

Chapter 4

Two-way contingency tables

{ch:twoway}

The analysis of two-way frequency tables concerns the association between two variables. A variety of specialized graphical displays help to visualize the pattern of association, using area of some region to represent the frequency in a cell. Some of these methods are focused on visualizing an odds ratio (for 2×2 tables), or the general pattern of association, or the agreement between row and column categories in square tables.

4.1 Introduction

{sec:twoway-intro}

Tables are like cobwebs, like the sieve of Danaides; beautifully reticulated, orderly to look upon, but which will hold no conclusion. Tables are abstractions, and the object a most concrete one, so difficult to read the essence of.

From *Chartism* by Thomas Carlyle (1840), Chapter II, Statistics

Most methods of statistical analysis are concerned with understanding relationships or dependence among variables. With categorical variables, these relationships are often studied from data which has been summarized by a **contingency table** in table form or frequency form, giving the frequencies of observations cross-classified by two or more such variables. As Thomas Carlyle said, it is often difficult to appreciate the message conveyed in numerical tables.

This chapter is concerned with simple graphical methods for understanding the association between two categorical variables. Some examples are also presented which involve a third, **stratifying variable**, where we wish to determine if the relationship between two primary variables is the same or different for all levels of the stratifying variable. More general methods for fitting models and displaying associations for three-way and larger tables are described in Chapter 5.

In Section 4.2, We describe briefly some numerical and statistical methods for testing whether an association exists between two variables, and measures for quantifying the strength of this association. In Section 4.3 we extend these ideas to situations where the relation between two variables is of primary interest, but there are one or more background variables to be controlled.

The main emphasis, however, is on graphical methods which help to describe the *pattern* of an association between variables. Section 4.4 presents the fourfold display, designed to portray the odds ratio in 2×2 tables or a set of k such tables. **Sieve diagrams** (Section 4.5) and **association**

plots (Section 4.6) are more general methods for depicting the pattern of associations in any two-way tables. When the row and column variables represent the classifications of different raters, specialized measures and visual displays for *inter-rater agreement* (Section 4.7) are particularly useful. Another specialized display, a *trilinear plot* or *ternary plot*, described in Section 4.8, is designed for three-column frequency tables or compositional data. In order to make clear some of the distinctions which occur in contingency table analysis, we begin with several examples.

{ex:berkeley1}

EXAMPLE 4.1: Berkeley admissions

Table 4.1 shows aggregate data on applicants to graduate school at Berkeley for the six largest departments in 1973 classified by admission and gender (Bickel *et al.*, 1975). See `UCBAdmissions` for the complete data set. For such data we might wish to study whether there is an association between admission and gender. Are male (or female) applicants more likely to be admitted? The presence of an association might be considered as evidence of sex bias in admission practices.

Table 4.1 is an example of the simplest kind of contingency table, a 2×2 classification of individuals according to two dichotomous (binary) variables. For such a table, the question of whether there is an association between admission and gender is equivalent to asking if the proportions of males and females who are admitted to graduate school are the same, or whether the difference in proportions admitted is zero. \triangle

{tab:berk22}

Table 4.1: Admissions to Berkeley graduate programs

	Admitted	Rejected	Total	% Admit	Odds(Admit)
Males	1198	1493	2691	44.52	0.802
Females	557	1278	1835	30.35	0.437
Total	1755	2771	4526	38.78	0.633

Although the methods for quantifying association in larger tables can be used for 2×2 tables, there are specialized measures (described in Section 4.2) and graphical methods for these simpler tables.

As we mentioned in Section 1.2.4 it is often useful to make a distinction between *response*, or outcome variables, on the one hand, and possible *explanatory* or predictor variables on the other. In Table 4.1, it is natural to consider admission as the outcome, and gender as the explanatory variable. In other tables, no variable may be clearly identified as *the* outcome, or there may be several response variables, giving a multivariate problem.

{ex:haireye1}

EXAMPLE 4.2: Hair color and eye color

Table 4.2 shows data collected by Snee (1974) on the relation between hair color and eye color among 592 students in a statistics course (a two-way margin of `HairEyeColor`). Neither hair color nor eye color is considered a response in relation to the other; our interest concerns whether an association exists between them. Hair color and eye color have both been classified into four categories. Although the categories used are among the most common, they are not the only categories possible.¹ Everyday observation suggests that there probably is an association between hair color and eye color, and we will describe tests and measures of associations for larger tables in Section 4.2.3.

¹If students had been asked to write down their hair and eye colors, it is likely that many more than four categories of each would appear in a sample of nearly 600.

Table 4.2: Hair-color eye-color data

Eye Color	Hair Color				Total
	Black	Brown	Red	Blond	
Green	5	29	14	16	64
Hazel	15	54	14	10	93
Blue	20	84	17	94	215
Brown	68	119	26	7	220
Total	108	286	71	127	592

△

If, as is suspected, hair color and eye color are associated, we would like to understand *how* they are associated. The graphical methods described later in this chapter and in Chapter 5 help reveal the pattern of associations present.

{ex:mental1}

EXAMPLE 4.3: Mental impairment and parents' SES

Srole *et al.* (1978, p. 289) gave the data in Table 4.3 on the mental health status of a sample of 1660 young New York residents in midtown Manhattan classified by their parents' socioeconomic status (SES); see `Mental` in the `vcdExtra` package. These data have also been analyzed by many authors, including Agresti (2013, §10.5.3), Goodman (1979), and Haberman (1979, p. 375).

There are five categories of SES and mental health is classified in the four categories “well”, “mild symptom formation”, “moderate symptom formation”, and “impaired”. It may be useful here to consider SES as explanatory and ask whether and how it predicts mental health status as a response.

Table 4.3: Mental impairment and parents' SES

SES	Mental impairment			
	Well	Mild	Moderate	Impaired
1	64	94	58	46
2	57	94	54	40
3	57	105	65	60
4	72	141	77	94
5	36	97	54	78
6	21	71	54	71

{tab:mental-tab}

Although there may be an overall association between these two variables, more powerful and focused tests are available when we treat these variables as *ordinal*, as we will see in Section 4.2.4.

△

{ex:arthrit1}

EXAMPLE 4.4: Arthritis treatment

The data in Table 4.4 compares an active treatment for rheumatoid arthritis to a placebo (Koch and Edwards, 1988), used in examples in Chapter 2 (Example 2.2). The outcome reflects

whether individuals showed no improvement, some improvement, or marked improvement. Here, the outcome variable is an ordinal one, and it is probably important to determine if the relation between treatment and outcome is the same for males and females. The data set is given in case form in *Arthritis*.

Table 4.4: Arthritis treatment data

		Improvement			
Treatment	Sex	None	Some	Marked	Total
Active	Female	6	5	16	27
	Male	7	2	5	14
Placebo	Female	19	7	6	32
	Male	10	0	1	11
Total		42	14	28	84

This is, of course, a three-way table, with factors Treatment, Sex, and Improvement. If the relation between treatment and outcome is the same for both genders, an analysis of the Treatment by Improvement table (collapsed over sex) could be carried out. Otherwise we could perform separate analyses for men and women, or treat the combinations of treatment and sex as four levels of a “population” variable, giving a 4×3 two-way table. These simplified approaches each ignore certain information available in an analysis of the full three-way table. \triangle

4.2 Tests of association for two-way tables

4.2.1 Notation and terminology

To establish notation, let $\mathbf{N} = \{n_{ij}\}$ be the observed frequency table of variables A and B with r rows and c columns, as shown in Table 4.5. In what follows, a subscript is replaced by a “+” when summed over the corresponding variable, so $n_{i+} = \sum_j n_{ij}$ gives the total frequency in row i , $n_{+j} = \sum_i n_{ij}$ gives the total frequency in column j , and $n_{++} = \sum_i \sum_j n_{ij}$ is the grand total; for convenience, n_{++} is also symbolized by n .

Table 4.5: The $r \times c$ contingency table

Row Category	Column category				Total
	1	2	\dots	c	
1	n_{11}	n_{12}	\dots	n_{1c}	n_{1+}
2	n_{21}	n_{22}	\dots	n_{2c}	n_{2+}
\vdots	\vdots	\vdots	\dots	\vdots	\vdots
r	n_{r1}	n_{r2}	\dots	n_{rc}	n_{r+}
Total	n_{+1}	n_{+2}	\dots	n_{+c}	n_{++}

When each observation is randomly sampled from some population and classified on two

categorical variables, A and B , we refer to the *joint distribution* of these variables, and let $\pi_{ij} = \Pr(A = i, B = j)$ denote the population probability that an observation is classified in row i , column j (or cell (ij)) in the table. Corresponding to these population joint probabilities, the cell proportions, $p_{ij} = n_{ij}/n$, give the sample joint distribution.

The row totals n_{i+} and column totals n_{+j} are called *marginal frequencies* for variables A and B respectively. These describe the distribution of each variable *ignoring* the other. For the population probabilities, the *marginal distributions* are defined analogously as the row and column totals of the joint probabilities, $\pi_{i+} = \sum_j \pi_{ij}$, and $\pi_{+j} = \sum_i \pi_{ij}$. The sample marginal proportions are, correspondingly, $p_{i+} = \sum_j p_{ij} = n_{i+}/n$, and $p_{+j} = \sum_i p_{ij} = n_{+j}/n$.

When one variable (the column variable, B , for example) is a response variable, and the other (A) is an explanatory variable, it is most often useful to examine the distribution of the response B for *each* level of A separately. These define the *conditional distributions* of B , given the level of A , and are defined for the population as $\pi_{j|i} = \pi_{ij}/\pi_{i+}$.

These definitions are illustrated for the Berkeley data (Table 4.1) below, using the function `CrossTable()`.

```
Berkeley <- margin.table(UCBAdmissions, 2:1)
library(gmodels)
CrossTable(Berkeley, prop.chisq=FALSE, prop.c=FALSE, format="SPSS")
```

```
##
##      Cell Contents
## |-----|
## |                Count |
## |                Row Percent |
## |                Total Percent |
## |-----|
##
## Total Observations in Table:  4526
##
##      | Admit
##      | Admitted | Rejected | Row Total |
## -----|-----|-----|-----|
##      Male |    1198 |    1493 |    2691 |
##      |    44.519% |    55.481% |    59.456% |
##      |    26.469% |    32.987% |
## -----|-----|-----|
##      Female |    557 |    1278 |    1835 |
##      |    30.354% |    69.646% |    40.544% |
##      |    12.307% |    28.237% |
## -----|-----|-----|
## Column Total |    1755 |    2771 |    4526 |
## -----|-----|-----|
##
##
##
```

The output shows the joint frequencies, n_{ij} , and joint sample percentages, $100 \times p_{ij}$, in the first row within each table cell. The second row in each cell (“Row percent”) gives the conditional percentage of admission or rejection, $100 \times p_{j|i}$ for males and females separately. The row and column labelled “Total” give the marginal frequencies, n_{i+} and n_{+j} , and marginal percentages, p_{i+} and p_{+j} .

4.2.2 2 by 2 tables

{sec:twoway}

The 2×2 contingency table of applicants to Berkeley graduate programs in Table 4.1 may be regarded as an example of a *cross-sectional study*. The total of $n = 4,526$ applicants in 1973 has been classified by both gender and admission status. Here, we would probably consider the total n to be fixed, and the cell frequencies n_{ij} , $i = 1, 2; j = 1, 2$ would then represent a single *multinomial sample* for the cross-classification by two binary variables, with probabilities cell p_{ij} , $i = 1, 2; j = 1, 2$ such that

$$p_{11} + p_{12} + p_{21} + p_{22} = 1 .$$

The basic null hypothesis of interest for a multinomial sample is that of independence. Are admission and gender independent of each other?

