

Direct and indirect effects in a logit model

Maarten L. Buis
Department of Sociology
Tübingen University
Tübingen, Germany
maarten.buis@uni-tuebingen.de

Abstract. This article discusses a method by Erikson et al. (2005) for decomposing a total effect in a logit model into direct and indirect effects. Moreover, this article extends this method in three ways. First, in the original method the variable through which the indirect effect occurs is assumed to be normally distributed. In this article the method is generalized by allowing this variable to have any distribution. Second, the original method did not provide standard errors for the estimates. In this article the bootstrap is proposed as a method of providing those. Third, I show how to include control variables in this decomposition, which was not allowed in the original method. The original method and these extensions are implemented in the `ldecomp` package.

Keywords: st0001, ldecomp, mediation, intervening variable, logit

The aim of this article is to study direct, indirect and total effects in a logit model, and in particular a generalization of a method by Erikson et al. (2005) for computing those when the variable whose effect we want to decompose is a categorical variable. Direct, indirect and total effects are studied in order to reveal a mechanism through which one variable affects another variable. The idea is illustrated using figure 1: There is a variable X which has an effect on a variable Y , but part of this effect occurs through another variable Z . This indirect effect occurs because X influences Z which in turn influences Y . Within this framework, the effect of X on Y while controlling for Z is called the direct effect. The indirect and direct effect together form the total effect of X on Y . The indirect effect is the part of the effect of X on Y that can be explained, while the direct effect is the residual or unexplained part of the effect. The aim of such an analysis is to try to explain why a variable X influences a variable Y , by specifying a mechanism: the effect occurs through the third variable Z . Z is called an intervening variable. It differs from a confounding variable in the direction of the effect between X and Z . If Z was a confounding variable then it would affect X rather than the other way round.

X , Y , and Z could be many things. For example, within political science there is literature starting with (Campbell et al. 1960) which explains the effect of party identification (X) on voting for a party (Y) in part through how one perceives the issues and the parties (Z). The idea is that someone's party identification is a relatively stable characteristic of a person, almost like a personality trait, which can color/distort the perceptions of the issues and the positions of the candidates or parties¹, which

1. A striking example of this is reported by Bartels (2002), who found that in a survey held in 1988

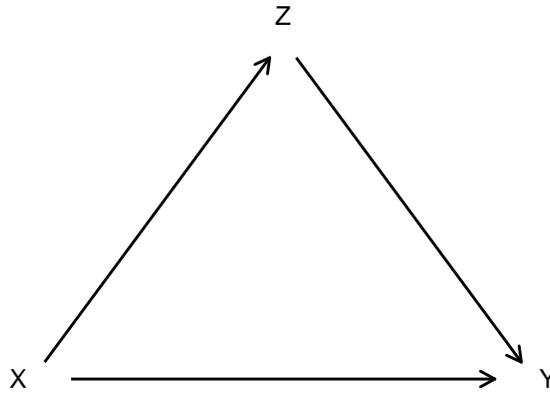


Figure 1: Direct, indirect, and total effects

in turn influences voting behavior. This indirect effect thus represents a mechanism through which party identification influences voting behavior, and we want to know how important that mechanism is relative to the total effect of party identification.

Alternatively, within sociology there is a literature starting with (Boudon 1974) that looks at the degree to which children from different social backgrounds (X) have different probabilities of attending college (Y), and the part this effect that can be ascribed to an indirect effect through academic performance in high school (Z): Children from a higher class background do well at high school and those children that do well at high school are more likely to attend college. Possible explanations for this could be that higher class parents have better genes (nature) or are able to provide a more intellectually stimulating environment for their children (nurture) or both. The aim is again to estimate how important this mechanism is relative to the total effect of family background.

More generally, the common practise of comparing regression coefficients before and after controlling for a set of other covariates is implicitly an attempt at decomposing a total effect into a direct and an indirect effect. Within linear regression ([R] **regress**) this is an easy and convenient way of estimating the direct, indirect and total effect. In this case the total effect of X is estimated using a regression which does not control for Z. The direct effect of X is the effect of X in a regression with both X and Z as explanatory variables, and the indirect effect is the difference between these two, since

more than 50% of the respondents who identified themselves as a "strong Democrat" thought that the inflation got worse or a lot worse during the Reagan presidency, while only 13% of the "strong Republicans" thought that this was the case. In actual fact the inflation rate in consumer prices fell from 13.5% in 1980 to 4.1% in 1988.

the total effect is the sum of the direct and indirect effect. However, this method will not work when dealing with non-linear models like logistic regression ([R] **logit**). In this article I will show why this is the case, and I will discuss a method for solving this problem proposed by Erikson et al. (2005), I will propose a generalization of this method, and introduce the `ldecomp` package that implements these methods in Stata.

1 The problem of indirect effects in a logit model

The key problem when estimating the direct, indirect, and total effects is that the standard method of estimating them — comparing estimates from models that do and do not control for Z — will not work in non-linear models like logistic regression. The easiest way to see that is in an example where there is no indirect effect.

