

# Quantifying the Bias Due to Observed Individual Confounders in Causal Treatment Effect Estimates

---

A Tutorial for the Selection Bias Decomposition (SBdecomp) Package

LAYLA PARAST, BETH ANN GRIFFIN

Sponsored by the National Institutes of Health



SOCIAL AND ECONOMIC WELL-BEING

For more information on this publication, visit [www.rand.org/t/TLA570-3](http://www.rand.org/t/TLA570-3)

Published by the RAND Corporation, Santa Monica, Calif.

© Copyright 2020 RAND Corporation

**RAND**® is a registered trademark.

#### Limited Print and Electronic Distribution Rights

This document and trademark(s) contained herein are protected by law. This representation of RAND intellectual property is provided for noncommercial use only. Unauthorized posting of this publication online is prohibited. Permission is given to duplicate this document for personal use only, as long as it is unaltered and complete. Permission is required from RAND to reproduce, or reuse in another form, any of its research documents for commercial use. For information on reprint and linking permissions, please visit [www.rand.org/pubs/permissions](http://www.rand.org/pubs/permissions).

The RAND Corporation is a research organization that develops solutions to public policy challenges to help make communities throughout the world safer and more secure, healthier and more prosperous. RAND is nonprofit, nonpartisan, and committed to the public interest.

RAND's publications do not necessarily reflect the opinions of its research clients and sponsors.

#### Support RAND

Make a tax-deductible charitable contribution at  
[www.rand.org/giving/contribute](http://www.rand.org/giving/contribute)

[www.rand.org](http://www.rand.org)

# Quantifying the Bias Due to Observed Individual Confounders in Causal Treatment Effect Estimates: A Tutorial for the Selection Bias Decomposition (`SBdecomp`) Package

Layla Parast and Beth Ann Griffin  
RAND Corporation

December 2020

## 1 Introduction

The Selection Bias decomposition (`SBdecomp`) package was designed to allow researchers to quantify the proportion of the selection bias explained by observed individual confounders when obtaining causal treatment effect estimates. The purpose of this tool is to describe the methodology behind the primary function of the `SBdecomp` package, its features and syntax, how to implement this function, and illustrate its use with an example (Parast and Griffin, 2020).

This package will be useful for settings where one would like to identify the most important observed confounder(s) in accounting for the observed selection bias in a naive treatment effect estimate relative to a propensity score weighted adjusted treatment effect estimate. Specifically, when estimating a treatment effect, it is essential to appropriately adjust for selection bias, using, for example, propensity score weighting. Selection bias occurs in an unadjusted treatment effect when individuals who are treated are substantially different from those who are untreated, with respect to covariates that are also associated with the outcome. A comparison of the unadjusted, naive treatment effect estimate with the propensity score adjusted treatment effect estimate provides an estimate of the selection bias caused by observed confounders. We propose two approaches to quantify the proportion of the selection bias explained by each observed confounder—a single confounder removal approach and a single confounder inclusion approach (Parast and Griffin, 2020). At the time of writing, these methods are applicable to settings with a binary treatment and a continuous outcome only.

## 2 Pets Data Example

To demonstrate the package, we use a random subset of the publicly available 2003 California Health Interview Survey (CHIS) data (note that results might differ if the full dataset is used; CHIS, 2005). This truncated dataset is called `petsdata`, and is included in the package; the data consists of responses from 2,102 adults. We assume that the main goal of the study is to investigate the effect of

dog ownership on general health. Dog ownership was assessed with the question “Do you have any dogs that you allow inside your home?”; 29.0 percent of respondents owned a dog. General health status of the individual was measured as the self-reported response to the question “Would you say that in general your health is excellent, very good, good, fair or poor?” Responses were coded from 1 through 5 with 5 indicating “Excellent.” Previous work has shown that such an analysis requires one to account for the differences in characteristics between those who do and do not own a dog (Saunders et al., 2017; Miles et al., 2017; Parast et al., 2017). Therefore, adjustment for such characteristics is necessary to ensure an unbiased estimate of the effect of dog ownership on general health. The following individual characteristics are available in this dataset: age, gender, race/ethnicity, household size, marriage status, whether the individual received TANF (Temporary Assistance for Needy Families), household annual income, whether the individual worked full time, whether the individual had a spouse that worked full time, whether the individual lived in a house, and a rural/urban measure or the individual’s address. Quantifying the extent to which each variable explains the selection bias here is an important step toward understanding what variables are necessary to adjust for in future studies examining the association between pet ownership and health outcomes. This is especially true if one or more of these variables is expensive or burdensome to measure, so that one can ensure unbiased treatment effect estimation while also using resources efficiently.

