Universität Zürich^UZH

Epidemiology, Biostatistics
and Prevention Institute,
Epidemiology Department
www.ebpi.uzh.ch

Interpreting Interactions in R

Sarah R Haile, PD PhD (sarah.haile@uzh.ch)

Version 1.1 of May 6, 2025

Contents

1 Introduction

Regression models are often used to explore associations between different variances, sometimes including interactions. Unfortunately, interactions are sometimes hard to interpret. Here we explain the interpretation of three different kinds of interactions

1. nominal (sometimes called categorical, or binary if there are only two categories) by nominal;
2. nominal by continuous; and
3. continuous by continuous.

Code examples in R are provided using the `birthwt` dataset `?MASS::birthwt`).

```
library(tidyverse) # for data cleaning
library(broom) # for coefficient tables
data(birthwt)
birthwt <- MASS::birthwt %>%
  mutate(smoke = factor(smoke, 0:1, c("non-smoker", "smoker")),
         race = factor(race, 1:3, c("white", "black", "other")),
         nonwhite = as.numeric(race != "white"),
         nonwhite = factor(nonwhite, 0:1, c("W", "NW")))

head(birthwt)
```

```
##    low age lwt  race      smoke ptl ht ui ftv  bwt nonwhite
## 85   0  19 182 black non-smoker   0  0  1   0 2523       NW
## 86   0  33 155 other non-smoker   0  0  0   3 2551       NW
## 87   0  20 105 white     smoker   0  0  0   1 2557        W
## 88   0  21 108 white     smoker   0  0  1   2 2594        W
## 89   0  18 107 white     smoker   0  0  1   0 2600        W
## 91   0  21 124 other non-smoker   0  0  0   0 2622       NW
```

2 Linear regression

2.1 Nominal by nominal

Without interaction With only main effects, we assume that the mean difference between categories of one variable is the same, regardless of the value of the 2nd variable, and vice versa.

With interaction Including an interaction term, we assume that the mean difference between categories of one variable differs according to the 2nd variable, and vice versa.

Interpretation of Interaction Coefficient The interaction term gives additional difference in means for non-reference levels of the two categorical variables.

Interpretation The reference category for `smoke` is non-smoking mothers, and for `nonwhite` is white mothers. Babies of smokers have on average -602g lower birthweights than non-smokers. Babies of non-white mothers have -604g lower birthweights than those of whites. However, the association with birthweight is not as strong as expected in non-white smokers, as they have on average 419g higher birthweights than would be expected considering the main effects only.

Interpretation for each group

Non-smoking, white mothers This is the reference group, with an average birthweight given by the intercept: 3429g.

Smoking, white mothers White mothers who smoke have babies with on average -602g lower birthweights than white mothers who do not smoke.

Non-smoking, non-white mothers Non-white mothers who do not smoke have babies with on average -604g lower birthweights than white mothers who do not smoke.

Smoking, non-white mothers Non-white mothers who do smoke have babies with on average $-602 + -604 + 419 = -787$g lower birthweights than white mothers who do not smoke.

```
m1 <- lm(bwt ~ smoke * nonwhite, data = birthwt)
tidy(m1, conf.int = TRUE) %>%
  mutate(p.value = format.pval(p.value, eps = 0.0001, sci = FALSE, digits = 2))
```

```
##                     term estimate std.error statistic p.value conf.low conf.high
## 1            (Intercept)     3429       103      33.4 <0.0001   3226.1      3631
## 2            smokesmoker     -602       140      -4.3 <0.0001   -877.3      -327
## 3             nonwhiteNW     -604       131      -4.6 <0.0001   -862.2      -346
## 4 smokesmoker:nonwhiteNW      419       217       1.9   0.055     -8.8       848
```

2.2 Nominal by continuous

Without interaction With only main effects, we assume that the slope of y over the continuous variable, x is the same regardless of the category of the nominal variable, $z = 0$ or $z = 1$. That is, the regression lines for each group in z are parallel.

With interaction Including an interaction term, we assume that the slope of y over x differs according to $z = 0$ or $z = 1$. The regression lines for each group in z no longer are assumed to be parallel.