To illustrate the problem, I create artificial data where I know that there cannot be an indirect effect of X through Z on Y , and show that the naïve method finds a substantial indirect effect. The variable Z is created so that it can take three values (0, 1, and 2), and the variable X is created so that it can take two values (0 and 1). There is no relationship between X and Z (a low X individual are as likely to be a high value on Z as a high X individual). So in this example, there is no indirect effect of X on Y through Z . We start by creating the variables X and Z .

```
. drop _all
. set obs 60000
obs was 0, now 60000
. gen z = ceil(_n / 20000) - 1
. bys z: gen x = ceil(_n / 10000) - 1
. tab x z
```

x	z			Total
	0	1	2	
0	10,000	10,000	10,000	30,000
1	10,000	10,000	10,000	30,000
Total	20,000	20,000	20,000	60,000

Next, we create the dependent variable Y according to the logistic regression equation $\frac{P(y=1)}{1-P(y=1)} = -4 + 4X + 2Z$ as discussed in (Buis 2007b)

```
. set seed 12345
. gen y = runiform() < invlogit(-4 + 4*x + 2*z)
```

Next we compute the naïve estimates of the indirect effect

(Continued on next page)

```

. qui logit y x z
. est store direct
. local direct = _b[x]
.
. qui logit y x
. est store total
. local total _b[x]
.
. est tab direct total

```

Variable	direct	total
x	4.0391332	2.6256242
z	2.026339	
_cons	-4.0452305	-1.3123133

```

. di "naive indirect effect = " `total' - `direct'
naive indirect effect = -1.413509

```

Someone who did not know that the data were created such that the indirect effect is exactly zero, and used the naïve estimate of the indirect effect would conclude that the indirect effect is about 54% of the total effect, and the fact that it has the opposite sign from the total effect would suggest that this indirect effect has a noticeable dampening influence on the effect of X on Y. This is not due to sampling variation, but due to a structural bias in the naïve estimator.

The reason for this bias is that a logistic regression is comparison of proportions that have first been transformed into log-odds ratios². This is illustrated in Figure 2. It consists of 4 vertical lines, on the two outer lines are plotted the probabilities of success on Y, while the equivalent log odds are plotted on the two inner lines. The two left lines represent the high X category, while the two right lines represent the low X group. The open symbols and the solid arrows show how the probabilities are transformed into log odds, and how, within each category of Z, the log odds of the high X group is compared with the log odds of the low X group. This represents what happens when we estimate a `logit` model with both X and Z as explanatory variables. When we leave the variable Z out of the model — for example, because we want to estimate the total effect — we are in effect first computing the average of the proportions and then transform them into log odds. This is represented by the closed circles and the dashed arrows. However, the more extreme values are less extreme in the probability metric than in the log odds metric, that is, the probabilities close to either 0 or 1 are more bunched together than their log odds counterparts. Therefore, computing the average proportion before transforming the proportions into log odds means that the extreme values are less influential than they would have been if the means were computed in

2. The odds and the log odds contain exactly the same information as the proportion, it is just presented differently. The proportions times a hundred tells you how many people out of a hundred are expected to attend college, while the odds tells you how many people are expected to attend college for every person that does not attend college. The odds (O) can be derived from the proportion (p) in the following way: $O = \frac{p}{1-p}$.

the log odds metric, so the average is being pulled towards the less extreme categories. As a consequence the effect in terms of log odds will be less when Z is left out of the model, even if there is no indirect effect. This problem is very similar to the problems discussed in (Bartus 2005) and (Buis 2007a).