Alternatively, if we consider admission the response variable, and gender an explanatory variable, we would treat the numbers of male and female applicants as fixed and consider the cell frequencies to represent two independent *binomial samples* for a binary response. In this case, the null hypothesis is described as that of homogeneity of the response proportions across the levels of the explanatory variable.

Odds and odds ratios

{sec:twoway-odds}

Measures of association are used to quantify the strength of association between variables. Among the many measures of association for contingency tables, the *odds ratio* is particularly useful for 2×2 tables, and is a fundamental parameter in several graphical displays and models described later. Other measures of strength of association for 2×2 tables are described in Stokes *et al.* (2000, Chapter 2) and Agresti (1996, §2.2).

For a binary response, where the probability of a “success” is π , the *odds* of a success is defined as

$$\text{odds} = \frac{\pi}{1 - \pi} .$$

Hence, odds = 1 corresponds to $\pi = 0.5$, or success and failure equally likely. When success is more likely than failure $\pi > 0.5$, and the odds > 1 ; for instance, when $\pi = 0.75$, odds = $.75/.25 = 3$, so a success is three times as likely as a failure. When failure is more likely, $\pi < 0.5$, and the odds < 1 ; for instance, when $\pi = 0.25$, odds = $.25/.75 = \frac{1}{3}$.

The odds of success thus vary *multiplicatively* around 1. Taking logarithms gives an equivalent measure which varies *additively* around 0, called the *log odds* or *logit*:

{eq:logit}

$$\text{logit}(\pi) \equiv \log(\text{odds}) = \log\left(\frac{\pi}{1 - \pi}\right) . \quad (4.1)$$

The logit is symmetric about $\pi = 0.5$, in that $\text{logit}(\pi) = -\text{logit}(1 - \pi)$. The following lines calculate the odds and log odds for a range of probabilities. As you will see in Chapter 7, the logit transformation of a probability is fundamental in logistic regression.

```
p <- c(0.05, .1, .25, .50, .75, .9, .95)
odds <- p / (1-p)
logodds <- log(odds)
data.frame(p, odds, logodds)
```

##	p	odds	logodds
## 1	0.05	0.052632	-2.9444
## 2	0.10	0.111111	-2.1972
## 3	0.25	0.333333	-1.0986
## 4	0.50	1.000000	0.0000
## 5	0.75	3.000000	1.0986
## 6	0.90	9.000000	2.1972
## 7	0.95	19.000000	2.9444

A binary response for two groups gives a 2×2 table, with Group as the row variable, say. Let π_1 and π_2 be the success probabilities for Group 1 and Group 2. The **odds ratio**, θ , is just the ratio of the odds for the two groups:

$$\text{odds ratio} \equiv \theta = \frac{\text{odds}_1}{\text{odds}_2} = \frac{\pi_1/(1 - \pi_1)}{\pi_2/(1 - \pi_2)} .$$

Like the odds itself, the odds ratio is always non-negative, between 0 and ∞ . When $\theta = 1$, the distributions of success and failure are the same for both groups (so $\pi_1 = \pi_2$); there is no association between row and column variables, or the response is independent of group. When $\theta > 1$, Group 1 has a greater success probability; when $\theta < 1$, Group 2 has a greater success probability.

Similarly, the odds ratio may be transformed to a log scale, to give a measure which is symmetric about 0. The **log odds ratio**, symbolized by ψ , is just the difference between the logits for Groups 1 and 2:

$$\log \text{ odds ratio} \equiv \psi = \log(\theta) = \log \left[\frac{\pi_1/(1 - \pi_1)}{\pi_2/(1 - \pi_2)} \right] = \text{logit}(\pi_1) - \text{logit}(\pi_2) .$$

Independence corresponds to $\psi = 0$, and reversing the rows or columns of the table merely changes the sign of ψ .

For sample data, the **sample odds ratio** is the ratio of the sample odds for the two groups:

$$\hat{\theta} = \frac{p_1/(1 - p_1)}{p_2/(1 - p_2)} = \frac{n_{11}/n_{12}}{n_{21}/n_{22}} = \frac{n_{11}n_{22}}{n_{12}n_{21}} . \quad (4.2) \quad \{\text{eq:soddsratio}\}$$

The sample estimate $\hat{\theta}$ in Eqn. (4.2) is the maximum likelihood estimator of the true θ . The sampling distribution of $\hat{\theta}$ is asymptotically normal as $n \rightarrow \infty$, but may be highly skewed in small to moderate samples.

Consequently, inference for the odds ratio is more conveniently carried out in terms of the log odds ratio, whose sampling distribution is more closely normal, with mean $\psi = \log(\theta)$, and asymptotic standard error (ASE)

$$\text{ASE}_{\log(\theta)} \equiv \hat{s}(\hat{\psi}) = \left\{ \frac{1}{n_{11}} + \frac{1}{n_{12}} + \frac{1}{n_{21}} + \frac{1}{n_{22}} \right\}^{1/2} = \left\{ \sum \sum n_{ij}^{-1} \right\}^{1/2} \quad (4.3) \quad \{\text{eq:aselogtheta}\}$$

A large-sample $100(1 - \alpha)\%$ confidence interval for $\log(\theta)$ may therefore be calculated as

$$\log(\theta) \pm z_{1-\alpha/2} \text{ASE}_{\log(\theta)} = \hat{\psi} \pm z_{1-\alpha/2} \hat{s}(\hat{\psi})$$

where $z_{1-\alpha/2}$ is the cumulative normal quantile with $1 - \alpha/2$ in the lower tail. Confidence intervals for θ itself are obtained by exponentiating the end points of the interval for $\psi = \log(\theta)$,²

$$\exp\left(\hat{\psi} \pm z_{1-\alpha/2}\hat{s}(\hat{\psi})\right) .$$

4.2.3 Larger tables: Overall analysis

For two-way tables overall tests of association can be carried out using `assocstats()`. If the data set has more than two factors (as in the Arthritis Treatment data), the other factors will be ignored (and collapsed) if not included when the table is constructed. This simplified analysis may be misleading if the excluded factors interact with the factors used in the analysis.

EXAMPLE 4.5: Arthritis treatment

Since the main interest is in the relation between *Treatment* and *Improved*, an overall analysis (which ignores *Sex*) can be carried out by creating a two-way table with `xtabs()` as shown below.

```
data("Arthritis", package="vcd")
Art.tab <- xtabs(~Treatment + Improved, data=Arthritis)
Art.tab

##           Improved
## Treatment None Some Marked
## Placebo    29     7     7
## Treated   13     7    21

round(100*prop.table(Art.tab, margin=1), 2)

##           Improved
## Treatment None Some Marked
## Placebo  67.44 16.28 16.28
## Treated  31.71 17.07 51.22
```

The row proportions show a clear difference in the outcome for the two groups: For those given the placebo, 67% reported no improvement; in the treated group, 51% reported marked improvement. χ^2 tests and measures of association are provided by `assocstats()` as shown below:

```
assocstats(Art.tab)

## Error in eval(expr, envir, enclos): could not find function "assocstats"
```

△

²Note that $\hat{\theta}$ is 0 or ∞ if any $n_{ij} = 0$. Haldane (1955) and Gart and Zweifl (1967) showed that improved estimators of θ and $\psi = \log(\theta)$ are obtained by replacing each n_{ij} by $[n_{ij} + \frac{1}{2}]$ in Eqn. (4.2) and Eqn. (4.3). This adjustment is preferred in small samples, and required if any zero cells occur. In large samples, the effect of adding 0.5 to each cell becomes negligible.

4.2.4 Tests for ordinal variables

For $r \times c$ tables, more sensitive tests than the test for general association (independence) are available if either or both of the row and column variables are ordinal. Generalized **Cochran-Mantel-Haenszel tests** (Landis *et al.*, 1978) which take the ordinal nature of a variable into account are provided by the `CMHtest()` in `vcdExtra`. These tests are based on assigning numerical scores to the table categories; the default (table) scores treat the levels as equally spaced. They generally have higher power when the pattern of association is determined by the order of an ordinal variable.

{ex:mental2}

EXAMPLE 4.6: Mental impairment and parents' SES

We illustrate these tests using the data on mental impairment and SES introduced in Example 4.3, where both variables can be considered ordinal.

```
data(Mental, package="vcdExtra")
mental.tab <- xtabs(Freq ~ ses + mental, data=Mental)
assocstats(mental.tab)      # standard chisq tests

## Error in eval(expr, envir, enclos): could not find function "assocstats"

CMHtest(mental.tab)        # CMH tests

## Error in eval(expr, envir, enclos): could not find function "CMHtest"
```

In this data set, all four tests show a highly significant association. However, the `cor` test for nonzero correlation uses only one degree of freedom, whereas the test of general association requires 15 df. △

The four tests differ in the types of departure from independence they are sensitive to:

General Association When the row and column variables are both nominal (unordered) the only alternative hypothesis of interest is that there is *some* association between the row and column variables. The CMH test statistic is similar to the (Pearson) Chi-Square and Likelihood Ratio Chi-Square in the result from `assocstats()`; all have $(r - 1)(c - 1)$ df.

Row Mean Scores Differ If the column variable is ordinal, assigning scores to the column variable produces a mean for each row. The association between row and column variables can be expressed as a test of whether these means differ over the rows of the table, with $r - 1$ df. This is analogous to the Kruskal-Wallis non-parametric test (ANOVA based on rank scores).

Column Mean Scores Differ Same as the above, assigning scores to the row variable.

Nonzero Correlation (Linear association) When *both* row and column variables are ordinal, we could assign scores to both variables and compute the correlation (r). The CMH χ^2 is equal to $(N - 1)r^2$, where N is the total sample size. The test is most sensitive to a pattern where the row mean score changes linearly over the rows.

4.2.5 Sample CMH Profiles

{sec:Sample

Two contrived examples may make the differences among these tests more apparent. Visualizations of the patterns of association reinforces the aspects to which the tests are most sensitive, and introduces the sieve diagram described more fully in Section 4.5.

General Association

The table below exhibits a general association between variables A and B , but no difference in row means or linear association. The row means are calculated by assigning integer scores, $b_i = i$ to the column categories. Figure 4.1(left) shows the pattern of association in this table graphically, as a sieve diagram (described in Section 4.5).

	b1	b2	b3	b4	b5	Total	Mean
a1	0	15	25	15	0	55	3.0
a2	5	20	5	20	5	55	3.0
a3	20	5	5	5	20	55	3.0
Total	25	40	35	40	25	165	3.0

This is reflected in the `CMHtest()` output shown below. **TODO: Something wrong here: does `CMHtest()` get rows/cols mixed up? Would be nice to calculate col means also.**

```
CMHtest(cmhdemo1)

## Error in eval(expr, envir, enclos): could not find function "CMHtest"
```

The chi-square values for non-zero correlation and different row mean scores are exactly zero because the row means are all equal. Only the general association test shows that A and B are associated.

```
sieve(cmhdemo1, shade=TRUE, main="General association",
      gp = shading_sieve(interpolate = 0, lty = c("solid", "longdash")))

## Error in eval(expr, envir, enclos): could not find function "sieve"

sieve(cmhdemo2, shade=TRUE, main="Linear association",
      gp = shading_sieve(interpolate = 0, lty = c("solid", "longdash")))

## Error in eval(expr, envir, enclos): could not find function "sieve"
```

Linear Association

The table below contains a weak, non-significant general association, but significant row mean differences and linear associations. The unstructured test of general association would therefore lead to the conclusion that no association exists, while the tests taking ordinal factors into account would conclude otherwise. Note that the largest frequencies shift towards lower levels of B as the level of variable A increases. See Figure 4.1(right) for a visual representation of this pattern.

	b1	b2	b3	b4	b5	Total	Mean
a1	2	5	8	8	8	31	3.48
a2	2	8	8	8	5	31	3.19
a3	5	8	8	8	2	31	2.81
a4	8	8	8	5	2	31	2.52
Total	17	29	32	29	17	124	3.00

Note that the χ^2 -values for the row-means and non-zero correlation tests from `CMHtest()` are very similar, but the correlation test is more highly significant since it is based on just one degree of freedom.

```
CMHtest(cmhdemo2)
```

```
## Error in eval(expr, envir, enclos): could not find function "CMHtest"
```

The difference in sensitivity and power among these tests for categorical data is analogous to the difference between general ANOVA tests and tests for linear trend (contrasts) in experimental designs with quantitative factors: The more specific test has greater power, but is sensitive to a narrower range of departures from the null hypothesis. The more focused tests for ordinal factors are a better bet when we believe that the association depends on the ordered nature of the factor levels.