We begin by loading the package and the data.

```
library(SBdecomp)
data(petsdata)
set.seed(1)
names(petsdata)
```

A seed is set here to ensure reproducible results. The variable **gotdog** contains the dog ownership information (i.e., the treatment/exposure of interest), which takes values 1 for dog ownership and 0 for those who do not own a dog. The variable **genhealth** is the self-reported general health of the individual (i.e., the outcome variable of interest), and takes numeric values 1 through 5. The available characteristics are:

- **age** - age of the individual (numeric)
- **ismale** - indicates the sex of the individual where 1 is male and 0 is female (binary)
- **race\_coll** - race/ethnicity of the individual, collapsed into categories where 1 is Latino, 4 is Asian, 6 is White, and 7 is a combined category of Pacific Islander/Native American/African American/other race/multiple race (categorical/factor)
- **hhsiz** - household size (numeric count)
- **ownhome** - indicates whether the individual owns the home they live in where 1 is yes and 0 is no (binary)
- **married** - indicates whether the individual is married where 1 is yes and 0 is no (binary)
- **ontanf** - indicates whether the individual receives TANF where 1 is yes and 0 is no (binary)

- **hhincome** - household income, log-transformed (numeric)
- **fulltime** - indicates whether the individual works full time where 1 is yes and 0 is no (binary)
- **spouse\_fulltime** - indicates whether the individual has a spouse that works full time where 1 is yes and 0 is no (binary)
- **liveinhouse** - indicates whether the individual lives in a house (as opposed to an apartment, duplex, or mobile home) where 1 is yes and 0 is no (binary)
- **ruralurban** - the rural/urban measure of the individual's home address where 1 = urban, 2 = 2nd city, 3 = suburban, 4 = town and rural (numeric)

Here, we look at a few rows of data.

```
> petsdata[1:5,]
  gotdog age ismale race_coll hhszize ownhome married ontanf hhincome fulltime
1      0  59      1        4       1       1       0       0 11.00212         1
2      0  45      0        6       3       1       1       0 11.69526         1
3      0  24      1        6       3       0       0       0 10.85902         1
4      0  33      0        6       2       0       0       0 10.81980         1
5      0  43      1        6       1       1       0       0 11.44037         1

  spouse_fulltime liveinhouse ruralurban genhealth
1                0           1           2           4
2                1           1           4           5
3                0           0           2           5
4                0           0           3           4
5                0           1           2           5
```

The naive estimate of the treatment effect (i.e., the effect of dog ownership on general health), is simply the average general health among dog owners minus the average general health among those who do not own a dog, calculated to be 0.22. A simple t-test indicates that this difference is significant with  $p < 0.001$ . That is, dog owners appear to have significantly better general health than those who do not own a dog.

```
> group_mean <- aggregate(genhealth ~ gotdog, petsdata, mean)
> group_mean
  gotdog genhealth
1      0  3.462517
2      1  3.685855
> mean(petsdata$genhealth[petsdata$gotdog==1]) - mean(petsdata$genhealth[petsdata$gotdog
+ ==0])
[1] 0.2233385

> t.test(petsdata$genhealth[petsdata$gotdog==1],
+        petsdata$genhealth[petsdata$gotdog==0])
```

Welch Two Sample t-test

```
data: petsdata$genhealth[petsdata$gotdog == 1] and petsdata$genhealth[petsdata$
gotdog == 0]
t = 4.0788, df = 1142.6, p-value = 4.841e-05
alternative hypothesis: true difference in means is not equal to 0
95 percent confidence interval:
 0.1159054 0.3307717
sample estimates:
mean of x mean of y
3.685855  3.462517
```

In the next two sections, we calculate the adjusted treatment effect estimate after accounting for differences between dog owners and those who do not own a dog with respect to the available (observed) characteristics. In addition, we examine the proportion of the estimated selection bias that is explained by each observed individual confounder.

## 3 Single Confounder Removal

### 3.1 Main Arguments

The main arguments for the `sbdecomp` function are the `outcome`, the `treatment`, and the `confounders`. There are two options for inputting these key variables. With the first option, the `outcome` is a numeric vector, the `treatment` is a vector of 0/1 values, and the `confounders` argument is a dataframe of the confounding variables. Other key arguments are the `type` of approach, which can be either “removal” or “inclusion”, and `estimation`, which can be either “nonparametric” or “parametric.”

```
> output = sbdecomp(outcome = petsdata$genhealth, treatment = petsdata$gotdog,
+ confounders = as.data.frame(petsdata[,c(2:13)]), type = "removal", estimation =
+ "parametric", balance = F)
```

Alternatively, the user can provide (1) the names of the variables for the `outcome`, `treatment`, and `confounders` arguments and (2) the dataset using the `data` argument, as illustrated below.

```
> output = sbdecomp(outcome = "genhealth", treatment = "gotdog", confounders =
+ c("age", "ismale", "race_coll", "hhsz", "ownhome", "married", "ontanf", "hhincome",
+ "fulltime", "spouse_fulltime", "liveinhouse", "ruralurban"), data = petsdata,
+ type = "removal", estimation = "parametric", balance = F)
```

It is important that confounders which are factors be coded as factors when supplied as arguments to the `sbdecomp` function. The function will check for factors and will state which variables are identified as factors. For example, when running the code above, the following message is displayed:

```
[1] "The following variable(s) were identified as factors: race_coll"
```

Let's take a look at the output from the code above.

```
>output
$delta.naive
[1] 0.2233385

$p.value.delta.naive
[1] 4.8406e-05

$conf.int.delta.naive
[1] 0.1159054 0.3307717

$delta.fully.adjusted
[1] -0.05602075

$p.value.delta.fully.adjusted
[1] 0.471058

$conf.int.delta.fully.adjusted
[1] -0.2083308 0.0962893

$B
      age      ismale race_coll      hhsz      ownhome      married      ontanf
B 0.3050521 0.02903118  0.35924 0.02629668 0.008339928 0.0006250488 0.01649
      hhincome      fulltime spouse_fulltime liveinhouse ruralurban
B  0.1483924 0.006128138      0.01656658  0.06822543 0.01561248
```