Interpretation of Interaction Coefficient The interaction term gives additional change in slope of y over x for the non-reference level of the nominal variable, $z = 1$. The slopes are given by:

$z = 0$: $\hat{\beta}_x$

$z = 1$: $\hat{\beta}_x + \hat{\beta}_{x:z}$

Interpretation For non-smokers, average birthweight increases by 28g per year of age of the mother. For smokers, the average birthweight actually decreases by -19g ($28 + -47$) per year of age of the mother.

```
m2 <- lm(bwt ~ smoke * age, data = birthwt)
tidy(m2, conf.int = TRUE) %>%
  mutate(p.value = format.pval(p.value, eps = 0.0001, sci = FALSE, digits = 2))

##               term estimate std.error statistic p.value conf.low conf.high
## 1      (Intercept)     2406       292       8.2 <0.0001   1829.6    2982.5
## 2      smokesmoker      798       484       1.6   0.101   -157.4    1753.7
## 3              age       28        12       2.3   0.024      3.8      51.7
## 4 smokesmoker:age      -47        20      -2.3   0.024    -86.9      -6.2
```

Tip Note that the main effect of smoking here gives the mean difference between smokers and non-smokers *for age = 0*. It may be easier to interpret models with nominal by continuous interactions if you first center the continuous variable (at mean, median or other relevant value).

```
summary(birthwt$age)

##    Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
##   14.00   19.00   23.00   23.24   26.00   45.00

birthwt <- birthwt %>%
  mutate(agec = age - 23)
m2c <- lm(bwt ~ smoke * agec, data = birthwt)
tidy(m2c, conf.int = TRUE) %>%
  mutate(p.value = format.pval(p.value, eps = 0.0001, sci = FALSE, digits = 2))

##                term estimate std.error statistic p.value conf.low conf.high
## 1       (Intercept)     3044        66      45.9 <0.0001   2913.0    3174.8
## 2       smokesmoker     -273       106      -2.6   0.011   -481.8     -64.2
## 3              agec       28        12       2.3   0.024      3.8      51.7
## 4 smokesmoker:agec      -47        20      -2.3   0.024    -86.9      -6.2
```

2.3 Continuous by continuous

Without interaction With only main effects, we assume that the slope of y over the continuous variable $x1$ is the same regardless of $x2$ and vice versa.

With interaction Including an interaction term, we assume that the slope of y over the continuous variable $x1$ differs with respect to $x2$, and vice versa.

Interpretation of Interaction Coefficient The interaction term gives the change in slope of y over $x1$ for each unit of $x2$, and the change in slope of y over $x2$ for each unit of $x1$. The actual slopes are given by:

slope over $x1$: $\hat{\beta}_{x1} + x_2\hat{\beta}_{x1:x2}$

slope over $x2$: $\hat{\beta}_{x2} + x_1\hat{\beta}_{x1:x2}$

Interpretation Average birthweight increases by on average 11.7g for every year of the mother's age, and 4.4g for each pound of the mother's weight. Increasing age and weight of the mother make these associations slight less pronounced (-0.3g per year of age and pound).

Tip Unless $x1 = 0$ and $x2 = 0$ are meaningful in your dataset, you may end up with strange values for the intercept or other main effect estimates. If this happens, try centering continuous variables. Don't forget that this will change the calculation of the predicted values:

$$\hat{y} = \hat{\beta}_{(Intercept)} + \hat{\beta}_{agec}(age - 23) + \hat{\beta}_{lwtc}(lwt - 121) + \hat{\beta}_{agec:lwtc}(age - 23)(lwt - 121)$$

```
summary(birthwt$lwt)

##    Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
##    80.0   110.0   121.0   129.8   140.0   250.0

birthwt <- birthwt %>%
  mutate(lwtc = lwt - 121)
m3 <- lm(bwt ~ agec * lwtc, data = birthwt)
tidy(m3, conf.int = TRUE) %>%
  mutate(p.value = format.pval(p.value, eps = 0.0001, sci = FALSE, digits = 2))

##          term estimate std.error statistic p.value conf.low conf.high
## 1 (Intercept)   2912.1     54.89     53.05 <0.0001  2803.82   3020.40
## 2        agec     11.7     10.81      1.09   0.279    -9.59     33.06
## 3        lwtc      4.4      1.76      2.51   0.013     0.94      7.90
## 4   agec:lwtc     -0.3      0.32     -0.93   0.355    -0.94      0.34
```

