



The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search

<http://ageconsearch.umn.edu>

aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

No endorsement of AgEcon Search or its fundraising activities by the author(s) of the following work or their employer(s) is intended or implied.

Multinomial goodness-of-fit: Large-sample tests with survey design correction and exact tests for small samples

Ben Jann
ETH Zürich
Zürich, Switzerland
ben.jann@soz.gess.ethz.ch

Abstract. I introduce the new `mgof` command to compute distributional tests for discrete (categorical, multinomial) variables. The command supports large-sample tests for complex survey designs and exact tests for small samples as well as classic large-sample χ^2 -approximation tests based on Pearson's X^2 , the likelihood ratio, or any other statistic from the power-divergence family (Cressie and Read, 1984, *Journal of the Royal Statistical Society, Series B (Methodological)* 46: 440–464). The complex survey correction is based on the approach by Rao and Scott (1981, *Journal of the American Statistical Association* 76: 221–230) and parallels the survey design correction used for independence tests in `svy: tabulate`. `mgof` computes the exact tests by using Monte Carlo methods or exhaustive enumeration. `mgof` also provides an exact one-sample Kolmogorov–Smirnov test for discrete data.

Keywords: st0142, `mgof`, `mgofi`, multinomial, goodness-of-fit, chi-squared, categorical data, exact tests, Monte Carlo, exhaustive enumeration, combinatorial algorithms, complex survey correction, power-divergence statistic, Kolmogorov–Smirnov, Benford's law

1 Introduction

A fundamental task in statistics is to test whether an observed distribution differs from a hypothetical null distribution, yet support for such tests is incomplete in the standard release of Stata. There are tools to test continuous distributions, such as normality tests ([R] `sktest`, [R] `swilk`) and a one-sample Kolmogorov–Smirnov test ([R] `ksmirnov`). However, distributional tests for discrete variables are missing.

This lack of tests for discrete variables might not seem to be an issue of serious concern because the classic multinomial goodness-of-fit test is a simple χ^2 test based on Pearson's X^2 , which is easy to compute from the values of a frequency table ([R] `tabulate`). Furthermore, several user implementations are available for this test (for example, [Weesie \[1997\]](#)). However, the classic test is only valid in simple random samples and cannot be used with sampling weights or other complex survey features (e.g., clusters and strata). Furthermore, the classic test is only asymptotic and may be biased in small samples or when the null distribution is very uneven.

I therefore present a new command, called **mgof**, that performs goodness-of-fit tests for discrete variables. The command supports complex survey designs and also offers exact tests for small samples. The complex survey tests are based on the approach by [Rao and Scott \(1981\)](#) and parallel the survey design correction used for independence tests in **tabulate** ([SVY] **svy: tabulate twoway**). The exact tests are computed by sampling from the null distribution (the Monte Carlo method) or by enumerating all possible data configurations (the exhaustive enumeration technique). Supported test statistics include Pearson's X^2 , the likelihood ratio, and any other statistic from the [Cressie and Read \(1984\)](#) family. The exact multinomial probability test and the exact discrete Kolmogorov–Smirnov test are also offered.

2 The mgof and mgofi commands

2.1 Syntax

```
mgof varname [=exp] [if] [in] [weight] [, options]
```

```
mgofi f1 f2 ... [/ p1 p2 ...] [, options]
```

<i>options</i>	description
Method 1	
<code>approx[(nfit)]</code>	compute large-sample χ^2 tests; the default
<code>* svy[(vcetype, svy_options)]</code>	adjust tests for survey design
<code>* vce(vcetype)</code>	adjust tests by using proportion variance estimate
<code>* cluster(varname)</code>	adjust tests for intragroup correlation
<code>noisily</code>	show output from proportion
Method 2	
<code>mc</code>	compute Monte Carlo exact tests
<code>reps(#)</code>	number of replications for mc ; default is <code>reps(10000)</code>
<code>level(#)</code>	set confidence level for mc ; default is <code>level(99)</code>
<code>cotype(type)</code>	set confidence-interval type for mc ; default is <code>cotype(exact)</code>
Method 3	
<code>ee</code>	compute exhaustive enumeration exact tests
Test statistics	
<code>nox2</code>	suppress Pearson's X^2 statistic
<code>nolr</code>	suppress likelihood-ratio statistic
<code>cr[(lambda)]</code>	include Cressie–Read statistic; <i>lambda</i> defaults to 2/3
<code>mlnp</code>	include outcome-probability statistic (mc and ee only)
<code>ksmirnov</code>	include Kolmogorov–Smirnov statistic (mc and ee only)

Other options

<u>freq</u>	display frequency table
<u>percent</u>	display frequency table in percentages
* <u>matrix</u> (<i>name</i>)	provide matrix of observed and expected counts
<u>expected</u> (<i>name</i>)	provide matrix (column vector) containing expected counts
<u>nodots</u>	suppress progress dots (mc and ee only)

* **svy**, **vce()**, **cluster()**, and **matrix()** are not allowed with **mgofi**.

by is allowed (unless **svy** is specified); see [D] **by**.

fweights, **pweights**, and **iweights** are allowed (see [U] **11.1.6 weight**), but **pweights** are not allowed with **ee** or **mc**, and **iweights** are not allowed with **ee** and not allowed with **mc** if the **mlnp** option is specified.

The null distribution (the distribution against which *varname* is tested) is specified by *exp*. *exp* is assumed to evaluate to the hypothesized probabilities of the categories of *varname* or to quantities proportional to these probabilities (e.g., expected counts; the scale does not matter). If *exp* is omitted, the uniform (geometric, equiprobable) distribution is used as the theoretical distribution.

mgofi is the immediate form of **mgof** ([U] **19 Immediate commands**) where *f1*, *f2*, etc., specify the observed counts and, optionally, *p1*, *p2*, etc., specify the theoretical probabilities or expected counts.

2.2 Options

Method 1 options

approx[(*nfit*)], the default method, computes classical large-sample χ^2 -approximation tests based on Pearson's X^2 and the likelihood-ratio statistic (see, for example, Horn [1977], Cressie and Read [1989], and Sokal and Rohlf [1995, chapter 17]). The degrees of freedom for χ^2 tests are determined as $k - nfit - 1$, where k is the number of categories, and *nfit*, provided by the user, indicates the number of fitted parameters (imposed restrictions) (*nfit* defaults to 0). If **pweights** are specified, the tests are corrected as outlined in section 4.

svy[(*vcetype*, *svy_options*)] specifies that the test results be adjusted for survey design effects according to the **svyset** specifications (see [SVY] **svyset**). *vcetype* and *svy_options* are as described in [SVY] **svy**. The correction procedure is described in section 4. The **svy** option is not allowed with **mgofi**.

vce(*vcetype*) specifies that the variance-covariance matrix of the proportions be estimated by using the **proportion** command (see [R] **proportion**) and that the tests be adjusted based on this estimate (see section 4). *vcetype* can be **analytic**, **cluster** *clustvar*, **bootstrap**, or **jackknife** (plus possible suboptions as described in [R] **vce_option**). **analytic** and **cluster** *clustvar* are not allowed in Stata 9. The **vce()** option is not allowed with **mgofi**.