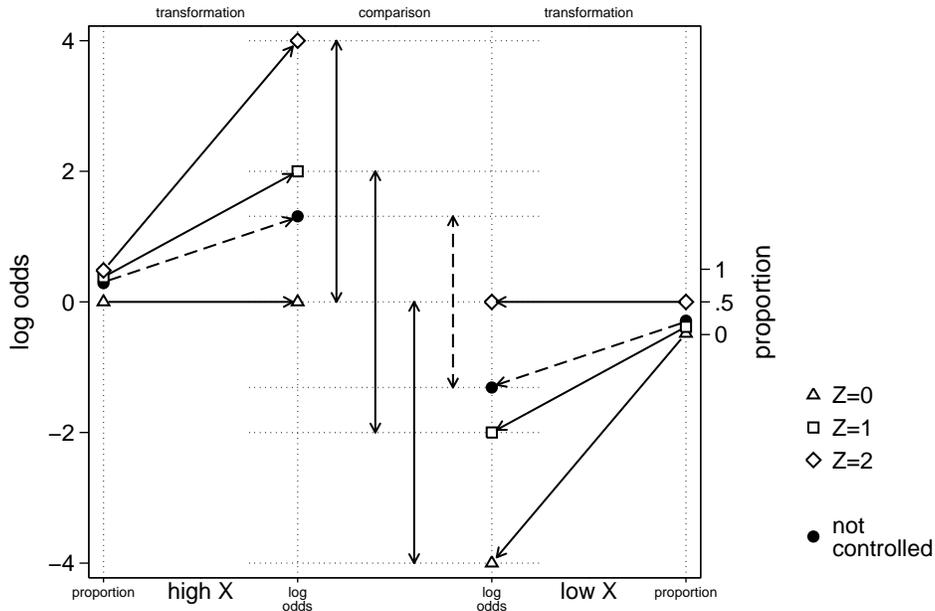


Figure 2: The effect of X with and without controlling for Z

2 A solution

2.1 Outline

In order for there to be an indirect effect of X on Y through Z, X needs to have an effect on Z, and Z has to have an effect on Y. For example, children from higher classes (X) are likely to do better at high school (Z) than children from lower classes, and those who have done better in high school (Z) are more likely to enter college (Y). In Figure 2 there was no indirect effect because the distribution of Z was the same for both categories of X. In figure 3 this is changed to include an indirect effect. The distribution of Z is represented by the size of the symbols. So, in this new example there are within the high X group more high Z individuals than medium Z individuals, and more medium Z individuals than low Z individuals. The distribution of Z for the low X group is exactly the opposite. Now there are two reasons why high X individuals are more likely to belong to the high Y group: 1) they are more likely to belong to the high Z group and

those who belong to the high Z group are more likely to belong to the high Y group. This is the indirect effect of X through Z on Y. Figure 3 shows this in the following way: For the high X group the high Z group is more influential when computing the average probability because they are more numerous, while for the low X group the low Z group are more influential, and this leads to a larger difference between the high and low status group. 2) they are more likely to belong to the high Y group even when compared with a low X group that has the same value on Z. Figure 3 shows this in the following way: For each level of Z the high X group has a higher proportion belonging to the high Y group than the low X group. This is the direct effect.

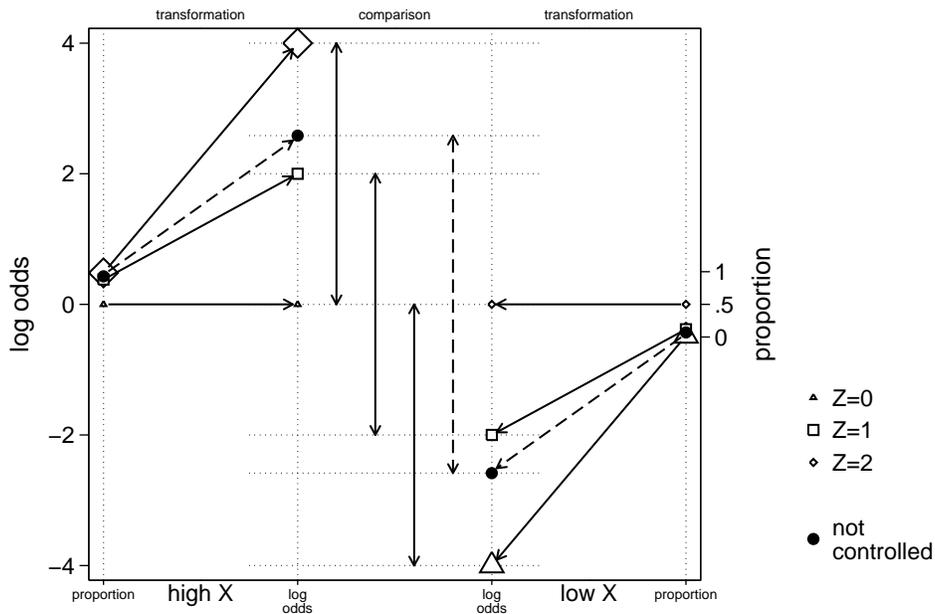


Figure 3: Indirect and direct effects

Erikson et al. (2005) propose two related methods for estimating a direct, indirect, and total effect such that the direct and indirect effects add up to the total effect. The first method, is illustrated in figure 4, and consists of the following steps:

1. Estimate a logistic regression using both X and Z, and optionally the interactions between X and Z.
2. Predict for each individual the log odds of success on Y (this is the linear predictor), and transform these to probabilities.
3. Compute within each group of X the average predicted probability and transform these to log odds. The difference in these log odds between the groups of X

represents the total effect.

4. Assume that the low X individuals are actually high X individuals, predict for each low X individual the log odds in this counterfactual scenario, and transform these to probabilities.
5. Compute the average of these probabilities. This is the counterfactual probability of success on Y for high X individuals if they had the distribution of Z of the low X individuals. These are then transformed into log odds.
6. The high X individuals and this counterfactual group differ with respect to the distribution of Z, but the probabilities conditional on X and Z are kept constant. Therefore, comparing these groups gives the effect of X due to the differences in the distribution of Z, that is, the indirect effect.
7. The low X individuals and the counterfactual group differ with respect to the probabilities conditional on X and Z, but the distribution of Z is kept constant. Therefore, comparing these groups gives the effect of X while controlling for the distribution of Z, that is, the direct effect.

