

ABSORB: A computer program for Assessing Bias using Sensitivity-analysis for Outcome Reporting Biases

Dan Jackson

July 1, 2010

Abstract

Despite investigators efforts, and the inception of trial registration, the concern that the study outcomes included in a meta-analysis may not be representative of the population often remains. This is because outcome reporting bias, commonly referred to as publication bias, results in a biased evidence base. Modelling this bias using weight functions has received considerable attention in the literature but has hitherto been difficult to implement in practice. Problems have also been identified with using weight functions in an entirely estimation based context. In order to overcome these difficulties, the *ABSORB* computer program, which performs sensitivity analyses using weight functions and the conventional random effects model for meta-analysis, has been developed. This is implemented in the freely available *R* programming language. The only additional programming required by the advanced user is to code the weight function. Three established weight functions for more general use are provided. *ABSORB* performs an entire sensitivity analysis based on the user specified weight function and provides both numerical and graphical output to aid interpretation when assessing the potential impact of outcome reporting bias.

Keywords: Meta-analysis, publication bias, sensitivity analysis, statistical software.

1 Introduction

Meta-analysis and evidence synthesis have fundamental roles in evidence based medicine. The standard methods assume that the data are representative of all studies that have been performed but this may not be true. This is because outcome reporting bias, more commonly referred to as publication bias [1, 2], may be present. If this has occurred then typically the concern is that outcomes providing weaker evidence that the treatment is effective are suppressed and hence the effectiveness of the treatment is overestimated. The term ‘outcome reporting bias’ is used here because it is recognised that studies can be published where the particular outcome required for the meta-analysis is not reported and therefore cannot be included in the meta-analysis.

Common methods for detecting outcome reporting bias involve some form of funnel plot, which may be assessed informally or using regression [3, 4]. Plots such as these, coupled with the investigators’ knowledge of the subject area, may result in the concern that outcome reporting bias is present. We assume here that, for whatever reason, there is the suspicion that outcome reporting bias may be present.

Various methods have been suggested for modelling publication and related biases in meta-analysis. In particular, weight functions are an established approach [5, 6, 7, 8, 9]. First we assume a model for the studies’ outcomes before selection. We then further posit a parametric model for the selection process, which results in the outcomes either being ‘published’, ie present in the meta-analysis, or ‘unpublished’. Although weight functions

are a natural way to model publication bias, their application is fraught with difficulties in practice. First of all, since the distribution of the outcomes follow non-standard distributions which typically contain integrals that cannot be evaluated analytically, the likelihood of the data cannot be written down in closed form in a way amenable to standard algorithms and computing packages. This has prevented practitioners adopting this kind of approach to assessing and quantifying the bias. A second difficulty is identifying parameters that describe the form of the bias. These are very poorly estimated [9, 10] and they have very drastic implications for the other estimates, such as the treatment effect. Similar observations have been made in the context of using selection models in the missing data literature [11]. Hence we advocate a sensitivity approach as previously recommended [9, 10, 12].

Alternative methods have been proposed but these are not without their problems. Trim and Fill [13] is easily implemented but makes very strong symmetry assumptions and does not allow for the uncertainty in the imputed outcomes, or indeed the number of these. Copas' selection selection model [10] usually performs well, and an *S-plus* program is in place to implement this [14], but this method encounters problems in some datasets [15].

The contribution of this paper is to introduce the *ABSORB* computer program, which is implemented in the *R* programming language, and allows the user to specify the form of their weight function and automatically produce a sensitivity analysis under the assumption that the standard random effects model applies to the population of outcomes. This allows a much greater degree of flexibility than has hitherto been possible and facilitates the routine use of weight functions to assess outcome reporting bias. The

rest of the paper is set out as follows. In section 2 the use of weight functions to model biases in meta-analysis is briefly explained and in section 3 the *ABSORB* program and the methods it uses are described in detail. In section 4 results from *ABSORB* for a familiar example are presented and the article concludes with a brief discussion in section 5.

2 Using weight functions to model publication bias

We assume the conventional random effects model [16, 17] for the studies' outcomes before selection. The methods are therefore widely applicable, because this is most common model in applications. Denoting these by Y_i , $i = 1, 2, \dots, n$, we assume $Y_i | \sigma_i^2 \sim N(\mu, \sigma_i^2 + \tau^2)$, where μ denotes the overall or average treatment effect and τ^2 denotes the between-study variance. The σ_i^2 are estimated in practice but assumed known and are conditioned on in application, a suitable approximation provided the studies are large enough. We further assume that the outcomes are independent.

