1.3.3. Analysis of a substantial dataset – US accident data

Each year the National Highway Traffic Safety Administration (NHTSA) in the USA collects, using a random sampling method, data from all police-reported crashes in which there is a harmful event (people or property), and from which at least one vehicle is towed. The data frame `nassCDS` (*DAAG*) is derived from NHTSA data for the years 1997 – 2002.[5]

The use of a complex sampling scheme has the consequence that the sampling fraction differs between observations. Each point has to be multiplied by the relevant sampling fraction, in order to get a proper estimate of its contribution to the total number of accidents. The column `weight` (`national` = *national inflation factor* in the SAS dataset) gives the relevant multiplier.

Meyer (2006) argues that on balance (over the period when their data were collected) airbags cost lives. In order to obtain a fair comparison, it is necessary to adjust, not only for the effects of seatbelt use, but also for speed of impact. When this is done, airbags appear on balance to be dangerous, with the most serious effects in high impact accidents, but the effect is at the level of statistical error.

Strictly, the conclusion is that, conditional on involvement in an accident that was sufficiently serious to be included in the database (at least one vehicle towed away from the scene), and conditioning also on `seatbelt` (seatbelt use or not) and `dvcat` (force of impact) there is a suggestion that airbags are harmful. Conditional on the airbag failing to prevent an outcome that is somewhat serious, there is a suggestion that airbags are harmful!

Farmer (2006) argued that these data have too many uncertainties and potential sources of bias to give reliable results when analyzed as will be done here. Additionally, there are other factors on which the effects of airbag use could and perhaps should be conditioned. Farmer presented a different analysis, based on the use of front seat passenger mortality as a standard against which to compare driver mortality, and limited to cars without passenger airbags. In the absence of any effect from airbags, the ratio of driver mortality to passenger mortality should be the same, irrespective of whether or not there was a driver airbag. In fact the ratio of driver fatalities to passenger fatalities was 11% lower in the cars with driver airbags.

From Highly to Mildly Misleading Analyses

The analyses presented here will be for a subset of the data that are further restricted. The oldest vehicles with airbags, represented in these data, were from 1986. In an analysis that does not allow for age of vehicle, this risks biasing results for vehicles without airbags towards results for older vehicles. If there is an adjustment for age of vehicle, vehicles that are older than 1986 do not contribute useful information, for purposes of assessing the effectiveness of airbags. In addition to omitting vehicles older than 1986, observations with weight 0, and one observation where the year of vehicle was unknown. This omits 2726 records out of the total of 26217, leaving 23491 records.

```
> library(DAAG)
> nassnew <- subset(nassCDS, !is.na(yearVeh) & yearVeh>=1986 & weight>0)
```

[5] They hold a subset of the columns from a corrected version of the data analyzed in the Meyer (2005) paper that is referenced on the help page for `nassCDS`. More complete data are available from one of the web pages
`http://www.stat.colostate.edu/~meyer/airbags.htm` (SAS transport file)
or `http://www.maths.anu.edu.au/~johnm/datasets/airbags/` (R image file).

Survival rates, according to airbag use: The following estimates numbers of front seat passengers alive and dead, classified by airbag use:

```
> library(DAAG)
> (abtab <- xtabs(weight ~ dead + airbag, data=nassnew))
```

```
        airbag
dead        none  airbag
  alive  4357430 6614169
  dead     29897   25919
```

Now use the function `prop.table()` can then be used to obtain the proportions in margin 1, i.e., the proportions dead, according to airbag use:

```
> round(prop.table(abtab, margin=2)["dead", ], 4)
```

```
  none airbag
0.0068 0.0039
```

```
> ## Alternatively, the following gives proportions alive & dead
> ## round(prop.table(abtab, margin=2), 4)
```

The above might suggest that the deployment of an airbag substantially reduces the risk of mortality.

```
> abSBtab <- xtabs(weight ~ dead + seatbelt + airbag, data=nassnew)
> ## Take proportions, retain margins 2 & 3, i.e. airbag & seatbelt
> round(prop.table(abSBtab, margin=2:3)["dead", , ], 4)
```

```
         airbag
seatbelt    none airbag
  none    0.0180 0.0155
  belted  0.0039 0.0021
```

The results are now much less favorable to airbags. To see why, consider:

```
> margin.table(abSBtab, margin=2:3)  # Add over margin 1
```

```
         airbag
seatbelt     none  airbag
  none     916169  885635
  belted  3471158 5754453
```

In the overall table, the results without airbags are mildly skewed (~4.12:1.37) to the results for `belted`, while with airbags they are highly skewed (~57.6:8.86) to the results for `belted`.

Taking Account of Estimated Force of Impact: Now take account, additionally, of estimated force of impact (`dvcat`):

```
> ASdvtab <- xtabs(weight ~ dead + seatbelt + airbag + dvcat,
                   data=nassnew)
> ## Use ftable to get a compact, flattened version of the table
> round(ftable(prop.table(ASdvtab, margin=2:4)["dead", , , ]), 6)
```

		dvcat	1-9km/h	10-24	25-39	40-54	55+
seatbelt	airbag						
none	none		0.000000	0.002583	0.020300	0.040323	0.204534
	airbag		0.004023	0.004873	0.010982	0.075990	0.269959
belted	none		0.000000	0.000380	0.005743	0.028141	0.139204
	airbag		0.000000	0.000195	0.003331	0.022666	0.157394

It will be apparent that differences between `none` and `airbag` are now below any reasonable threshold of statistical detectability.

More Variables Still

There are at least two other variables that may affect the risk of death. These are the year of manufacture of the vehicle, and the age of the occupant. Possibly also the year of the accident might be important, but the data do not have enough information to allow this effect to be modeled in addition to all the others. Subsection 1.3.3 uses a generalized linear model (GLM) to analyse these data.

1.3.4. Summary of continuous outcome data - an example

Unequal subgroup weights create exactly the same potential, as with binary (or categorical) outcome data, for misleading summary.

Unequal subgroup weights with continuous data – an example

Figure 1.3.4 relates to data collected in an experiment on the use of painkillers.[6]. Notice that the overall comparison (average for baclofen versus average for no baclofen) goes in a different direction from the comparison for the two sexes separately.

Researchers had been looking for a difference between the two analgesic treatments, without and with baclofen. When the paper was first submitted for publication, an alert reviewer spotted that some of the treatment groups contained more women than men, and proposed a re-analysis to determine whether this accounted for the results.[7] When the data were analysed to take account of the gender effect, it turned out that the main effect was a gender effect, with a much smaller difference between treatments.

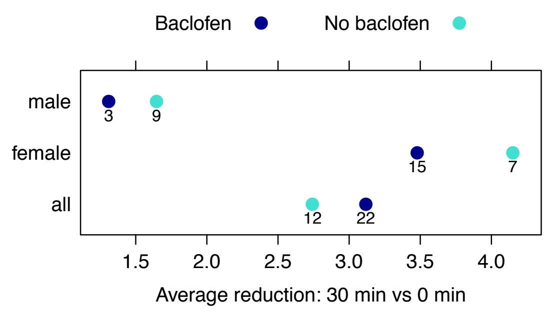

Figure 1.11.: Does baclofen, following operation (additional to earlier painkiller), reduce pain? Subgroup numbers, shown below each point in the graph, weight the overall averages when sex is ignored.

[6]Gordon, N. C. et al.(1995): "Enhancement of Morphine Analgesia by the GABAB against Baclofen". Neuroscience 69: 345-349

[7]Cohen, P. 1996. Pain discriminates between the sexes. New Scientist, 2 November, p. 16.