

Marginal odds ratios

What they are, how to compute them,
and why social scientists might want to use them

Ben Jann and Kristian Bernt Karlson

Alan Walters Lecture series
Department of Economics, University of Birmingham
April 26, 2023

Outline

- 1 Background
- 2 Marginal (log) odds ratios
- 3 Two illustrations
- 4 Estimation
- 5 Applied example
- 6 Discussion

Acknowledgements: This presentation is based on work which is part of a project by Kristian Bernt Karlson that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 851293).

Background

- Odds ratios form the backbone of much quantitative research in social sciences and epidemiology.
- Close to a hallmark of these disciplines!
- But: Falling out of favor!
 - ▶ Magnitude of odds ratios depends on unmeasured covariates orthogonal to the predictor of interest.
 - ▶ Noncollapsibility (rescaling bias).
 - ▶ Invalidates cross-model and subgroup coefficient comparisons.

Background

- Solutions?
- KHB for cross-model comparisons (Karlson et al. 2012)
- Compare sign not magnitude
- Average marginal effects based on nonlinear probability model (AME)
- Linear probability models (LPM)

Background

- At least in social sciences, AMEs are now the standard. Some even suggest just using LPM, as the difference between results from LPM and the AMEs are typically small.

- ▶ For illustration, here is a quick analysis of papers published in the *American Sociological Review*:

	2010–2015	2016–2021
“odds ratio”	32	9
“marginal effect”	2	11
“linear probability model”	3	16

- But this might be throwing out the baby with the bathwater, because ...

... magnitudes depend on the margin

... AMEs focus on absolute probability differences, not relative differences, which are key to much theory and research.

Background

- What we suggest:

Use **marginal (log) odds ratios**, which . . .

. . . behave like AME but retain the (relative) odds ratio interpretation!

- ✓ unaffected by noncollapsibility
- ✓ an average effect (population-averaged)
- ✓ comparable across populations/studies

- 1 Background
- 2 Marginal (log) odds ratios
- 3 Two illustrations
- 4 Estimation
- 5 Applied example
- 6 Discussion

Marginal odds ratio

- Following Zhang (2008) and Daniel et al. (2021) we use potential outcomes notation to define the marginal odds ratio.
- Y_t : Potential outcome that would realize if treatment T was set to level t by manipulation (i.e., without changing anything else).
- We focus on *binary* outcomes only, that is, $Y_t \in \{0, 1\}$ (failure or success).
- Thus:

$\Pr(Y_t = 1) = E[Y_t]$ is the (marginal) probability that Y_t will be equal to 1 (probability of success).

Marginal odds ratio

- Consider a binary treatment $T \in \{0, 1\}$.
- The marginal odds ratio (MOR) of the alternative treatment ($T = 1$) versus the standard treatment ($T = 0$) is defined as

$$\text{MOR} = \frac{\text{odds}(\Pr[Y_1 = 1])}{\text{odds}(\Pr[Y_0 = 1])}$$

where $\text{odds}(p)$ stands for $p/(1 - p)$.

- Interpretation of MOR: The ratio of the odds of success if everyone would receive the alternative treatment versus the odds of success if everyone would receive the standard treatment (assuming that there are no general equilibrium effects, i.e., SUTVA holds).

“Marginal” refers to how a predictor affects the “marginal distribution” of an outcome (i.e., not to a marginal change in a predictor). “Unconditional” would be another term but we use “marginal” because the term is established in the literature (Stampf et al. 2010; Karlson, Popham, and Holm 2021).

Adjusting for covariates

- The probability of success may not only depend on T , but also on other factors \mathbf{X} .
- Assume that \mathbf{X} has a specific distribution in the population and let $\Pr(Y_t = 1|\mathbf{X} = \mathbf{x}) = E[Y_t|\mathbf{X} = \mathbf{x}]$ be the conditional success probability given $\mathbf{X} = \mathbf{x}$.
- By the law of iterated expectations,

$$\Pr(Y_t = 1) = E_{\mathbf{X}}[\Pr(Y_t = 1|\mathbf{X} = \mathbf{x})]$$

where $E_{\mathbf{X}}$ is the expectation over the distribution of \mathbf{X} .

