

THE STATA JOURNAL

Editors

H. JOSEPH NEWTON
Department of Statistics
Texas A&M University
College Station, Texas
editors@stata-journal.com

NICHOLAS J. COX
Department of Geography
Durham University
Durham, UK
editors@stata-journal.com

Associate Editors

CHRISTOPHER F. BAUM, Boston College
NATHANIEL BECK, New York University
RINO BELLOCCO, Karolinska Institutet, Sweden, and
University of Milano-Bicocca, Italy
MAARTEN L. BUIS, WZB, Germany
A. COLIN CAMERON, University of California–Davis
MARIO A. CLEVES, University of Arkansas for
Medical Sciences
WILLIAM D. DUPONT, Vanderbilt University
PHILIP ENDER, University of California–Los Angeles
DAVID EPSTEIN, Columbia University
ALLAN GREGORY, Queen's University
JAMES HARDIN, University of South Carolina
BEN JANN, University of Bern, Switzerland
STEPHEN JENKINS, London School of Economics and
Political Science
ULRICH KOHLER, University of Potsdam, Germany

FRAUKE KREUTER, Univ. of Maryland–College Park
PETER A. LACHENBRUCH, Oregon State University
JENS LAURITSEN, Odense University Hospital
STANLEY LEMESHOW, Ohio State University
J. SCOTT LONG, Indiana University
ROGER NEWSON, Imperial College, London
AUSTIN NICHOLS, Urban Institute, Washington DC
MARCELLO PAGANO, Harvard School of Public Health
SOPHIA RABE-HESKETH, Univ. of California–Berkeley
J. PATRICK ROYSTON, MRC Clinical Trials Unit,
London
PHILIP RYAN, University of Adelaide
MARK E. SCHAFFER, Heriot-Watt Univ., Edinburgh
JEROEN WEESIE, Utrecht University
IAN WHITE, MRC Biostatistics Unit, Cambridge
NICHOLAS J. G. WINTER, University of Virginia
JEFFREY WOOLDRIDGE, Michigan State University

Stata Press Editorial Manager

LISA GILMORE

Stata Press Copy Editors

DAVID CULWELL and DEIRDRE SKAGGS

The *Stata Journal* publishes reviewed papers together with shorter notes or comments, regular columns, book reviews, and other material of interest to Stata users. Examples of the types of papers include 1) expository papers that link the use of Stata commands or programs to associated principles, such as those that will serve as tutorials for users first encountering a new field of statistics or a major new technique; 2) papers that go “beyond the Stata manual” in explaining key features or uses of Stata that are of interest to intermediate or advanced users of Stata; 3) papers that discuss new commands or Stata programs of interest either to a wide spectrum of users (e.g., in data management or graphics) or to some large segment of Stata users (e.g., in survey statistics, survival analysis, panel analysis, or limited dependent variable modeling); 4) papers analyzing the statistical properties of new or existing estimators and tests in Stata; 5) papers that could be of interest or usefulness to researchers, especially in fields that are of practical importance but are not often included in texts or other journals, such as the use of Stata in managing datasets, especially large datasets, with advice from hard-won experience; and 6) papers of interest to those who teach, including Stata with topics such as extended examples of techniques and interpretation of results, simulations of statistical concepts, and overviews of subject areas.

The *Stata Journal* is indexed and abstracted by *CompuMath Citation Index*, *Current Contents/Social and Behavioral Sciences*, *RePEc: Research Papers in Economics*, *Science Citation Index Expanded* (also known as *SciSearch*, *Scopus*, and *Social Sciences Citation Index*).

For more information on the *Stata Journal*, including information for authors, see the webpage

<http://www.stata-journal.com>

Subscriptions are available from StataCorp, 4905 Lakeway Drive, College Station, Texas 77845, telephone 979-696-4600 or 800-STATA-PC, fax 979-696-4601, or online at

<http://www.stata.com/bookstore/sj.html>

Subscription rates listed below include both a printed and an electronic copy unless otherwise mentioned.

| U.S. and Canada | | Elsewhere | |
|--|-------|--|-------|
| Printed & electronic | | Printed & electronic | |
| 1-year subscription | \$ 98 | 1-year subscription | \$138 |
| 2-year subscription | \$165 | 2-year subscription | \$245 |
| 3-year subscription | \$225 | 3-year subscription | \$345 |
| 1-year student subscription | \$ 75 | 1-year student subscription | \$ 99 |
| 1-year university library subscription | \$125 | 1-year university library subscription | \$165 |
| 2-year university library subscription | \$215 | 2-year university library subscription | \$295 |
| 3-year university library subscription | \$315 | 3-year university library subscription | \$435 |
| 1-year institutional subscription | \$245 | 1-year institutional subscription | \$285 |
| 2-year institutional subscription | \$445 | 2-year institutional subscription | \$525 |
| 3-year institutional subscription | \$645 | 3-year institutional subscription | \$765 |
| Electronic only | | Electronic only | |
| 1-year subscription | \$ 75 | 1-year subscription | \$ 75 |
| 2-year subscription | \$125 | 2-year subscription | \$125 |
| 3-year subscription | \$165 | 3-year subscription | \$165 |
| 1-year student subscription | \$ 45 | 1-year student subscription | \$ 45 |

Back issues of the *Stata Journal* may be ordered online at

<http://www.stata.com/bookstore/sjj.html>

Individual articles three or more years old may be accessed online without charge. More recent articles may be ordered online.

