

Replication Code for Figures in “A Practical Guide to Dealing with Attrition in Political Science Experiments”

This version: January 2023

```
rm(list=ls())
library(viridis)
require(tidyverse)
require(waffle)
require(grid)
require(ggpattern)
require(ggpubr)
require(devtools)

#install our attritevis package
install_github("lbassan/attritevis", dependencies = TRUE)
library(attritevis)
set.seed(12)

#12
```

Code to replicate figures from the paper A Practical Guide to Dealing with Attrition in Political Science Experiments by Lo, Renshon, and Bassan-Nygate (2023).

1 Figure 1:

Data frame `lit_review.csv` summarizes experimental paper published in JEPS and their discussion of attrition. Below we create the waffle figure.

```
#Reading in CSV Data
attrition <- read_csv("lit_review.csv")
```

```

#Functions to remove "*" and change "Yes" to 1 and "No" to 0
remove_star <- function(x) {
  return(str_extract(x, "Yes|No"))
}

yesno_onezero <- function(x) {
  return(case_when(x == "Yes" ~ 1,
                  x == "No" ~ 0))
}

attrition <- attrition %>%
  mutate_at(c(7:14), remove_star) %>%
  mutate_at(c(7:14), yesno_onezero)

#Creating table of proportions
prop_att <- mean(attrition$Attrition)
prop_noatt <- mean(attrition$`0 Attrition`[attrition$Attrition == 1])
prop_attdv <- mean(attrition$`Response Rate DV`[attrition$Attrition == 1])
prop_quan <- mean(attrition$`Quantified Attrition`[attrition$Attrition == 1 & attrition$`0 Attrition` == 0 & attrition$`Response Rate DV` == 0])
prop_adj <- mean(attrition$`Sample Adjustments`[attrition$Attrition == 1 & attrition$`0 Attrition` == 0 & attrition$`Response Rate DV` == 0])

attrition_summary <- as_tibble(data.frame(
  c("Measurement",
    "Proportion that mention attrition",
    "Proportion \"no attrition\"",
    "Proportion DV",
    "Proportion quantify",
    "Proportion adjust"),
  c("Value",
    prop_att,
    prop_noatt,
    prop_attdv,
    prop_quan,
    prop_adj)

```

```

))

#Creating variable for the waffle plot
count <- attrition %>%
  mutate(waffle = case_when(`Sample Adjustments` == 1 ~ "Attrition mentioned, quantified,
analyzed",
                             Attrition == 1 & `0 Attrition` == 0 & `Response Rate DV` == 0
& `Sample Adjustments` == 0 & `Quantified Attrition` == 1 ~ "Attrition mentioned and quant
ified",
                             Attrition == 1 & `0 Attrition` == 0 & `Response Rate DV` == 0
& `Sample Adjustments` == 0 & `Quantified Attrition` == 0 ~ "Attrition mentioned only",
                             `Response Rate DV` == 1 ~ "Attrition is DV",
                             `0 Attrition` == 1 ~ "Attrition mentioned - none in study",
                             Attrition == 0 ~ "No mention of attrition")) %>%
  group_by(waffle) %>%
  summarise(n = n())

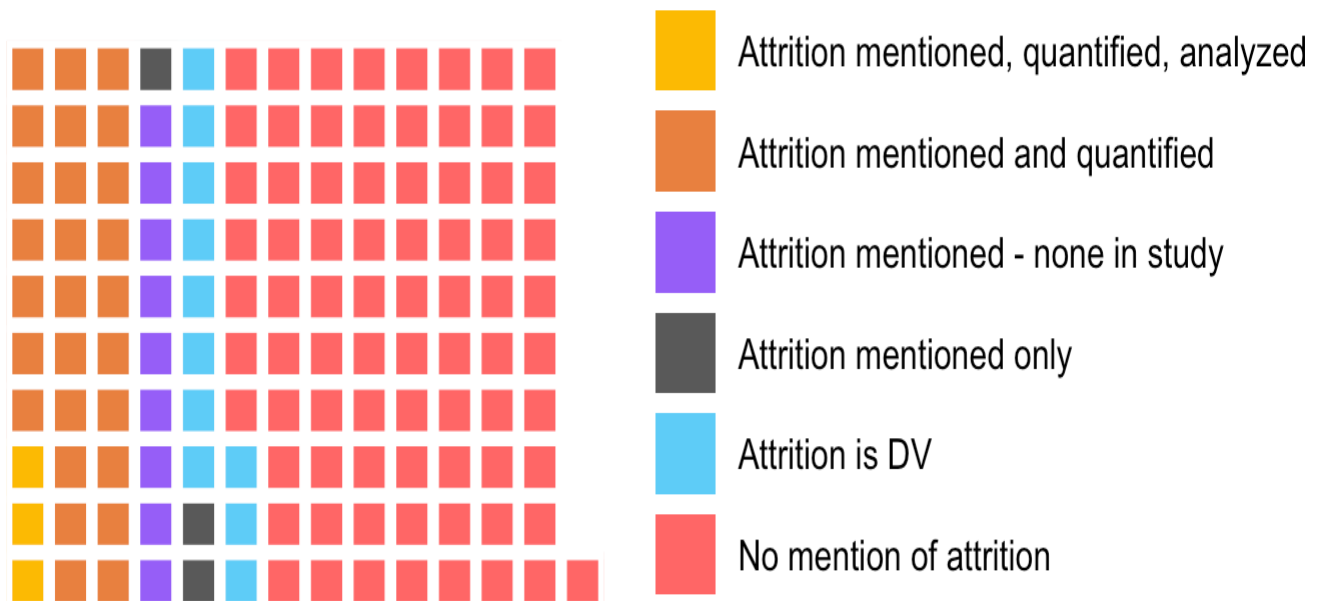
#Reordering to make legend easier to read and plot look better
count <- count[c(5,3,2,4,1,6),]

#Creating waffle plot
case_counts <- count$n
names(case_counts) <- count$waffle

plot1<-waffle(case_counts, colors = c(
  "#fcba03", #For Attrition mentioned, quantified, analyzed
  "#e8803f", #For Attrition mentioned and quantified
  "#965ef7", #For Attrition mentioned, none in study
  "#595959", #For Attrition mentioned only
  "#5eccf7", #For Attrition is DV
  "#ff6666" #For No mention of attrition
)) +
  theme(legend.key.size = unit(10, "mm"), legend.text = element_text(size = 12))

plot1

```



```
ggsave("Figures/plot1.png",
       width = 7, height = 5,
       plot1)
```

2 Figure 3:

Attrition timeline visualization, we rely on the `attritevis` package, presented in the paper.

```
#Make toy plots for paper
```

```

require(ggpattern)

#(a) Low Level Attrition
#Attrition post-treatment (throughout survey)
n <- 1000
df <- data.frame(
  Q1 = sample(c("Treatment", "Control"), n, rep = TRUE), #we will assume conditions are assigned when entering survey
  Q2 = sample(c(18:90), n, rep = TRUE), #age
  Q3 = sample(c("m", "f"), n, rep = TRUE, prob = c(0.55, 0.45)), #sex
  Q4 = sample(c(0,1), n, rep = TRUE)#other general pre-treatment questions
  df$Q5 = df$Q1 #at Q5 respondents are presented with treatment (say, vignette)
  df$Q6 = sample(c(0,1), n, rep = TRUE) #post treatment questions
  df$Q7 = sample(c(0,1), n, rep = TRUE)
  df$Q8 = sample(c(0,1), n, rep = TRUE)
  df$Q9 = sample(c(0,1), n, rep = TRUE)
  df$Q10 = sample(c(0,1), n, rep = TRUE)

  df_a<-df

#Generate attrition post
invisible(
  sapply(sample(1:nrow(df_a), 200),function(x) {
    a <- sample(2:10,1)
    df_a[x,a:ncol(df_a)] <<- NA
  })

#generate plot (a)
a<- attritevis::plot_attrition(data=df_a,
  freq = FALSE,
  treatment_q = "Q1",
  outcome_q = c("Q6", "Q7"),
  title = "(a) Low Level Attrition",
  mycolors = c(Control="#000066",
    Treatment = "#CC0033"),

```

```

        total = FALSE,
        tline = FALSE)

#note that treatment was administered in Q1 but `given` in Q5, so we manually input this

a<-a + geom_vline(xintercept = 5,
                  color = "black",
                  size = 1)+

  annotate(geom = "text",
          label = "Treatment Given",
          x = 5,
          y = 0.5,
          color = "black",
          angle = 90,
          vjust = 1.5)

a

#(b) Pre-treatment Attrition
df_b<-df
#Generate attrition pre-treatment
invisible(
  sapply(sample(1:nrow(df_b), 700),function(x) {
    a <- sample(2:4,1)
    df_b[x,a:ncol(df_b)] <<- NA
  }
))

#generate plot (b)
b<-attritevis::plot_attrition(data=df_b,
                              freq = FALSE,
                              treatment_q = "Q1",
                              outcome_q = c("Q6", "Q7"),
                              title = "(b) Pre-treatment Attrition",
                              mycolors = c(Control="#000066",
                                             Treatment = "#CC0033"),

```

```

        tline = FALSE,
        total = FALSE
    )

#note that treatment was administered in Q1 but `given` in Q5, so we manually input this

b<-b + geom_vline(xintercept = 5,
                 color = "black",
                 size = 1) +

    annotate(geom = "text",
           label = "Treatment Given",
           x = 5,
           y = 0.5,
           color = "black",
           angle = 90,
           vjust = 1.5)

b

#(c) Post-treatment Attrition (immediate)
df_c<-df

#First, we generate some general attrition at treatment
invisible(
  sapply(sample(1:nrow(df_c), 410, 0.8*nrow(df_c)),function(x) {
    a <- sample(5:6,1)
    df_c[x,a:ncol(df_c)] <<- NA
  }
))

#second, we add some attrition that's correlated with the treatment
#specifically, we want to demonstrate attrition that happens at a certain time
#to do so, we add a running var that will demonstrate time
df_c$no<-rownames(df_c)
df_c$Q6<-ifelse(df_c$Q1=="Treatment"&(df_c$no>115&df_c$no<373), NA,df_c$Q6)
df_c$Q7<-ifelse(is.na(df_c$Q6),NA,df_c$Q7)
df_c$Q8<-ifelse(is.na(df_c$Q6),NA,df_c$Q8)

```

```

df_c$Q9<-ifelse(is.na(df_c$Q6),NA,df_c$Q9)
df_c$Q10<-ifelse(is.na(df_c$Q6),NA,df_c$Q10)

df_c$no<-NULL

c<-attritevis::plot_attrition(data=df_c,
                              freq = FALSE,
                              treatment_q = "Q1",
                              outcome_q = c("Q6", "Q7"),
                              title = "(c) Post-treatment Attrition (immediate)",
                              mycolors = c(Control="#000066",
                                             Treatment = "#CC0033"),
                              tline = FALSE,
                              total = FALSE)

#note that treatment was administered in Q1 but `given` in Q5, so we manually input this

c<-c + geom_vline(xintercept = 5,
                  color = "black",
                  size = 1) +

  annotate(geom = "text",
          label = "Treatment Given",
          x = 5,
          y = 0.5,
          color = "black",
          angle = 90,
          vjust = 1.5)

c

#(d) Post-treatment Attrition (prolonged)
df_d<-df
#Generate attrition at DV + after
invisible(
  sapply(sample(1:nrow(df_d), 700),function(x) {
    a <- sample(6:10,1)
    df_d[x,a:ncol(df_d)] <<- NA
  })
)

```



```

}
))

d<-attritevis::plot_attrition(data=df_d,
  freq = FALSE,
  treatment_q = "Q1",
  outcome_q = c("Q6", "Q7"),
  title = "(d) Post-treatment Attrition (prolonged)",
  mycolors = c(Control="#000066",
    Treatment = "#CC0033"),
  tline = FALSE,
  total = FALSE
)

#note that treatment was administered in Q1 but `given` in Q5, so we manually input this

d<-d + geom_vline(xintercept = 5,
  color = "black",
  size = 1) +

  annotate(geom = "text",
    label = "Treatment Given",
    x = 5,
    y = 0.5,
    color = "black",
    angle = 90,
    vjust = 1.5)

d

require(grid)
#save all plots in one figure

plot2 <- ggarrange(a + rremove("ylab") + rremove("xlab"), b + rremove("ylab")
  + rremove("xlab"), c + rremove("ylab") + rremove("xlab"),
  d + rremove("ylab") + rremove("xlab"), # remove axis labels from plots

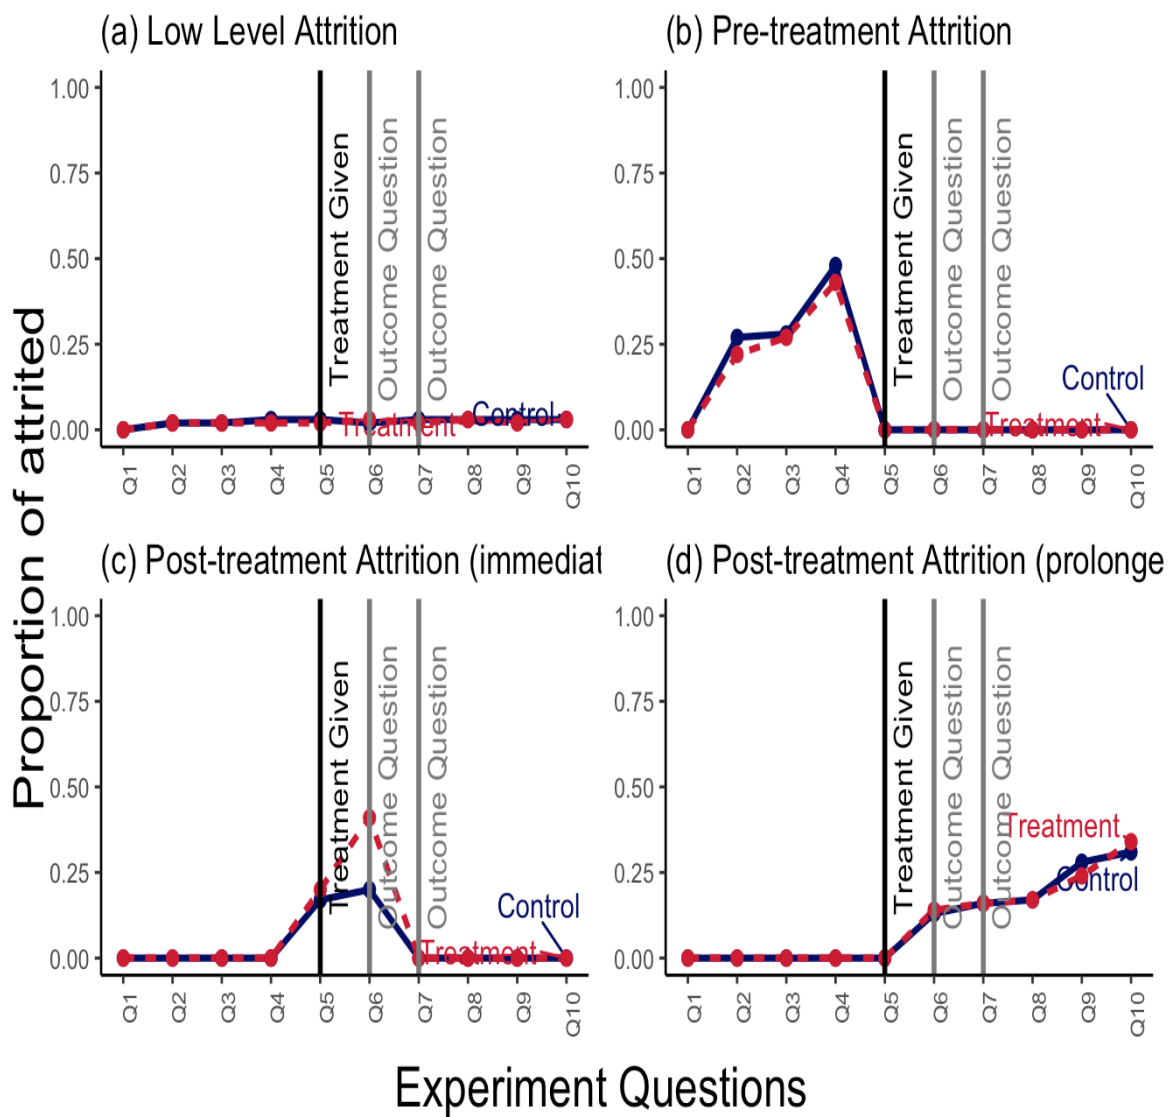
```

```

labels = NULL,
ncol = 2, nrow = 2,
common.legend = TRUE, legend = "top",
align = "hv",
font.label = list(size = 10, color = "black", face = "bold",
                  family = NULL, position = "top"))

annotate_figure(plot2, left = textGrob("Proportion of attrited", rot = 90, vjust = 1, gp
= gpar(cex = 1.5)),
              bottom = textGrob("Experiment Questions", gp = gpar(cex = 1.5)))