Adjusting for covariates

- The marginal odds ratio, adjusting for \mathbf{X} , can then be written as

$$\text{MOR} = \frac{\text{odds}\{E_{\mathbf{X}}[\text{Pr}(Y_1 = 1|\mathbf{X} = \mathbf{x})]\}}{\text{odds}\{E_{\mathbf{X}}[\text{Pr}(Y_0 = 1|\mathbf{X} = \mathbf{x})]\}}$$

- We term this the *adjusted* MOR.
- Note:
 - ▶ The adjusted MOR is the same as the unadjusted MOR by definition (i.e., same estimand)!
 - ▶ However, estimation based on the adjusted MOR formulation can be used to address confounding bias in observational data. It can also be used to increase efficiency in analysis of randomized experiments.
 - ▶ Close relationship to AME, which is defined as

$$\text{AME} = E_{\mathbf{X}}[\text{Pr}(Y_1 = 1|\mathbf{X} = \mathbf{x})] - E_{\mathbf{X}}[\text{Pr}(Y_0 = 1|\mathbf{X} = \mathbf{x})]$$

Continuous treatment

- In case of a continuous treatment, the MOR may depend on the level of the treatment (i.e., MOR may not be constant).
- We define the level-specific marginal log odds ratio as the derivative of the marginal log odds by the treatment:

$$\begin{aligned}\ln \text{MOR}(t) &= \lim_{\epsilon \rightarrow 0} \frac{\ln \text{odds}[\Pr(Y_{t+\epsilon} = 1)] - \ln \text{odds}[\Pr(Y_t = 1)]}{\epsilon} \\ &= \lim_{\epsilon \rightarrow 0} \frac{\ln \text{odds}\{E_{\mathbf{X}}[\Pr(Y_{t+\epsilon} = 1 | \mathbf{X} = \mathbf{x})]\} - \ln \text{odds}\{E_{\mathbf{X}}[\Pr(Y_t = 1 | \mathbf{X} = \mathbf{x})]\}}{\epsilon}\end{aligned}$$

- We can then obtain the average MOR by integrating over the distribution of T :

$$\overline{\text{MOR}} = \exp\{E_T[\ln \text{MOR}(t)]\}$$

Continuous treatment

- Another possibility is to integrate over T when obtaining the population-averaged probabilities, that is,

$$\begin{aligned}\ln \text{MOR}' &= \lim_{\epsilon \rightarrow 0} \frac{\ln \text{odds}\{E_T[\Pr(Y_{t+\epsilon} = 1)]\} - \ln \text{odds}\{E_T[\Pr(Y_t = 1)]\}}{\epsilon} \\ &= \lim_{\epsilon \rightarrow 0} \frac{\ln \text{odds}\{E_{T,\mathbf{X}}[\Pr(Y_{t+\epsilon} | \mathbf{X} = \mathbf{x})]\} - \ln \text{odds}\{E_{T,\mathbf{X}}[\Pr(Y_t | \mathbf{X} = \mathbf{x})]\}}{\epsilon}\end{aligned}$$

- This corresponds to the marginal odds ratio that is obtained if treatment is slightly increased for each population member, given the member's existing values for T and \mathbf{X} .

Relationship to the logistic model

- Consider a simple logistic model

$$\Pr(Y_t = 1) = \text{logistic}(\alpha + \delta t) \quad \text{where} \quad \text{logistic}(z) = \frac{\exp(z)}{1 + \exp(z)}$$

which implies

$$\ln \text{odds}\{\Pr(Y_t = 1)\} = \alpha + \delta t$$

- Assume T is binary. We then recover the MOR as

$$\text{MOR} = \exp\{(\alpha + \delta) - (\alpha)\} = \exp(\delta)$$

- Meaning: the (exponent of the) slope coefficient in a simple logistic regression estimates the MOR
(The same also holds in case of a continuous treatment, which is easy to show.)