<http://www.stata-journal.com/archives.html>

The *Stata Journal* is published quarterly by the Stata Press, College Station, Texas, USA.

Address changes should be sent to the *Stata Journal*, StataCorp, 4905 Lakeway Drive, College Station, TX 77845, USA, or emailed to sj@stata.com.



Copyright © 2013 by StataCorp LP

Copyright Statement: The *Stata Journal* and the contents of the supporting files (programs, datasets, and help files) are copyright © by StataCorp LP. The contents of the supporting files (programs, datasets, and help files) may be copied or reproduced by any means whatsoever, in whole or in part, as long as any copy or reproduction includes attribution to both (1) the author and (2) the *Stata Journal*.

The articles appearing in the *Stata Journal* may be copied or reproduced as printed copies, in whole or in part, as long as any copy or reproduction includes attribution to both (1) the author and (2) the *Stata Journal*.

Written permission must be obtained from StataCorp if you wish to make electronic copies of the insertions. This precludes placing electronic copies of the *Stata Journal*, in whole or in part, on publicly accessible websites, file servers, or other locations where the copy may be accessed by anyone other than the subscriber.

Users of any of the software, ideas, data, or other materials published in the *Stata Journal* or the supporting files understand that such use is made without warranty of any kind, by either the *Stata Journal*, the author, or StataCorp. In particular, there is no warranty of fitness of purpose or merchantability, nor for special, incidental, or consequential damages such as loss of profits. The purpose of the *Stata Journal* is to promote free communication among Stata users.

The *Stata Journal* (ISSN 1536-867X) is a publication of Stata Press. Stata, **STATA**, Stata Press, Mata, **MATA**, and NetCourse are registered trademarks of StataCorp LP.

A short guide and a forest plot command (`ipdforest`) for one-stage meta-analysis

Evangelos Kontopantelis
NIHR School for Primary Care Research
Institute of Population Health
University of Manchester
Manchester, UK
e.kontopantelis@manchester.ac.uk

David Reeves
Institute of Population Health
University of Manchester
Manchester, UK
david.reeves@manchester.ac.uk

Abstract. In this article, we describe a new individual patient data meta-analysis postestimation command, `ipdforest`. The command produces a forest plot following a one-stage meta-analysis with `xtmixed` or `xtmelogit`. (These commands have been renamed in Stata 13 to `mixed` and `meqrlogit`, respectively; `ipdforest` is currently not compatible with the new names.) The overall effect is obtained from the preceding mixed-effects regression and the study effects from linear or logistic regressions on each study, which are executed within `ipdforest`. Individual patient data meta-analysis models with Stata are discussed.

Keywords: `st0309`, `ipdforest`, meta-analysis, forest plot, individual patient data, IPD, one-stage

1 Introduction

Meta-analysis, the methodology that allows results from independent studies to be combined, is usually a two-stage process. First, the relevant summary effect statistics are extracted from published articles on the included studies. These are then combined into an overall effect estimate using a suitable meta-analysis model (Harris et al. 2008; Kontopantelis and Reeves 2010). However, problems often arise when an article does not report all the statistical information required as input for the meta-analysis (for example, it fails to provide a variance estimate for the outcome measure); reports a statistic other than the effect size (such as a t -value or p -value) that needs to be transformed with a loss of precision; or provides a sample too clinically heterogeneous for the study to be included in the meta-analysis (Kontopantelis and Reeves 2009).

When individual patient data (IPD) from each study are available, meta-analysts can avoid these problems when estimating study effects; outcomes can be easily standardized, while clinical heterogeneity can be addressed, at least partially, with subgroup analyses and patient-level covariate control. Furthermore, when IPD data are available, meta-analysts can use a mixed-effects regression model to combine information across studies in a single stage. This is recognized as the best approach for performing an IPD meta-analysis, with the two-stage method being at best equivalent in certain scenarios (Mathew and Nordström 2010).

Despite these advantages of the one-stage approach, one obvious advantage of two-stage meta-analysis is the ability to convey information graphically through a forest plot. Because study effects have been calculated or extracted in the first stage of the process, they and their respective confidence intervals can be used to demonstrate the relative strength of the intervention in each study and across all the studies. Forest plots are informative, easy to follow, and particularly useful for readers with little or no experience in meta-analysis methods. It is not surprising, then, that they have become a key feature of meta-analysis and are always presented when two-stage meta-analyses are performed.