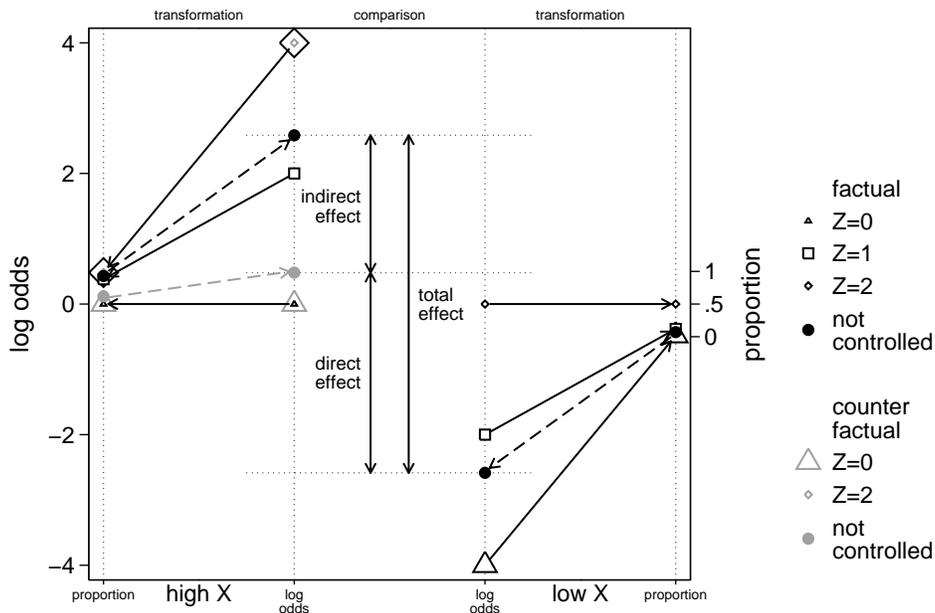


Figure 4: Method 1

Figure 4 and equation 1 both show that this way the total effect is always the sum of the direct and indirect effect. In equation 1 the O is the odds of success on Y, the

first subscript represents the logistic regression coefficients, and the second subscript represents the distribution of Z^3 . So $O_{x=1,z|x=1}$ is the odds of success for the high X group, while $O_{x=1,z|x=0}$ is the counterfactual odds of success for the high X group if it had the distribution of Z of the low X group.

$$\underbrace{\ln(O_{x=1,z|x=1}) - \ln(O_{x=0,z|x=0})}_{\text{total}} = \underbrace{\ln(O_{x=0,z|x=1}) - \ln(O_{x=0,z|x=0})}_{\text{indirect}} + \underbrace{\ln(O_{x=1,z|x=1}) - \ln(O_{x=0,z|x=1})}_{\text{direct}} \quad (1)$$

Using the rule that $\ln(a) - \ln(b) = \ln(a/b)$ it can be shown that these effects are actually log odds ratios:

$$\ln \left(\underbrace{\frac{O_{x=1,z|x=1}}{O_{x=0,z|x=0}}}_{\text{total}} \right) = \ln \left(\underbrace{\frac{O_{x=0,z|x=1}}{O_{x=0,z|x=0}}}_{\text{indirect}} \right) + \ln \left(\underbrace{\frac{O_{x=1,z|x=1}}{O_{x=0,z|x=1}}}_{\text{direct}} \right) \quad (2)$$

This means that this decomposition can also be presented in terms of odds ratios, by exponentiating both sides of equation 2. Because of the rule that $\exp(a + b) = \exp(a) \times \exp(b)$ the total effect is now the product of the direct and indirect effects:

$$\underbrace{\frac{O_{x=1,z|x=1}}{O_{x=0,z|x=0}}}_{\text{total}} = \underbrace{\frac{O_{x=0,z|x=1}}{O_{x=0,z|x=0}}}_{\text{indirect}} \times \underbrace{\frac{O_{x=1,z|x=1}}{O_{x=0,z|x=1}}}_{\text{direct}} \quad (3)$$

The second method of decomposing the total effect into a direct and indirect effect proposed by (Erikson et al. 2005) is exactly the same as the first method, except that it uses the counterfactual probability of success on Y for the low X individuals assuming that they have the distribution of Z of the high X individuals, as is illustrated in figure 5. The logic behind these two methods is exactly the same, but they do not have to result in exactly the same estimates for the direct and indirect effects, though they are often very close. Jackson et al. (2007) propose to solve this problem by computing the size of the indirect effect relative to the total effect using both methods and report the average of the two.

Another difficulty with this method is that the variable whose effect we want to decompose (X) must be categorical, and that the number of effects that must be decomposed increases very quickly with the number of categories in X. The reason is that the decomposition will be different for all pairwise comparisons of categories. Normally

3. In fact, as can be seen in the description of this technique, these are averaged estimated log odds, so some people would prefer to incorporate that in the notation by adding hats and bars. However, I think that this notation is complicated enough as is and it is not crucial for the points that these equations make, so I chose not to include those.