```



```

ggsave("Figures/plot2.png",
       width = 10,

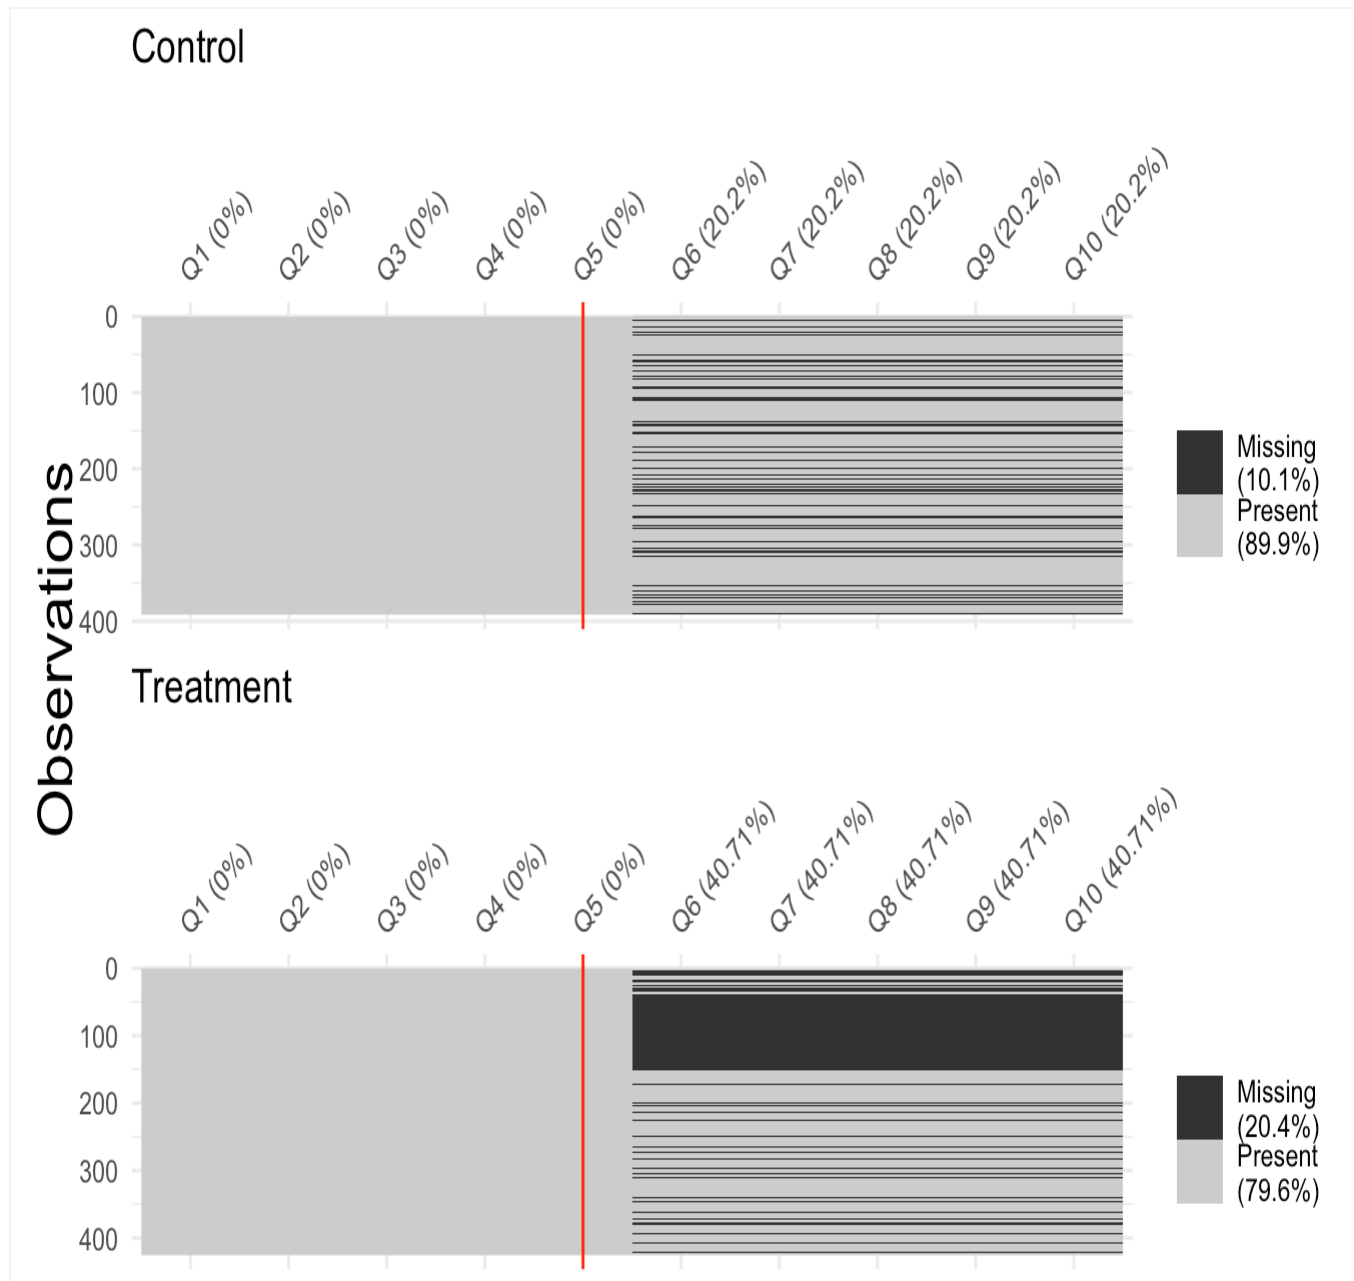
```

```
height = 7,  
plot2)
```

3 Figure 4:

Vis miss plot.

```
plot3<- attritevis::vis_miss_treat(data=df_c,  
treatment = "Q5")
```



```
ggsave("Figures/plot3.png", plot3)
```

`attritevis` Package: An R Vignette

This version: January 2023

1 Introduction

Attrition, the loss of study units from a sample, can often occur throughout an experimental study and at times pose a threat to inference. There are several studies, and accompanying R packages, that provide ex-post solutions to missingness such as double-sampling or extreme bounds. We provide a visually based guidance to assessing the types of missingness a study may have with a particular eye towards experimental and design adjustments a researcher can make after piloting a study.

2 Usage

- Visualizing survey attrition across treatment condition and over-time.
- Utilizing and comparing balance tests at precise moments in the survey.
- Incorporating estimation and visualization of Manski bounds for studies suffering from problematic attrition.

3 Assumptions

- Data must be *ordered by survey questions*, i.e. if respondents answered Q1 before Q2, the variable Q1 must appear before Q2 (i.e. in an earlier column) in the dataframe.
- When attrition is defined as completely leaving the survey. Hence, when attrition is reported in the package it does not count skippers, i.e. respondents who skipped a question(s) but continued later in the survey, as attrited. In the function `plot_attrition` users can count skippers by setting `y = "responded"`.
- For balance tests, treatment and control conditions must be defined.

4 Functions

4.1 `attrition()`

4.1.1 Description

- Converts survey data into a frame that includes:
 - `attrited` – how many respondents attrited (left the survey) at each question.
 - `proportion` – number of attrited respondents / number of respondents who entered survey.
 - `prop_q` – number of attrited respondents / number of respondents entering into the question. - `questions` – question names.
 - `responded` – how many respondents responded in each question.
 - `prop_r` – number of respondents who responded / number of respondents who entered survey.

4.1.2 Arguments

- `data` - a `data.frame` where variables are ordered by survey questions, such that earlier survey questions appear in smaller valued columns.

4.2 `attrition_table()`

4.2.1 Description

- Yields same `data.frame` as function `attrition`, but converts it into a table. Allows to subset table by treatment and control groups, which yields several tables by condition.

4.2.2 Arguments

- `data` - a `data.frame` where variables are ordered by survey questions, such that earlier survey questions appear in smaller valued columns.
- `treatment_q` - a `string` character that corresponds to treatment variable. When specified, the function yields several tables by condition.

4.3 `plot_attrition()`

4.3.1 Description

- Plots attrition or response in survey data over time.

4.3.2 Arguments

*`data` must be `data.frame`. Note that this function works only if the order of variables = order of questions in the survey. Users must remove irrelevant observations, for instance individuals who did not meet quotas, so as not to confuse them with attrited respondents. Note that using the `qualtrics`

embedded data feature users can note which respondents failed an attention check or were removed since they did not meet a quota.

*`y` is a character that corresponds to the Y axis. When `y = attrited`, attrition is plotted. When `y = responded` responses are plotted. Default is `y = attrited`.

*`freq` is a logical argument that notes whether Y axis is a raw number or a proportion. Default is `freq=TRUE`, which is the frequency of attrited OR responded respondents. When `freq=FALSE` Y axis is the proportion of total N (attrited OR responded), calculated as number of attrited OR responded divided by the number of respondents entering into the question.

*`treatment_q` is a character of name(s) of question(s) in which treatments were administered. Marked in the plot with a red vertical line.

*`outcome_q` is a character of name(s) of outcome question(s). Marked in the plot with a blue vertical line.

*`mycolors` is a character of color names to be used as values in `scale_colour_manual` argument in ggplot. Default is `mycolors=NULL`, which defaults to greyscale. `mycolors` must be == length of the unique values of the `treatment_q` variable. To use this argument, users should specify which color corresponds to which factor (for example, `treatment = "red"`).

*`title` is a character to be used for plot title.

*`total` is a logical argument that notes whether the total # of attrited/responded is plotted. Default is `TRUE`. Argument can be changed to `FALSE` only when `treatment_q` is full.

*`tline` is a logical argument that allows users to remove treatment vline, default is `tline=TRUE`.

*`outcomeline` is a logical argument that allows users to remove outcome vlins, default is `outcomeline=TRUE`.

4.4 `balance_cov()`

4.4.1 Description

- Tests whether specified covariates are balanced across specified treatment and control groups. Output is a t-test if covariate is a numeric or integer, and a 2-sample proportion test if covariate is a factor.

4.4.2 Arguments

- `data` - a `data.frame`, from which `treatment` and `question` are taken.
- `treatment` - a `string` character that corresponds to the name of the treatment variable. Note that values of said variable must be specified as `treatment` and `control`.
- `question` - a `string` character that corresponds to the name of the point in survey (question), for which balance test is required.
- `factor` - logical argument that specifies whether `question` is a factor. Default is `factor = FALSE` (i.e. question is a numeric or integer).
- `factor_name` - character that corresponds to specific factor (i.e. female), if question is a factor (i.e. sex).

- `p_adjust` - Vector of numbers that correspond to p-values obtained in all tests. Use this to adjust for p-values if running multiple tests.

4.5 `balance_attrite()`

4.5.1 Description

- Tests whether specified treatment causes attrition in a specified question. Output is a logistic regression, regressing attrition (remain in survey=0, attrited=1) over specified treatment.

4.5.2 Arguments

- `data` - a data.frame, from which `treatment` and `question` are taken.
- `treatment` - a `string` character that corresponds to the name of the treatment variable. Note that values of said variable must be specified as `treatment` and `control`.
- `question` - a `string` character that corresponds to the name of the point in survey (question), for which balance test is required.

4.6 `bounds()`

4.6.1 Description

- Yields extreme (Manski) bounds or trimming (Lee) bounds, using the `attrition` package by Alex Coppock.

4.6.2 Arguments

- `data` - a data.frame, from which `treatment` and `DV` are taken.
- `treatment` - a `string` character that corresponds to the name of the treatment variable. Note that values of said variable must be specified as `treatment` and `control`.
- `DV` - a `string` character that corresponds to the name of the outcome variable.
- `type` - character that corresponds to the type of bounds required ("`Manski`" or "`Lee`"). Default is `type = "Manski"`.

4.7 `vis_miss_treat()`

4.7.1 Description

- Calls the `vis_miss` function from `visdat` package. We allow users to facet missingness by conditions, creating several missingness maps per condition.

4.7.2 Arguments

- `data` - a `data.frame`.
- `treatment_q` - a `string` character that corresponds to the name of the treatment variable. If `treatment_q = NULL`, missingness map appears for all data, when `treatment_q` is not `NULL`, missingness is faceted by condition.

5 Example

Let's begin demonstrating the uses of `attritevis`, with a working example. We begin by installing the package.

We load test data from Lo, Renshon, and Bassan-Nygate 2021 (study 5B) which is an experimental survey study on whether peer-praise can encourage respondents to choose an empathy task.

The experiment manipulates peer-praise and measures empathy in a behavioral task. There are two arms in the peer-praise randomization: peer-praise and no praise (control). In the first arm, a word cloud of praise, drawn from real praise collected in a pilot study, is given for people who behave empathetically, with a line of text about peer group average thermometer ratings towards people who are empathetic – "Peers of yours on this platform have said they hold favorable feelings towards people who engage in empathetic behavior, with an average of 7.9, on a scale of 0 (least favorable) to 10 (most favorable), That same peer group provided real feedback for empathetic behavior which is pictured in the word cloud below". The word cloud is presented in figure 1. Respondents in the control condition do not receive any additional information.



Figure 1: Word cloud of real praise presented to treated respondents.

Our outcome of interest is choosing to empathize with an image in a behavioral task. In the task, subjects choose between two “cards” a FEEL and a DESCRIBE task, that correspond to an empathy or objective task, in which they empathize/describe an image of a man. The cards are presented in figure 2. Below is a description of the survey, with information on the various variables collected.

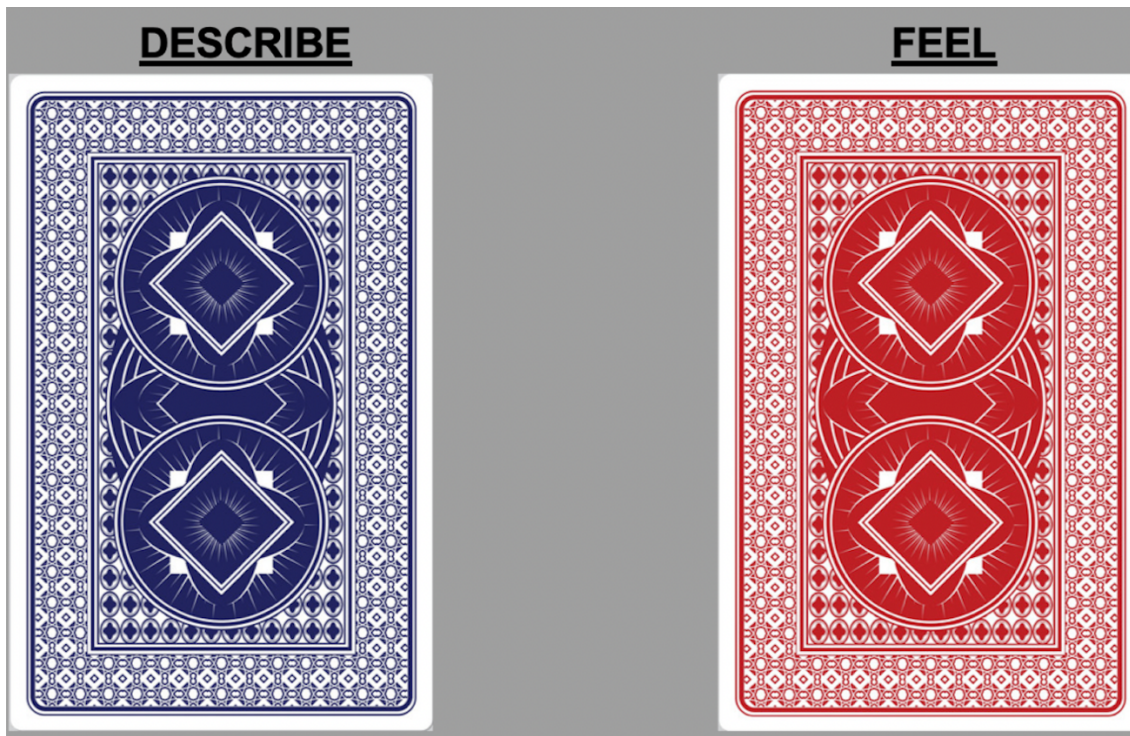


Figure 2: Choice task FEEL and DESCRIBE cards.