However, under a one-stage meta-analysis model, only the overall effect is calculated, not individual study effects; thus creating a forest plot is not straightforward. A search by the authors failed to identify one-stage meta-analysis forest-plot modules in any general or meta-analysis specialist statistical package. We attempt to address this gap in Stata with the `ipdforest` command.

This article is divided into two sections. In the first section, we describe IPD meta-analysis models and their implementation in Stata with available mixed-effects models. In the second section, we describe the `ipdforest` command in detail and provide an example.

2 Individual patient data meta-analysis

A description of IPD meta-analysis methods for continuous and binary outcomes has been provided by Higgins et al. (2003) and Turner et al. (2000), respectively. Although we will only explore a representative selection of linear random-effects models in Stata (using the `xtmixed` command), application to the logistic case using `xtmelogit` should be straightforward. Let us assume IPD from a group of studies. For each trial, we have the exposure variable, which is continuous or binary (for example, control or intervention group membership), and baseline and follow-up data for the continuous outcome and covariates. We will also assume that both the outcome measure and any covariates have been measured in the same way across studies and that, therefore, standardization is not required. In the models that follow, in general, we denote a fixed effect by γ and a random effect by β .

Possibly the simplest approach is to assume that there is a common intercept across studies and that baseline is a fixed effect but to allow the treatment effect to vary at random across studies. Thus we have

$$\begin{aligned} \dot{Y}_{ij} &= \gamma_0 + \beta_{1j}\text{group}_{ij} + \gamma_2 Y_{ij} + \epsilon_{ij} \\ \beta_{1j} &= \gamma_1 + u_{1j} \end{aligned} \tag{1a}$$

and

$$\begin{aligned} \epsilon_{ij} &\sim N(0, \sigma_j^2) \\ u_{1j} &\sim N(0, \tau_1^2) \end{aligned} \tag{1b}$$

where i is the patient; j is the study; \hat{Y}_{ij} is the outcome for patient i in study j ; γ_0 is the fixed common intercept; β_{1j} is the random treatment effect for study j ; γ_1 is the mean treatment effect; group_{ij} is the exposure for patient i in study j ; γ_2 is the fixed baseline effect; Y_{ij} is the baseline score for patient i in study j ; u_{1j} is the random treatment effect for study j (shifting the regression line up or down by study); τ_1^2 is the between-study variance; ϵ_{ij} is the error term for patient i in study j ; and σ_j^2 is the within-study variance for study j .

However, the common intercept and fixed baseline assumptions are difficult to justify, and such a model should be approached with caution—if at all. A more accepted model allows for different fixed intercepts and fixed baseline effects for each study:

$$\begin{aligned}\hat{Y}_{ij} &= \gamma_{0j} + \beta_{1j}\text{group}_{ij} + \gamma_{2j}Y_{ij} + \epsilon_{ij} \\ \beta_{1j} &= \gamma_1 + u_{1j}\end{aligned}\tag{2}$$

where γ_{0j} is the fixed intercept for study j and γ_{2j} is the fixed baseline effect for study j .

Another possibility, although contentious (Whitehead 2002), is to assume that study intercepts are random, as in a multicenter study; for example,

$$\begin{aligned}\hat{Y}_{ij} &= \beta_{0j} + \beta_{1j}\text{group}_{ij} + \gamma_{2j}Y_{ij} + \epsilon_{ij} \\ \beta_{0j} &= \gamma_0 + u_{0j} \\ \beta_{1j} &= \gamma_1 + u_{1j}\end{aligned}\tag{3a}$$

In this case, it is probably wiser to assume a nonzero correlation ρ between the random effects:

$$\begin{aligned}\epsilon_{ij} &\sim N(0, \sigma_j^2) \\ u_{0j} &\sim N(0, \tau_0^2) \\ u_{1j} &\sim N(0, \tau_1^2) \\ \text{cov}(u_{0j}, u_{1j}) &= \rho\tau_0\tau_1\end{aligned}\tag{3b}$$