We will denote the weight function by $P(\text{published} | y_i, \sigma_i^2) = p(y_i, \sigma_i^2)$. This emphasises that the probability of publication can depend on both the studies' outcomes and their precisions but we make the, perhaps strong, assumption that these alone affects the outcomes' probabilities of publication. Since the studies present in the meta-analysis are published, they follow the distribution $Y_i | (\text{published}, \sigma_i^2)$, which can be obtained from the assumed form of the weight function as

$$f(y_i | \text{published}, \sigma_i^2) = \frac{f(y_i | \sigma_i^2) P(\text{published} | y_i, \sigma_i^2)}{\int_{-\infty}^{\infty} f(y_i | \sigma_i^2) P(\text{published} | y_i, \sigma_i^2) dy_i}$$

which, from the random effects model for $Y_i|\sigma_i^2$ and the weight function as expressed above, gives

$$f(y_i|\text{published}, \sigma_i^2) = \frac{\phi(y_i/\sqrt{\sigma_i^2 + \tau^2})p(y_i, \sigma_i^2)}{\int_{-\infty}^{\infty} \phi(y_i/\sqrt{\sigma_i^2 + \tau^2})p(y_i, \sigma_i^2)dy_i} \quad (1)$$

where $\phi(\cdot)$ denotes the standard normal density function. The integrals in the denominator of (1) are evaluated numerically if necessary.

3 The ABSORB program

I have developed the program *ABSORB*, for **A**ssessing **B**ias using **S**ensitivity-analysis for **O**utcome **R**eporting **B**ias, which facilitates the routine use of weight functions to model publication bias in the context of standard meta-analysis data. Here the user defines the type of weight function that they would like to consider and with a single command produce a sensitivity analysis.

3.1 Dealing with the uncertainty in the between-study variance

When applying the random effects model in practice, it is conventional to use the estimated between-study variance $\hat{\tau}^2$ as if it were the true value when pooling the studies' outcomes and making inferences about the overall treatment effect μ . This is justified as a suitable approximation under the random effects model because the scale parameter of a location-scale model may be approximated by a consistent estimate when making inferences about the location. This is not so when the studies' outcomes are subjected to a selection process because μ and τ^2 do not retain their interpretation as

location and scale parameters. Hence uncertainty in τ^2 is not ignored by *ABSORB* unless $\hat{\tau}^2 = 0$. This adopts the convention of collapsing the random effects model to a fixed effects model when $\hat{\tau}^2 = 0$. Since τ^2 is consistently estimated, asymptotically this provides a suitable approximation because only if a fixed effects model truly applies do we have $\hat{\tau}^2 \rightarrow 0$, and we can safely assume the between-study variance away when making inferences about the treatment effect.

3.2 Applying the ABSORB function in practice

We require the user to specify their weight function in terms of two parameters, which are referred to as the ‘family member’ parameter, which we denote by γ , and the ‘severity of the bias’ parameter, which we denote by β , where $\beta \geq 0$. Hence we write the weight function in (1) as $p(y_i, \sigma_i^2; \gamma, \beta)$. In order to provide interpretable results, a further requirement that we make is that if $\beta = 0$ then the weight function is unity, ie $p(y_i, \sigma_i^2; \gamma, 0) = 1$. Hence from (1) the random effects model applies if $\beta = 0$ and there is no outcome reporting bias. More generally, if $p(y_i, \sigma_i^2; \gamma, 0)$ is a function of σ_i^2 , but not y_i , then there is no bias. In particular, if $p(y_i, \sigma_i^2; \gamma, 0) = c$, where c is any constant, then $\beta = 0$ implies no bias; as previously noted the weight function need only be specified within a constant of proportionality [5, 12]. However we take $p(y_i, \sigma_i^2; \gamma, 0) = 1$ so that the no bias model provides the largest probability allowable and hence the smallest number of missing outcomes as explained below.

As advocated by Jackson [12], we assume that as β increases so does, in some sense, the severity of the bias. Hence we make the requirement that

β is easily interpretable but we are flexible regarding the precise manner in which it models bias.

3.3 Three example weight functions for use with ABSORB

In order to make the ideas concrete, and provide a rich array of weight functions for sensitivity analyses in practice, we will define three possibilities. These are simplified versions of those used by Bowden *et al.* [18]. We assume here that negative estimates of effect show a positive effect and are more likely to be published, as is the case in the common scenario where the outcomes are log odds ratios where the event of interest is death. This issue is returned to in the discussion.

3.3.1 Weight function 1

We assume that $p(y_i, \sigma_i^2; \gamma, \beta) = 1$ if $y_i < \gamma$ and $p(y_i, \sigma_i^2; \gamma, \beta) = 1 - \beta$ otherwise. Outcomes that fail to meet some magnitude threshold γ are less likely to be published than other outcomes. This weight function may be coded in *R* as

```
wf=function(x, sigma2, beta, family)
{
r=sign(x-family)
power=0.5*(r^r+1)
(1-beta)^power
}
```

3.3.2 Weight function 2

We assume that $p(y_i, \sigma_i^2; \gamma, \beta) = 1$ if $y_i < \gamma\sigma_i$ and $p(y_i, \sigma_i^2; \gamma, \beta) = 1 - \beta$ otherwise. Outcomes that fail to meet some threshold in terms of their one-sided statistical significance, determined by γ , are less likely to be published than other outcomes. This weight function may be coded in *R* by replacing the line

```
r=sign(x-family)
```

with

```
r=sign(x-family * sigma2^0.5)
```

in the code given above for weight function 1.