`cluster(varname)` is equivalent to Stata 10's `vce(cluster clustvar)`. The `cluster()` option is not allowed with `mgofi`.

`noisily` displays the output from the `proportion` command, which is used to estimate the variances of the proportions if `svy`, `vce()`, or `cluster()` is specified or if `pweights` are applied.

Method 2 options

`mc` causes the exact p -values to be approximated by sampling from the null distribution (Monte Carlo simulation). The default number of replications for the simulation is 10,000; see the `reps()` option (the same set of samples is used for all test statistics). Confidence intervals of 99% are displayed for the estimated p -values.

`reps(#)` sets the number of replications for the `mc` method. The default is `reps(10000)`.

`level(#)` sets the level for the confidence intervals of the p -values computed by the `mc` method. The default is `level(99)`. Unlike many other Stata commands, `mgof` does not depend on `set level` (see [R] `level`).

`citype(type)` specifies how the binomial confidence intervals for the p -values from the `mc` method are to be calculated. Available types are `exact`, `wald`, `wilson`, `agresti`, and `jeffreys`. See [R] `ci`. `citype(exact)` is the default.

Method 3 option

`ee` causes the exact p -values to be computed by cycling through all possible data compositions given the sample size and the number of categories. Because the number of compositions grows very fast—it is equal to $(n + k - 1)! / ((k - 1)!n!)$, where n is the sample size and k is the number of categories—the `ee` method is only feasible for very small samples and few categories. An important exception is when the null distribution is uniform (and `ksmirnov` is not specified). In this case, the tests are based on enumerating partitions, which are much fewer in number than compositions.

Test statistics options

`nox2` suppresses Pearson's X^2 statistic.

`nolr` suppresses the likelihood-ratio statistic.

`cr[(lambda)]` specifies that the Cressie–Read statistic, with parameter $\lambda = \text{lambda}$, be included (Cressie and Read [1984]; also see Weesie [1997]). The default for `lambda` is $2/3$.

`mlnp` requests that a test based on the (minus log) multinomial probability of the observed outcome be included (see Horn [1977]). `mlnp` is not allowed with `approx`.

`ksmirnov` requests that the two-sided Kolmogorov–Smirnov statistic be included. The Kolmogorov–Smirnov statistic is sensitive to the order of the categories and should

be used only with variables that have a natural order (i.e., ordinal or discrete metric data). The Kolmogorov–Smirnov test implemented in Stata’s `ksmirnov` is conservative in the case of discrete data (see, for example, [Conover \[1972\]](#)). The methods implemented here are exact. `ksmirnov` is not allowed with `approx`.

Other options

`freq` displays a table containing observed and expected frequencies.

`percent` displays a table containing observed and expected percentages.

`matrix(name)` specifies that the observed and expected counts are to be taken from matrix *name* (see [P] `matrix`). The first column of the matrix provides the observed counts, and the second column, if present, provides the expected counts or theoretical probabilities. The uniform distribution is used if the matrix does not contain a second column. Do not provide noninteger observed counts with the `ee` or `mc` method. The `matrix()` option is not allowed with `mgofi`.

`expected(name)` specifies that the expected counts or theoretical probabilities are to be taken from column vector *name* (see [P] `matrix`). `mgof` aborts if the number of elements in the vector does not match the number of outcomes.

`nodots` suppresses the progress dots for the `ee` and `mc` methods. The default is to display a dot per 2% of completed computations.

2.3 Saved results

Scalars

<code>r(N)</code>	number of observations	<code>r(N_pop)</code>	population size
<code>r(N_strata)</code>	number of strata	<code>r(N_psu)</code>	number of primary sampling units (PSUs)
<code>r(N_clust)</code>	number of clusters	<code>r(df_r)</code>	design degrees of freedom
<code>r(df)</code>	degrees of freedom for χ^2	<code>r(df1)</code>	numerator df for F
<code>r(df2)</code>	denominator df for F	<code>r(delta)</code>	mean-generalized design effect
<code>r(a2)</code>	squared variation coefficient of generalized design effects	<code>r(reps)</code>	number of replications
<code>r(partitions)</code>	number of partitions	<code>r(compositions)</code>	number of compositions
<code>r(stat)</code>	value of test statistic	<code>r(F_stat)</code>	design-corrected F
<code>r(p_stat)</code>	(design-corrected) p -value	<code>r(p_stat_srs)</code>	uncorrected p -value
<code>r(p_stat_lb)</code>	lower CI bound for p -value	<code>r(p_stat_ub)</code>	upper CI bound for p -value

where *stat* is `x2` (Pearson’s X^2), `lr` (likelihood ratio), `cr` (Cressie–Read statistic), `mlnp` (minus log outcome-probability), or `ksmirnov` (Kolmogorov–Smirnov D)

Macros

<code>r(depvar)</code>	name of tabulated variable	<code>r(h0)</code>	definition of the theoretical distribution
<code>r(method)</code>	test method	<code>r(stats)</code>	list of test statistics
<code>r(lambda)</code>	Cressie–Read λ	<code>r(citype)</code>	Monte Carlo CI type
<code>r(cilevel)</code>	Monte Carlo confidence level		

Matrix

<code>r(count)</code>	observed and expected counts
-----------------------	------------------------------

3 Classic large-sample tests

The classic large-sample goodness-of-fit tests for discrete data are based on the result that statistics such as Pearson's X^2 ,

$$X^2 = \sum_{j=1}^k \frac{(f_j - h_j)^2}{h_j}$$

where f_j and h_j are the observed and expected (theoretical) counts for the categories $j = 1, \dots, k$, the likelihood-ratio statistic,

$$G^2 = 2 \sum_{j=1}^k f_j \ln \left(\frac{f_j}{h_j} \right)$$

or, generally, the power-divergence statistic ([Cressie and Read 1984](#)),

$$D^2(\lambda) = \frac{2}{\lambda(\lambda + 1)} \sum_{j=1}^k f_j \left\{ \left(\frac{f_j}{h_j} \right)^\lambda - 1 \right\}$$

are asymptotically $\chi^2(k - 1)$ distributed.¹ With significance level α given, an observed distribution is considered significantly different from the null distribution if the chosen test statistic exceeds the $(1 - \alpha)$ quantile of the χ^2 distribution with $(k - 1)$ degrees of freedom. For guidelines on choosing a test statistic, see the “Which Test Statistic?” sections in [Read and Cressie \(1988\)](#). One result, for example, is that in small samples the approximation of the χ^2 distribution is much better for Pearson's X^2 than for the likelihood ratio. Based on various simulations, [Read and Cressie \(1988\)](#) propose the $\lambda = 2/3$ power-divergence statistic as a good compromise in a wide range of situations. $\lambda = 2/3$ is the default for $D^2(\lambda)$ in `mgof`.