4.3 Stratified analysis

{sec:twoway-strat}

An overall analysis ignores other variables (like sex), by collapsing over them. In the `Arthritis` data, it is possible that the treatment is effective only for one gender, or even that the treatment has opposite effects for men and women. If so, pooling over the ignored variable(s) can be seriously misleading.

A *stratified analysis* controls for the effects of one or more background variables. This is similar to the use of a blocking variable in an ANOVA design. Tests for association can be obtained by applying a function (`assocstats()`, `CMHtest()`) over the levels of the stratifying variables.

{ex:arthrit3}

EXAMPLE 4.7: Arthritis treatment

The statements below request a stratified analysis of the arthritis treatment data with CMH tests, controlling for sex. Essentially, the analysis is carried out separately for males and females.

The table `Art.tab2` is constructed as a three-way table, with `sex` as the last dimension.

```
Art.tab2 <- xtabs(~Treatment + Improved + Sex, data=Arthritis)
Art.tab2

## , , Sex = Female
##
##      Improved
## Treatment None Some Marked
##  Placebo    19    7    6
##   Treated    6    5   16
##
## , , Sex = Male
```

```
##
##           Improved
## Treatment None Some Marked
## Placebo    10    0     1
## Treated     7    2     5
```

`assocstats()` only applies to two-way tables, so we use `apply()` to run it for each level of *Sex*. `CMHtest()` is designed for such stratified tables, and uses all dimensions after the first two as strata.

```
apply(Art.tab2, MARGIN=3, FUN=assocstats)
```

```
## Error in match.fun(FUN): object 'assocstats' not found
```

Note that even though the strength of association (ϕ -coefficient) is similar in the two groups, the χ^2 tests show significance for females, but not for males. This is true even using the more powerful CMH tests below, treating *Treatment* as ordinal. The reason is that there were more than twice as many females as males in this sample.

```
CMHtest(Art.tab2)
```

```
## Error in eval(expr, envir, enclos): could not find function "CMHtest"
```

```
apply(Art.tab2, 3, sum)
```

```
## Female    Male
##      59     25
```

△

4.3.1 Assessing homogeneity of association

{sec:twoway-homog}

In a stratified analysis it is often crucial to know if the association between the primary table variables is the same over all strata. For $2 \times 2 \times k$ tables this question reduces to whether the odds ratio is the same in all k strata. The `vcd` package implements Woolf's test (Woolf, 1995) in `woolf_test()` for this purpose.

For larger n -way tables, this question is equivalent to testing whether the association between the primary variables, A and B , say, is the same for all levels of the stratifying variables, C , D , ...

In the case of a 3-way table, this can be stated as the *loglinear model* of no three-way association, $[AB][AC][BC]$. This notation (described in Section 8.2) lists only the high-order association terms in a linear model for log frequency.

{ex:berkeley1a}

EXAMPLE 4.8: Berkeley admissions

Here we illustrate the use of Woolf's test for the `UCBAdmissions` data. The test is significant, indicating that the odds ratios cannot be considered equal across departments. We will see why when we visualize the data by department in the next section.

```
woolf_test(UCBAdmissions)
```

```
## Error in eval(expr, envir, enclos): could not find function "woolf_test"
```



x:arthrit4}

EXAMPLE 4.9: Arthritis treatment

For the arthritis data, homogeneity means that there is no three-way Treatment * Improved * Sex association. That is, the association between treatment and outcome (`improve`) is the same for both men and women. This hypothesis can be stated as the loglinear model,

$$[\text{SexTreatment}] [\text{SexImproved}] [\text{TreatmentImproved}] . \quad (4.4) \quad \{\text{eq:ST02}\}$$

Such tests can be carried out most conveniently using `loglm()` in the MASS package. The model formula uses the standard R notation $()^2$ to specify all terms of order 2.

```
library(MASS)
loglm(~ (Treatment + Improved + Sex)^2, data=Art.tab2)

## Call:
## loglm(formula = ~(Treatment + Improved + Sex)^2, data = Art.tab2)
##
## Statistics:
##
##              X^2 df P(> X^2)
## Likelihood Ratio 1.7037  2  0.42663
## Pearson          1.1336  2  0.56735
```

Even though we found in the CMH analysis above that the association between *Treatment* and *Improved* was stronger for females than males, the analysis using `loglm()` is clearly non-significant, so we cannot reject homogeneity of association. 

4.4 Fourfold display for 2 x 2 tables

{sec:twoway-fourfold}

The *fourfold display* is a special case of a *radial diagram* (or “polar area chart”) designed for the display of 2×2 (or $2 \times 2 \times k$) tables (Fienberg, 1975, Friendly, 1994a,b). In this display the frequency n_{ij} in each cell of a fourfold table is shown by a quarter circle, whose radius is proportional to $\sqrt{n_{ij}}$, so the area is proportional to the cell count. The fourfold display is similar to a pie chart in using segments of a circle to show frequencies. It differs from a pie chart in that it keeps the angles of the segments constant and varies the radius, whereas the pie chart varies the angles and keeps the radius constant.

The main purpose of this display is to depict the sample odds ratio, $\hat{\theta} = (n_{11}/n_{12}) \div (n_{21}/n_{22})$. An association between the variables ($\theta \neq 1$) is shown by the tendency of diagonally opposite cells in one direction to differ in size from those in the opposite direction, and the display uses color or shading to show this direction. Confidence rings for the observed θ allow a visual test of the hypothesis of independence, $H_0 : \theta = 1$. They have the property that (in a standardized display) the rings for adjacent quadrants overlap *iff* the observed counts are consistent with the null hypothesis. 

{ex:berkeley2}

EXAMPLE 4.10: Berkeley admissions

Figure 4.2(left) shows the basic, unstandardized fourfold display for the Berkeley admissions data (Table 4.1). Here, the area of each quadrant is proportional to the cell frequency, shown numerically in each corner. The odds ratio is proportional to the product of the areas shaded dark,

divided by the product of the areas shaded light. The sample odds ratio, $\text{Odds}(\text{Admit}|\text{Male}) / \text{Odds}(\text{Admit}|\text{Female})$ is 1.84 (see Example 4.8) indicating that males were nearly twice as likely to be admitted.

```
fourfold(Berkeley, std="ind.max") # unstandardized
## Error in eval(expr, envir, enclos): could not find function "fourfold"
fourfold(Berkeley, margin=1) # equating gender
## Error in eval(expr, envir, enclos): could not find function "fourfold"
```

However, it is difficult to make these visual comparisons because there are more men than women, and because the proportions admitted and rejected are unequal. In the unstandardized display the confidence bands have no interpretation as a test of $H_0 : \theta = 1$.

{tab:berkrow}

Table 4.6: Admissions to Berkeley graduate programs, Frequencies and Row Percentages

	Frequencies		Row Percents	
	Admitted	Rejected	Admitted	Rejected
Males	1198	1493	44.52	55.48
Females	557	1278	30.35	69.65

The data in a 2×2 table can be standardized to make these visual comparisons easier. Table 4.6 shows the Berkeley data with the addition of row percentages (which equate for the number of men and women applicants) indicating the proportion of each gender accepted and rejected. We see that 44.52% of males were admitted, while only 30.35% of females were admitted. Moreover, the row percentages have the same odds ratio as the raw data: $44.52 \times 69.65 / 30.35 \times 55.48 = 1.84$. Figure 4.2(right) shows the fourfold display where the area of each quarter circle is proportional to these row percentages.

With this standardization, the confidence rings have the property that the confidence rings for each upper quadrant will overlap with those for the quadrant below it if the odds ratio does not differ from 1.0. (Details of the calculation of confidence rings are described in the next section.) No similar statement can be made about the corresponding left and right quadrants, however, because the overall rate of admission has not been standardized.

As a final step, we can standardize the data so that *both* table margins are equal, while preserving the odds ratio. Each quarter circle is then drawn to have an area proportional to this standardized cell frequency. This makes it easier to see the association between admission and sex without being influenced by the overall admission rate or the differential tendency of males and females to apply. With this standardization, the four quadrants will align (overlap) horizontally and vertically when the odds ratio is 1, regardless of the marginal frequencies. The fully standardized display, which is usually the most useful form, is shown in Figure 4.3.

```
fourfold(Berkeley) # standardize both margins
## Error in eval(expr, envir, enclos): could not find function "fourfold"
```

These displays also use color (blue) and diagonal tick marks to show the direction of positive association. The visual interpretation (also conveyed by area) is that males are more likely to be accepted, females more likely to be rejected.

The quadrants in Figure 4.3 do not align and the 95% confidence rings around each quadrant do not overlap, indicating that the odds ratio differs significantly from 1—putative evidence of gender bias. The very narrow width of the confidence rings gives a visual indication of the precision of the data—if we stopped here, we might feel quite confident of this conclusion.

4.4.1 Confidence rings for odds ratio

Confidence rings for the fourfold display are computed from a confidence interval for θ , whose endpoints can each be mapped into a 2×2 table. Each such table is then drawn in the same way as the data.

The interval for θ is most easily found by considering the distribution of $\hat{\psi} = \log \hat{\theta}$, whose standard error may be estimated by Eqn. (4.3). Then an approximate $1 - \alpha$ confidence interval for ψ is given by

$$\hat{\psi} \pm \hat{s}(\hat{\psi}) z_{1-\alpha/2} = \{\hat{\psi}_l, \hat{\psi}_u\},$$

as described in Section 4.2.2. The corresponding limits for the odds ratio θ are $\{\exp(\hat{\psi}_l), \exp(\hat{\psi}_u)\}$. For the data shown in Figure 4.3, $\hat{\psi} = \log \hat{\theta} = .6104$, and $\hat{s}(\hat{\psi}) = 0.0639$, so the 95% limits for θ are $\{1.624, 2.087\}$, as shown by the calculations below. The same result is returned by `confint()` for an "oddsratio" object.

```
summary(oddsratio(Berkeley))

## Error in summary(oddsratio(Berkeley)): could not find function "oddsratio"

exp(.6103 + c(-1, 1) * qnorm(.975) * 0.06398)

## [1] 1.6240 2.0869

confint(oddsratio(Berkeley, log=FALSE))

## Error in confint(oddsratio(Berkeley, log = FALSE)): could not find function
"oddsratio"
```

Now consider how to find a 2×2 table whose frequencies correspond to the odds ratios at the limits of the confidence interval. A table standardized to equal row and column margins can be represented by the 2×2 matrix with entries

$$\begin{bmatrix} p & (1-p) \\ (1-p) & p \end{bmatrix},$$

whose odds ratio is $\theta = p^2/(1-p)^2$. Solving for p gives $p = \sqrt{\theta}/(1 + \sqrt{\theta})$. The corresponding frequencies can then be found by adjusting the standardized table to have the same row and column margins as the data. The results of these computations which generate the confidence rings in Figure 4.3 are shown in Table 4.7.