3.3.3 Weight function 3

We assume that $p(y_i, \sigma_i^2; \gamma, \beta) = \exp(-\beta(\Phi(y_i/\sigma_i))^\gamma)$, where $\gamma > 0$ and $\Phi(\cdot)$ denotes the standard normal cumulative distribution function. As β increases from zero, studies which achieve smaller one-sided p-values, $\Phi(y_i/\sigma_i)$, become increasingly less likely to be published than those with more significant findings. Unlike weight function 2, here the probability of publication varies continuously over the range of statistical significance, rather than depending solely on a particular threshold or significance level. The parameter γ determines how sharply the relative probability of publication drops for less significant results compared to those with more significant findings. This weight function can be coded as

```
wf=function (x, sigma2, beta, family)
```

```

{
pval=pnorm(x/((sigma2)^0.5))
exp(-beta*pval^family)
}

```

3.3.4 Other weight functions

The user is free to code any weight function that they would like to consider in their sensitivity analysis, subject to the above requirements. The only potential complication is that the weight function must be coded so that it accepts vector values and returns a vector whose entries are the evaluated weight function values because this is required by R's *integrate* command which is called by *ABSORB*. For two of our weight functions, the necessary integrals can be evaluated analytically but all quantities are obtained numerically by *ABSORB* so that it is completely generic and can be used with any weight function, subject to the requirements stated above.

The user must define their weight function using the program name *wf* because this is the function that *ABSORB* calls when evaluating the weight function. This must be specified in terms of the variables *x* (the outcome), the within-study variance (*sigma2*), the bias (*beta*) and family (*family*) parameters, as shown in the code for the three illustrative weight functions above. *ABSORB* also calls the functions *likelihood* and *denominators*.

3.4 Other arguments that can be passed to ABSORB

The *ABSORB* function has four variables that must be given by the user and, with the defaults for the others, and is called as *ABSORB(ys, sigma2s, maxbeta, family)*. The variables *ys* and *sigma2s* are vectors containing the

outcomes and within-study variances, and so pass the data to *ABSORB*. The variable *maxbeta* is the maximum value of β that the user wishes to consider; *ABSORB* will automatically consider a range of values of β from zero up to the maximum value specified. The variable *family* is a vector of length 2, which provides the range of values for the family member parameter γ that the user wishes to investigate.

3.5 The output provided by ABSORB

A standard DerSimonian and Laird [16] random effects meta-analysis is initially performed and then the data are re-analysed using maximum likelihood. The log-likelihood is the sum of the logarithm of terms given by (1), subject to the constraint that $\beta = 0$, and this is maximised using the the DerSimonian and Laird estimates as starting values. This maximisation is performed in terms of μ and $\log(\tau^2)$. Numerical difficulties were occasionally encountered with some weight functions when integrating over the whole real line in the denominator of (1), as an integral was sometimes thought incorrectly not to converge, so these integrations were performed over the intervals containing points within 6 standard deviations of the mean, ie the intervals $(\mu - 6\sqrt{\sigma_i^2 + \tau^2}, \mu + 6\sqrt{\sigma_i^2 + \tau^2})$.

The standard errors of $\hat{\mu}$ from all likelihood based analyses are obtained from the observed information matrix. If $\hat{\tau}^2$ is determined to be zero (less than 0.005 is the default) then uncertainty in τ^2 is ignored as explained above and the second derivative with respect to μ alone is inverted to provide standard errors of $\hat{\mu}$, but otherwise the inverse of the entire observed information matrix, parameterised in terms of μ and $\log(\tau^2)$, is used. The *fdHess* com-

mand, which is part of the *nlme* package, was found to be a reliable method for evaluating the Hessian, so this package is required.

The *ABSORB* function then uses these maximum likelihood estimates to provide starting values for a further maximum likelihood analysis with γ set to the minimum value specified by the user and a small value of β . The resulting estimates are then used as starting values in another analysis with a slightly larger value of β and so on until the maximum β specified by the user is reached. The default is that 10 equally spaced steps between zero and the maximum are used. In addition to showing the results graphically, all likelihood based estimates $\hat{\mu}$ and $\hat{\tau}^2$ and the standard errors of $\hat{\mu}$ are stored by the program and returned to the user.

This process is then repeated for three other values of γ : the maximum value specified by the range given by the user and two values equally spaced between the two extremes. In order to interpret the various models, the weight function can roughly be converted to the number of unpublished outcomes using the formula provided by Bowden *et al.* [18]

$$N = \sum_{i=1}^n \frac{1}{p(y_i, \sigma_i^2; \gamma, \beta)} - n. \quad (2)$$

Equation (2) sums the expected number of unpublished outcomes corresponding to those published. The form of equation (2) shows why we have taken $p(y_i, \sigma_i^2; \gamma, 0) = 1$, rather than $p(y_i, \sigma_i^2; \gamma, 0) = c$ for some other c , to indicate no bias. If $c < 1$ were used, more unpublished outcomes would be implied by any weight function that is proportional to the corresponding weight function with $p(y_i, \sigma_i^2; \gamma, 0) = 1$ and hence this approach provides a minimum number of unpublished study outcomes as explained by Bowden

et al. [18].