After answering pre-treatment covariates, respondents in the study were asked to complete two practice rounds of the main empathy task. After completing the practice rounds, respondents complete three trials of the above mentioned tasks. Before each task, respondents are randomized into treatment and control groups. Treated respondents received the light-touch peer-praise treatment. During each trial, before respondents select between the FEEL and DESCRIBE tasks, happiness, the hypothesized mechanism, is measured. Treatment variables are labeled `treat1`, `treat2`, etc. Outcome variables, which are the choice-task card questions, are labeled `card1`, `card2`, etc. Mediators, which are measures of the emotion happiness, are labeled `Happy_1_1`, `Happy_1_2`... `Happy_2_1`, `Happy_2_2`... `Happy_3_1`, `Happy_3_2`, etc. After respondents complete all three trials post-task and post-treatment covariates are collected. Importantly, the dataframe `test_data` is organized based on the survey questions order. That is, if Q1 came before Q2 in the survey, the variable Q1 comes before the variable Q2 in the dataframe.

After loading the test data and ensuring that variables are ordered by survey questions, we may want to transform our dataframe to an attrition dataframe, using the function `attrition`.

5.1 Attrition dataframe

```
attrition_data <- attritevis::attrition(test_data)
```

This function creates a frame that indicates, per question:

```

- `attrited` -- how many respondents attrited (left the survey) at each question.
- `proportion` -- number of attrited respondents / number of respondents who entered survey.
- `prop_q` -- number of attrited respondents / number of respondents entering into the question.
- `questions` -- question names.
- `responded` -- how many respondents responded in each question.
- `prop_r` -- number of respondents who responded / number of respondents who entered survey.

```

Using base R we can explore how many people attrited overall, and what proportion of the general population this is.

```

sum(attrition_data$attrited) #How many respondents attrited overall?
## [1] 129

sum(attrition_data$attrited)/nrow(test_data) #What proportion of the overall sample is this? (0.21)
## [1] 0.2067308

```

Next, we can look at specific variables, and learn whether respondents attrited. Let's choose the variable `cards_a` to demonstrate. This is a variable that notes whether respondents clicked the "FEEL" or "DESCRIBE" button during their first practice round. Using base R we can extract the number of attrited respondents, as well as the proportion of total N attrited, for this question.

```

attrition_data[attrition_data$questions == 'cards_a', 'attrited']
## [1] 37

attrition_data[attrition_data$questions == 'cards_a', 'proportion']
## [1] 0.06

```

We learn that at the question `cards_a` 37 respondents attrited from the survey. This is equivalent to 6% of the number of respondents who entered the survey at this question. Is this a lot though? Where else do we see attrition in the study? To assess, we visualize attrition across the survey timeline.

5.2 Attrition table

We can further create that of this dataframe using the function `attrition_table`.

```

attritevis::attrition_table(test_data)

```

attrited	prop_q	proportion	questions	responded	prop_r
0	0.00	0.00	consent	624	1.00

attrited	prop_q	proportion	questions	responded	prop_r
3	0.00	0.00	age	621	1.00
0	0.00	0.00	sex	618	0.99
0	0.00	0.00	education	621	1.00
0	0.00	0.00	state	621	1.00
0	0.00	0.00	income	621	1.00
0	0.00	0.00	part_id	621	1.00
0	0.00	0.00	race	621	1.00
0	0.00	0.00	religion	621	1.00
1	0.00	0.00	attrition_1	620	0.99
6	0.01	0.01	attrition_2	614	0.98
37	0.06	0.06	cards_a	577	0.92
0	0.00	0.00	pa	553	0.89
0	0.00	0.00	pb_1	536	0.86
0	0.00	0.00	pb_2	536	0.86
0	0.00	0.00	pb_3	536	0.86
0	0.00	0.00	pc	534	0.86

attrited	prop_q	proportion	questions	responded	prop_r
0	0.00	0.00	cards_b	534	0.86
0	0.00	0.00	p2a	522	0.84
0	0.00	0.00	p2b_1	515	0.83
0	0.00	0.00	p2b_2	515	0.83
0	0.00	0.00	p2b_3	515	0.83
0	0.00	0.00	p2c	515	0.83
0	0.00	0.00	treat1	577	0.92
0	0.00	0.00	Happy_1_1	514	0.82
0	0.00	0.00	Happy_1_2	514	0.82
0	0.00	0.00	Happy_1_3	514	0.82
0	0.00	0.00	cards1	505	0.81
0	0.00	0.00	X1a	502	0.80
0	0.00	0.00	X1b_1	502	0.80
0	0.00	0.00	X1b_2	502	0.80
0	0.00	0.00	X1b_3	502	0.80
0	0.00	0.00	X1c	502	0.80

attrited	prop_q	proportion	questions	responded	prop_r
0	0.00	0.00	treat2	577	0.92
0	0.00	0.00	Happy_2_1	502	0.80
0	0.00	0.00	Happy_2_2	502	0.80
0	0.00	0.00	Happy_2_3	502	0.80
0	0.00	0.00	cards2	500	0.80
0	0.00	0.00	X2a	498	0.80
0	0.00	0.00	X2b_1	497	0.80
0	0.00	0.00	X2b_2	497	0.80
0	0.00	0.00	X2b_3	497	0.80
0	0.00	0.00	X2c	497	0.80
0	0.00	0.00	treat3	577	0.92
80	0.14	0.13	Happy_3_1	497	0.80
0	0.00	0.00	Happy_3_2	497	0.80
0	0.00	0.00	Happy_3_3	497	0.80
0	0.00	0.00	cards3	497	0.80
0	0.00	0.00	X3a	497	0.80

attrited	prop_q	proportion	questions	responded	prop_r
0	0.00	0.00	X3b_1	497	0.80
0	0.00	0.00	X3b_2	497	0.80
0	0.00	0.00	X3b_3	497	0.80
0	0.00	0.00	post1	491	0.79
0	0.00	0.00	post2_7	497	0.80
0	0.00	0.00	post3	497	0.80
0	0.00	0.00	post4	497	0.80
0	0.00	0.00	post5	496	0.79
0	0.00	0.00	post6	497	0.80
1	0.00	0.00	post7	496	0.79
0	0.00	0.00	post8	496	0.79
0	0.00	0.00	post9	496	0.79
0	0.00	0.00	post10	496	0.79
0	0.00	0.00	post11_1	496	0.79
0	0.00	0.00	post11_8	496	0.79
1	0.00	0.00	post13_1	495	0.79

attrited	prop_q	proportion	questions	responded	prop_r
0	0.00	0.00	post14_1	495	0.79
0	0.00	0.00	post15_1	493	0.79
0	0.00	0.00	post16_1	494	0.79
0	0.00	0.00	post17	495	0.79
0	0.00	0.00	ideology	495	0.79
0	0.00	0.00	trump_approval	495	0.79
0	0.00	0.00	pres_approval	495	0.79

We can also use the argument `treatment_q` to facet attrition table by condition. This is a character that corresponds to a specific variable, which is where the treatment conditions were administered.

```
attritevis::attrition_table(data= test_data,
                             treatment_q = "treat1"
                             )
```

[[1]]

attrited	prop_q	proportion	questions	responded	prop_r
0	0.00	0.00	control.consent	289	1.00
0	0.00	0.00	control.age	289	1.00
0	0.00	0.00	control.sex	287	0.99
0	0.00	0.00	control.education	289	1.00
0	0.00	0.00	control.state	289	1.00

attrited	prop_q	proportion	questions	responded	prop_r
0	0.00	0.00	control.income	289	1.00
0	0.00	0.00	control.part_id	289	1.00
0	0.00	0.00	control.race	289	1.00
0	0.00	0.00	control.religion	289	1.00
0	0.00	0.00	control.attrition_1	289	1.00
0	0.00	0.00	control.attrition_2	289	1.00
0	0.00	0.00	control.cards_a	289	1.00
0	0.00	0.00	control.pa	278	0.96
0	0.00	0.00	control.pb_1	270	0.93
0	0.00	0.00	control.pb_2	270	0.93
0	0.00	0.00	control.pb_3	270	0.93
0	0.00	0.00	control.pc	268	0.93
0	0.00	0.00	control.cards_b	268	0.93
0	0.00	0.00	control.p2a	262	0.91
0	0.00	0.00	control.p2b_1	260	0.90
0	0.00	0.00	control.p2b_2	260	0.90

attrited	prop_q	proportion	questions	responded	prop_r
0	0.00	0.00	control.p2b_3	260	0.90
0	0.00	0.00	control.p2c	260	0.90
0	0.00	0.00	control.cond_new	289	1.00
0	0.00	0.00	control.Happy_1_1	260	0.90
0	0.00	0.00	control.Happy_1_2	260	0.90
0	0.00	0.00	control.Happy_1_3	260	0.90
0	0.00	0.00	control.cards1	251	0.87
0	0.00	0.00	control.X1a	249	0.86
0	0.00	0.00	control.X1b_1	249	0.86
0	0.00	0.00	control.X1b_2	249	0.86
0	0.00	0.00	control.X1b_3	249	0.86
0	0.00	0.00	control.X1c	249	0.86
0	0.00	0.00	control.treat2	289	1.00
0	0.00	0.00	control.Happy_2_1	249	0.86
0	0.00	0.00	control.Happy_2_2	249	0.86
0	0.00	0.00	control.Happy_2_3	249	0.86

attrited	prop_q	proportion	questions	responded	prop_r
0	0.00	0.00	control.cards2	248	0.86
0	0.00	0.00	control.X2a	246	0.85
0	0.00	0.00	control.X2b_1	246	0.85
0	0.00	0.00	control.X2b_2	246	0.85
0	0.00	0.00	control.X2b_3	246	0.85
0	0.00	0.00	control.X2c	246	0.85
0	0.00	0.00	control.treat3	289	1.00
43	0.15	0.15	control.Happy_3_1	246	0.85
0	0.00	0.00	control.Happy_3_2	246	0.85
0	0.00	0.00	control.Happy_3_3	246	0.85
0	0.00	0.00	control.cards3	246	0.85
0	0.00	0.00	control.X3a	246	0.85
0	0.00	0.00	control.X3b_1	246	0.85
0	0.00	0.00	control.X3b_2	246	0.85
0	0.00	0.00	control.X3b_3	246	0.85
0	0.00	0.00	control.post1	244	0.84

attrited	prop_q	proportion	questions	responded	prop_r
0	0.00	0.00	control.post2_7	246	0.85
0	0.00	0.00	control.post3	246	0.85
0	0.00	0.00	control.post4	246	0.85
0	0.00	0.00	control.post5	246	0.85
0	0.00	0.00	control.post6	246	0.85
0	0.00	0.00	control.post7	246	0.85
0	0.00	0.00	control.post8	246	0.85
0	0.00	0.00	control.post9	246	0.85
0	0.00	0.00	control.post10	246	0.85
0	0.00	0.00	control.post11_1	246	0.85
0	0.00	0.00	control.post11_8	246	0.85
0	0.00	0.00	control.post13_1	246	0.85
0	0.00	0.00	control.post14_1	246	0.85
0	0.00	0.00	control.post15_1	244	0.84
0	0.00	0.00	control.post16_1	245	0.85
0	0.00	0.00	control.post17	246	0.85

attrited	prop_q	proportion	questions	responded	prop_r
0	0.00	0.00	control.ideology	246	0.85
0	0.00	0.00	control.trump_approval	246	0.85
0	0.00	0.00	control.pres_approval	246	0.85

[[2]]

attrited	prop_q	proportion	questions	responded	prop_r
0	0.00	0.00	treatment.consent	288	1.00
0	0.00	0.00	treatment.age	288	1.00
0	0.00	0.00	treatment.sex	288	1.00
0	0.00	0.00	treatment.education	288	1.00
0	0.00	0.00	treatment.state	288	1.00
0	0.00	0.00	treatment.income	288	1.00
0	0.00	0.00	treatment.part_id	288	1.00
0	0.00	0.00	treatment.race	288	1.00
0	0.00	0.00	treatment.religion	288	1.00
0	0.00	0.00	treatment.attrition_1	288	1.00
0	0.00	0.00	treatment.attrition_2	288	1.00
0	0.00	0.00	treatment.cards_a	288	1.00

attrited	prop_q	proportion	questions	responded	prop_r
0	0.00	0.00	treatment.pa	275	0.95
0	0.00	0.00	treatment.pb_1	266	0.92
0	0.00	0.00	treatment.pb_2	266	0.92
0	0.00	0.00	treatment.pb_3	266	0.92
0	0.00	0.00	treatment.pc	266	0.92
0	0.00	0.00	treatment.cards_b	266	0.92
0	0.00	0.00	treatment.p2a	260	0.90
0	0.00	0.00	treatment.p2b_1	255	0.89
0	0.00	0.00	treatment.p2b_2	255	0.89
0	0.00	0.00	treatment.p2b_3	255	0.89
0	0.00	0.00	treatment.p2c	255	0.89
0	0.00	0.00	treatment.cond_new	288	1.00
0	0.00	0.00	treatment.Happy_1_1	254	0.88
0	0.00	0.00	treatment.Happy_1_2	254	0.88
0	0.00	0.00	treatment.Happy_1_3	254	0.88
0	0.00	0.00	treatment.cards1	254	0.88

attrited	prop_q	proportion	questions	responded	prop_r
0	0.00	0.00	treatment.X1a	253	0.88
0	0.00	0.00	treatment.X1b_1	253	0.88
0	0.00	0.00	treatment.X1b_2	253	0.88
0	0.00	0.00	treatment.X1b_3	253	0.88
0	0.00	0.00	treatment.X1c	253	0.88
0	0.00	0.00	treatment.treat2	288	1.00
0	0.00	0.00	treatment.Happy_2_1	253	0.88
0	0.00	0.00	treatment.Happy_2_2	253	0.88
0	0.00	0.00	treatment.Happy_2_3	253	0.88
0	0.00	0.00	treatment.cards2	252	0.88
0	0.00	0.00	treatment.X2a	252	0.88
0	0.00	0.00	treatment.X2b_1	251	0.87
0	0.00	0.00	treatment.X2b_2	251	0.87
0	0.00	0.00	treatment.X2b_3	251	0.87
0	0.00	0.00	treatment.X2c	251	0.87
0	0.00	0.00	treatment.treat3	288	1.00

attrited	prop_q	proportion	questions	responded	prop_r
37	0.13	0.13	treatment.Happy_3_1	251	0.87
0	0.00	0.00	treatment.Happy_3_2	251	0.87
0	0.00	0.00	treatment.Happy_3_3	251	0.87
0	0.00	0.00	treatment.cards3	251	0.87
0	0.00	0.00	treatment.X3a	251	0.87
0	0.00	0.00	treatment.X3b_1	251	0.87
0	0.00	0.00	treatment.X3b_2	251	0.87
0	0.00	0.00	treatment.X3b_3	251	0.87
0	0.00	0.00	treatment.post1	247	0.86
0	0.00	0.00	treatment.post2_7	251	0.87
0	0.00	0.00	treatment.post3	251	0.87
0	0.00	0.00	treatment.post4	251	0.87
0	0.00	0.00	treatment.post5	250	0.87
0	0.00	0.00	treatment.post6	251	0.87
1	0.00	0.00	treatment.post7	250	0.87
0	0.00	0.00	treatment.post8	250	0.87

attrited	prop_q	proportion	questions	responded	prop_r
0	0.00	0.00	treatment.post9	250	0.87
0	0.00	0.00	treatment.post10	250	0.87
0	0.00	0.00	treatment.post11_1	250	0.87
0	0.00	0.00	treatment.post11_8	250	0.87
1	0.00	0.00	treatment.post13_1	249	0.86
0	0.00	0.00	treatment.post14_1	249	0.86
0	0.00	0.00	treatment.post15_1	249	0.86
0	0.00	0.00	treatment.post16_1	249	0.86
0	0.00	0.00	treatment.post17	249	0.86
0	0.00	0.00	treatment.ideology	249	0.86
0	0.00	0.00	treatment.trump_approval	249	0.86
0	0.00	0.00	treatment.pres_approval	249	0.86