The baseline could also have been modeled as a random effect, and we could have allowed for nonzero correlations between the three random effects, thus complicating (3) further:

$$\begin{aligned}\hat{Y}_{ij} &= \beta_{0j} + \beta_{1j}\text{group}_{ij} + \beta_{2j}Y_{ij} + \epsilon_{ij} \\ \beta_{0j} &= \gamma_0 + u_{0j} \\ \beta_{1j} &= \gamma_1 + u_{1j} \\ \beta_{2j} &= \gamma_2 + u_{2j}\end{aligned}\tag{4a}$$

with effects

$$\begin{aligned}
 \epsilon_{ij} &\sim N(0, \sigma_j^2) \\
 u_{0j} &\sim N(0, \tau_0^2) \\
 u_{1j} &\sim N(0, \tau_1^2) \\
 u_{2j} &\sim N(0, \tau_2^2) \\
 \text{cov}(u_{0j}, u_{1j}) &= \rho_1 \tau_0 \tau_1 \\
 \text{cov}(u_{0j}, u_{2j}) &= \rho_2 \tau_0 \tau_2 \\
 \text{cov}(u_{1j}, u_{2j}) &= \rho_3 \tau_1 \tau_2
 \end{aligned} \tag{4b}$$

In some cases, the focus might be on interactions. For example, if we assume a continuous and standardized variable X , we can expand (2) to include fixed effects, in this instance, for both X and its interaction with the treatment:

$$\begin{aligned}
 \hat{Y}_{ij} &= \gamma_{0j} + \beta_{1j} \text{group}_{ij} + \gamma_{2j} Y_{ij} + \gamma_3 X_{ij} + \gamma_4 \text{group}_{ij} X_{ij} + \epsilon_{ij} \\
 \beta_{1j} &= \gamma_1 + u_{1j}
 \end{aligned} \tag{5}$$

If we consider `Yfin` and `Ybas` as representing the outcome and baseline, respectively, the exposure variable `group`, and the study identifier `studyid` for four studies, we can implement the models described above by using `xtmixed`.

Model (1): Fixed common intercept; random treatment effect; fixed effect for baseline.

```
. xtmixed Yfin i.group Ybas || studyid:group, nocons
```

The `nocons` option suppresses estimation of the intercept as a random effect.

Model (2): Fixed study-specific intercepts; random treatment effect; fixed study-specific effects for baseline (where `Ybas'i'`=`Ybas` if `studyid='i'` and equals 0 otherwise).

```
. xtmixed Yfin i.group i.studyid Ybas1 Ybas2 Ybas3 Ybas4 || studyid:group,
> nocons
```

Model (3): Random study intercept; random treatment effect; fixed study-specific effects for baseline.

```
. xtmixed Yfin i.group Ybas1 Ybas2 Ybas3 Ybas4 || studyid:group, cov(uns)
```

Model (4): Random study intercept; random treatment effect; random effect for baseline.

```
. xtmixed Yfin i.group Ybas || studyid:group Ybas, cov(uns)
```

In general, a covariate (or an interaction term) can be modeled as a fixed or random effect, but in the latter case, the complexity of the model increases and nonconvergence issues are more likely to be encountered. If we also consider patient covariate `age` and its interaction with the treatment effect, then (5) will be

```
. xtmixed Yfin i.group i.studyid Ybas1 Ybas2 Ybas3 Ybas4 age i.group#c.age
> || studyid: group, nocons
```

Or alternatively, `age` can be modeled as a random effect:

```
. xtmixed Yfin i.group i.studyid Ybas1 Ybas2 Ybas3 Ybas4 age i.group#c.age
> || studyid: group age, nocons
```

3 The ipdforest command

3.1 Syntax

```
ipdforest varname [, re(varlist) fe(varlist) fets(namelist) ia(varname) auto
label(varlist) or gsavedir(string) gsavename(string) eps gph
export(string) ]
```

where *varname* is the exposure variable, continuous or binary (for example, intervention or control).

3.2 Options

`re(varlist)` specifies covariates to be included as random factors. For each covariate specified, a different regression coefficient is estimated for each study.

`fe(varlist)` specifies covariates to be included as fixed factors. For each covariate specified, the respective coefficient in the study-specific regressions is fixed to the value returned by the multilevel regression.

`fets(namelist)` specifies covariates to be included as study-specific fixed factors (that is, by using the estimated study fixed effects from the main regression in all individual study regressions). Only baseline scores and study identifiers can be included. For each covariate specified, the respective coefficient in the study-specific regressions is fixed to the value returned by the multilevel regression for the specific study. For study-specific intercepts, the study identifier (not in factor-variable format, for example, `studyid`) or the *stub* of the dummy variables (for example, `studyid_` when dummy study identifiers are `studyid_1 studyid_3`, etc.) would be included. For study-specific baseline scores, only the *stub* of the dummy variables is accepted (for example, `dept0s_` when dummy study baseline scores are `dept0s_1 dept0s_3`, etc.).

`ia(varname)` specifies covariates for which the interaction with the exposure variable will be calculated and displayed. The covariate should also be specified as a fixed, random, or study-specific fixed effect. If binary, the command will provide two sets of results, one for each group. If categorical, it will provide as many sets of results as there are categories. If continuous, it will provide one set of results for the main effect and one for the interaction. Although the command will allow a variable to be interacted with the exposure variable as a fixed or study-specific fixed effect, the variable necessarily will be included as a random effect in the individual regressions (it will not run a regression with the interaction term only; the main effects must be included as well). Therefore, although the overall effect will differ between a model with a fixed-effects interacted variable and a random-effects one, the individual study effects will be identical across the two approaches.