Table 4.7: Odds ratios and equivalent tables for 95% confidence rings for the Berkeley data.

{tab:berkod}

	Odds Ratio	Standardized Table		Equivalent Frequencies	
Lower limit	1.624	0.560	0.440	1167.1	587.9
		0.440	0.560	1523.9	1247.1
Data	1.841	0.576	0.424	1198.0	557.0
		0.424	0.576	1493.0	1278.0
Upper limit	2.087	0.591	0.409	1228.4	526.6
		0.409	0.591	1462.6	1308.4

4.4.2 Stratified analysis for $2 \times 2 \times k$ tables

{sec:twoway}

In a $2 \times 2 \times k$ table, the last dimension often corresponds to “strata” or populations, and it is typically of interest to see if the association between the first two variables is homogeneous across strata. For such tables, simply make one fourfold panel for each stratum. The standardization of marginal frequencies is designed to allow easy visual comparison of the pattern of association when the marginal frequencies vary across two or more populations

The admissions data shown in Figure 4.2 and Figure 4.3 were actually obtained from six departments —the six largest at Berkeley (Bickel *et al.*, 1975). To determine the source of the apparent sex bias in favor of males, we make a new plot, Figure 4.4, stratified by department.

```
# fourfold display
UCB <- aperm(UCBAdmissions, c(2, 1, 3))
fourfold(UCB, mfrow=c(2, 3))

## Error in eval(expr, envir, enclos): could not find function "fourfold"
```

Surprisingly, Figure 4.4 shows that, for five of the six departments, the odds of admission is approximately the same for both men and women applicants. Department A appears to differ from the others, with women approximately 2.86 ($= (313/19)/(512/89)$) times as likely to gain admission. This appearance is confirmed by the confidence rings, which in Figure 4.4 are joint³ 95% intervals for θ_c , $c = 1, \dots, k$.

This result, which contradicts the display for the aggregate data in Figure 4.2, is a nice example of *Simpson’s paradox*⁴, and illustrates clearly why an overall analysis of a three- (or higher-) way table can be misleading. The resolution of this contradiction can be found in the large differences in admission rates among departments. Men and women apply to different departments differentially, and in these data women happen to apply in larger numbers to departments that have a low acceptance rate. The aggregate results are misleading because they falsely assume

³For multiple-strata plots, `fourfold()` by default adjusts the significance level for multiple testing, using Holm’s (1979) method provided by `p.adjust()`.

⁴Simpson’s paradox (Simpson, 1951) occurs in a three-way table, $[A, B, C]$, when the marginal association between two variables, A, B collapsing over C differs in *direction* from the partial association $A, B|C = c_k$ at the separate levels of C . Strictly speaking, Simpson’s paradox would require that for all departments separately the odds ratio $\theta_k < 1$ (which occurs for Departments A, B, D, and F in Figure 4.4) while in the aggregate data $\theta > 1$.

men and women are equally likely to apply in each field.⁵

Visualization principles

TODO: Move this to Ch. 1 An important principle in the display of large, complex data sets is *controlled comparison*—we want to make comparisons against a clear standard, with other things held constant. The fourfold display differs from a pie chart in that it holds the angles of the segments constant and varies the radius. An important consequence is that we can quite easily compare a series of fourfold displays for different strata, since corresponding cells of the table are always in the same position. As a result, an array of fourfold displays serve the goals of comparison and detection better than an array of pie charts.

Moreover, it allows the observed frequencies to be standardized by equating either the row or column totals, while preserving the design goal for this display—the odds ratio. In Figure 4.4, for example, the proportion of men and women, and the proportion of accepted applicants were equated visually in each department. This provides a clear standard which also greatly facilitates controlled comparison.

Another principle is *visual impact*—we want the important features of the display to be easily distinguished from the less important (Tukey, 1993). Figure 4.4 distinguishes the one department for which the odds ratio differs significantly from 1 by shading intensity, even though the same information can be found by inspection of the confidence rings.

{ex:wheeze1}

EXAMPLE 4.11: Breathlessness and wheeze in coal miners

The various ways of standardizing a collection of 2×2 tables allows visualizing relations with different factors (row percentages, column percentages, strata totals) controlled. However, different kinds of graphs can speak more eloquently to other questions by focusing more directly on the odds ratio.

Agresti (2002, Table 9.8) cites data from Ashford and Sowden (1970) on the association between two pulmonary conditions, breathlessness and wheeze, in a large sample of coal miners. The miners are classified into age groups, and the question treated by Agresti is whether the association between these two symptoms is homogeneous over age. These data are available in the `CoalMiners` data in `vcd`, a $2 \times 2 \times 9$ frequency table. The first group, aged 20-24 has been omitted from these analyses.

```
data("CoalMiners", package="vcd")
CM <- CoalMiners[, , 2:9]
ftable(CM, row.vars = 3)
```

	Breathlessness	B		NoB	
	Wheeze	W	NoW	W	NoW
Age					
25-29		23	9	105	1654
30-34		54	19	177	1863
35-39		121	48	257	2357
40-44		169	54	273	1778
45-49		269	88	324	1712
50-54		404	117	245	1324
55-59		406	152	225	967
60-64		372	106	132	526

⁵This explanation ignores the possibility of structural bias against women, e.g., lack of resources allocated to departments that attract women applicants.

The question of interest can be addressed by displaying the odds ratio in the 2×2 tables with the margins of breathlessness and wheeze equated (i.e., with the default `std='margins'` option), which gives the graph shown in Figure 4.5. Although the panels for all age groups show an overwhelmingly positive association between these two symptoms, one can also (by looking carefully) see that the strength of this association declines with increasing age.

```
fourfold(CM, mfc01 = c(2,4))

## Error in eval(expr, envir, enclos): could not find function "fourfold"
```

However, note that the pattern of change over age is somewhat subtle compared to the dominant positive association within each panel. When the goal is to display how the odds ratio varies with a quantitative factor such as age, it is often better to simply calculate and plot the odds ratio directly.

The `oddsratio()` function in `vcd` calculates odds ratios for $2 \times 2 (\times k)$ tables. By default, it returns the log odds. Use the option `log=FALSE` to get the odds ratios themselves. It is easy to see that the (log) odds ratios decline with age.

```
oddsratio(CM)

## Error in eval(expr, envir, enclos): could not find function "oddsratio"

oddsratio(CM, log=FALSE)

## Error in eval(expr, envir, enclos): could not find function "oddsratio"
```

When the analysis goal is to understand how the odds ratio varies with a stratifying factor (which could be a quantitative variable), it is often better to plot the odds ratio directly.

The lines below use the `plot()` method for "oddsratio" objects. This produces a line graph of the log odds ratio against the stratum variable, together with confidence interval error bars. In addition, because age is a quantitative variable, we can calculate, and display the fitted relation for a linear model relating `lodds` to age. Here, we try using a quadratic model (`poly(age, 2)`) mainly to see if the trend is nonlinear.

```
lodds <- oddsratio(CM)

## Error in eval(expr, envir, enclos): could not find function "oddsratio"

plot(lodds, lwd=2, cex=1.25, pch=16,
     xlab = "Age Group",
     main = "Breathlessness and Wheeze in Coal Miners")

## Error in plot(lodds, lwd = 2, cex = 1.25, pch = 16, xlab = "Age Group",
## : object 'lodds' not found

age <- seq(25, 60, by = 5)
mod <- lm(lodds ~ poly(age, 2))

## Error in eval(expr, envir, enclos): object 'lodds' not found

lines(fitted(mod), col = "red", lwd=2)

## Error in fitted(mod): object 'mod' not found
```

In Figure 4.6, it appears that the decline in the log odds ratio levels off with increasing age. One virtue of fitting the model in this way is that we can test the additional contribution of the quadratic term, which turns out to be insignificant.

```
summary(mod)
```

```
## Error in summary(mod): object 'mod' not found
```

△

4.5 Sieve diagrams

{sec:twoway-sieve}

The wise ones fashioned speech with their thought, sifting it as grain is sifted through a sieve.

Buddha

For two- (and higher-) way contingency tables, the design principles of perception, detection, and comparison (see Chapter 1) suggest that we should try to show the observed frequencies in relation to what we would expect those frequencies to be under a reasonable null model—for example, the hypothesis that the row and column variables are unassociated.

To this end, several schemes for representing contingency tables graphically are based on the fact that when the row and column variables are independent, the estimated expected frequencies, m_{ij} , are products of the row and column totals (divided by the grand total).

$$m_{ij} = \frac{n_{i+}n_{+j}}{n_{++}} .$$

Then, each cell can be represented by a rectangle whose area shows the observed cell frequency, n_{ij} , expected frequency, m_{ij} , or deviation (residual) from independence, $n_{ij} - m_{ij}$. Visual attributes (color, shading) of the rectangles can be used to highlight the pattern of association.

For example, for any two-way table, the expected frequencies under independence can be represented by rectangles whose widths are proportional to the total frequency in each column, n_{+j} , and whose heights are proportional to the total frequency in each row, n_{i+} ; the area of each rectangle is then proportional to m_{ij} . Figure 4.7 (left) shows the expected frequencies for the hair and eye color data (Table 4.2), calculated using `independence_table()` in `vcd`.

```
haireye <- margin.table(HairEyeColor, 1:2)
expected = independence_table(haireye)
```

```
## Error in eval(expr, envir, enclos): could not find function "independence_table"
```

```
round(expected, 1)
```

```
## Error in eval(expr, envir, enclos): object 'expected' not found
```

```
## Error in eval(expr, envir, enclos): could not find function "sieve"
```

```
## Error in eval(expr, envir, enclos): could not find function "sieve"
```

Figure 4.7 (left) simply represents the model—what the frequencies would be if hair color and eye color were independent—not the data. Note, however, that the rectangles are cross-ruled so that the number of boxes in each (counting up the fractional bits) equals the expected frequency with which the cell is labeled, and moreover, the rulings are equally spaced in all cells. Hence, cross-ruling the cells to show the observed frequency would give a data display which implicitly compares observed and expected frequencies as shown in Figure 4.7 (right).

Riedwyl and Schüpbach 1983, 1994 proposed a *sieve diagram* (later called a *parquet diagram*) based on this principle. In this display the area of each rectangle is always proportional to expected frequency but observed frequency is shown by the number of squares in each rectangle, as in Figure 4.7 (right).

Hence, the difference between observed and expected frequency appears as variations in the density of shading. Cells whose observed frequency n_{ij} exceeds the expected m_{ij} appear denser than average. The pattern of positive and negative deviations from independence can be more easily seen by using color, say, red for negative deviations, and blue for positive.⁶

{ex:haireye2}

EXAMPLE 4.12: Hair color and eye color

The sieve diagram for hair color and eye color shown in Figure 4.7 (right) can be interpreted as follows: The pattern of color and shading shows the high frequency of blue-eyed blonds and people with brown eyes and dark hair. People with hazel eyes are also more likely to have red or brown hair, and those with green eyes more likely to have red or blond hair, than would be observed under independence. △

{ex:vision1}

EXAMPLE 4.13: Visual acuity

In World War II, all workers in the U.K. Royal Ordnance factories were given test of visual acuity (unaided distance vision) of their left and right eyes on a 1 (high) to 4 (low) scale. The dataset `VisualAcuity` in `vcd` gives the results for 10,719 workers (3242 men, 7477 women) aged 30-39.