As an example, assume we are testing a sequence of numbers against Benford's law. The law states that under certain conditions the first digit of numbers in the base-10 system follows a probability distribution given as

$$\Pr(j) = \log_{10}(1 + 1/j), \quad j = 1, \dots, 9$$

([Newcomb 1881](#); [Benford 1938](#); [Hill 1998](#)). In a small Swiss mail survey, respondents were asked to indicate the first digit of the street number of an acquaintance.² The distribution of the indicated digits is astonishingly close to Benford's law, as the results from `mgof` reveal:

-
1. The Cressie–Read statistic is equal to Pearson's X^2 if $\lambda = 1$ and, as the limiting value, to the likelihood ratio if $\lambda = 0$. Further special cases are the Freeman–Tukey statistic with $\lambda = -1/2$, the Kullback–Leibler information with $\lambda = -1$, and Neyman's modified X^2 statistic with $\lambda = -2$ (see [Cressie and Read \[1984\]](#) and [Weesie \[1997\]](#)).
 2. The survey was conducted in December 2006 and January 2007 by the sociology department of the ETH Zürich. Respondents were sampled from the residents of the German-speaking part of Switzerland. The street number question was included toward the end of a very short questionnaire about income inequality; it had no relation to the other questions. The response rate of the survey was 41%.

```

. mata mata mlib index
.mlib libraries to be searched are now
    lmatabase;lmataado;lmataopt;lmgof;lfreduse
. use digits, clear
(2007 Swiss Street Number Data)
. mgof firstdigit = log10(1+1/firstdigit), cr percent

```

Number of obs = 313
 N of outcomes = 9
 Chi2 df = 8

Goodness-of-fit	Coef.	P-value
Pearson's X2	6.226606	0.6219
Log likelihood ratio	6.475677	0.5941
Cressie-Read (2/3)	6.303507	0.6133

firstdigit	observed	expected
1	32.59	30.10
2	17.57	17.61
3	14.70	12.49
4	10.86	9.69
5	6.39	7.92
6	6.07	6.69
7	4.47	5.80
8	4.15	5.12
9	3.19	4.58
Total	100.00	100.00

The p -values of the tests based on Pearson's X^2 , the Cressie–Read statistic, and the likelihood-ratio statistic suggest that the null hypothesis of Benford-distributed digits cannot be rejected.

4 Survey design correction for large-sample tests

The results of the standard χ^2 tests are only valid for simple random samples. In the case of nonidentical sampling probabilities, nonindependence, or stratification, the tests may be considerably biased. `mgof` therefore offers a survey design correction that is based on [Rao and Scott \(1981\)](#) and is analogous to the default independence-test correction used in `svy: tabulate twoway` (see [SVY] `svy: tabulate twoway` and the references therein). The procedure determines the “design effects” for the variances of the proportion estimates for the single outcomes and then corrects the χ^2 test statistic for the level and variation of these design effects. [Rao and Scott \(1981\)](#) call this a second-order correction; a first-order correction would ignore the variation in design effects. Finally, the corrected statistic is converted into an F statistic to adjust for the degrees of freedom of the employed variance estimates.

More precisely, let $\hat{V}/(n-1)$ be a consistent estimate of the variance–covariance matrix of the proportion estimates \hat{p}_j , $j = 1, \dots, k$, where n is the number of observa-

tions. Furthermore, let \hat{v}_{ij} denote an element of \hat{V} , let m be the number of PSUs or clusters, and let L be the number of strata. The correction then assumes

$$F = \frac{\chi^2}{\hat{\delta}(\hat{a}^2 + 1)} d^{-1} = \frac{\chi^2}{\hat{\delta}} (k - 1)^{-1}$$

to be $F(d, (m - L)d)$ distributed, where χ^2 stands for X^2 , G^2 , or $D^2(\lambda)$, and where

$$\hat{\delta} = (k - 1)^{-1} \sum_{i=1}^k \hat{v}_{ii} / \hat{p}_i, \quad \hat{a}^2 = (\hat{\delta}^2 - 1)^{-1} \left\{ (k - 1)^{-1} \sum_{i=1}^k \sum_{j=1}^k \hat{v}_{ij}^2 / (\hat{p}_i \hat{p}_j) \right\}$$

and $d = (k - 1)/(1 + \hat{a}^2)$. $\hat{\delta}$ is the mean, and \hat{a}^2 is the squared variation coefficient of the “generalized design effects” for the proportions (Rao and Scott 1981). `proportion` is used to estimate $\hat{V}/(n - 1)$, taking into account `pweights`, clusters, or other complex survey design settings (see [R] **proportion** for details).

Sribney ([1998]; also see Thomas, Singh, and Roberts [1996]) provides simulation evidence indicating that the F -based variant of the second-order Rao–Scott correction works well for independence tests in two-way contingency tables. Although it appears reasonable to assume that these results translate to goodness-of-fit tests in one-way tables, it seems important to perform at least a few brief checks. I therefore ran the following simulations.³

Simulation 1

Simulation 1 parallels the simulation reported by Sribney (1998). In each replication, a sample was initialized by drawing a number of cluster sizes from a uniform distribution and expanding the clusters to individual observations. Then two sets of variables with categorical values $j \in \{1, \dots, k\}$ and k varying between 2 and 9 were generated from an underlying continuous variable y , which was $N(0, 1)$ distributed and had an intraclass (cluster) correlation of 0.25 (see Sribney [1998] for details on how to generate such a variable). For the first set of categorical variables, y was categorized using normal quantiles with equally spaced probabilities, so that the variables had a geometric (uniform) distribution. For the second set, the cutpoints were chosen according to Benford’s first-digit law (see references in section 3), where the base of the number system was set to $k + 1$. Hence, the probabilities of the categories of the variables in the second set were given as

$$\Pr(j) = \log(1 + 1/j) / \log(k + 1), \quad j = 1, \dots, k$$

3. See Thomas and Rao (1987) and Rai, Srivastava, and Gupta (2001) for additional results.

For a variable with two categories, for example, the probabilities were (0.631, 0.369); for a variable with nine categories, they were (0.301, 0.176, ..., 0.046).