Relationship to the logistic model

- If we condition on \mathbf{X} , then

$$\ln \text{odds}\{\Pr(Y_t = 1 | \mathbf{X} = \mathbf{x})\} = \alpha + \delta t + \mathbf{x}\beta$$

- Here $\exp(\delta)$ is the *conditional* odds ratio (i.e., the odds ratio within a subgroup defined by a specific value of \mathbf{X}). For a binary treatment:

$$\text{COR} = \frac{\text{odds}\{\text{logistic}(\alpha + \delta + \mathbf{x}\beta)\}}{\text{odds}\{\text{logistic}(\alpha + \mathbf{x}\beta)\}} = \exp(\delta)$$

- The conditional odds ratio (COR) is different from the MOR, which has a more involved form. For example, in case of a binary treatment:

$$\text{MOR} = \frac{\text{odds}\{E_{\mathbf{X}}[\text{logistic}(\alpha + \delta + \mathbf{x}\beta)]\}}{\text{odds}\{E_{\mathbf{X}}[\text{logistic}(\alpha + \mathbf{x}\beta)]\}}$$

This will be different from COR when $\beta \neq \mathbf{0}$.

Relationship to the logistic model

- The difference between MOR and COR is referred to as *noncollapsibility* or *rescaling bias*.
- “Noncollapsibility of the OR derives from the fact that when the expected probability of outcome is modeled as a nonlinear function of the exposure, the marginal effect cannot be expressed as a weighted average of the conditional effects” (Pang et al. 2016).
- MOR will be attenuated compared to COR (what is commonly referred to as rescaling effects).
 - ▶ For example, if there is just a single covariate, the relationship between MOR and COR can be approximated by

$$\ln \text{MOR} = \frac{\ln \text{COR}}{\sqrt{1 + 0.35\beta^2 \text{var}(X)}}$$

Main message

MOR and COR correspond to different estimands!
They are conceptually different.

- Both MOR and COR are valid estimands. Why should we choose one over the other?
 1. While there exists only one MOR, there are many CORs, as the latter depend on the conditioning set \mathbf{X} . That is, the interpretation of the COR depends on the covariates included in the regression equation.
 2. The MOR has an interpretation similar to an AME on the probability scale: it quantifies the “population response” to a treatment.
 3. Because the MOR is unaffected by noncollapsibility, it can be used to compare results from same-sample models including different covariates (e.g. effect decompositions in mediation analysis).
 4. The MOR is straightforward to compare across different studies or populations as it does not depend in arbitrary ways on the conditioning set.

- 1 Background
- 2 Marginal (log) odds ratios
- 3 Two illustrations**
- 4 Estimation
- 5 Applied example
- 6 Discussion

Academic ability and intergenerational college mobility

- Comparison of “secondary effects” of family background on educational attainment between the United States (National Longitudinal Survey of Youth, 1979 cohort) and Denmark (Danish Longitudinal Survey of Youth).
- How much does educational attainment (here: going to college) depend on whether parents have college education, once we control for academic ability (the “primary effect”)?

Academic ability and intergenerational college mobility

Odds ratios and average marginal effects of parental college attainment gap in college attainment unadjusted and adjusted for academic ability. The United States and Denmark. Standard errors in parentheses.

	USA (N = 10,068)	DNK (N = 2,185)	USA/DNK
1: MOR: Unadjusted	7.7 (0.46)	3.8 (0.55)	2.03*
2: MOR: Adjusted	2.5 (0.13)	2.5 (0.33)	1.00
3: COR: Adjusted	3.4 (0.23)	2.9 (0.45)	1.17
4: COR: Unadjusted (khh)	14.2 (1.05)	4.9 (0.78)	2.90*
5: AME: Unadjusted	0.43 (0.01)	0.31 (0.03)	1.38*
6: AME: Adjusted	0.17 (0.01)	0.20 (0.20)	0.86

Note: MOR is marginal odds ratio; COR is conditional odds ratio; AME is average marginal effect; khb is the Karlson-Holm-Breen decomposition method (using orthogonalized predictors). US data are from the NLSY79; the Danish data are from the Danish Longitudinal Survey of Youth. * indicates that the country difference in log odds ratios is statistically significant at a five percent level.