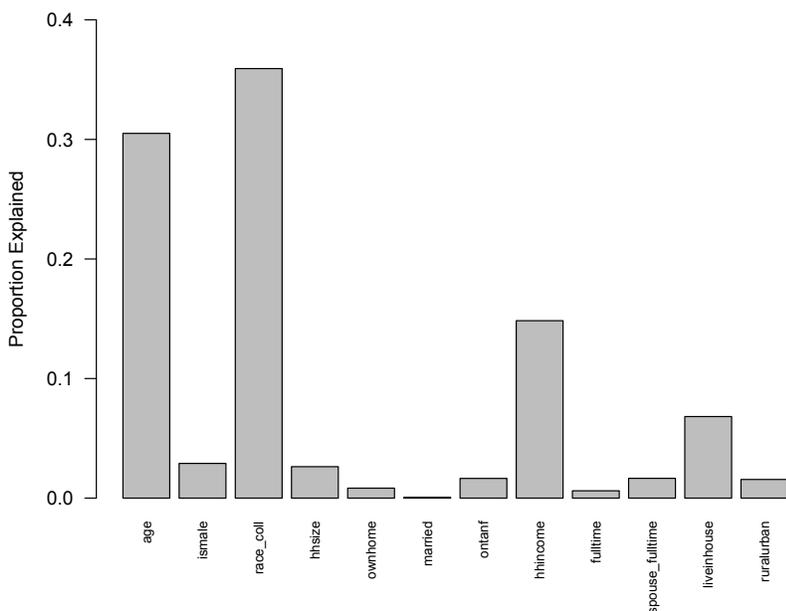
The `delta.naive` quantity is the naive treatment effect (i.e., the simple difference between the general health of individuals with a dog and the general health of individuals without a dog); this value matches the difference in means calculated in the previous section. The quantity `p.value.delta.naive` is the associated p-value and `conf.int.delta.naive` is the associated 95 percent confidence interval. The `delta.fully.adjusted` is the adjusted estimate of the treatment effect accounting for differences in the two ownership groups with the provided confounders. The quantity `p.value.delta.fully.adjusted` is the associated p-value and `conf.int.delta.fully.adjusted` is the associated 95 percent confidence interval. This adjusted estimate uses propensity score weighting where the propensity score weights are estimated using a simple logistic regression, because the `estimation` argument is “parametric” in this illustration. This adjusted estimate is much lower and in the opposite direction of the naive estimate: -0.056 (adjusted) versus 0.223 (naive). In addition, the magnitude of the standardized effect size of the adjusted estimate is substantially lower than that for the naive estimate: 0.04 versus 0.19 (not shown); for more on effect size, please see Cohen (2013). This means that adjusting for these confounders substantially changes the estimated treatment effect, an indication of substantial potential selection bias in the naive, unadjusted treatment effect of dog ownership on general health. The difference in these two estimates, 0.279 (standardized effect size = 0.24), is the estimated selection bias due to the observed covariates used in our propensity score model. This is not unexpected, because previous work has shown that dog owners significantly differ from those who do not own a dog on

several characteristics that are also related to health (e.g., dog-owners are more likely to be older, own a home, have a higher income, and be white; Saunders et al., 2017).

The quantities provided in `B` are the estimated proportions of the observed selection bias explained by each confounder using the single confounder removal approach, because the `type` argument is “removal” in this illustration. For example, differences in age between the two groups account for 30.5 percent of the estimated selection bias. These values can also be examined visually using a bar plot as shown in Figure 1. (Note that if you wish to create more-detailed plots with these results, the needed `B` quantities and confounder names can be extracted using `output$B` and `names(output$B)`.)

```
> bar.sbdecomp(output)
```

These results show that race/ethnicity explains the largest proportion of the estimated selection bias due to the observed confounders used in our propensity score model, followed by age and household income. The least important confounder in this analysis is the variable indicating whether the individual is married or not. These results align with previous work demonstrating significant and large differences between dog owners and those who do not own a dog with respect to race/ethnicity, age, and household income, characteristics known to be associated with health. For example, individuals of certain racial/ethnic minorities are less likely to own a dog and more likely to be in poorer health (Saunders et al., 2017; Bahls, 2011; Baciu et al., 2017). Missing data are handled by analyzing complete cases only throughout this package.



**Figure 1: Proportion of Estimated Selection Bias Explained by Each Confounder Using the Single Confounder Removal Approach and Parametric Estimation**

