On the central role of Somers' *D*

Roger Newson
Imperial College London, UK
r.newson@imperial.ac.uk
http://www.imperial.ac.uk/nhli/r.newson/

The Avon Longitudinal Study of Parents and Children (ALSPAC)
http://www.alspac.bris.ac.uk/

Presented at the 12th UK Stata Users' Group Meeting on 11-12 September 2006

This presentation can be downloaded from the conference website at
http://ideas.repec.org/s/boc/usug06.html

What is Somers' D?

What is Somers' D?

- We assume a population of (X, Y)–pairs, and a sampling scheme for sampling pairs of pairs (X_i, Y_i) and (X_j, Y_j) from that population.

What is Somers' D?

- We assume a population of (X, Y)–pairs, and a sampling scheme for sampling pairs of pairs (X_i, Y_i) and (X_j, Y_j) from that population.
- **Kendall's** τ_a is defined as the expectation

$$\tau_{XY} = \mathrm{E}[\mathrm{sign}(X_i - X_j)\mathrm{sign}(Y_i - Y_j)]$$

  or as the difference between the probabilities of concordance and discordance between the two (X, Y)–pairs.

What is Somers' D?

- We assume a population of (X, Y)–pairs, and a sampling scheme for sampling pairs of pairs (X_i, Y_i) and (X_j, Y_j) from that population.
- **Kendall's** τ_a is defined as the expectation

$$\tau_{XY} = \mathrm{E}[\mathrm{sign}(X_i - X_j)\mathrm{sign}(Y_i - Y_j)]$$

  or as the difference between the probabilities of concordance and discordance between the two (X, Y)–pairs.

- **Somers'** D is defined as the ratio

$$D_{YX} = \tau_{XY}/\tau_{XX}$$

  or as the difference between the two corresponding *conditional* probabilities, given that one X–value is known to be larger than the other X–value.

What is Somers' D?

- We assume a population of (X, Y)–pairs, and a sampling scheme for sampling pairs of pairs (X_i, Y_i) and (X_j, Y_j) from that population.
- **Kendall's** τ_a is defined as the expectation

$$\tau_{XY} = \mathrm{E}[\mathrm{sign}(X_i - X_j)\mathrm{sign}(Y_i - Y_j)]$$

  or as the difference between the probabilities of concordance and discordance between the two (X, Y)–pairs.
- **Somers'** D is defined as the ratio

$$D_{YX} = \tau_{XY}/\tau_{XX}$$

  or as the difference between the two corresponding *conditional* probabilities, given that one X–value is known to be larger than the other X–value.
- These definitions can be extended to cases where the X–values and/or the Y–values may be weighted and/or left–censored and/or right–censored.

You have already met Somers' D

You have already met Somers' D

- If X and Y are *both* binary, then Somers' D is the difference between proportions:

$$D_{YX} = \Pr(Y = 1|X = 1) - \Pr(Y = 1|X = 0)$$

You have already met Somers' D

- If X and Y are *both* binary, then Somers' D is the difference between proportions:

$$D_{YX} = \Pr(Y = 1|X = 1) - \Pr(Y = 1|X = 0)$$

- If X is binary, and Y_1 and Y_0 are sampled from groups $X = 1$ and $X = 0$, then

$$D_{YX} = \Pr(Y_1 > Y_0) - \Pr(Y_0 > Y_1)$$

You have already met Somers' D

- If X and Y are *both* binary, then Somers' D is the difference between proportions:

$$D_{YX} = \Pr(Y = 1|X = 1) - \Pr(Y = 1|X = 0)$$

- If X is binary, and Y_1 and Y_0 are sampled from groups $X = 1$ and $X = 0$, then

$$D_{YX} = \Pr(Y_1 > Y_0) - \Pr(Y_0 > Y_1)$$

- The two groups may be treatment groups, subpopulations, or different scenarios in the same population.

You have already met Somers' D

- If X and Y are *both* binary, then Somers' D is the difference between proportions:

$$D_{YX} = \Pr(Y = 1|X = 1) - \Pr(Y = 1|X = 0)$$

- If X is binary, and Y_1 and Y_0 are sampled from groups $X = 1$ and $X = 0$, then

$$D_{YX} = \Pr(Y_1 > Y_0) - \Pr(Y_0 > Y_1)$$

- The two groups may be treatment groups, subpopulations, or different scenarios in the same population.
- Special cases include the population attributable risk, the ROC area, Harrell's c index, the Gini inequality index, and the parameters behind the "non–parametric" sign test and Wilcoxon and Gehan–Breslow ranksum tests.

You have already met Somers' D

- If X and Y are *both* binary, then Somers' D is the difference between proportions:

$$D_{YX} = \Pr(Y = 1|X = 1) - \Pr(Y = 1|X = 0)$$

- If X is binary, and Y_1 and Y_0 are sampled from groups $X = 1$ and $X = 0$, then

$$D_{YX} = \Pr(Y_1 > Y_0) - \Pr(Y_0 > Y_1)$$

- The two groups may be treatment groups, subpopulations, or different scenarios in the same population.
- Special cases include the population attributable risk, the ROC area, Harrell's c index, the Gini inequality index, and the parameters behind the "non–parametric" sign test and Wilcoxon and Gehan–Breslow ranksum tests.
- *However*, D_{YX} exists whether or not X is binary, and is used to define...

Median differences and slopes

Median differences and slopes

- Somers' D and Kendall's τ_a measure associations between X and Y in terms of differences between proportions.

Median differences and slopes

- Somers' D and Kendall's τ_a measure associations between X and Y in terms of differences between proportions.
- To make monetary or other practical decisions, we may need to know other parameters, such as a between–treatment difference expressed in Y–units, or a treatment effect in Y–units per X–unit.

Median differences and slopes

- Somers' D and Kendall's τ_a measure associations between X and Y in terms of differences between proportions.
- To make monetary or other practical decisions, we may need to know other parameters, such as a between–treatment difference expressed in Y–units, or a treatment effect in Y–units per X–unit.
- The **Theil–Sen median slope** of Y with respect to X is defined as a solution in β to the equation

$$D_{Y-\beta X,X} = 0$$

or (in words) as a linear effect of X on Y sufficient to explain the observed Somers' D.

Median differences and slopes

- Somers' D and Kendall's τ_a measure associations between X and Y in terms of differences between proportions.
- To make monetary or other practical decisions, we may need to know other parameters, such as a between–treatment difference expressed in Y–units, or a treatment effect in Y–units per X–unit.
- The **Theil–Sen median slope** of Y with respect to X is defined as a solution in β to the equation

$$D_{Y-\beta X,X} = 0$$

  or (in words) as a linear effect of X on Y sufficient to explain the observed Somers' D.
- If X is binary, then the Theil–Sen median slope is known as the **Hodges–Lehmann median difference** between groups $X = 1$ and $X = 0$.

The Stata 9 version of the `somersd` package

The Stata 9 version of the `somersd` package

The `somersd` package, downloadable from SSC, has 3 modules to calculate confidence intervals for a large family of rank statistics:

The Stata 9 version of the `somersd` package

The `somersd` package, downloadable from SSC, has 3 modules to calculate confidence intervals for a large family of rank statistics:

- The module `somersd` estimates Somers' D, Harrell's c or Kendall's τ_a, saving the results as estimation results.

The Stata 9 version of the `somersd` package

The `somersd` package, downloadable from SSC, has 3 modules to calculate confidence intervals for a large family of rank statistics:

- The module `somersd` estimates Somers' D, Harrell's c or Kendall's τ_a, saving the results as estimation results.
- The module `censlope` estimates Somers' D, and then estimates the corresponding Theil–Sen median slope.

The Stata 9 version of the `somersd` package

The `somersd` package, downloadable from SSC, has 3 modules to calculate confidence intervals for a large family of rank statistics:

- The module `somersd` estimates Somers' D, Harrell's c or Kendall's τ_a, saving the results as estimation results.
- The module `censlope` estimates Somers' D, and then estimates the corresponding Theil–Sen median slope.
- The module `cendif` estimates a restricted range of Hodges–Lehmann median differences, mostly for small samples.

The Stata 9 version of the `somersd` package

The `somersd` package, downloadable from SSC, has 3 modules to calculate confidence intervals for a large family of rank statistics:

- The module `somersd` estimates Somers' D, Harrell's c or Kendall's τ_a, saving the results as estimation results.
- The module `censlope` estimates Somers' D, and then estimates the corresponding Theil–Sen median slope.
- The module `cendif` estimates a restricted range of Hodges–Lehmann median differences, mostly for small samples.

All of these rank parameters have multiple versions for multiple sampling designs, with data weighted and/or censored and/or clustered and/or stratified.

Example: Prenatal paracetamol exposure and IgE

Example: Prenatal paracetamol exposure and IgE

- The Avon Longitudinal Study of Parents and Children (ALSPAC) is a birth cohort study, based at Bristol University, UK.

