf{B}} = \mathbf{P}\Lambda\mathbf{P}' = \sum_{i=1}^{k} \lambda_i \mathbf{p}_i \mathbf{p}_i',$$

where $\lambda_1 > \cdots > \lambda_k$ and the $\{\mathbf{p}_i\}$ are normalized eigenvectors. Then

$$\begin{aligned} \hat{y} &= \hat{y}_s + \mathbf{z}'\hat{\mathbf{B}}\mathbf{z} \\ &= \hat{y}_s + \sum_{i=1}^{k} \lambda_i \left(\mathbf{z}'\mathbf{p}_i\right)^2 \end{aligned}$$

is maximized, among all $\mathbf{z}$ with norm c, by $\mathbf{z} = c\mathbf{p}_1$ (hence $\mathbf{z}'\mathbf{p}_i = 0$ if $i > 1$) and then

$$\hat{y} = \hat{y}_s + \lambda_1 c^2.$$

So one would move to $\mathbf{x} = \mathbf{x}_s + c\mathbf{p}_1$ for a suitable c – perhaps chosen experimentally, and increased as long as y increases, or chosen as large as possible with the independent variables remaining within the experimental region.

- **Central composite designs** are the most common second-order designs in RSM. When there are k factors, and with $\mathbf{x} = (x_1, ..., x_k)'$ in coded form, they consist of

  1. n_f 'corner points' in which each $x_i \in \{-1, 1\}$ (the 'factorial portion' – a 2^{k-p} design),

  2. n_c 'centre points' $\mathbf{x} = \mathbf{0}$,

  3. $2k$ 'star points' or 'axial points' $\mathbf{x} = (0, ..., 0, \pm\alpha, 0, ..., 0)'$ with $\alpha \in \left[1, \sqrt{k}\right]$.

  – Motivation - the axial points allow for efficient estimation of the pure quadratic terms, while the corner points allow for the estimation of the interactions.

One can implement such a design by augmenting the first-order designs (and data) of the previous section by axial points and possibly additional centre points.

- Choice of the factorial portion:

  - How small can the factorial portion be if the CCD is to allow for the estimation of all parameters in a full quadratic model? There are $1+k+k+\frac{k(k-1)}{2}$ parameters and n_f+1+2k distinct points, and so it is necessary to have

    $$n_f \geq 1+2k+\frac{k\,(k-1)}{2}-(1+2k)=\frac{k\,(k-1)}{2}.$$

  - Recall that in a 2^{k-p} design, there are $2^{k-p}-1$ sets of aliased factors (not counting the constant term). (Example – (12.1)).

  - **Theorem**: In any CCD whose factorial portion is a 2^{k-p} design that does not use any main effect as a defining relation, one can estimate the parameters β_0, $\{\beta_i\}_{i=1}^k$, $\{\beta_{ii}\}_{i=1}^k$ and exactly one β_{ij} $(i<j)$ from each set of aliased effects. (**Proof**: 'straightforward but tedious'.)

- **Examples**:
  (i) $k = 2$, can use a 2^{2-1}_{II} with $I = AB$. Or a full 2^2 factorial.
  (ii) $k = 3$, can use a 2^{3-1}_{III} with $I = ABC$. [**Details**: the $2^{3-1} - 1 = 3$ sets of aliases are $A = BC, B = AC, C = AB$, so estimating exactly one interaction from each of these gives us the ability to estimate all of them.] Or a full 2^3 factorial.
  (iii) A resolution III design whose defining relation does not contain words of length four is 'resolution III^*', and satisfies the conditions of the theorem.
  Possible designs are tabulated in the text.

- Choice of α:

  - $\alpha = \sqrt{k}$: all noncentral points are equidistant from $\mathbf{0}$ ($\|\mathbf{x}\| = \sqrt{k}$). 'Spherical' design. Most suitable for a spherical design region; perhaps impractical for large k.

- $\alpha = 1$: need only 3 levels $\{-1, 0, 1\}$ for each factor. 'Face-centred' design.

- $\alpha = \sqrt[4]{n_f}$: 'rotatable' design - var$[\hat{y}(\mathbf{x})]$ is a function of $\mathbf{x}$ only through $\|\mathbf{x}\|$. But this is in the coded variables and is not scale invariant. For the theory see Box & Hunter (1957) on course web site. Example: If $k = 2, n_f = 4$ and $\alpha = \sqrt{2}$ then the design is both spherical and rotatable.

- Number of centre points:

  - Subjective; when α is near $\sqrt{k}$ then typically $n_c = 3, 4, 5$; when α is near 1 then typically $n_c = 1, 2$. These choices provide reasonably stable prediction variances.

  - If $\alpha = \sqrt{k}$ then it is necessary to have $n_c > 0$ in order that $\mathbf{X}'\mathbf{X}$ be invertible. Example: suppose $k = 2$, the factorial portion is 2^2,

$\alpha = \sqrt{k}$ but $n_c = 0$. The model matrix is

$$\mathbf{X} = \left(\begin{array}{cccccc} 1 & x_1 & x_2 & x_1^2 & x_2^2 & x_1x_2 \\ \hline 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & -1 & 1 & 1 & -1 \\ 1 & -1 & 1 & 1 & 1 & -1 \\ 1 & -1 & -1 & 1 & 1 & 1 \\ 1 & \alpha & 0 & \alpha^2 & 0 & 0 \\ 1 & 0 & \alpha & 0 & \alpha^2 & 0 \\ 1 & -\alpha & 0 & \alpha^2 & 0 & 0 \\ 1 & 0 & -\alpha & 0 & \alpha^2 & 0 \end{array}\right)$$

and '$x_1^2 + x_2^2 \equiv 2$' – the sum of the x_1^2 and x_2^2 columns is twice the first column. (If $n_c > 0$ then this is avoided.) The generalization is obvious (?).

- There are various other possibilities, in particular 'uniform shell' designs, consisting of a number of points uniformly spaced on the k-dimensional unit sphere and one or more centre points. For instance when $k = 2$ there are n_1 points on the unit circle (which then form a regular polygon) and n_2 centre points. Some are tabulated in the text.

Part VII

ROBUST PARAMETER DESIGN

19. Robust parameter design I

- General idea: Noise factors (hard-to-control, perhaps random in practice); choose control factor settings to reduce effects of variation in noise factors.

- **Example 1** (will return to this): Epitaxial layer growth experiment. There are 8 control factors (the 4 considered earlier and 4 more such as $H =$ nozzle position). Recall that the wafers are mounted on the top or the bottom of one of four 'facets' on a 'susceptor' (cylinder) which is spun inside a jar. The two locations are the levels of factor L, and the four facets form the levels of factor M. These are the 'noise' factors – the aim is to attain uniform wafer thickness (y) over facets and locations. There will be $2 \times 4 = 8$ observations for each combination of control factor settings (i.e. a full factorial for the noise factors); these in turn

form a 2^{8-4}_{IV} fractional factorial with defining relations

$$D = -ABC, F = ABE, G = ACE, H = BCE.$$

Formally, we aim to minimize the variation in y resulting from varying levels of L and M, while maintaining an average y of 14.5 ($\pm$.5 μm).

- **Example 2**: Leaf spring experiment. Recall five factors $B, C, D, E, Q =$ 'quench oil temperature'; this latter factor is not controllable in normal production. With some difficulty, for experimental purposes it can be varied over two levels (130-150 and 150-170^oF). The control factors form a 2^{4-1}_{IV} half factorial with $I = BCDE$; for each combination of these levels both noise factor settings are used. This experiment is replicated 3 times.

- Strategies for reducing variation

  - Read §11.2, 11.3; various strategies discussed including in particular 'reducing the variation in the noise factors'.

  - At fixed (optimal) levels of the control factors the variation in the response will result from random variation and from variation in the noise factors, and so one might aim to reduce the latter.

  - This approach can be difficult; an alternate is to find levels of the control factors for which the response is less sensitive to the noise variation. This is the idea behind Robust Parameter Design (RPD).

- Variance reduction through RPD (Taguchi)

  - Control factors $\mathbf{x}$, noise factors $\mathbf{z}$, response

$$E\left[y|\mathbf{x},\mathbf{z}\right] = f\left(\mathbf{x},\mathbf{z}\right).$$

– Variation arises through $\mathbf{z}$, but also through the interaction of $\mathbf{z}$ with $\mathbf{x}$; e.g.

$$\begin{aligned} y &= \mu + \alpha x_1 + \beta z + \gamma x_2 z + \varepsilon \\ &= \mu + \alpha x_1 + (\beta + \gamma x_2) z + \varepsilon; \end{aligned}$$

one might aim to set x_2 so as to minimize $|\beta + \gamma x_2|$, then y will change only slowly with changes in z.

– More generally, choose $\mathbf{x}$ so that $f(\mathbf{x}, \mathbf{z})$ is 'flat' as a function of $\mathbf{z}$.
Reason: If $y_{|\mathbf{x}_0,\mathbf{z}} = f(\mathbf{x}_0, \mathbf{z}) + \varepsilon$ with $\mathbf{z}$ varying around $\boldsymbol{\mu}_0$ then the delta method leads to

$$\sigma_Y^2 \approx \sum \left(\frac{\partial f(\mathbf{x}_0, \mathbf{z})}{\partial z_i}_{|\boldsymbol{\mu}_0} \right)^2 \text{var}[z_i]$$

and so one might aim to minimize some of the var$[z_i]$ (the first strategy above; this is called 'tolerance design') or by choosing values of $\mathbf{x}_0$ at which the gradient is small in magnitude (f is 'flat'); the latter is Taguchi's notion of RPD.

- **Cross array strategy**: here there are two design matrices - a control 'array' and a noise array; the levels of each are crossed (as in Example 1 - all 2^{8-4} combinations of levels in the control array are run will all 2×4 combinations of the noise array.

- **Approach 1: location and dispersion modelling**. Recall two step procedures discussed earlier. In this 'nominal-the-best' application we first model the dispersion and select levels of the important dispersion factors to minimize dispersion; we then model the mean and choose levels of the important 'adjustment' factors to achieve the desired mean. This was implemented in Lectures 10 (epitaxial layer growth experiment) and 13 (leaf spring experiment), with $\bar{y}$ and $\log S^2$ computed from the replicates (6 and 3 respectively) at each combination of levels of the control factors. In the current approach they are instead computed, at each combination of levels of the control factors,

from the combinations of levels (perhaps replicated) of the noise factors ($2 \times 4 = 8$ combinations and 2 combinations $\times$ 3 replicates $= 6$, respectively).

In other words the noise factors and replicates together are treated as replicates of the levels of the control factors; the rest of the analysis is as was done previously.

- **Example 1 continued**: Location/dispersion modelling in epitaxial layer growth experiment.

  - 2^{8-4}_{IV} fractional factorial with defining relations $D = -ABC, F = ABE, G = ACE, H = BCE$. We aim to minimize the variation in y resulting from varying levels of L (2 levels) and M (4 levels), while maintaining an average y of 14.5 ($\pm$.5 μm). These are treated as 8 replicates.

  - R code on course web site. Fit the 8 main effects and 7 two-factor interactions with A to both y and $\log S^2$. Look at the half-normal plots:

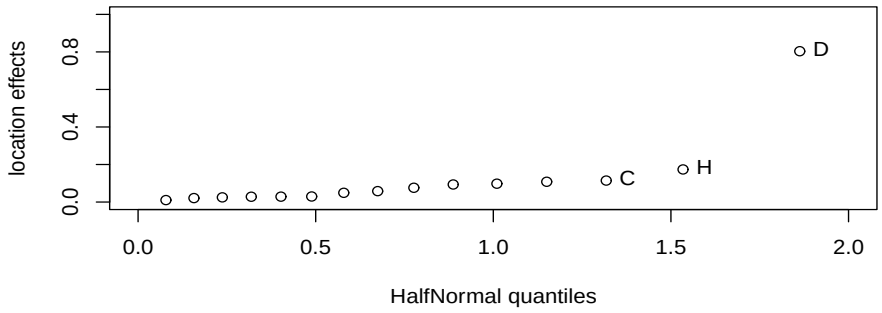

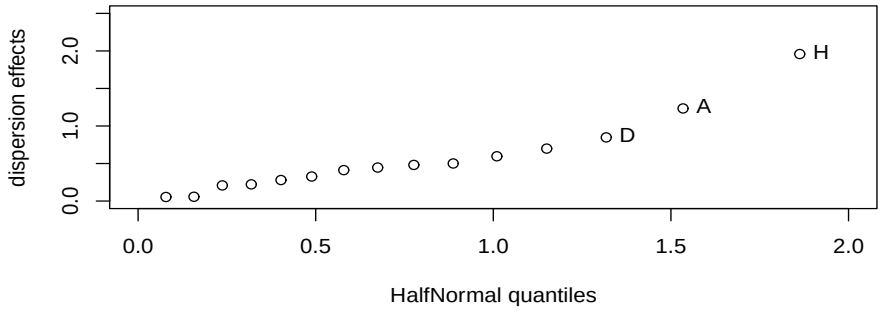

Significant location effect is xD; significant dispersion effects are xA and xH.