`auto` allows `ipdforest` to automatically detect the specification of the preceding model. This option cannot be issued along with options `re()`, `fe()`, `fets()`, or `ia()`. The `auto` option will work in most situations, but it comes with certain limitations. It uses the returned command string of the preceding command, which is effectively constrained to 244 characters; therefore, the `auto` option will return an error if `ipdforest` follows a very wide regression model—in such a situation, only the manual specification can be used. In addition, the variable names used in the preceding model must follow certain rules: 1) fixed-effects covariates (manually with option `fe()`) must not contain underscores; 2) for study-specific intercepts (manually with option `fets()`), factor-variable format is allowed or a *varlist* (for example, `cons_2–cons_16`), but each variable must contain a single underscore followed by the study number (not necessarily sequential); and 3) for study-specific baseline scores (manually with option `fets()`), each variable must contain a single underscore followed by the study number (again, not necessarily sequential). There are no restrictions for random-effects covariates (manually with option `re()`). For interactions (manually with option `ia()`), the factor-variable notation should be preferred (for example, `i.group#c.age`) and, alternatively, the older `xi:` notation. Interactions expanded to dummy variables cannot be identified with the `auto` option, and only the manual specification should be used in this case. Variables whose names start with an `_I` and contain a capital `X` will be assumed to be expanded interaction terms, and if detected in the last model, `ipdforest` will terminate with a syntax error.

`label(varlist)` specifies labels for the studies. Up to two variables can be specified and converted to strings. If two variables are specified, they will be separated by a comma. Usually, the author names and the year of study are selected as labels. If `label()` is not specified, the command automatically uses the value labels of the numeric cluster variable, if any, to label the forest plot. Either way, the final string is truncated to 30 characters.

`or` reports odds ratios instead of coefficients. It can only be used following the execution of `xtmelogit`.

`gsavedir(string)` specifies the directory in which to save the graph, if different from the active directory.

`gsavename(string)` specifies the optional name prefix for the graph. Graphs are saved as `gsavename_graphname.gph` or `gsavename_graphname.eps`, where `graphname` includes a description of the summary effect (for example, `main_group` for the main effect if `group` is the exposure variable).

`eps` saves the graph in `.eps` format instead of the default `.gph`.

`gph` saves the graph in `.gph` format. `gph` is the default. Use it to save in both formats: including only the `eps` option will save the graph in `.eps` format only.

`export(string)` exports the study identifiers, weights, effects, and standard errors in a Stata dataset (named after `string`). It is provided for users who wish to use other commands or software to draw the forest plots.

3.3 Stored results

`ipdforest` stores the following in `r()`:

Scalars

| | | | |
|-------------------------|--|----------------------------|--|
| <code>r(Isq)</code> | heterogeneity measure I^2 | <code>r(eff1pe_ov)</code> | overall effect estimate |
| <code>r(Hsq)</code> | heterogeneity measure H_M^2 | <code>r(eff1se_ov)</code> | standard error of the overall effect |
| <code>r(tausq)</code> | $\hat{\tau}^2$, between-study variance estimate | <code>r(eff1pe_sti)</code> | effect estimate for study i |
| <code>r(tausqlo)</code> | $\hat{\tau}^2$, lower 95% confidence interval | <code>r(eff1se_sti)</code> | standard error of the effect for study i |
| <code>r(tausqup)</code> | $\hat{\tau}^2$, upper 95% confidence interval | | |

If an interaction with a continuous variable is included in the model, it also stores the following:

Scalars

| | | | |
|---------------------------|--|----------------------------|---|
| <code>r(eff2pe_ov)</code> | overall interaction effect estimate | <code>r(eff2pe_sti)</code> | interaction effect estimate for study i |
| <code>r(eff2se_ov)</code> | standard error of the overall interaction effect | <code>r(eff2se_sti)</code> | interaction effect standard error for study i |

If the interaction variable is binary, the first set of results corresponds to the effects for the first category of the binary (for example, `sex = 0`) and the second set for the second category (for example, `sex = 1`). If the variable is categorical, the command returns as many sets of effect results as there are categories (with each set corresponding to one category). Estimation results from `xtmixed` or `xtmelogit` in `e()` are restored after the execution of `ipdforest`.

3.4 Methods

The `ipdforest` command is issued following a random-effects IPD meta-analysis conducted using a linear (`xtmixed`) or logistic (`xtmelogit`) two-level regression with patients nested within studies. The command provides a meta-analysis summary table

and a forest plot. Study effects are calculated within `ipdforest`, while the overall effect and variance estimates are extracted from the preceding regression. The default estimation methods for `xtmixed` and `xtmelogit` are restricted maximum likelihood and maximum likelihood, respectively. A description of these methods is beyond the scope of this article.