Example: Prenatal paracetamol exposure and IgE

- The Avon Longitudinal Study of Parents and Children (ALSPAC) is a birth cohort study, based at Bristol University, UK.
- The mothers of 12127 children were asked whether they ever used paracetamol (acetaminophen) in weeks 20–32 of pregnancy.

Example: Prenatal paracetamol exposure and IgE

- The Avon Longitudinal Study of Parents and Children (ALSPAC) is a birth cohort study, based at Bristol University, UK.
- The mothers of 12127 children were asked whether they ever used paracetamol (acetaminophen) in weeks 20–32 of pregnancy.
- At 7 years of age, immunoglobulin E (IgE) was measured in the blood of 4848 of these children.

Example: Prenatal paracetamol exposure and IgE

- The Avon Longitudinal Study of Parents and Children (ALSPAC) is a birth cohort study, based at Bristol University, UK.
- The mothers of 12127 children were asked whether they ever used paracetamol (acetaminophen) in weeks 20–32 of pregnancy.
- At 7 years of age, immunoglobulin E (IgE) was measured in the blood of 4848 of these children.
- Shaheen *et al.* (2005) found (using geometric mean ratios) that the children of paracetamol users typically had slightly higher IgE levels than children of paracetamol non–users.

Example: Prenatal paracetamol exposure and IgE

- The Avon Longitudinal Study of Parents and Children (ALSPAC) is a birth cohort study, based at Bristol University, UK.
- The mothers of 12127 children were asked whether they ever used paracetamol (acetaminophen) in weeks 20–32 of pregnancy.
- At 7 years of age, immunoglobulin E (IgE) was measured in the blood of 4848 of these children.
- Shaheen *et al.* (2005) found (using geometric mean ratios) that the children of paracetamol users typically had slightly higher IgE levels than children of paracetamol non–users.
- We will re–measure this association, using `censlope` to estimate Somers' D and Hodges–Lehmann median ratios.

Distribution of IgE in the 4848 children with IgE and paracetamol data

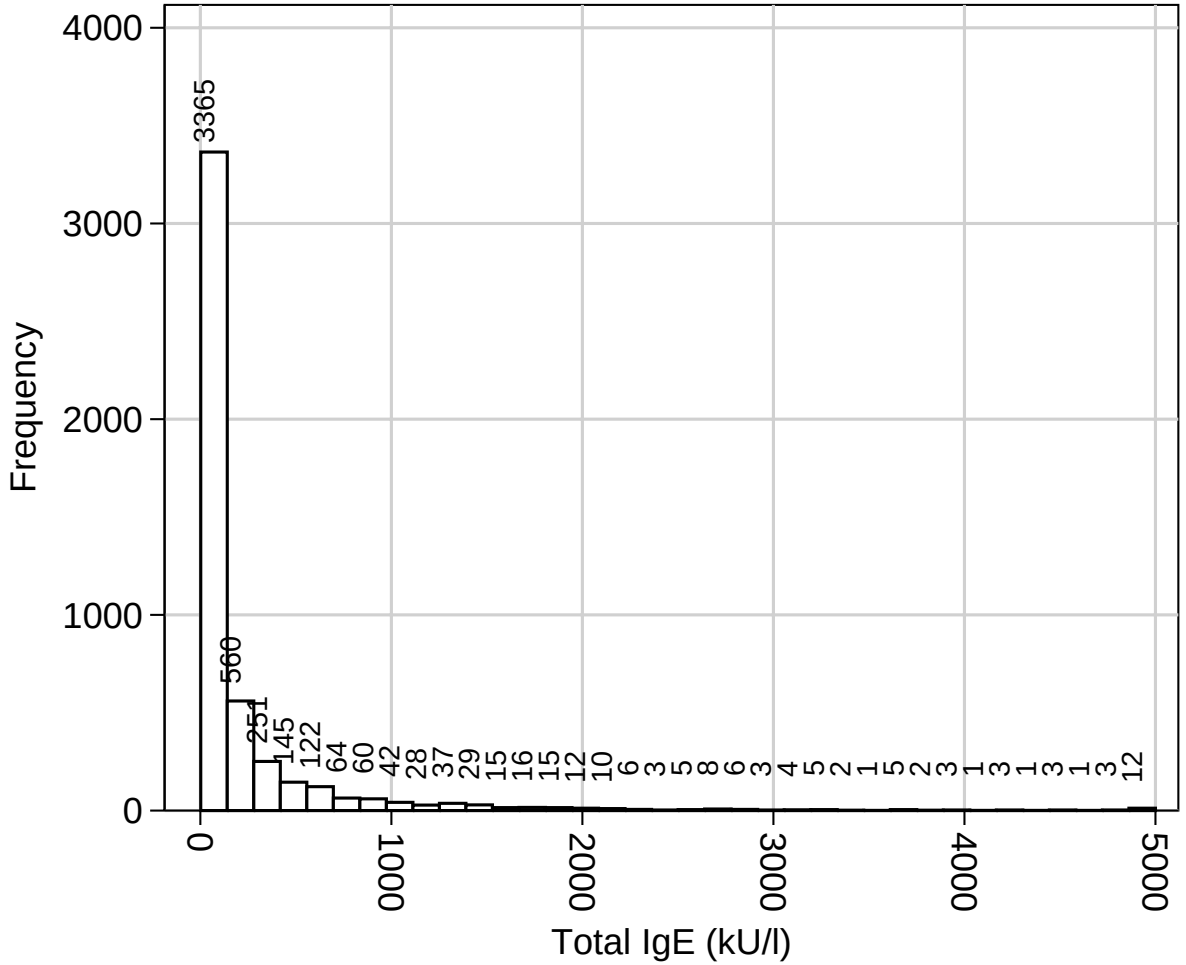

Distribution of IgE in the 4848 children with IgE and paracetamol data

- Total IgE, measured in kilounits/litre (kU/l), is raised in individuals with allergic diseases such as asthma.

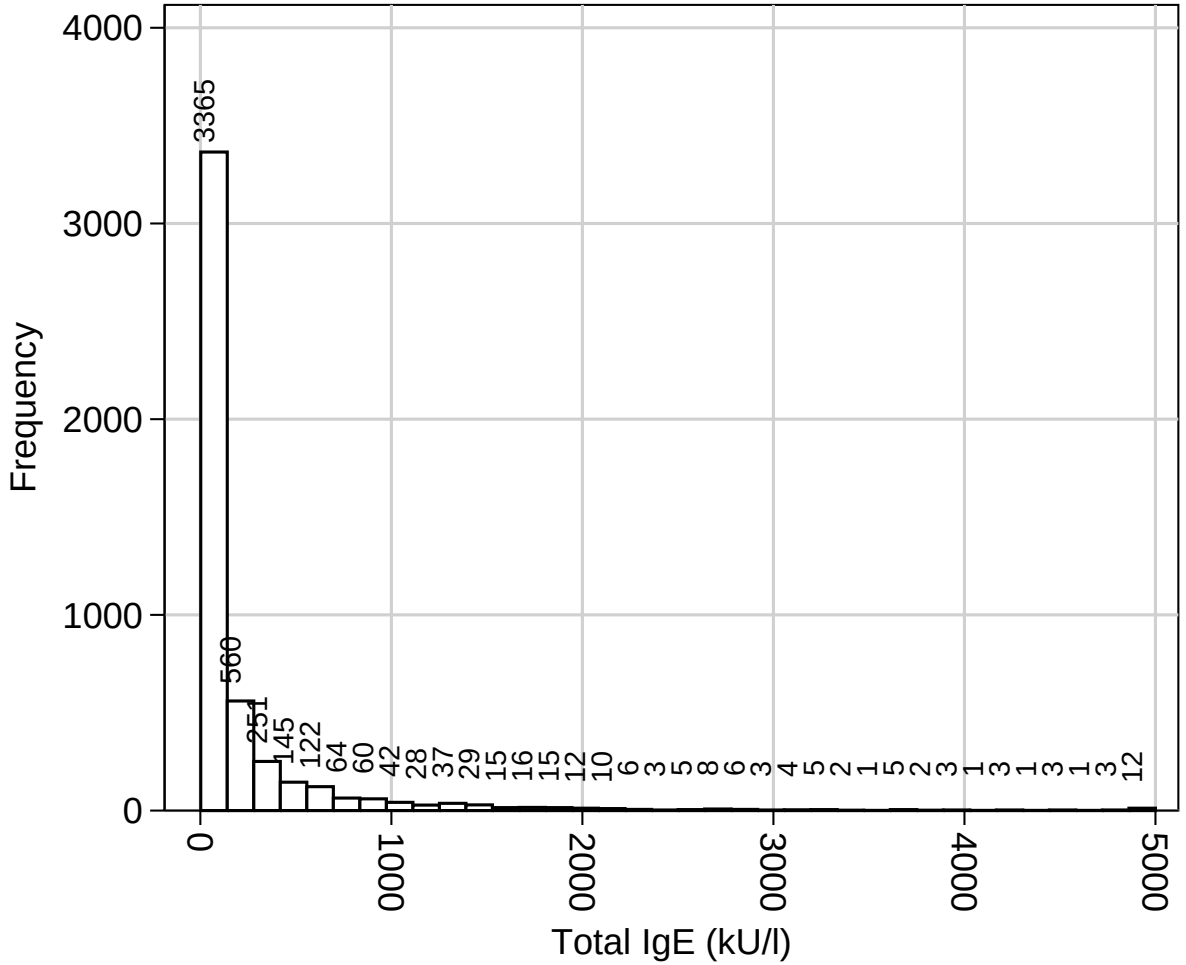

Distribution of IgE in the 4848 children with IgE and paracetamol data

- Total IgE, measured in kilounits/litre (kU/l), is raised in individuals with allergic diseases such as asthma.
- In the 4848 children with IgE and paracetamol data, its overall distribution is non–Normal.