Fit the significant effects:

```
> fit.location$coef
(Intercept)          xD
 14.3519477   0.4019523
> fit.dispersion$coef
(Intercept)          xA          xH
 -1.8199427   0.6169592  -0.9794596
```

Best settings are $A-$ and $H+$; now

$$\hat{y} = 14.352 + .402xD = 14.5$$

at

$$xD = .3683,$$

from which an optimal time D can be obtained by interpolation.

- A variation on this is to replace the minimization of $\log S^2$, at the first step by the maximization of the sample signal-to-noise (SN) ratio $\hat{\eta} = \log\left(\frac{\bar{y}^2}{S^2}\right)$, estimating $\eta = \log\left(\frac{\mu^2}{\sigma^2}\right)$ - details in §11.9. Not always an appropriate approach, since large values of η can arise not only from small values of σ^2 (the desired aim) but also large values of μ.

20. Robust parameter design II

- **Approach 2: response modelling**. Which interactions between control and noise factors are increasing the variability of y due to noise? The previous approach cannot answer this. Here we model $\hat{y}$ as a function of both control and noise factors. This 'response model' is then analyzed in one of two ways:

  - Control-by-noise interaction plots – allow one to choose control factor settings at which y is more constant as the noise factors change;

  - Transmitted variance model – model var$[\hat{y}]$ as a function of the control variables, by treating the noise variables as random variables; then minimize over the control levels.

    These two approaches can lead to different settings.

- **Example 1 continued**: Response modelling in epitaxial layer growth experiment.

  - Data in text (Table 11.2, in an unusual format – getting it into R is a project in itself).

  - 2_{IV}^{8-4} fractional factorial with defining relations $D = -ABC, F = ABE, G = ACE, H = BCE$. We aim to minimize the variation in y resulting from varying levels of L (location of wafer on facet; at 2 levels) and M (location of facet on cylinder; at 4 levels), while maintaining an average y of 14.5 ($\pm$.5 μm).

  - Three orthogonal contrasts for M:

$$\begin{aligned} xMl &= I(M \in \{1,2\}) - I(M \in \{3,4\}) \\ xMq &= I(M \in \{1,4\}) - I(M \in \{2,3\}) \\ xMc &= I(M \in \{1,3\}) - I(M \in \{2,4\}) \end{aligned}$$

Analysis 1: Control-by-noise interaction plots.

Fit all: `fit1 = lm(y ~(xA + xB + xC + xD + xE + xF + xG + xH + xL + xMl + xMq + xMc)^3)` and isolate those 107 terms which can be fitted. Identify the significant ones from a half-normal plot.

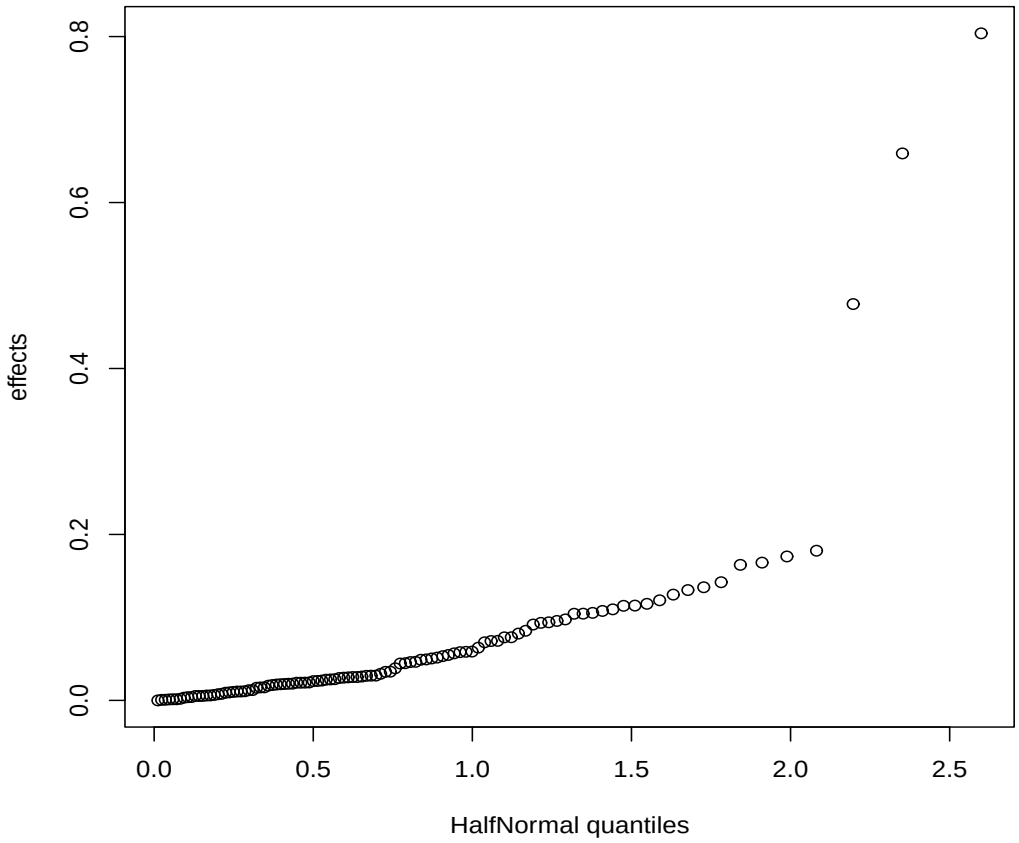

Seven most significant effects, in increasing order, are xA:xH:xMq, xC:xMl, xH, xMl, xH:xL, xL, xD.

Fit these seven most significant effects; they include three control×noise interactions:

```
lm(formula = y ~xD + xH + xL + xMl +
          xH:xL + xC:xMl + xA:xH:xMq)

Coefficients:
             Estimate Std. Error t value Pr(>|t|)
(Intercept) 14.35195    0.02650 541.629  < 2e-16
xD           0.40195    0.02650  15.169  < 2e-16
xH           0.08671    0.02650   3.272 0.001393
xL           0.32957    0.02650  12.438  < 2e-16
xMl         -0.09020    0.02650  -3.404 0.000903
xH:xL       -0.23880    0.02650  -9.012 3.84e-15
xMl:xC      -0.08304    0.02650  -3.134 0.002169
xH:xA:xMq   -0.08166    0.02650  -3.082 0.002552
```

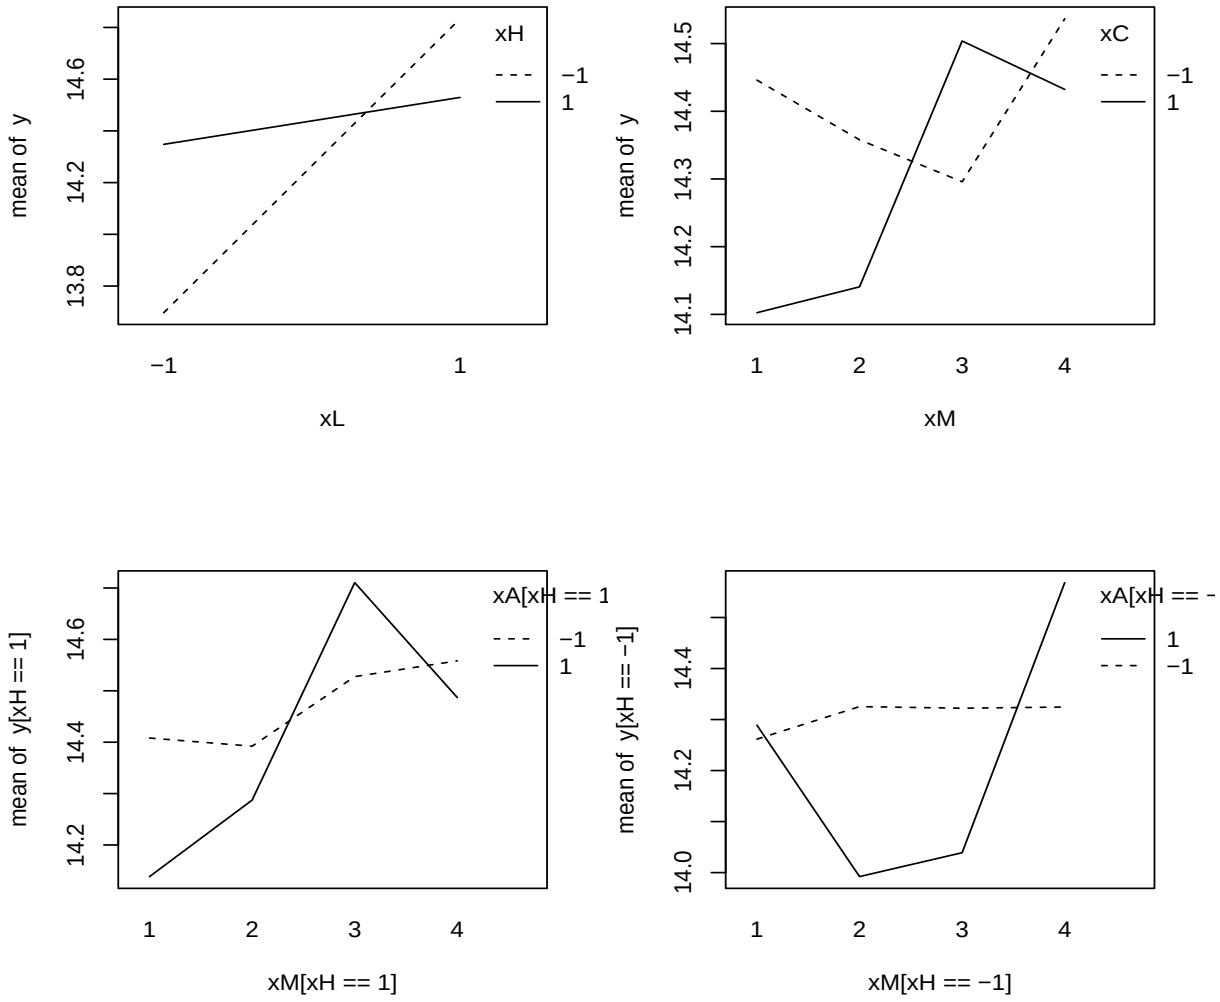

Significant control×noise interactions: best (variance minimizing) settings seem to be $H+, C-, A-$. Confirms previous analysis and adds information about C.

Analysis 2: Transmitted variance model.

$$\begin{aligned}\hat{y} &= 14.352 + 0.402xD + 0.087xH + 0.330xL \\ &- 0.090xMl - 0.239xH \cdot xL - 0.083xC \cdot xMl \\ &\quad -0.082xA \cdot xH \cdot xMq \\ &= 14.352 + (0.330 - 0.239xH)\, xL \\ &\quad - (0.090 + 0.083xC)\, xMl - (0.082xA \cdot xH)\, xMq.\end{aligned}$$

View the noise variables as independent r.v.s, each with $P(X = \pm 1) = .5$. Then $E[X] = 0$, $\text{var}[X] = E\left[X^2\right] = 1$ and so

$$\begin{aligned}\text{var}[\hat{y}] &= (0.330 - 0.239xH)^2 + (0.090 + 0.083xC)^2 \\ &\quad + (0.082xA \cdot xH)^2 .\end{aligned}$$

Since $xH^2 = xC^2 = xA^2 = 1$, this becomes

$$\begin{aligned}\text{var}[\hat{y}] &= const. - 2 \cdot 0.330 \cdot 0.239xH \\ &\quad +2 \cdot 0.090 \cdot 0.083xC \\ &= const - 0.1574xH + 0.0150xC;\end{aligned}$$

minimized with $H+$ and $C-$, as in the first analysis.

- An alternative strategy: **Single array and response modelling**. Using a single array can result in considerable savings.

- **Example 2 continued**: Consider a simplified version of the leaf spring experiment, with 3 control factors $A =$ heating time, $B =$ transfer time, $C =$ hold-down time and 2 noise factors $a =$ high heat temperature, $b =$ quench oil temperature, all at 2 levels. A crossed array, with a 2^3 factorial for the control factors and a 2^2 factorial for the noise factors requires $8 \times 4 = 32$ runs. Arrays of this size are necessary to estimate the two-factor interactions – for instance a 2^{3-1} design for the control factors would have main effects aliased with two-factor interactions.