Figure 4.8 shows the sieve diagram for data from the larger sample of women (Kendall and Stuart (1961, Table 33.5), Bishop *et al.* (1975, p. 284)). The `VisualAcuity` data is a frequency data frame and we first convert it to table form (`VA.tab`), a $4 \times 4 \times 2$ table to re-label the variables and levels. **Done:** [Make this an exercise in Ch. 2](#)

```
# re-assign names/dimnames
data("VisualAcuity", package="vcd")
VA.tab <- xtabs(Freq ~ right + left + gender, data=VisualAcuity)
dimnames(VA.tab)[1:2] <- list(c("high", 2, 3, "low"))
names(dimnames(VA.tab))[1:2] <- paste(c("Right", "Left"), "eye grade")
#str(VA.tab)
```

```
sieve(VA.tab[,,"female"], shade=TRUE)
```

```
## Error in eval(expr, envir, enclos): could not find function "sieve"
```

The diagonal cells show the obvious: people tend to have the same visual acuity in both eyes, and there is strong lack of independence. The off diagonal cells show a more subtle pattern that

⁶Positive residuals are also shown by solid lines, negative residuals by broken lines, so that they may still be distinguished in monochrome versions.

suggests symmetry—the cells below the diagonal are approximately equally dense as the corresponding cells above the diagonal. Moreover, the relatively consistent pattern on the diagonals $\pm 1, \pm 2, \dots$ away from the main diagonals suggests that the association may be explained in terms of the *difference* in visual acuity between the two eyes.

These suggestions can be tested by fitting intermediate models between the null model of independence (which fits terribly) and the saturated model (which fits perfectly), as we shall see later in this book. A model of *quasi-independence*, for example (see Example 8.8 in Chapter 8) ignores the diagonal cells and tests whether independence holds for the remainder of the table. The *symmetry* model for a square table allows association, but constrains the expected frequencies above and below the main diagonal to be equal. Such models provide a way of testing *specific* explanatory models that relate to substantive hypotheses and what we observe in our visualizations. These and other models for square tables are discussed further in Section 8.7. \triangle

4.5.1 Larger tables: The strucplot framework

{sec:twoway-sieve-lar

The implementation of sieve diagrams in `vcd` is far more general than illustrated in the examples above. For one thing, the `sieve` function has a formula method, which allows one to specify the variables in the display as a model formula. For example, for the `VisualAcuity` data, a plot of the (marginal) frequencies for left and right eye grades pooling over gender can be obtained with the call below (this plot is not shown).

```
sieve(Freq ~ right + left, data = VisualAcuity, shade=TRUE)
```

More importantly, sieve diagrams are just one example of the *strucplot framework*, a general system for visualizing n -way frequency tables in a hierarchical way. We describe this framework in more detail in Section 5.3 in context of mosaic displays. For now, we just illustrate the extension of the formula method to provide for conditioning variables. In the call below, the formula `Freq ~ right + left | gender` means to produce a separate block in the plot for the levels of *gender*.⁷

```
sieve(Freq ~ right + left | gender, data = VisualAcuity, shade=TRUE)
## Error in eval(expr, envir, enclos): could not find function "sieve"
```

In Figure 4.9, the relative sizes of the blocks for the conditioning variable (*gender*) show the much larger number of women than men in this data. Within each block, color and density of the box rules shows the association of left and right acuity, and it appears that the pattern for men is similar to that observed for women. The methods described in Section 4.3.1 can be used to test the hypothesis of homogeneity of association, and loglinear models described in Chapter 8 provide specific tests of hypotheses of *symmetry*, *quasi-independence* and other models for structured associations.

{ex:berkeley3}

EXAMPLE 4.14: Berkeley admissions

This example illustrates some additional flexibility of sieve plots with the strucplot framework, using the Berkeley admissions data. The left panel of Figure 4.10 shows the sieve diagrams for the relation between department and admission, conditioned by gender. It can easily be seen

⁷An equivalent plot, but one labeled more nicely, as in Figure 4.8 can be produced from the `VA.tab` table using `sieve(VA.tab, shade=TRUE, condvar='gender')`.

that (a) overall, there were more male applicants than female; (b) there is a moderately similar pattern of observed > expected (blue) for males and females.

```
# conditioned on gender
sieve(UCBAdmissions, shade=TRUE, condvar='Gender')

## Error in eval(expr, envir, enclos): could not find function "sieve"

# three-way table, Department first, with cell labels
UCB <- aperm(UCBAdmissions, c(3,1,2))
dimnames(UCB)[[3]] <- c("M", "F") # abbreviate for display
sieve(UCB, shade=TRUE, pop=FALSE)

## Error in eval(expr, envir, enclos): could not find function "sieve"

labeling_cells(text = UCB, gp_text = gpar(fontface = 2))(UCB)

## Error in eval(expr, envir, enclos): could not find function "labeling_cells"
```

In the right panel of Figure 4.10, the three-way table was first permuted to make *Dept* the first splitting variable. Each 2×2 table of *Admit* by *Gender* then appears, giving a sieve diagram version of what we showed earlier in fourfold displays (Figure 4.4). The function `labeling_cells()` is used here to write the cell frequency in each rectangle.

Finally, for tables of more than two dimensions, there is a variety of different models for “independence,” and the strucplot framework allows these to be specified with the `expected` argument, either as an array of numbers conforming to the `data` argument, or as a model formula for `loglm()`.

For example, a sieve diagram may be used to determine if the association between gender and department is the same across departments by fitting the model `~Admit*Gender + Dept`, which says that *Dept* is independent of the combinations of *Admit* and *Gender*. This is done as shown below, giving the plot in Figure 4.11.

```
UCB2 <- aperm(UCBAdmissions, c(3,2,1))
sieve(UCB2, shade=TRUE, expected=~Admit*Gender + Dept,
      split_vertical=c(FALSE,TRUE,TRUE))

## Error in eval(expr, envir, enclos): could not find function "sieve"
```

In terms of the loglinear models discussed in Chapter 5, this is equivalent to fitting the model of *joint independence*, $[AdmitGender][Dept]$. Figure 4.11 shows the greater numbers of male applicants in departments A and B (whose overall rate of admission is high) and greater numbers of female applicants in the remaining departments (where the admission rate is low).

△

4.6 Association plots

{sec:twoway-assoc}

In the sieve diagram the foreground (rectangles) shows expected frequencies; deviations from independence are shown by color and density of shading. The *association plot* (Cohen, 1980,

Friendly, 1991) puts deviations from independence in the foreground: the area of each box is made proportional to the (observed – expected) frequency.

For a two-way contingency table, the signed contribution to Pearson χ^2 for cell i, j is

$$d_{ij} = \frac{n_{ij} - m_{ij}}{\sqrt{m_{ij}}} = \text{Pearson residual}, \quad \chi^2 = \sum_{ij} (d_{ij})^2 \quad (4.5)$$

In the association plot, each cell is shown by a rectangle, having:

- (signed) height $\sim d_{ij}$
- width = $\sqrt{m_{ij}}$.

so, the area of each cell is proportional to the raw residual, $n_{ij} - m_{ij}$. The rectangles for each row in the table are positioned relative to a baseline representing independence ($d_{ij} = 0$) shown by a dotted line. Cells with observed $>$ expected frequency rise above the line (and are colored blue); cells that contain less than the expected frequency fall below it (and are shaded red).

```
haireye <- margin.table(HairEyeColor, 1:2)
assoc(haireye, shade=TRUE)

## Error in eval(expr, envir, enclos): could not find function "assoc"
```

Figure 4.12 shows the association plot for the data on hair color and eye color. In constructing this plot, each rectangle is shaded according to the value of the Pearson residual Eqn. (4.5), using a simple scale shown in the legend, where residuals $|d_{ij}| > 2$ are shaded blue or red depending on their sign and residuals $|d_{ij}| > 4$ are shaded with a more saturated color.

One virtue of the association plot is that it is quite simple to interpret in terms of the pattern of positive and negative d_{ij} values. Bertin (1981) uses similar graphics to display large complex contingency tables. Like the sieve diagram, however, patterns of association are most apparent when the rows and columns of the display are ordered in a sensible way.

```
assoc(HairEyeColor, shade=TRUE)

## Error in eval(expr, envir, enclos): could not find function "assoc"
```

We note here that the association plot also belongs to the strucplot framework and thus extends to higher-way tables. For example, the full `HairEyeColor` table is also classified by `Sex`. The plot for the three-way table is shown in Figure 4.13. **TODO: Perhaps combine these two figures into two panels, side by side.** In this plot the third table variable (`Sex` here) is shown nested within the first two, allowing easy comparison of the profiles of hair and eye color for males and females.

4.7 Observer agreement

When the row and column variables represent different observers rating the same subjects or objects, interest is focused on *observer agreement* rather than mere association. In this case, measures and tests of agreement provide a method of assessing the reliability of a subjective classification or assessment procedure.

{sec:twoway-agree}

For example, two (or more) clinical psychologists might classify patients on a scale with categories (a) normal, (b) mildly impaired, (c) severely impaired. Or, ethologists might classify the behavior of animals in categories of cooperation, dominance and so forth, or paleologists might classify pottery fragments according to categories of antiquity or cultural groups. As these examples suggest, the rating categories are often ordered, but not always.

For two raters, a contingency table can be formed classifying all the subjects/objects rated according to the rating categories used by the two observers. In most cases, the same categories are used by both raters, so the contingency table is square, and the entries in the diagonal cells are the cases where the raters agree.

In this section we describe some measures of the strength of agreement and then a method for visualizing the pattern of agreement. But first, the following examples show some typical agreement data.

{ex:sexisfun1}

EXAMPLE 4.15: Sex is fun

The `SexualFun` table in `vcd` (Agresti (1990, Table 2.10), from Hout *et al.* (1987)) summarizes the responses of 91 married couples to a questionnaire item: “Sex is fun for me and my partner: (a) Never or occasionally, (b) fairly often, (c) very often, (d) almost always. ”

```
data("SexualFun", package="vcd")
SexualFun

##           Wife
## Husband   Never Fun Fairly Often Very Often Always fun
## Never Fun           7           7           2           3
## Fairly Often       2           8           3           7
## Very Often         1           5           4           9
## Always fun         2           8           9          14
```

In each row the diagonal entry is not always the largest, though it appears that the partners tend to agree more often when either responds “almost always”. △

{ex:MS1}

EXAMPLE 4.16: Diagnosis of MS patients

Landis and Koch (1977) gave data on the diagnostic classification of multiple sclerosis (MS) patients by two neurologists, one from Winnipeg and one from New Orleans. There were two samples of patients, 149 from Winnipeg and 69 from New Orleans, and each neurologist classified all patients into one of four diagnostic categories: (a) Certain MS, (b) Probable MS, (c) Possible MS, (d) Doubtful, unlikely, or definitely not MS

These data are available in `MSPatients`, a $4 \times 4 \times 2$ table, as shown below. It is convenient to show the data in separate slices for the Winnipeg and New Orleans patients:

```
MSPatients[,,"Winnipeg"]

## Error in eval(expr, envir, enclos): object 'MSPatients' not found

MSPatients[,,"New Orleans"]

## Error in eval(expr, envir, enclos): object 'MSPatients' not found

apply(MSPatients, 3, sum)           # show sample sizes

## Error in apply(MSPatients, 3, sum): object 'MSPatients' not found
```

In this example, note that the distribution of degree of severity of MS may differ between the two patient samples. As well, for a given sample, the two neurologists may be more or less strict about the boundaries between the rating categories.