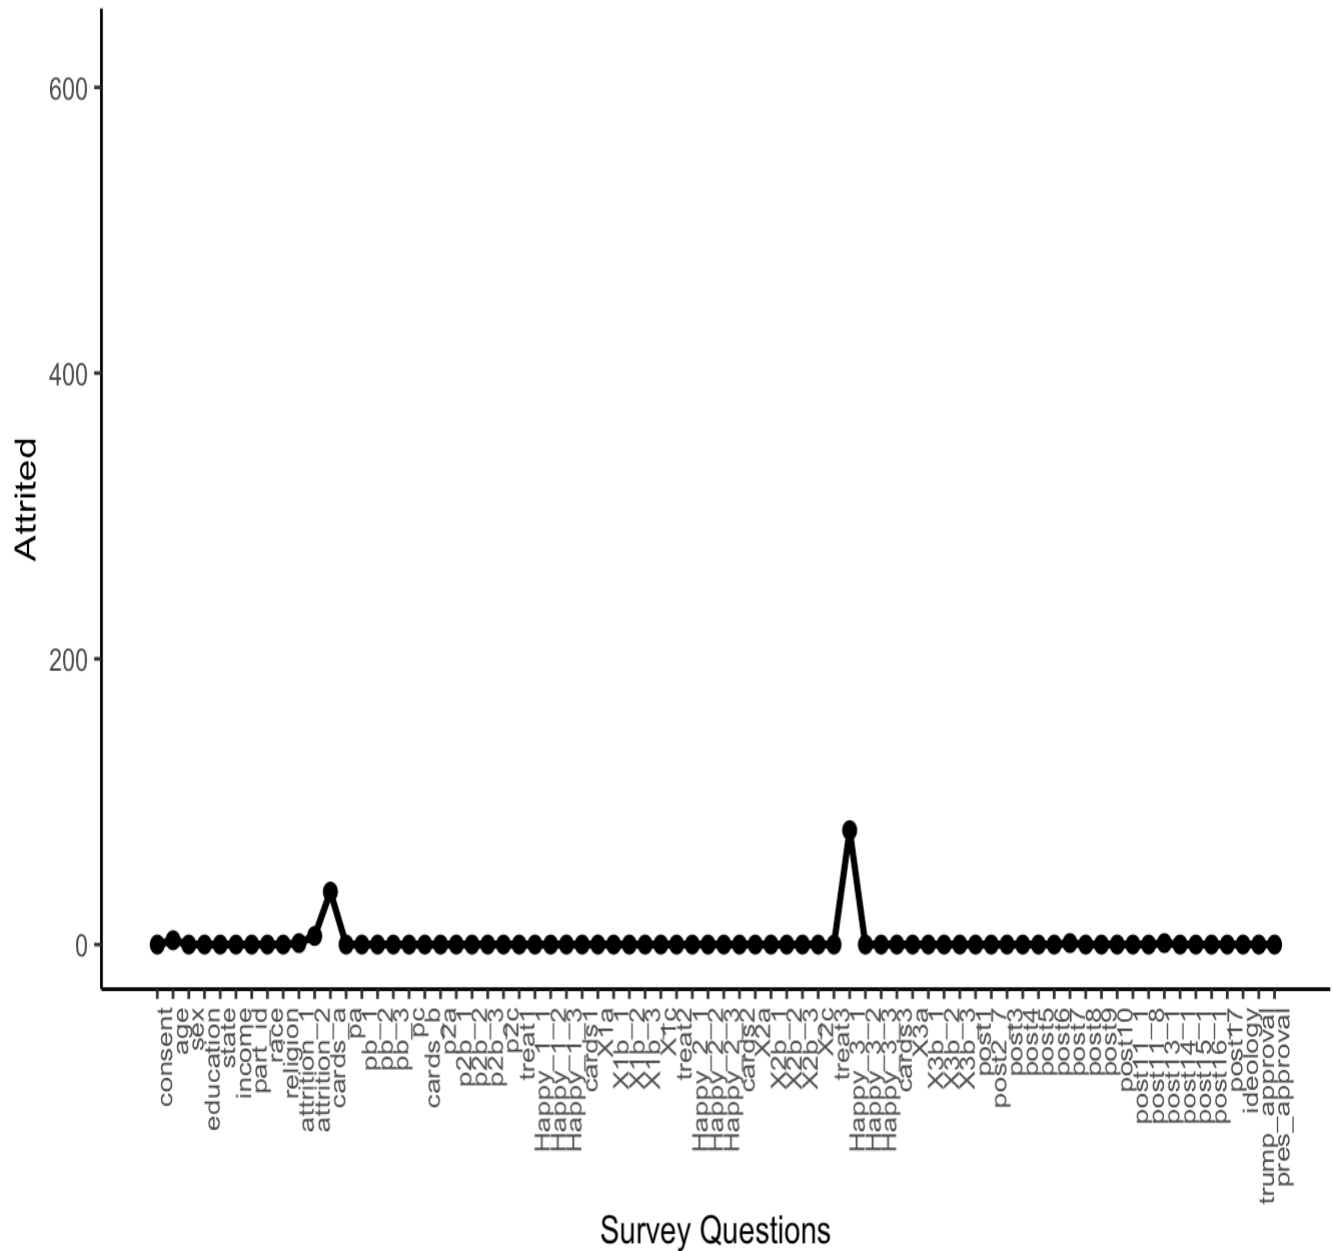
5.3 Visualizing attrition

5.3.1 Attrition timeline

We may want to visualize attrition across the survey, to look at all the survey questions at once. The function `plot_attrition` allows us to plot attrition across survey questions, indicating where treatment and outcome questions were collected.

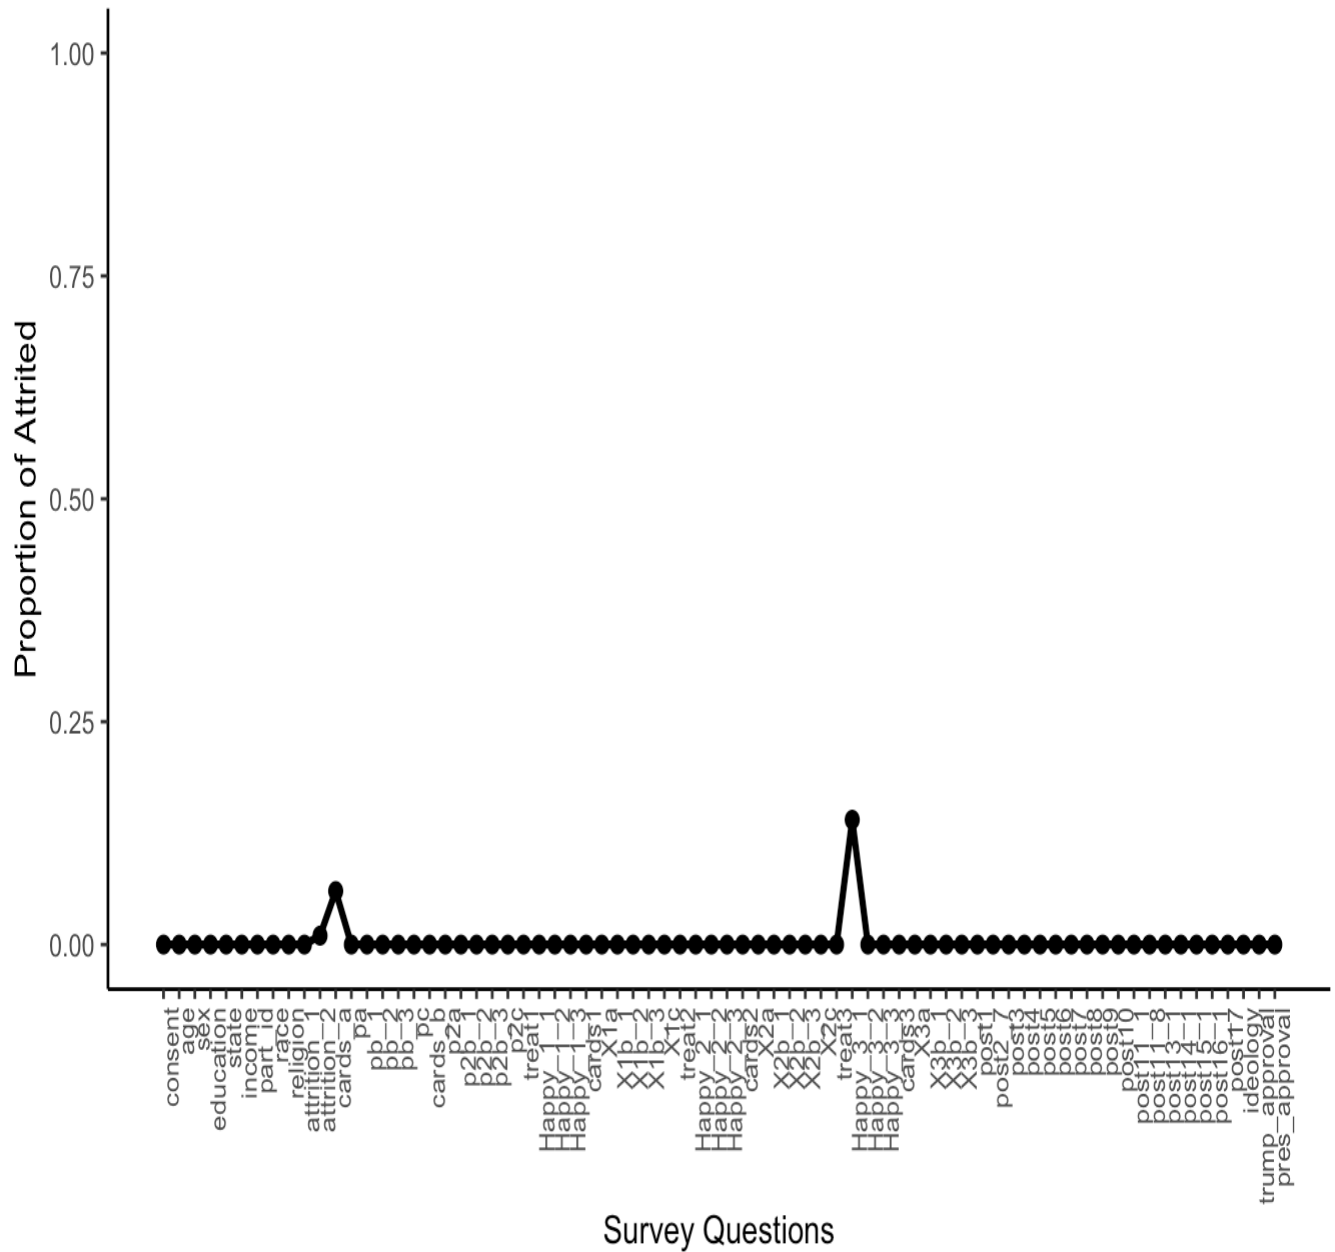
There are several ways in which users can use this function. Simply plugging in the dataset into the function yields a figure that plots the number of respondents that attrited (left the survey completely) over each question in the study.


```
attritevis::plot_attrition(test_data)
```



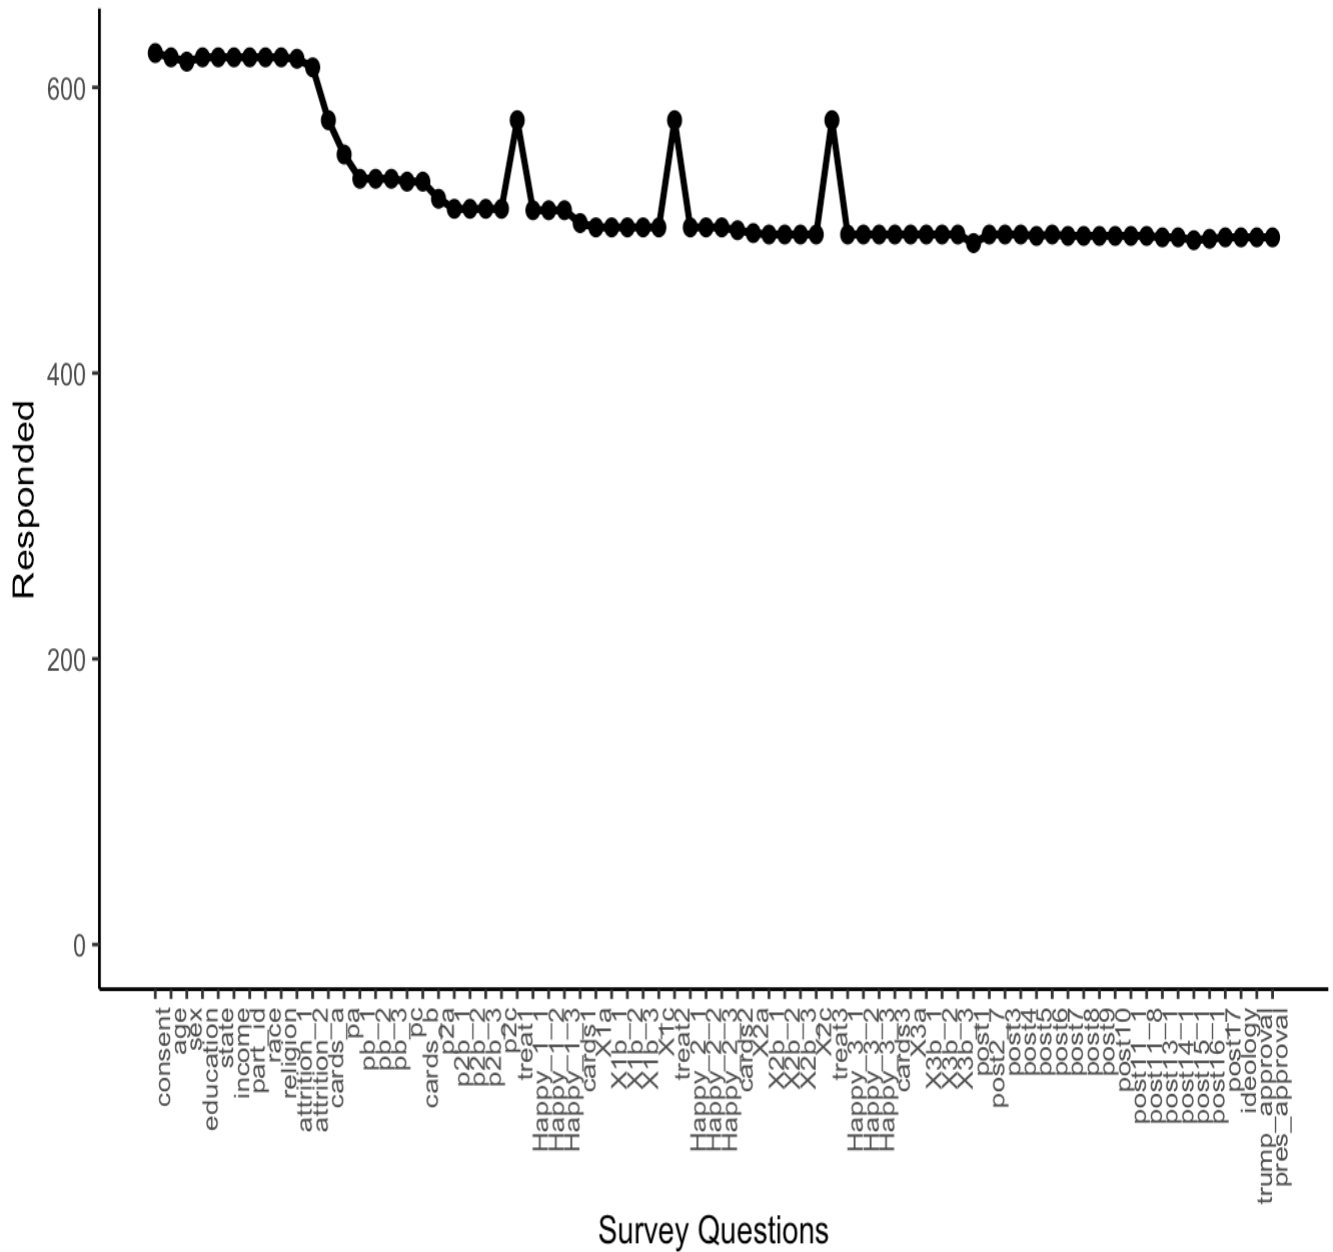
When users specify `freq=FALSE`, the y axis plots the proportion of attrited.

```
attritevis::plot_attrition(test_data,  
                             freq=FALSE)
```