Jackson [12] showed how the derivatives of estimates with respect to β at the origin provide insight and evaluated these analytically when focussing on the between-study variance τ^2 . Here the derivatives $d\hat{\mu}/d\beta$ are evaluated numerically at $\beta = 0$ as five times the difference between the maximum likelihood estimate of μ at $\beta = 0.2$ and the maximum likelihood estimate under the random effects model, $\beta = 0$. This rather large default difference of 0.2 when computing this derivative is often desirable as smaller differences can provide misleading numerical derivatives due to numerical error. The gradients are only approximate, but give an indication of how sensitive inferences are across the whole range of γ . Finally, $d\hat{\mu}/dN$, where N is obtained from (2), is obtained as $(d\hat{\mu}/d\beta)/(dN/d\beta)$, where $d\hat{\mu}/d\beta$ is obtained as described above and $dN/d\beta$ is calculated numerically using a much smaller default difference of 0.01.

In addition to the graphical output shown in the next section, *ABSORB* also returns seven variables. *DL* is the DerSimonian and Laird estimate of treatment effect and its standard error. The variables *estimates*, *standard errors* and *tau2ests* are 4 by $(resolution+1)$ matrices containing the estimates of treatment effect, its standard errors and the estimates of τ^2 from the sensitivity analysis. The rows correspond to the four values of the family parameters used, in ascending order, and the columns correspond to the bias parameters. The variables *gradbeta* and *gradN* are the gradients calculated to produce the gradient plots. Finally, *beta* and *family* are the values of β used in the sensitivity analysis, and the γ values used when drawing the gradient plots, respectively.

3.6 Default variables

The *ABSORB* program has a number of default values which may be changed by the user: *relStep*, *resolution*, *diff*, *diffN*, *round*, and *seed*. The variable *relStep* is the relative step size used in the finite differences by the *fdHess* command and is defined in terms of multiples of the DerSimonian and Laird standard error for the treatment effect. Its default value is 1, so that the relative step sizes used to obtain information matrices are around a standard error of the treatment effect. This is a relatively large value to those typically used in conjunction with *fdHess* but ensures that any numerical error when evaluating the likelihood numerically is small compared to the differences in the function values used to obtain derivatives.

The variable *resolution* refers to the number of steps along horizontal axes used when plotting the points from the sensitivity analyses and its default is 10. The variables *diff* and *diffN* are the differences used to obtain derivatives as described above, and have default values of 0.2 and 0.01. The variable *round* is the point where estimates of the logarithm of τ^2 are considered to be zero when determining whether to collapse to a fixed effects model when obtaining standard errors for μ . Its default is $\log(0.005)$.

Finally, despite all the care in choosing appropriate values when performing the numerical methods, difficulties can occur. Hence, before performing any numerical maximisation, rather than using the initial values as described for clarity of exposition above, a small random normal perturbation along the μ axis is added to the starting value before performing every numerical maximisation, where this perturbation has standard deviation equal to one twentieth of the DerSimonian and Laird standard deviation. In order

to check that the numerical methods have been successful, we suggest that the *ABSORB* function is run with alternative values of the *seed* variable, whose default value is 1. *R*'s random seed is set to *seed* prior to simulating any random perturbations, and if different perturbations provide the same results then the user may be confident that the numerical methods have been successful.

4 Application

In order to illustrate how *ABSORB* may be used in practice, we will use the data in Table 1, taken from Copas and Jackson [7], the results of 14 randomized clinical trials concerning the use of prophylactic corticosteroids in cases of premature birth. Briefly, if a birth is anticipated to be premature, the treatment is administered to the mother in order to improve the chance of the infant's survival. We summarize the results of these trials in Table 1, giving y_i , the estimated log-odds ratio, and the within-study variances, σ_i^2 . A negative log-odds ratio indicates that the treatment is beneficial. We also show the 'test statistics', y_i/σ_i , in table 1. As explained by Copas and Jackson [7], there is the concern that publication bias may be present and they performed a sensitivity analysis using their bound and their 'worst case weight function'.

In order to perform a sensitivity analysis using less extreme weight functions, all three types of weight functions described in section 3.3 will be used. The first type was used with maximum $\beta = 0.9$ and γ in the range $(-1.6, 0)$. This range of γ ensures that at least one study is in each of the 'more likely to be published' and 'less likely to be published' categories . If all studies

Table 1: The corticosteroids data.