- If instead one uses a single array – a 2^{5-1}_V half-factorial with

$$I = ABCab,$$

then all main effects are strongly clear, and all two-factor interactions are clear. In the single array strategy only 16 runs are required.

- The ensuing analysis is carried out by the response modelling approach.

Part VIII

COMPUTER EXPERIMENTS

Source materials on course web site:

- JSM 2006 presentation (kindly provided) by Thomas Santner

- JSM 2006 presentation (kindly provided) by William Notz

- 1989 Statistical Science paper by Sacks, Welch, Mitchell & Wynn

As well, some material in the next two lectures is from *The Design and Analysis of Computer Experiments* (Springer, 2003), by T. J. Santner, B. J. Williams and W. I. Notz.

21. Computer experiments - introduction & modelling

- In some situations performing a physical experiment is not feasible: physical process is technically too difficult to study; number of variables is too large; too expensive, etc.

- It might be feasible to conduct a computer experiment, if the physical process relating the inputs $\mathbf{x}$ to the response $y(\mathbf{x})$ can be described by a mathematical model amenable to efficient numerical methods. (Examples: weather modelling, design of prosthetic devices, ... See Santner slides pp. 14-19.)

- Interest is in experimental regimes requiring few runs (since the mathematical code can be very time consuming); note that $y(\mathbf{x})$ is deterministic - *randomization, blocking, replication play no role*.

- Sometimes one can combine physical and computer experiments.

- Some possible problems of interest:

  1. Interpolation: given 'training' inputs and outputs $\left\{\mathbf{x}_i^t, y_i^t = y\left(\mathbf{x}_i^t\right)\right\}_{i=1}^n$, predict the output at a new input $\mathbf{x}_0$.

  2. Design: determine a set of inputs at which to carry out the code. The appropriate design will depend on the objective:

     – exploratory designs ('space-filling');

     – prediction-based designs;

     – optimization-based designs (akin to RSM).

  3. Uncertainty analysis - output might depend on environmental (random) variables $\mathbf{X}_e$ as well as design variables $\mathbf{x}_c$; we might aim to study the distribution of $y\left(\mathbf{x}_c, \mathbf{X}_e\right)$, perhaps with an aim to reducing the variability (akin to RPD?).

- Gaussian Stochastic Process (GaSP) Models: although the response is non-random, its behaviour as a function of $\mathbf{x}$ is modelled as a realization of a random function $Y(\mathbf{x})$, correlated with, hence allowing the use of, those at nearby values of $\mathbf{x}$:

$$Y(\mathbf{x}) = \mathbf{f}'(\mathbf{x})\boldsymbol{\beta} + Z(\mathbf{x}).$$

  Here $\mathbf{f}(\mathbf{x}) = (f_1(\mathbf{x}), ..., f_k(\mathbf{x}))$ is a vector of known regression functions, $\boldsymbol{\beta}$ an unknown regression vector of regression parameters and $Z(\mathbf{x})$ a GaSP:

  - $\mathbf{x} \in \boldsymbol{\chi}$ = 'design space';
  - $E[Z(\mathbf{x})] = 0$ (so $E[Y(\mathbf{x})] = \mathbf{f}'(\mathbf{x})\boldsymbol{\beta}$);
  - var$[Z(\mathbf{x})] = \sigma_Z^2$;
  - cov$[Z(\mathbf{x}_1), Z(\mathbf{x}_2)] = \sigma_Z^2 R(\mathbf{x}_1 - \mathbf{x}_2)$ where $R(\mathbf{0}) = 1$, R is symmetric.
  - Often $R = R(\|\mathbf{x}_1 - \mathbf{x}_2\|)$; $R(\cdot)$ may also depend on unknown parameters $\boldsymbol{\theta}$.

  - All finite dimensional marginals are multivariate normal.

See Santner slides pp. 36, 37.

- If corr$[Z(\mathbf{x}_1), Z(\mathbf{x}_1+\mathbf{h})] = R(\mathbf{h})$ then the process is (weakly, or second order) stationary: corr$[Z(\mathbf{x}_1+\mathbf{t}), Z(\mathbf{x}_2+\mathbf{t})] =$ corr$[Z(\mathbf{x}_1), Z(\mathbf{x}_2)]$ for any $\mathbf{t}$. It is strongly stationary if all marginal distributions are invariant in this sense. In the presence of normality these notions are equivalent (why?).

- If corr$[Z(\mathbf{x}_1), Z(\mathbf{x}_2)] = R(\|\mathbf{x}_1-\mathbf{x}_2\|)$ then the process is stationary and isotropic - the correlation depends only on how far apart the inputs are.

- One way to generate correlation functions is

$$\text{corr}\left[Z(\mathbf{x}_1), Z(\mathbf{x}_1+\mathbf{h})\right] = E_{\mathbf{w}}\left[\cos\left(\mathbf{h}'\mathbf{w}\right)\right]$$

for a random vector $\mathbf{w}$; the distribution of $\mathbf{w}$ is the associated spectral distribution (or spectral density, if there is one).

– **Example**: If x is one dimensional and the spectral density is $N\left(0, 2/\theta^2\right)$, i.e.

$$f(w) = \frac{1}{\sqrt{2\pi}}\frac{\theta}{\sqrt{2}}e^{-\frac{w^2\theta^2}{4}},$$

then

$$\begin{aligned} R(h) &= \text{corr}\left[Z(x_1), Z(x_1+h)\right] \\ &= \int_{-\infty}^{\infty} \cos(hw) f(w)\, dw \\ (and &= \int_{-\infty}^{\infty} e^{ihw} f(w)\, dw) \\ &= e^{-\left(\frac{h}{\theta}\right)^2}. \end{aligned}$$

– **Example**: If x is one dimensional and the spectral density is the t-density then $R(\cdot)$ is the Matérn correlation function, with a much more complicated form (although one case is $R(h) = e^{-\frac{\sqrt{2}|h|}{\theta}}$).

- **Prediction**: Given training data $\left\{\mathbf{x}_i^t, y_i^t\right\}_{i=1}^n$ and an untried new input $\mathbf{x}_0$, the minimum MSE predictor of $Y(\mathbf{x}_0)$ is (assigned)

$$\hat{y}(\mathbf{x}_0) = E\left[Y(\mathbf{x}_0) \,|\, \mathbf{Y}^n = \mathbf{y}^n\right],$$

where

$$\begin{aligned} \mathbf{y}^n &= \left(y\left(\mathbf{x}_1^t\right), ..., y\left(\mathbf{x}_n^t\right)\right), \\ \mathbf{Y}^n &= \left(Y\left(\mathbf{x}_1^t\right), ..., Y\left(\mathbf{x}_n^t\right)\right). \end{aligned}$$

For a GaSP model, the joint distribution of $\begin{pmatrix} Y(\mathbf{x}_0) \\ \mathbf{Y}^n \end{pmatrix}$ is

$$N_{n+1}\left(\begin{pmatrix} \mathbf{f}'(\mathbf{x}_0)\boldsymbol{\beta} \\ \mathbf{F}\boldsymbol{\beta} \end{pmatrix}, \sigma_Z^2 \begin{pmatrix} 1 & \mathbf{r}'(\mathbf{x}_0) \\ \mathbf{r}(\mathbf{x}_0) & \mathbf{R} \end{pmatrix}\right),$$

where

$$\begin{aligned} \mathbf{R} &= R\left(\mathbf{x}_i - \mathbf{x}_j\right) : n \times n, \\ \mathbf{r}'(\mathbf{x}_0) &= \left(R(\mathbf{x}_0 - \mathbf{x}_1), ..., R(\mathbf{x}_0 - \mathbf{x}_n)\right), \\ \mathbf{F} &= \left(\mathbf{f}(\mathbf{x}_1), ..., \mathbf{f}(\mathbf{x}_n)\right)' : n \times k. \end{aligned}$$

It follows from standard Normal theory (STAT 575 lec. 3, for instance) that

$$\hat{y}(\mathbf{x}_0) = \mathbf{f}'(\mathbf{x}_0)\boldsymbol{\beta} + \mathbf{r}'(\mathbf{x}_0)\mathbf{R}^{-1}\left(\mathbf{y}^n - \mathbf{F}\boldsymbol{\beta}\right).$$

This is evaluated at the GLS estimate (i.e. the BLUE)

$$\hat{\boldsymbol{\beta}} = \left(\mathbf{F}'\mathbf{R}^{-1}\mathbf{F}\right)^{-1}\mathbf{F}'\mathbf{R}^{-1}\mathbf{y}^{n},$$

resulting in

$$\hat{y}\left(\mathbf{x}_0\right) = \mathbf{f}'\left(\mathbf{x}_0\right)\hat{\boldsymbol{\beta}} + \mathbf{r}'\left(\mathbf{x}_0\right)\mathbf{R}^{-1}\left(\mathbf{y}^{n} - \mathbf{F}\hat{\boldsymbol{\beta}}\right) = \mathbf{a}'\mathbf{y}^{n}, \tag{21.1}$$

where

$$\mathbf{a} = \mathbf{R}^{-1}\left[\begin{array}{c} \mathbf{F}\left(\mathbf{F}'\mathbf{R}^{-1}\mathbf{F}\right)^{-1}\mathbf{f}\left(\mathbf{x}_0\right) \\ + \left(\mathbf{I}_n - \mathbf{F}\left(\mathbf{F}'\mathbf{R}^{-1}\mathbf{F}\right)^{-1}\mathbf{F}'\mathbf{R}^{-1}\right)\mathbf{r}\left(\mathbf{x}_0\right) \end{array}\right].$$

Note that $\hat{y}\left(\mathbf{x}_0\right)$ is linear in the data and is unbiased: $E\left[\hat{y}\left(\mathbf{x}_0\right)\right] = E\left[Y\left(\mathbf{x}_0\right)\right]$. The correction term in (21.1) (correcting the naive predictor $\mathbf{f}'\left(\mathbf{x}_0\right)\hat{\boldsymbol{\beta}}$, whose expectation is $E\left[Y\left(\mathbf{x}_0\right)\right]$) is a linear combination of the residuals in the training set.

- Without the assumption of normality, and using only the mean and covariance functions given above, this is the linear function of the data which minimizes the prediction variance, among all linear unbiased predictors, i.e. is BLUP (assigned).

- Example: If $\mathbf{f}'(\mathbf{x})\boldsymbol{\beta} = \beta_0$ (location only) then

$$\hat{y}(\mathbf{x}_0) = \hat{\beta}_0 + \mathbf{r}'(\mathbf{x}_0)\mathbf{R}^{-1}\left(\mathbf{y}^n - \hat{\beta}_0\mathbf{1}_n\right),$$

where

$$\hat{\beta}_0 = wlse = \frac{\mathbf{1}_n'\mathbf{R}^{-1}\mathbf{y}^n}{\mathbf{1}_n'\mathbf{R}^{-1}\mathbf{1}_n}.$$

- Unknown correlation parameters $\boldsymbol{\theta}$ are typically replaced by estimates, or merely given assumed values. OR, they can be given prior distributions and then a further averaging of $\hat{y}(\mathbf{x}_0|\boldsymbol{\theta})$ is done to obtain the 'Fully Bayesian' predictor. If instead one averages over a distribution of $\hat{\boldsymbol{\theta}}$ after computing $\hat{y}\left(\mathbf{x}_0|\hat{\boldsymbol{\theta}}\right)$ then one obtains the 'empirical Bayes' predictor.