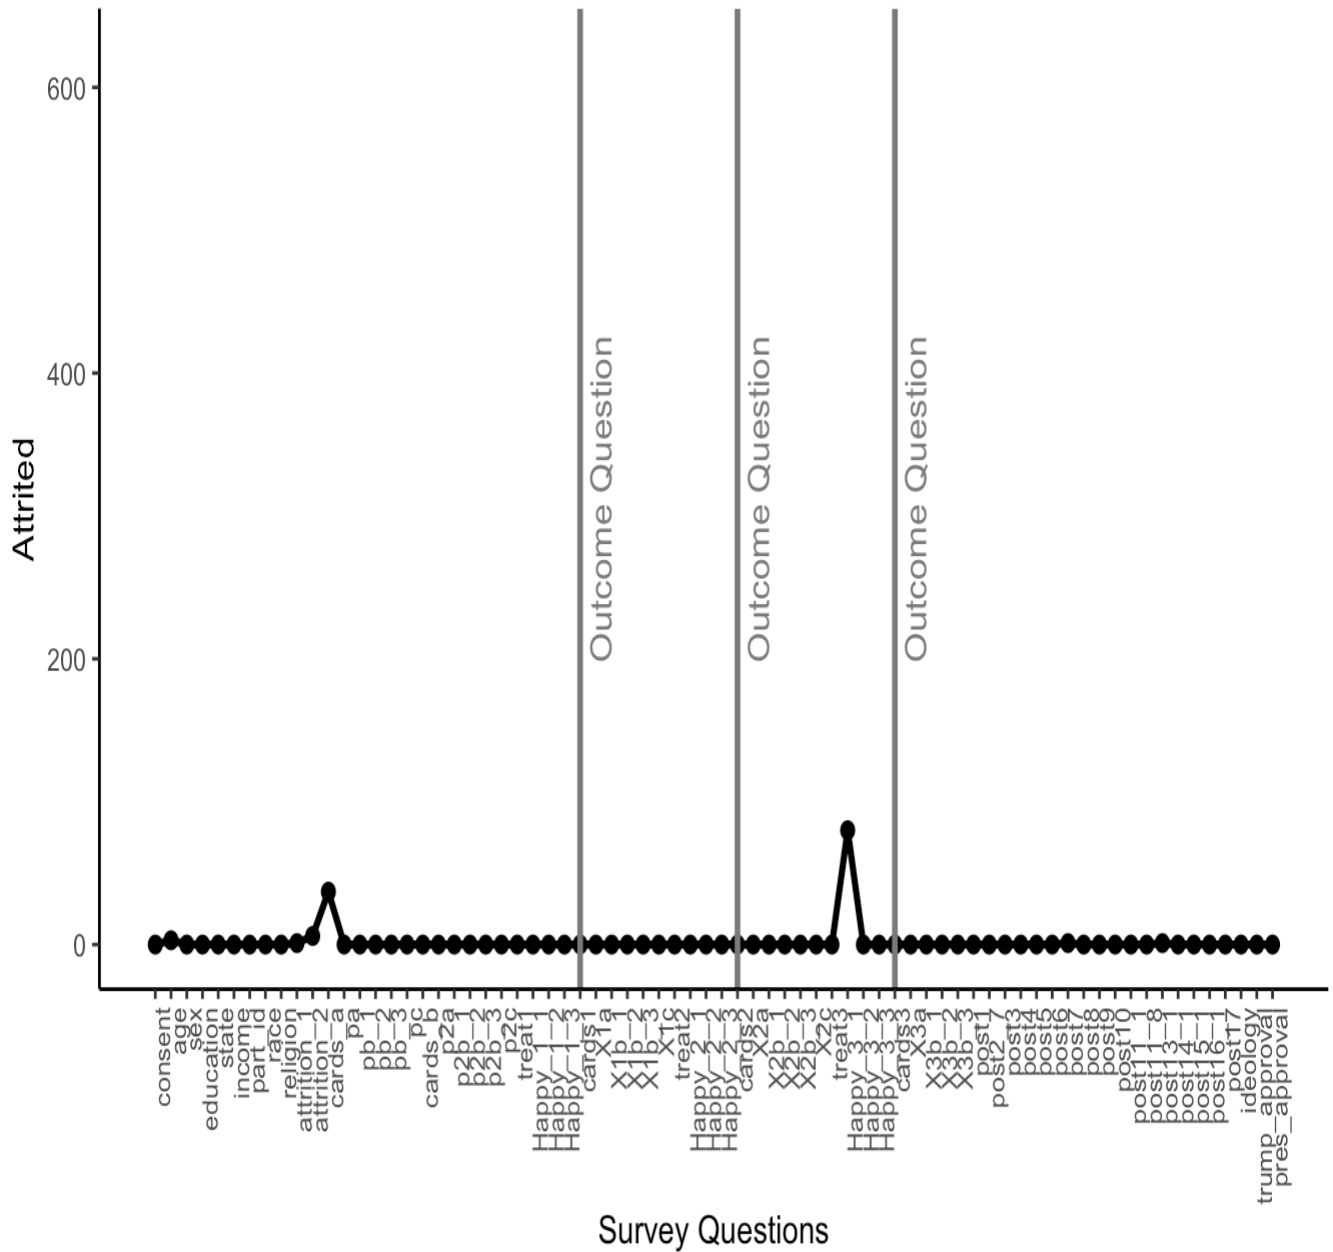
Users can further specify `y="resonded"` to account for response, rather than attrition. This argument can be used with either `freq=TRUE` (default), or `freq=FALSE`, plotting response or proportion of responded, accordingly.

```
attritevis::plot_attrition(test_data,
                           y="resonded")
```



Using the `outcome_q` argument, users can specify where outcome questions were measure. These are noted with gray vertical lines.

```
attritevis::plot_attrition(test_data,
  outcome_q = c("cards1", "cards2", "cards3"))
```

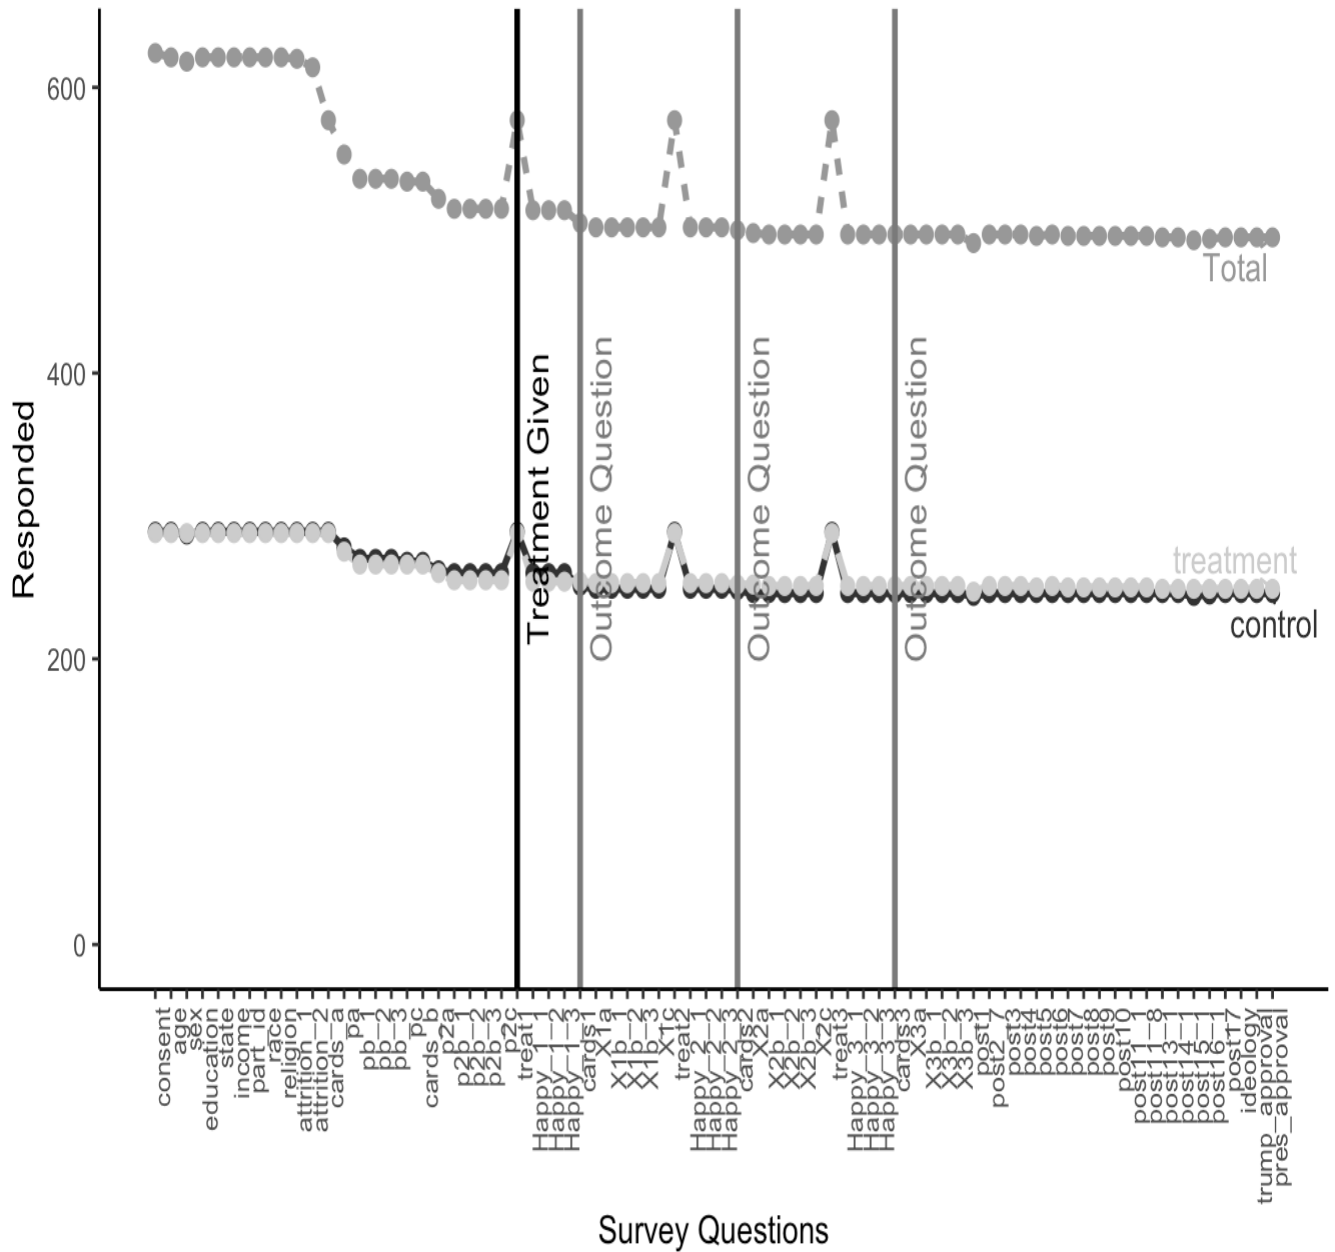


When `treatment_q`, which corresponds to treatment variable, is not NULL, the plot both notes where treatment was collected with a vertical line, and breaks down attrition by treatment conditions.

```

attritevis::plot_attrition(test_data,
  y = "responded",
  outcome_q = c("cards1", "cards2", "cards3"),
  treatment_q = "treat1")

```



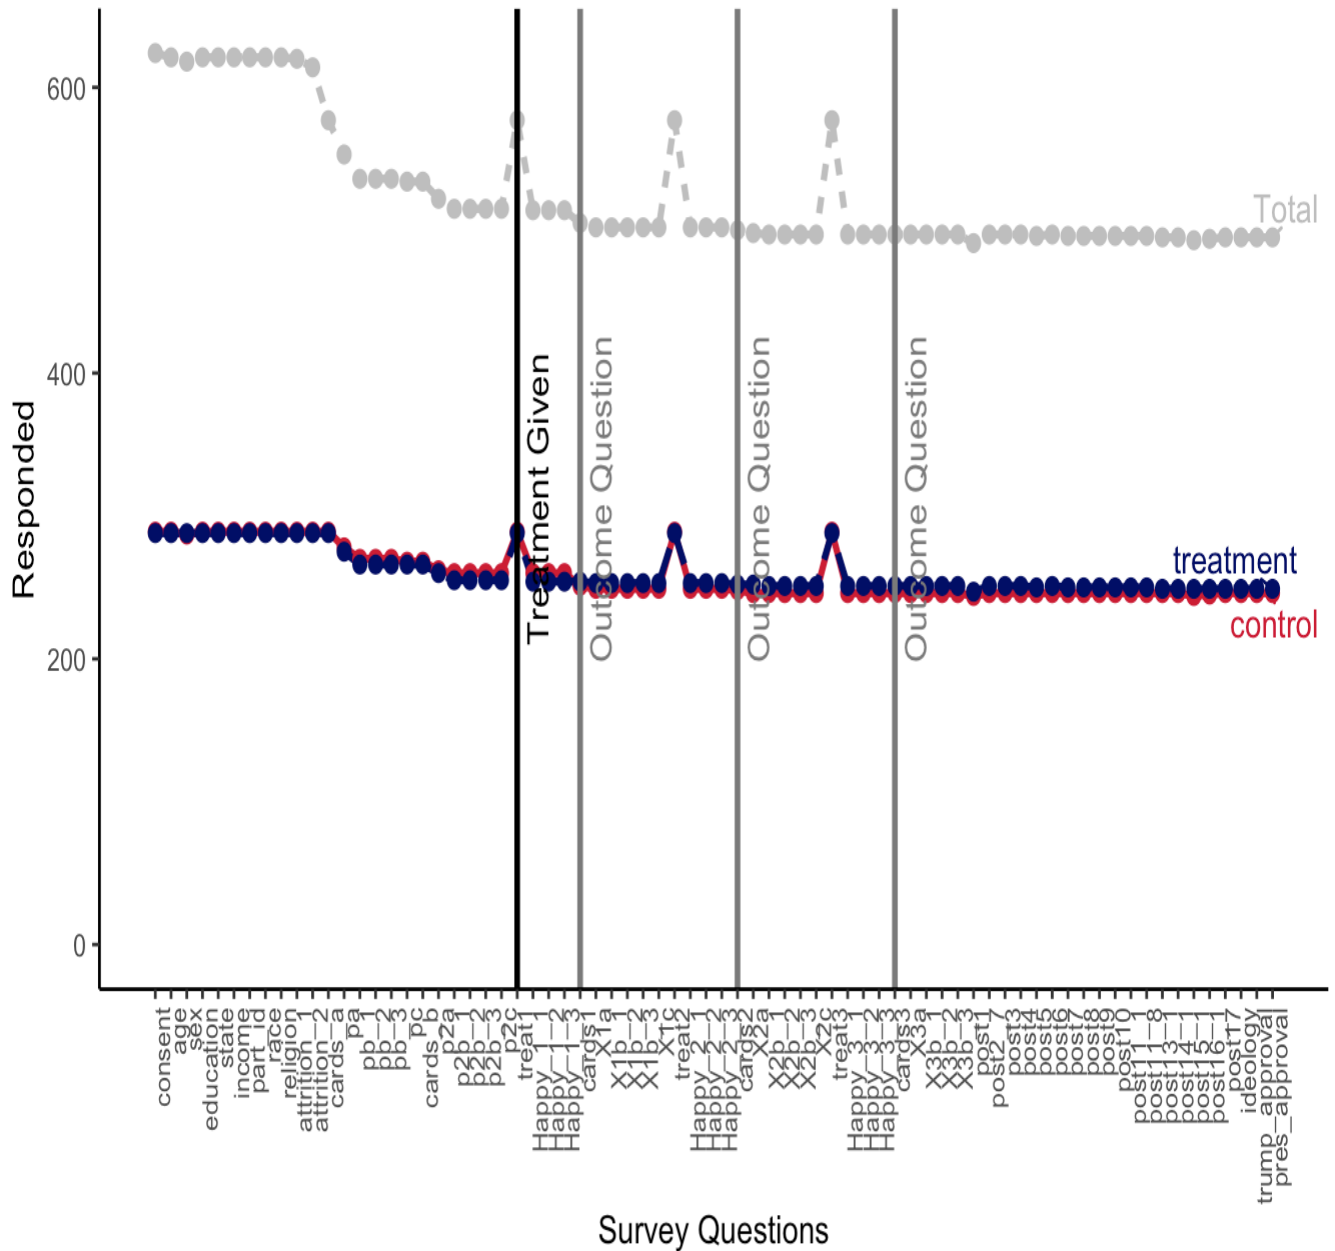
Color default is greyscale, but users can use the `mycolors` argument to specify which colors they want to use to mark each conditions' geom_line. The length of `mycolors` must be equal to the length of `unique(treatment_q)`. To use this argument, users should specify which color corresponds to which factor. See below the running example:

```

attritevis::plot_attrition(test_data,
  y = "responded",
  outcome_q = c("cards1", "cards2", "cards3"),
  treatment_q = "treat1",
  mycolors = c(treatment = "#000066",