Study	y_i	σ_i^2	y_i/σ_i
1	-1.55	0.41	-2.43
2	-1.49	0.87	-1.59
3	-1.33	0.34	-2.27
4	-0.35	0.14	-0.95
5	-0.19	0.26	-0.38
6	-0.43	0.28	-0.81
7	-0.61	0.22	-1.29
8	-0.97	0.46	-1.44
9	-1.64	0.83	-1.80
10	-1.19	2.78	-0.71
11	-0.28	0.46	-0.41
12	0.03	0.25	0.06
13	-0.06	0.07	-0.23
14	-0.54	0.05	-2.41

are in just one of these categories then weight functions types 1 and 2 simply provide the same inferences as the random effects model. With the bias parameter set to a maximum of $\beta = 0.9$, less encouraging results only have a probability of 0.1 of being published.

Having defined wf using the first set of illustrative code in section 3.3, the command `ABSORB(ys, sigma2s, betamax=0.9, family=c(-1.6,0))` performs the sensitivity analysis, where ys and $sigma2s$ are vectors containing the data, and the results are shown in figure 1. *ABSORB* shows the results of the sensitivity analyses for the four values of γ as described above in the plots in the first two rows. Here 95% intervals are shown using the estimates $\hat{\mu}$, their standard errors and the usual normal approximations. The DerSimonian and Laird estimate and interval are shown with solid points and a line at $\beta = 0$ and show that these and the maximum likelihood inferences are in agreement under the random effects model, both of which collapse to a fixed

effects model.

The weight function is converted to the number of unpublished results for four values of β using (2) and these are shown at the top of the plots to aid interpretation. The derivatives of $\hat{\mu}$ with respect to β and N are also shown, over the entire range of the family values specified when calling *ABSORB*, in the third and final row. These derivatives are not continuous because as the family value changes there are step changes in weight functions of types 1 and 2. Hence these plots, which are produced by plotting the gradients over 11 points and connecting these with lines, can only attempt to to an indication of the magnitudes of the gradients for these kinds of weight functions.

From the first two rows of figure 1 we see that the null hypothesis $H_0 : \mu = 0$ is only brought into question γ lies in the middle of the data and β is around 0.9. The robustness of the inference of a treatment effect is confirmed from the gradient plots. The first of these shows that the family value $\gamma = -0.53$ is one of the most sensitive models and even this needs a very large $\beta \approx 0.9$, which corresponds to around 54 unpublished outcomes, to threaten the null. Unless the family parameter $\gamma \approx 0$, the second derivative plot further shows that a few additional outcomes can only increase the treatment effect by around 0.02 to 0.03 per outcome.

In order to perform a sensitivity analysis using the second type of weight function, the definition of *wf* was redefined as described in section 3.2 and *ABSORB(ys, sigma2s, betamax=0.9, family=c(-2.4,0))* was called. The results are shown in figure 2 and similar conclusions are obtained to those using the first type of weight function.

Finally a continuous weight function was considered using the third type

of weight function by redefining wf again and calling $ABSORB(ys, sigma2s, betamax=6, family=c(0.5,1.5))$; $\gamma = 1$ has been a common choice and a range of possibilities around this was desired. For $\gamma = 1$ and $\beta = 6$, the probability of publishing a ‘null outcome’ (one with a p-value of 0.5, and therefore an effect of 0) is only around 0.05 so $\beta = 6$ can be interpreted as a very large bias parameter. The results are shown in figure 3 which again show that the inference that there is a treatment effect is robust to the possibility of bias, because for each γ shown we require $\beta \approx 3$ to threaten the null hypothesis. For $\gamma = 1$ for example, $\beta = 3$ means that a ‘null outcome’ only has a probability of 0.22 of being published and the probability of publishing any studies with outcomes that indicate the treatment is harmful quickly decays. For $\gamma = 1.5$ relatively few studies appear to correspond to a model which provides a confidence interval that includes the null but we emphasise that this is a rough guide and because so many of the outcomes in this sample are statistically significant, these outcomes are very likely to be published under this model. Hence the number of unpublished studies corresponding to those that have been observed is necessarily small. The simplistic nature of the calculation of N should be borne in mind when considering this quantity.

It is also interesting to note that, for this weight function, as β becomes larger the random effects model no longer collapses to a fixed effects model and the confidence intervals widen when this occurs. Publication bias also has implications for estimates of between-study variance [12, 19] and this is emphasised by this example.

To summarise, figures 1, 2 and 3 suggest that quite severe bias is needed to call the null hypothesis of no treatment effect into question. In partic-

ular, more unpublished results are needed than in the worst case scenario considered by Copas and Jackson [7]. Smaller degrees of bias and numbers of unpublished outcomes are needed to bring the clinical significance of the treatment into question but the results shown can be used to assess all the various questions the investigator may have about the impact of outcome reporting bias.

5 Discussion

The aim of *ABSORB* is to overcome the practical difficulties encountered when using weight functions to model outcome reporting bias and to facilitate their routine use. It is hoped that this program will prove useful for those concerned that this type of bias may be present and wish to assess this quantitatively. The weight functions provided can easily be amended to correspond to analogous weight functions where positive outcomes indicate a positive treatment effect and are more publishable. Alternatively the user could change the signs of their estimates, use the weight functions provided, and obtain the results from the sensitivity analysis by symmetry.