## 3.2 Additional Quantities

An additional argument in the `sbdecomp` function is `Bonly` which indicates whether the user would like to see only the B estimates or would like to see additional estimates beyond the B estimates. By default, `Bonly` is set to `TRUE` so that only the estimates of B are shown. If `Bonly` is set to `FALSE`, several additional quantities are provided.

```
> sbdecomp(outcome = petsdata$genhealth, treatment = petsdata$gotdog, confounders =  
+ as.data.frame(petsdata[,c(2:13)]), type = "removal", estimation = "parametric",  
+ Bonly = F, balance = F)
```

```
$delta.naive
```

```
[1] 0.2233385
```

```
$p.value.delta.naive
```

```
[1] 4.8406e-05
```

```
$conf.int.delta.naive
```

```
[1] 0.1159054 0.3307717
```

```
$delta.fully.adjusted
```

```
[1] -0.05602075
```

```
$p.value.delta.fully.adjusted
```

```
[1] 0.471058
```

```
$conf.int.delta.fully.adjusted
```

```
[1] -0.2083308 0.0962893
```

```
$B
```

	age	ismale	race_coll	hhsz	ownhome	married	ontanf
B	0.3050521	0.02903118	0.35924	0.02629668	0.008339928	0.0006250488	0.01649
	hhincome	fulltime	spouse_fulltime	liveinhouse	ruralurban		
B	0.1483924	0.006128138	0.01656658	0.06822543	0.01561248		

```
$estimated.selection.bias
```

```
[1] 0.2793593
```

```
$lambda
```

	age	ismale	race_coll	hhsz	ownhome	married
lambda	-0.05813741	-0.005532818	-0.06846466	0.005011672	-0.00158944	-0.000119123
	ontanf	hhincome	fulltime	spouse_fulltime	liveinhouse	ruralurban
lambda	-0.003142696	-0.02828092	0.001167912	-0.003157292	-0.01300253	-0.002975457

```
$delta.each
```

	age	ismale	race_coll	hhsz	ownhome	married
--	-----	--------	-----------	------	---------	---------

```

delta 0.002116667 -0.05048793 0.01244392 -0.06103242 -0.05443131 -0.05590162
      ontanf      hhincome      fulltime spouse_fulltime liveinhouse ruralurban
delta -0.05287805 -0.02773983 -0.05718866      -0.05286345 -0.04301821 -0.05304529

```

The `estimated.selection.bias` is the estimated selection bias, the difference between the naive estimate and the fully adjusted estimate, which we calculated in the previous section. The `lambda` value for each confounder is the difference between the fully adjusted treatment effect estimate and the adjusted treatment effect without adjusting for the particular confounder (i.e., removing that confounder). The `delta` value for each confounder is the resulting adjusted treatment effect when the treatment effect does *not* adjust for the particular confounder. These quantities inform the calculation of B.

Bootstrap calculated standard error estimates are available for B and `lambda` if the user sets the `standard.error` argument to TRUE. The argument `boot.rep` indicates the number of bootstrap replications to be used when estimating the standard error. By default, `standard.error` is FALSE and `boot.rep` is 500. Note that this procedure can be very time intensive (particularly for large datasets and/or when using nonparametric estimation).

```

> set.seed(1)
> sbdecomp(outcome = petsdata$genhealth, treatment = petsdata$gotdog, confounders =
+ as.data.frame(petsdata[,c(2:13)]), type = "removal", estimation = "parametric",
+ Bonly = T, balance = F, standard.error = T )

[1] "The following variable(s) were identified as factors: race_coll"
$delta.naive
[1] 0.2233385

$p.value.delta.naive
[1] 4.8406e-05

$conf.int.delta.naive
[1] 0.1159054 0.3307717

$delta.fully.adjusted
[1] -0.05602075

$p.value.delta.fully.adjusted
[1] 0.471058

$conf.int.delta.fully.adjusted
[1] -0.2083308 0.0962893

$B
      age      ismale race_coll      hhsiz      ownhome      married ontanf hhincome

```

```

B 0.3050521 0.02903118 0.35924 0.02629668 0.008339928 0.0006250488 0.01649 0.1483924
  fulltime spouse_fulltime liveinhouse ruralurban
B 0.006128138 0.01656658 0.06822543 0.01561248

```

```
$B.standard.error
```

```

      age      ismale      race_coll      hhsize      ownhome
0.10421764 0.02694145 0.13762867 0.02555004 0.04075178
  married      ontanf      hhincome      fulltime spouse_fulltime
0.02560001 0.03078032 0.11339237 0.02045496 0.02029934
liveinhouse      ruralurban
0.07391494 0.03399472

```

Such standard error estimates are useful because they reflect the variability of the estimated B and lambda quantities and might inform a user's confidence in making decisions using these quantities. Note that standard error estimates for lambda are only provided if Bonly is TRUE.

### 3.3 Balance Assessment

Another argument is `balance`; by default, `balance` is TRUE. When `balance` is TRUE, several additional quantities related to the balance assessments are provided. To assess balance, the effect size (ES) for the difference between the two groups for each confounder included in the adjustment is calculated. For each confounder, the ES is calculated by taking the absolute standardized difference in the means (weighted means, if weighted) of that confounder between the two groups. For categorical confounders, the ES is calculated for each category of the confounder. Both the mean and the maximum ES difference across all confounders included in the adjustment are calculated. First, `balance.naive.mean` and `balance.naive.max` show the raw (unweighted) mean and max ES across all the variables being considered. Then, `balance.fully.adjusted.mean` and `balance.fully.adjusted.max` show the mean and max ES for the variables when using the fully adjusted propensity score weights.

```

> sbdecomp(outcome = petsdata$genhealth, treatment = petsdata$gotdog, confounders =
+ as.data.frame(petsdata[,c(2:13)]), type = "removal", estimation = "parametric",
+ balance = T)