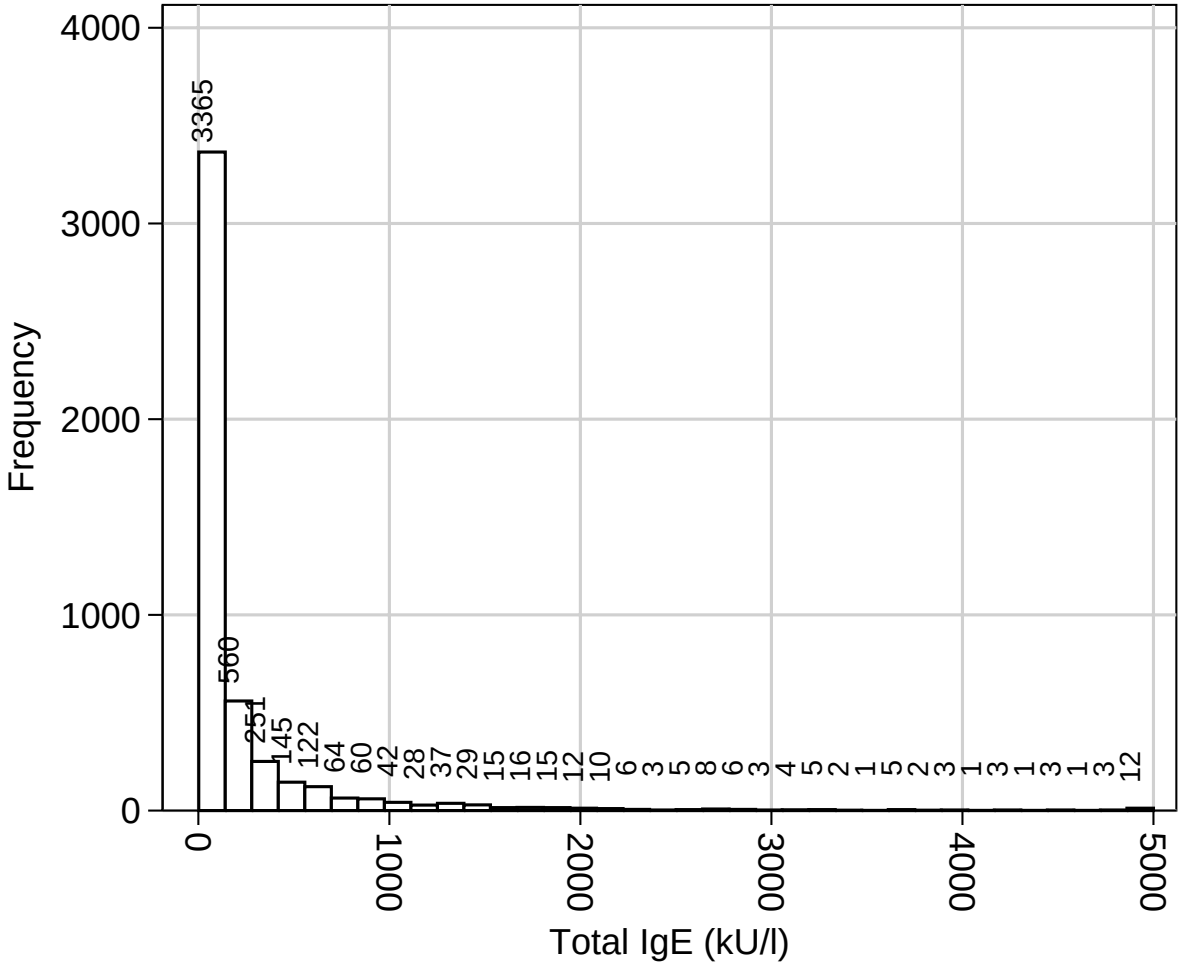

Distribution of IgE in the 4848 children with IgE and paracetamol data

- Total IgE, measured in kilounits/litre (kU/l), is raised in individuals with allergic diseases such as asthma.
- In the 4848 children with IgE and paracetamol data, its overall distribution is non–Normal.
- We wish to compare typical levels in the children of paracetamol users and non–users.

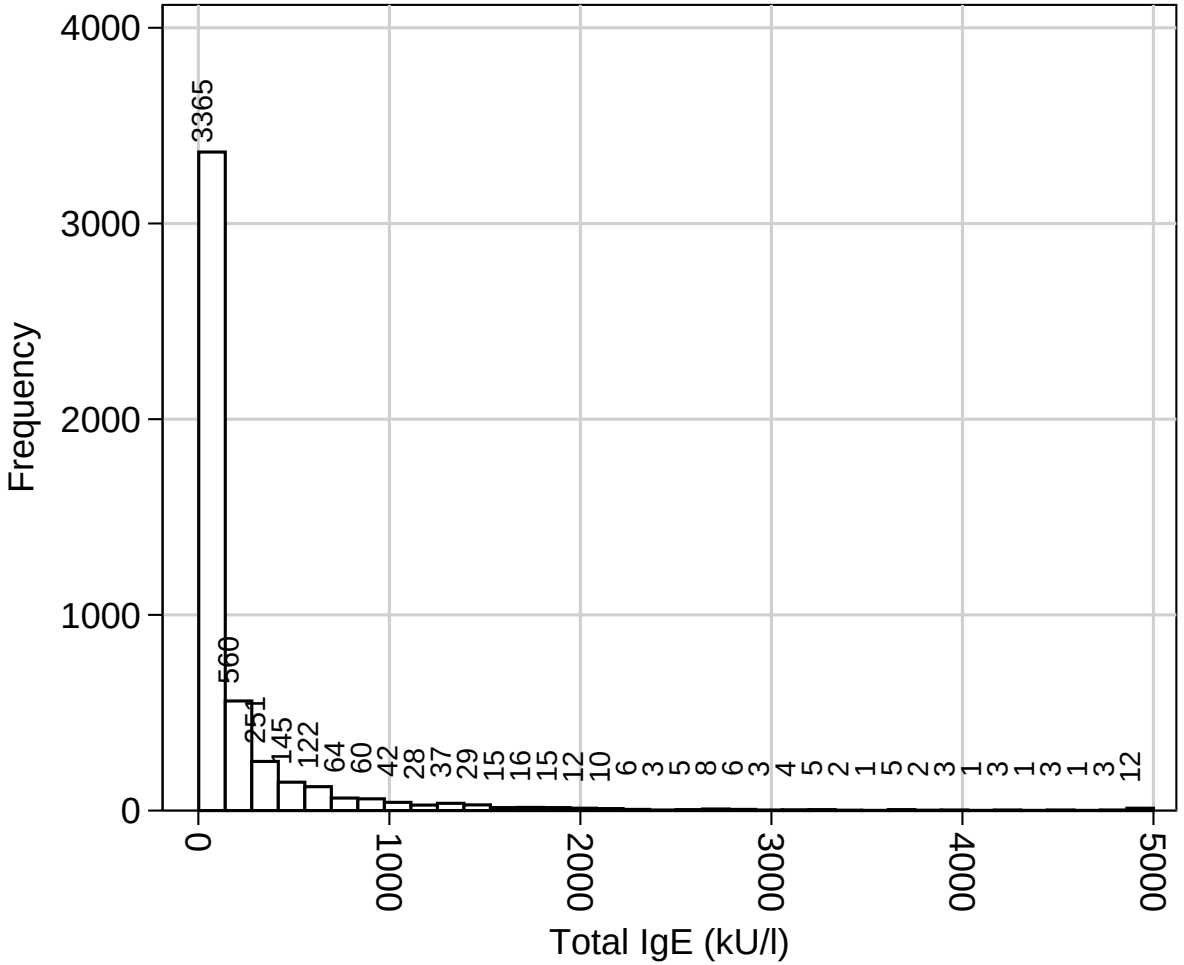

Comparing IgE levels using `censlope`

Comparing IgE levels using `censlope`

- Of the 4848 children, 2051 had mothers who ever used paracetamol during weeks 20–32 of pregnancy.