△

4.7.1 Measuring agreement

{sec:agreemeas}

In assessing the strength of *agreement* we usually have a more stringent criterion than in measuring the strength of *association*, because observers ratings can be strongly associated without strong agreement. For example, one rater could use a more stringent criterion and thus consistently rate subjects one category lower (on an ordinal scale) than another rater.

More generally, measures of agreement must take account of the marginal frequencies with which two raters use the categories. If observers tend to use the categories with different frequency, this will affect measures of agreement.

Here we describe some simple indices that summarize agreement with a single score (and associated standard errors or confidence intervals). Von Eye and Mun (2006) treat this topic from the perspective of loglinear models.

Intraclass correlation

An analysis of variance framework leads to the **intraclass correlation** as a measure of inter-rater reliability, particularly when there are more than two raters. This approach is not covered here, but various applications are described by Shrout and Fleiss (1979), and implemented in R in `ICC()` in the `psych` package.

Cohen's Kappa

Cohen's kappa (κ) (Cohen, 1960, 1968) is a commonly used measure of agreement that compares the observed agreement to agreement expected by chance if the two observer's ratings were independent. If p_{ij} is the probability that a randomly selected subject is rated in category i by the first observer and in category j by the other, then the observed agreement is the sum of the diagonal entries, $P_o = \sum_i p_{ii}$. If the ratings were independent, this probability of agreement (by chance) would be $P_c = \sum_i p_{i+} p_{+i}$. Cohen's κ is then the ratio of the difference between actual agreement and chance agreement, $P_o - P_c$, to the maximum value this difference could obtain:

$$\kappa = \frac{P_o - P_c}{1 - P_c} . \quad (4.6) \quad \{\text{eq:kappa}\}$$

When agreement is perfect, $\kappa = 1$; when agreement is no better than would be obtained from statistically independent ratings, $\kappa = 0$. κ could conceivably be negative, but this rarely occurs in practice. The minimum possible value depends on the marginal totals.

For large samples (n_{++}), κ has an approximate normal distribution when $H_0 : \kappa = 0$ is true and its standard error (Fleiss, 1973, Fleiss *et al.*, 1969) is given by

$$\hat{\sigma}(\kappa) = \frac{P_c + P_c^2 - \sum_i p_{i+} p_{+i} (p_{i+} + p_{+i})}{n_{++} (1 - P_c)^2} .$$

Hence, it is common to conduct a test of $H_0 : \kappa = 0$ by referring $z = \kappa/\hat{\sigma}(\kappa)$ to a unit normal distribution. The hypothesis of agreement no better than chance is rarely of much interest, however. It is preferable to estimate and report a confidence interval for κ .

Weighted Kappa

The original (unweighted) κ only counts strict agreement (the same category is assigned by both observers). A weighted version of κ (Cohen, 1968) may be used when one wishes to allow for *partial* agreement. For example, exact agreements might be given full weight, one-category difference given weight 1/2. This typically makes sense only when the categories are *ordered*, as in severity of diagnosis.

Weighted κ uses weights, $0 \leq w_{ij} \leq 1$ for each cell in the table, with $w_{ii} = 1$ for the diagonal cells. In this case P_o and P_c are defined as weighted sums

$$P_o = \sum_i \sum_j w_{ij} p_{ij}$$

$$P_c = \sum_i \sum_j w_{ij} p_{i+p+j}$$

and these weighted sums are used in Eqn. (4.6).

For an $r \times r$ table, two commonly-used pattern of weights are those based on equal spacing of weights (Cicchetti and Allison, 1971) for a near-match, and *Fleiss-Cohen weights* (Fleiss and Cohen, 1972), based on an inverse-square spacing,

$$w_{ij} = 1 - \frac{|i-j|}{r-1} \quad \text{equal spacing}$$

$$w_{ij} = 1 - \frac{|i-j|^2}{(r-1)^2} \quad \text{Fleiss-Cohen}$$

The Fleiss-Cohen weights attach greater importance to near disagreements, as you can see below for a 4×4 table. These weights also provide a measure equivalent to the intraclass correlation.

Integer Spacing Cicchetti Allison weights				Inverse Square Spacing Fleiss-Cohen weights			
-----				-----			
1	2/3	1/3	0	1	8/9	5/9	0
2/3	1	2/3	1/3	8/9	1	8/9	5/9
1/3	2/3	1	2/3	5/9	8/9	1	8/9
0	1/3	2/3	1	0	5/9	8/9	1

Computing Kappa

The function `Kappa()` in `vcd` calculates unweighted and weighted **Kappa**. The `weights` argument can be used to specify the weighting scheme as either "Equal-Spacing" or "Fleiss-Cohen". The function returns a "Kappa" object, for which there is a `confint.Kappa()` method, providing confidence intervals. The `summary.Kappa()` method also prints the weights.

The lines below illustrate `Kappa` for the `SexualFun` data.

```
Kappa (SexualFun)
## Error in eval(expr, envir, enclos): could not find function "Kappa"
confint (Kappa (SexualFun))
## Error in confint (Kappa (SexualFun)): could not find function "Kappa"
```

4.7.2 Observer Agreement Chart

{sec:twoway-Bangdiwal

The observer agreement chart proposed by Bangdiwala (1985, 1987) provides a simple graphic representation of the strength of agreement in a contingency table, and alternative measures of strength of agreement with an intuitive interpretation. More importantly, it shows the *pattern* of disagreement when agreement is less than perfect.

The agreement chart is constructed as an $n \times n$ square, where $n = n_{++}$ is the total sample size. Black squares, each of size $n_{ii} \times n_{ii}$, show observed agreement. These are positioned within k larger rectangles, each of size $n_{i+} \times n_{+i}$ as shown in the left panel of Figure 4.14. The large rectangle shows the maximum possible agreement, given the marginal totals. Thus, a visual impression of the strength of agreement is given by

$$B = \frac{\text{area of dark squares}}{\text{area of rectangles}} = \frac{\sum_i^k n_{ii}^2}{\sum_i^k n_{i+} n_{+i}} \tag{4.7} \quad \text{{eq:bangb}}$$

When there is perfect agreement, the k rectangles determined by the marginal totals are all squares, completely filled by the shaded squares reflecting the diagonal n_{ii} entries, and $B = 1$.

```
agreementplot (SexualFun, main="Unweighted", weights=1)
## Error in eval(expr, envir, enclos): could not find function "agreementplot"
agreementplot (SexualFun, main="Unweighted")
## Error in eval(expr, envir, enclos): could not find function "agreementplot"
```

Partial agreement

Partial agreement is allowed by including a weighted contribution from off-diagonal cells, b steps from the main diagonal. For a given cell frequency, n_{ij} , a pattern of weights, w_1, w_2, \dots, w_b is applied to the cell frequencies as shown schematically below:

$$\begin{array}{ccccccc}
 & & n_{i-b,i} & & & & w_b \\
 & & \vdots & & & & \vdots \\
 n_{i,i-b} & \cdots & n_{i,i} & \cdots & n_{i,i+b} & \Leftarrow & w_b \cdots 1 \cdots w_b \\
 & & \vdots & & & & \vdots \\
 & & n_{i-b,i} & & & & w_b
 \end{array}$$

These weights are incorporated in the agreement chart (right panel of Figure 4.14) by successively lighter shaded rectangles whose size is proportional to the sum of the cell frequencies,

denoted A_{bi} , shown above. A_{1i} allows 1-step disagreements, using weights 1 and w_1 ; A_{2i} includes 2-step disagreements, etc. From this, one can define a weighted measure of agreement, B^w , analogous to weighted κ :

$$B^w = \frac{\text{weighted sum of areas of agreement}}{\text{area of rectangles}} = 1 - \frac{\sum_i^k [n_{i+}n_{+i} - n_{ii}^2 - \sum_{b=1}^q w_b A_{bi}]}{\sum_i^k n_{i+}n_{+i}}$$

where w_b is the weight for A_{bi} , the shaded area b steps away from the main diagonal, and q is the furthest level of partial disagreement to be considered.

The function `agreementplot()` actually calculates both B and B^w and returns them invisibly as the result of the call. The results, $B = 0.146$, and $B^w = 0.498$, indicate a stronger degree of agreement when 1-step disagreements are included.

```
B <-agreementplot(SexualFun)
## Error in eval(expr, envir, enclos): could not find function "agreementplot"
unlist(B) [1:2]
## Error in unlist(B): object 'B' not found
```

{ex:mammograms}

EXAMPLE 4.17: Mammogram ratings

The Mammograms data in `vcdExtra` gives a 4×4 table of (probably contrived) ratings of 110 mammograms by two raters from Kundel and Polansky (2003), used to illustrate the calculation and interpretation of agreement measures in this context.⁸

```
data("Mammograms", package="vcdExtra")
B <- agreementplot(Mammograms, main="Mammogram ratings")
## Error in eval(expr, envir, enclos): could not find function "agreementplot"
```

The agreement plot in Figure 4.15 shows substantial agreement among the two raters, particularly when one-step disagreements are taken into account. Careful study of this graph shows that the two raters more often agree exactly for the extreme categories of “Absent” and “Severe.” The amounts of unweighted and weighted agreement are shown numerically in the B and B^w statistics.

```
unlist(B) [1:2]
## Error in unlist(B): object 'B' not found
```

△

4.7.3 Observer bias in agreement

With an ordered scale, it may happen that one observer consistently tends to classify the objects into higher or lower categories than the other, perhaps due to using stricter thresholds for the

⁸In practice, of course, rater agreement on severity of diagnosis from radiology images varies with many factors. See Antonio and Crespi (2010) for a meta-analytic study concerning agreement in breast cancer diagnosis.

boundaries between adjacent categories. This bias produces differences in the marginal totals, n_{i+} , and n_{+i} and decreases the maximum possible agreement. While special tests exist for *marginal homogeneity*, the observer agreement chart shows this directly by the relation of the dark squares to the diagonal line: When the marginal totals are the same, the squares fall along the diagonal. The measures of agreement, κ and B , cannot determine whether lack of agreement is due to such bias, but the agreement chart can detect this.

{ex:MS2}

EXAMPLE 4.18: Diagnosis of MS patients

Agreement charts for both patient samples in the `MSPatients` data are shown in Figure 4.16. The `agreementplot()` function only handles two-way tables, so we do these separately by indexing on the last dimension (*Patients*).

```
agreementplot(MSPatients[,,"Winnipeg"], main="Winnipeg patients")
## Error in eval(expr, envir, enclos): could not find function "agreementplot"
agreementplot(MSPatients[,,"New Orleans"], main="New Orleans patients")
## Error in eval(expr, envir, enclos): could not find function "agreementplot"
```

It can be seen that, for both groups of patients, the rectangles for the two intermediate categories lie largely below the diagonal line (representing equality). This indicates that the Winnipeg neurologist tends to classify patients into more severe diagnostic categories. The departure from the diagonal is greater for the Winnipeg patients, for whom the Winnipeg neurologist uses the two most severe diagnostic categories very often, as can also be seen from the marginal totals printed in the plot margins.

Nevertheless there is a reasonable amount of agreement if one-step disagreements are allowed, as can be seen in `figreffig:MS-agree` and quantified in the B^w statistics below. The agreement charts also serve to explain why the B measures for exact agreement are so much lower.

```
agr1 <- agreementplot(MSPatients[,,"Winnipeg"])
## Error in eval(expr, envir, enclos): could not find function "agreementplot"
agr2 <- agreementplot(MSPatients[,,"New Orleans"])
## Error in eval(expr, envir, enclos): could not find function "agreementplot"
rbind(Winnipeg=unlist(agr1), NewOrleans=unlist(agr2))[,1:2]
## Error in unlist(agr1): object 'agr1' not found
```

△

4.8 Trilinear plots

{sec:twoway-trilinear}

The *trilinear plot* (also called a *ternary diagram* or *trinomial plot*) is a specialized display for a 3-column contingency table or for three variables whose relative proportions are to be displayed.