As in [Sribney \(1998\)](#), two types of simulations were conducted: one with small-variance degrees of freedom (few PSUs) and one with large-variance degrees of freedom (many PSUs). For the simulations with few PSUs, 20 clusters were generated with sizes between 30 and 70 observations, resulting in a sample size of approximately 1,000 observations. For the simulations with many PSUs, 200 clusters were generated with sizes between 3 and 10 observations, resulting in a sample size of approximately 1,200 observations.⁴ Ten thousand replications were computed for both types, and the nominal significance level was set to $\alpha = 0.05$. In comparison to the simulation by [Sribney \(1998\)](#), the clusters were regenerated in each replication.

Tests based on the first- and second-order Rao–Scott corrections (RS1 and RS2, respectively) according to the following definitions were evaluated:

$$\begin{aligned} X_{\text{RS1}}^2 &= X^2 / \hat{\delta} && \overset{a}{\sim} \chi^2(k-1) \\ F_{\text{RS1}} &= X_{\text{RS1}}^2 / (k-1) && \overset{a}{\sim} F\{k-1, (m-L)(k-1)\} \\ X_{\text{RS2}}^2 &= X_{\text{RS1}}^2 / (1 + \hat{a}^2) && \overset{a}{\sim} \chi^2(d) \\ F_{\text{RS2}} &= X_{\text{RS2}}^2 / d && \overset{a}{\sim} F\{d, (m-L)d\} \end{aligned}$$

F_{RS2} corresponds to the correction method implemented in `mgof` as outlined above. Furthermore, an adjusted Wald F test was considered, with

$$F_{\text{adj}}^W = W \frac{(m-L) - (k-1) + 1}{(m-L)(k-1)} \overset{a}{\sim} F\{k-1, (m-L) - (k-1) + 1\}$$

where

$$W = (\hat{p}^* - p^*)^T \{\hat{V}^* / (n-1)\}^{-1} (\hat{p}^* - p^*)$$

and where $\hat{p} = (\hat{p}_1, \dots, \hat{p}_k)$ and $p = (p_1, \dots, p_k)$ are the vectors of observed and expected probabilities, and the asterisk indicates that one of the categories is left out (see, for example, [Thomas and Rao \[1987\]](#)).⁵

The results of the simulations are depicted in figures 1 and 2. Displayed are the rejection rates for the various test statistics, that is, the proportion of times they rejected the (true) null hypothesis at the $\alpha = 0.05$ level. In the case of small-variance degrees

4. The number of 1,300 given in [Sribney \(1998\)](#) seems to be incorrect.

5. Rao–Scott-corrected tests based on the likelihood-ratio statistic were also conducted. Results were similar to the tests based on Pearson's X^2 .

of freedom (figure 1), both of the first-order Rao–Scott-corrected tests (RS1) and the adjusted Wald F test are highly anticonservative (i.e., rejecting the null hypothesis too often) as the number of categories increases. The second-order Rao–Scott-corrected X^2 test is also anticonservative, but the bias does not depend on the number of categories. For the second-order Rao–Scott-corrected F test, however, the simulated rejection rates match the nominal 5% very well.




Figure 1. Rejection rates in small-variance degrees of freedom simulations

In the case of large-variance degrees of freedom (figure 2), both of the first-order Rao–Scott-corrected tests and the adjusted Wald F test are still slightly anticonservative. The second-order Rao–Scott correction again performs well, at least for the uniform variables. With the Benford-distributed variables, the second-order Rao–Scott correction seems to be slightly conservative.⁶

6. In both simulations, the uncorrected tests were highly anticonservative (51–57% and 11–16% rejection, respectively).




Figure 2. Rejection rates in large-variance degrees of freedom simulations

Simulation 2

Simulation 1 is model-based and somewhat artificial. To get an impression of the performance of the Rao–Scott correction in a more realistic setting, I conducted a second simulation based on sampling from a “real” population, the Swiss 2000 Census.⁷

The sampling plan for the simulation was as follows: First, a total of 100 municipalities was sampled proportional to size (with replacement) from 15 strata (the Swiss cantons, with some aggregations of small cantons). Second, 10 households were drawn with replacement from each sampled municipality, and one individual was selected from each sampled household. The a priori sampling probabilities for the individuals varied depending on household size and because of moderate oversampling of small strata. The `gsample` user-written command was used to draw the samples (Jann 2006).

A number of categorical variables were tested with the true population distribution as the null hypothesis. The variables included sex (1 = male, 2 = female), nationality (Swiss, foreign), marital status (single, married, divorced, widowed), education (6 levels), and socioeconomic status (12 levels). The distributions, representing a broad mix of patterns, are listed in table 1.

7. The Swiss Census was conducted by the Swiss Federal Statistical Office and covers all residents of Switzerland in December 2000. For the purpose of the simulation, I restricted the population to individuals age 15 or older. The population then consisted of 6,043,350 individuals, 3,179,246 households, and 2,896 municipalities. The strata sizes varied between 200,000 and 1,000,000 individuals.

Table 1. Population distributions (in percent)

	1	2	3	4	5	6	7	8	9	10	11	12
Sex	48.5	51.5										
Nationality	80.2	19.8										
Marital status	30.1	56.3	6.9	6.8								
Education	13.0	25.6	36.7	7.9	9.1	7.6						
Socioeconomic status	2.1	5.6	5.1	10.1	12.0	4.9	7.4	15.5	7.8	14.2	10.8	4.5

The rejection rates from 5,000 replications with a nominal $\alpha = 0.05$ significance level are displayed in table 2. The test statistics are the same as in simulation 1, namely, first- and second-order Rao–Scott corrections of Pearson’s X^2 , with and without F conversion, and the adjusted Wald F .⁸ In addition, rejection rates for Pearson’s X^2 without correction (X_{SRS}^2) are reported. As is evident from the rejection rates in the last column in table 2 (up to 40%), the uncorrected tests yield intolerably anticonservative results in such a sample. Among the considered correction procedures, the second-order Rao–Scott F performed best. The rejection rates for F_{RS2} closely match the nominal 5% for sex, education, and socioeconomic status, but they are somewhat anticonservative for nationality and marital status, possibly because of the strongly skewed distributions of the two variables. The degree of anticonservatism, however, is not dramatic and is smallest among all evaluated procedures. As in simulation 1, the adjusted Wald F is considerably anticonservative in some cases.