Comparing IgE levels using `censlope`

- Of the 4848 children, 2051 had mothers who ever used paracetamol during weeks 20–32 of pregnancy.
- Given a randomly–chosen paracetamol–exposed child and a randomly–chosen paracetamol–unexposed child, Somers' D is the difference between the probability that the exposed child has the higher IgE and the probability that the unexposed child has the higher IgE.

Comparing IgE levels using `censlope`

- Of the 4848 children, 2051 had mothers who ever used paracetamol during weeks 20–32 of pregnancy.
- Given a randomly–chosen paracetamol–exposed child and a randomly–chosen paracetamol–unexposed child, Somers' D is the difference between the probability that the exposed child has the higher IgE and the probability that the unexposed child has the higher IgE.
- The Hodges–Lehmann median ratio is the median ratio of IgE levels between two such randomly–chosen children.

Comparing IgE levels using `censlope`

- Of the 4848 children, 2051 had mothers who ever used paracetamol during weeks 20–32 of pregnancy.
- Given a randomly–chosen paracetamol–exposed child and a randomly–chosen paracetamol–unexposed child, Somers' D is the difference between the probability that the exposed child has the higher IgE and the probability that the unexposed child has the higher IgE.
- The Hodges–Lehmann median ratio is the median ratio of IgE levels between two such randomly–chosen children.
- (It is defined as the exponential of the Hodges–Lehmann median difference between the logged IgE values.)

Comparing IgE levels using `censlope`

- Of the 4848 children, 2051 had mothers who ever used paracetamol during weeks 20–32 of pregnancy.
- Given a randomly–chosen paracetamol–exposed child and a randomly–chosen paracetamol–unexposed child, Somers' D is the difference between the probability that the exposed child has the higher IgE and the probability that the unexposed child has the higher IgE.
- The Hodges–Lehmann median ratio is the median ratio of IgE levels between two such randomly–chosen children.
- (It is defined as the exponential of the Hodges–Lehmann median difference between the logged IgE values.)
- We will calculate confidence intervals for these two parameters, using `censlope` with Fisher's z transform.

```
. censlope lnigetot para32g, transf(z) eform;
Outcome variable: lnigetot
Somers' D with variable: para32g
Transformation: Fisher's z
Valid observations: 4848

Symmetric 95% CI for transformed Somers' D
------------------------------------------------------------------------------
             |             Jackknife
    para32g |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
------------+-----------------------------------------------------------------
   lnigetot |   .0533954   .0168421     3.17   0.002     .0203856    .0864053
------------------------------------------------------------------------------

Asymmetric 95% CI for untransformed Somers' D
               Somers_D     Minimum     Maximum
   lnigetot   .05334475   .02038276    .0861909

95% CI(s) for percentile ratio(s)
   Percent  Pctl_Ratio     Minimum     Maximum
        50    1.172549   1.0616111   1.2944986
```

How to adjust for confounders?

How to adjust for confounders?

- To adjust for confounders, we used a propensity score (Rosenbaum and Rubin, 1983).

How to adjust for confounders?

- To adjust for confounders, we used a propensity score (Rosenbaum and Rubin, 1983).
- We fitted a logistic regression model to data from the 12127 children with data on maternal paracetamol use in late pregnancy.

How to adjust for confounders?

- To adjust for confounders, we used a propensity score (Rosenbaum and Rubin, 1983).
- We fitted a logistic regression model to data from the 12127 children with data on maternal paracetamol use in late pregnancy.
- Paracetamol exposure was regressed with respect to the following confounders: gender, maternal age, prenatal tobacco exposure, mother's education, housing tenure, parity, maternal anxiety, maternal ethnic origin, multiple pregnancy, birth weight, gestational age at birth, head circumference, antibiotics in pregnancy, alcohol intake in pregnancy, maternal disease and infection history, younger siblings, presence of pets, breast feeding, day care, dampness problems, passive smoking exposure after birth, obesity index at 7 years.

How to adjust for confounders?

- To adjust for confounders, we used a propensity score (Rosenbaum and Rubin, 1983).
- We fitted a logistic regression model to data from the 12127 children with data on maternal paracetamol use in late pregnancy.
- Paracetamol exposure was regressed with respect to the following confounders: gender, maternal age, prenatal tobacco exposure, mother's education, housing tenure, parity, maternal anxiety, maternal ethnic origin, multiple pregnancy, birth weight, gestational age at birth, head circumference, antibiotics in pregnancy, alcohol intake in pregnancy, maternal disease and infection history, younger siblings, presence of pets, breast feeding, day care, dampness problems, passive smoking exposure after birth, obesity index at 7 years.
- The predicted log paracetamol odds, or propensity score, was grouped into 32 propensity strata, using `xtile`.