`ipdforest` estimates individual study effects and their standard errors by using one-level linear or logistic regression analyses. Following `xtmixed`, `regress` is used, and following `xtmelogit`, `logit` is used for each study in the meta-analysis. The `ipdforest` command controls these regressions for fixed- or random-effects covariates that were specified in the preceding two-level regression. The user has full control over the covariates to be included in the `ipdforest` command, including their specification as fixed or random effects. However, we strongly recommend using the same specification as in the preceding `xtmixed` or `xtmelogit` command because the reported overall effect and its confidence interval is taken from that model.

In the estimation of individual study effects, `ipdforest` controls for a random-effects covariate (that is, allowing the regression coefficient to vary by study) by including the covariate as an independent variable in each regression. Control for a fixed-effects covariate (where the regression coefficient is assumed constant across studies and is given by the coefficient estimated under `xtmixed` or `xtmelogit`) is a little more complex. Because it is not possible to specify a fixed value for a regression coefficient under `regress`, the continuous outcome variable is adjusted by subtracting the contribution of the fixed covariates to its values prior to analysis. For a binary outcome, the equivalent is achieved using the `offset` option in `logit`. Patient weights are uniform; therefore, each study's weight is the ratio of its participants over the total number of participants across all studies.

Between-study variability in the treatment effect, known as heterogeneity, arises from differences in study design, quality, outcomes, or populations and needs to be accounted for in the meta-analysis model when present. Heterogeneity is usually reported in the form of measures or tests that compare the between- and within-study variance estimates. For continuous outcomes, `ipdforest` reports two heterogeneity measures, I^2 and H_M^2 , based on the `xtmixed` output. I^2 values of 25%, 50%, and 75% correspond to low, moderate, and high heterogeneity, respectively (Higgins et al. 2003), while H_M^2 takes values in the $[0, +\infty)$ range with 0 indicating perfect homogeneity (Mittlböck and Heinzl 2006). We have not attempted to calculate an IPD version of Cochran's Q , the orthodox χ_{k-1}^2 homogeneity test, considering its poor performance when the number of studies k is small (Hardy and Thompson 1998). For binary outcomes, an estimate of the within-study variance is not reported under `xtmelogit`, and hence, heterogeneity measures cannot be computed. The between-study variance estimate $\hat{\tau}^2$ and its confidence interval are reported under both models.

Fixed-effects meta-analysis models are widely used when heterogeneity is very low or 0. However, a more conservative approach is to take account of even low levels of between-study variability by adopting a random-effects model (Hunter and Schmidt 2000). When between-study variance is estimated to be close to 0, results with the two

approaches converge. Therefore, although `ipdforest` is a postestimation command for random-effects IPD meta-analysis, output is close to that for a fixed-effects model when $\hat{\tau}^2 \approx 0$.

3.5 Example

As an example, we apply the `ipdforest` command to a dataset of four depression intervention studies. Data were provided by the authors of the studies, and we had complete information in terms of age, gender, exposure (control and intervention group membership), continuous outcome baseline, and endpoint values for 518 patients. Because the findings of the IPD meta-analysis had not been published when this article was being prepared, we used fake author names and generated random continuous and binary outcome variables for the purposes of this example while keeping the covariates at their actual values. We introduced correlation between baseline and endpoint scores and between-study variability, although the exact specification of the data generation is unimportant.

Using the semiartificial dataset, we perform a logistic IPD meta-analysis, followed by the `ipdforest` command.

```
. use ipdforest_example
. describe
Contains data from ipdforest_example.dta
  obs:          518
  vars:          17          6 Feb 2012 11:35
  size:         20,202
```

| variable name | storage type | display format | value label | variable label |
|---------------|--------------|----------------|-------------|--------------------------------|
| studyid | byte | %22.0g | stid | Study identifier |
| patid | int | %8.0g | | Patient identifier |
| group | byte | %20.0g | grplbl | Intervention/control group |
| sex | byte | %10.0g | sexlbl | Gender |
| age | float | %10.0g | | Age in years |
| depB | byte | %9.0g | | Binary outcome, endpoint |
| depBbas | byte | %9.0g | | Binary outcome, baseline |
| depBbas1 | byte | %9.0g | | Bin outcome baseline, trial 1 |
| depBbas2 | byte | %9.0g | | Bin outcome baseline, trial 2 |
| depBbas5 | byte | %9.0g | | Bin outcome baseline, trial 5 |
| depBbas9 | byte | %9.0g | | Bin outcome baseline, trial 9 |
| depC | float | %9.0g | | Continuous outcome, endpoint |
| depCbas | float | %9.0g | | Continuous outcome, baseline |
| depCbas1 | float | %9.0g | | Cont outcome baseline, trial 1 |
| depCbas2 | float | %9.0g | | Cont outcome baseline, trial 2 |
| depCbas5 | float | %9.0g | | Cont outcome baseline, trial 5 |
| depCbas9 | float | %9.0g | | Cont outcome baseline, trial 9 |

```
Sorted by:  studyid  patid
```