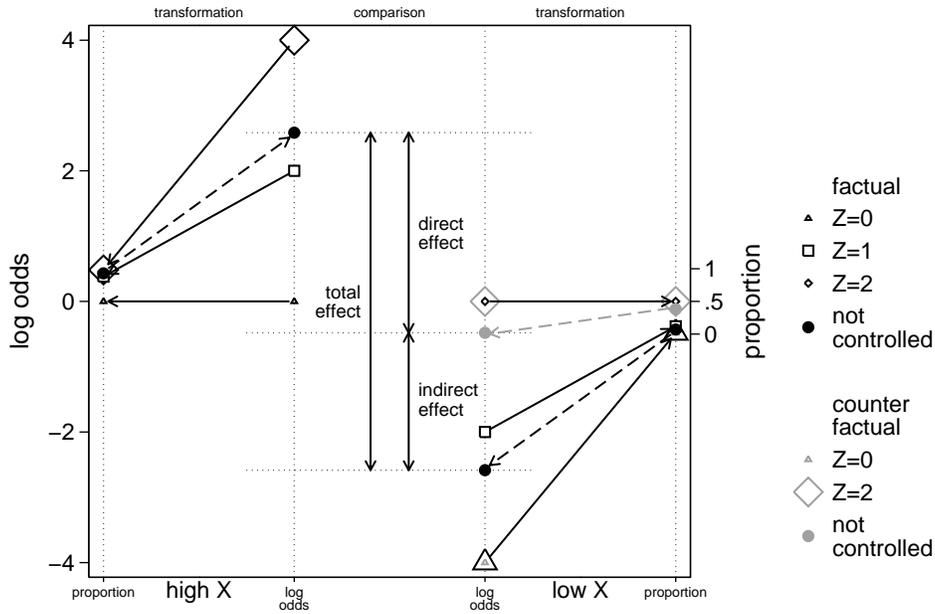


Figure 5: Method 2

a categorical explanatory variable consisting of C categories can be entered as $C - 1$ dummies, each dummy representing a comparison between the reference category and one of the other categories, leading to $C - 1$ effects. All other pairwise comparisons can be directly derived from those effects. Consider, a logit model with a categorical variable with three categories, low ($x = 0$), middle ($x = 1$), and high ($x = 2$), and the category low was chosen to be the reference category. This will result in two odds ratios: the odds ratio comparing the the medium and the low group and the odds ratio comparing the high and the low group: $O_{x=1}/O_{x=0}$ and $O_{x=2}/O_{x=0}$. This excludes the third possible comparison: the medium versus the high group. This comparison can be derived directly from the other two odds ratios by dividing the two odds ratios:

$$\frac{O_{x=1}/O_{x=0}}{O_{x=2}/O_{x=0}} = \frac{O_{x=1}}{O_{x=0}} \times \frac{O_{x=0}}{O_{x=2}} = \frac{O_{x=1}}{O_{x=2}} \quad (4)$$

This same reasoning will however not work for the decomposition. Consider the three indirect effects using the first method, $O_{x=0,z|x=1}/O_{x=0,z|x=0}$, $O_{x=0,z|x=2}/O_{x=0,z|x=0}$, and $O_{x=1,z|x=2}/O_{x=1,z|x=1}$:

$$\frac{O_{x=0,z|x=2}/O_{x=0,z|x=0}}{O_{x=0,z|x=1}/O_{x=0,z|x=0}} \neq \frac{O_{x=1,z|x=2}}{O_{x=1,z|x=1}} \quad (5)$$

So with this decomposition we cannot get away with only presenting all comparisons with a single reference category (leading to $C - 1$ effects), but we will have to display the decomposition for all pairwise comparisons (leading to $\binom{C}{2}$ effects).

2.2 Generalization

In the original formulation of this method, the variable through which the indirect effect occurs (Z) has to be normally distributed. This is because the way Erikson et al. (2005) propose to compute the average predicted probabilities. As discussed in steps 3 and 5, these averaged predicted probabilities play a key role in this decomposition. Erikson et al. (2005) compute these averages by assuming that Z follows a normal distribution, and integrate over this distribution. This method is generalized in this article in order to allow Z to follow any distribution by computing the average predicted probability by using the arithmetic mean of the predicted and counterfactual probabilities, this way integrating over the empirical distribution of Z instead of over a normal distribution. As an added bonus this method is also easier to compute as the integration over the normal distribution has to be done numerically, because there is no closed form solution for this integral. For these reasons the generalized method is the default method in the `ldecomp` package, which implements both decompositions in Stata.

2.3 Standard errors

In order to get an idea about the degree of uncertainty around these estimates one would usually also estimate standard errors. These are not provided by Erikson et al. (2005) and Jackson et al. (2007), but can be easily computed using the bootstrap (Efron and Tibshirani 1993). The bootstrap uses the idea that the standard error is the result of the following thought experiment: Assume we could draw many samples from the population and compute a statistic in each of these samples. As these samples will slightly differ from one another, so will the estimates of the statistic. The standard error is the standard deviation of these different estimates and indicates how much variation one can expect due to the fact that the estimate is based on a random sample. Drawing many random samples from the population is not practical, but we do have a good estimate of the population, the sample, and we can without any difficulty draw (with replacement) many samples from this 'estimate of the population'. So, when using the bootstrap many samples are drawn with replacement from the observed sample, the statistic is computed within each sample, and the estimate of the standard error is the standard deviation of these statistics. Within Stata this process has been automated with the `bootstrap` command, see: [R] `bootstrap`.