Tip Graph the predicted values in order to make sense of continuous by continuous interactions.

```
nd <- crossing(age = seq(15, 35, 5),
               lwt = seq(75, 200, 25)) %>%
  mutate(agec = age - 23,
         lwtc = lwt - 121)
nd$pred <- predict(m3, newdata = nd)
qplot(age, pred, data = nd, color = factor(lwt), geom = "line") + ylim(0, 5000)
qplot(lwt, pred, data = nd, color = factor(age), geom = "line") + ylim(0, 5000)
```

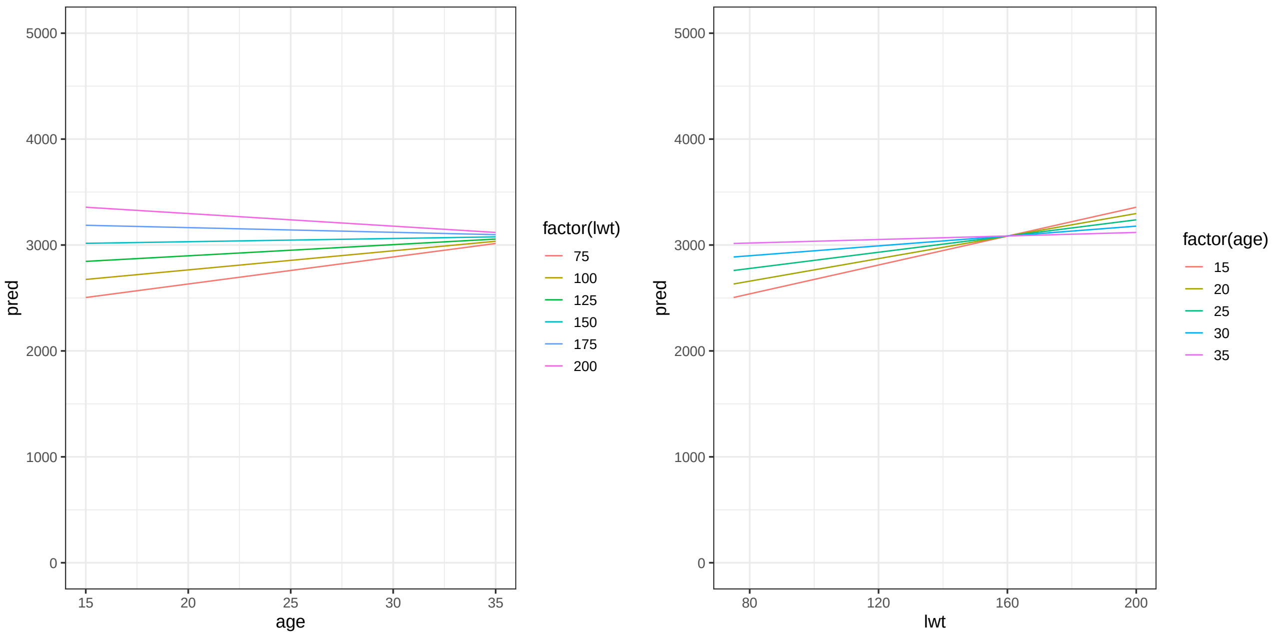

3 Logistic regression

The interpretations given in this section apply equally to

- logistic regression for binary outcomes ($e^{\hat{\beta}}$ = odds ratio (OR)),
- poisson regression for count outcomes ($e^{\hat{\beta}}$ = incidence rate ratio (IRR)),
- Cox proportional hazards regression for survival outcomes ($e^{\hat{\beta}}$ = hazard ratio (HR)),
- and other regression models where relevant coefficients are intepreted as $e^{\hat{\beta}}$, not $\hat{\beta}$.

3.1 Nominal by nominal

Without interaction With only main effects, we assume that the odds ratio comparing categories of one variable is the same, regardless of the value of the 2nd variable, and vice versa.

With interaction Including an interaction term, we assume that the odds ratio comparing categories of one variable differs according to the 2nd variable, and vice versa. An OR < 1 for the interaction, indicates the association is less strong than expected when considering only the main effects, while OR > 1 indicates the association is stronger than expected.

Interpretation of Interaction Coefficient The interaction term gives multiplicative effect of non-reference levels of the two categorical variables.