Paracetamol exposure prevalence in the 32 propensity groups

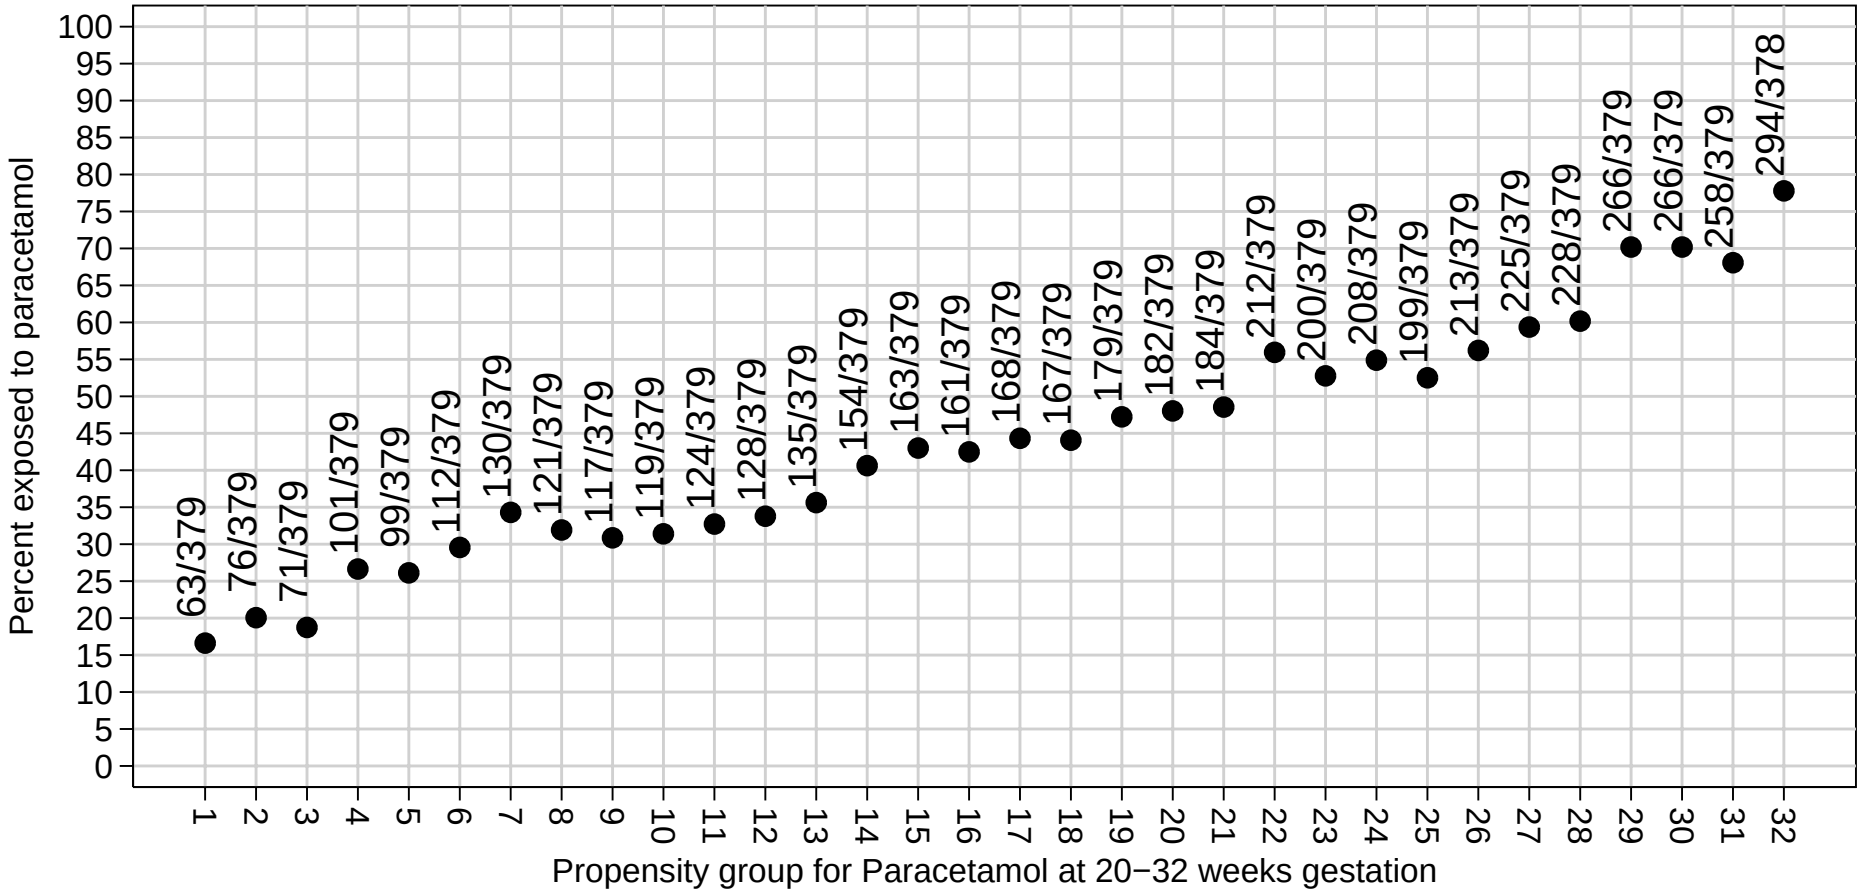

Paracetamol exposure prevalence in the 32 propensity groups

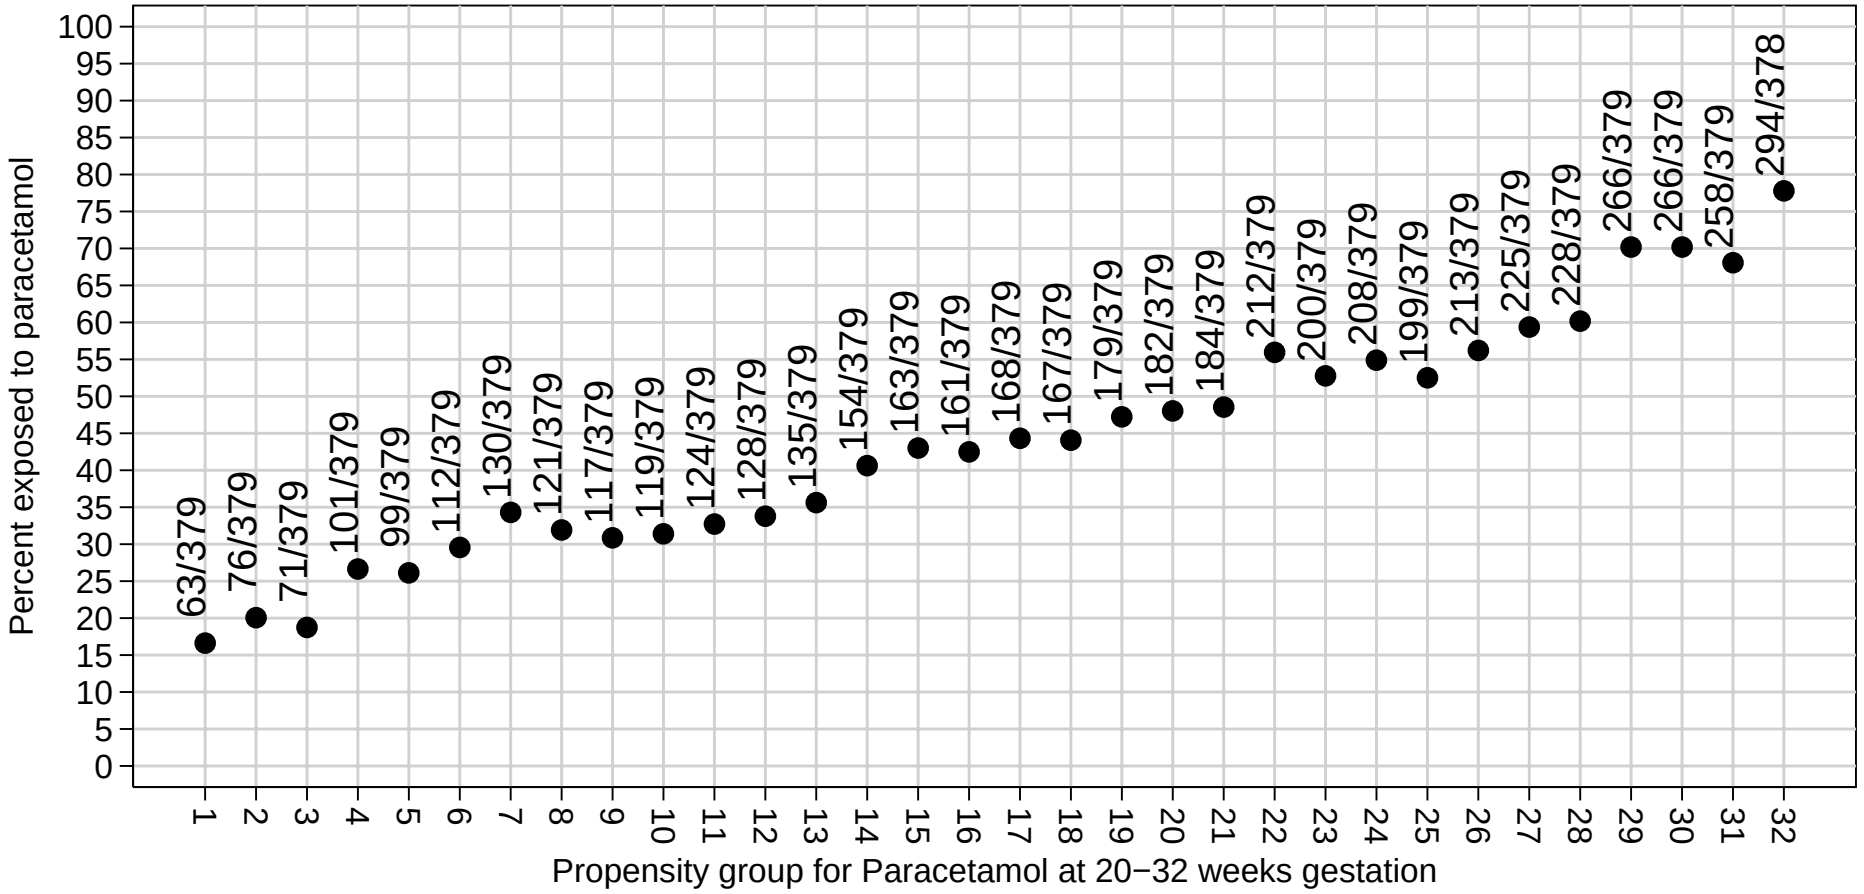

Paracetamol propensity predicts paracetamol exposure, but not *too* well!

Within–strata rank statistics using `somersd`

Within–strata rank statistics using `somersd`

- Kendall's τ_a and Somers' D can be restricted to comparisons within strata, using the `wstrata()` option of `somersd`.

Within–strata rank statistics using `somersd`

- Kendall's τ_a and Somers' D can be restricted to comparisons within strata, using the `wstrata()` option of `somersd`.
- *Therefore*, so can median slopes, differences and ratios.

Within–strata rank statistics using `somersd`

- Kendall's τ_a and Somers' D can be restricted to comparisons within strata, using the `wstrata()` option of `somersd`.
- *Therefore*, so can median slopes, differences and ratios.
- We can therefore adjust our rank statistics for confounders by restricting to comparisons within the 32 propensity groups.

Within–strata rank statistics using `somersd`

- Kendall's τ_a and Somers' D can be restricted to comparisons within strata, using the `wstrata()` option of `somersd`.
- *Therefore*, so can median slopes, differences and ratios.
- We can therefore adjust our rank statistics for confounders by restricting to comparisons within the 32 propensity groups.
- We will now estimate a propensity–adjusted Somers' D and median ratio, using `censlope`.

```
. censlope lnigetot para32g, transf(z) eform wstrata(pg_para32g);
Outcome variable: lnigetot
Somers' D with variable: para32g
Transformation: Fisher's z
Within strata defined by: pg_para32g
Valid observations: 4848

Symmetric 95% CI for transformed Somers' D
------------------------------------------------------------------------------
             |                 Jackknife
     para32g |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
    lnigetot |   .0416191    .018089     2.30   0.021     .0061653    .0770729
------------------------------------------------------------------------------

Asymmetric 95% CI for untransformed Somers' D
                Somers_D     Minimum     Maximum
    lnigetot   .04159508   .00616518   .07692067

95% CI(s) for percentile ratio(s)
   Percent  Pctl_Ratio     Minimum     Maximum
        50   1.1256541   1.0165742   1.2556066
```

Is 32 propensity groups enough?

Is 32 propensity groups enough?

- 32 propensity groups is more than most statisticians use most of the time (typically 5).

Is 32 propensity groups enough?

- 32 propensity groups is more than most statisticians use most of the time (typically 5).
- *However*, children in the same stratum have the same discrete propensity *group*, not the same continuous propensity *score*.

Is 32 propensity groups enough?