(Karlson and Jann 2023)

Academic ability and intergenerational college mobility

- Unadjusted or gross marginal odds ratio is about twice as large in the US as in Denmark, meaning that Denmark is significantly more educationally mobile.
- However, the “secondary effects” of family background are of similar magnitude in the two countries (adjusted MOR). (This means that academic ability “mediates” a significantly larger portion of the overall effect in the US than in Denmark.)
- Had we been using the adjusted COR, we would have concluded that, net of academic ability, Denmark is a (albeit only slightly) more educationally mobile country.
- This is because the attenuating impact of noncollapsibility is more pronounced in the US than in Denmark (since academic ability is a much stronger predictor of college attainment in the US).

College gap in attitudes toward racial segregation

- Analysis of how the gap in attitudes toward racial segregation between respondents with and without a four-year college degree changed over time.
- General Social Survey cumulative file, 1976–1996
- Outcome: Agreement with claim that white people have a right to keep black people out of their neighborhoods.
- Predictor of interest: college attainment
- Controls: survey year, age, gender, race, marital status, liberal–conservative scale
- Joint model across time points with survey year fully interacted with college attainment and all other covariates; we obtain AMEs and MORs from the model at different years.

College gap in attitudes toward racial segregation

Average marginal effects and marginal odds ratios of the college gap in the attitude toward racial segregation in 1976, 1981, 1986, 1991, and 1996. Standard errors in parenthesis.

	AME_{LPM}	AME_{LOGIT}	MOR	lnMOR
1976	-0.215 (0.019)	-0.213 (0.021)	0.366 (0.043)	-1.006 (0.118)
1981	-0.180 (0.013)	-0.181 (0.012)	0.374 (0.030)	-0.983 (0.079)
1986	-0.146 (0.009)	-0.147 (0.008)	0.383 (0.025)	-0.959 (0.065)
1991	-0.111 (0.011)	-0.115 (0.009)	0.393 (0.035)	-0.934 (0.089)
1996	-0.076 (0.017)	-0.088 (0.010)	0.403 (0.053)	-0.909 (0.130)
<i>1976–1996 difference</i>	0.138 (0.031)	0.125 (0.028)	-	0.097 (0.211)
<i>1976–1996 prop. reduction</i>	64.4% (9.6)	58.5% (7.7)	-	9.7% (20.0)


Note: MOR is marginal odds ratio; LPM is linear probability model; AME is average marginal effect. Estimates are adjusted for gender, race, age, marital status, and overall political view. Data are from General Social Surveys Cumulative File, N = 12,239.

(Karlson and Jann 2023)

College gap in attitudes toward racial segregation

- Main finding is that the absolute college gap (as measured by the AME) in the attitude toward racial segregation has declined significantly over the 20-year period, whereas the relative gap (as measured by the MOR) remained practically unchanged.
- The absolute gap reduced over time because there is a general decline in support of racial segregation, and this decline is steeper among the non-college educated in absolute terms because they start at a higher level than the college educated. However, the relative difference does not change much (see figure).
- The example illustrates how the MOR can be an informative complement to the AME.

College gap in attitudes toward racial segregation



(Karlson and Jann 2023)

- 1 Background
- 2 Marginal (log) odds ratios
- 3 Two illustrations
- 4 Estimation**
- 5 Applied example
- 6 Discussion

Estimation

- Estimand \Rightarrow Estimation
- There are several approaches how we can estimate the MOR.
 - ▶ G-computation (using predictions from a model)
 - ▶ Inverse probability weighting
 - ▶ Unconditional logistic regression (RIF regression)
- All are discussed in detail in Jann and Karlson (2023) (for binary/categorical as well as continuous treatments; including formulas for analytic standard errors based on influence functions).
- Here we focus on G-computation as it closely resembles the formulation of the adjusted MOR above. That is, G-computation obtains the MOR that is *implied* by the chosen logit model. The other methods follow a somewhat different logic.