- The most common way to estimate correlation parameters is by maximum likelihood, assuming a Normal likelihood. In that case the likelihood for $\left(\boldsymbol{\beta},\boldsymbol{\theta},\sigma_Z^2\right)$ is

$$p\left(\mathbf{y}^n|\boldsymbol{\beta},\boldsymbol{\theta},\sigma_Z^2\right)$$
$$\propto \left(\begin{array}{c} \left|\sigma_Z^2\mathbf{R}\right|^{-1/2} \\ \cdot \exp\left\{-\frac{1}{2\sigma_Z^2}\left(\mathbf{y}^n-\mathbf{F}\boldsymbol{\beta}\right)'\mathbf{R}^{-1}\left(\mathbf{y}^n-\mathbf{F}\boldsymbol{\beta}\right)\right\} \end{array} \right),$$

with log likelihood

$$l\left(\boldsymbol{\beta},\boldsymbol{\theta},\sigma_Z^2\right) = -\frac{1}{2}\left\{ \begin{array}{c} n\log\sigma_Z^2+\log\left|\mathbf{R}\left(\boldsymbol{\theta}\right)\right| \\ +\frac{\left(\mathbf{y}^n-\mathbf{F}\boldsymbol{\beta}\right)'\mathbf{R}^{-1}\left(\mathbf{y}^n-\mathbf{F}\boldsymbol{\beta}\right)}{\sigma_Z^2} \end{array} \right\}.$$

The mle of $\boldsymbol{\beta}$ is the GLSE $\hat{\boldsymbol{\beta}}\left(\boldsymbol{\theta}\right)$, and then

$$l\left(\hat{\boldsymbol{\beta}},\boldsymbol{\theta},\sigma_Z^2\right) = -\frac{1}{2}\left\{ \begin{array}{c} n\log\sigma_Z^2+\log\left|\mathbf{R}\left(\boldsymbol{\theta}\right)\right| \\ +\frac{\left(\mathbf{y}^n-\mathbf{F}\hat{\boldsymbol{\beta}}\right)'\mathbf{R}^{-1}\left(\mathbf{y}^n-\mathbf{F}\hat{\boldsymbol{\beta}}\right)}{\sigma_Z^2} \end{array} \right\}.$$

The mle of σ_Z^2 is

$$\begin{aligned}\hat{\sigma}_Z^2(\boldsymbol{\theta}) &= \frac{\left(\mathbf{y}^n - \mathbf{F}\hat{\boldsymbol{\beta}}\right)' \mathbf{R}^{-1} \left(\mathbf{y}^n - \mathbf{F}\hat{\boldsymbol{\beta}}\right)}{n} \\ &= \frac{\mathbf{y}^{n\prime} \begin{bmatrix} \mathbf{R}^{-1} \\ -\mathbf{R}^{-1}\mathbf{F}\left(\mathbf{F}'\mathbf{R}^{-1}\mathbf{F}\right)^{-1}\mathbf{F}'\mathbf{R}^{-1} \end{bmatrix} \mathbf{y}^n}{n}\end{aligned}$$

and then

$$l\left(\hat{\boldsymbol{\beta}}, \boldsymbol{\theta}, \hat{\sigma}_Z^2\right) = -\frac{1}{2}\left\{n \log \hat{\sigma}_Z^2(\boldsymbol{\theta}) + \log|\mathbf{R}(\boldsymbol{\theta})| + n\right\}.$$

The mle of $\boldsymbol{\theta}$ is then

$$\hat{\boldsymbol{\theta}} = \arg\min\left\{n \log \hat{\sigma}_Z^2(\boldsymbol{\theta}) + \log|\mathbf{R}(\boldsymbol{\theta})|\right\},$$

- Example and summary - Santner slides pp. 52-54.

22. Computer experiments - design

- Framework: we seek points $\mathbf{x}_1, ..., \mathbf{x}_n \in \chi \subset \mathbb{R}^d$ at which deterministic outputs $y(\mathbf{x}_i)$ will be produced by a computer code.

- No replication (no randomness!), fewer runs is better, and the design should cover χ reasonably well (at least for an exploratory design, when we know little about the true response).

- The RSM designs we have looked at violate this last point, and so might not be suitable for computer experiments (or others where the true response might differ from that which is being fitted - more on this later).

- In what follows, assume $\chi = [0, 1]^d$. Possible 'space-filling' strategies:

- randomly sample n points from χ (good 'on average'),
- stratified sampling,
- Latin Hypercube design of n points in χ.

- Constructing a Latin Hypercube design:
  - Divide the j^{th} margin [0, 1] into n equal subintervals, randomly choose a point from each (or use the midpoints). This yields a 'uniform' design $x_{1j}, ..., x_{nj}$ for each margin $j = 1, ..., d$.
  - Form the array

$$\begin{array}{cccc} x_{11} & x_{12} & \cdots & x_{1d} \\ x_{21} & x_{22} & \cdots & x_{2d} \\ & & \vdots & \\ x_{n1} & x_{n2} & \cdots & x_{nd} \end{array}$$

    and then randomly permute the elements in each column (independently). The n rows so obtained define the LHS.

  - The projections on the coordinate axes remain uniform, in the sense used above.

  - See Notz slides pp. 18-21, p. 27.

- There is an optimality property: if the purpose is to estimate

$$\int_{\chi} G\left(y\left(\mathbf{x}\right)\right) d\mathbf{x}$$

  for some function G, then the estimate

$$\frac{1}{n}\sum_{i=1}^{n} G\left(y\left(\mathbf{x}_i\right)\right)$$

  typically (more precisely, if $G \circ y$ is monotonic in each coordinate) has a smaller variance than than the sample mean of a SRS. However, more often the goal is prediction of $y\left(\mathbf{x}\right)$ at previously unobserved inputs.

- See the Handbook chapter 'Latin Hypercubes and Space-filling Designs' (course web site) for many more details.

- Minimax and Maximin Designs.

  - Goal is to minimize the maximum distance between $\{\mathbf{x}_1, ..., \mathbf{x}_n\}$ and all other points in χ: if $\xi = \{\mathbf{x}_1, ..., \mathbf{x}_n\}$ is the design, and if

$$\phi(\mathbf{x};\xi) = \min_{i=1,...,n} d(\mathbf{x}_i, \mathbf{x})$$

    is the distance between $\{\mathbf{x}_1, ..., \mathbf{x}_n\}$ and $\mathbf{x}$, then a minimax design ξ_* has

$$\max_{\mathbf{x}} \phi(\mathbf{x};\xi_*) \leq \max_{\mathbf{x}} \phi(\mathbf{x};\xi)$$

    for any other design ξ. A maximin design ξ_{**} has

$$\min_{\mathbf{x}} \phi(\mathbf{x};\xi_{**}) \geq \min_{\mathbf{x}} \phi(\mathbf{x};\xi)$$

    for any other design ξ. Both aim to 'spread out' the design points in a uniform-like manner.

  - The latter seem to be much easier to find, and software is available to do this. See Notz slides pp. 35, 36.

- Although the designs looked at so far aim for a kind of uniformity, this has not been rigorously required. An optimal Uniform Design (UD) will minimize some measure of discrepancy from uniformity. Specifically, let

$$D_{\mathbf{x}} = [0, x_1] \times \cdots \times [0, x_d]$$

be the d-dimensional rectangle with 'corner' $\mathbf{x}$, and for a design ξ let

$$\begin{aligned} F_n(\mathbf{x}) &= \frac{1}{n}\sum_{i=1}^{n} I(\mathbf{x}_i \in D_{\mathbf{x}}) \\ &= \text{ proportion of the design points } \mathbf{X} \text{ with} \\ &\quad X_1 \leq x_1 \text{ and } \cdots \text{ and } X_d \leq x_d \end{aligned}$$

be the e.d.f. of the design (example: if $d = 1$ then ...). The L_p discrepancy of the design is

$$\left(\int_{\chi} |F_n(\mathbf{x}) - F(\mathbf{x})|^p \, d\mathbf{x}\right)^{\frac{1}{p}},$$

where $F(\mathbf{x})$ is the uniform d.f. on χ. (If $p = \infty$ this is $\sup_{\mathbf{x}} |F_n(\mathbf{x}) - F(\mathbf{x})|$, the 'star' discrepancy).

In one dimension the design on $[0, 1]$ which minimizes the star discrepancy [and which, as is not hard to show, is a minimax design] is

$$x_i = \frac{i - .5}{n}, i = 1, ..., n.$$

Proof: For any n-point design with ordered design points t_i, the design measure is a step function:

$$\sup_x |F_n(x) - F(x)| = \sup_x |F_n(x) - x|$$

$$= \max_{1 \le i \le n} \max \left\{ \left| t_i - \frac{i-1}{n} \right|, \left| t_i - \frac{i}{n} \right| \right\} \stackrel{??}{\ge} \frac{.5}{n}.$$

Why? The inner max exceeds its average, which exceeds $.5/n$ (equality if $t_i \in \left[\frac{i-1}{n}, \frac{i}{n}\right]$, strict inequality otherwise). The lower bound is attained if $t_i = x_i$. $\square$

These are the margins of a Latin Hypercube design as described above, if the midpoints of the intervals are used.

See Notz slides p. 43 for a design minimizing the star discrepancy in two dimensions.

- An appealing class of space filling designs is given by Sobol sequences. These have the property that a design of size n is obtained by adding one point to a design of size $n-1$; thus they are attractive for sequential experiments. See Notz slides pp. 61-65.

- Other criteria: if $\hat{y}(\mathbf{x})$ is the predictor and

$$\text{mse}\left[\hat{y}(\mathbf{x})\right] = E\left[\left\{\hat{y}(\mathbf{x}) - y(\mathbf{x})\right\}^2\right]$$

we might:

  - aim to minimize the integrated mean squared error (IMSE)

$$\text{imse} = \int_{\chi} \text{mse}\left[\hat{y}(\mathbf{x})\right] w(\mathbf{x})\, d\mathbf{x},$$

    where $w(\mathbf{x})$ is a weight function (a probability density on χ);

  - aim to minimize

$$\max_{\mathbf{x}\in\chi} \text{mse}\left[\hat{y}(\mathbf{x})\right]$$

    (less popular, in my experience).

- The mse of the BLUP $\hat{y}(\mathbf{x}) = \mathbf{a}'(\mathbf{x})\,\mathbf{y}^n$ is (assigned)

$$\text{mse}\left[\hat{y}(\mathbf{x})\right] = \sigma_Z^2 \left\{ 1 - \left(\begin{pmatrix} \mathbf{f}(\mathbf{x}) \\ \mathbf{r}(\mathbf{x}) \end{pmatrix}' \begin{pmatrix} \mathbf{0}_{k\times k} & \mathbf{F}' \\ \mathbf{F} & \mathbf{R}_{n\times n} \end{pmatrix}^{-1} \begin{pmatrix} \mathbf{f}(\mathbf{x}) \\ \mathbf{r}(\mathbf{x}) \end{pmatrix} \right) \right\}$$

  Thus the imse is

$$\int_\chi \text{mse}\left[\hat{y}(\mathbf{x})\right] w(\mathbf{x})\, d\mathbf{x} = \sigma_Z^2 \text{ times } \{1 \text{ minus } \text{tr}\left[\begin{pmatrix} \mathbf{0} & \mathbf{F}' \\ \mathbf{F} & \mathbf{R} \end{pmatrix}^{-1} \times \int_\chi \begin{pmatrix} \mathbf{f}(\mathbf{x})\,\mathbf{f}'(\mathbf{x}) & \mathbf{f}(\mathbf{x})\,\mathbf{r}'(\mathbf{x}) \\ \mathbf{r}(\mathbf{x})\,\mathbf{f}'(\mathbf{x}) & \mathbf{r}(\mathbf{x})\,\mathbf{r}'(\mathbf{x}) \end{pmatrix} w(\mathbf{x})\, d\mathbf{x} \right] \},$$

  with $\mathbf{r}(\mathbf{x})$, $\mathbf{R}$ and $\mathbf{F}$ depending on the design.

- One might use an **exchange algorithm** to minimize this loss - this is an iterative method where, at each iteration, one design point is replaced by another which decreases the loss. Or one can proceed **sequentially** – given an n-point design, evaluate the loss resulting from adding one more point

$\mathbf{x}_*$; choose the minimizing $\mathbf{x}_*$ as the $(n+1)^{th}$ point. This is a natural approach since it allows one to re-estimate the model at each stage.

- Summary - Notz slides pp. 67-69.

- Simulated annealing? - will be discussed later. Genetic algorithms? - will discuss if there is time.

- **Genetic algorithm**: Suppose the design space consists of N points, and a design of size n is to be chosen. Generate a random 'population' of 40 designs; compute the loss $\mathbb{L}_j$, $j = 1, ..., 40$. This is the first 'generation'.

- Assign a 'fitness level' to each; normalize to get a probability distribution:

$$\begin{aligned} \text{fitness}_j &= \frac{1}{\left(\mathbb{L}_j - .999 \min_k \mathbb{L}_j\right)^2}; \\ \psi_j &= \frac{\text{fitness}_j}{\sum_j \text{fitness}_j}. \end{aligned}$$

- Form the next generation:
  - The best (fittest) two in the current generation always survive to the next.
  - Choose pairs of designs ('parents') from the current generation – $P\left(\text{choose } j^{th}\right) = \psi_j$ – and combine them to form a 'child'. Continue until a new generation of size 40 has been generated.