```

```
control = "#CC0033"))
```



Finally, users can make manual changes to the plot by removing the `Total` line (using the argument `total = FALSE`), and by removing the treatment and/or outcome vertical lines (using `tline=FALSE`, and/or `outcomeline=FALSE`, respectively). Generally, we recommend keeping the `Total` line unless treatment is administered at with entrance to study.

```
attritevis::plot_attrition(test_data,  
  y = "responded",  
  outcome_q = c("cards1", "cards2", "cards3"),  
  treatment_q = "treat1",
```

```

mycolors = c(treatment = "#000066",
             control = "#CC0033"),

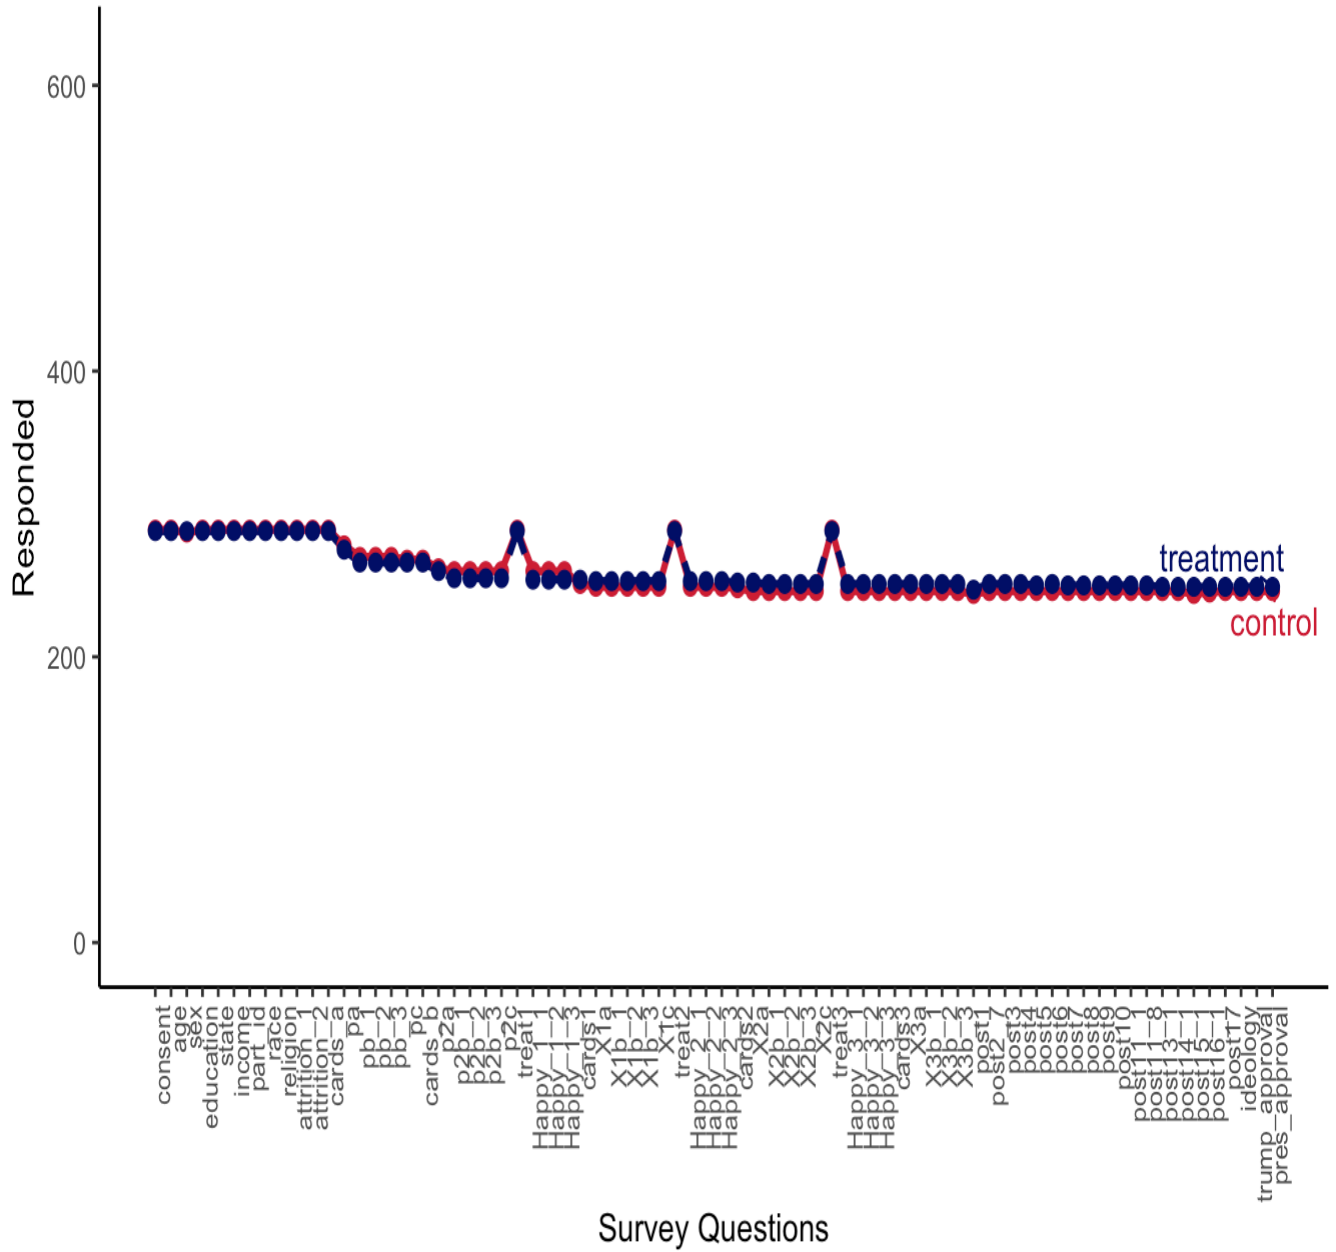
total = FALSE,

outcomeline = FALSE,

tline = FALSE

)

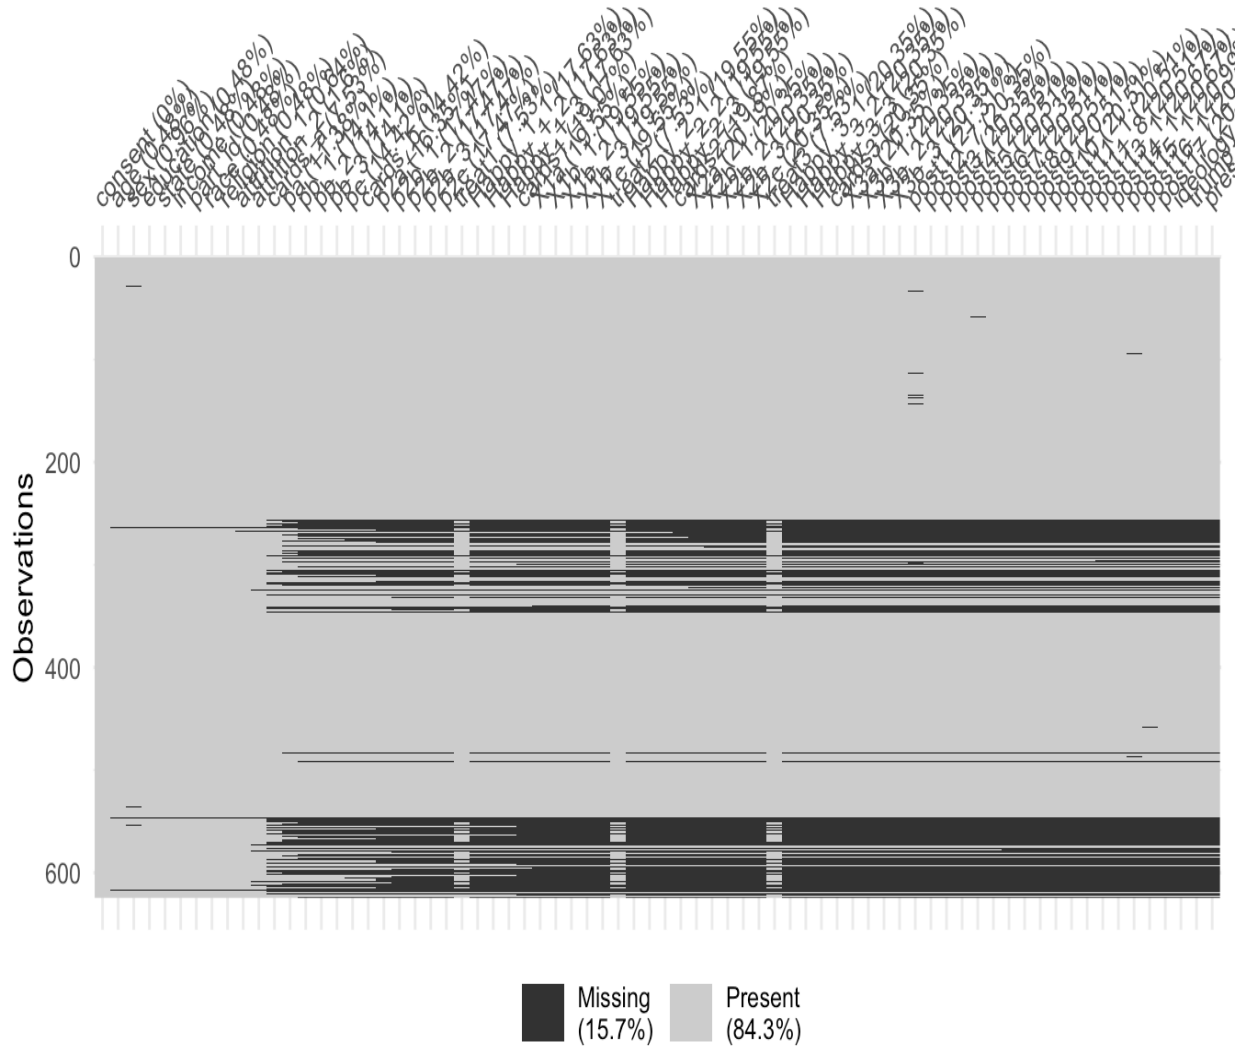
```



5.3.2 Vis miss

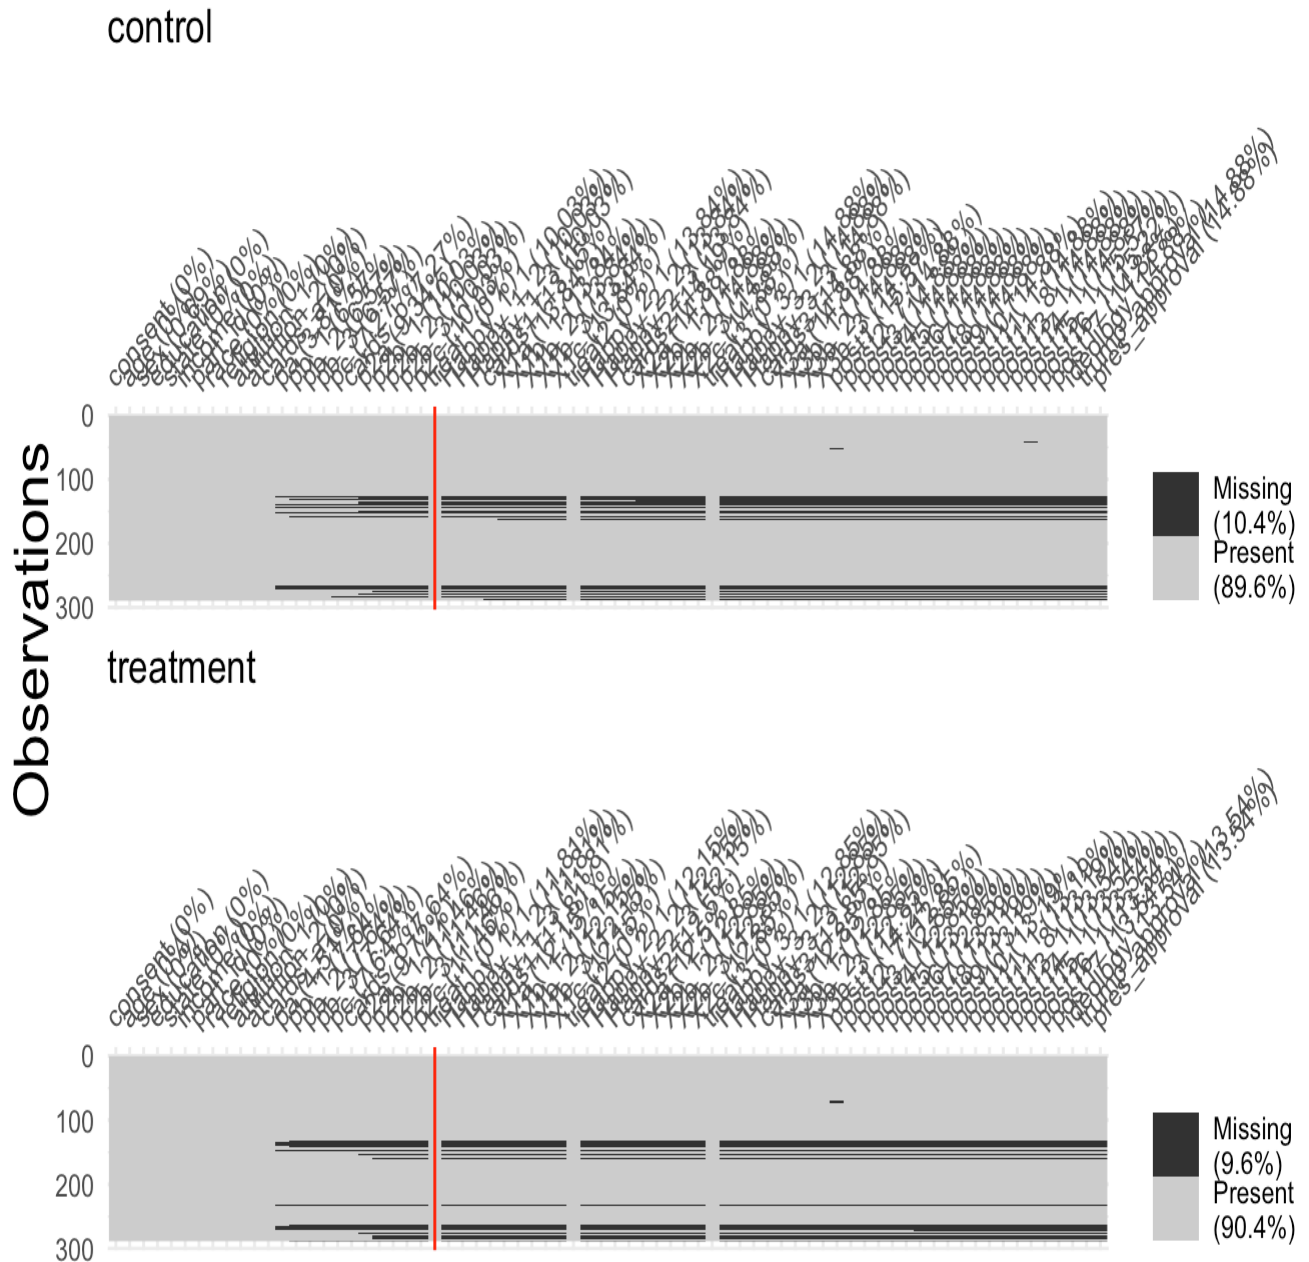
Users can also visualize missingness with the `vis_miss_treat()` function that calls the `vis_miss` function from the `visdat` package.

```
attritevis::vis_miss_treat(test_data)
```



`attritevis` allows users to facet missingness by conditions, creating several missingness maps per condition, and marks treatment variable with a red vertical line.

```
attritevis::vis_miss_treat(test_data,  
  treatment_q = "treat1")
```

5.4 Balance tests

Once we have identified the specific survey points where attrition takes place, we want to conduct balance tests at these specific points to ensure balance across treatment and control, and learn if (and when) balance became an issue. We can do this using the functions `balance_cov()` and `balance_attrite()`.