G-computation

- G-computation estimates the MOR using counterfactual predictions from a logit model (or any other model in principle).
- For example, for a binary treatment, the procedure is as follows.
 1. Regress Y on T and \mathbf{X} using logistic regression (or, in principle, any other model).
 2. Use the model estimates to generate two predictions of $\Pr(Y = 1)$ for each observation, one with T set to 0 and one with T set to 1.
 3. Predictions are then averaged across the sample to obtain estimates of the population-averaged success probability by treatment level.
 4. These average predictions can then be plugged into the formula for the MOR:

$$\ln \widehat{\text{MOR}} = \ln \text{odds}(\bar{p}^{T=1}) - \ln \text{odds}(\bar{p}^{T=0})$$

- Note that, in Stata, `margins` followed by `nlcom` can be used to do the above computations in case of a binary treatment.

G-computation

- For continuous treatments we evaluate level-specific MORs (using analytic derivatives) at each level of the treatment (possibly using an approximation grid) and then average over the treatment distribution (not directly possible with `margins`).
- An alternative approach is based on applying fractional logit to averaged counterfactual predictions at each value of T . For binary/categorical treatments this leads to the same results as the procedure above. For continuous treatments results slightly differ (due to the different implicit averaging). Nonetheless we prefer this approach due to its generality and flexibility.

Software

- We provide three Stata packages for the estimation of marginal odds ratios, each implementing one of the three approaches. All packages provide consistent standard errors and support complex survey estimation.
 - ▶ `lnmor`: G-computation
(<https://github.com/benjann/lnmor>)
 - ▶ `ipwlogit`: Inverse probability weighting
(<https://github.com/benjann/ipwlogit>)
 - ▶ `riflogit`: Unconditional logistic regression
(<https://github.com/benjann/riflogit>)
- Installation:
 - ▶ Type

```
ssc install name, replace
```

where *name* is the package name.

- 1 Background
- 2 Marginal (log) odds ratios
- 3 Two illustrations
- 4 Estimation
- 5 Applied example**
- 6 Discussion

Example

- Application: gender gap in STEM

```
. use stem, clear
```

```
(Excerpt from TREE cohort 2)
```

```
. describe
```

```
Contains data from stem.dta
```

```
Observations:      6,809
```

```
Variables:         7
```

```
Excerpt from TREE cohort 2
```

```
11 Jan 2023 11:05
```

Variable name	Storage type	Display format	Value label	Variable label
stem	byte	%8.0g		Is in STEM training
male	byte	%8.0g		Is male
mathscore	double	%10.0g		Math score
repeat	byte	%8.0g		Ever repeated a grade
books	byte	%19.0g	books	Number of books at home
wt	double	%10.0g		Sampling weight
psu	int	%8.0g		Sampling unit

Sorted by:

Example

- Probability difference

```
. mean stem [pw=wt], over(male) cluster(psu)
```

```
Mean estimation                               Number of obs = 6,809  
(Std. err. adjusted for 800 clusters in psu)
```

	Mean	Robust std. err.	[95% conf. interval]	
c.stem@male				
0	.1632341	.0093646	.1448519	.1816163
1	.2748702	.014516	.2463762	.3033643

```
. regress stem i.male [pw=wt], cluster(psu) noheader  
(sum of wgt is 78,600.15)
```

```
(Std. err. adjusted for 800 clusters in psu)
```

stem	Coefficient	Robust std. err.	t	P> t	[95% conf. interval]	
1.male	.1116361	.0142969	7.81	0.000	.0835721	.1397
_cons	.1632341	.0093653	17.43	0.000	.1448506	.1816177