- Parents combine by 'crossover' with probability $P_{crossover}$ $(= .9)$; with probability $1 - P_{crossover}$ the child is identical to the fittest parent.

- Example of crossover: represent a design by a vector of n ones and $N - n$ zeros; e.g. $n = 3, N = 5$; parents $\{2, 3, 4\}$ and $\{1, 3, 5\}$:

$$\begin{pmatrix} 0 \\ 1 \\ 1 \\ 1 \\ 0 \end{pmatrix} + \begin{pmatrix} 1 \\ 0 \\ 1 \\ 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \\ 2 \\ 1 \\ 1 \end{pmatrix} \overset{random}{\rightarrow} \begin{pmatrix} 1 \\ 0 \\ 2 \\ 1 \\ 0 \end{pmatrix} \overset{2\rightarrow 1}{\rightarrow} \begin{pmatrix} 1 \\ 0 \\ 1 \\ 1 \\ 0 \end{pmatrix}.$$

- Mutation: In each child, randomly choose a one and a zero, with probability $P_{mutation}$ ($= .05$) swap them. Repeat n times with each child.

- Repeat, forming new generations and evaluating their fitnesses until the best design has not changed for 200 consecutive generations.

Part IX

OPTIMAL DESIGN

Source materials on course web site:

1. Robustness of Design (D. Wiens)

2. Optimal Designs for Non-Linear Models: Introduction and Historical Overview (D. Wiens)

3. A Basis for the Selection of a Response Surface Design (Box & Draper 1959)

4. Robust Discrimination Designs (D. Wiens)

References 1 and 2 are chapters of an upcoming Handbook on Design and Analysis of Experiments.

23. Concepts of classical optimal design

- An experimenter anticipates making observations at various values $\{\mathbf{x}_i\}_{i=1}^n$ of vector-valued predictor variables chosen from a 'design space' χ, and then observing the corresponding values of a dependent variable Y. The predictors affect $E[Y]$ through a vector $\mathbf{f}(\mathbf{x})$ of regressors:

$$Y_i = \mathbf{f}'(\mathbf{x}_i)\,\boldsymbol{\theta} + \varepsilon_i.$$

- The errors are i.i.d. with mean zero. Normality plays only a limited role here – the measures of performance typically depend only on the first two moments of the estimates.

- Let ξ be the 'design measure'. This is sometimes viewed as a p.m.f.: if the design places n_i out of a total of n observations at the point $\mathbf{x}_i \in \chi$

then $\xi(\mathbf{x}_i) = n_i/n$. Sometimes as a probability measure on χ, placing mass $\xi(\mathbf{x}_i)$ at $\mathbf{x}_i$: for instance

$$\begin{aligned}\frac{1}{n}\mathbf{X}'\mathbf{X} &= \sum \frac{n_i}{n}\mathbf{f}(\mathbf{x}_i)\mathbf{f}'(\mathbf{x}_i) \\ &= \sum \xi(\mathbf{x}_i)\mathbf{f}(\mathbf{x}_i)\mathbf{f}'(\mathbf{x}_i) \\ &= \int_\chi \mathbf{f}(\mathbf{x})\mathbf{f}'(\mathbf{x})\xi(d\mathbf{x});\end{aligned}$$

we call this matrix $\mathbf{M}(\xi)$. The final form – think of the integral as $E_\xi[\mathbf{f}(\mathbf{x})\mathbf{f}'(\mathbf{x})]$ if it helps – makes it convenient to think of approximating a design by something with a measure ξ whose masses are not necessarily multiples of n^{-1} – perhaps even a continuous measure.

- Some common optimality principles (all following from $\text{cov}\left[\hat{\boldsymbol{\theta}}\right] \propto \mathbf{M}^{-1}(\xi)$:
  - D-optimality – minimize $\det\left(\mathbf{M}^{-1}(\xi)\right)$. ($\propto$ volume of a confidence ellipsoid on regression parameters)

- A-optimality – minimize $tr\left(\mathbf{M}^{-1}(\xi)\right)$. ($\propto$ average variance of regression parameters)

- E-optimality – minimize $ch_{\max}\mathbf{M}^{-1}(\xi)$. ($\propto$ maximum variance of a normalized linear combination $\mathbf{c}'\hat{\boldsymbol{\theta}}$)

- G-optimality – minimize

$$\bar{d}(\xi) = \sup_{\mathbf{x}\in\chi} \mathbf{f}'(\mathbf{x})\,\mathbf{M}^{-1}(\xi)\,\mathbf{f}(\mathbf{x}).$$

($\propto$ maximum variance of a prediction $\hat{Y}(\mathbf{x}) = \mathbf{f}'(\mathbf{x})\,\hat{\boldsymbol{\theta}}$)

- **Theorem** (General Equivalence Theorem): In a p-parameter regression model, a design ξ_0 is D-optimal iff it is also G-optimal, and then $\bar{d}(\xi_0) = p$.
  Proof: For $t \in [0,1]$ and any design ξ_1 define

$$\xi_t = (1-t)\,\xi_0 + t\xi_1.$$

  Then ξ_0 is D-optimal iff the function

$$\phi(t) = -\log|\mathbf{M}(\xi_t)|$$

is minimized at $t = 0$, for every ξ_1. We need two preliminary results; each is proven in the STAT 766 notes on my web site.

- **Lemma 1**: For p.s.d. matrices $\mathbf{A}$ and $\mathbf{B}$ and for $t \in [0, 1]$ we have that

$$\begin{aligned} & -\log|(1-t)\mathbf{A} + t\mathbf{B}| \\ \leq\ & (1-t)(-\log|\mathbf{A}|) + t(-\log|\mathbf{B}|). \end{aligned}$$

- **Lemma 2**: For a p.d. matrix $\mathbf{A}$ the matrix of partial derivatives of $\log|\mathbf{A}|$ with respect to the elements of $\mathbf{A}$ is $\mathbf{A}'^{-1}$.

By Lemma 2,

$$\begin{aligned} \phi'(t) &= \sum_{i,j} \frac{\partial - \log|\mathbf{M}(\xi_t)|}{\partial [\mathbf{M}(\xi_t)]_{ij}} \frac{d[\mathbf{M}(\xi_t)]_{ij}}{dt} \\ &= \sum_{i,j} -[\mathbf{M}(\xi_t)]^{ji} \int_{\mathcal{X}} \left[\mathbf{f}(\mathbf{x})\mathbf{f}'(\mathbf{x})\right]_{ij} (\xi_1 - \xi_0)(d\mathbf{x}) \\ &= -tr\left\{\mathbf{M}^{-1}(\xi_t) \int_{\mathcal{X}} \left[\mathbf{f}(\mathbf{x})\mathbf{f}'(\mathbf{x})\right](\xi_1 - \xi_0)(d\mathbf{x})\right\} \\ &= -\int_{\mathcal{X}} \mathbf{f}'(\mathbf{x})\mathbf{M}^{-1}(\xi_t)\mathbf{f}(\mathbf{x})(\xi_1 - \xi_0)(d\mathbf{x}). \end{aligned}$$

By Lemma 1, for any $\mathbf{M}(\xi_0) = \mathbf{A}$ and $\mathbf{M}(\xi_1) = \mathbf{B}$ the function $\phi(t) = -\log|\mathbf{M}(\xi_t)|$ is convex, and so is minimized at $t = 0$ iff

$$0 \leq \phi'(0) = -\int_{\chi} \mathbf{f}'(\mathbf{x})\,\mathbf{M}^{-1}(\xi_0)\,\mathbf{f}(\mathbf{x})\,(\xi_1 - \xi_0)\,(d\mathbf{x})\,. \tag{23.1}$$

We always have (*how?*)

$$p = \int_{\chi} \mathbf{f}'(\mathbf{x})\,\mathbf{M}^{-1}(\xi)\,\mathbf{f}(\mathbf{x})\,\xi\,(d\mathbf{x}) \leq \bar{d}(\xi)\,;$$

this in (23.1) shows that ξ_0 is D-optimal iff

$$\forall \xi_1 : \int_{\chi} \mathbf{f}'(\mathbf{x})\,\mathbf{M}^{-1}(\xi_0)\,\mathbf{f}(\mathbf{x})\,\xi_1\,(d\mathbf{x}) \leq p. \tag{23.2}$$

If (23.2) holds then $\mathbf{f}'(\mathbf{x})\,\mathbf{M}^{-1}(\xi_0)\,\mathbf{f}(\mathbf{x}) \leq p$ for all $\mathbf{x} \in \chi$, so that $\bar{d}(\xi_0) \leq p$; thus a D-optimal design ξ_0 has

$$\bar{d}(\xi_0) = p \leq \bar{d}(\xi_1)\,,$$

for any other design ξ_1, and so is G-optimal as well.

Conversely, suppose that ξ_0 is G-optimal. Then $p \leq \bar{d}(\xi_0) \leq \bar{d}(\xi_1)$ for any other design ξ_1; in particular, with $\xi_1 = \xi_D$ (= the D-optimal design, which always exists) we obtain $\bar{d}(\xi_0) = p$. Thus for *any* ξ_1,

$$\int_{\chi} \mathbf{f}'(\mathbf{x}) \mathbf{M}^{-1}(\xi_0) \mathbf{f}(\mathbf{x}) \xi_1(d\mathbf{x}) \leq \bar{d}(\xi_0) = p,$$

so that (23.2) holds and ξ_0 is D-optimal. □

- **A sufficient condition**: This derivation brings out that a design ξ_0 is D- or G-optimal if it places all of its mass where $\mathbf{f}'(\mathbf{x}) \mathbf{M}^{-1}(\xi_0) \mathbf{f}(\mathbf{x}) = p$, and at all other points we have $\mathbf{f}'(\mathbf{x}) \mathbf{M}^{-1}(\xi_0) \mathbf{f}(\mathbf{x}) < p$. (Then $\bar{d}(\xi_0) = p$ and so (23.2) holds.) This is used to check for optimality.

- Suppose that a design ξ_0 has exactly p points of support $\{\mathbf{x}_1, ..., \mathbf{x}_p\}$. If it is to be D-optimal, then it is necessary to have $\xi_0(\mathbf{x}_i) = 1/p$ for $i = 1, ..., p$. [Thus, to construct a p-point D-optimal design we choose p points and put mass $1/p$ at

each of them; if the sufficient condition above holds then we have chosen the right p points.]

Proof: For a p-point design we have

$$\mathbf{M}(\xi) = \sum_{i=1}^{p} \xi(\mathbf{x}_i)\,\mathbf{f}(\mathbf{x}_i)\,\mathbf{f}'(\mathbf{x}_i) = \mathbf{F}'\mathbf{D}_\xi\mathbf{F},$$

where $\mathbf{F}_{p\times p}$ has rows $\mathbf{f}'(\mathbf{x}_i)$ and

$$\mathbf{D}_\xi = diag\left(\xi(\mathbf{x}_1), ..., \xi(\mathbf{x}_p)\right).$$

Thus

$$\begin{aligned} -\log|\mathbf{M}(\xi)| &= -\log|\mathbf{F}|^2 - \sum_{i=1}^{p}\log\xi(\mathbf{x}_i) \\ &= -\log|\mathbf{F}|^2 + p\cdot aver\left\{-\log\xi(\mathbf{x}_i)\right\}. \end{aligned}$$

By Jensens' Inequality

$$aver\left\{-\log\xi(\mathbf{x}_i)\right\} \geq -\log\left(aver\left\{\xi(\mathbf{x}_i)\right\}\right) = \log p,$$

with equality iff $\xi(\mathbf{x}_i) = 1/p$ for $i = 1, ..., p$. $\square$

- **Example**: (Assigned) For quadratic regression on $[-1, 1]$ the design with mass $1/3$ at each of $0, \pm 1$ is D- and G-optimal.

- In the situation as above, is a p-point design desirable?

- There are similar 'equivalence theorems' in other design situations.

- **Computing**. Many possibilities; one simple one is as follows. Start with any design; at the k^{th} step of the algorithm, if the current design is ξ_k then find $\mathbf{x}_{(k)}$ maximizing $q_k(\mathbf{x}) = \mathbf{f}'(\mathbf{x})\,\mathbf{M}^{-1}(\xi_k)\,\mathbf{f}(\mathbf{x})$ and put

$$\xi_{k+1} = \frac{k}{k+1}\xi_k + \frac{1}{k+1}\delta_{\mathbf{x}_{(k)}}.$$

  This algorithm can be shown to converge to a D-optimal design.

  - The interpretation is that, given a design with k points (not necessarily distinct), we add one more point $\mathbf{x}_{(k)} = \arg\max q_k(\mathbf{x})$ to get a $(k+1)$-point design; then iterate.