For nominal by nominal interactions, we examine the effects of two covariates simultaneously by multiplying the odds ratios. To see the effect of covariates $x1$ and $x2$, we multiply $e^{\hat{\beta}_{x1}}$ with $e^{\hat{\beta}_{x2}}$ to get $e^{\hat{\beta}_{x1}}e^{\hat{\beta}_{x2}} = e^{\hat{\beta}_{x1}+\hat{\beta}_{x2}}$. (Note that we can either a) first add the coefficients and then exponentiate, or b) first exponentiate to get odds ratios, and then multiply.) With interaction, we calculate the odds ratio as follows:

$$OR_{x1,x2} = e^{\hat{\beta}_{x1}}e^{\hat{\beta}_{x2}}e^{\hat{\beta}_{x1:x2}}.$$

```
m4 <- glm(low ~ smoke * nonwhite, data = birthwt, family = binomial)
tidy(m4, conf.int = TRUE, exponentiate = TRUE) %>%
  mutate(p.value = format.pval(p.value, eps = 0.0001, sci = FALSE, digits = 2))

##                      term estimate std.error statistic p.value conf.low conf.high
## 1             (Intercept)     0.10      0.52      -4.4 <0.0001    0.030      0.25
## 2             smokesmoker     5.76      0.60       2.9  0.0034    1.939     21.37
## 3              nonwhiteNW     5.43      0.58       2.9  0.0035    1.911     19.63
## 4 smokesmoker:nonwhiteNW     0.32      0.78      -1.5  0.1413    0.064      1.39
```

Interpretation In this example, we look at the odds of having birthweight less than 2.5kg. Smokers have 5.76 higher odds of having a baby with low birthweight compared to non-smokers. Similarly, nonwhite mothers have a 5.43 higher odds of having a baby with low birthweight compared to white mothers. Nonwhite mothers who smoke however have a 10 times higher odds of having a baby with low birthweight than white mothers who do not smoke.

```
exp(coef(m4)["smokesmoker"]) * exp(coef(m4)["nonwhiteNW"]) *
      exp(coef(m4)["smokesmoker:nonwhiteNW"])

## smokesmoker
##    9.999999
```

3.2 Nominal by continuous

Without interaction With only main effects, we assume that the odds ratio increases the same amount per unit of the continuous variable, x, is the same regardless of the category of the nominal variable, $z = 0$ or $z = 1$.

With interaction Including an interaction term, we assume that the change in odds ratio over the continuous variable differs according the value of z

Interpretation of Interaction Coefficient The interaction term gives additional change in odds for the non-reference level of the nominal variable, $z = 1$. The ORs are given by:

$z = 0$: e^{β_x}

$z = 1$: $e^{\beta_x} e^{\beta_{x:z}}$

Interpretation In this example, the odds of having a baby with low birthweight decreases by a factor of 0.92 per every year of the mother's age if the mother doesn't smoke, and by a factor of $0.92 * 1.08 = 0.99$ for every year if she does smoke.

```
m5 <- glm(low ~ smoke * agec, data = birthwt, family = binomial)
tidy(m5, conf.int = TRUE, exponentiate = TRUE) %>%
  mutate(p.value = format.pval(p.value, eps = 0.0001, sci = FALSE, digits = 2))

##                term estimate std.error statistic p.value conf.low conf.high
## 1       (Intercept)     0.33     0.221      -5.0 <0.0001     0.21       0.5
## 2       smokesmoker     2.05     0.324       2.2   0.027     1.09       3.9
## 3              agec     0.92     0.045      -1.8   0.065     0.84       1.0
## 4 smokesmoker:agec     1.08     0.065       1.1   0.263     0.95       1.2
```

3.3 Continuous by continuous

Without interaction With only main effects, we assume that the change in OR over the continuous variable $x1$ is the same regardless of $x2$ and vice versa.

With interaction Including an interaction term, we assume that the change in OR over the continuous variable $x1$ differs with respect to $x2$, and vice versa.

Interpretation of Interaction Coefficient The interaction term gives the change in OR over $x1$ for each unit of $x2$, and the change in slope of y over $x2$ for each unit of $x1$. The actual slopes are given by:

slope over $x1$: $e^{\beta_{x1}+x_2\beta_{x1:x2}}$

slope over $x2$: $e^{\beta_{x2}+x_1\beta_{x1:x2}}$

Interpretation The odds ratios considering an interaction between age and weight are *very* slightly lower (99.9% of the odds ratio considering only main effects [99.7 - 1.002%] per year of age and pound in weight), but this difference is not statistically significant.