We generate a centered age variable, interacted with the exposure variable in a mixed-effects logistic regression model. The model includes fixed study-specific intercepts and fixed study-specific effects for baseline and random treatment and age effects. The `ipdforest` command follows the regression model, requesting outcomes for both the main effect and the interaction.

```
. quietly summarize agec
. quietly generate agec = age-r(mean)
. xtlogit depB group agec sex i.studyid depBbas1 depBbas2 depBbas5 depBbas9
> i.group#c.agec || studyid:group agec, var nocons or

Refining starting values:
Iteration 0:  log likelihood = -347.40378 (not concave)
Iteration 1:  log likelihood = -336.07882 (not concave)
Iteration 2:  log likelihood = -329.28268

Performing gradient-based optimization:
Iteration 0:  log likelihood = -329.28268 (not concave)
Iteration 1:  log likelihood = -326.79754
Iteration 2:  log likelihood = -326.5689
Iteration 3:  log likelihood = -326.55747
Iteration 4:  log likelihood = -326.55747

Mixed-effects logistic regression
Group variable: studyid

Number of obs      =      518
Number of groups   =         4
Obs per group: min =         42
                  avg =      129.5
                  max =        214

Integration points =       7
Log likelihood = -326.55747

Wald chi2(11)     =      42.06
Prob > chi2       =      0.0000
```

| depB | Odds Ratio | Std. Err. | z | P> z | [95% Conf. Interval] | |
|--------------|------------|-----------|-------|-------|----------------------|----------|
| group | 1.840804 | .3666167 | 3.06 | 0.002 | 1.245894 | 2.71978 |
| agec | .9867902 | .0119059 | -1.10 | 0.270 | .9637288 | 1.010403 |
| sex | .7117592 | .1540753 | -1.57 | 0.116 | .4656639 | 1.087912 |
| studyid | | | | | | |
| 2 | 1.050007 | .5725515 | 0.09 | 0.929 | .3606168 | 3.057302 |
| 5 | .8014552 | .5894511 | -0.30 | 0.763 | .1896011 | 3.387799 |
| 9 | 1.281413 | .6886055 | 0.46 | 0.644 | .4469621 | 3.673734 |
| depBbas1 | 3.152909 | 1.49528 | 2.42 | 0.015 | 1.244587 | 7.987251 |
| depBbas2 | 4.480302 | 1.863908 | 3.60 | 0.000 | 1.982385 | 10.12573 |
| depBbas5 | 2.387336 | 1.722993 | 1.21 | 0.228 | .5802064 | 9.823007 |
| depBbas9 | 1.881203 | .7086506 | 1.68 | 0.093 | .8990571 | 3.936261 |
| group#c.agec | | | | | | |
| 1 | 1.011776 | .0163748 | 0.72 | 0.469 | .9801858 | 1.044385 |
| _cons | .5533714 | .2398341 | -1.37 | 0.172 | .2366473 | 1.293993 |

| Random-effects Parameters | Estimate | Std. Err. | [95% Conf. Interval] | |
|---------------------------|----------|-----------|----------------------|---|
| studyid: Independent | | | | |
| var(group) | 6.15e-21 | 2.03e-11 | 0 | . |
| var(agec) | 6.03e-18 | 4.41e-11 | 0 | . |

LR test vs. logistic regression: chi2(2) = 0.00 Prob > chi2 = 1.0000

Note: LR test is conservative and provided only for reference.

. ipdforest group, fe(sex) re(agec) ia(agec) or

One-stage meta-analysis results using xtmelogit (ML method) and ipdforest

Main effect (group)

| Study | Effect | [95% Conf. Interval] | | % Weight |
|----------------|--------|----------------------|-------|----------|
| Hart 2005 | 2.118 | 0.942 | 4.765 | 19.88 |
| Richards 2004 | 2.722 | 1.336 | 5.545 | 30.69 |
| Silva 2008 | 2.690 | 0.748 | 9.676 | 8.11 |
| Kompany 2009 | 1.895 | 0.969 | 3.707 | 41.31 |
| Overall effect | 1.841 | 1.246 | 2.720 | 100.00 |