```

```

$delta.naive
[1] 0.2233385

```

```

$p.value.delta.naive
[1] 4.8406e-05

```

```

$conf.int.delta.naive
[1] 0.1159054 0.3307717

```

```

$delta.fully.adjusted
[1] -0.05602075

```

```
$p.value.delta.fully.adjusted
```

```
[1] 0.471058
```

```
$conf.int.delta.fully.adjusted
```

```
[1] -0.2083308 0.0962893
```

```
$B
```

```
      age      ismale race_coll   hhsiz  ownhome   married  ontanf  
B 0.3050521 0.02903118 0.35924 0.02629668 0.008339928 0.0006250488 0.01649  
      hhincome   fulltime spouse_fulltime liveinhouse ruralurban  
B 0.1483924 0.006128138 0.01656658 0.06822543 0.01561248
```

```
$balance.naive.mean
```

```
[1] 0.2420667
```

```
$balance.naive.max
```

```
[1] 0.525
```

```
$balance.fully.adjusted.mean
```

```
[1] 0.0374
```

```
$balance.fully.adjusted.max
```

```
[1] 0.069
```

```
$balance.mean
```

```
      age      ismale race_coll   hhsiz  ownhome   married  
balance.mean 0.04342857 0.03814286 0.05163636 0.03542857 0.04471429 0.03957143  
      ontanf  hhincome   fulltime spouse_fulltime liveinhouse ruralurban  
balance.mean 0.03928571 0.03728571 0.03828571 0.03735714 0.035 0.03764286
```

```
$balance.max
```

```
      age ismale race_coll hhsiz  ownhome married  ontanf  hhincome  fulltime  
balance.max 0.135 0.063 0.287 0.062 0.118 0.068 0.073 0.071 0.079  
      spouse_fulltime liveinhouse ruralurban  
balance.max 0.069 0.14 0.073
```

In the pets study data, prior to adjustment, the group of individuals who owned pets were notably different from those who did not with `balance.naive.mean` = 0.242 and `balance.naive.max` = 0.525. After weighting by the fully adjusted propensity score, we see the groups are highly similar with `balance.fully.adjusted.mean` = 0.0374 and `balance.fully.adjusted.max` = 0.069. A threshold of < 0.1 is generally used to indicate acceptable balance, meaning that we have good balance in this

example when using the fully adjusted estimate.

When using the single confounder removal approach, balance is also assessed for each adjusted estimate (i.e., each estimate that results from removing a particular confounder from adjustment). Results are provided for each variable in `balance.mean` and `balance.max`. For example, when age is removed from adjustment, the mean ES difference across the remaining confounders is 0.043 while the max is 0.135, indicating relatively good, though borderline balance. In contrast, when race/ethnicity is removed from the adjustment, the maximum ES difference is 0.287 indicating poor balance.

### 3.4 Nonparametric Estimation

When the `estimation` argument is “nonparametric,” a generalized boosted model (GBM) is used to estimate the fully adjusted treatment effect, and the adjusted treatment effects when each confounder is removed individually (McCaffrey et al., 2004; Ridgeway et al., 2014). If this option is used, the `sbdecomp` function will accept additional arguments that will be passed to the `ps` function from the `twang` package, specifically `n.trees`, `interaction.depth`, `shrinkage`, `verbose`, `stop.method`, and `cv.folds`. (The `twang` package is the Toolkit for Weighting and Analysis of Nonequivalent Groups and provides functions for propensity score estimating and weighting, nonresponse weighting, and diagnosis of the weights.) The default values for these arguments are 20,000, 4, 0.005, FALSE, `es.max`, and 0, respectively. For more information on appropriate values for these arguments, see the help file and tutorial for `twang` (Ridgeway et al., 2014). For small datasets in particular, users might prefer to set the `interaction.depth` to 3 rather than use the default of 4.

Here is the code for using the nonparametric estimation option using the default settings.

```
> output = sbdecomp(outcome = petsdata$genhealth, treatment = petsdata$gotdog,  
+ confounders = as.data.frame(petsdata[,c(2:13)]), type = "removal", estimation =  
+ "nonparametric", balance = T)
```

```
> output
```

```
$delta.naive
```

```
[1] 0.2233385
```

```
$p.value.delta.naive
```

```
[1] 4.8406e-05
```

```
$conf.int.delta.naive
```

```
[1] 0.1159054 0.3307717
```

```
$delta.fully.adjusted
```

```
[1] 0.01930506
```

```
$p.value.delta.fully.adjusted
```

```
[1] 0.7614851
```

```
$conf.int.delta.fully.adjusted
```

```
[1] -0.1053355 0.1439457
```

```
$B
```

```
      age      ismale  race_coll   hhsz   ownhome   married  ontanf  hhincome
B 0.1687302 0.009022499 0.08181766 0.02063108 0.06165633 0.002939019      0 0.3926911
      fulltime spouse_fulltime liveinhouse ruralurban
B 0.0002918589      0.001902147      0.206686 0.0536321
```

```
$balance.naive.mean
```

```
[1] 0.2420667
```

```
$balance.naive.max
```

```
[1] 0.525
```

```
$balance.fully.adjusted.mean
```

```
[1] 0.05773333
```

```
$balance.fully.adjusted.max
```

```
[1] 0.14
```

```
$balance.mean
```

```
      age      ismale  race_coll   hhsz   ownhome  married
balance.mean 0.06107143 0.05821429 0.04636364 0.06328571 0.05564286      0.06
      ontanf  hhincome  fulltime spouse_fulltime liveinhouse ruralurban
balance.mean 0.06092857 0.05478571 0.05971429      0.05885714      0.056 0.05285714
```

```
$balance.max
```

```
      age  ismale  race_coll  hhsz  ownhome  married  ontanf  hhincome  fulltime
balance.max 0.122 0.138      0.127 0.135 0.131 0.141 0.14 0.115 0.14
      spouse_fulltime liveinhouse ruralurban
balance.max      0.141      0.133      0.129
```