5.4.1 Balance across covariates

Once we've identified whether (and when) attrition occurs in our survey, we want to know that our treatment and control groups are balanced across covariates throughout the survey, to detect

differential attrition. We can do this using the function `balance_cov()`, which we will demonstrate with three covariates: `age`, `sex`, and `ideology`.

We begin with the covariate `age`, which was collected pretreatment and is a numeric variable. In order to use the function `balance_cov()` we must define `treatment` and `control` arms under the `treatment` variables. We define `treat1` as the treatment variable, and `age` as the question.

```
unique(test_data$treat1)
## [1] "treatment" "control"   NA
attritevis::balance_cov(data = test_data,
  treatment = "treat1",
  question = "age")
##
## Welch Two Sample t-test
##
## data:  treat_data$question1 and control_data$question1
## t = -0.32688, df = 568.57, p-value = 0.7439
## alternative hypothesis: true difference in means is not equal to 0
## 95 percent confidence interval:
## -2.002600  1.431137
## sample estimates:
## mean of x mean of y
##  37.42361  37.70934
```

The output is a t-test that determines whether there is a difference between the average age of the control group and the treatment group. We learn that age is balanced across treatment and control groups, with a mean of approximately 37.4 years old in treated respondents and 37.7 in controlled respondents ($p=0.7$).

We can also use the function `balance_cov()` when the covariate (`question`) is a factor, but we must specify which factor we are interested in. For example, let's say we want to test whether at the question `sex` in the survey missingness created observable differences across treatment and control groups. Sex is a factor variable with two factors: female and male. We can look at whether the proportion of female still remains similar across groups. To do so, we must determine that `factor = TRUE` and specify the `factor_name` (in this case, female).

```
attritevis::balance_cov(data = test_data,
  treatment = "treat1",
  question = "sex",
  factor = TRUE,
  factor_name = "female")
##
## 2-sample test for equality of proportions with continuity correction
```

```
##
## data:  x out of n
## X-squared = 1.1305, df = 1, p-value = 0.2877
## alternative hypothesis: two.sided
## 95 percent confidence interval:
##  -0.12931498  0.03623038
## sample estimates:
##   prop 1    prop 2
## 0.3576389 0.4041812
```

The output is a 2-sample proportion test. We learn that sex is also balanced between treatment and control, with similar proportions of females across the groups ($p=0.3$).

There are certain post-treatment variables for which we may want to ensure balance across treatment and control as well. Note, however, that these should be variables that we hypothesize would stay stable after treatment. For example, we occasionally include demographic questions at the end of the survey to avoid survey fatigue before treatments. In our running example, the `ideology` question was collected post-treatment, but we expect it to stay stable across treatment and control.

```
attritevis::balance_cov(data = test_data,
                        treatment = "treat1",
                        question = "ideology")
##
## Welch Two Sample t-test
##
## data:  treat_data$question1 and control_data$question1
## t = 1.023, df = 492.91, p-value = 0.3068
## alternative hypothesis: true difference in means is not equal to 0
## 95 percent confidence interval:
##  -0.1660012  0.5266633
## sample estimates:
## mean of x mean of y
##  3.879518  3.699187
```

If users run several balance tests, we recommend adjusting p-values. `balance_cov` allows users to do so by specifying the p-values of the tests they ran (`p_adjust`).

In our running example, since we ran three balance tests (with ideology, sex, and age), we want to adjust our p-values:

```
attritevis::balance_cov(data = test_data,
                        treatment = "treat1",
```

```

    question = "ideology",
    p_adjust = c(0.7,0.3,0.3))
##
## Welch Two Sample t-test
##
## data:  treat_data$question1 and control_data$question1
## t = 1.023, df = 492.91, p-value = 0.3068
## alternative hypothesis: true difference in means is not equal to 0
## 95 percent confidence interval:
##  -0.1660012  0.5266633
## sample estimates:
## mean of x mean of y
##  3.879518  3.699187

```

Original_p....p_adjust	Adjusted_p....a
0.7	0.70
0.3	0.45
0.3	0.45

5.4.2 Balance across attrition

Next, we can check whether our treatment is correlated with attrition at any moment in the survey. The `balance_attrite()` function converts the specified `question` into a binary variable such that `attrition = 1`, and `remaining in survey = 0`, and runs a logistic regression (regressing the specified `question` over the specified `treatment`) to examine whether treatment affects attrition.

Using our visualization, we identified that attrition occurs at the post-treatment question `Happy_3_1`. We can use the function `balance_attrite()`, to examine whether our treatment caused attrition at this point in the survey:

```

attritevis::balance_attrite(data = test_data,
    treatment = "treat1",
    question = "Happy_3_1")
##
## Call:
## glm(formula = question1 ~ treatment1, family = binomial(link = "logit"),

```

```
##      data = data2)
##
## Deviance Residuals:
##      Min        1Q    Median        3Q        Max
## -0.5676  -0.5676  -0.5244  -0.5244   2.0259
##
## Coefficients:
##              Estimate Std. Error z value Pr(>|z|)
## (Intercept)      -1.7441     0.1653 -10.552  <2e-16 ***
## treatment1treatment  -0.1704     0.2415  -0.706    0.48
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## (Dispersion parameter for binomial family taken to be 1)
##
##      Null deviance: 464.49  on 576  degrees of freedom
## Residual deviance: 463.99  on 575  degrees of freedom
## (47 observations deleted due to missingness)
## AIC: 467.99
##
## Number of Fisher Scoring iterations: 4
```

We learn that our `treat1` does not affect attrition in variable `Happy_3_1`.

6 Simulated data

As we demonstrated above, attrition doesn't seem to pose a threat to inference in our dataset. But what does it look like when attrition is an issue? We simulate attrition on `test_data` to demonstrate what this would look like.

6.1 Treatment causes attrition

In a toy example (`test_sim`), suppose respondents enter a survey (Q1-Q10), where treatment is assigned at "Q5". We generate attrition such that treatment is causing respondents to drop out of the survey right after treatment. We might see something like this if respondents are particularly taxed by a treatment in the survey and therefore more likely to drop out after receiving treatment.

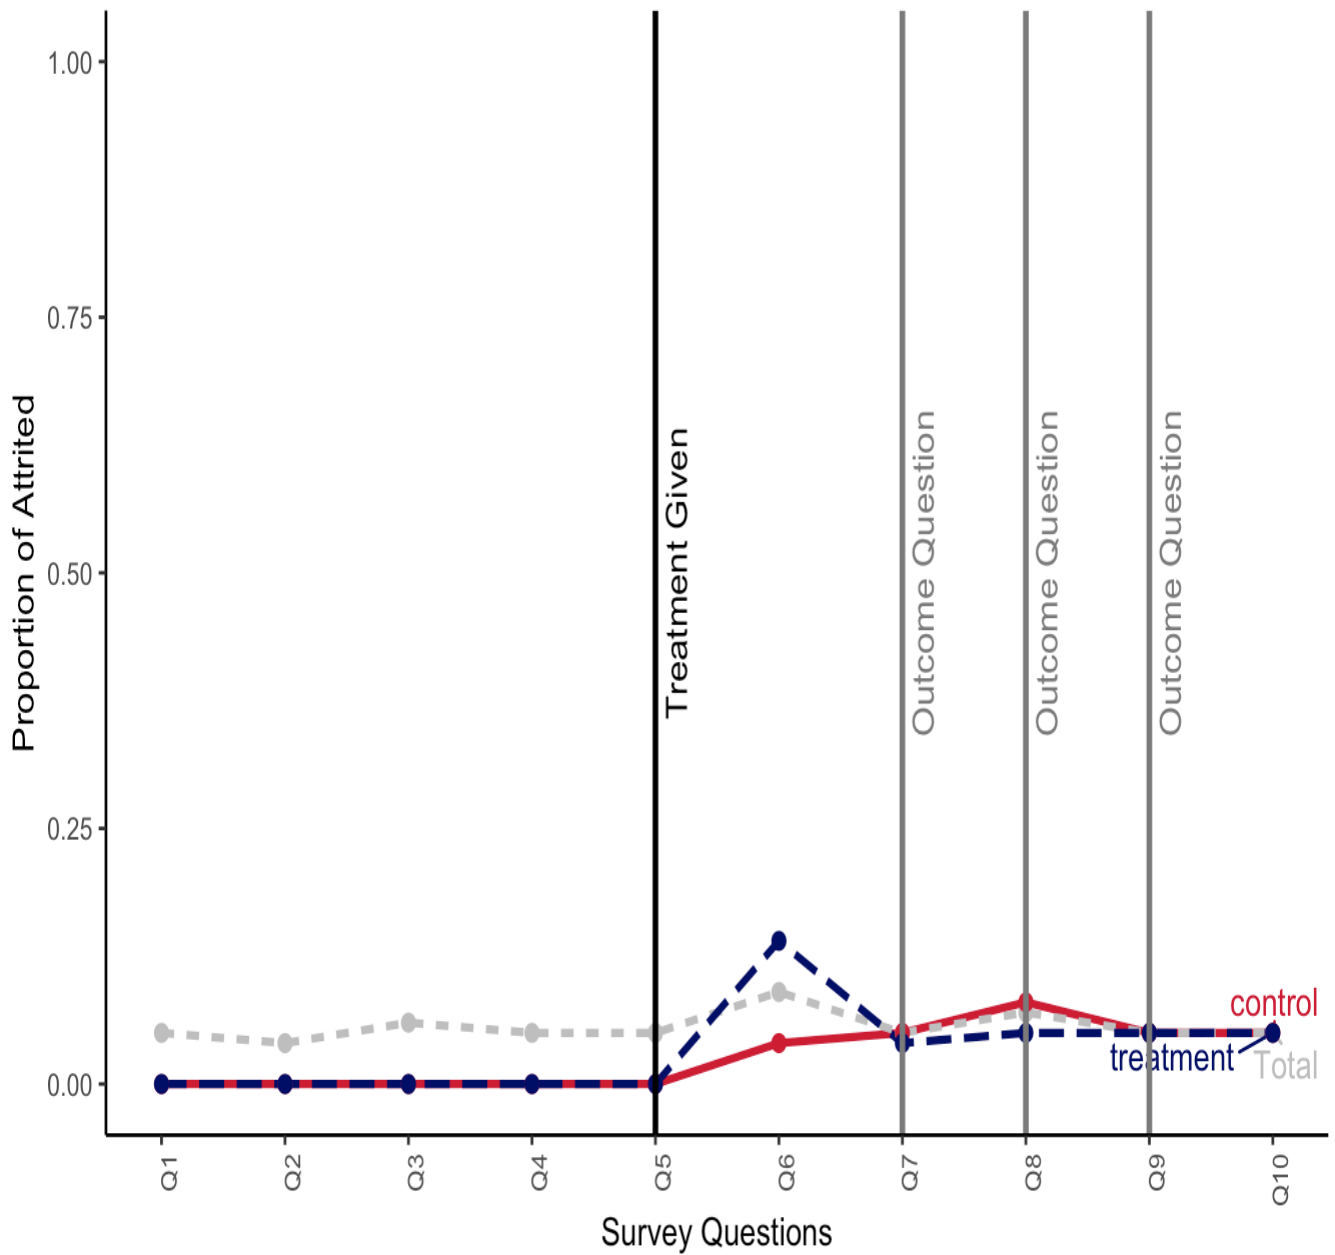
6.1.1 Plot

We visualize attrition using the `plot_attrition()` and `vis_miss_treat()` functions.

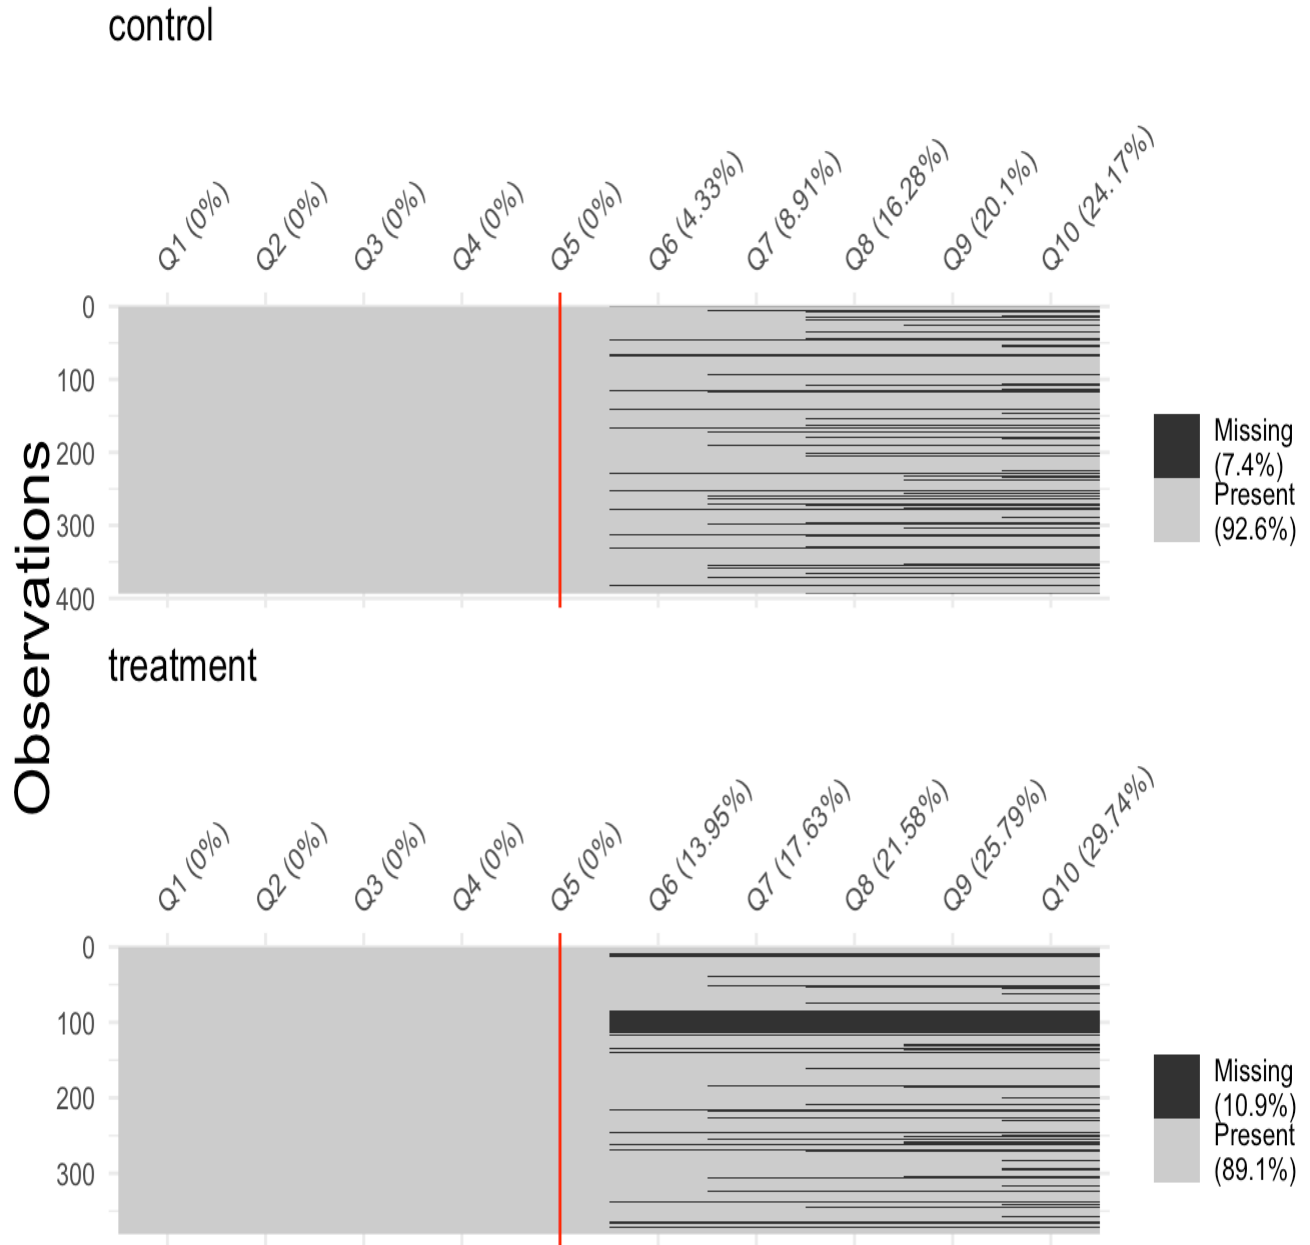
```

attritevis::plot_attrition(test_sim,
  treatment_q = "Q5",
  outcome_q = c("Q7", "Q8", "Q9"),
  freq = FALSE,
  mycolors = c(treatment = "#000066",
               control = "#CC0033")
)

```



```
attritevis::vis_miss_treat(test_sim, treatment_q = "Q5")
```



We learn that attrition mostly occurs after Q6, and that treated respondents seem to be attriting more.