Example

- Unadjusted (gross) OR

```
. logit stem i.male [pw=wt], or cluster(psu) nolog
Logistic regression
```

```
Number of obs = 6,809
Wald chi2(1) = 67.37
Prob > chi2 = 0.0000
Pseudo R2 = 0.0172
```

```
Log pseudolikelihood = -40949.271
```

(Std. err. adjusted for 800 clusters in psu)

stem	Odds ratio	Robust std. err.	z	P> z	[95% conf. interval]	
1.male	1.943143	.1572675	8.21	0.000	1.658109	2.277176
_cons	.1950775	.0133747	-23.84	0.000	.1705485	.2231342

Note: _cons estimates baseline odds.

Example

- Conventional approach: “conditional” OR

```
. logit stem i.male mathscore i.repeat books [pw=wt], or cluster(psu) nolog
Logistic regression                               Number of obs = 6,809
                                                    Wald chi2(4) = 596.06
                                                    Prob > chi2 = 0.0000
Log pseudolikelihood = -31906.84                 Pseudo R2 = 0.2342
                                                    (Std. err. adjusted for 800 clusters in psu)
```

stem	Odds ratio	Robust std. err.	z	P> z	[95% conf. interval]	
1.male	1.959347	.167555	7.87	0.000	1.656991	2.316874
mathscore	2.605975	.1251954	19.94	0.000	2.371795	2.863278
1.repeat	.6564493	.0965051	-2.86	0.004	.4921137	.8756628
books	1.086996	.0341203	2.66	0.008	1.022137	1.155971
_cons	.1058298	.0166896	-14.24	0.000	.0776911	.1441599

Note: _cons estimates baseline odds.

Example

- G-computation approach (`lnmor` is a post-estimation command, i.e. first estimate the model, then apply `lnmor`)

```
. lnmor i.male, or
```

```
Enumerating predictions: male..done
```

```
Marginal odds ratio
```

```
Number of obs      =      6,809
```

```
Command            =      logit
```

```
(Std. err. adjusted for 800 clusters in psu)
```

stem	Odds Ratio	Robust std. err.	t	P> t	[95% conf. interval]	
1.male	1.677102	.1103147	7.86	0.000	1.473958	1.908244

Example

- Compare results (SEs in parentheses)

ln(MOR)	Unadjusted	Conditional	Adjusted
1.male	0.664 (0.0809)	0.673 (0.0855)	0.517 (0.0658)

MOR	Unadjusted	Conditional	Adjusted
1.male	1.943 (0.157)	1.959 (0.168)	1.677 (0.110)

Example

- Could also do a statistical test for confounding!

```
. quietly logit stem i.male [pw=wt], cluster(psu)
. quietly lnmor i.male, nodots rif(RIF*)
. quietly logit stem i.male mathscore i.repeat books [pw=wt], cluster(psu)
. quietly lnmor i.male, nodots rif(RIFadj*)
. quietly total RIF2 RIFadj2 [pw=wt], cluster(psu)
. lincom RIFadj2 - RIF2
( 1)  - RIF2 + RIFadj2 = 0
```

Total	Coefficient	Std. err.	t	P> t	[95% conf. interval]	
(1)	-.1472398	.042049	-3.50	0.000	-.2297794	-.0647002

```
. drop RIF*
```

Some further functionality

- Obtain results for several predictors in one call

```
. logit stem i.male mathscore i.repeat books [pw=wt], cluster(psu)
```

```
(output omitted)
```

```
. lnmod i.male mathscore i.repeat books, or  
(mathscore has 380 levels; using 82 binned levels)
```

```
Enumerating predictions: male..mathscore.....
```

```
.....repeat..books.....done
```

```
Marginal odds ratio                               Number of obs    =    6,809  
                                                    Command          =    logit
```

```
(Std. err. adjusted for 800 clusters in psu)
```

stem	Odds Ratio	Robust std. err.	t	P> t	[95% conf. interval]	
1.male	1.677102	.1103147	7.86	0.000	1.473958	1.908244
mathscore	2.544088	.1227496	19.35	0.000	2.314196	2.796817
1.repeat	.7244886	.0838845	-2.78	0.006	.5771998	.9093623
books	1.065503	.0258046	2.62	0.009	1.016036	1.11738