Table 2. Rejection rates (in percent; based on 5,000 replications)

	X_{RS1}^2	F_{RS1}	X_{RS2}^2	F_{RS2}	F_{adj}^W	X_{SRS}^2
Sex	4.9	4.6	4.9	4.6	4.7	16.6
Nationality	7.3	6.9	7.3	6.9	8.3	23.3
Marital status	9.3	9.0	7.1	6.7	13.8	29.1
Education	7.1	6.8	5.0	4.8	8.5	32.8
Socioeconomic status	10.7	10.4	5.8	5.6	12.0	41.2

Overall, the results from simulations 1 and 2 suggest that the second-order Rao–Scott-corrected F test outperforms the other considered tests and is a good default choice.

Example

In the example in section 3, the digit distribution of house numbers of acquaintances was analyzed. The survey from which the data was taken is based on a simple ran-

8. Results were similar for the Rao–Scott corrections based on the likelihood-ratio statistic.

dom sample of households. Because only one adult was selected per household, the individual-level sampling probabilities depended on the household size: individuals from larger households had smaller sampling probabilities. Therefore one should apply probability weights (inverse to the number of adult household members) to make the data representative of the population of individuals instead of the population of households. Because the standard tests would not be valid, `mgof` applies the complex survey design correction if `pweights` are specified. In addition to the uncorrected test statistics, corrected F statistics and associated p -values are reported. The results from `mgof`, applied to the weighted example data, are as follows:

```
. use digits, clear
(2007 Swiss Street Number Data)
. mgof firstdigit = log10(1+1/firstdigit) [pw=w]
```

Number of obs =	313
N of outcomes =	9
F df1	= 7.93123
F df2	= 2474.54

Goodness-of-fit	Coef.	F-value	P-value
Pearson's X2	6.144661	0.6435	0.7402
Log likelihood ratio	6.425913	0.6730	0.7145

Again there is virtually no evidence against the null hypothesis that the data are distributed according to Benford's law.

A more general syntax to obtain results with design correction is to set the survey properties by using `svyset` (see [SVY] `svyset`) and then apply the `svy` option in `mgof`, as in the following example:⁹

```
. svyset [pw=w]
      pweight: w
      VCE: linearized
      Single unit: missing
      Strata 1: <one>
      SU 1: <observations>
      FPC 1: <zero>
. mgof firstdigit = log10(1+1/firstdigit), svy
```

Number of strata =	1	Number of obs =	313
Number of PSUs =	313	Pop size =	583
		Design df =	312
		N of outcomes =	9
		F df1	= 7.93123
		F df2	= 2474.54

Goodness-of-fit	Coef.	F-value	P-value
Pearson's X2	6.144661	0.6435	0.7402
Log likelihood ratio	6.425913	0.6730	0.7145

9. Because of some technical difficulties, `svy` cannot be used as a prefix command (see [U] 11.1.10 **Prefix commands**) with `mgof`.

5 Exact small-sample tests

The χ^2 distribution of statistics such as Pearson's X^2 is only asymptotic, and the p -values of the standard χ^2 goodness-of-fit tests may be biased when the sample is very small or the null distribution is highly uneven. In such cases, it is desirable to compute the exact p -values.

Exhaustive enumeration

The most straightforward approach to compute an exact goodness-of-fit test is to construct all possible data combinations given the number of observations and the number of categories and then to sum up the probabilities of all configurations that are at least as distant from the null hypothesis as the observed data (for example, [Radlow and Alf, Jr. \[1975\]](#)). The steps of the procedure can be summarized as follows:

1. Calculate the probability of each data configuration, given the null hypothesis. Under simple random sampling, the probability of a specific configuration $f = (f_1, \dots, f_k)$ given expected cell probabilities $p = (p_1, \dots, p_k)$ (the null hypothesis) is

$$\Pr(f | p) = \frac{n!}{(f_1! f_2! \dots f_k!)} \cdot p_1^{f_1} p_2^{f_2} \dots p_k^{f_k}$$

where $n = \sum f$.

2. Calculate the value of the test statistic—say, Pearson's X^2 or the likelihood-ratio statistic, G^2 —for each data configuration, given the null hypothesis.
3. Compute the exact p -value as the sum of the probabilities of all configurations in which the test statistic is at least as large as in the observed data.

A natural variant to using Pearson's X^2 or the likelihood ratio as the test criterion is to directly use $\Pr(f | p)$ to determine whether a configuration adds to the p -value or not (see, for example, [Horn \[1977\]](#)). `mgof` offers such a multinomial probability test via the `mlnp` option, where the reported outcome-probability statistic is parameterized as $-\ln(\Pr(f | p))$ so that large values indicate departure from the null distribution and the scaling is similar to a χ^2 statistic. Little evidence exists for whether directly using the multinomial probability as the fit statistic is superior to using a statistic such as Pearson's X^2 , but both procedures yield “exact” p -values. The only difference is in how discrepancies between the null hypothesis and the data are valued. While statistics such as Pearson's X^2 are defined in terms of differences (or ratios) between expected and observed counts, the direct approach uses the multinomial probability as the measure of discrepancy. In general, different fit statistics give weight to different patterns of deviations from the null distribution, and the usefulness of a specific statistic may depend on the situation ([Read and Cressie 1988](#), 136–137).

The method outlined above—whichever fit statistic is employed—is called the “exhaustive enumeration” method because all possible data configurations are enumerated. It can be implemented by using algorithms to construct k -part compositions of n

(Reingold, Nievergelt, and Deo 1977; Nijenhuis and Wilf 1978). For example, for $n = 3$ and $k = 2$, the possible compositions are as follows:¹⁰

```
. mata
----- mata (type end to exit) -----
: f = mm_compositions(3,2)
: f
      1   2   3   4
1  

|   |   |   |   |
|---|---|---|---|
| 3 | 2 | 1 | 0 |
| 0 | 1 | 2 | 3 |


2
: end
-----
```

Using the sex distribution in a sample with three individuals as an example, possible sample compositions are the following: 3 females and 0 males, 2 females and 1 male, 1 female and 2 males, or 0 females and 3 males. Assume that the second outcome, 2 females and 1 male, is the observed outcome. To compute the p -value of the test that the ratio of females to males is, say, 1 to 2, one would determine for each composition Pearson's X^2 and the probability of the composition given the null hypothesis, and then sum up the probabilities of all compositions for which X^2 is at least as large as the observed X^2 . For example,¹¹

```
. mata
----- mata (type end to exit) -----
: h = (1, 2)'
: x2 = colsum((f :- h):^2 ./ h)
: x2
      1   2   3   4
1  

|   |     |   |     |
|---|-----|---|-----|
| 6 | 1.5 | 0 | 1.5 |
|---|-----|---|-----|


: n = 3
: pr = exp(lnfactorial(n) :- colsum(lnfactorial(f)) :+ colsum(f :* ln(h ./ n)))
: pr
      1           2           3           4
1  

|            |            |            |            |
|------------|------------|------------|------------|
| .037037037 | .222222222 | .444444444 | .296296296 |
|------------|------------|------------|------------|


: pvalue = sum(pr :* (x2 >= x2[2]))
: pvalue
      .555555556
: end
-----
```

10. The `mm_compositions()` Mata function is part of the `moremata` package, which provides a number of combinatorial algorithms (among many other functions; Jann [2005]).