2.4 Control variables

Control variables can be included in this method by including these by estimating the logit model from step 1 including these control variables, but afterwards, at steps 2 and 4, fix their value at one specific value, for example their mean.

3 Implementation in Stata

Both the original method proposed by Erikson et al. (2005), and the generalization proposed in this paper have been implemented in Stata as the `ldcomp` package. Furthermore, `ldcomp` by default uses the bootstrap to compute standard errors and confidence intervals.

3.1 Syntax

```
ldcomp depvar [control_var1] [...] [if] [in] [weight] , direct(varname)
      indirect(varlist) [ at(control_var1 #[; control_var2 #] [...]) obspr
      predpr predodds or rindirect normal range(##) nip(#) interactions
      nolegend nodecomp nobootstrap bootstrap_options ]
```

`fweights`, `pweights`, and `iwweights` are allowed when the `nobootstrap` option is specified.

3.2 Options

`direct`(*varname*) specifies the variable whose effect we want to decompose into a direct, indirect and total effect. This has to be a categorical variable, each value of *varname* is assumed to represent a group.

`indirect`(*varlist*) specifies the variable(s) through which the indirect effect occurs. By default, multiple variables are allowed and these can follow any distribution. If the `normal` option is specified only one variable can be entered, and this variable is assumed to be normally distributed.

`at`(*control_var1* #[; *control_var2* #] [...]) specifies the values at which the control variables are to be fixed. The default is to fix the control variables at their mean value.

`rindirect` specifies that the size of the indirect effects relative to the total effect are to be displayed.

`or` specifies that the decomposition is to be displayed in terms of odds-ratios instead of log odds-ratios.

`obspr` specifies that a table of the observed proportions is to be displayed.

`predpr` specifies that a table of predicted and counterfactual proportions is to be displayed. If the `normal` option is not specified and there are no control variables then the diagonal elements of this table will be exactly the same as the observed proportions.

`predodds` specifies that a table of predicted and counterfactual odds is to be displayed.

normal specifies that the predicted and counterfactual proportions are to be computed according to the method specified by Erikson et al. (2005). This means that the variable in `indirect()` is assumed to be normally distributed. By default `ldecomp` uses the more flexible method described in this article.

range(##) specifies the range over which the numerical integration of variable specified in `indirect()` is to be performed. The default is the minimum of that variable minus 10% of the range of the variable and the maximum of the variable plus 10% of the range of the variable. This option can only be specified with the **normal** option.

nip(##) specifies the number of integration points used in the numerical integration of the variable specified in `indirect()`. The default 1000. This option can only be specified with the **normal** option.

interactions specifies that interactions between the categories of the variable specified in `direct()` and the variable(s) specified in `indirect()` are to be included. In other words the effects of the variables specified in `indirect()` on the dependent variable are allowed to differ from one another for each category of the variable specified in `direct()`. This option was primarily added for compatibility with Erikson et al. (2005).

nolegend suppresses the legend that is by default displayed at the bottom of the main table.

nodecomp prevents `ldecomp` from displaying the table of decompositions, which can be useful in combination with the `obspr`, `predpr`, and/or `predodds` options.

nobootstrap prevents `ldecomp` from using `bootstrap` to calculate standard errors.

bootstrap_options The following options of `bootstrap` are allowed: `reps`, `strata`, `size`, `cluster`, `idcluster`, `saving`, `bca`, `mse`, `nodots`, `seed`, and `jackknifeopts`. See: [R] `bootstrap`.

3.3 Example

The use of `ldecomp` is illustrated using data from the Wisconsin Longitudinal Study (Hauser and Sewell 1957–1977), which contains data on a random sample of 10,317 men and women who graduated from Wisconsin high schools in 1957. In this example we want to study the part of the effect of social class on the probability of entering college that can be explained by performance during high school. The dependent variable is whether a respondent ever attended college (`college`). Class is measured by the type of occupation of the father (`ocf57`). The original data contained five classes, which would lead to $\binom{5}{2} = 10$ effects that are to be decomposed. To keep the number of effects manageable, the number of classes has been reduced to three: a lower class (unskilled workers and farmers), a middle class (skilled workers), and a higher class (white collar workers, professionals, and executives). The performance during high school is measured with the percentile rank of high school grades (`hsrankq`). This means that computing the counterfactual proportions using the method proposed by Erikson et al. (2005) will

be problematic, as percentile rank scores will follow a uniform distribution instead of a normal distribution. However, the default method can be used without problem as it does not make any assumption about the distribution of performance. Relevant descriptive statistics are shown below:

```
. use wisconsin.dta, clear
(Wisconsin Longitudinal Study)
. recode ocf57 2=1 3=2 4=3 5=3
(ocf57: 6196 changes made)
. label define ocf57 1 "lower" 2 "middle" 3 "higher", modify
. label value ocf57 ocf57
.
. table ocf57 if !missing(hsrankq, college) , ///
> contents(mean college mean hsrankq freq)    ///
> format(%9.3g) stubwidth(15)
```

occupation of r father in 1957	mean(college)	mean(hsrankq)	Freq.
lower	.284	48.2	5218
middle	.38	50.6	868
higher	.619	56.2	2837