- 32 propensity groups is more than most statisticians use most of the time (typically 5).
- *However*, children in the same stratum have the same discrete propensity *group*, not the same continuous propensity *score*.
- *Therefore*, the association between paracetamol exposure and IgE within paracetamol propensity groups *might possibly* be due to a residual association of both variables with the paracetamol propensity score.

Is 32 propensity groups enough?

- 32 propensity groups is more than most statisticians use most of the time (typically 5).
- *However*, children in the same stratum have the same discrete propensity *group*, not the same continuous propensity *score*.
- *Therefore*, the association between paracetamol exposure and IgE within paracetamol propensity groups *might possibly* be due to a residual association of both variables with the paracetamol propensity score.
- Fortunately, `somersd` can help us to check this possibility.

The two interpretations of Somers' D

The two interpretations of Somers' D

Given an outcome variable Y and a predictor variable X, interpretations of Somers' D fall into two classes:

The two interpretations of Somers' D

Given an outcome variable Y and a predictor variable X, interpretations of Somers' D fall into two classes:

- We may interpret D_{YX} as a measure of the **effect** of X on Y, especially if X is binary, as in the examples so far.

The two interpretations of Somers' D

Given an outcome variable Y and a predictor variable X, interpretations of Somers' D fall into two classes:

- We may interpret D_{YX} as a measure of the **effect** of X on Y, especially if X is binary, as in the examples so far.
- Alternatively, we may interpret D_{XY} as a **performance indicator** for X as a predictor of Y, for comparison with another predictor W.

The two interpretations of Somers' D

Given an outcome variable Y and a predictor variable X, interpretations of Somers' D fall into two classes:

- We may interpret D_{YX} as a measure of the **effect** of X on Y, especially if X is binary, as in the examples so far.
- Alternatively, we may interpret D_{XY} as a **performance indicator** for X as a predictor of Y, for comparison with another predictor W.

The second interpretation is possible because, *if* a positive association of Y with X is caused entirely by a positive association of both variables with a third variable W, *then* we must have the inequality

$$D_{XY} \leq D_{WY}$$

(see Newson (2002) and Newson (2006)), and we can test this inequality using `somersd` and `lincom`.

Comparing Somers' D parameters for paracetamol and paracetamol propensity

Comparing Somers' D parameters for paracetamol and paracetamol propensity

- In the present example, Y is IgE, X is paracetamol exposure, and W is paracetamol propensity.

Comparing Somers' D parameters for paracetamol and paracetamol propensity

- In the present example, Y is IgE, X is paracetamol exposure, and W is paracetamol propensity.
- We use `somersd` to estimate D_{XY} and D_{WY}.

Comparing Somers' D parameters for paracetamol and paracetamol propensity

- In the present example, Y is IgE, X is paracetamol exposure, and W is paracetamol propensity.
- We use `somersd` to estimate D_{XY} and D_{WY}.
- Again, we use the options `wstrata(pg_para32g)` to compare children in the same propensity group, and `transf(z)` to use Fisher's z–transform.

Comparing Somers' D parameters for paracetamol and paracetamol propensity

- In the present example, Y is IgE, X is paracetamol exposure, and W is paracetamol propensity.
- We use `somersd` to estimate D_{XY} and D_{WY}.
- Again, we use the options `wstrata(pg_para32g)` to compare children in the same propensity group, and `transf(z)` to use Fisher's z–transform.
- We then compare the z-transformed D_{XY} and D_{WY}, using `lincom`.

```
. somersd lnigetot para32g ps_para32g, transf(z) wstrata(pg_para32g);
Somers' D with variable: lnigetot
Transformation: Fisher's z
Within strata defined by: pg_para32g
Valid observations: 4848

Symmetric 95% CI for transformed Somers' D
------------------------------------------------------------------------------
             |              Jackknife
    lnigetot |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
     para32g |   .0181683   .0078918     2.30   0.021      .0027006     .033636
  ps_para32g |  -.0082111   .0099832    -0.82   0.411     -.0277777    .0113556
------------------------------------------------------------------------------

Asymmetric 95% CI for untransformed Somers' D
                Somers_D     Minimum     Maximum
     para32g    .0181663   .00270058   .03362334
  ps_para32g  -.00821087   -.0277706   .01135515
```

```
. somersd lnigetot para32g ps_para32g, transf(z) wstrata(pg_para32g);
Somers' D with variable: lnigetot
Transformation: Fisher's z
Within strata defined by: pg_para32g
Valid observations: 4848

Symmetric 95% CI for transformed Somers' D
------------------------------------------------------------------------------
             |              Jackknife
    lnigetot |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
     para32g |   .0181683   .0078918     2.30   0.021      .0027006     .033636
  ps_para32g |  -.0082111   .0099832    -0.82   0.411     -.0277777    .0113556
------------------------------------------------------------------------------

Asymmetric 95% CI for untransformed Somers' D
                Somers_D     Minimum     Maximum
     para32g    .0181663   .00270058   .03362334
  ps_para32g  -.00821087   -.0277706   .01135515
```

Paracetamol exposure (para32g) is a significant positive predictor, and paracetamol propensity (ps_para32g) is a non–significant negative predictor.

However, to test the inequality, we use `lincom` to define a confidence interval and a P–value for half the difference between the two z–transformed Somers' D parameters, as follows:

However, to test the inequality, we use `lincom` to define a confidence interval and a P–value for half the difference between the two z–transformed Somers' D parameters, as follows:

```
. lincom (para32g-ps_para32g)/2;

 ( 1)  .5 para32g - .5 ps_para32g = 0

------------------------------------------------------------------------------
    lnigetot |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
         (1) |   .0131897   .0063639     2.07   0.038     .0007167    .0256626
------------------------------------------------------------------------------
```

However, to test the inequality, we use `lincom` to define a confidence interval and a P–value for half the difference between the two z–transformed Somers' D parameters, as follows:

```
. lincom (para32g-ps_para32g)/2;

 ( 1)  .5 para32g - .5 ps_para32g = 0

------------------------------------------------------------------------------
    lnigetot |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
         (1) |   .0131897   .0063639     2.07   0.038     .0007167    .0256626
------------------------------------------------------------------------------
```

We see that the difference is (just) significantly positive. So the positive association between IgE and paracetamol exposure within paracetamol propensity *groups* is probably *not* due to a residual positive association of both variables with paracetamol propensity *score*.

IgE and prenatal paracetamol exposure: summary

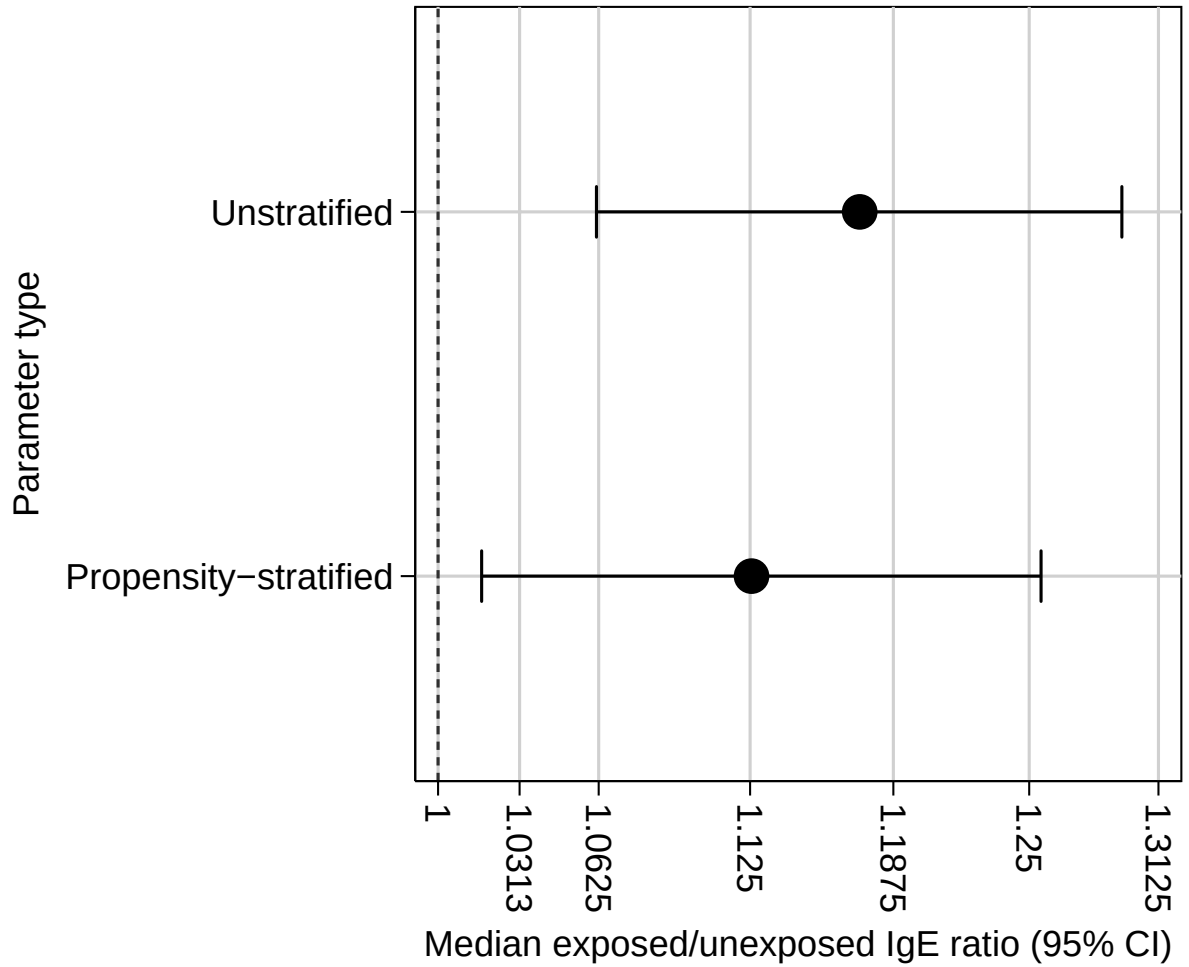

IgE and prenatal paracetamol exposure: summary

- A random exposed child typically has 6% to 29% more IgE than a random unexposed child.