One-stage meta-analysis results using xtmelogit (ML method) and ipdforest

Interaction effect (group x agec)

| Study | Effect | [95% Conf. Interval] | | % Weight |
|----------------|--------|----------------------|-------|----------|
| Hart 2005 | 0.972 | 0.901 | 1.049 | 19.88 |
| Richards 2004 | 0.995 | 0.937 | 1.055 | 30.69 |
| Silva 2008 | 0.987 | 0.888 | 1.098 | 8.11 |
| Kompany 2009 | 1.077 | 1.015 | 1.144 | 41.31 |
| Overall effect | 1.012 | 0.980 | 1.044 | 100.00 |

Heterogeneity Measures

| | value | [95% Conf. Interval] | |
|---------|-------|----------------------|---|
| I (%) | . | | |
| H | . | | |
| tau est | 0.000 | 0.000 | . |

Maximum likelihood converged successfully in this example, and the between-study variance estimate $\hat{\tau}^2$ was practically 0. Note that the intercept for the reference study (`studyid = 1`) was estimated in `_cons`. The reported coefficients under `studyid` are the differences in intercept compared with the first study. I^2 and H_M^2 could not be estimated because residual variability is not reported under `xtmelogit`. The overall treatment effect was significant at the 95% level, but the overall effect for the interaction of treatment and age was not. The forest plots created by `ipdforest` are displayed in figures 1 and 2.




Figure 1. Main-effect IPD forest plot reporting odds ratios




Figure 2. Interaction-effect IPD forest plot reporting odds ratios

4 Discussion

The aim of this article was to provide a practical guide for conducting one-stage IPD meta-analysis and to present `ipdforest`, a new forest-plot command. `ipdforest` aims to help meta-analysts better communicate their results through the familiar and distinctive forest plot—a graphical output not previously available in one-stage IPD meta-analysis software routines.

Although only binary or continuous exposure variables can be modeled, categorical exposures can also be investigated with the use of dummy variables and a focus on the comparison of interest through one of these. In addition, `ipdforest` is fully compatible with the estimates produced by the multiple-imputation estimation command `mi estimate: xtmixed` or `mi estimate: xtmelogit`.

Note that these commands were renamed in Stata 13: `xtmixed` to `mixed` and `xtmelogit` to `meqrlogit`. `ipdforest` is not yet compatible with the new commands, but users of version 13 can still use the older commands before calling `ipdforest`.

5 Acknowledgments

We would like to thank Isabel Canette, senior statistician at StataCorp, for her help with advanced aspects of `xtmixed` and `xtmelogit` and the anonymous reviewer whose comments and suggestions improved the command significantly. Evan Kontopantelis is on an NIHR School for Primary Care fellowship.

6 References

- Hardy, R. J., and S. G. Thompson. 1998. Detecting and describing heterogeneity in meta-analysis. *Statistics in Medicine* 17: 841–856.
- Harris, R. J., M. J. Bradburn, J. J. Deeks, R. M. Harbord, D. G. Altman, and J. A. C. Sterne. 2008. `metan`: Fixed- and random-effects meta-analysis. *Stata Journal* 8: 3–28.
- Higgins, J. P. T., S. G. Thompson, J. J. Deeks, and D. G. Altman. 2003. Measuring inconsistency in meta-analyses. *British Medical Journal* 327: 557–560.
- Hunter, J. E., and F. L. Schmidt. 2000. Fixed effects vs. random effects meta-analysis models: Implications for cumulative research knowledge. *International Journal of Selection and Assessment* 8: 275–292.
- Kontopantelis, E., and D. Reeves. 2009. MetaEasy: A meta-analysis add-in for Microsoft Excel. *Journal of Statistical Software* 30: 1–25.
- . 2010. `metaan`: Random-effects meta-analysis. *Stata Journal* 10: 395–407.
- Mathew, T., and K. Nordström. 2010. Comparison of one-step and two-step meta-analysis models using individual patient data. *Biometrical Journal* 52: 271–287.
- Mittlböck, M., and H. Heinzl. 2006. A simulation study comparing properties of heterogeneity measures in meta-analyses. *Statistics in Medicine* 25: 4321–4333.
- Turner, R. M., R. Z. Omar, M. Yang, H. Goldstein, and S. G. Thompson. 2000. A multilevel model framework for meta-analysis of clinical trials with binary outcomes. *Statistics in Medicine* 19: 3417–3432.
- Whitehead, A. 2002. *Meta-Analysis of Controlled Clinical Trials*. Chichester, UK: Wiley.

About the authors

Evangelos (Evan) Kontopantelis is a research fellow in statistics at the Centre for Biostatistics, Institute of Population Health at the University of Manchester in England. His research interests include statistical methods in health sciences with a focus on meta-analysis, longitudinal data modeling, and large clinical databases. For this article, he was partly supported by an NIHR School for Primary Care Research fellowship in primary health care.

David Reeves is a senior research fellow in statistics at the Centre for Biostatistics, Institute of Population Health at the University of Manchester in England. His methodological research interests include the robustness of statistical methods, the analysis of observational studies, and the application of social network analysis methods to health systems.