Some further functionality

- Using `at()` to evaluate interactions

```
. logit stem i.male##c.mathscore##c.mathscore##i.repeat##c.books [pw=wt], ///  
>     cluster(psu)  
  (output omitted)  
. lnmor i.male, nodots or at(repeat)  
Marginal odds ratio                               Number of obs   =    6,809  
                                                    Command          =    logit  
  
Evaluated at:  
  1: repeat = 0  
  2: repeat = 1  
  
                (Std. err. adjusted for 800 clusters in psu)
```

	stem	Odds Ratio	Robust std. err.	t	P> t	[95% conf. interval]	
1	1.male	1.700848	.1254404	7.20	0.000	1.47161	1.965795
2	1.male	1.514373	.3346268	1.88	0.061	.9814315	2.336716

Some further functionality

- Using `at()` to evaluate interactions

```
. lnmor i.male, nodots or at(mathscore = -2(2)2)
```

```
Marginal odds ratio      Number of obs      =      6,809
                          Command                =      logit
```

```
Evaluated at:
```

```
1: mathscore = -2
2: mathscore = 0
3: mathscore = 2
```

(Std. err. adjusted for 800 clusters in psu)

	stem	Odds Ratio	Robust std. err.	t	P> t	[95% conf. interval]	
1	1.male	1.697881	.6743511	1.33	0.183	.7786114	3.702489
2	1.male	1.890371	.2008559	5.99	0.000	1.534504	2.328768
3	1.male	1.992419	.356629	3.85	0.000	1.402137	2.831202

Some further functionality

- Analyze nonlinear effects:

```
. logit stem i.male c.mathscore##c.mathscore i.repeat c.books [pw=wt], ///  
>     cluster(psu)  
    (output omitted)  
. lnmor c.mathscore##c.mathscore  
(mathscore has 380 levels; using 82 binned levels)  
Enumerating predictions: mathscore.....  
.....done  
Marginal log odds ratio                               Number of obs    =    6,809  
                                                       Command           =    logit  
                                                       (Std. err. adjusted for 800 clusters in psu)
```

stem	Coefficient	Robust std. err.	t	P> t	[95% conf. interval]	
mathscore	1.0229	.0698468	14.64	0.000	.8857948	1.160005
c.mathscore# c.mathscore	-.0757576	.0276335	-2.74	0.006	-.1300004	-.0215148

Some further functionality

- Analyze nonlinear effects: level-specific MORs using option `dx()`

```
. lnmor mathscore, or dx(-3(1)3)
Enumerating predictions: mathscore.....done
Marginal odds ratio                               Number of obs   =    6,809
                                                    Command         =    logit
                                                    Type of dx()    =    levels
                                                    (Std. err. adjusted for 800 clusters in psu)
```

stem	Odds Ratio	Robust std. err.	t	P> t	[95% conf. interval]	
mathscore@11	4.424869	1.003165	6.56	0.000	2.835509	6.905096
mathscore@12	3.805942	.6538704	7.78	0.000	2.71645	5.332398
mathscore@13	3.259636	.385185	10.00	0.000	2.584828	4.110613
mathscore@14	2.779552	.1943118	14.62	0.000	2.423143	3.188383
mathscore@15	2.378328	.1147795	17.95	0.000	2.163366	2.61465
mathscore@16	2.056595	.1576843	9.40	0.000	1.769237	2.390626
mathscore@17	1.788865	.2248208	4.63	0.000	1.397778	2.289376

```
Terms affected by dx(): mathscore
Levels of dx(): -3 -2 -1 0 1 2 3
```

- 1 Background
- 2 Marginal (log) odds ratios
- 3 Two illustrations
- 4 Estimation
- 5 Applied example
- 6 Discussion**

Discussion

- We provide a clear definition of the marginal OR (clarification of estimand).
- We illustrate the advantages of the marginal odds ratio over the conditional odds ratio; we illustrate the value of the marginal odds ratio as a complement to AMEs.
- We provide flexible software that can estimate the marginal OR for categorical as well as continuous predictors (including support for complex surveys).
- But ...