Other types of bias are also possible in meta-analysis, for example those due to the methodological quality of studies and the populations they sample. *ABSORB* makes no attempt to address these other biases. A further limitation is that only one bias parameter is permitted and a family parameter is required. This places some restrictions on what may be used for the weight function but a wide variety forms of this function can be handled. *ABSORB* provides no assistance with the problem of assessing whether bias may be present, or the degree of this, from evidence such as the funnel plot

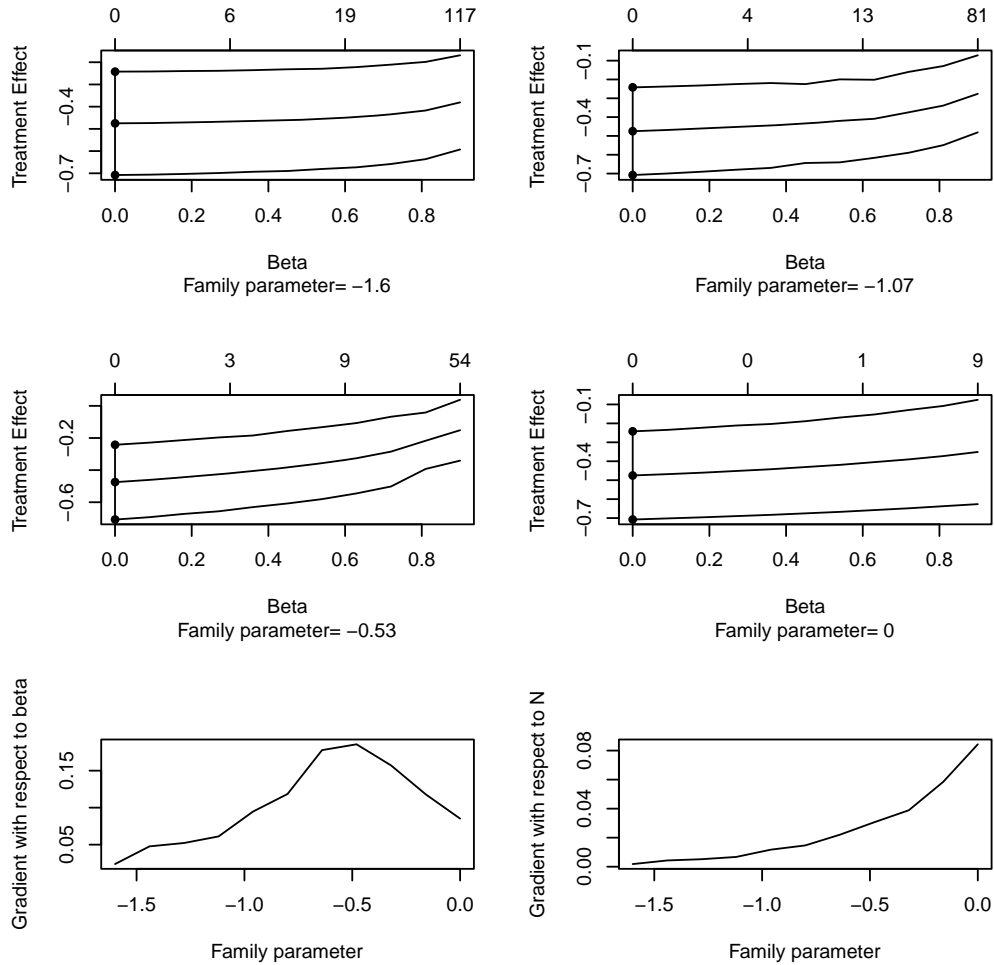


Figure 1: Sensitivity analysis for the corticosteroids data. Output using the first type of weight function and $ABSORB(y_s, \sigma_{2s}, \beta_{max}=0.9, family=c(-1.6, 0))$. The rough conversion to the number of unpublished studies is shown at the top of the first four plots.