There are big differences between the classes in the proportion of respondents that attend college. Moreover, those classes with low proportion attending college also tend to have done worse during high school. So, part of the differences in proportions attending college could be due to differences in performance. `ldecomp` is used to estimate these direct, indirect, and total effects. In the example below, the effects are presented as odds ratios, so the total effect is the product of the indirect and direct effect. Consider the decomposition according to the first method of the difference between high class (3) and low class (1) students, that is, the first three coefficients of the panel labeled "3/1". Overall the odds of attending college for high class students is 4.10 times as large as that odds for low class students (the total effect). Low class students would have a 1.23 times higher odds of attending college if they had the same performance as high class students (indirect effect according to method 1), while the high class students would have a 3.34 times higher odds of attending college than low class students when we keep the performance constant at the level of high class students (direct effect according to method 1).

(Continued on next page)

of the indirect effect relative to the size of the total effect can help avoid confusion.

In addition to the decomposition itself, `ldecomp` can produce a number of tables that together illustrate step by step how it builds this decomposition. The first of these tables is the table of the predicted and counterfactual proportions shown below.

```
. ldecomp college , direct(ocf57) indirect(hsrankq) predpr nodecomp
predicted and counterfactual proportions
```

distribution	association		
	lower	middle	higher
lower	.284	.366	.571
middle	.296	.38	.585
higher	.327	.415	.619

The diagonal elements in this table represent the predicted proportions, both factual distribution of performance (the rows) and factual conditional probabilities (the columns), while the off-diagonal elements represent the counterfactual proportions. For example, 28.4% of the children from lower class fathers will attend college. If these children had the same performance as the children of higher class fathers then 32.7% of them would attend college. If they had the same conditional probabilities as the children of higher class fathers 57.1% would attend college. This indicates that the direct effect is stronger than the indirect effect. The next step in the computation is to transform these proportions into odds, which is done in the next table.

```
. ldecomp college , direct(ocf57) indirect(hsrankq) predodds nodecomp
predicted and counterfactual odds
```

distribution	association		
	lower	middle	higher
lower	.396	.578	1.33
middle	.421	.613	1.41
higher	.487	.71	1.62

The proportions are transformed into odds by dividing the proportion by one minus the proportion, so the odds of attending college for children from lower class fathers is $\frac{.284}{1-.284} = .397$. The difference between this number and the number in the table is due to rounding. This odds is interpreted as for every child of a lower class father that does not go to college there are only .397 children of lower class fathers that do go to college.

The results that were obtained at the beginning of this example can be computed using the predicted and counterfactual odds from the previous table. For example, if we return to the decomposition of the total effect of having a higher class father rather than a father from the lower class using method 1, we can compute it by filling in equation 3 with the predicted and counterfactual odds from the previous table:

$$\underbrace{\frac{1.62}{.396}}_{total} = \underbrace{\frac{.487}{.396}}_{indirect} \times \underbrace{\frac{1.62}{.487}}_{direct}$$

$$\underbrace{4.09}_{total} = \underbrace{1.23}_{indirect} \times \underbrace{3.33}_{direct}$$

This example also makes it possible to look if the naïve method is really that bad as I claim, and if the generalization that I proposed in this article is really an improvement. We know that `hsrankq` deviates considerable from a normal/Gaussian distribution⁴, so if the generalized method is an improvement on the original method then in this example it should yield noticeably different estimates. The estimates using the original method by Erikson et al. (2005) are obtained by specifying the `normal` option in `ldecomp`. The naïve estimate was computed as follows: first a `logit` model of `college` on `ocf57` was estimated. The effect of `ocf57` is the naïve estimate of the total effect. Then a `logit` model of `college` on `ocf57` and `hsrankq` was estimated. The effect of `ocf57` is the naïve estimate of the direct effect. The naïve estimate of the indirect effect relative to the total effect is (total effect - direct effect) / total effect. The results are shown in Table 1. This table clearly illustrates the major under-estimation of the indirect effect when the naïve method is used. Moreover, the method by Erikson et al. (2005) also leads to a considerable underestimation of about a quarter of the effect obtained using the generalized method. This underestimation is the result of the incorrect assumption made by the method by Erikson et al. (2005) that `hsrankq` is normally distributed. Since the generalized method makes no such assumption, it seems to be the safest method of the three for computing the decomposition of total effects into direct and indirect effects after a logistic regression model.