... is it worth the hassle? How much do applied researchers love odds ratios?

... will it change practice?

More information

- Jann, B., K.B. Karlson. 2023. Estimation of marginal odds ratios. University of Bern Social Sciences Working Papers 44. <https://ideas.repec.org/p/bss/wpaper/44.html>
- Karlson, K.B., B. Jann. 2023. Marginal Odds Ratios: What They Are, How to Compute Them, and Why Sociologists Might Want to Use Them. University of Bern Social Sciences Working Papers 45. <https://ideas.repec.org/p/bss/wpaper/45.html> (forthcoming in *Sociological Science*)

References

- Daniel, R., J. Zhang, and D. Farewell. 2021. Making apples from oranges: Comparing noncollapsible effect estimators and their standard errors after adjustment for different covariate sets. *Biometrical Journal* 63(3): 528–557.
- Jann, Ben. 2022. VIOLINPLOT: Stata module to draw violin plots. Available from <https://ideas.repec.org/c/boc/bocode/s459132.html>.
- Karlson, K. B., A. Holm, R. Breen. 2012. Comparing Regression Coefficients Between Models using Logit and Probit: A New Method. *Sociological Methodology* 42: 286–313.
- Karlson, K. B., F. Popham, and A. Holm. 2021. Marginal and Conditional Confounding Using Logits. *Sociological Methods & Research*.
- Pang, M., J. S. Kaufman, and R. W. Platt. 2016. Studying noncollapsibility of the odds ratio with marginal structural and logistic regression models. *Statistical Methods in Medical Research* 25(5): 1925–1937.
- Stampf, S., E. Graf, C. Schmoor, and M. Schumacher. 2010. Estimators and confidence intervals for the marginal odds ratio using logistic regression and propensity score stratification. *Statistics in Medicine* 29: 760–769.
- Zhang, Z. 2008. Estimating a Marginal Causal Odds Ratio Subject to Confounding. *Communications in Statistics – Theory and Methods* 38(3): 309–321.


Appendix: Some simulation results

- Setup

- ▶ Binary outcome Y depends on treatment T and control variable X through a logistic model.
- ▶ The effects of T and X on Y (the conditional log odds ratios) are set to 1 in all simulations (intercept is 0).
- ▶ X has a standard normal distribution.
- ▶ T is either binary or continuous.
- ▶ Two scenarios:
 1. unconfounded: T is independent from X (X has an even distribution if binary and a standard normal distribution if continuous)
 2. confounded: T depends on X (binary: logistic model with slope 0.5; continuous: linear model with slope 0.5 and standard normal errors)
- ▶ 10'000 replications.
- ▶ Using `violinplot` (Jann 2022) to display results.


Appendix: Some simulation results

- Distribution of effect estimates for binary treatment




Appendix: Some simulation results

- Distribution of standard errors for binary treatment




Appendix: Some simulation results

- Distribution of effect estimates for continuous treatment



Appendix: Some simulation results

- Distribution of standard errors for continuous treatment



Appendix: Some simulation results

- Binary treatment:
 - ▶ All estimators appear to work well.
 - ▶ However, note that the treatment has an even distribution in these simulations; may need to do more simulations with uneven distributions.
- Continuous treatment:
 - ▶ IPW does not fully remove confounding. Furthermore, stability of IPW becomes problematic. Truncation helps somewhat but also increases bias.
 - ▶ MOR' ("observed") is a different estimand than \overline{MOR} ("averaged"). RIF logit appears to approximate MOR' , not \overline{MOR} .