– R-code for construction of a design for SLR on $[0, 1]$ is on the course web site.

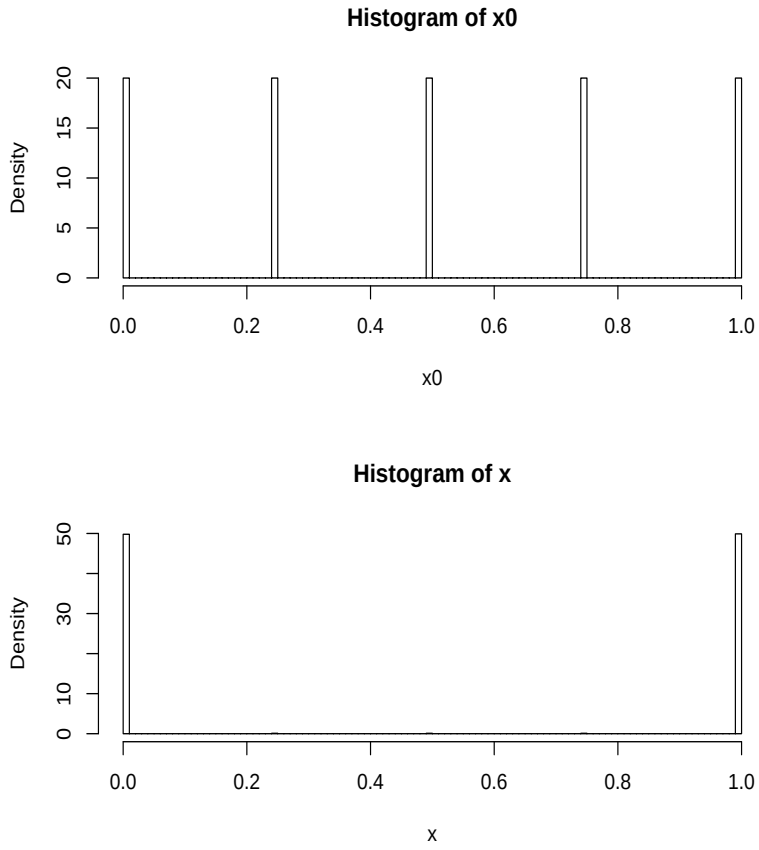

Start with 5 points equally spaced; apply algorithm to generate a total of 1000 points. The points added were all either 0 or 1.

24. Model-robust design

- Why robustness of design?
  - The models for which we construct designs are generally at best approximations.
  - The 'best' design for a slightly wrong model can be much more than slightly sub-optimal.
  - Although we will fit the assumed, 'ideal' model, we should design for protection against biases arising from any of a range ('neighbourhood') of alternate models.
  - We look at mse, not merely variance. (Maximum mse, as the model varies.)

- Seminal paper – Box and Draper (1959). Design for a polynomial or multinomial response when the class of alternatives consists of all such functions of a given higher degree.
  - Example: Fit a straight line in an independent variable $x \in [-1, 1]$.
  - Mean response $E[Y(x)] = \phi_0 + \phi_1 x + \phi_2 x^2$ for parameters ϕ.
  - The lses are biased.

George Edward Pelham Box FRS (18 October 1919 – 28 March 2013)

‘ ... all models are wrong, but some are useful.’

- ‘... the optimal design in typical situations in which both variance and bias occur is very nearly the same as would be obtained if *variance were ignored completely* and the experiment designed so as to *minimize the bias alone*.’ (Box and Draper 1959, p. 622.)

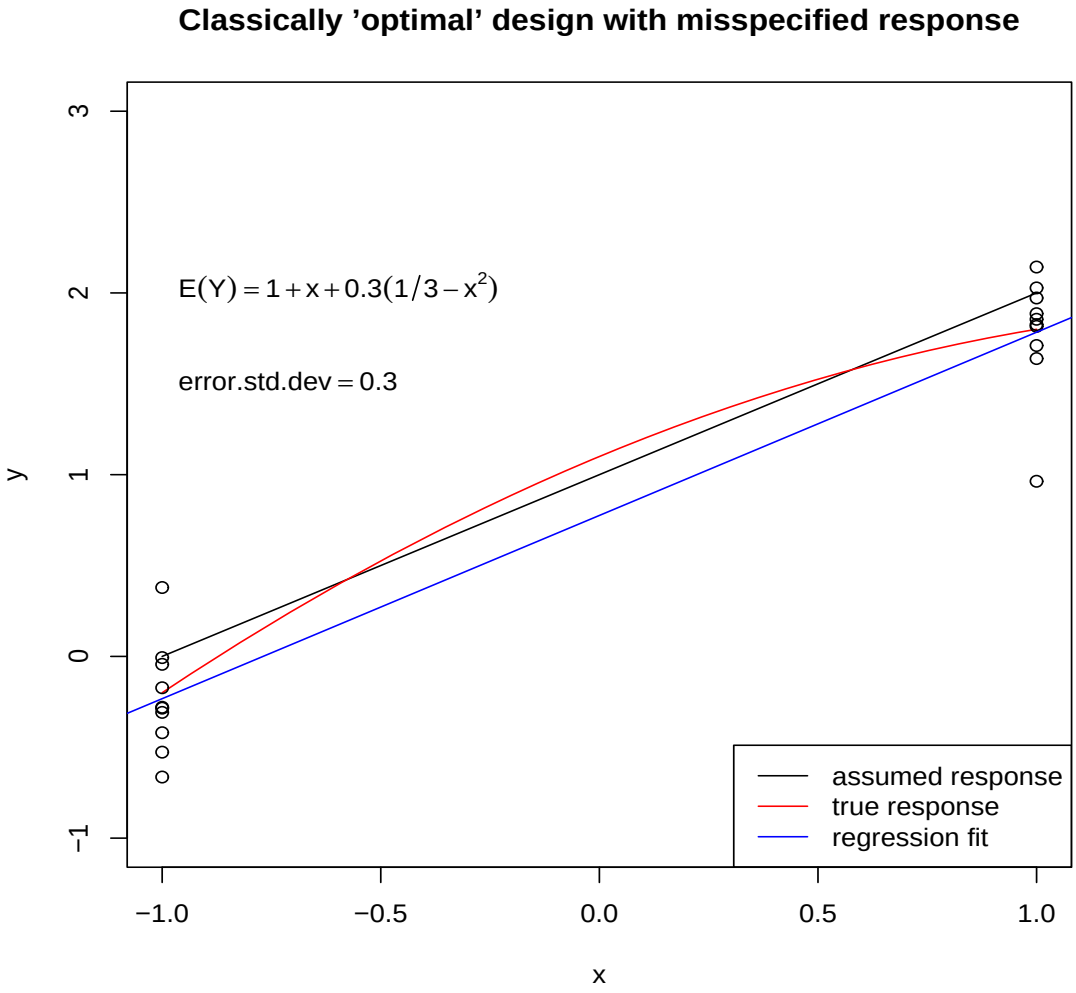

Assumed response linear in x but in fact $E[Y] = 1 + x + 0.3\left(1/3 - x^2\right)$. Parameters of assumed response $\beta_0 + \beta_1 x$ estimated by least squares.

- Generalize: Fit the model $E\left[Y\left(\mathbf{x}\right)\right] = \mathbf{f}'\left(\mathbf{x}\right)\boldsymbol{\theta}$, for a p-vector $\mathbf{f}\left(\mathbf{x}\right)$ of regressors, each element of which is a function of q independent variables $\mathbf{x} = \left(x_1, ..., x_q\right)'$, with $\mathbf{x}$ to be chosen from a *design space* χ. Here we take a *finite* design space

$$\chi = \left\{\mathbf{x}_1, ..., \mathbf{x}_N\right\}.$$

  Entertain alternatives

$$E\left[Y\left(\mathbf{x}\right)\right] = \mathbf{f}'\left(\mathbf{x}\right)\boldsymbol{\theta} + \psi\left(\mathbf{x}\right),$$

  for some function ψ. But now what does $\boldsymbol{\theta}$ mean? It is not identifiable: in the alternate models one might equally well write

$$E\left[Y\left(\mathbf{x}\right)\right] = \mathbf{f}'\left(\mathbf{x}\right)\left(\boldsymbol{\theta} + \boldsymbol{\phi}\right) + \left(\psi\left(\mathbf{x}\right) - \mathbf{f}'\left(\mathbf{x}\right)\boldsymbol{\phi}\right)$$

  for arbitrary $\boldsymbol{\phi}$. So *define* the target parameter to be that which makes the fitted response most accurate:

$$\boldsymbol{\theta} = \arg\min_{\eta} \sum_{i=1}^{N} \left(E\left[Y\left(\mathbf{x}_i\right)\right] - \mathbf{f}'\left(\mathbf{x}_i\right)\eta\right)^2,$$

and then *define* $\psi(\mathbf{x}) = E[Y(\mathbf{x})] - \mathbf{f}'(\mathbf{x})\boldsymbol{\theta}$. This leads (*how?*) to the orthogonality requirement

$$\sum_{i=1}^{N} \mathbf{f}(\mathbf{x}_i)\psi(\mathbf{x}_i) = \mathbf{0}. \tag{24.1}$$

- Let ξ be the design, placing mass $\xi_i = \xi(\mathbf{x}_i) = n_i/n$ at $\mathbf{x}_i$. Define

$$\begin{aligned}
\mathbf{M}_\xi &= \sum_{i=1}^{N} \xi_i \mathbf{f}(\mathbf{x}_i)\mathbf{f}'(\mathbf{x}_i), \\
\mathbf{b}_{\psi,\xi} &= \sum_{i=1}^{N} \xi_i \mathbf{f}(\mathbf{x}_i)\psi(\mathbf{x}_i), \\
\mathbf{A} &= \frac{1}{N}\sum_{i=1}^{N} \mathbf{f}(\mathbf{x}_i)\mathbf{f}'(\mathbf{x}_i)
\end{aligned}$$

and assume that $\mathbf{M}_\xi$ is invertible. The covariance matrix of the *lse* $\hat{\boldsymbol{\theta}}$ is $\left(\sigma_\varepsilon^2/n\right)\mathbf{M}_\xi^{-1}$, the bias is

$$E\left[\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\right] = \mathbf{M}_\xi^{-1}\mathbf{b}_{\psi,\xi}, \tag{24.2}$$

and the loss, taken to be **average** mse, is

$$\begin{aligned}\text{amse} &= \frac{1}{N}\sum_{i=1}^{N} E\left[\left\{\hat{Y}\left(\mathbf{x}_i\right) - E\left[Y\left(\mathbf{x}_i\right)\right]\right\}^2\right] \\ &= \frac{\sigma_\varepsilon^2}{n}\text{trace}\left(\mathbf{A}\mathbf{M}_\xi^{-1}\right) + \mathbf{b}'_{\psi,\xi}\mathbf{M}_\xi^{-1}\mathbf{A}\mathbf{M}_\xi^{-1}\mathbf{b}_{\psi,\xi} \\ &\quad + \frac{1}{N}\sum_{i=1}^{N}\psi^2\left(\mathbf{x}_i\right). \qquad (24.3)\end{aligned}$$

- We will find a **minimax** design, i.e. we maximize amse over all ψ satisfying (24.1) and

$$\sum_{i=1}^{N}\psi^2\left(\mathbf{x}_i\right) \le \frac{\tau^2}{n} \qquad (24.4)$$

for a constant τ^2, and then find a design minimizing this maximum loss.

- It is convenient to temporarily transform to an orthogonal basis. Let $\mathbf{Q}_1$ be an $N \times p$ matrix whose columns form an orthogonal basis for the column space of the matrix $\mathbf{F}$ with rows $\{\mathbf{f}'\left(\mathbf{x}_i\right)\}_{i=1}^N$.

Augment this by $\mathbf{Q}_2 : N \times (N-p)$ whose columns form an orthogonal basis for the orthogonal complement of this space. (What are these in R terminology?) Then $(\mathbf{Q}_1 \vdots \mathbf{Q}_2)$ is an orthogonal matrix and $\boldsymbol{\psi} = (\psi(\mathbf{x}_1), ..., \psi(\mathbf{x}_N))'$ is orthogonal to the column space of $\mathbf{Q}_1$, hence is necessarily of the form

$$\boldsymbol{\psi} = \left(\tau/\sqrt{n}\right) \mathbf{Q}_2 \mathbf{c},$$

where (by (24.4)) $\|\mathbf{c}\| \leq 1$.