Table 1: Comparing different estimates of the size of the (average) indirect effect relative to the size of the total effect

	generalization	(Erikson et al. 2005)	naïve
middle v. low	0.135	0.110	0.009
high v. low	0.144	0.105	0.014
high v. middle	0.149	0.102	0.017

4 Summary and discussion

This article started by showing why getting estimates of the direct and indirect effects in a logistic regression is hard and presented a method by Erikson et al. (2005) to estimate these effects, and proposes a generalization of this method. The idea is that a categorical variable `X` can have a direct effect on a variable `Y`, and an indirect effect through `Z`. One can find the direct effect of `X` by comparing the log odds of successes in

4. As noted before, it is a percentile rank score, so it follows a uniform distribution

one category of X with the counterfactual log odds of successes of another category of X given that they have the distribution of Z of the first category. This way the factual and counterfactual group only differ with respect to X and not with respect to the distribution of Z, thus controlling for the distribution of Z. Similarly, the indirect effect can be found by comparing the log odds of success within one category of X with the counterfactual log odds of success within that same category of X with a distribution of Z of another category of X. This way the factual and counterfactual groups differ only with respect to the distribution of Z. In its original form this method assumed that the variable through which the indirect effect occurs has a normal distribution. In this article this method was generalized by allowing the variable to have any distribution. Moreover, the use of the bootstrap is proposed in order to estimate standard errors and I added the possibility to include control variables.

This is a relatively new methodology, and like any new methodology there are still some loose ends to tie up. First, the fact that there are two different estimates of the direct and two different estimates of the indirect effects is less than elegant. Second, the fact that a separate decomposition needs to be estimated for all pairwise comparisons class instead of all comparisons with a single reference category can quickly lead to a very large number of estimates. Third, this method is not the only way of attaining estimates of direct and indirect effects. There are for instance the methods proposed by Gomulka and Stern (1990), Even and Macpherson (1990), Fairlie (2005), Yun (2004), and Bauer and Sinning (2008). Two of these have been implemented in Stata: The method by Fairlie (2005) in the `fairlie` package (Jann 2006), and the method by Bauer and Sinning (2008) in the `nldecompose` package (Sinning et al. 2008). How these alternatives compare to the method discussed in this article needs to be explored.

5 Acknowledgments

I thank the members of the Comparative Stratification Research Seminar at the Vrije Universiteit Amsterdam, Michelle Jackson, William Gould, and an anonymous referee for their helpful comments.

6 References

- Bartels, L. M. 2002. Beyond the Running Tally: Partisan Bias in Political Perceptions. *Political Behavior* 24: 117–150.
- Bartus, T. 2005. Estimation of Marginal Effects using `margeff`. *Stata Journal* 5: 309–329.
- Bauer, T., and M. Sinning. 2008. An extension of the Blinder-Oaxaca decomposition to nonlinear models. *Advances in Statistical Analysis* 92: 197–206.
- Boudon, R. 1974. *Education, Opportunity and Social Inequality*. New York: Wiley.
- Buis, M. L. 2007a. Predict and adjust with logistic regression. *Stata Journal* 7: 221–226.

- . 2007b. Stata tip 48: Discrete uses for uniform(). *Stata Journal* 7: 434–435.
- Campbell, A., P. E. Converse, W. E. Miller, and D. E. Stokes. 1960. *The American Voter*. New York: Wiley.
- Efron, B., and R. J. Tibshirani. 1993. *An Introduction to the Bootstrap*. Boca Raton: Chapman & Hall / CRC.
- Erikson, R., J. H. Goldthorpe, M. Jackson, M. Yaish, and D. R. Cox. 2005. On class differentials in educational attainment. *Proceedings of the National Academy of Science* 102(27): 9730–9733.
- Even, W., and D. Macpherson. 1990. Plant size and the decline of unionism. *Economic Letters* 32: 393–398.
- Fairlie, R. W. 2005. An extension of the Blinder-Oaxaca decomposition technique to logit and probit models. *Journal of Economic and Social Measurement* 30: 305–316.
- Gomulka, J., and N. Stern. 1990. The employment of married women in the United Kingdom 1970–83. *Economica* 57: 171–199.
- Hauser, R. M., and W. H. Sewell. 1957–1977. *Wisconsin Longitudinal Study (WLS) [graduates and siblings]: 1957–1977*. Madison, WI: University of Wisconsin-Madison. <http://www.ssc.wisc.edu/~wls/documentation/>.
- Jackson, M., R. Erikson, J. H. Goldthorpe, and M. Yaish. 2007. Primary and secondary effects in class differentials in educational attainment: The transition to A-level courses in England and Wales. *Acta Sociologica* 50(3): 211–229.
- Jann, B. 2006. FAIRLIE: Stata module to generate nonlinear decomposition of binary outcome differentials. Available from <http://ideas.repec.org/c/boc/bocode/s456727.html>.
- Sinning, M., M. Hahn, and T. Bauer. 2008. The Blinder-Oaxaca decomposition for nonlinear regression models. *The Stata Journal* 8: 480–492.
- Yun, M.-S. 2004. Decomposing differences in the first moment. *Economic Letters* 82: 275–280.

About the author

Maarten L. Buis is affiliated with the Department of Sociology of the University Tübingen, where he studies the interaction between demographic processes and inequality of educational opportunity. He is also a frequent contributor to *statalist*.