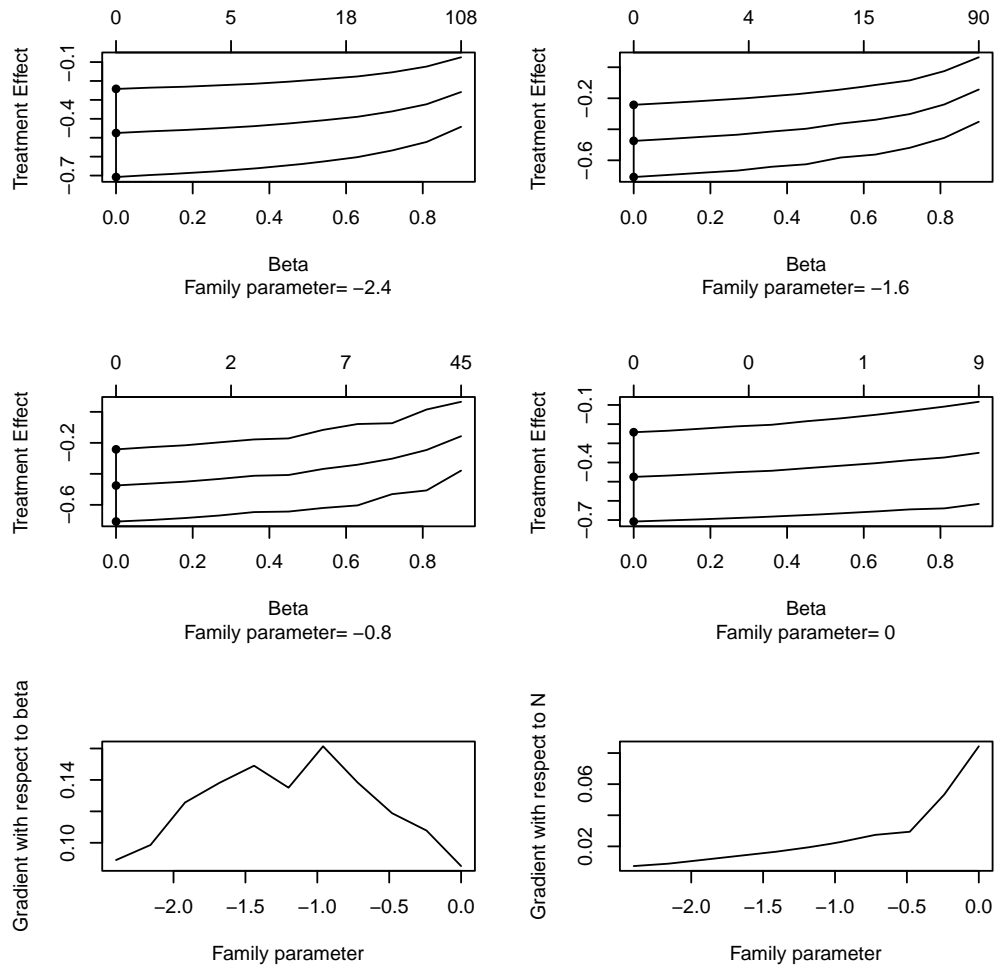


Figure 2: Sensitivity analysis for the corticosteroids data. Output using the second type of weight function and $ABSORB(ys, sigma2s, betamax=0.9, family=c(-2.4, 0))$. The rough conversion to the number of unpublished studies is shown at the top of the first four plots.

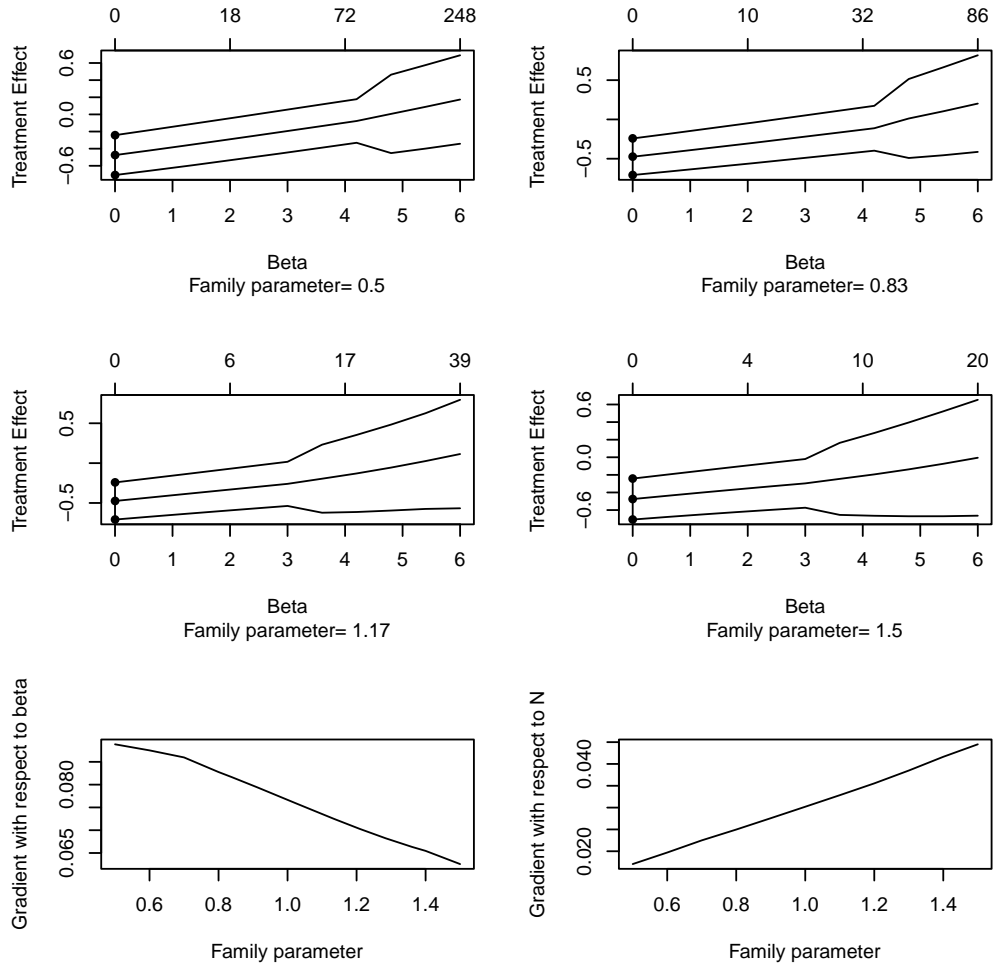


Figure 3: Sensitivity analysis for the corticosteroids data. Output using the third type of weight function and $ABSORB(ys, sigma2s, betamax=6, family=c(0.5,1.5))$. The rough conversion to the number of unpublished studies is shown at the top of the first four plots.