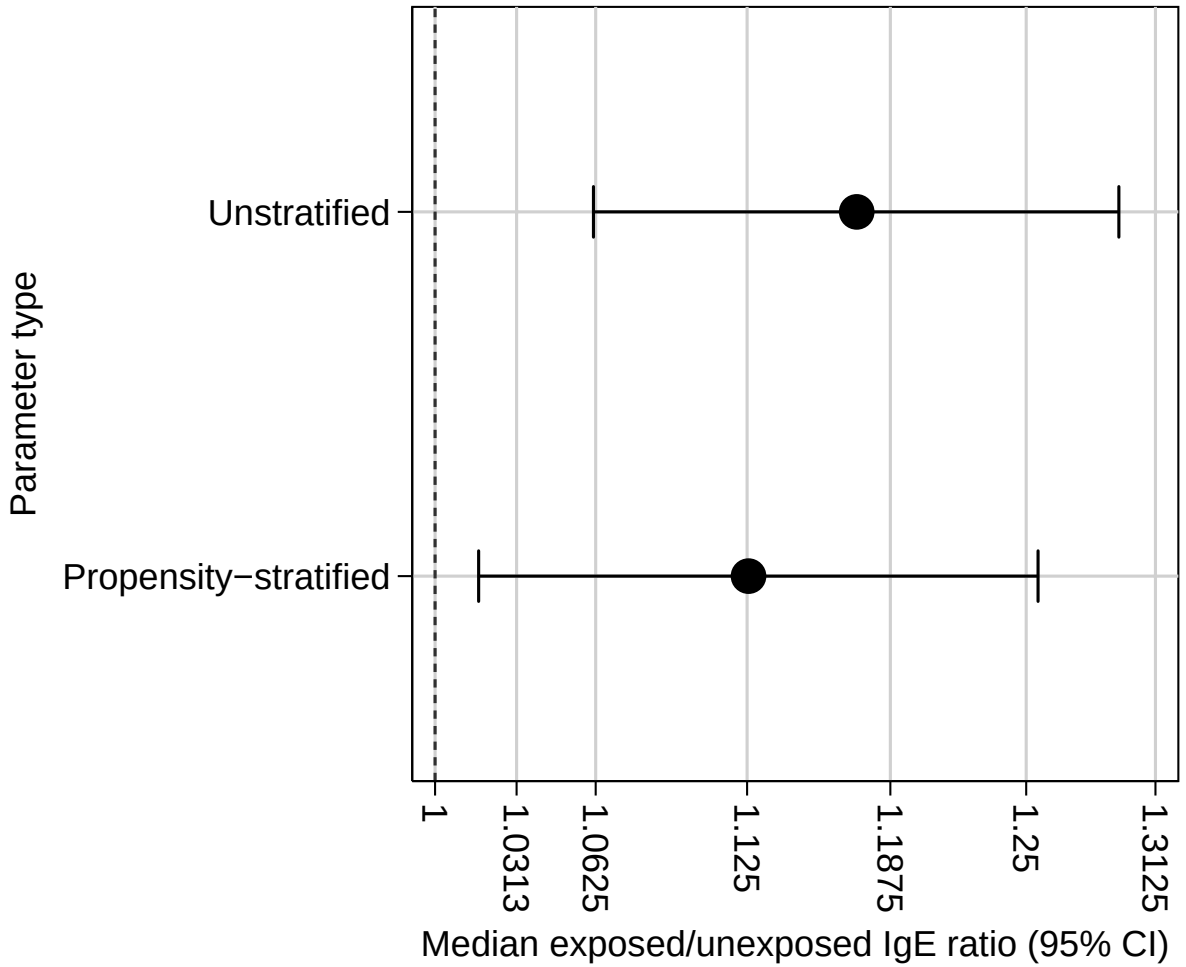

IgE and prenatal paracetamol exposure: summary

- A random exposed child typically has 6% to 29% more IgE than a random unexposed child.
- If they are in the same paracetamol propensity group, then the exposed child typically has 2% to 26% more IgE.

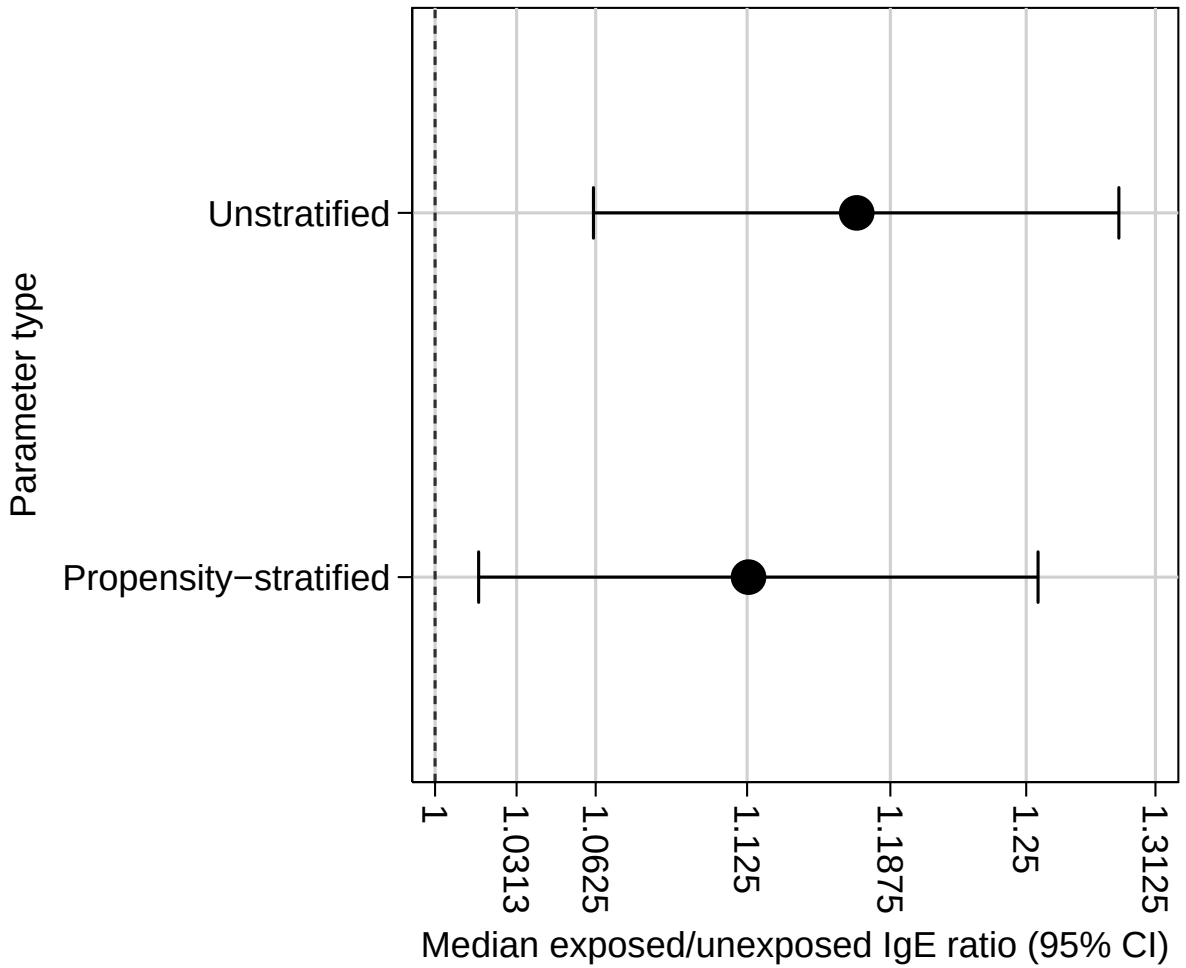

IgE and prenatal paracetamol exposure: summary

- A random exposed child typically has 6% to 29% more IgE than a random unexposed child.
- If they are in the same paracetamol propensity group, then the exposed child typically has 2% to 26% more IgE.
- This relative difference is probably *not* caused by paracetamol propensity (as defined here).