11. The logarithmic form of $\Pr(f | p)$ is used here to turn products into sums and thus make computations easier. The logarithmic form is computationally more robust because `lnfactorial()` can be evaluated for much larger numbers than `factorial()` can.

The p -value is 0.556 and, hence, we cannot reject the null hypothesis (which is not surprising given only three observations). The same result can be obtained by `mgofi` as follows:

```
. mgofi 2 1 / 1 2, ee nodots
                        Number of obs =      3
                        N of outcomes =      2
                        Compositions  =      4
```

Goodness-of-fit	Coef.	Exact P-value
Pearson's X2	1.5	0.5556
Log likelihood ratio	1.386294	0.5556

The number of possible data configurations increases rapidly with additional observations and categories (the combinatorial explosion), which imposes restrictions for the application of the exhaustive enumeration method. The formula for the number of k -fold compositions of n is $(n+k-1)!/((k-1)!n!)$. For example, with $n = 50$ and $k = 5$, the problem size is 316,251, which can be handled (`mgof` takes about 7 seconds on my computer for a problem of this size). If k is increased to 6, the number of compositions becomes 3,478,761, taking about 80 seconds to compute. With $k = 7$, the number further increases to 32,468,436, and with $k = 10$, it is 12,565,671,261 (which would take three to four days to compute).

The examples suggest that the exhaustive enumeration method is only feasible for very small problems. However, in the case of a uniform null distribution the amount of computations can be reduced a great deal because of redundancies among the compositions. If all elements of h (and hence of p) are equal, then the order of the elements in a composition does not affect the value of $\Pr(f|p)$ or X^2 . In the toy example above, for instance, the compositions $\{2, 1\}$ and $\{1, 2\}$ would yield identical values for $\Pr(f|p)$ and X^2 if p is uniform, as would $\{3, 0\}$ and $\{0, 3\}$. Hence, the exact p -values can be computed by enumerating only compositions with unique sets of elements. These compositions are equivalent to the (zero-padded) integer partitions of n into k or fewer addends ([Andrews 1984](#); [Andrews and Eriksson 2004](#)). For example, in the case of $n = 4$ and $k = 3$, possible partitions are as follows:

```
. mata
----- mata (type end to exit) -----
: f = mm_partitions(4,3)
: f
      1   2   3   4
1  [ 4   3   2   2 ]
2  [ 0   1   2   1 ]
3  [ 0   0   0   1 ]

: end
-----
```

As is immediately clear, the number of partitions grows much slower with n and k than the number of compositions. For example, for $n = 50$ and $k = 10$, as above, the number of partitions is only 62,740, compared to 12.6 billion compositions.¹² In the case of a uniform null distribution, it is therefore much more efficient to compute the p -value based on partitions, where each partition is weighted by the number of possible permutations (see, for example, Hirji [1997]).¹³ Table 3 gives an overview of the approximate computing times on a standard personal computer for different problem sizes depending on whether the null distribution is uniform.

Table 3. Approximate computing times for the exhaustive enumeration algorithm depending on sample size n , number of categories k , and type of null hypothesis

Nonuniform null distribution								Uniform null distribution							
$n \backslash k$	3	4	5	6	8	10	15	$n \backslash k$	3	4	5	6	8	10	15
10	<1s	<1s	<1s	<1s	<1s	2s	4s	50	<1s	<1s	<1s	<1s	2s	3s	7s
20	<1s	<1s	<1s	1s	20s	4m	9h	75	<1s	<1s	<1s	3s	14s	42m	3m
30	<1s	<1s	1s	7s	4m	1h	>7d	100	<1s	<1s	2s	9s	1m	5m	36m
40	<1s	<1s	3s	27s	23m	13h	>7d	150	<1s	1s	11s	1m	15m	2h	1d
50	<1s	<1s	7s	1m	2h	3d	>7d	200	<1s	3s	32s	4m	2h	16h	>7d
75	<1s	2s	33s	9m	23h	>7d	>7d	300	<1s	10s	3m	28m	22h	>7d	>7d
100	<1s	4s	2m	36m	7d	>7d	>7d	500	1s	45s	19m	6h	>7d	>7d	>7d
150	<1s	13s	8m	4h	>7d	>7d	>7d	1000	4s	6m	5m	7d	>7d	>7d	>7d

Notes: Assuming a base speed of 45,000 compositions and 20,000 partitions per second; s = seconds, m = minutes, h = hours, d = days.

In the example in section 3, there were 313 observations and 9 categories. This is too large a problem for the exhaustive enumeration method ($2.56 \cdot 10^{15}$ compositions or $1.09 \cdot 10^{10}$ partitions, respectively). To illustrate, I apply the exhaustive enumeration tests by using a 5% sample of the data:¹⁴

12. There is no simple formula for the number of partitions as there is for the number of compositions. However, the number of partitions can be computed by using the `mm.npartitions()` function from the `moremata` package (Jann 2005).
13. The number of permutations equals $k!/(r_1!, \dots, r_i!)$, where r_i denotes the number of repetitions of the i th distinct addend in the partition. The algorithm used by `mgof` to enumerate the partitions is based on Algorithm ZS1 by Zoghbi and Stojmenovic (1998), with some modifications to generate restricted partitions.
14. The `gsample` command is provided by Jann (2006).

```
. set seed 12
. use digits, clear
(2007 Swiss Street Number Data)
. gsample 5, percent generate(freq)
. mgof firstdigit = log10(1+1/firstdigit) [fw=freq], ee
Percent completed (735471 compositions)
0 ----- 20 ----- 40 ----- 60 ----- 80 ----- 100
.....
                        Number of obs =      16
                        N of outcomes =       9
                        Compositions  =  735471
```

Goodness-of-fit	Coef.	Exact P-value
Pearson's X2	7.902401	0.4207
Log likelihood ratio	9.157266	0.4564

```
. mgof firstdigit [fw=freq], ee
Percent completed (201 partitions)
0 ----- 20 ----- 40 ----- 60 ----- 80 ----- 100
.....
                        Number of obs =      16
                        N of outcomes =       9
                        Partitions    =   201
```

Goodness-of-fit	Coef.	Exact P-value
Pearson's X2	15.5	0.0557
Log likelihood ratio	18.13734	0.0348

The results indicate that the data in the sample are closer to Benford's law than they are to the uniform distribution. Especially for the likelihood ratio, the exact p -values can be considerably different from the approximate p -values in such small samples. The approximate p -values are 0.443 and 0.329 for the tests against Benford's law and 0.050 and 0.020 for the equal-probability case.