Using nonparametric estimation, based on the output and as shown in Figure 2, the fully adjusted treatment effect estimate is 0.019 and household income is identified as explaining the highest proportion of the estimated selection bias, followed by the variable indicating whether the individual lives in a house. It is notable that the inferences from the nonparametric and parametric approaches differ for the single confounder removal case. As noted in the output, balance is not well achieved for the nonparametric models with max ES ranging from 0.115 to 0.14 for each case. For this case study data, the parametric approach proves to have better balancing properties and thus would be the preferred approach. Every application will be different and the balance statistics provided here are useful metrics for determining which approach is optimal.

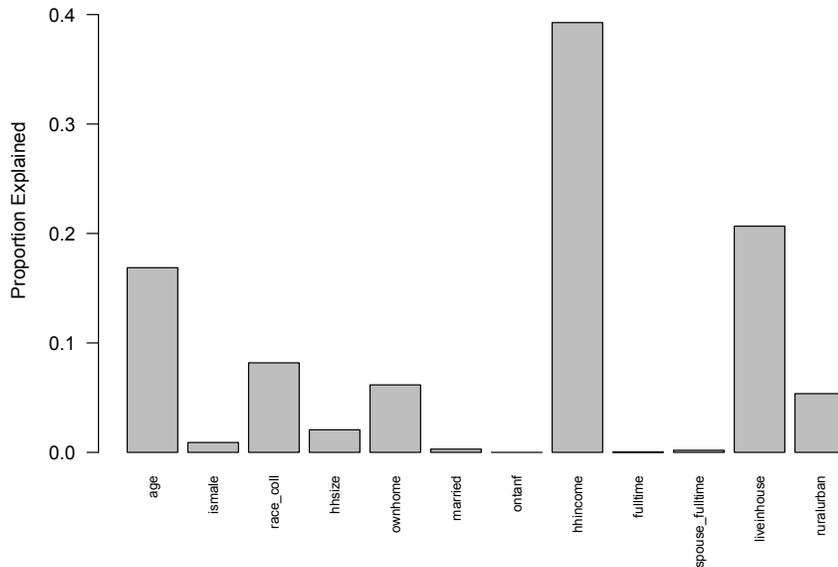


Figure 2: Proportion of Estimated Selection Bias Explained by Each Confounder Using the Single Confounder Removal Approach and Nonparametric Estimation

## 4 Single Confounder Inclusion

### 4.1 Parametric Estimation

To estimate the proportion of the estimated selection bias explained by each confounder using the single confounder inclusion approach, the `type` argument should be identified as “inclusion.” As with the single confounder removal approach, users have the option of parametric or nonparametric estimation, and can request that additional quantities such as the `lambda` values and `delta` values be included in the output. We first examine results using parametric estimation below.

```
> output = sbdecomp(outcome = petsdata$genhealth, treatment = petsdata$gotdog,
+ confounders = as.data.frame(petsdata[,c(2:13)]), type = "inclusion", estimation =
+ "parametric", balance = T)
```

```
> output
$delta.naive
[1] 0.2233385
```

```
$p.value.delta.naive
[1] 4.8406e-05
```

```
$conf.int.delta.naive
[1] 0.1159054 0.3307717
```

```
$delta.fully.adjusted
```

```
[1] -0.05602075
```

```
$p.value.delta.fully.adjusted
```

```
[1] 0.471058
```

```
$conf.int.delta.fully.adjusted
```

```
[1] -0.2083308 0.0962893
```

```
$B
```

```
      age      ismale race_coll      hhsiz  ownhome  married  
B 0.04398239 0.001172353 0.1105127 0.005517636 0.1259806 0.02103899  
      ontanf  hhincome  fulltime spouse_fulltime liveinhouse  ruralurban  
B 0.01518807 0.3389299 0.07206396      0.06791753 0.1890645 0.008631345
```

```
$balance.naive.mean
```

```
[1] 0.2420667
```

```
$balance.naive.max
```

```
[1] 0.525
```

```
$balance.fully.adjusted.mean
```

```
[1] 0.0374
```

```
$balance.fully.adjusted.max
```

```
[1] 0.069
```

```
$balance.mean
```

```
      age  ismale  race_coll  hhsiz  ownhome  married  ontanf  hhincome  fulltime  
balance.mean 0.009      0      0  0.01      0      0      0      1.679      0  
      spouse_fulltime liveinhouse ruralurban  
balance.mean      0      0      0.013
```