Tip Plotting predicted odds ratios or probabilities for these models will make the models easier to understand.

```
m6 <- glm(low ~ agec * lwtc, data = birthwt, family = binomial)
tidy(m6, conf.int = TRUE, exponentiate = TRUE) %>%
  mutate(p.value = format.pval(p.value, eps = 0.0001, sci = FALSE, digits = 2))

##          term estimate std.error statistic p.value conf.low conf.high
## 1 (Intercept)     0.49    0.1638    -4.345 <0.0001     0.35      0.67
## 2        agec     0.96    0.0332    -1.195    0.23     0.90      1.02
## 3        lwtc     0.99    0.0062    -2.055    0.04     0.97      1.00
## 4   agec:lwtc     1.00    0.0012    -0.015    0.99     1.00      1.00
```

```
nd <- crossing(age = seq(15, 35, 5),
               lwt = seq(75, 200, 25)) %>%
  mutate(agec = age - 23,
         lwtc = lwt - 121)
nd$pred <- predict(m6, newdata = nd, type = "response")
qplot(age, pred, data = nd, color = factor(lwt), geom = "line") +
  ylim(0, 1) + ylab("Pr(Birthweight < 2500g)")
qplot(lwt, pred, data = nd, color = factor(age), geom = "line") +
  ylim(0, 1) + ylab("Pr(Birthweight < 2500g)")
```

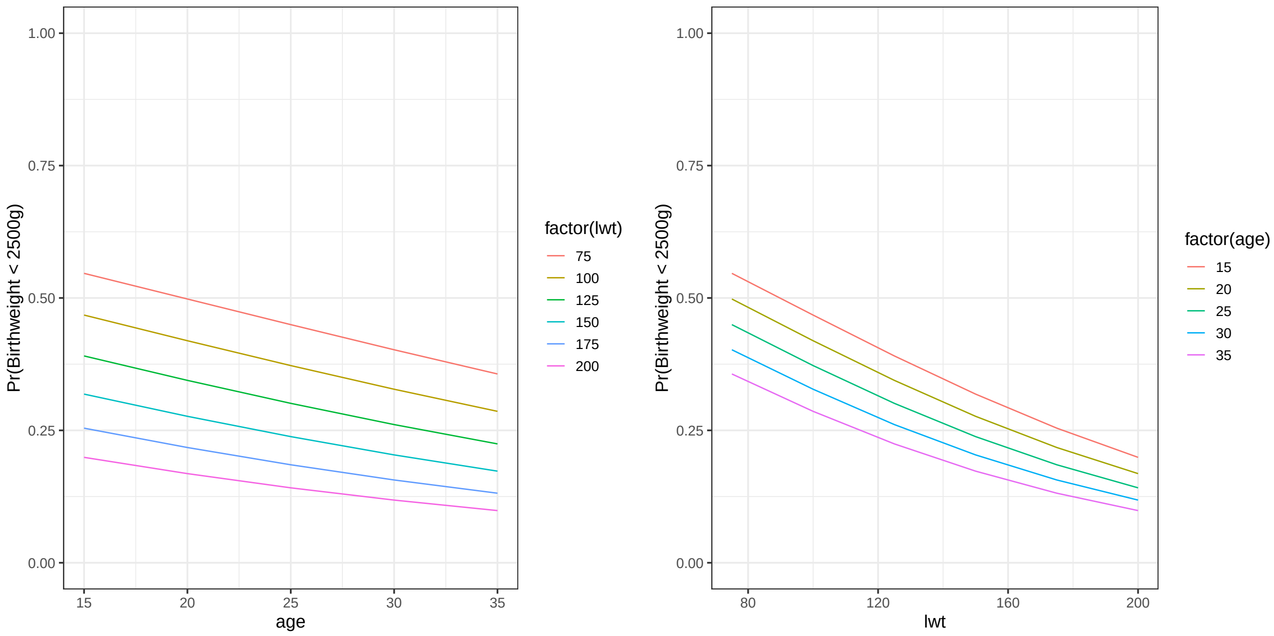

4 Recent changes

1.0 Original version

1.1 Updated UZH logo, updated R code to include `tidyverse` (e.g. `dplyr`) and `broom` packages, removed Stata examples.

Software used to generate this report

- R version: R version 4.5.0 (2025-04-11)
- Base packages: stats, graphics, grDevices, utils, datasets, methods, base
- Other packages: broom, lubridate, forcats, stringr, dplyr, purrr, readr, tidyr, tibble, tidyverse, ggplot2, knitr

This document was generated on 2025-05-06 at 14:25.