□ Technical note

In the example above, the sample was marked by using frequency weights instead of by constructing a dataset containing the sampled observations. The rationale behind this was to ensure that all original categories remained visible to `mgof` even if some of them were missing in the sample. In general, a goodness-of-fit test will be biased if categories with theoretical probabilities greater than zero are omitted from the test because of lack of corresponding observations. Such "missing" categories can be introduced to `mgof` by adding extra observations to the dataset and assigning zero weights to them. Alternatively, zero observed counts can be specified by using the immediate `mgofi` or the `matrix()` option.

Monte Carlo method

Many problems are too large for exhaustive enumeration, but one might still want to compute the exact p -values. An approach that is easy to implement and can handle larger problems in reasonable time is to approximate the exact p -values by Monte Carlo simulation. The approach is made by sampling from the null distribution¹⁵ and computing the p -value as the fraction of replications in which the test statistic is at least as large as in the observed data. A drawback of the procedure is that the computed p -values are subject to random variation (which, however, can be made arbitrarily small by increasing the number of replications). Ninety-nine percent confidence intervals (computed by using `cii`; see [R] `ci`) are therefore reported by `mgof`.

With the Benford example from section 3, the Monte Carlo method can simulate the exact p -values in a few seconds:

```
. use digits, clear
(2007 Swiss Street Number Data)
. mgof firstdigit = log10(1+1/firstdigit), mc
Percent completed (10000 replications)
0 — 20 — 40 — 60 — 80 — 100
.....
                                Number of obs =    313
                                N of outcomes =     9
                                Replications =  10000
```

Goodness-of-fit	Coef.	Exact P-value	[99% Conf. Interval]	
Pearson's X2	6.226606	0.6147	0.6021	0.6272
Log likelihood ratio	6.475677	0.5921	0.5793	0.6048

As can be expected for a sample of this size, the approximate p -values (0.6219 and 0.5941; see section 3) are very close to the simulated exact p -values and lie within the computed 99% confidence limits.

Discrete Kolmogorov–Smirnov test

In addition to the multinomial tests based on, say, Pearson's X^2 or the likelihood ratio, `mgof` also offers an exact (two-sided) Kolmogorov–Smirnov test for ordered discrete data. The Kolmogorov–Smirnov test has higher power than the multinomial tests but is only appropriate if the categories have a natural order.

The (two-sided) Kolmogorov–Smirnov test statistic is defined as the supremum (least upper bound) of the absolute difference between the theoretical and the empirical distribution function. In the discrete case, the statistic can be expressed as

15. Sampling from a null distribution is equivalent to sampling n units with replacement from a population of k elements, where the sampling weights for the elements correspond to the theoretical probabilities of the categories.

$$D = \max_j |H(j) - F(j)|, \quad j = 1, \dots, k$$

with

$$H(j) = \frac{1}{n} \sum_{i=1}^j h_i \quad \text{and} \quad F(j) = \frac{1}{n} \sum_{i=1}^j f_i$$

where h_j and f_j denote the expected and observed counts for category j (also see, for example, [Wood and Altavela \[1978\]](#)).

The distribution of the Kolmogorov–Smirnov statistic for continuous data is well known (see [R] **ksmirnov**) but does not hold for discrete data (for example, [Conover \[1972\]](#)). **mgof** therefore performs the discrete Kolmogorov–Smirnov test without making assumptions about the distribution of D using Monte Carlo simulation or exhaustive enumeration.

In the Benford example in section 3, an argument could be put forward that the Kolmogorov–Smirnov test is more appropriate than the multinomial tests because the digits do have a natural order. Indeed, the p -value from the Kolmogorov–Smirnov test is considerably lower than the p -values based on Pearson’s X^2 or the likelihood ratio:

```
. mgof firstdigit = log10(1+1/firstdigit), mc ksmirnov
Percent completed (10000 replications)
0 _____ 20 _____ 40 _____ 60 _____ 80 _____ 100
.....
                                Number of obs =      313
                                N of outcomes =       9
                                Replications =    10000
```

Goodness-of-fit	Coef.	Exact P-value	[99% Conf. Interval]	
Pearson’s X2	6.226606	0.6311	0.6186	0.6435
Log likelihood ratio	6.475677	0.6093	0.5966	0.6219
Kolmogorov-Smirnov D	.0582185	0.0967	0.0892	0.1046

6 Concluding remarks

In this article, I introduced a new and flexible command for multinomial goodness-of-fit tests. The main features of the command are that it can be used with complex survey designs and that it offers methods to determine exact p -values in small samples.

The second-order Rao–Scott correction, which is used by the command to account for survey design, is an improvement over performing uncorrected tests or using the Wald statistic, as I illustrated in the simulations in section 4. However, there is also some

evidence that the correction is not always optimal. For example, [Magnussen and Köhl \(2006\)](#) present a study in which the second-order Rao–Scott correction did not perform as well as certain other procedures in the context of single-stage cluster sampling. A comprehensive simulation in which different procedures are systematically evaluated under various survey designs and data structures would be valuable.

Further, computational speed is a major concern when computing exact p -values. The exhaustive enumeration method is slow or even unfeasible unless the sample is very small. Some speed gains could be made if the underlying combinatorial algorithms, which are currently implemented in Mata, were ported to C, although this would not greatly increase the range of feasible applications. A more promising approach would be to implement fast algorithms for exact p -values, extending the work of [Mehta and Patel \(1983\)](#); [Baglivo, Olivier, and Pagano \(1992\)](#); or [Hirji \(1997, 2006\)](#). However, the returns on this appeared limited to me given the availability of the Monte Carlo approximation method and given that, for Pearson’s X^2 or Cressie–Read’s $D^2(\lambda = 2/3)$ (but not for G^2), the χ^2 approximation is usually quite good, even with relatively small samples (for example, [Read \[1984\]](#); for a somewhat more conservative view, see [Formann \[1995\]](#)). An exception may be if several very small p -values have to be estimated with great accuracy ([Keich and Nagarajan 2006](#)). Although currently not planned, a future extension of the command to cover such fast algorithms might therefore be desirable.