- With $\mathbf{D}_\xi = diag(\xi_1, \cdots, \xi_N)$ we find that

$$\begin{aligned} \mathbf{M}_\xi &= \mathbf{F}' \mathbf{D}_{\boldsymbol{\xi}} \mathbf{F}, \\ \mathbf{A} &= \frac{1}{N} \mathbf{F}' \mathbf{F}, \\ \mathbf{b}_{\psi,\xi} &= \mathbf{F}' \mathbf{D}_{\boldsymbol{\xi}} \boldsymbol{\psi}, \end{aligned}$$

and then

$$\begin{aligned} \text{amse} &= \frac{\sigma_\varepsilon^2}{n} \text{tr}\left(\mathbf{A}\mathbf{M}_\xi^{-1}\right) \\ &+ \frac{\tau^2}{nN} \mathbf{c}' \left[\mathbf{Q}_2' \mathbf{D}_\xi \mathbf{Q}_1 \left(\mathbf{Q}_1' \mathbf{D}_\xi \mathbf{Q}_1\right)^{-2} \mathbf{Q}_1' \mathbf{D}_\xi \mathbf{Q}_2 + \mathbf{I}_{N-p}\right] \mathbf{c}. \end{aligned}$$

Carrying out the maximization over $\mathbf{c}$ and returning to the original notation (assigned) gives that, in terms of

$$K_{\boldsymbol{\xi}} = \sum_{i=1}^{N} \xi_i^2 \mathbf{f}\left(\mathbf{x}_i\right) \mathbf{f}'\left(\mathbf{x}_i\right)$$

and

$$H_{\boldsymbol{\xi}} = M_{\xi} \mathbf{A}_{\xi}^{-1} M_{\xi}$$

the maximum loss is $\left(\sigma_{\varepsilon}^2 + \tau^2\right)/n$ times

$$\mathbb{L}_{\nu}\left(\xi\right) = \left(1-\nu\right) \mathsf{tr}\left(\mathbf{A}\mathbf{M}_{\xi}^{-1}\right) + \nu ch_{\mathsf{max}}\left(K_{\boldsymbol{\xi}} {H_{\boldsymbol{\xi}}}^{-1}\right), \tag{24.5}$$

where $\nu = \tau^2/\left(\sigma_{\varepsilon}^2 + \tau^2\right)$ and ch_{max} denotes the maximum eigenvalue.

- Note that $\nu \in [0,1]$ is the only parameter that needs to be specified; it can be chosen by the experimenter according to the weight he/she places on errors due to model misspecification rather than to random variation.

- The minimization is carried out numerically. The figure gives an example, with $N = 40$, $n = 20$ and $\nu = .5$; this illustrates a recurring theme – a naive yet sensible, and 'near optimal' method of robustifying a design is to take the replicates in the classically optimal design for the model in question and spread these out into nearby clusters.

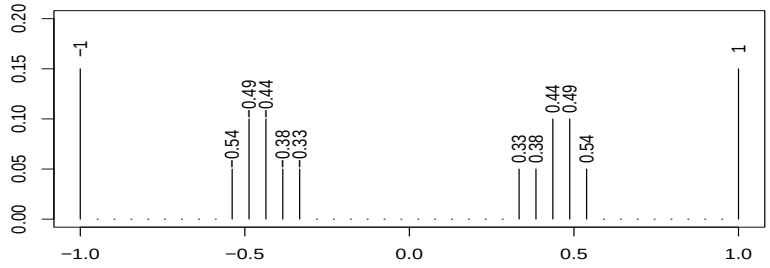

Minimax design ($n = 20$) for approximate cubic regression when bias and variance receive equal weight.

25. Nonlinear design - an example

- Experimenter wishes to discriminate between two regression models/densities:

$$\text{Model 0: } Y|\mathbf{x} \sim f_0\left(y|\mathbf{x}, \mu_0, \tau\right)$$
$$\text{Model 1: } Y|\mathbf{x} \sim f_1\left(y|\mathbf{x}, \mu_1, \tau\right)$$

  with

$$\mu_j\left(\mathbf{x}\right) = E_{F_j}\left[Y|\mathbf{x}\right]$$

  and τ represents nuisance parameters common to both densities.

- Example: Michaelis-Menten (MM) vs. Exponential (Exp) models for output (y) given concentration (x); horizontal asymptote V_j:

$$\mu_0\left(x\right) = \eta_0\left(x|\theta_0\right) = \frac{V_0 x}{K_0 + x},$$
$$\text{with } \theta_0 = \left(V_0, K_0\right)^T;$$
$$\mu_1\left(x\right) = \eta_1\left(x|\theta_1\right) = V_1\left(1 - e^{-K_1 x}\right),$$
$$\text{with } \theta_1 = \left(V_1, K_1\right)^T.$$

- If all parameters specified under each hypothesis H_0: Model 0 vs. H_1: Model 1, the Neyman - Pearson test is most powerful. Reject for large values of $R = \sum R_i$, where

$$R_i = 2 \log \left\{ \frac{f_1 \left(y_i | \mathbf{x}_i, \mu_1 \right)}{f_0 \left(y_i | \mathbf{x}_i, \mu_0 \right)} \right\}.$$

- Choose $\mathbf{x}_1, \cdots, \mathbf{x}_n$ to maximize power $\Rightarrow$ maximize

$$E_{H_1} [R] = 2 \sum_{\text{design}} \mathbb{I} \left(\mu_0 \left(\mathbf{x}_i \right), \mu_1 \left(\mathbf{x}_i \right) \right),$$

where

$$\begin{aligned} &\mathbb{I} \left(\mu_0 \left(\mathbf{x} \right), \mu_1 \left(\mathbf{x} \right) \right) \\ &= \int_{-\infty}^{\infty} f_1 \left(y | \mathbf{x}, \mu_1 \right) \log \left\{ \frac{f_1 \left(y | \mathbf{x}, \mu_1 \right)}{f_0 \left(y | \mathbf{x}, \mu_0 \right)} \right\} dy \end{aligned}$$

is the Kullback-Leibler divergence, measuring the information which is lost when f_0 is used to approximate f_1.

- Example: If, under each model, the data are normally distributed with means $\mu_j(\mathbf{x}) = \eta_j\left(\mathbf{x}|\theta_j\right)$ then KL divergence is

$$\mathbb{I}\left(\mu_0(\mathbf{x}), \mu_1(\mathbf{x})\right) = \left\{\eta_1(\mathbf{x}|\theta_1) - \eta_0(\mathbf{x}|\theta_0)\right\}^2$$

and an approach to designing is to choose the design, for n observations, to maximize

$$\sum \left\{\eta_1(\mathbf{x}_i|\theta_1) - \eta_0(\mathbf{x}_i|\theta_0)\right\}^2.$$

(Atkinson & Fedorov, 1975 - 'T-optimality'.)

  - First minimize over θ_1, θ_0, then maximize over design; or estimate/design sequentially.

- f_j log-normal, gamma, etc. - $\mathbb{I}(\cdot,\cdot)$ more complicated.

- That $E[Y|\mathbf{x}] = \eta_j\left(\mathbf{x}|\theta_j\right)$ for a known $\eta_j(\cdot)$ is typically only a convenient approximation; this suggests taking a robust approach to the design problem.

- Finite design space $\chi = \{\mathbf{x}_i\}_{i=1}^{N}$.

- Suppose that $E[Y|\mathbf{x}]$ is approximated by $\eta_j\left(\mathbf{x}|\theta_j\right)$ in Model j. For identifiability, *define*

$$\theta_j = \arg\min_{\theta} \sum_{\mathbf{x}\in\chi} \left\{E[Y|\mathbf{x}] - \eta_j(\mathbf{x}|\theta)\right\}^2, \qquad (*)$$

  and $\psi_j(\mathbf{x}) = E[Y|\mathbf{x}] - \eta_j\left(\mathbf{x}|\theta_j\right)$. Condition (*) implies that the 'residuals' $\boldsymbol{\psi}_j = \left(\psi_j(\mathbf{x}_1), ..., \psi_j(\mathbf{x}_N)\right)'$ are orthogonal to the 'regressors' - the columns of

$$\mathbf{U}_j = \frac{\partial \boldsymbol{\eta}_j\left(\theta_j\right)}{\partial \theta_j},$$

  i.e. $\mathbf{U}_j'\boldsymbol{\psi}_j = \mathbf{0}$.

- Thus the means are 'contaminated':

  $\mu_j(\mathbf{x}) = \eta_j\left(\mathbf{x}|\theta_j\right) + \psi_j(\mathbf{x}) =$ *fitted + model error.*

  Define neighbourhoods of the $\eta_j(\cdot)$ by

$$\mathbb{M}_j = \left\{\eta_j\left(\cdot|\theta_j\right) + \psi_j(\cdot) \,\middle|\, \mathbf{U}_j'\boldsymbol{\psi}_j = \mathbf{0}, \left\|\boldsymbol{\psi}_j\right\| \leq \tau_j\right\}.$$

- **Example**: Normal data. Let $\xi = (n_1/n, \cdots, n_N/n)'$, where $n_i \geq 0$ is the number of observations to be made at $\mathbf{x}_i$, and $n = \sum_{i=1}^{N} n_i$. The *robust T-optimal design* problem is that of determining

$$\xi^* = \arg\max_{\xi} \min_{\psi_1, \psi_0} \sum_{i=1}^{N} \left\{ \mu_1(\mathbf{x}_i) - \mu_0(\mathbf{x}_i) \right\}^2 \xi_i$$

with $\mu_j(\mathbf{x}) = \eta_j\left(\mathbf{x}|\theta_j\right) + \psi_j(\mathbf{x})$,

subject to $\mathbf{U}_j'\boldsymbol{\psi}_j = \mathbf{0}$ and $\left\|\boldsymbol{\psi}_j\right\| \leq \tau_j$ for $j = 0, 1$.

- The orthogonality constraints state that $\boldsymbol{\psi}_j = \mathbf{V}_j \mathbf{c}_j$, where the columns of $\mathbf{V}_j$ are an orthonormal basis for col $\left(\mathbf{U}_j\right)^{\perp}$. This reduces the minimization to one of minimizing over $\mathbf{c}_0, \mathbf{c}_1$ subject to the constraints $\left\|\mathbf{c}_j\right\| \leq \tau_j$ (equality is attained; introduce Lagrange multipliers λ_0, λ_1 to handle these two remaining constraints).

- Minimization can be handled analytically: Minimum divergence (information loss) is

$$\mathbb{D}\left(\xi|\lambda_1, \lambda_0\right) =$$
$$\left\|\left(\mathbf{I}_N + \mathbf{D}_{\boldsymbol{\xi}}^{1/2}\mathbf{P}\mathbf{D}_{\boldsymbol{\xi}}^{1/2}\right)^{-1}\mathbf{D}_{\boldsymbol{\xi}}^{1/2}\left(\boldsymbol{\eta}_1 - \boldsymbol{\eta}_0\right)\right\|^2,$$

where

$$\mathbf{D}_{\boldsymbol{\xi}} = diag\left(\xi_1, ..., \xi_N\right),$$
$$\mathbf{P} = \frac{\mathbf{V}_0\mathbf{V}_0'}{\lambda_0} + \frac{\mathbf{V}_1\mathbf{V}_1'}{\lambda_1}.$$

- We maximize this over designs $\boldsymbol{\xi}$ for fixed λ_0, λ_1 and then calculate $\tau_j = \left\|\mathbf{c}_j\right\|$ in terms of the lambdas, which then parameterize the solutions.

Simulated annealing:

Step 1 Choose an initial n - point design.

Step 2 Carry out the following, until L new designs have been tested without any improvement:

2a Make a small random change in ξ. Let ξ' be the resulting design. Compute $\mathbb{D}(\xi'; \lambda_1, \lambda_0)$.

2b If $\nabla\mathbb{D} = \mathbb{D}(\xi'; \lambda_1, \lambda_0) - \mathbb{D}(\xi; \lambda_1, \lambda_0) > 0$ then the new design ξ' is accepted (and relabelled ξ). Otherwise, it is accepted with probability $e^{\nabla\mathbb{D}/T} < 1$ for a 'temperature' parameter T.

Step 3 Lower the temperature: $T \leftarrow .95T$; repeat Steps 1 and 2.

Continue until the fraction of improved designs found, at a fixed temperature, drops below $1/L$. ($L = 500$ in examples.) Stop.

- Recall the two models:

$$\mu_0(x) \approx \eta_0(x|\theta_0) = \frac{V_0 x}{K_0 + x},$$
$$\mu_1(x) \approx \eta_1(x|\theta_1) = V_1\left(1 - e^{-K_1 x}\right).$$

All examples use $\chi = .1(.1)5$.