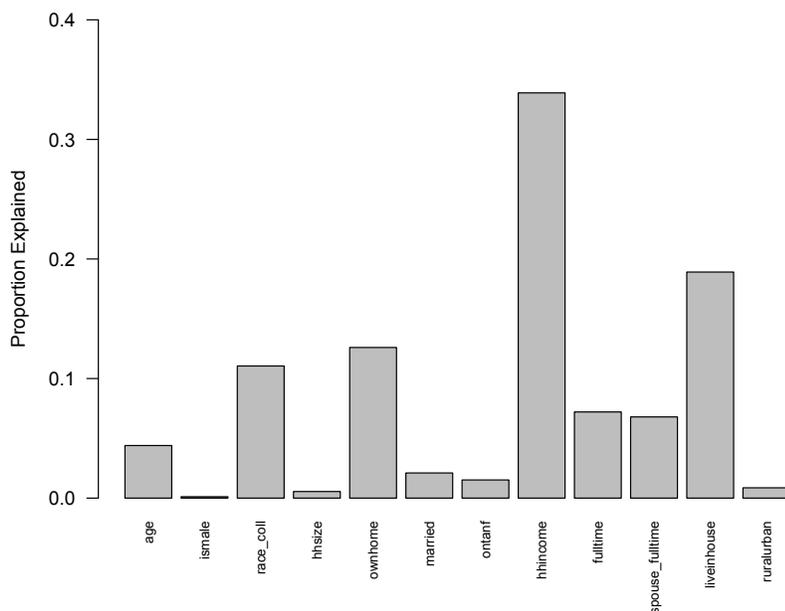
```
$balance.max
```

```
      age  ismale  race_coll  hhsiz  ownhome  married  ontanf  hhincome  fulltime  
balance.max 0.009      0      0  0.01      0      0      0      1.679      0  
      spouse_fulltime liveinhouse ruralurban  
balance.max      0      0      0.013
```

```
> bar.sbdecomp(output)
```

The `delta.naive` and `delta.fully.adjusted` estimates are the same as when the removal option is used; the B values are different. As in the previous section, these values can also be examined visually using a bar plot (shown in Figure 3). These results show that using this approach, household income

explains the largest proportion of the estimated selection bias, followed by the variable indicating whether the individual lives in a house.



**Figure 3: Proportion of Estimated Selection Bias Explained by Each Confounder Using the Single Confounder Inclusion Approach and Parametric Estimation**

## 4.2 Balance Assessment

When using the single confounder inclusion approach, balance is assessed for only the variable that gets included in the propensity score model (e.g., the single included variable). Because there is only a single confounder in the adjustment model at a time, the mean and maximum are with respect to only a single variable and are thus the same (with the exception of categorical variables). Results are provided for each variable in `balance.mean` and `balance.max` (these will only be potentially different for categorical variables, where the mean or max across the different levels of the variable is calculated). For all variables except for `hhincome`, good balance with respect to the difference in the variable between the two groups can be achieved. Given the poor balance obtained for household income, the findings from the parametric approach for the single confounder inclusion method should be taken with serious caution. Next we will see if nonparametric estimation can provide better balance.

## 4.3 Nonparametric Estimation

Results using the single confounder inclusion approach with nonparametric estimation are shown below and in Figure 4, and show that household income again explains the highest proportion (39.5 percent) of the estimated selection bias, followed by the variable indicating whether the individual lives in a house (17.0 percent), similar to when parametric estimation is used. Balance looks much better here,

with all ES quantities less than 0.10. Therefore, findings from the nonparametric approach should be considered as optimal for the single confounder inclusion approach for our case study data. Again, every application will behave differently and it is important to assess balance to determine whether results from the parametric or nonparametric approach are more robust.

```
> output = sbdecomp(outcome = petsdata$genhealth, treatment = petsdata$gotdog,
+ confounders = as.data.frame(petsdata[,c(2:13)]), type = "inclusion", estimation =
+ "nonparametric", balance = T)

> bar.sbdecomp(output)

> output
$delta.naive
[1] 0.2233385

$p.value.delta.naive
[1] 4.8406e-05

$conf.int.delta.naive
[1] 0.1159054 0.3307717

$delta.fully.adjusted
[1] 0.01930506

$p.value.delta.fully.adjusted
[1] 0.7614851

$conf.int.delta.fully.adjusted
[1] -0.1053355 0.1439457

$B
      age      ismale  race_coll    hhsiz  ownhome  married  ontanf  hhincome
B 0.03538716 0.001056745 0.09961484 0.02419814 0.1135574 0.0189643 0.01369035 0.3953232
  fulltime spouse_fulltime liveinhouse  ruralurban
B 0.0649576      0.06122005   0.1704205 0.001609669

$balance.naive.mean
[1] 0.2420667

$balance.naive.max
[1] 0.525

$balance.fully.adjusted.mean
[1] 0.05773333
```

```
$balance.fully.adjusted.max
[1] 0.14
```

```
$balance.mean
      age ismale race_coll hhsiz e ownhome married ontanf hhincome fulltime
balance.mean 0.016      0      0 0.005      0      0      0      0.023      0
      spouse_fulltime liveinhouse ruralurban
balance.mean      0      0      0
```

```
$balance.max
      age ismale race_coll hhsiz e ownhome married ontanf hhincome fulltime
balance.max 0.016      0      0 0.005      0      0      0      0.023      0
      spouse_fulltime liveinhouse ruralurban
balance.max      0      0      0
```