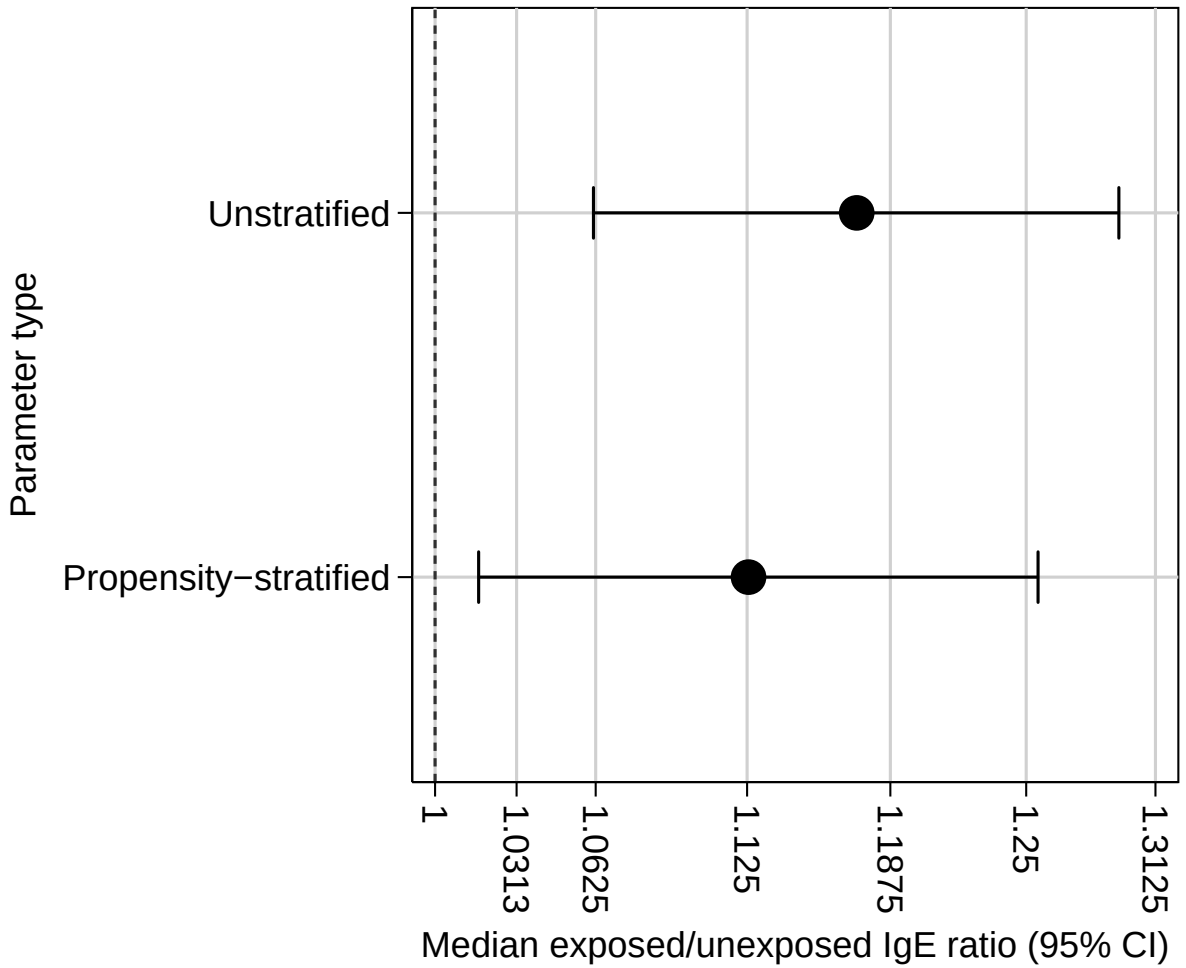

The case for rank methods

The case for rank methods

- Somers' D and Kendall's τ_a have "democratic" influence functions, based on the principle "one comparison, one vote".

The case for rank methods

- Somers' D and Kendall's τ_a have "democratic" influence functions, based on the principle "one comparison, one vote".
- This ensures that minorities of extreme values do not have too much influence.

The case for rank methods

- Somers' D and Kendall's τ_a have "democratic" influence functions, based on the principle "one comparison, one vote".
- This ensures that minorities of extreme values do not have too much influence.
- This in turn implies that the Central Limit Theorem *typically* works faster for rank parameters than for regression parameters.

The case for rank methods

- Somers' D and Kendall's τ_a have "democratic" influence functions, based on the principle "one comparison, one vote".
- This ensures that minorities of extreme values do not have too much influence.
- This in turn implies that the Central Limit Theorem *typically* works faster for rank parameters than for regression parameters.
- *Also*, rank parameters are often easier to interpret (as differences between proportions, or as median differences or ratios).

The case for rank methods

- Somers' D and Kendall's τ_a have "democratic" influence functions, based on the principle "one comparison, one vote".
- This ensures that minorities of extreme values do not have too much influence.
- This in turn implies that the Central Limit Theorem *typically* works faster for rank parameters than for regression parameters.
- *Also*, rank parameters are often easier to interpret (as differences between proportions, or as median differences or ratios).
- By contrast, an arithmetic mean difference is *usually* a proxy for a median difference, and *may* be expressed in incomprehensible units, such as a symptom score after a Normalizing transformation.

The case *against* rank methods

The case *against* rank methods

- *Some* people still think that they cannot produce confidence intervals.

The case *against* rank methods

- *Some* people still think that they cannot produce confidence intervals.
- More people think that they cannot be adjusted for confounding variables.

The case *against* rank methods

- *Some* people still think that they cannot produce confidence intervals.
- More people think that they cannot be adjusted for confounding variables.
- (They can, but we needed to use regression methods to define the propensity score.)

The case *against* rank methods

- *Some* people still think that they cannot produce confidence intervals.
- More people think that they cannot be adjusted for confounding variables.
- (They can, but we needed to use regression methods to define the propensity score.)
- A more valid argument is that of Fisher (1935), which implies that, *if* we know the distributional family *a priori*, *then* we can define narrower confidence intervals using maximum–likelihood methods than using rank methods.

The case *against* rank methods

- *Some* people still think that they cannot produce confidence intervals.
- More people think that they cannot be adjusted for confounding variables.
- (They can, but we needed to use regression methods to define the propensity score.)
- A more valid argument is that of Fisher (1935), which implies that, *if* we know the distributional family *a priori*, *then* we can define narrower confidence intervals using maximum–likelihood methods than using rank methods.
- *For instance*, using a t–test instead of `censlope` may reduce the minimum detectable difference by a modest 5%, when comparing 2 samples of 40. Or from infinity to a finite difference, when comparing 2 samples of 3.

Summary

Summary

- The `somersd` package computes confidence intervals for the "Somers' D family" of rank parameters.

Summary

- The `somersd` package computes confidence intervals for the "Somers' D family" of rank parameters.
- These confidence intervals are robust to distributional assumptions.

Summary

- The `somersd` package computes confidence intervals for the "Somers' D family" of rank parameters.
- These confidence intervals are robust to distributional assumptions.
- *However*, they are less robust to small sample numbers.

Summary

- The `somersd` package computes confidence intervals for the "Somers' D family" of rank parameters.
- These confidence intervals are robust to distributional assumptions.
- *However*, they are less robust to small sample numbers.
- More work is needed (and is in progress) to find more quantitative information about these tradeoffs.

Summary

- The `somersd` package computes confidence intervals for the "Somers' D family" of rank parameters.
- These confidence intervals are robust to distributional assumptions.
- *However*, they are less robust to small sample numbers.
- More work is needed (and is in progress) to find more quantitative information about these tradeoffs.
- *Meanwhile*, I would like to thank StataCorp for the Mata programming language, which made `somersd` possible in its present form.

References

Fisher, R. A. 1935. The logic of inductive inference. *Journal of the Royal Statistical Society* **98(1)**: 39–82.

Newson, R. 2002. Parameters behind "nonparametric" statistics: Kendall's tau, Somers' *D* and median differences. *The Stata Journal* **2(1)**: 45–64.

Newson, R. 2006. Confidence intervals for rank statistics: Somers' *D* and extensions. *The Stata Journal* **6(2)**: 309–334.

Rosenbaum, P. R. and D. B. Rubin. 1983. The central role of the propensity score in observational studies for causal effects. *Biometrika* **70(1)**: 41–55.

Shaheen, S. O., R. B. Newson, A. J. Henderson, J. E. Headley, F. D. Stratton, R. W. Jones, D. P. Strachan and the ALSPAC Study Team. 2005. Prenatal paracetamol exposure and risk of asthma and elevated immunoglobulin E in childhood. *Clinical and Experimental Allergy* **35**: 18–25.