or the investigators' subject area knowledge. *ABSORB* must be used with other methods and is an addition to the armoury, rather than the solution to, the problems publication bias presents for meta-analysis. In particular, it is a poor substitute to finding all the outcomes, but is potentially useful if this is not feasible.

The use of weight functions is fairly computationally intensive, and a few minutes is needed for *ABSORB* to produce its output for typical meta-analyses, but compared to the time taken to search for trials and outcomes it may be regarded as fast. Every effort has been made to make the program as efficient as possible and yet ensure that it performs well. For example, the program could be made faster by using quasi-Newton maximisation algorithms but the slower Nelder-Mead has been preferred because it is more robust.

The *R* code for the *ABSORB* program is available at

<http://www.mrc-bsu.cam.ac.uk/Software/download.html#Rpackages>.

Acknowledgement

DJ is employed by the UK Medical Research Council (grant code U.1052.00.006).

References

- [1] P. Nieminen, G. Rucker, J. Miettunen, J.R. Carpenter, and M Schumacher. Statistically significant papers in psychiatry were cited more often than others. *Journal of Clinical Epidemiology*, 60:939–946, 2007.

- [2] A.O. Carter, G.H. Griffin, and T.P. Catter. A survey identified publication bias in the secondary literature. *Journal of Clinical Epidemiology*, 59:241–245, 2006.
- [3] M. Egger, G.D. Smith, M. Schneider, and C. Minder. Bias in meta-analysis detected by a simple, graphical test. *British Medical Journal*, 315:629–634, 1997.
- [4] D. R. Jones J. L. Peters, A. J. Sutton and K. R. Abrams. Assessing publication bias in meta-analyses in the presence of between-study heterogeneity. *Journal of the Royal Statistical Society Series A*, 173:575–591, 2010.
- [5] L. V. Hedges. Modeling publication selection effects in meta-analysis. *Statistical Science*, 7:246–255, 1992.
- [6] S. Iyengar and J. Greenhouse. Selection models and the file drawer problem. *Statistical Science*, 3(1):109–117, 1988.
- [7] J. Copas and D. Jackson. A bound for publication bias based on the fraction of unpublished studies. *Biometrics*, 60:146–153, 2004.
- [8] G.H. Givens and G. Smith. Publication bias in meta-analysis: a Bayesian approach to account for issues exemplified in the passive smoking debate. *Statistical Science*, 12:221–250, 1997.
- [9] C. Preston, D. Ashby, and R. Smyth. Adjusting for publication bias: Modelling the selection process. *Journal of Evaluation in Clinical Practice*, 10:313–322, 2004.

- [10] J. Copas and J.Q. Shi. Meta-analysis, funnel plots and sensitivity analysis. *Biostatistics*, 1:247–262, 2000.
- [11] M. G. Kenward. Selection models for repeated measurements with non-random dropout: an illustration of sensitivity. *Statistics in Medicine*, 17:2723–2732, 1998.
- [12] D. Jackson. Assessing the implications of publication bias for two popular estimates of between-study variance in meta-analysis. *Biometrics*, 63:187–193, 2006.
- [13] S. Duval and R. Tweedie. Trim and fill a simple funnel-plot based method for testing and adjusting for publication bias meta-analysis. *Biometrics*, 56:455–463, 2000.
- [14] J. Copas and J.Q. Shi. A sensitivity analysis for publication bias in systematic reviews. *Statistical methods in medical research*, 10:251–265, 2001.
- [15] J.R. Carpenter, G. Schwartzer, and G. Rucker. Empirical evaluation showed that the copas selection model provided a useful summary in 80% of meta-analyses. *Journal of clinical epidemiology*, 62:624–631, 2009.
- [16] R. DerSimonian and N. Laird. Meta-analysis in clinical trials. *Controlled Clinical Trials*, 7:177–188, 1986.
- [17] R.J. Hardy and S.G. Thompson. A likelihood approach to meta-analysis with random effects. *Statistics in Medicine*, 15:619–629, 1996.

- [18] J. Bowden, D. Jackson, and S.G. Thompson. Modelling multiple sources of dissemination bias in meta-analysis. *Statistics in Medicine*, 29:945–955, 2010.
- [19] D. Jackson. The implications of publication bias for meta-analysis other parameter. *Statistics in Medicine*, 25:2911–2921, 2006.