Individuals may be assigned to one of three diagnostic categories, for example, or a chemical process may yield three constituents in varying proportions, or we may look at the division of votes among three parties in a parliamentary election. This display is useful, therefore, for both frequencies and proportions.

Trilinear plots are featured prominently in Aitchison (1986), who describes statistical models for this type of *compositional data*. Upton (1976, 1994) uses them in detailed analyses of spatial and temporal changes in British general elections. Wainer (1996) reviews a variety of other uses of trilinear plots and applies them to aid in understanding the distributions of students achievement in the National Assessment of Educational Progress, making some aesthetic improvements to the traditional form of these plots along the way.

A trilinear plot displays each observation as a point inside an equilateral triangle whose coordinate corresponds to the relative proportions in each column. The three vertices represent the three extremes when 100% occurs in one of the three columns; a point in the exact center corresponds to equal proportions of $\frac{1}{3}$ in all three columns. For instance, Figure 4.1 shows three points whose compositions of three variables, A, B, and C are given in the data frame DATA below.

```
library(ggtern)
DATA <- data.frame(
  A = c(40, 20, 10),
  B = c(30, 60, 10),
  C = c(30, 20, 80),
  id = c("1", "2", "3"))
ggtern(data = DATA,
       mapping = aes(x=C, y=A, z=B,
                    label=id, colour=id)) +
  geom_point(aes(size=2)) +
  geom_text(vjust=-.5, size=8) +
  theme_tern_rgbw() +
  theme(plot.margin=unit(c(0, 0, 0, 0), "mm")) +
  guides(size = "none")
```

Note that each apex corresponds to 100% of the labeled variable, and the percentage of this variable decrease linearly along a line to the midpoint of the opposite baseline. The grid lines in the figure show the percentage value along each axis.

The construction of trilinear plots is described in detail in http://en.wikipedia.org/wiki/Ternary_plot. Briefly, let $P(a, b, c)$ represent the three components normalized so that $a+b+c = 1.0$. If the apex corresponding to Point A in Figure 4.1 is given (x, y) coordinates of $(x_A, y_A) = (0, 0)$, and those at apex B are $(x_B, y_B) = (100, 0)$, then the coordinates of apex C are $(x_C, y_C) = (50, 50\sqrt{3})$. The cartesian coordinates (x_P, y_P) of point P are then calculated as

$$y_P = c y_C$$

$$x_P = y_P \left(\frac{y_C - y_B}{x_C - x_B} \right) + \frac{\sqrt{3}}{2} y_C (1 - a)$$

In R, trilinear plots are implemented in the `triplot()` function in the `TeachingDemos` package, and also in the `ggtern` package, an extension of the `ggplot2` framework. The latter is much more flexible, because it inherits all of the capabilities of `ggplot2` for plot annotations, faceting, and layers. In essence, the function `ggtern()` is just a wrapper for `ggplot(...)` which adds a change in the coordinate system from cartesian (x, y) coordinates to the ternary coordinate system with `coord_tern()`.




Figure 4.1: A trilinear plot showing three points, for variables A, B, C. ^{fig:tripdemo2}

For example, the following code⁹ creates a data frame DATA containing 100 uniformly distributed random points. It uses `stat_density2d()` to draw contours of the densities of the points in the trilinear space.

```
set.seed(1)
DATA <- data.frame(x = runif(100),
                  y = runif(100),
                  z = runif(100))
plot <- ggtern(data = DATA,
              aes(x, y, z))
plot + stat_density2d(method = "lm", fullrange = T,
                    n = 200, geom = "polygon",
                    aes(fill = ..level..,
                       alpha = ..level..)) +
  geom_point() +
  theme_tern_rgbw() +
  labs(title = "Uniform data with density contours") +
  scale_fill_gradient(low = "blue", high = "red") +
  guides(color = "none", fill = "none", alpha = "none")
```

{ex:lifeboat1}

EXAMPLE 4.19: Lifeboats on the *Titanic*

We examine the question of who survived and why in the sinking of the *RMS Titanic* in Section 5.4 (Example ??), where we analyze a four-way table, *Titanic*, of the 2201 people on board (1316 passengers and 885 crew), classified by Class, Sex, Age, and Survival. A related data set, *Lifeboats* in *vcd* tabulates the survivors according to the life boats on which they were loaded. This data sheds some additional light on the issue of survival and provides a nice illustration of trilinear plots.

⁹This example was taken from the ggtern web site, <http://ggtern.com/2013/12/12/patched-density-functions-2/>.




Figure 4.2: A trilinear plot with density contours^{fig:ggterm-demo}

A bit of background: after the disaster, the British Board of Trade launched several inquiries, the most comprehensive of which resulted in the *Report on the Loss of the “Titanic” (S.S.)* by Lord Mersey (Mersey, 1912).¹⁰ The data frame `Lifeboats` in `vcd` contains the data listed on p. 38 of that report.¹¹

Of interest here is the composition of the boats by the three categories, men, women and children and crew, and according to the launching of the boats from the port or starboard side. This can be shown in a trilinear display using the following statements. The plot, shown in Figure 4.3, has most of the points near the top, corresponding to a high percentage of women and children. We create a variable, `id`, used to label those boats with more than 10% male passengers. In the `ggplot2` framework, plot aesthetics, such as color and shape can be mapped to variables in the data set, and here we map these both to `side` of the boat.

```
data("Lifeboats", package="vcd")
# label boats with more than 10% men
Lifeboats$id <- ifelse(Lifeboats$men/Lifeboats$total > .1,
                      as.character(Lifeboats$boat), "")
ggtern(data = Lifeboats,
        mapping = aes(x = women, y = men, z = crew,
                      colour=side, shape=side, label=id)) +
  theme_tern_rgbw() +
  theme(plot.margin=unit(c(0, 0, 0, 0), "mm")) +
  geom_point(aes(size=2)) +
```

¹⁰The *Titanic* was outfitted with 20 boats, half on each of the port and starboard sides, of which 14 were large lifeboats with a capacity of 65, two were emergency boats designed for 40 persons, and the remaining four were collapsible boats capable of holding 47, a total capacity of 1178 (considered adequate at that time). Two of the collapsible boats, lashed to the roof of the officers quarters, were ineffectively launched and utilized as rafts after the ship sunk. The report lists the time of launch and composition of the remaining 18 boats according to male passengers, women and children, and “men of crew”, as reported by witnesses.

¹¹The “data” lists a total of 854 in 18 boats, although only 712 were in fact saved. Mersey notes “it is obvious that these figures are quite unreliable”.

```
labs(title = "Lifeboats on the Titanic") +
labs(T="Women and children") +
guides(size = "none") +
geom_smooth(method="lm", size=1.5, aes(fill=side)) +
geom_text(vjust=1, color="black")
```




Figure 4.3: Lifeboats on the *Titanic*, showing the composition of each boat. Boats with more than 10% male passengers are labeled. fig:lifeboats1

The resulting plot in Figure 4.3, makes it immediately apparent that many of the boats launched from the port side differ substantially from the remaining boats, whose passengers were almost entirely women and children. Boat 1 had only 20% (2 out of 10) women and children, while the percentage for boat 3 was only 50% (25 out of 50). We highlight the difference in composition of the boats launched from the two sides by adding a linear regression smooth for the relation $\text{men} \sim \text{women}$.

The trilinear plot scales the numbers for each observation to sum to 1.0, so differences in the total number of people on each boat cannot be seen in Figure 4.3. The total number reported loaded is plotted against launch time in Figure 4.4, with a separate regression line and loess smooth fit to the data for the port and starboard sides.

```
ggplot(data = Lifeboats,
       aes(x=launch, y=total, colour=side, label=boat)) +
geom_smooth(method="lm", aes(fill=side), size=1.5) +
geom_smooth(method="loess", aes(fill=side), se=FALSE, size=1.2) +
geom_point() + ylim(c(0,100)) +
geom_text(vjust=-.5, color="black") +
labs(y="Total loaded", x="Launch time")
```




Figure 4.4: Number of people loaded on lifeboats on the Titanic vs. time of launch, by side of boat. The plot annotations show the linear regression and loess smooth.

From the linear regression lines in Figure 4.4, it seems that the rescue effort began in panic on the port side, with relatively small numbers loaded, and (from Figure 4.3), small proportions of women and children. But the loading regime on that side improved steadily over time. The procedures began more efficiently on the starboard side but the numbers loaded increased only slightly. The smoothed loess curves indicate that over time, for each side, there was still a large variability from boat to boat.

△

4.9 Chapter summary

{sec:toway-summary}

- A contingency table gives the frequencies of observations cross-classified by two or more categorical variables. With such data we are typically interested in testing whether associations exist, quantifying the strength of association, and understanding the nature of the association among these variables.
- For 2×2 tables, association is easily summarized in terms of the odds ratio or its logarithm. This measure can be extended to stratified $2 \times 2 \times k$ tables, where we can also assess whether the odds ratios are equal across strata or how they vary.
- For $r \times c$ tables, measures and tests of general association between two categorical variables are most typically carried out using the Pearson's chi-square or likelihood-ratio tests provided by `assocstats()`. Stratified tests controlling for one or more background variables, and tests for ordinal categories are provided by the Cochran-Mantel-Haenszel tests given by `CMHtest()`.
- For 2×2 tables, the fourfold display provides a visualization of the association between variables in terms of the odds ratio. Confidence rings provide a visual test of whether the odds ratio differs significantly from 1. Stratified plots for $2 \times 2 \times k$ tables are also provided by `fourfold()`.

- Sieve diagrams and association plots provide other useful displays of the pattern of association in $r \times c$ tables. These also extend to higher-way tables as part of the strucplot framework.
- When the row and column variables represent different observers rating the same subjects, interest is focused on agreement rather than mere association. Cohen's κ is one measure of strength of agreement. The observer agreement chart provides a visual display of how the observers agree and disagree.
- Another specialized display, the trilinear plot is useful for three-column frequency tables or compositional data.

4.10 Further reading

{sec:twoway-reading}

4.11 Lab exercises

{lab:twoway-lab}

Exercise 4.1 The data set `fat`, created below, gives a 2×2 table recording the level of cholesterol in diet and the presence of symptoms of heart disease for a sample of 23 people.

```
fat <- matrix( c(6, 4, 2, 11), 2, 2)
dimnames(fat) <- list(diet=c("LoChol", "HiChol"),
                     disease=c("No", "Yes"))
```

- Use `chisq.test(fat)` to test for association between diet and disease. Is there any indication that this test may not be appropriate here?
- Use a fourfold display to test this association visually. Experiment with the different options for standardizing the margins, using the `margin` argument to `fourfold()`. What evidence is shown in different displays regarding whether the odds ratio differs significantly from 1?
- `oddsratio(fat, log=FALSE)` will give you a numerical answer. How does this compare to your visual impression from fourfold displays?
- With such a small sample, Fisher's exact test may be more reliable for statistical inference. Use `fisher.test(fat)`, and compare these results to what you have observed before.
- Write a one-paragraph summary of your findings and conclusions for this data set.