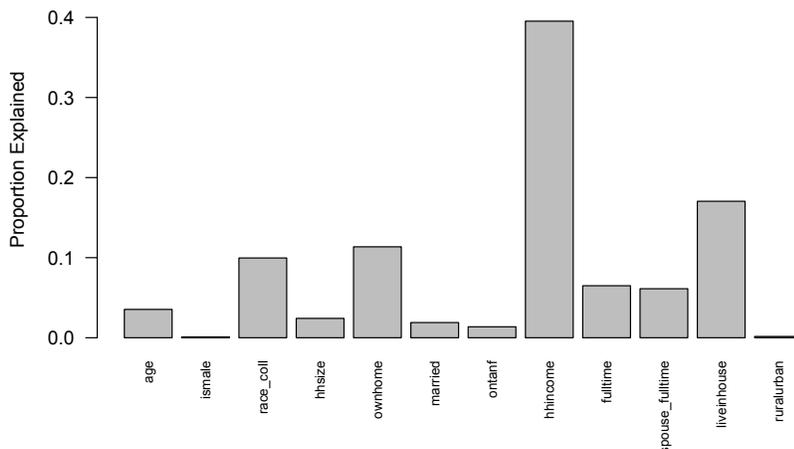


Figure 4: Proportion of Estimated Selection Bias Explained by Each Confounder Using the Single Confounder Inclusion Approach and Nonparametric Estimation

## 5 Conclusion

To our knowledge, prior to this work, there were no readily available methods to decompose the selection bias and quantify the proportion of the estimated selection bias explained by each confounder when estimating causal effects using propensity score weights. We hope that users will find this tool helpful in analyzing data when there is a substantive interest in identifying the variable or variables that explains the largest proportion of the estimated selection bias.

## About This Tutorial

This tutorial and the `SBdecomp` package were supported by funding from grant 1R01DA034065 from the National Institute on Drug Abuse. The overarching goal of the grant is to develop statistical methods and tools that will provide addiction health services researchers and others with the tools and training they need to study the effectiveness of treatments using observational data. The work is an extension of the Toolkit for Weighting and Analysis of Nonequivalent Groups, or TWANG, which contains a set of functions to support causal modeling of observational data through the estimation and evaluation of propensity score weights. The TWANG package was first developed in 2004 by RAND researchers for the R statistical computing language and environment and has since been expanded to include tools for SAS, Stata, and Shiny. For more information about TWANG and other causal tools being developed, see [www.rand.org/statistics/twang](http://www.rand.org/statistics/twang).

RAND Social and Economic Well-Being is a division of the RAND Corporation that seeks to actively improve the health and social and economic well-being of populations and communities throughout the world. This research was conducted in the Social and Behavioral Policy Program within RAND Social and Economic Well-Being. The program focuses on such topics as risk factors and prevention programs, social safety net programs and other social supports, poverty, aging, disability, child and youth health and well-being, and quality of life, as well as other policy concerns that are influenced by social and behavioral actions and systems that affect well-being. For more information, email [sbp@rand.org](mailto:sbp@rand.org).

## Acknowledgments

We thank Nelson Lim at the RAND Corporation for helpful discussions on this topic and the UCLA Center for Health Policy Research for providing the California Health Interview Survey data. Note that this tutorial uses a random subset of 2003 survey data for illustration; results might differ if the full dataset is used. In addition, we thank the beta testers for this tutorial and package: Irineo Cabrerros, Maria DeYoreo, and Joseph Pane.

## References

- Baciu, A., Negussie, Y., Geller, A., and Weinstein, J. (2017). *Communities in Action: Pathways to Health Equity*. Washington, D.C.: National Academies Press.
- Bahls, C. (2011). *Health Policy Brief: Achieving Equity in Health*. Health Affairs, Robert Wood Johnson Foundation.
- CHIS (2005). California Health Interview Survey (CHIS) 2003 Methodology Series. Los Angeles: UCLA Center for Health Policy Research.
- Cohen, J. (2013). *Statistical Power Analysis for the Behavioral Sciences, 2nd ed.* New York: Routledge.

- McCaffrey, D. F., Ridgeway, G., and Morral, A. R. (2004). Propensity Score Estimation with Boosted Regression for Evaluating Causal Effects in Observational Studies. *Psychological Methods* **9(4)**, 403–425.
- Miles, J. N., Parast, L., Babey, S. H., Griffin, B. A., and Saunders, J. M. (2017). A Propensity-Score-Weighted Population-Based Study of the Health Benefits of Dogs and Cats for Children. *Anthrozoös* **30(3)**, 429–440.
- Parast, L. and Griffin, B. A. (2020). Quantifying the Bias Due to Individual Confounders in Causal Treatment Effect Estimates. *Statistics in Medicine* **39(18)**, 2447–2476.
- Parast, L., McCaffrey, D. F., Burgette, L. F., de la Guardia, F. H., Golinelli, D., Miles, J. N., and Griffin, B. A. (2017). Optimizing Variance-Bias Trade-Off in the TWANG Package for Estimation of Propensity Scores. *Health Services and Outcomes Research Methodology* **17**, 175–197.
- Ridgeway, G., McCaffrey, D., Morral, A., Burgette, L., and Griffin, B. A. (2014). *Toolkit for Weighting and Analysis of Nonequivalent Groups: A Tutorial for the TWANG Package*. Santa Monica, C.A.: RAND Corporation, TL-136/1-NIDA.
- Saunders, J., Parast, L., Babey, S. H., and Miles, J. V. (2017). Exploring the Differences Between Pet and Non-Pet Owners: Implications for Human-Animal Interaction Research and Policy. *PloS One* **12(6)**, e0179494.