7 Acknowledgments

I would like to thank Bill Gould, Debra Hevenstone, and two anonymous reviewers for their helpful comments.

This research was supported by ETH Research Grant Number TH-8/04-3 on “Empirical studies in scientific misconduct: Prevalence and methods for detecting fraud and errors using the Benford distribution”.

8 References

- Andrews, G. E. 1984. *The Theory of Partitions*. Cambridge: Cambridge University Press.
- Andrews, G. E., and K. Eriksson. 2004. *Integer Partitions*. 2nd ed. Cambridge: Cambridge University Press.
- Baglivo, J., D. Olivier, and M. Pagano. 1992. Methods for exact goodness-of-fit tests. *Journal of the American Statistical Association* 87: 464–469.
- Benford, F. 1938. The law of anomalous numbers. *Proceedings of the American Philosophical Society* 78: 551–572.
- Conover, W. J. 1972. A Kolmogorov goodness-of-fit test for discontinuous distributions. *Journal of the American Statistical Association* 67: 591–596.

- Cressie, N., and T. R. C. Read. 1984. Multinomial goodness-of-fit tests. *Journal of the Royal Statistical Society. Series B (Methodological)* 46: 440–464.
- . 1989. Pearson's X^2 and the loglikelihood ratio statistic G^2 : A comparative review. *International Statistical Review* 57: 19–43.
- Formann, A. K. 1995. Small-sample comparison of the exact and asymptotic upper tail probabilities of chi-squared goodness-of-fit statistics: Pearson's X^2 , likelihood ratio, and power-divergence statistic ($\lambda = 2/3$). *Journal of Statistical Computation and Simulation* 51: 369–384.
- Hill, T. P. 1998. The first digit phenomenon. *American Scientist* 86: 358.
- Hirji, K. F. 1997. A comparison of algorithms for exact goodness-of-fit tests for multinomial data. *Communications in Statistics—Simulation and Computation* 26: 1197–1227.
- . 2006. *Exact Analysis of Discrete Data*. Boca Raton, FL: Chapman & Hall/CRC.
- Horn, S. D. 1977. Goodness-of-fit tests for discrete data: A review and an application to a health impairment scale. *Biometrics* 33: 237–247.
- Jann, B. 2005. moremata: Stata module (Mata) to provide various Mata functions. Statistical Software Components S455001, Boston College Department of Economics. Downloadable from <http://ideas.repec.org/c/boc/bocode/s455001.html>.
- . 2006. gsample: Stata module to draw a random sample. Statistical Software Components S456716, Boston College Department of Economics. Downloadable from <http://ideas.repec.org/c/boc/bocode/s456716.html>.
- Keich, U., and N. Nagarajan. 2006. A fast and numerically robust method for exact multinomial goodness-of-fit test. *Journal of Computational & Graphical Statistics* 15: 779–802.
- Magnussen, S., and M. Köhl. 2006. A better alternative to Wald's test-statistic for simple goodness-of-fit tests under one-stage cluster sampling. *Forest Ecology and Management* 221: 123–132.
- Mehta, C. R., and N. R. Patel. 1983. A network algorithm for performing Fisher's exact test in $r \times c$ contingency tables. *Journal of the American Statistical Association* 78: 427–434.
- Newcomb, S. 1881. Note on the frequency of use of the different digits in natural numbers. *American Journal of Mathematics* 4: 39–40.
- Nijenhuis, A., and H. S. Wilf. 1978. *Combinatorial Algorithms for Computers and Calculators*. 2nd ed. New York: Academic Press.
- Radlow, R., and E. F. Alf, Jr. 1975. An alternate multinomial assessment of the accuracy of the χ^2 test of goodness of fit. *Journal of the American Statistical Association* 70: 811–813.

- Rai, A., A. K. Srivastava, and H. C. Gupta. 2001. Small sample comparison of modified chi-square test statistics for survey data. *Biometrical Journal* 43: 483–495.
- Rao, J. N. K., and A. J. Scott. 1981. The analysis of categorical data from complex sample surveys: Chi-squared tests for goodness of fit and independence in two-way tables. *Journal of the American Statistical Association* 76: 221–230.
- Read, T. R. C. 1984. Small-sample comparisons for the power divergence goodness-of-fit statistics. *Journal of the American Statistical Association* 79: 929–935.
- Read, T. R. C., and N. A. C. Cressie. 1988. *Goodness-of-Fit Statistics for Discrete Multivariate Data*. New York: Springer.
- Reingold, E. M., J. Nievergelt, and N. Deo. 1977. *Combinatorial Algorithms: Theory and Practice*. Englewood Cliffs, NJ: Prentice Hall.
- Sokal, R. R., and F. J. Rohlf. 1995. *Biometry*. 3rd ed. New York: Freeman.
- Sribney, W. M. 1998. svy7: Two-way contingency tables for survey or clustered data. *Stata Technical Bulletin* 45: 33–49. Reprinted in *Stata Technical Bulletin Reprints*, vol. 8, pp. 297–322. College Station, TX: Stata Press.
- Thomas, D. R., and J. N. K. Rao. 1987. Small-sample comparisons of level and power for simple goodness-of-fit statistics under cluster sampling. *Journal of the American Statistical Association* 82: 630–636.
- Thomas, D. R., A. C. Singh, and G. R. Roberts. 1996. Tests of independence on two-way tables under cluster sampling: An evaluation. *International Statistical Review* 64: 295–311.
- Weesie, J. 1997. sg68: Goodness-of-fit statistics for multinomial distributions. *Stata Technical Bulletin* 36: 26–28. Reprinted in *Stata Technical Bulletin Reprints*, vol. 6, pp. 183–186. College Station, TX: Stata Press.
- Wood, C. L., and M. M. Altavela. 1978. Large-sample results for Kolmogorov–Smirnov statistics for discrete distributions. *Biometrika* 65: 235–239.
- Zoghbi, A., and I. Stojmenovic. 1998. Fast algorithms for generating integer partitions. *International Journal of Computer Mathematics* 70: 319–332.

About the author

Ben Jann is a sociologist at ETH Zürich in Zürich, Switzerland.