- Example: Assume Model 1 is correct: $E[Y|x] = \eta_1(x|\theta_1 = (1,1)')$. Then

$$\theta_0 = \arg\min_{\theta} \sum_{\mathbf{x}\in\chi} \{E[Y|\mathbf{x}] - \eta_0(\mathbf{x}|\theta)\}^2 = (1.22, .91)'.$$

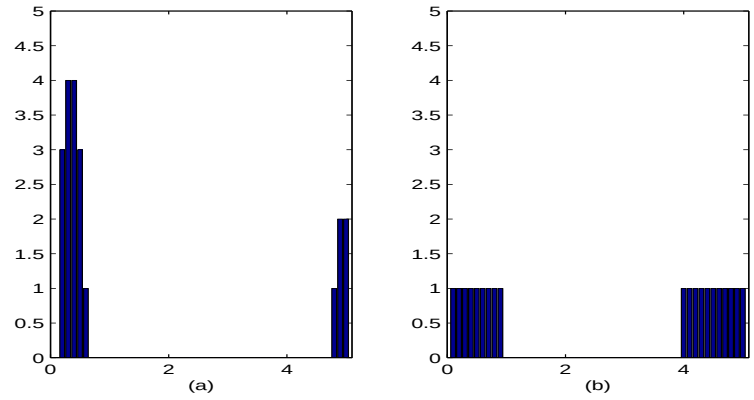

Design frequencies ($n = 20$):

(a) $\lambda_1 = \lambda_0 = 2$ ($\tau_1 = .0064, \tau_0 = .0060$),

(b) $\lambda_1 = \lambda_0 = .05$ ($\tau_1 = .0442, \tau_0 = .0433$).

- Example: Suppose neither model is correct, and that $E[Y|x] = .35 + .12x$ is linear. The closest MM and Exponential responses are plotted below; they have $\theta_0 = (1.02, 1.13)', \theta_1 = (.85, .73)'$.

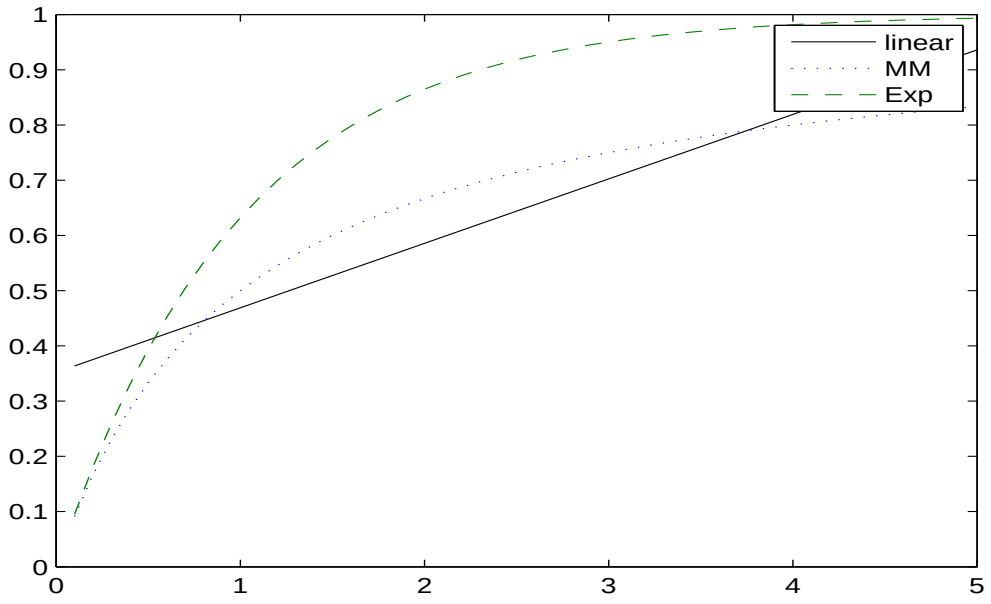

Michaelis-Menten and Exponential responses with linear approximation to MM.

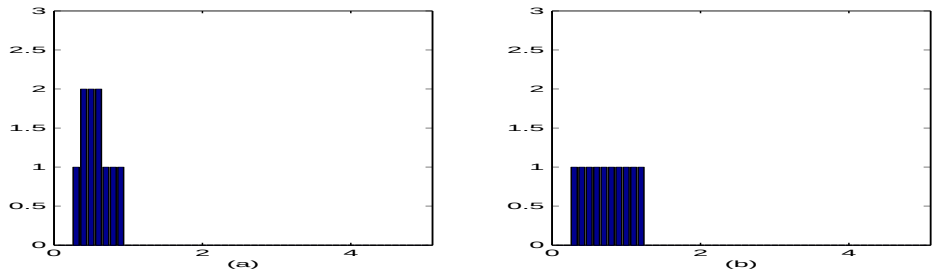

Design frequencies ($n = 10$); $E[Y|x]$ linear:
(a) $\lambda_1 = \lambda_0 = 2$ ($\tau_1 = .0079, \tau_0 = .0071$);
(b) $\lambda_1 = \lambda_0 = .05$ ($\tau_1 = .0739, \tau_0 = .0622$).

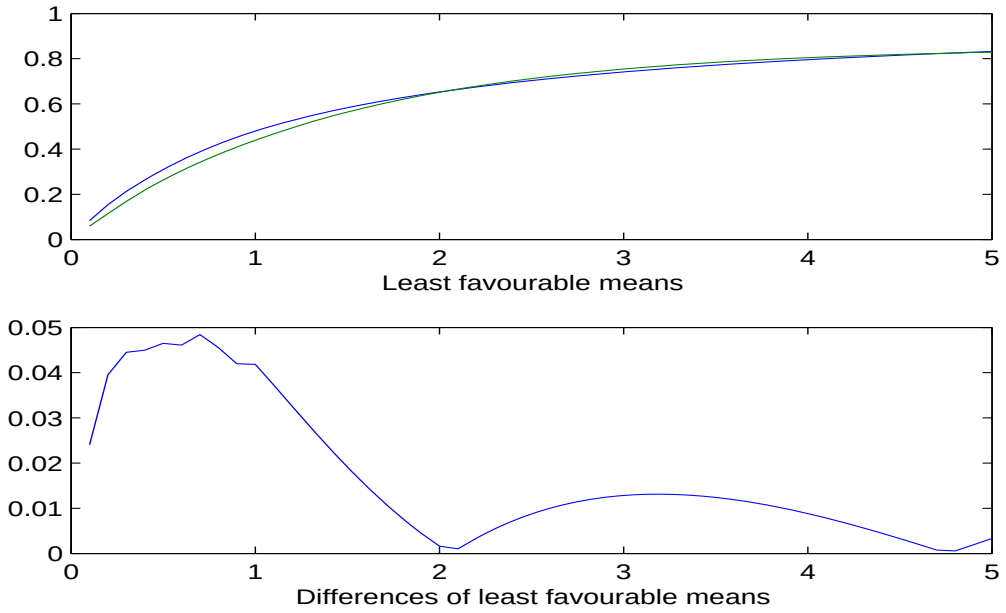

Top: Least favourable $\mu_0^*(x)$, $\mu_1^*(x)$.
Bottom: $\left|\mu_0^*(x) - \mu_1^*(x)\right|$. Both for case (a).

Another possibility is to design sequentially:

Step 1 Choose a small, initial design ξ.

Step 2 Minimize over $\psi_0(\cdot), \psi_1(\cdot)$ as in the previous algorithms. The parameters required in the evaluation of $\mu_j^*(\mathbf{x}_i) = \eta_j\left(\mathbf{x}_i|\theta_j\right) + \psi_j^*(\mathbf{x}_i)$ are replaced by estimates $\hat{\theta}_j$, rather than being defined as minimizers w.r.t. $E[Y|\mathbf{x}]$.

Step 3 Make the next observation at $\mathbf{x}_{\text{new}} = \arg\max_{\mathbf{x}\in\chi} \mathbb{I}\left(\mu_0^*(\mathbf{x}), \mu_1^*(\mathbf{x})\right)$.

Steps 2 and 3 are repeated until an n-point design is obtained.

Example: Sequential designs, $n = 10, 40$. Sampling was done from a log-normal population with mean $\mu(x) = .35 + .12x$ and coefficient of variation $cv^2 = var/mean^2 = 1$. Starting design places one observation at each of $x = .4, 1.2$ - the extremes of the design of previous example.

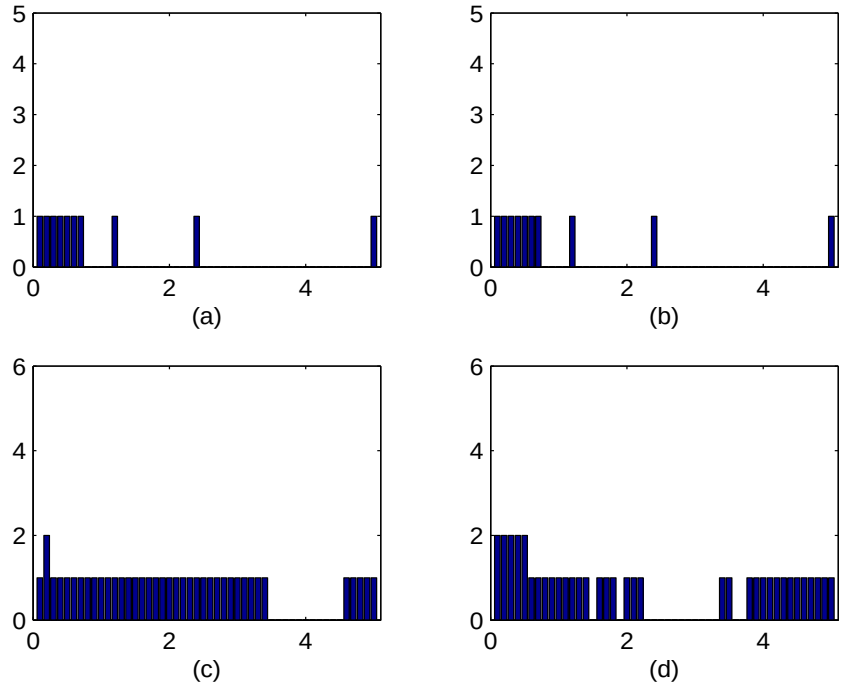

(a): $n = 10, \tau_0 = \tau_1 = .01$;
(b): $n = 10, \tau_0 = \tau_1 = .05$;
(c): $n = 40, \tau_0 = \tau_1 = .01$;
(d): $n = 40, \tau_0 = \tau_1 = .05$.

- Designs optimal for discrimination are typically poor for estimation and/or prediction. We might aim to combine high discriminatory power ("KL-optimality") with small mse of the fitted values ("I-optimality") by shifting the emphasis from the former to the latter, as evidence accrues in favour of one of the two models.

- Given an n-point design, choose the next design point according to

$$\mathbf{x}_{\text{new}} = \arg\max_{\mathbf{x}\in\chi} \min_{\psi_{n,0},\psi_{n,1}} \Delta_n(\mathbf{x}), \text{ where}$$

$$\Delta_n(\mathbf{x}) = (1-\kappa_n)\,\Delta_n^{(1)}(\mathbf{x}) + \kappa_n \Delta_{n,j_n}^{(2)}(\mathbf{x}).$$

Here (1): κ_n increases with $|r|$ (observed log-l'hood ratio for discrimination);

(2): $\Delta_n^{(1)}(\mathbf{x}) =$

$$\frac{\mathbb{I}\left(\eta_0\left(\mathbf{x}|\hat{\boldsymbol{\theta}}_{n,0}\right)+\psi_{n,0}(\mathbf{x}),\eta_1\left(\mathbf{x}|\hat{\boldsymbol{\theta}}_{n,1}\right)+\psi_{n,1}(\mathbf{x})\right)}{n}$$

measures discriminatory power, and (3): $\Delta_{n,j_n}^{(2)}(\mathbf{x})$ measures the drop in mse of the fitted values (averaged overχ) in the favoured model, as a result of adding a design point $\mathbf{x}$. ($j_n = j_n(r) \in \{0,1\}$.)

Simulation: True response exponential; $\tau_0 = \tau_1 = .01$. Top: $cv^2 = .1$. Bottom: $cv^2 = 1$. Initial 10-point uniform design. (a) sequentially chosen design points. (b) final design ($n = 40$). (c) favoured model j_n (exponential = 1, MM = 0) and weights κ_n placed on estimation/prediction in the favoured model.