6.1.2 Balance

We learn that most respondents attrite at the post-treatment question Q6, and conduct a balance test. Note that Q6 is an outcome, and we expect our treatment to affect it. It thus does not make sense to use the `balance_cov()` function. Instead, we want to examine whether our treatment caused attrition, and thus use the function `balance_attrite()`:

```

attritevis::balance_attrite(data = test_sim,
  treatment = "Q5",
  question = "Q6")

##
## Call:
## glm(formula = question1 ~ treatment1, family = binomial(link = "logit"),
## data = data2)
##
## Deviance Residuals:
##      Min       1Q   Median       3Q      Max
## -0.5481  -0.5481  -0.2974  -0.2974   2.5062
##
## Coefficients:
##              Estimate Std. Error z value Pr(>|z|)
## (Intercept)      -3.0964     0.2479 -12.489 < 2e-16 ***
## treatment1treatment  1.2767     0.2888   4.421 9.82e-06 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## (Dispersion parameter for binomial family taken to be 1)
##
##      Null deviance: 469.71  on 772  degrees of freedom
## Residual deviance: 447.08  on 771  degrees of freedom
## (227 observations deleted due to missingness)
## AIC: 451.08
##
## Number of Fisher Scoring iterations: 5

```

We learn that treated respondents are more likely to attrite, treatment is positively associated with attrition and is statistically significant.

6.1.3 Bounds

Next, we use the function `bounds()` to get extreme value (Manski) bounds. This function calls the function `estimator_ev` from the `attrition` package by Alex Coppock. `treatment` is the assignment indicator (Z). `DV` is the outcome of interest (Y). Our `bounds()` function removes respondents who attrited pre-treatment and calculates R (the response indicator variable) based on missingness on the DV (missing=0, response=1), based on the assumptions drawn by Manski.

The default for the bounds type is `type = "Manski"`, but we can also specify the type of bounds such that `type = "Lee"` to get Trimming (Lee) bounds. Since we cannot defy the monotonicity assumption, Lee bounds cannot be yielded here, however we demonstrate the use of `type = "Lee"` in the next section.

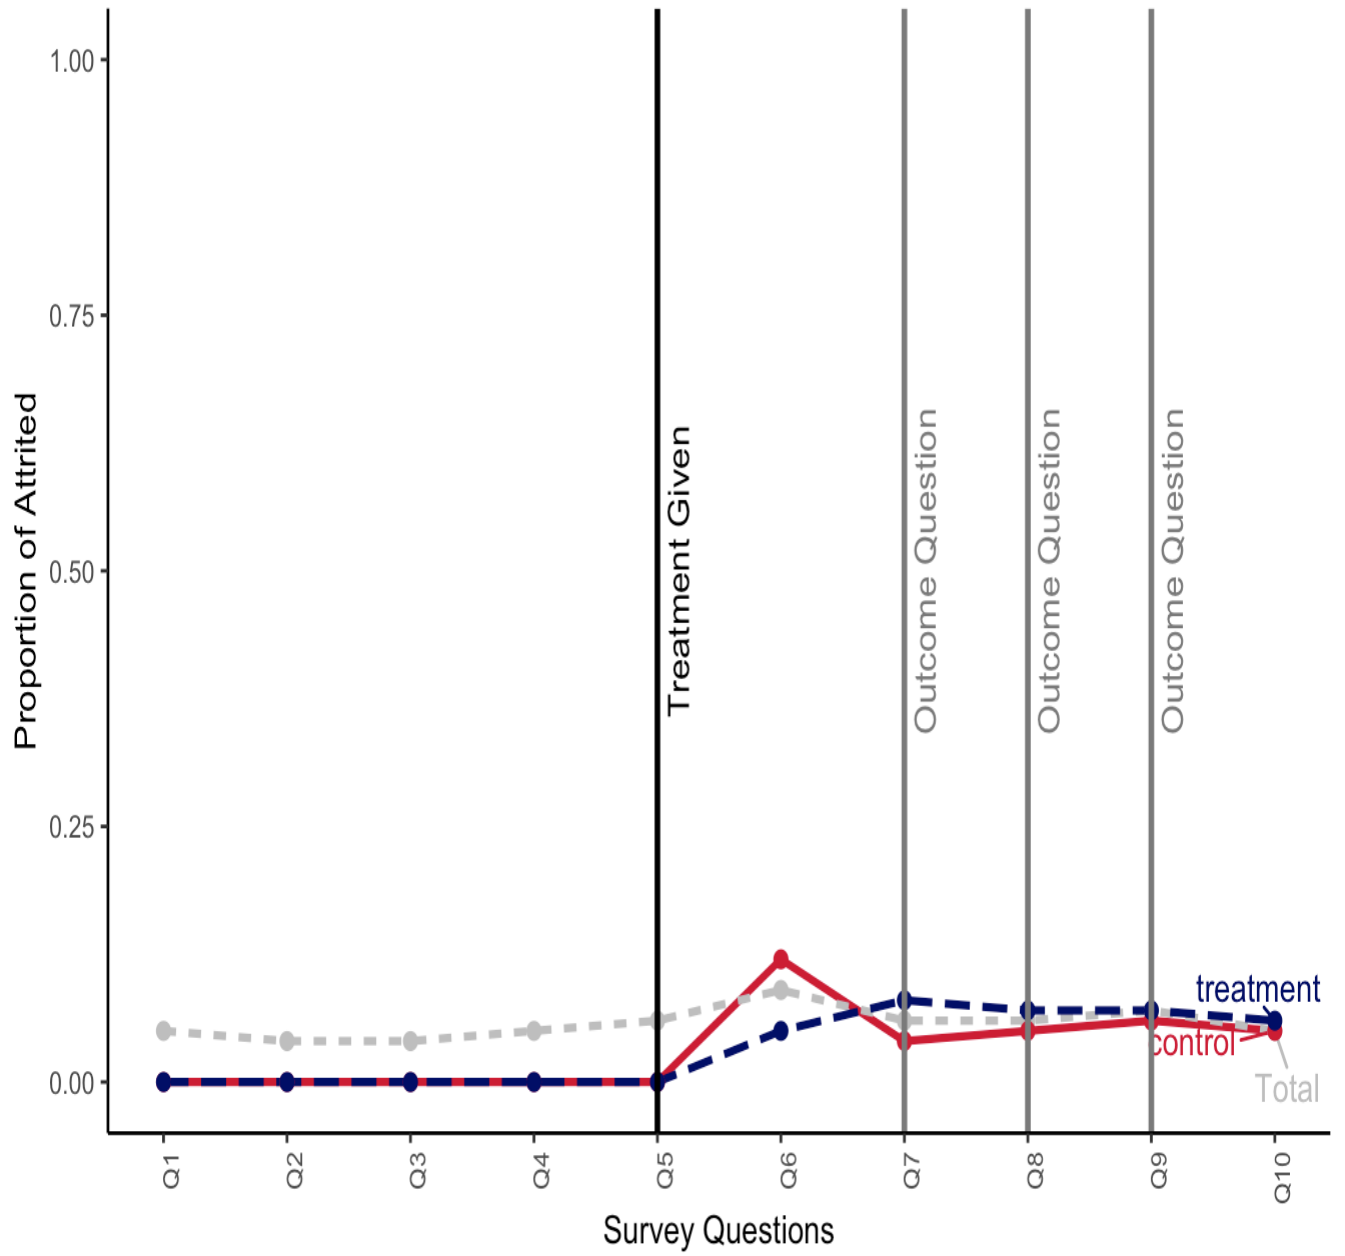
6.2 Control causes attrition

We repeat this process, but instead we look at a case where control causes attrition in `test_sim2`. We might see something like this if positive emotions (like happiness) are ramped up with treatment, making attrition less likely.

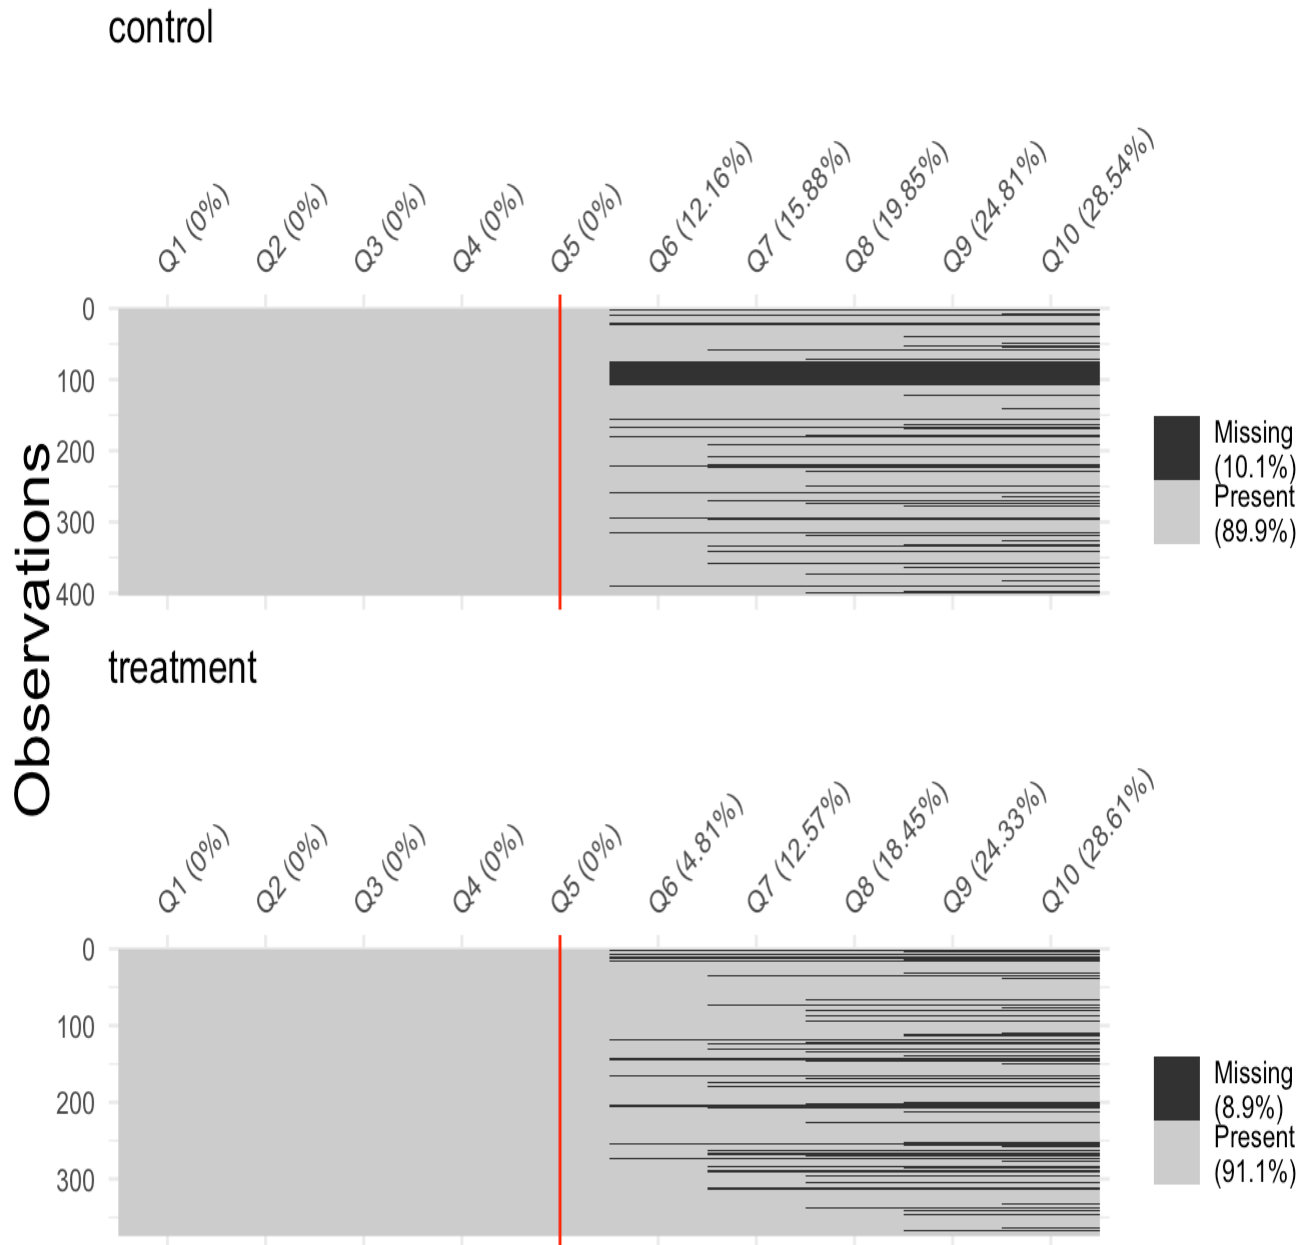
6.2.1 Plot

We visualize attrition using the `plot_attrition()` and `vis_miss_treat()` functions.

```
attritevis::plot_attrition(test_sim2,
  treatment_q = "Q5",
  outcome_q = c("Q7", "Q8", "Q9"),
  freq = FALSE,
  mycolors = c(treatment = "#000066",
               control = "#CC0033")
)
```



```
attritevis::vis_miss_treat(test_sim2, treatment_q = "Q5")
```



We learn that attrition mostly occurs after Q6, and that treated respondents seem to be attriting more.

6.2.2 Balance

We learn that most respondents attrite at the post-treatment question Q6, and conduct a balance test. Note that Q6 is an outcome, and we expect our treatment to affect it. It thus does not make sense to use the `balance_cov()` function. Instead, we want to examine whether our treatment caused attrition, and thus use the function `balance_attrite()`:

```
attritevis::balance_attrite(data = test_sim2,
  treatment = "Q5",
```

```

question = "Q6")

##
## Call:
## glm(formula = question1 ~ treatment1, family = binomial(link = "logit"),
##      data = data2)
##
## Deviance Residuals:
##      Min        1Q    Median        3Q        Max
## -0.5092  -0.5092  -0.3141  -0.3141   2.4633
##
## Coefficients:
##              Estimate Std. Error z value Pr(>|z|)
## (Intercept)      -1.9775     0.1524 -12.974 < 2e-16 ***
## treatment1treatment -1.0071     0.2856  -3.526 0.000422 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## (Dispersion parameter for binomial family taken to be 1)
##
##      Null deviance: 456.45  on 776  degrees of freedom
## Residual deviance: 442.62  on 775  degrees of freedom
## (223 observations deleted due to missingness)
## AIC: 446.62
##
## Number of Fisher Scoring iterations: 5

```

We learn that treated respondents are more likely to attrite, treatment is positively associated with attrition and is statistically significant.

6.2.3 Bounds

Next, we use the function `bounds()` to to get extreme value (Manski) bounds and lee sharp bounds.

```

attritevis::bounds(data = test_sim2,
  treatment = "Q5",
  DV = "Q8")
##      ci_lower      ci_upper      low_est      upp_est      low_var      upp_var
## -0.296727542  0.310514515  0.006893486  0.006893486  0.023997584  0.023997584
##lee sharp

```

```

attritevis::bounds(data = test_sim2,
  treatment = "Q5",
  DV = "Q7",
  type = "Lee")
##      upper_bound      lower_bound      Out0_mono      Out1L_mono      Out1U_mono
##      0.34041298      0.10240873      3.92625369      4.02866242      4.26666667
## control_group_N      treat_group_N      Q      f1      f0
##      339.00000000      327.00000000      0.03790379      0.12566845      0.15880893
##      pi_r_1      pi_r_0
##      0.87433155      0.84119107

```