{lab:4.2}

Exercise 4.2 The data set `Abortion` in `vcdExtra` gives a $2 \times 2 \times 2$ table of opinions regarding abortion in relation to sex and status of the respondent. This table has the following structure:

```
data("Abortion", package="vcdExtra")
str(Abortion)

## table [1:2, 1:2, 1:2] 171 152 138 167 79 148 112 133
## - attr(*, "dimnames")=List of 3
## ..$ Sex : chr [1:2] "Female" "Male"
## ..$ Status : chr [1:2] "Lo" "Hi"
## ..$ Support_Abortion: chr [1:2] "Yes" "No"
```

- Taking support for abortion as the outcome variable, produce fourfold displays showing the association with sex, stratified by status.

- (b) Do the same for the association of support for abortion with status, stratified by sex.
 (c) For each of the problems above, use `oddsratio()` to calculate the numerical values of the odds ratio, as stratified in the question.

{lab:4.3}

Exercise 4.3 The `JobSat` table on income and job satisfaction created in Example 2.5 is contained in the `vcdExtra` package.

- (a) Carry out a standard χ^2 test for association between income and job satisfaction. Is there any indication that this test might not be appropriate? Repeat this test using `simulate.p.value = TRUE` to obtain a Monte Carlo test that does not depend on large sample size. Does this change your conclusion?
 (b) Both variables are ordinal, so CMH tests may be more powerful here. Carry out that analysis. What do you conclude?

{lab:4.4}

Exercise 4.4 The `Hospital` data in `vcd` gives a 3×3 table relating the length of stay (in years) of 132 long-term schizophrenic patients in two London mental hospitals with the frequency of visits by family and friends.

- (a) Carry out a χ^2 test for association between the two variables.
 (b) Use `assocstats()` to compute association statistics. How would you describe the strength of association here?
 (c) Produce an association plot for these data, with visit frequency as the vertical variable. Describe the pattern of the relation you see here.
 (d) Both variables can be considered ordinal, so `CMHtest()` may be useful here. Carry out that analysis. Do any of the tests lead to different conclusions?

{lab:4.5}

Exercise 4.5 The two-way table `Mammograms` in `vcdExtra` gives ratings on the severity of diagnosis of 110 mammograms by two raters.

- (a) Assess the strength of agreement between the raters using Cohen's κ , both unweighted and weighted.
 (b) Use `agreementplot()` for a graphical display of agreement here.

{lab:4.6}

Exercise 4.6 Agresti and Winner (1997) gave the data in Table 4.8 on the ratings of 160 movies by the reviewers Gene Siskel and Roger Ebert for the period from April 1995 through September 1996. The rating categories were Con (“thumbs down”), Mixed and Pro (“thumbs up”).

Table 4.8: Movie ratings by Siskel & Ebert, April 1995–September 1996. *Source:* Agresti and Winner (1997)

{tab:siskel-ebert}

		Ebert			total
		Con	Mixed	Pro	
Siskel	Con	24	8	13	45
	Mixed	8	13	11	32
	Pro	10	9	64	83
total		42	30	88	160

- (a) Assess the strength of agreement between the raters using Cohen's κ , both unweighted and weighted.
 (b) Use `agreementplot()` for a graphical display of agreement here.

- (c) Assess the hypothesis that the ratings are *symmetric* around the main diagonal. *Hint*: Symmetry for a square table \mathbf{T} means that $t_{ij} = t_{ji}$ for $i \neq j$. The expected frequencies under the hypothesis of symmetry are the average of the off-diagonal cells, $\mathbf{E} = (\mathbf{T} + \mathbf{T}^T)/2$.

{lab:4.7}

Exercise 4.7 For the `VisualAcuity` data set:

- Use the code shown in the text to create the table form, `VA.tab`.
- Perform the CMH tests for this table.
- Use `loglm()` method described in Section 4.3.1 to test whether the association between left and right eye acuity can be considered the same for men and women.

{lab:4.8}

Exercise 4.8 The graph in Figure 4.4 may be misleading, in that it doesn't take account of the differing capacities of the 18 life boats on the *Titanic*, given in the variable `cap` in the `Lifeboats` data.

- Calculate a new variable, `pctloaded` as the percentage loaded relative to the boat capacity.
- Produce a plot similar to Figure 4.4, showing the changes over time in this measure.

```
detach(package:ggtern)
.locals$ch04 <- setdiff(ls(), .globals)
remove(list=locals$ch04[sapply(.locals$ch04, function(n) { !is.function(get(n)) })])
```


References

- Agresti, A. (1990). *Categorical Data Analysis*. New York: Wiley-Interscience.
- Agresti, A. (1996). *An Introduction to Categorical Data Analysis*. New York: Wiley Interscience.
- Agresti, A. (2002). *Categorical Data Analysis*. Wiley Series in Probability and Statistics. New York: Wiley-Interscience [John Wiley & Sons], 2nd edn.
- Agresti, A. (2013). *Categorical Data Analysis*. Wiley Series in Probability and Statistics. New York: Wiley-Interscience [John Wiley & Sons], 3rd edn.
- Agresti, A. and Winner, L. (1997). Evaluating agreement and disagreement among movie reviewers. *CHANCE*, 10(2), 10–14.
- Aitchison, J. (1986). *The Statistical Analysis of Compositional Data*. London: Chapman and Hall.
- Antonio, A. L. M. and Crespi, C. M. (2010). Predictors of interobserver agreement in breast imaging using the breast imaging reporting and data system. *Breast Cancer Research and Treatment*, 120(3), 539–546.
- Ashford, J. R. and Sowden, R. D. (1970). Multivariate probit analysis. *Biometrics*, 26, 535–546.
- Bangdiwala, S. I. (1985). A graphical test for observer agreement. In *Proceeding of the International Statistics Institute*, vol. 1, (pp. 307–308). Amsterdam: ISI.
- Bangdiwala, S. I. (1987). Using SAS software graphical procedures for the observer agreement chart. *Proceedings of the SAS User's Group International Conference*, 12, 1083–1088.
- Bertin, J. (1981). *Graphics and Graphic Information-processing*. New York: de Gruyter. (trans. W. Berg and P. Scott).
- Bickel, P. J., Hammel, J. W., and O'Connell, J. W. (1975). Sex bias in graduate admissions: Data from Berkeley. *Science*, 187, 398–403.
- Bishop, Y. M. M., Fienberg, S. E., and Holland, P. W. (1975). *Discrete Multivariate Analysis: Theory and Practice*. Cambridge, MA: MIT Press.
- Carlyle, T. (1840). *Chartism*. London: J. Fraser.
- Cicchetti, D. V. and Allison, T. (1971). A new procedure for assessing reliability of scoring EEG sleep recordings. *American Journal of EEG Technology*, 11, 101–109.
- Cohen, A. (1980). On the graphical display of the significant components in a two-way contingency table. *Communications in Statistics— Theory and Methods*, A9, 1025–1041.

- Cohen, J. (1960). A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20, 37–46.
- Cohen, J. (1968). Weighted kappa: Nominal scale agreement with provision for scaled disagreement or partial credit. *Psychological Bulletin*, 70, 213–220.
- Fienberg, S. E. (1975). Perspective Canada as a social report. *Social Indicators Research*, 2, 153–174.
- Fleiss, J. L. (1973). *Statistical Methods for Rates and Proportions*. New York: John Wiley and Sons.
- Fleiss, J. L. and Cohen, J. (1972). The equivalence of weighted kappa and the intraclass correlation coefficient as measures of reliability. *Educational and Psychological Measurement*, 33, 613–619.
- Fleiss, J. L., Cohen, J., and Everitt, B. S. (1969). Large sample standard errors of kappa and weighted kappa. *Psychological Bulletin*, 72, 332–327.
- Friendly, M. (1991). *SAS System for Statistical Graphics*. Cary, NC: SAS Institute, 1st edn.
- Friendly, M. (1994a). A fourfold display for 2 by 2 by K tables. Tech. Rep. 217, York University, Psychology Dept.
- Friendly, M. (1994b). SAS/IML graphics for fourfold displays. *Observations*, 3(4), 47–56.
- Gart, J. J. and Zweifel, J. R. (1967). On the bias of various estimators of the logit and its variance with applications to quantal bioassay. *Biometrika*, 54, 181–187.
- Goodman, L. A. (1979). Simple models for the analysis of association in cross-classifications having ordered categories. *Journal of the American Statistical Association*, 74, 537–552.
- Haberman, S. J. (1979). *The Analysis of Qualitative Data: New Developments*, vol. II. New York: Academic Press.
- Haldane, J. B. S. (1955). The estimation and significance of the logarithm of a ratio of frequencies. *Annals of Human Genetics*, 20, 309–311.
- Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, 6(2), 65–70.
- Hout, M., Duncan, O. D., and Sobel, M. E. (1987). Association and heterogeneity: Structural models of similarities and differences. *Sociological Methodology*, 17, 145–184.
- Kendall, M. G. and Stuart, A. (1961). *The Advanced Theory of Statistics*, vol. 2. London: Griffin.
- Koch, G. and Edwards, S. (1988). Clinical efficiency trials with categorical data. In K. E. Peace, ed., *Biopharmaceutical Statistics for Drug Development*, (pp. 403–451). New York: Marcel Dekker.
- Kundel, H. L. and Polansky, M. (2003). Measurement of observer agreement. *Radiology*, 228(2), 303–308.
- Landis, J. R. and Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, 33, 159–174.

- Landis, R. J., Heyman, E. R., and Koch, G. G. (1978). Average partial association in three-way contingency tables: A review and discussion of alternative tests. *International Statistical Review*, 46, 237–254.
- Mersey, L. (1912). Report on the loss of the “Titanic” (S. S.). Parliamentary command paper 6352.
- Riedwyl, H. and Schüpbach, M. (1983). Siebdiagramme: Graphische darstellung von kontingenztafeln. Tech. Rep. 12, Institute for Mathematical Statistics, University of Bern, Bern, Switzerland.
- Riedwyl, H. and Schüpbach, M. (1994). Parquet diagram to plot contingency tables. In F. Faulbaum, ed., *Softstat '93: Advances In Statistical Software*, (pp. 293–299). New York: Gustav Fischer.
- Shrout, P. E. and Fleiss, J. L. (1979). Intraclass correlations: Uses in assessing rater reliability. *Psychological Bulletin*, 86, 420–428.
- Simpson, E. H. (1951). The interpretation of interaction in contingency tables. *Journal of the Royal Statistical Society, Series B*, 30, 238–241.
- Snee, R. D. (1974). Graphical display of two-way contingency tables. *The American Statistician*, 28, 9–12.
- Srole, L., Langner, T. S., Michael, S. T., Kirkpatrick, P., Opler, M. K., and Rennie, T. A. C. (1978). *Mental Health in the Metropolis: The Midtown Manhattan Study*. New York: NYU Press.
- Stokes, M. E., Davis, C. S., and Koch, G. G. (2000). *Categorical Data Analysis Using the SAS System*. Cary, NC: SAS Institute, 2nd edn.
- Tukey, J. W. (1993). Graphic comparisons of several linked aspects: Alternative and suggested principles. *Journal of Computational and Graphical Statistics*, 2(1), 1–33.
- Upton, G. J. G. (1976). The diagrammatic representation of three-party contests. *Political Studies*, 24, 448–454.
- Upton, G. J. G. (1994). Picturing the 1992 British general election. *Journal of the Royal Statistical Society, Series A*, 157(Part 2), 231–252.
- Von Eye, A. and Mun, E. (2006). *Analyzing Rater Agreement: Manifest Variable Methods*. Taylor & Francis.
- Wainer, H. (1996). Using trilinear plots for NAEP state data. *Journal of Educational Measurement*, 33(1), 41–55.
- Wolf, B. (1995). On estimating the relation between blood group and disease. *Annals of Human Genetics*, 19, 251–253.